Time-Efficient Methods for Non-Invasive Brain Blood Flow Imaging Using Arterial Spin Labelling MRI

Joseph G. Woods

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Supervisors: Thomas W. Okell and Michael A. Chappell

Trinity 2019

Time-Efficient Methods for Non-Invasive Brain Blood Flow Imaging Using Arterial Spin Labelling MRI

Joseph G. Woods

Linacre College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Trinity Term 2019

Abstract

Accurate knowledge of brain tissue perfusion and clear visualisation of cerebral arteries are of great interest for the diagnosis and research of a range of diseases and conditions such as stroke, dementia, Moyamoya, and cancer, as well as in drug development. ^{15}O-H$_2$0 PET imaging is considered the gold standard for measuring cerebral blood flow (CBF), while contrast-enhanced x-ray digital subtraction angiography is the gold standard for visualising cerebral arteries. However, both methods use ionising radiation and injectable contrast agents so carry some risks to the patient and are not appropriate for longitudinal monitoring or research.

There is growing interest in arterial spin labelling (ASL) magnetic resonance imaging (MRI), which uses magnetically labelled blood as an endogenous and freely diffusible tracer and can be flexibly used to visualise the arterial tree as well as to measure CBF. However, ASL is an intrinsically low SNR technique, requiring signal averaging and long scan times.

The main aim of this thesis is to improve the time-efficiency of ASL MRI perfusion and angiographic scans, increasing the accuracy and SNR of CBF measurements and vessel depiction in clinically relevant scan times.

First, a statistical framework was developed and validated for optimising multi-delay ASL perfusion protocol timings to improve the accuracy of CBF and arterial transit time measurements. This flexible framework can be targeted to improve the accuracy of specific parameters and provides a more informed method for designing ASL protocols.

Next, this framework was used to optimise and objectively compare a range of existing pseudo-continuous ASL (PCASL) perfusion protocol designs, including time-encoded preparations, from the literature and to explore novel protocol designs. These comparisons allowed the advantages and disadvantages of each approach to be explored, and recommendations about the best ASL experimental designs to be made.

Finally, the benefits of time-encoded PCASL preparations for ASL angiography were investigated. Angiography visualises the pass-through of labelled blood, in contrast to perfusion imaging which relies on the accumulation of labelled blood in tissue, and so the short label durations used in time-encoded PCASL have a greatly reduced impact on the magnitude of the ASL signal. A variable flip-angle approach was also used to remove signal discontinuities and demonstrate a significant SNR improvement over conventional PCASL angiography.

Acknowledgements

There are few pursuits in which one person deserves all of the credit and my DPhil is certainly not one of these. I owe much to my colleagues, friends, and family, all of whom I greatly admire and who have helped, supported, and advised me during my DPhil. These acknowledgements are by no means exhaustive, but simply highlight a few of the many inspiring people who I have had the pleasure to have crossed paths with in Oxford.

Firstly, I would like to thank my supervisors, Thomas Okell and Michael Chappell. When I started my research, I had only rudimentary knowledge of MRI and knew next to nothing about arterial spin labelling. They have invested huge amounts of time and effort in guiding my research, providing interesting ideas, and answering my many questions. Michael has a real knack for clearly explaining complicated analysis topics and has, on many occasions, helped me condense which questions are most interesting to ask and how best to ask them. Tom is a superb scientist and supervisor and has undoubtedly given me far more of his time than I could have expected, for which I am extremely grateful. I have always looked forward to, and enjoyed, our weekly meetings, where we would discuss new ideas and I would quiz him on some aspect of MRI physics that I had recently come across. I believe the longest of these meetings lasted four hours! I hope that some of that time was also useful to him.

There are many people in FMRIB and from the ONBI course who have helped make my time studying in Oxford all the more enjoyable. In particular, from ONBI, Martin Hailstone for being a distracting but excellent desk neighbour, Jim Smith for instigating various great ideas such as ONBI assassin and MATLAB Snake and Tetris, Zobair Arya and Jack Allen for both moving into research at FMRIB with me and being highly supportive throughout our DPhils. At FMRIB, Robert Brand for his excellent sense of humour, culinary skills, and advice; Emma Lawrance for being the most approachable and welcoming person I've ever met, being one of two members of Sandwich Club, and for all of the adventures we've shared; Adam Berrington for putting up with me distracting him and always being great fun; Hongbae Jeong for his kindness and dancing; Mark Chiew for his excellent teaching style, generosity with his time, and good humour; Nicholas Blockley for always giving me a good laugh; Will Clarke for his help with Bloch simulations and Yuriko Suzuki for helping me with adiabatic inversion pulses. Other notable mentions include Alon Baram, Fritz Bayer, Naiara Demnitz, Nadine Grädel, Amy Howard, Ainslie Johnstone, Rick Lange, Alberto Lazari, Tim Muller, Caitlin O'Brien, Sophie Schauman, Yoko Spirig, Yan Tong, Olivia Viessmann, Thomas Wassenaar, and Wenchuan Wu. I would also like to thank Karla Miller and Peter Jezzard for their support and for being great role models for everyone in the lab, and all of the WIN and AVIC core staff for keeping everything running smoothly.

Outside of the lab, I have been very fortunate to befriend many kind and interesting people, which I have found to be one of the best aspects of studying in Oxford. My fellow Bradmore Road housemates made it easy to settle in: Luke Wonneck with his infectious excitement and interesting hallway discussions, Adam Formica with his friendliness and Americana music drifting through my open window, Carlos Rios Ocampo with his compassion and generosity, as well as Maro Dotulong, Arsalan Ihsan, Claire Michaud, Ndjodi Ndeunyema, Moé Suzuki, and Donata Wollensak. My housemates at my long-term home on Weldon Road made it easy to feel at home in Oxford: Kristiina Visakorpi with

her love of the outdoors and her empathy, Zobair Arya with his friendly competitiveness and humour, Chuor de Garang Alier, Gerardo Rodriguez Hernandez, Einar Magnússon, Julio Acosta Matos, Ndjodi Ndeunyema, Gladys Ngetich, and Carlos Rios Ocampo.

My friends in the Coffee Club and the Easter Hike provided endless hours of joy and happiness which supplied a welcome escape from work: Geoff Stanley with his excellent planning abilities and physical feats, James Mitchell with his indomitable good spirits and coffee addiction, Erica Bower with her exuberance and resilience, Beate Dirks, Ed and Imo Doddridge, Victoria Wyllie de Echeverria, Adam Formica, Rachel Friedman, Lea Kliem, Milan Kloewer, Emma Lawrance, Carlos Rios Ocampo, and Kristiina Visakorpi.

The people who I have met at Linacre College and through my friends here have also played a huge part in helping me through my DPhil. Amongst these include: Milton Barbosa with his kindness, Zali Fung with her adventurous nature, Nanna Rask with her warm-heartedness, Louis Spanias with his excellent sense of humour, Ashley Tsai with her gentleness, Megan Turnbull with her empathy, and Shimeng Zhou with her intuitive pensiveness. My bandmates, Julianne Viola and Adam Formica, gave me many evenings of joy during our practices and performances. The 25 Friends group lead to many joyous occasions and another musical outlet and I formed many good memories with my friends at the college rowing club: Katarzyna Anna, Will Budd, Giovanni Carù, Chris Cooper, Matthew Davey, Ben Eacott, Adam Formica, Tobias Højgaard, Conor Manders, Ask Neve, Dingeman Wiertz, Derek Xu, and Qifeng Yang.

My family have given me unwavering support along the way, for which I am very grateful. Sheltering and feeding me during holidays and weekends away, listening to me ramble on about my work, and providing much-needed hugs are just some of the things I am thankful for.

Last but not least, I owe my greatest debt of gratitude to Zeena-Britt Sanders. You picked me up when things were not going well and celebrated with me when they did. You came up with so many adventures for us to go on and endlessly went in the scanner for me so I could test new sequences. You listened to my presentations over and over again and helped make presenting less intimidating. Your love and support really did drag me past the finish line and I’m not sure how I would have done it without you. You’re one of the strongest and most caring people I know and I feel very lucky to have you in my life.

Contents

Preface

Preface

List of Publications Arising from this Thesis

Journal Articles

1. **Woods JG**, Chappell MA, Okell TW. A general framework for optimizing arterial spin labeling MRI experiments. *Magn Reson Med* 2019;81:2474–2488.

Conference Proceedings

1. **Woods JG**, Chappell MA, Okell TW. Optimizing Post Labeling Delays in Multiple-Delay Arterial Spin Labeling MRI for Cerebral Perfusion Imaging. In: *Proceedings of the 34th Annual Meeting of the ESMRMB, Barcelona, Spain*, 2017. p. 245.
2. **Woods JG**, Chappell MA, Okell TW. Comparison of optimized single-PLD, sequential multi-PLD and time-encoded PCASL for cerebral blood flow measurements. In: *Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France*, 2018. p. 2162.
3. **Woods JG**, Chappell MA, Okell TW. A novel hybrid of time-encoded and sequential multi-PLD PCASL for improved cerebral blood flow estimation. In: *Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France* ,2018. p. 2161.
4. **Woods JG**, Schauman SS, Chiew M, Chappell MA, Okell TW. Optimization of time-encoded pseudo-continuous ASL angiography with a variable flip-angle scheme. In: *Proceedings of the 27th Annual Meeting of the ISMRM, Montreal, Canada*, 2019. p. 844.

1
Introduction

1.1 Motivation

Although the brain only represents about 2% of body weight, it is the most perfused organ in the body and uses ~20% of the body's total resting oxygen consumption. Blood continually transports nutrients and oxygen to the brain via the internal carotid and vertebral arteries into a dense network of smaller arteries. Due to its large metabolic demands and its low level of energy reserves, if the supply of blood to the tissue is interrupted for even just a few minutes, irreversible catastrophic cell death will occur.[1]

The rate of blood flow to the cerebral tissue capillary bed, known as perfusion, is sensitive not only to a sudden interruption of the arterial supply,[2–4] but also to neuronal activity, response to drug administration,[5] and a range of diseases including: dementia,[6,7] cancer,[8] depression,[9] epilepsy,[10] arteriovenous malformations,[11] and Moyamoya disease.[12] Measurement of cerebral perfusion and visualisation of the arterial supply is therefore of great interest for basic neuroscience and clinical research studies, as well as in clinical practice.

1.1.1 Angiography

A common technique used for visualisation of the cerebral arteries is x-ray digital subtraction angiography (DSA).[13] A catheter is inserted into the femoral artery and guided to the artery of interest. An x-ray opaque contrast agent (such as iodine) is injected which can then be imaged from several angles. It can provide excellent spatial and temporal resolution with a high signal-to-noise ratio (SNR). However, its downsides include exposure to ionising radiation, possible reaction to the contrast agent, and an increased risk of stroke,[14] making its repeated use problematic.

Alternatively, arterial spin labelling[15,16] (ASL) MRI uses magnetically labelled hydrogen-1 (^{1}H) in blood water as a freely diffusible endogenous contrast agent, making it completely non-invasive. It can be used for angiography by imaging the labelled blood while still in the large arteries.[17,18] There are no safety concerns beyond the standard concerns for MRI (tissue heating, projectiles, dislodgement of metallic implants, and hearing damage),[19] ionising radiation is not used, and because there is no injection of contrast, there is no increased risk of an allergic reaction. This makes ASL particularly useful for longitudinal monitoring. However, due to a much lower intrinsic SNR and a tracer half-life of ~1.65 s at 3T,[20] ASL angiography cannot achieve the same spatial or temporal resolution as x-ray DSA, unless long scan times are used.

Other MR angiography techniques include contrast enhanced imaging,[21] time-of-flight[22] (TOF) imaging and phase-contrast imaging.[23] Contrast enhanced imaging uses an injected paramagnetic contrast agent (typically Gadolinium-based), resulting in relatively high SNR. However, because the first passage of the bolus must be captured, there is a trade-off between high temporal resolution and spatial resolution. In contrast, ASL angiography can repeat the preparation many times to achieve the desired temporal and spatial resolutions. There is also concern over the safety of Gadolinium-based contrast agents.[24] TOF and phase-contrast imaging are both non-invasive methods, using blood as an endogenous image contrast, however TOF imaging cannot provide dynamic information, both are insensitive to slow-flowing blood,[25] and phase-contrast imaging suffers from partial volume effects requiring high spatial resolutions which result in long scan times.

1.1.2 Perfusion

There is a growing interest in the absolute quantification of brain tissue perfusion, both in research and in the clinic, due to the correlation between local perfusion and tissue health in a range of diseases. This is evidenced by the publication of the ASL consensus paper in 2015, which sets out the standard recommended ASL implementation and quantification process, and the large number of citations it has attracted.[26]

Oxygen-15 (^{15}O) positron emission tomography[27] (PET) is considered the gold standard method for quantitative perfusion imaging. However, the patient is exposed to ionising radiation, arterial cannulation is required to measure an arterial input function, and due to the short half-life of ^{15}O (~2 minutes), an on-site cyclotron is required.[28] On the other hand, beyond the previously stated advantages of non-invasiveness and no ionising radiation, ASL MRI is able to provide quantitative perfusion measurements[15,16] with higher SNR and higher spatial resolution than ^{15}O PET due to signal averaging.[28] In comparison to ASL angiography, this is performed by imaging the labelled blood water once it has exchanged into the tissue.

Another MRI perfusion technique is dynamic susceptibility contrast (DSC) imaging.[29,30] DSC MRI requires an injection of a Gadolinium-based contrast agent, which is imaged during its first pass through the vasculature. Since the contrast agent largely remains intravascular during this time, the technique is mostly sensitive to blood volume, but perfusion can theoretically be calculated using the central volume principle.[31]

1.1.3 Summary

ASL MRI can be flexibly used to provide dynamic images of arterial vessels (which can be quantitatively modelled)[32,33] as well as quantitative perfusion estimates. However, the technique suffers from intrinsically low SNR. Therefore, methods for improving the SNR and accuracy of the resulting angiograms and perfusion measurements, while keeping scan times short, are important to improve the utility of ASL MRI in research and the clinic.

1.2 Thesis Outline

This thesis describes the development of several techniques to improve the time-efficiency of ASL-based brain perfusion and angiography scans, by improving the quantification accuracy of ASL perfusion imaging and the SNR of ASL dynamic angiography.

In Chapter 2, an overview of the background MRI principles and techniques used throughout this thesis are given.

Chapter 3 provides an overview of key components used in ASL perfusion and angiography sequences and modelling. This includes a description of several labelling techniques, dynamic sampling methods, image readouts, background suppression strategies, and a commonly used perfusion model.

Chapter 4 describes a novel protocol optimisation framework for improving the estimation of CBF and arterial transit time (ATT) with multi-delay ASL perfusion imaging. This framework is used to optimise a pseudo-continuous ASL (PCASL) multi-delay protocol in order to maximise the measurement accuracy of CBF and ATT or just CBF. The resulting improvements to measurement accuracy are then validated using Monte Carlo simulations and in vivo experiments.

In Chapter 5, several existing and novel PCASL perfusion protocol designs are optimised using the framework developed in Chapter 4 to maximise their CBF measurement accuracy. These accuracies are then compared using Monte Carlo simulations and in vivo experiments to identify which protocol design results in the highest CBF accuracy.

In Chapter 6, a time-encoded PCASL labelling scheme is combined with an angiographic readout. It is shown that this allows some of the temporal information to be generated from the preparation, allowing shorter readouts with larger flip-angles to be used. A variable flip-angle scheme is used to reduce signal discontinuities in the time-encoded data which also maximises SNR across the readout. The resulting increase in SNR is demonstrated using simulations and in vivo experiments. Further time-efficiency improvements are presented, including undersampling the acquisition of the time-encoded protocol and improving the suppression of background signal across the readout, for which preliminary in vivo results are presented.

Finally, Chapter 7 provides a summary of the work undertaken in this thesis and a conclusion relating to the major findings. Several possible directions for future research building on the work in this thesis are also outlined.

1.3 References

1 Saver JL. Time Is Brain—Quantified. *Stroke* 2006; **37**: 263–266.

2 Wang DJJ, Alger JR, Qiao JX, Gunther M, Pope WB, Saver JL *et al.* Multi-delay multi-parametric arterial spin-labeled perfusion MRI in acute ischemic stroke - Comparison with dynamic susceptibility contrast enhanced perfusion imaging. *NeuroImage Clin* 2013; **3**: 1–7.

3 Harston GWJ, Okell TW, Sheerin F, Schulz U, Mathieson P, Reckless I *et al.* Quantification of Serial Cerebral Blood Flow in Acute Stroke Using Arterial Spin Labeling. *Stroke* 2017; **48**: 123–130.

4 Chalela JA, Alsop DC, Gonzalez-Atavales JB, Maldjian JA, Kasner SE, Detre JA. Magnetic Resonance Perfusion Imaging in Acute Ischemic Stroke Using Continuous Arterial Spin Labeling. *Stroke* 2000; **31**: 680–687.

5 Macintosh BJ, Pattinson KTS, Gallichan D, Ahmad I, Miller KL, Feinberg DA *et al.* Measuring the Effects of Remifentanil on Cerebral Blood Flow and Arterial Arrival Time Using 3D Grase MRI with Pulsed Arterial Spin Labelling. *J Cereb Blood Flow Metab* 2008; **28**: 1514–1522.

6 Wierenga CE, Hays CC, Zlatar ZZ. Cerebral blood flow measured by arterial spin labeling MRI as a preclinical marker of alzheimer's disease. *J Alzheimer's Dis* 2014; **42**: S411–S419.

7 Gao Y-Z, Zhang J-J, Liu H, Wu G-Y, Xiong L, Shu M. Regional Cerebral Blood Flow and Cerebrovascular Reactivity in Alzheimer's Disease and Vascular Dementia Assessed by Arterial Spinlabeling Magnetic Resonance Imaging. *Curr Neurovasc Res* 2012; **10**: 49–53.

8 Wang P, Li J, Diao Q, Lin Y, Zhang J, Li L *et al.* Assessment of glioma response to radiotherapy using 3D pulsed-continuous arterial spin labeling and 3D segmented volume. *Eur J Radiol* 2016; **85**: 1987–1992.

9 Duhameau B, Ferré J-C, Jannin P, Gauvrit J-Y, Vérin M, Millet B *et al.* Chronic and treatment-resistant depression: A study using arterial spin labeling perfusion MRI at 3Tesla. *Psychiatry Res Neuroimaging* 2010; **182**: 111–116.

10 Leonhardt G, de Greiff A, Weber J, Ludwig T, Wiedemayer H, Forsting M *et al.* Brain Perfusion Following Single Seizures. *Epilepsia* 2005; **46**: 1943–1949.

11 Wolf RL, Wang J, Detre JA, Zager EL, Hurst RW. Arteriovenous Shunt Visualization in Arteriovenous Malformations with Arterial Spin-Labeling MR Imaging. *Am J Neuroradiol* 2008; **29**: 681–687.

12 Wang R, Yu S, Alger JR, Zuo Z, Chen J, Wang R *et al.* Multi-delay arterial spin labeling perfusion MRI in moyamoya disease–comparison with CT perfusion imaging. *Eur Radiol* 2014; **24**: 1135–1144.

13 Jeans WD. The development and use of digital subtraction angiography. *Br J Radiol* 1990; **63**: 161–168.

14 Cloft HJ, Lynn MJ, Feldmann E, Chimowitz M. Risk of Cerebral Angiography in Patients with Symptomatic Intracranial Atherosclerotic Stenosis. *Cerebrovasc Dis* 2011; **31**: 588–591.

15 Detre JA, Leigh JS, Williams DS, Koretsky AP. Perfusion imaging. *Magn Reson Med* 1992; **23**: 37–45.

16 Williams DS, Detre JA, Leigh JS, Koretsky AP. Magnetic resonance imaging of perfusion using spin inversion of arterial water. *Proc Natl Acad Sci U S A* 1992; **89**: 4220–4220.

17 Dixon WT, Du LN, Faul DD, Gado M, Rossnick S. Projection angiograms of blood labeled by adiabatic fast passage. *Magn Reson Med* 1986; **3**: 454–462.

18 Wang SJ, Nishimura DG, Macovski A. Multiple-readout selective inversion recovery angiography. *Magn Reson Med* 1991; **17**: 244–251.

19 Durbridge G. Magnetic Resonance Imaging: Fundamental Safety Issues. *J Orthop Sport Phys Ther* 2011; **41**: 820–828.

20 Lu H, Clingman C, Golay X, van Zijl PCM. Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn Reson Med* 2004; **52**: 679–682.

21 Zhang H, Maki JH, Prince MR. 3D contrast-enhanced MR angiography. *J Magn Reson Imaging* 2007; **25**: 13–25.

22 Masaryk TJ, Modic MT, Ross JS, Ruggieri PM, Laub GA, Lenz GW *et al.* Intracranial circulation: preliminary clinical results with three-dimensional (volume) MR angiography. *Radiology* 1989; **171**: 793–799.

23 Moran PR. A flow velocity zeugmatographic interlace for NMR imaging in humans. *Magn Reson Imaging* 1982; **1**: 197–203.

24 Kuo PH, Kanal E, Abu-Alfa AK, Cowper SE. Gadolinium-based MR Contrast Agents and Nephrogenic Systemic Fibrosis. *Radiology* 2007; **242**: 647–649.

25 Hendrikse J, Klijn CJM, van Huffelen AC, Kappelle LJ, van der Grond J. Diagnosing Cerebral Collateral Flow Patterns: Accuracy of Non-Invasive Testing. *Cerebrovasc Dis* 2008; **25**: 430–437.

26 Alsop DC, Detre JA, Golay X, Günther M, Hendrikse J, Hernandez-Garcia L *et al.* Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn Reson Med* 2015; **73**: 102–116.

27 Herscovitch P, Markham J, Raichle ME. Brain blood flow measured with intravenous H2(15)O. I. Theory and error analysis. *J Nucl Med Off Publ Soc Nucl*

Med 1983; **24**: 782–9.

28 Fan AP, Jahanian H, Holdsworth SJ, Zaharchuk G. Comparison of cerebral blood flow measurement with [15 O]-water positron emission tomography and arterial spin labeling magnetic resonance imaging: A systematic review. *J Cereb Blood Flow Metab* 2016; **36**: 842–861.

29 Østergaard L, Weisskoff RM, Chesler DA, Gyldensted C, Rosen BR. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part I: Mathematical approach and statistical analysis. *Magn Reson Med* 1996; **36**: 715–725.

30 Østergaard L, Sorensen AG, Kwong KK, Weisskoff RM, Gyldensted C, Rosen BR. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part II: Experimental comparison and preliminary results. *Magn Reson Med* 1996; **36**: 726–736.

31 Wong EC. An introduction to ASL labeling techniques. *J Magn Reson Imaging* 2014; **40**: 1–10.

32 Okell TW, Chappell M a., Schulz UG, Jezzard P. A kinetic model for vessel-encoded dynamic angiography with arterial spin labeling. *Magn Reson Med* 2012; **68**: 969–979.

33 Okell TW, Chappell MA, Jezzard P. A theoretical framework for quantifying blood volume flow rate from dynamic angiographic data and application to vessel-encoded arterial spin labeling MRI. *Med Image Anal* 2013; **17**: 1025–1036.

2

Principles of Magnetic Resonance Imaging

The aim of this chapter is to provide a brief introduction to the basic principles of MRI acquisition and image formation. Descriptions of the MRI signal source, excitation, relaxation, and spatial encoding using the *k*-space formalism, are given. Furthermore, a description of a number of key concepts relevant to this thesis are provided, including adiabatic inversion and several readout trajectories. A more in-depth description of the phenomena and techniques used in the broad field of MRI can be found in references [1–3] and the references therein, upon which this chapter is based.

A note on notation: throughout this thesis, boldface is used for matrices while an arrow accent is used for column vectors.

2.1 The NMR Signal

2.1.1 Magnetisation

All nucleons (protons and neutrons) possess an angular momentum, commonly known as spin, which is quantum mechanical in nature. Protons and neutrons possess a spin of ½. For a nucleus to be 'MR visible', it must have a non-zero net spin. This is the case for hydrogen-1 (^{1}H), which has a spin of ½ and is extremely naturally abundant in the human body. For these reasons, ^{1}H is the most commonly used nucleus in MRI.

When a spin-½ particle is placed within an external magnetic field, $\vec{B}$, according to the Zeeman effect, two energy levels are possible: a lower-energy state which is aligned with $\vec{B}$, 'spin-up', and a higher-energy state which is anti-aligned with $\vec{B}$, 'spin-down'. The energy difference, ΔE, between these two states is given by

$$\Delta E = \hbar\gamma B, \tag{2.1}$$

where $\hbar$ is Planck's constant ($6.62607 \cdot 10^{-34}$ J · s) divided by 2π, γ is the gyromagnetic ratio of the nucleus ($42.5775 \cdot 10^{6}$ s^{-1} · T^{-1} for ^{1}H), and B is the magnitude of the external magnetic field. Transitions of spins between these two energy states can, therefore, occur by the emission or absorption of a photon with an energy equal to ΔE and angular frequency, ω:

$$\omega = \gamma B, \tag{2.2}$$

known as the Larmor frequency.

For a large spin ensemble, the ratio between the number of spins occupying the low energy state, N_1, and the high energy state, N_2, can be described by the Boltzmann distribution:

$$\frac{N_1}{N_2} = e^{\left(\frac{\hbar\omega}{kT}\right)} \cong 1 + \frac{\hbar\omega}{kT}, \tag{2.3}$$

where k is Boltzmann's constant ($1.38065 \cdot 10^{-23}$ J · K^{-1}) and T is the temperature of the spin environment in units of Kelvin. Note, this first-order approximation is possible because $kT \gg \hbar\omega$ for body temperature (~310 K) and clinical magnetic field strengths (0.2T - 7T). For $T = 310$ K and $B = 3$ T (the field strength used throughout this thesis), this corresponds to an excess of only about 10 spins per million that align with $\vec{B}$, resulting in a small net magnetic moment, $\vec{M}$. However, since ^{1}H is extremely abundant in the body, a signal can be observed in practice.

For ^{1}H, the magnitude of the net magnetic moment for a large spin ensemble is therefore:

$$M \cong N\left(\frac{\hbar\gamma}{2}\right)^2 \frac{B}{kT}, \tag{2.4}$$

where *N* is the total number of spins in the ensemble. Perturbation of this spin ensemble away from $\vec{B}$ is the source of the signal in MRI. Importantly, the magnitude of this net magnetic moment is proportional to the field strength and is why much effort has been invested in building high field scanners in NMR and MRI.

In MRI, where a large group, or system, of spins is almost exclusively measured, it is possible to use classical models, rather than quantum mechanical models, to describe the motion of the net magnetisation vector, which provides a greater intuitive understanding. In these models, the net magnetisation vector can be thought of as a magnetic dipole.

2.1.2 The Bloch Equation

The Bloch equation, named after Felix Bloch, who with his colleagues first proposed its use,[4,5] is an empirical classical model that describes the motion of the net magnetisation vector belonging to a system of spins:

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B} + \frac{1}{T_1}\begin{pmatrix} 0 \\ 0 \\ M_0 - M_z \end{pmatrix} - \frac{1}{T_2}\begin{pmatrix} M_x \\ M_y \\ 0 \end{pmatrix}, \tag{2.5}$$

where '×' is the cross-product, $\vec{B}$ is the net magnetic field, which by convention is assumed to be approximately aligned with $\hat{z}$, γ is the gyromagnetic ratio introduced above, M_0 is the equilibrium magnetisation magnitude, M_x, M_y, and M_z are the components of $\vec{M}$ in the $\hat{x}$, $\hat{y}$, and $\hat{z}$ directions in the laboratory frame of reference, T_1 is the spin-lattice (or longitudinal) relaxation time, and T_2 is the spin-spin (or transverse) relaxation time.

If we first ignore the relaxation terms, the Bloch equation demonstrates that a static magnetic field, $\vec{B} = B_0\hat{z}$ (e.g. the static magnetic field in an MRI scanner), exerts a torque on the magnetisation vector, $\vec{M}$, causing it to precess around $\hat{z}$ with angular frequency equal to the Larmor frequency, $\omega_0 = \gamma B_0$. The rate of precession is therefore proportional to the energy difference between the two quantum energy states. The Bloch equation in vector form for this case is:

$$\begin{pmatrix} \frac{dM_x}{dt} \\ \frac{dM_y}{dt} \\ \frac{dM_z}{dt} \end{pmatrix} = \gamma \begin{pmatrix} B_0 M_y \\ -B_0 M_x \\ 0 \end{pmatrix}, \tag{2.6}$$

the solution to which is given by:

$$M_{xy}(t) = M_{xy}(0) e^{-i\omega_0 t}, \tag{2.7}$$

where $M_{xy}(t) = M_x(t) + \ iM_y(t)$.

2.1.3 Excitation and Signal Detection

If the magnetisation is at thermal equilibrium, aligned with $\vec{B}_0$, then according to Equation 2.6, the magnetisation vector will not move because $M_x = 0$ and $M_y = 0$. A signal can only be detected if there is a non-zero transverse component of the magnetisation vector. Any non-zero component of magnetisation in the transverse plane will precess at the Larmor frequency and will be measurable by inducing an electromotive force in a coil of wire (the receive coil) via Faraday's law of induction. The voltage in the coil is recorded and is proportional to the amount of magnetisation in the object being imaged.

If a 2nd magnetic field, $\vec{B}_1(t) = B_1(t)\cos(\omega_{RF}t)\,\hat{x} - B_1\,(t)\sin(\omega_{RF}t)\,\hat{y}$, is introduced, which is circularly polarised, is perpendicular to $\vec{B}_0$, and has a carrier frequency of ω_{RF}, then the Bloch equation (ignoring relaxation) becomes:

$$\frac{d\vec{M}}{dt} = \gamma \vec{M} \times \left(\vec{B}_0 + \vec{B}_1(t)\right). \tag{2.8}$$

Note, the oscillating $\vec{B}_1$ magnetic field is often called an 'RF pulse' because, for standard field strengths, the carrier frequency is typically within the radio-frequency range of the electromagnetic spectrum. Other terms include 'excitation pulse', 'inversion pulse', and ‘refocussing pulse’ depending on the angle that the magnetisation is moved by and the purpose of the pulse (see below).

So far, we have viewed $\vec{M}$ from the laboratory frame of reference. It is often convenient to move to a rotating frame of reference which rotates at the Larmor frequency. In this case, the static field solution simplifies to $M'_{xy}(t) = 0$ (where the prime symbol denotes the rotating frame of reference), effectively freezing the observed precession due to $\vec{B}_0$. If $\omega_{RF} = \omega_0$ and $\vec{M}'(0) = M\hat{z}$, Equation 2.8 simplifies to:

$$\frac{d\vec{M}\prime}{dt} = \gamma\vec{M}' \times B_1\hat{x}' = \gamma\begin{pmatrix} 0 \\ M_z B_1 \\ -M_y B_1 \end{pmatrix}, \qquad (2.9)$$

which demonstrates that at resonance (when $\omega_{RF} = \omega_0$), the magnetic vector, $\vec{M}$, will precess around $\vec{B}_1$ with frequency $\omega_1 = \gamma B_1$. In the above example where the magnetisation is initially aligned with $\hat{z}$ and $\vec{B}_1$ is applied along $\hat{x}'$, the magnetisation will rotate in the $\hat{y}'$-$\hat{z}$ plane (note $\hat{z} = \hat{z}'$), moving the magnetisation into the transverse ($\hat{x}'$-$\hat{y}'$) plane. Note that B_1 can be applied in any arbitrary direction perpendicular to $\hat{z}$.

The angle, α, by which the magnetisation vector is rotated by the applied RF pulse (the flip-angle) is equal to:

$$\alpha = \gamma \int B_1 dt. \qquad (2.10)$$

Any arbitrary flip-angle can be achieved by modulation of the magnitude of $\vec{B}_1$and the duration that $\vec{B}_1$ is applied for (the pulse duration).

If the RF pulse is applied off-resonance ($\omega_{RF} \neq \omega_0$), then the RF pulse will appear to precess in the rotating frame of reference. This will cause $\vec{M}$ to rotate around $\hat{z}$ as well as $\hat{x}'$. As $\vec{M}$ is rotated away from $\hat{z}$, the axis that $\vec{M}$ is being rotated around also changes, meaning that $\vec{M}$ never moves far from $\hat{z}$. The further off-resonance the RF pulse is, the less perturbed $\vec{M}$ will be.

2.1.4 Relaxation

So far, we have ignored the longitudinal and transverse relaxation terms. The longitudinal relaxation term, governed by spin-lattice interactions, describes the recovery of $\vec{M}$ along the $\hat{z}$ axis. If $\vec{M}$ is perturbed from thermal equilibrium, described by the Boltzmann distribution above, the spin system will gradually return to this equilibrium position. Assuming that $\vec{B} = B_0\hat{z}$, the regrowth of $\vec{M}$ along the $\hat{z}$ axis is classically described by an exponential recovery equation of the form:

$$M_z(t) = M_0 + (M_z(0) - M_0) \cdot e^{\frac{-t}{T_1}}. \tag{2.11}$$

The mechanisms for longitudinal relaxation (also called T_1-relaxation) are the random motion of the magnetic moments of neighbouring spins (the lattice). Some of this motion can be at or near the Larmor frequency, thus exerting a non-negligible influence on the net magnetisation. In the quantum mechanical sense, this causes transitions between the spin states.

The transverse relaxation time, governed by spin-spin as well as spin-lattice interactions, describes the decay process of $\vec{M}$ in the transverse plane. For a spin system, this can also be described using a classical empirical model, where the general solution to the Bloch equation is:

$$M'_{xy}(t) = M'_{xy}(0)e^{\frac{-t}{T_2}}. \tag{2.12}$$

The random motion of neighbouring spins causes molecular-level magnetic field variations. This means that over very small spatial scales, spins have different Larmor frequencies and so precess at slightly different rates. This leads to a loss of coherence in the phase of the spin system causing an increasing amount of the net magnetisation to be cancelled out. T_2 relaxation can occur without T_1 relaxation, but T_1 processes also lead to T_2 relaxation; therefore, $T_2 \leq T_1$.

The transverse relaxation that is typically observed is shorter than T_2, and is referred to as T_2^* relaxation. The manifestation of T_2^* relaxation is similar to T_2 relaxation, but whereas the latter is caused by molecular-scale field perturbations, the former is caused by field variations on a larger scale. The source of these are variations in susceptibility of the object being imaged, which lead to spins in different regions experiencing different field strengths. This causes different precession frequencies in the object which leads to the signal dephasing. Importantly, T_2^* effects can be reversed using a 180° refocussing pulse, but T_2 relaxation cannot.

2.1.5 Slice-Selective Excitation

The source of signal in MRI has been discussed, but we have not yet discussed how the magnetisation is spatially localised in order to create an image. A common first step in this process is to use a technique called slice-selective excitation. This typically involves creating a linear magnetic field gradient, $\vec{G}$, along the axis perpendicular to the desired slice. Note, whichever direction this gradient is applied along, it is designed to only modulate the $\hat{z}$ component of the B_0 field. For instance, if a transverse slice is desired, a gradient is applied along the $\hat{z}$ axis such that the effective static field, $B_0\hat{z}$, is decreased for $z < 0$ and increased for $z > 0$, with no change at $z = 0$. This causes spins to precess at different frequencies based on their position along the $\hat{z}$ axis. Mathematically, the magnetic field is given by:

$$\vec{B} = \vec{B}_0 + \vec{G} \cdot \vec{r} = B_0\hat{z} + (G_z \cdot z)\hat{z}, \tag{2.13}$$

where $\vec{r} = x\hat{x} + y\hat{y} + z\hat{z}$, $\vec{G} = \vec{\nabla} B_z$, and so $G_z = dB_z/dz$.

The next step is to play an RF pulse concurrently with the slice-select gradient. This RF pulse should only excite a specific range of frequencies, meaning that it has a finite bandwidth ($BW = \Delta\omega$). Spins precessing with frequencies outside this frequency

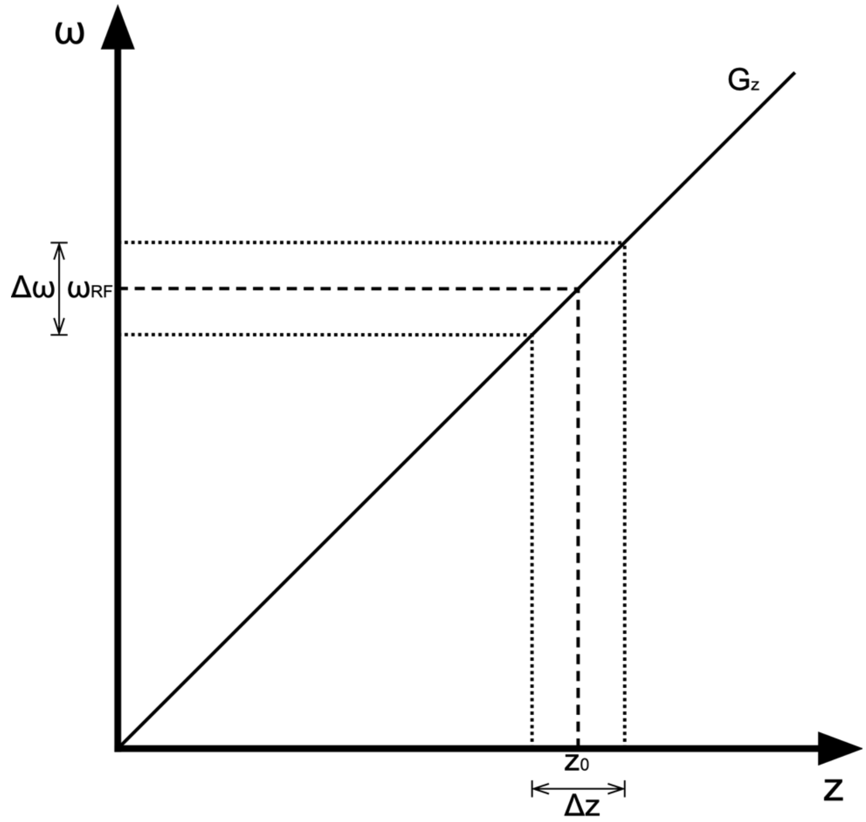

Figure 2.1: This figure demonstrates the effect of slice-selective excitation. When a bandwidth-limited (Δω) pulse is played out concurrently with a linear gradient field along $\hat{z}$ (Gz), the excitation of specific frequencies corresponds to exciting a specific spatial region (Δz).

range are off-resonance to the RF pulse and so are not significantly perturbed. The RF pulse has a base frequency (or carrier frequency), ω_{RF}, equal to the desired central frequency (the centre of the slice), which if at isocentre ($z = 0$) is simply $\omega = \gamma B_0$. The effect of playing a bandwidth-limited pulse during a slice-select gradient is demonstrated in Figure 2.1. A combination of the gradient strength (G_z) and RF pulse bandwidth ($\Delta\omega$) excite a slice with width Δz.

To ensure that the flip-angle across the frequency range (or slice) is approximately uniform, the pulse amplitude (the pulse envelope) is modulated over time to achieve the desired frequency response. A commonly used shaping function is the sinc pulse. For an infinitely long pulse duration, the frequency profile of a sinc pulse is a boxcar function, producing a perfect rectangular slice profile. However, in practice pulse durations are

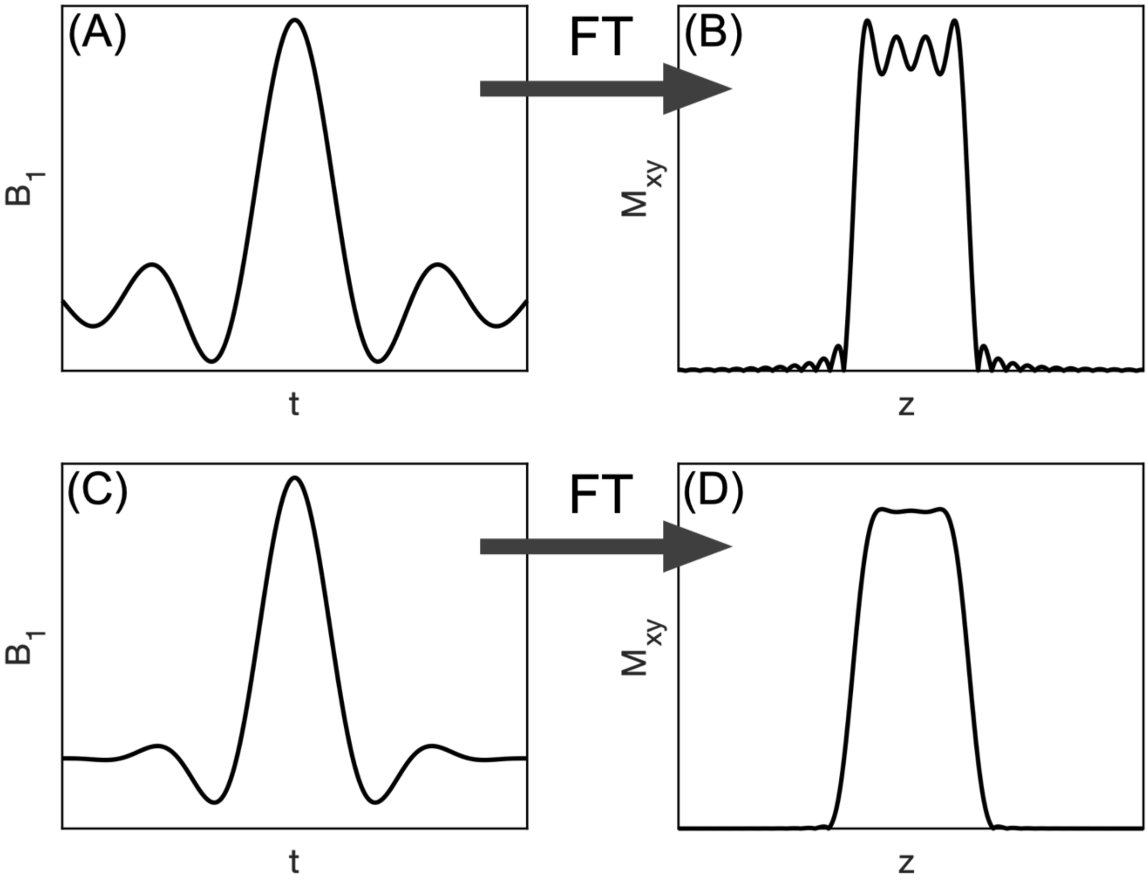

Figure 2.2: The slice profile of a sinc pulse with TBWP = 8. (A) A truncated since pulse, which modulates the amplitude of B_1 over time, (B) the Fourier transform of the truncated sinc pulse in (A), (C) the same truncated sinc pulse in (A) but which has been windowed with a Hanning window, (D) the Fourier transform of (C).

typically on the order of a few milliseconds, which leads to an imperfect slice profile. To reduce slice imperfections resulting from the finite pulse duration, the truncated sinc pulse is typically windowed. For small flip-angles (typically $\alpha < 90°$), the frequency profile of a pulse is well approximated by the Fourier transform of the pulse envelope (for large flip-angles, the slice profile will deviate from this approximation due to non-linearities in the Bloch equations). Note, for a sinc pulse, the time-bandwidth-product (TBWP) is equal to the number of zero-crossings, where TBWP $= T_p \cdot \Delta\omega$ (T_p is the pulse duration). A higher TBWP leads to an improved slice profile. The sinc pulse slice profile with and without a Hanning window is demonstrated in Figure 2.2. It is evident that windowing a truncated sinc pulse results in a much more uniform flip-angle across the slice and suppresses the side lobes.

2.1.6 Spatial-Encoding and *k*-Space

Now that the signal has been localised to a single finite slice of the object being imaged, further spatial encoding is required to map the spatial distribution of magnetisation within the slice. This is most commonly performed by Fourier-encoding the signal, using two linear gradient fields along the $\hat{x}$ and $\hat{y}$ axes, G_x and G_y (for a transverse slice, as above). These gradient fields also modulate B_0 along the $\hat{z}$ axis, but as a function of position along $\hat{x}$ and $\hat{y}$.

Ignoring relaxation effects once again, the transverse signal detected by a receive coil with uniform sensitivity across the object can be described as the complex sum of the magnetisation within the excited volume:

$$s \propto \iiint \rho(\vec{r}) \cdot e^{-i\phi(\vec{r},t)} d\vec{r}, \tag{2.14}$$

where it is assumed that the volume has been excited uniformly, $\rho(\vec{r})$ is the density of ^{1}H at the position $\vec{r}$, and $\phi(\vec{r}, t)$ is the spatial and time-dependent phase of the spins. Note that this equation describes the detected signal after demodulation with the Larmor frequency - essentially viewing the signal in the rotating frame of reference.

The phase, ϕ, is the time integral of the local precession frequency in the rotating frame of reference which is modulated by the gradient fields:

$$\phi(\vec{r}, t) = \int \gamma \vec{G}(t) \cdot \vec{r} dt. \tag{2.15}$$

If we set $\vec{k}(t) = \int \gamma \vec{G}(t)\, dt$, then Equation 2.14 becomes

$$s \propto \iiint \rho(\vec{r}) \cdot e^{-i\vec{k}(t)\cdot\vec{r}} d\vec{r}, \tag{2.16}$$

which is the Fourier transform of the spin density, $\rho(\vec{r})$. Once this Fourier space (also known as *k*-space) has been sufficiently sampled, an image of the spin density can be recovered by performing an inverse discrete Fourier transform of this spatial frequency data, though the fast Fourier transform is most often used due to the large reduction in the number of operations (($O(N\log(N))$ rather than $O(N^2)$).

G_x and G_y are used to navigate this 2D k-space since their time integral governs the spatial frequency pattern of phase, k, imposed upon the magnetisation, while the induced voltage is sampled at discrete intervals. It is common to sequentially acquire data in one k-space direction, such as $\hat{x}$, by keeping G_x at a constant magnitude. The remaining direction, $\hat{y}$, is navigated by applying G_y in short blips. Due to the constant gradient used along $\hat{x}$, this process is often referred to as frequency encoding, while the signal modulation along $\hat{y}$ is called phase encoding. Multiple slices can be acquired to create a 3D image of an object, or k-space sampling can be extended to three dimensions by also sampling along the third direction, k_z.

2.1.7 Spoiling

Spoiling is used to dephase unwanted transverse magnetisation, thus cancelling its signal. This is typically done in three ways. The simplest method is to wait for a period of time until the majority of the transverse magnetisation has decayed due to T_2 relaxation. However, this is not a time-efficient method. Alternatively, gradient spoiling uses a large gradient area to dephase the magnetisation along the direction of the gradient, cancelling the signal within a voxel. The last method is RF spoiling. This technique varies the RF phase according to a predetermined pattern, typically a quadratic scheme,[6] cancelling transverse magnetisation and preventing it from being refocussed. Usually, a combination of gradient and RF spoiling is used with cycled gradient amplitudes.

2.1.8 Effective Magnetic Field

It is sometimes convenient to define an effective magnetic field in the rotating frame of reference, $\vec{B}_{\text{eff}}$, which incorporates all of the separate magnetic fields: the RF field, $\vec{B}_1$, the gradient fields, $\vec{B}_G$, and other magnetic fields that act along $\hat{z}$, which are collectively referred to as off-resonance fields, $\vec{B}_{OR}$:

$$\vec{B}_{\text{eff}} = \vec{B}_1 + \vec{B}_G + \vec{B}_{OR} \tag{2.17}$$

where

$$\vec{B}_1 = B_1 \begin{pmatrix} \cos(\phi_1) \\ \sin(\phi_1) \\ 0 \end{pmatrix}, \tag{2.18}$$

$$\vec{B}_G = \begin{pmatrix} 0 \\ 0 \\ \vec{G} \cdot \vec{r} \end{pmatrix}, \tag{2.19}$$

$$\vec{B}_{OR} = \begin{pmatrix} 0 \\ 0 \\ \omega_{OR}/\gamma \end{pmatrix}, \tag{2.20}$$

ϕ_1 is the phase of the RF pulse and ω_{OR} is the off-resonance precession frequency. In the rotating frame of reference, the Bloch equation becomes:

$$\frac{d\vec{M}'}{dt} = \gamma \vec{M}' \times \vec{B}_{\text{eff}} + \frac{1}{T_1}\begin{pmatrix} 0 \\ 0 \\ M_0 - M_z \end{pmatrix} - \frac{1}{T_2}\begin{pmatrix} M'_x \\ M'_y \\ 0 \end{pmatrix}, \tag{2.21}$$

which describes the precession of spin about the effective magnetic field, $\vec{B}_{\text{eff}}$.

2.1.9 Numerical Bloch Simulations

The movement of $\vec{M}$ in the rotating frame of reference by the effects of excitation, relaxation, gradients, and off-resonance can be simulated in various ways. Two methods are used in this thesis, which are described here.

The first is the simplest and is used several times in this work. If we consider only: instantaneous RF pulses, individual spin isochromats (and so a homogeneous $\vec{B}_0$ field), perfect spoiling of residual transverse signal (M_{xy} is set to zero after each measurement), and ideal flip-angles, then the motion of $\vec{M}$ can be described analytically for many situations in which these assumptions are valid. For example, if a 60° single excitation pulse acts on an isochromat in thermal equilibrium, the resulting transverse signal immediately after the instantaneous RF pulse can be expressed as $M_{xy} = M_0 \sin(60°)$ and the remaining longitudinal magnetisation as $M_z = M_0 \cos(60°)$. The time evolution of M_{xy} after the pulse is simply $M_{xy}(t) = M_0 \sin(60°)\, e^{-t/T_2}$. Essentially, this technique limits us to situations for which there are analytical solutions to the Bloch equations. A weakness of this method is that more complicated interactions cannot be accurately described.

The second method that is used in this thesis is more complicated and is referred to as the '*hard pulse approximation*'. The Bloch equation can be formulated in such a way that the influence of an RF pulse, gradient field, and off-resonance field (collectively $\vec{B}_{\text{eff}}$) and relaxation effects can be written mathematically as rotation matrices which describe specific rotations of $\vec{M}$. Application of an RF pulse can then be simulated by approximating the time-varying $\vec{B}_1$ field as a series of N small flip-angle, constant B_1 (hard), RF pulses. These discrete hard pulses are combined with $\vec{B}_G$ and $\vec{B}_{OR}$ to form a discrete time-varying $\vec{B}_{\text{eff}}$ field. After application of the i^{th} discretised $\vec{B}_{\text{eff}}$ pulse, relaxation effects are simulated for a short period of time, dT, where $dT = T_p/N$ (T_p being the duration of the RF pulse). It is important to combine the effects of $\vec{B}_{\text{eff}}$, because these rotation matrices do not commute. This technique is used in Appendix D for simulating the slice profile of adiabatic inversion pulses. The implementation used in this thesis was originally written by Brian Hargreaves.[7]

2.2 Readout Methods

This section gives a brief overview of several methods for sampling *k*-space which are used within this thesis.

2.2.1 Spoiled Gradient Recalled (SPGR)

One of the simplest methods for acquiring data in *k*-space is the single-line FLASH sequence.[8] With the addition of gradient and RF spoiling, this sequence is referred to as an SPGR readout in this thesis. Many of the features in this sequence are used in more complicated readout methods, and so the concepts described here are useful for understanding the content within this thesis. An SPGR readout with a radial trajectory is also used in Chapter 6 for acquiring angiographic images.

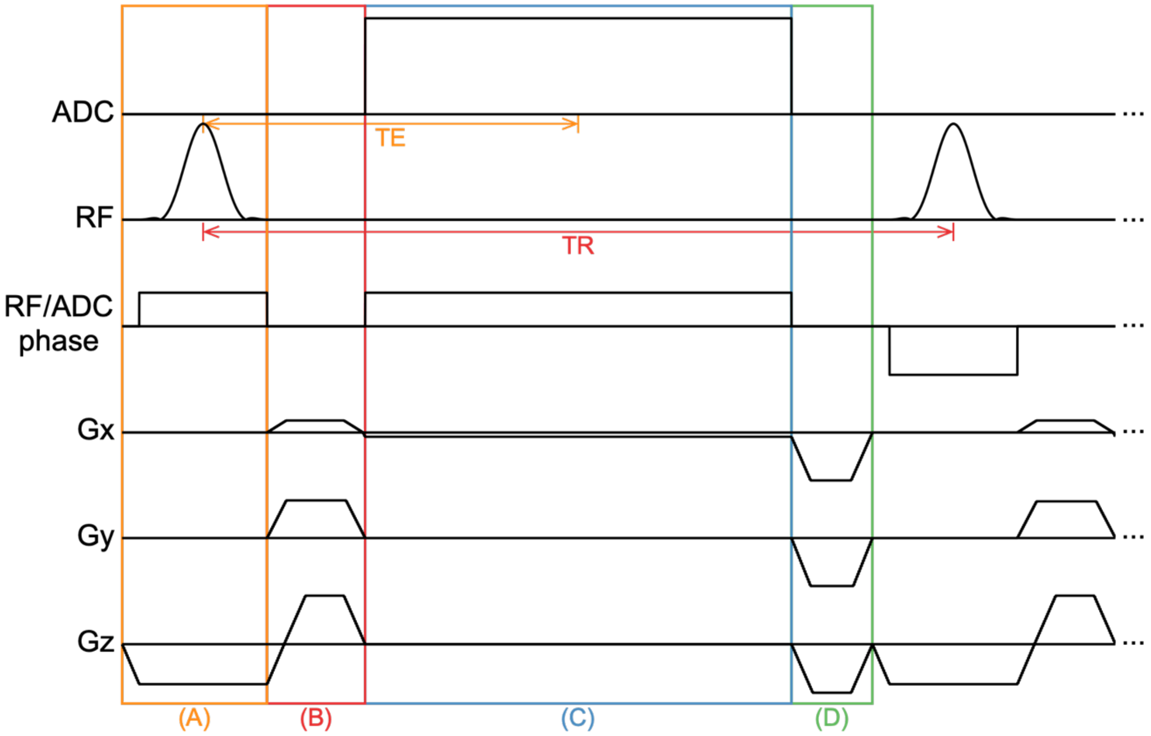

Figure 2.3: Sequence diagram for an SPGR readout with a Cartesian trajectory. The time between excitation pulses is referred to as the repetition time (TR). The echo time (TE) is the point where the centre of the readout line is.

2.2.1.1 Cartesian Trajectory

A portion of the SPGR readout sequence with a Cartesian trajectory is shown in Figure 2.3. In section (A), the windowed-sinc excitation pulse, with a varied phase, is played out with a gradient on the G_z axis for slice-selection; (B) the phase rewinding gradient on G_z, phase-encode gradient blip on G_y, and read-out prephasing gradient G_x are played out concurrently; (C) the readout gradient on G_x with signal sampling (analogue-to-digital converter: ADC) are performed; (D) a phase rewinding gradient on G_y and spoiling gradients on G_x and G_y are played out simultaneously. After (D), the next k-space line is read out after another excitation pulse. With both gradient and RF spoiling, after a small number of TRs, the transverse signal can be accurately predicted by assuming perfect spoiling.

Because the TR is usually on the order of 5-15 ms, flip-angles $\ll 90°$ are used in order to maximise the steady-state signal. This may seem counter-intuitive, since for a single excitation the acquired signal is maximised with a flip-angle of 90°. However, for a series of identical flip-angles (referred to as constant flip-angles), each RF pulse attenuates

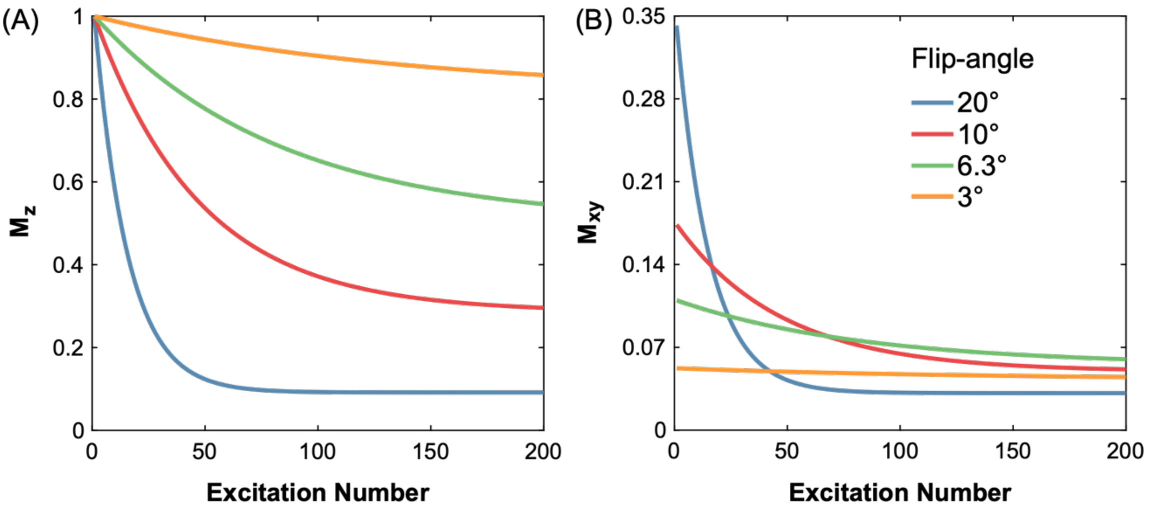

Figure 2.4: The effect of the SPGR excitation pulses on the longitudinal (A) and transverse (B) magnetisation for a range of flip-angles. Instantaneous RF pulses and perfect spoiling were assumed. Parameters: TR 10 ms, T_1 1.65 s, The Ernst angle is 6.3° for these parameters.

the longitudinal magnetisation by a factor of $\cos(\alpha)$, reducing the available magnetisation for the next TR. The optimal constant flip-angle can be found using the Ernst angle:

$$\alpha_{opt} = \cos^{-1}\left(e^{-\frac{TR}{T_1}}\right). \tag{2.22}$$

For an object with T_1 = 1.65 s (a typical T_1 value for arterial blood at 3T)[9] and a TR = 10 ms, the optimal flip-angle is 6.3°. The effect of the SGPR flip-angle on the longitudinal and transverse magnetisation is shown in Figure 2.4. Flip-angles larger than the Ernst angle result in higher signal at early excitations, but a lower steady-state signal. Flip-angles smaller than the Ernst angle do not make efficient use of the available magnetisation. In

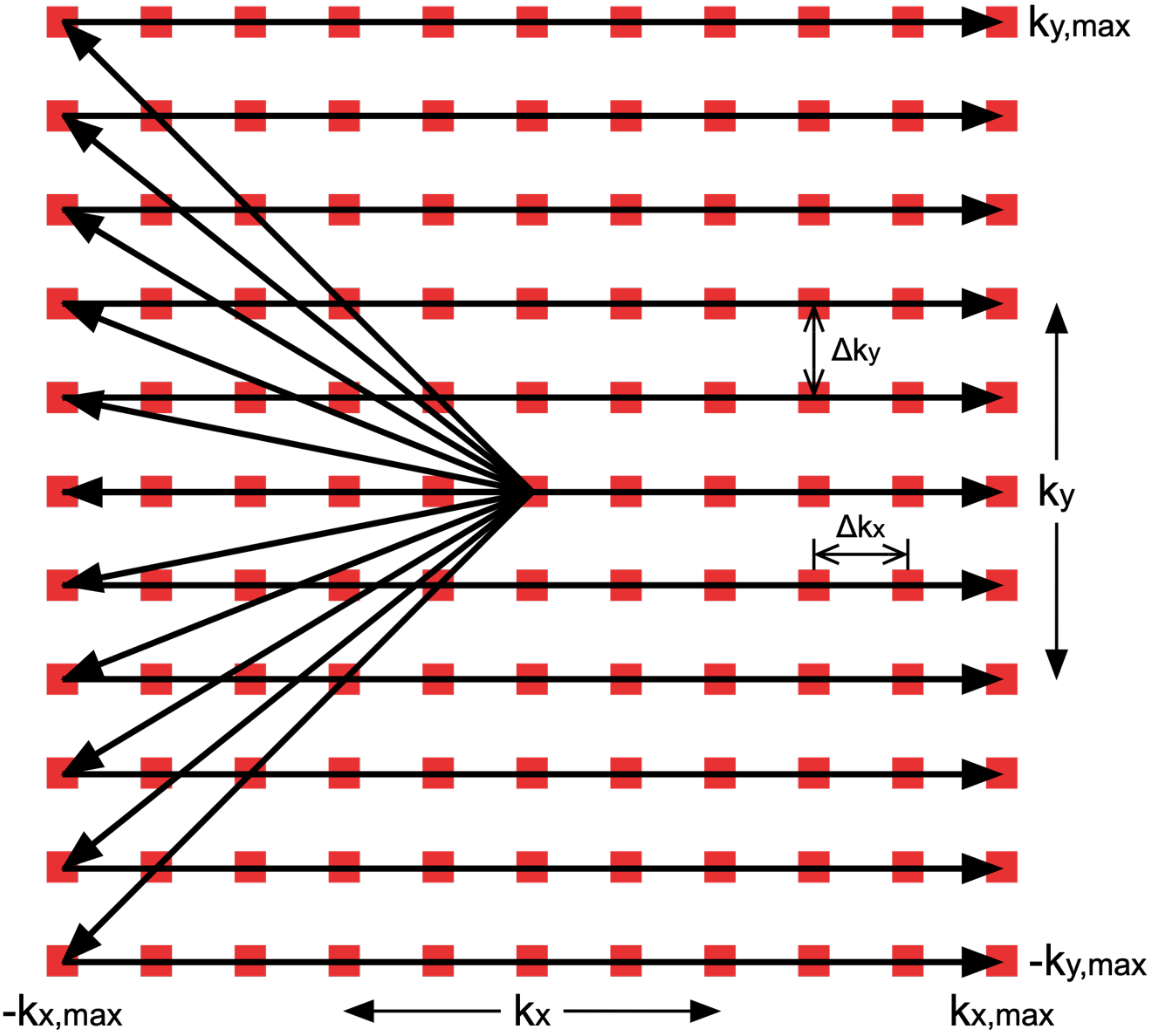

Figure 2.5: The 2D k-space trajectory corresponding to the sections (B) and (C) of the Cartesian SPGR sequence in Figure 2.3. $k_{x,max}$ and $k_{y,max}$ represent the highest spatial frequencies sampled in k_x and k_y. Δk_x and Δk_y are the step sizes between the k-space samples.

Chapter 6, where an SPGR readout is used and the signal of interest, the ASL signal, is a transient signal decaying with T_1, this equation no longer holds.

The 2D *k*-space trajectory traversed by the sequence in Figure 2.3 is shown in Figure 2.5. Each line that is traversed in this 2D *k*-space corresponds to the action of the G_x and G_y gradients in sections (B) and (C) of Figure 2.3, though the amplitude of G_y is varied for each TR to navigate to the beginning of each line. Two key *k*-space concepts are also highlighted: (1) the highest spatial frequencies sampled along k_x and k_y represent the finest detail that can be reconstructed in the imaged object and (2) the step size between *k*-space samples in each direction is inversely proportional to the image field-of-view (FOV) in each corresponding image direction.

2.2.1.2 Radial Trajectory

An alternative *k*-space trajectory to a Cartesian sampling pattern is the radial trajectory. This trajectory was used in the first demonstration of MRI, predating the *k*-space formalism, and was originally used by Paul Lauterbur in 1973 to produce a series of projection images which were then combined via the back-projection algorithm (commonly used for X-ray CT imaging).[10] However, the non-Cartesian *k*-space samples can also be reconstructed with an FFT (which requires equally spaced Cartesian samples) if the data is first resampled onto a Cartesian grid. This is usually done by a method called 'gridding'.[11]

The *k*-space trajectory traversed by a fully sampled radial readout is shown in Figure 2.6 on the same grid as in Figure 2.5. This trajectory can be achieved by rotating the readout line by an angle of π/N radians for each acquisition, where N is the total number of radial spokes.

A downside of radial readouts is that they require a $\pi/2$ (~57%) increase in the number of k-space lines (and a corresponding increase in the scan time) compared to an equivalent Cartesian acquisition to fully sample *k*-space. This is due to the requirement of maintaining a maximum distance of Δk between samples at the less densely sampled edge of *k*-space. However, one benefit of more densely sampling the centre of *k*-space (the low spatially frequencies) is that radial acquisitions are more robust to motion.[12]

Further issues with radial readouts are the more severe artefacts resulting from B_0 inhomogeneities and gradient nonlinearities. In Cartesian imaging, these phenomena only affect the location and amplitude of the point-spread function (PSF), causing geometric and intensity distortions. In radial imaging, the shape of the PSF is also distorted, resulting in blurring which is much more difficult to correct in post-processing. However, B_0 correction can be incorporated into the reconstruction if it is known, these effects are typically small for SPGR readouts, and modern scanners have much greater B_0 field

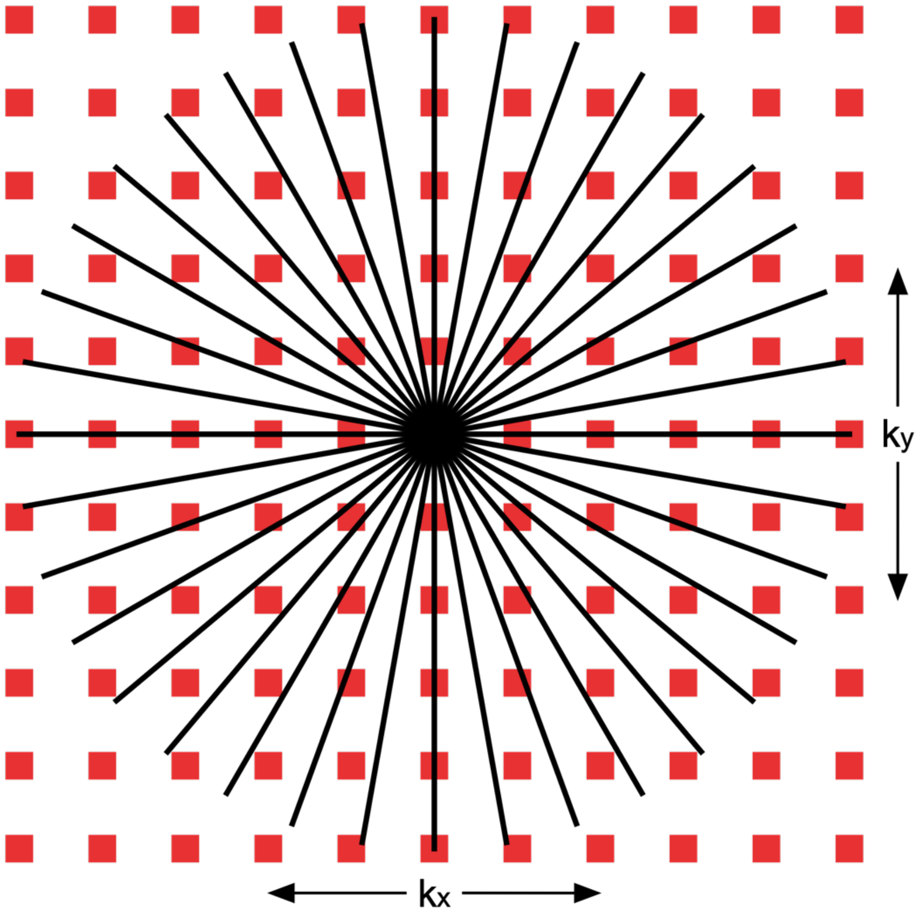

Figure 2.6: The 2D k-space trajectory of a radial readout overlaid on the same Cartesian grid as in Figure 2.5.

homogeneity and gradient linearity than when these artefacts were first identified, making the resulting artefacts much less severe.

2.2.2 2D Echo Planar Imaging (2D EPI)

A much faster Cartesian method for covering 2D *k*-space is with the single-shot EPI readout, first proposed by Peter Mansfield in 1977.[13] Unlike the single-line SPGR sequences described above, the entire 2D *k*-space can be acquired after a single acquisition. This is achieved by applying a phase-encode blip after each readout line is acquired to move directly to the next *k*-space line. This next line is then traversed by using a gradient with equal area but negative amplitude to the preceding readout line. The resulting *k*-space trajectory is shown in Figure 2.7.

Issues with this readout include the following. The echo time (defined as the time taken to reach the central *k*-space sample) is typically prolonged compared to the SPGR

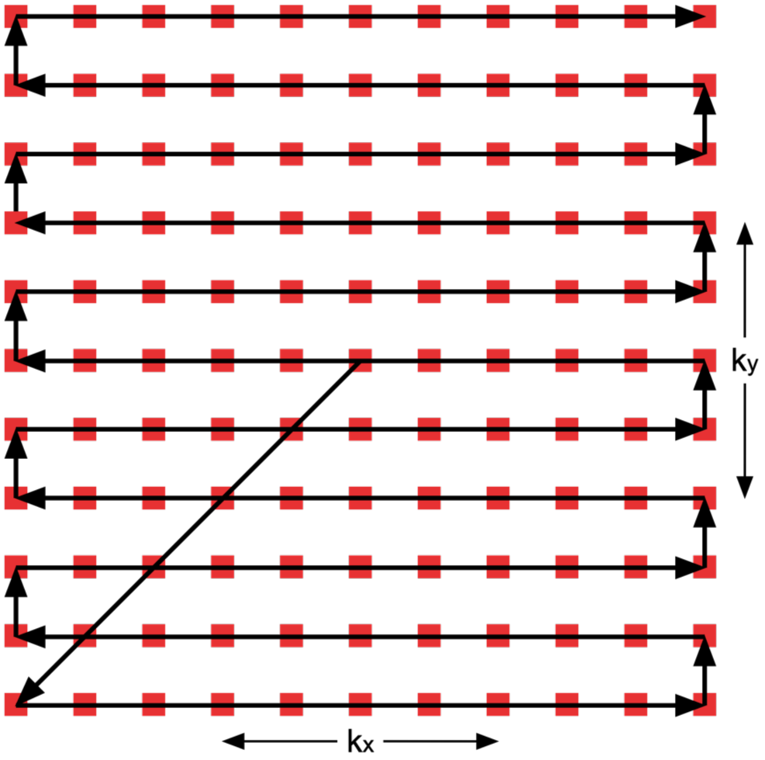

Figure 2.7: EPI k-space trajectory. The same samples are acquired as in the Cartesian SPGR readout, but much time is saved by continually moving through k-space after a single excitation.

readout, resulting in greater signal decay and lower SNR images. The T_2^* signal decay across k-space lines weights the signal intensity for different spatial frequencies, causing blurring in image space in the phase-encode direction. T_2^* decay also causes severe signal loss in areas of off-resonance, such as in the sinuses, due to intravoxel dephasing. Discrete ghosts in image space can form due to timing differences between odd and even lines of k-space, however, these artefacts are routinely minimised using a pre-scan of the central readout line to calibrate for this timing offset. Due to the low effective bandwidth in the phase-encode direction (the speed with which k-space data is read out), large chemical-shift artefacts can occur. Most prominently, the signal from fat around the skull is shifted along the direction of the phase-encode gradient but is often minimised by using a spectrally-selective pulse to saturate the fat signal immediately before the EPI excitation pulse. Large geometric distortions can also occur due to the low bandwidth in the phase-encode direction, but these can be largely corrected for in post-processing if a B_0 field map is also available.

Extremely fast 2D imaging can, therefore, be achieved with an EPI readout, but at the cost of increased artefacts. Much effort has been spent on robustly minimising these artefacts, however, making EPI a valuable and commonly used tool in neuroimaging. Furthermore, a large advantage of fast single-shot imaging over the SPGR readout described above is the improved robustness to motion because the entirety of the 2D k-space can be acquired in ~50 ms.

2.2.3 3D Gradient and Spin Echo Imaging (3D GRASE)

MRI is not restricted to slice-selection and 2D k-space sampling. It is also possible to excite a thick slab or use non-selective excitation and then spatially encode the signal along $\hat{z}$. This can be done with any of the previously described readouts: the Cartesian SPGR readout can be extended by also sampling in k_z; the radial SPGR readout by either the stack of stars approach[14] or by a true 3D radial trajectory;[15] and the EPI readout can be extended to 3D by use of echo-volumnar-imaging[16] or by 3D EPI.[17]

An alternative 3D readout strategy is 3D GRASE.[18,19] This readout technique uses an EPI trajectory to acquire each plane of k_z before an RF refocusing pulse is used to reverse the effects of T_2^* relaxation before moving to the next k_z plane which is acquired with another EPI trajectory. Its name derived from the fact that both gradient and spin echoes are formed during the readout.

A section of the 3D GRASE readout sequence used in this work is shown in Figure 2.8. The fat saturation and EPI phase correction scans are also used in the EPI readout. The TE for a 3D GRASE sequence is typically longer than for a gradient-recalled-echo (GRE) EPI readout due to the spin echo. The refocussing pulses are played out 90° out of phase to the excitation pulse in order to rephase spins which have dephased due to T_2^* effects; this is commonly known as a CPMG pulse train, named after Carr, Purcell, Meiboom, and Gill.[20] At the end of each EPI section, the gradient zeroth order moments on all axes are fully rewound in order to satisfy the conditions of the CPMG pulse train. Crusher gradients

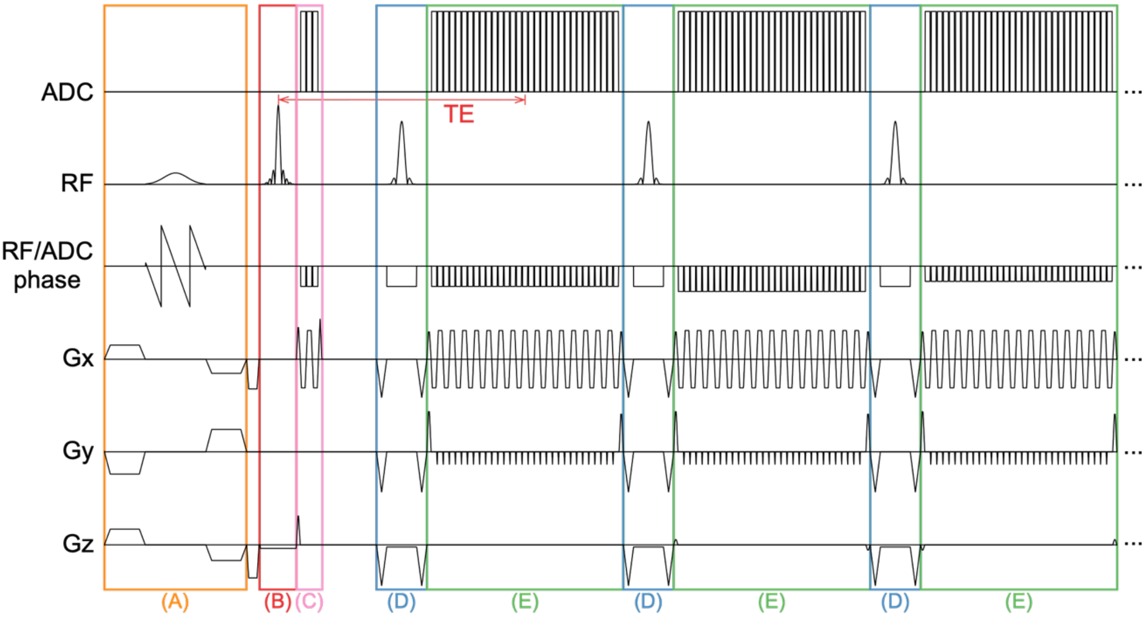

Figure 2.8: The 3D GRASE readout sequence with three k_z partitions. (A) Spectrally-selective fat saturation; (B) slab-selective excitation; (C) EPI phase correction scan; (D) refocussing pulses with associated crusher gradients; (E) EPI readout trajectory.

are placed before and after each refocussing pulse in order to spoil newly excited signal (the free-induction-decay: FID) and secondary spin echoes which do not originate from the original excitation pulse, while preserving signal from primary and stimulated echoes.

Centric ordering of the k_z partitions is most often used. This is where the central $k_z = 0$ plane is acquired first before acquiring the remaining partitions in the following order: k_z = 1, -1, 2, -2, 3, -3, etc. This minimises the TE and means low spatial frequencies are acquired first (which contain most of the image contrast information) before the signal has decayed significantly due to T_2 relaxation. However, for a single-shot acquisition with a

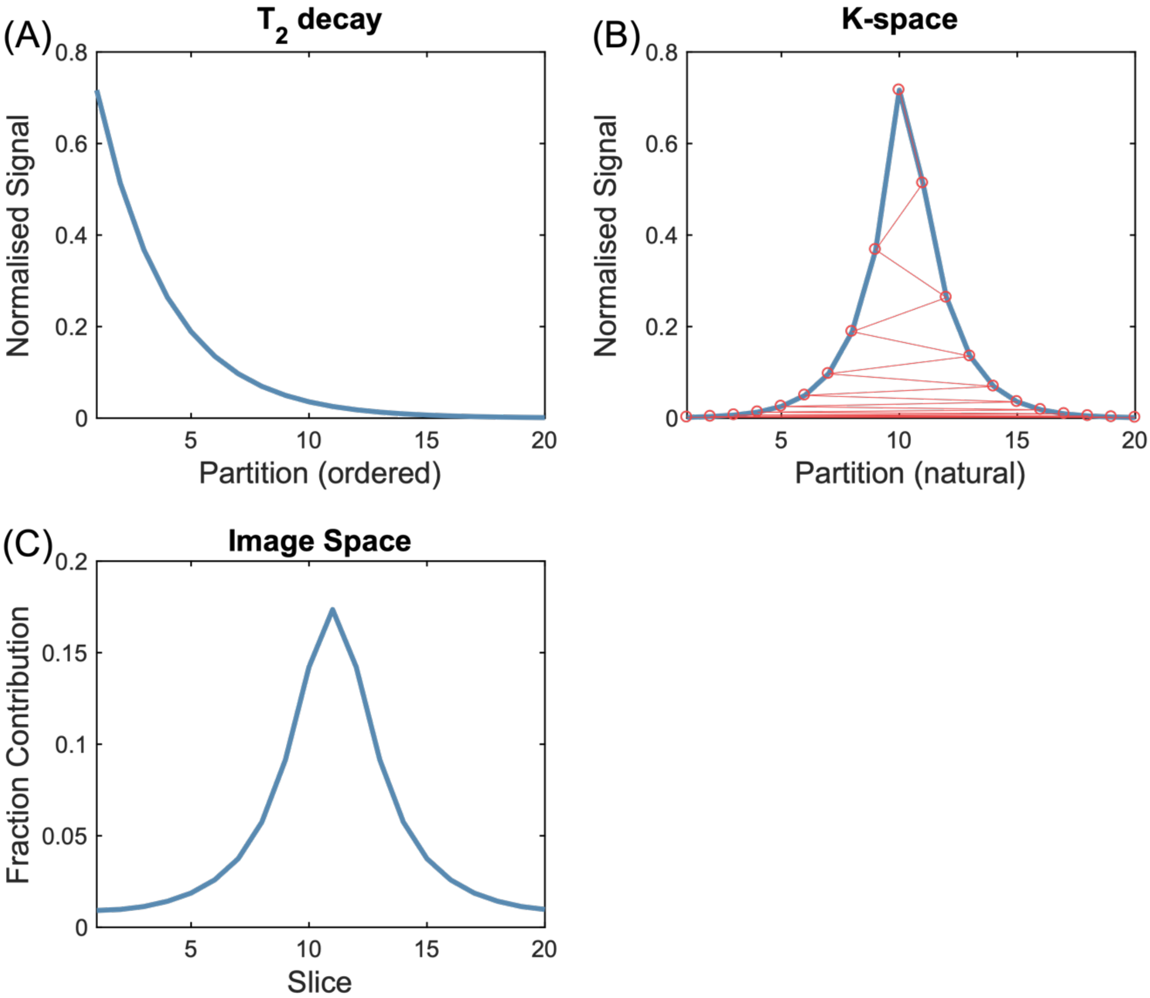

Figure 2.9: The simulated effect of T_2 decay across 3D GRASE partitions on the image space point-spread function. (A) The simulated signal for each partition in the order that they were acquired; (B) the signal for each partition ordered by k_z position with the acquisition order shown in red; (C) the IFFT (the PSF) of the k-space modulation function in (B). The area under the curve in (C) is normalised to equal 1. The simulation only considers the centre of k-space, assumes 180° refocussing pulses, TE = 50 ms, and T_2 = 150 ms.

long echo train length (ETL), this can lead to severe blurring in the $\hat{z}$ direction due to down-weighting of the high spatial frequency information. This effect is demonstrated in Figure 2.9 where the resulting PSF from the T_2 k-space filtering shows severe blurring across slices. Use of refocussing flip-angles < 180°[21] or variable flip-angles[22] have been proposed as methods for reducing T_2 blurring by storing signal in the longitudinal plane which acts to increase the signal at later partitions. Partial Fourier[23] and parallel imaging[24,25] can also be used to reduce the ETL, reducing T_2 induced blurring. However, these methods all come at the cost of reduced SNR, though by varying amounts.

Apart from through-plane T_2 blurring, 3D GRASE readouts also suffer from many of the same artefacts as 2D EPI, but the use of spin echoes does reduce intravoxel dephasing caused by off-resonance fields.

2.3 Adiabatic Inversion

An important class of RF pulses in MRI, particularly for ASL, are pulses which use a process called adiabatic fast passage. An excellent review of the key concepts of adiabatic pulses may be found in references [26] and [27]. The concept of adiabatic inversion is made use of several times throughout this thesis, so a brief introduction is given here.

The key difference between adiabatic pulses and those described above is that rather than only modulating the B_1 amplitude over time, the carrier frequency of the pulse is also modulated. For adiabatic inversion pulses, the pulse frequency is swept from one side of resonance to the other over the course of the pulse, with the centre of the pulse corresponding to $\omega_{RF} = \omega_c$, where ω_c is the central frequency of the slice ($\omega_c = \omega_0$ for non-selective on-resonance pulses or slice-selective pulses at iso-centre).

As described in Section 2.1.3, an off-resonance spin will experience a B_1 field that rotates around $\hat{z}$. Since spins precess around the effective field, it will only be rotated by a very small angle from the $\hat{z}$ axis. With adiabatic inversion pulses, the pulse frequency is swept from one side of ω_c to the other. This means that an on-resonance spin aligned with +$\hat{z}$ will at first experience a B_1 field that is rotating very fast in the Larmor rotating frame of reference, with the effective field being approximately aligned with +$\hat{z}$. As the pulse frequency approaches ω_c, the rotation of the B_1 field will appear to slow down and will begin to rotate the spin into the $\hat{x}$-$\hat{y}$ plane. The spin will pass through the $\hat{x}$-$\hat{y}$ plane when $\omega_{RF} = \omega_c$ and will continue to precess around the effective field towards the -$\hat{z}$ axis as B_1 moves far off-resonance.

Since it is the movement of the effective field that is responsible for inversion, the flip-angle is no longer equal to the integral of the B_1 amplitude over time (Equation 2.10). Care must be taken to ensure the adiabatic condition is met, which states that the effective field should rotate more slowly than the precession of the spins around the effective field. This can be achieved with the use of long pulse durations or high B_1 amplitudes. The key advantage of adiabatic inversion pulses is that as long as the adiabatic condition is satisfied, the flip-angle (or inversion efficiency) is relatively insensitive to variation in B_1 magnitude.

Another key concept is that spins precessing at different resonant frequencies, e.g. due to off-resonance fields, are inverted at different times during the pulse because ω_{RF} will be on-resonance at different times for these spins. Also, as long as the range of frequencies that are swept is large, adiabatic inversion pulses are relatively insensitive to variations in B_0.

2.3.1 Hyperbolic Secant Pulse

If a uniform gradient field is applied at the same time as an adiabatic inversion pulse, then the spin resonances are distributed linearly in space along the gradient axis (assuming a homogenous field before application of the gradient). Gradient fields in combination with adiabatic pulses can, therefore, be used to invert a slab of spins, with the characteristic that spins across the slab will be inverted at different times during the pulse.

The most commonly used spatially selective adiabatic inversion pulse is the hyperbolic secant pulse.[28] This pulse is governed by the following modulation functions:

$$\begin{aligned} B_1(t) &= \quad A_0 \operatorname{sech}(\beta t) \\ \Delta\omega(t) &= \quad -\mu\beta \tanh(\beta t)\text{'} \end{aligned} \tag{2.23}$$

where $\Delta\omega(t)$ is the frequency modulation function, A_0 is the maximum B_1 amplitude, β is the modulation angular frequency in units of $\mathrm{rad} \cdot s^{-1}$, and μ is a dimensionless shaping parameter (specifically it is the ratio of the slab width to the width of the transition zone). These shaping functions are demonstrated in Figure 2.10.

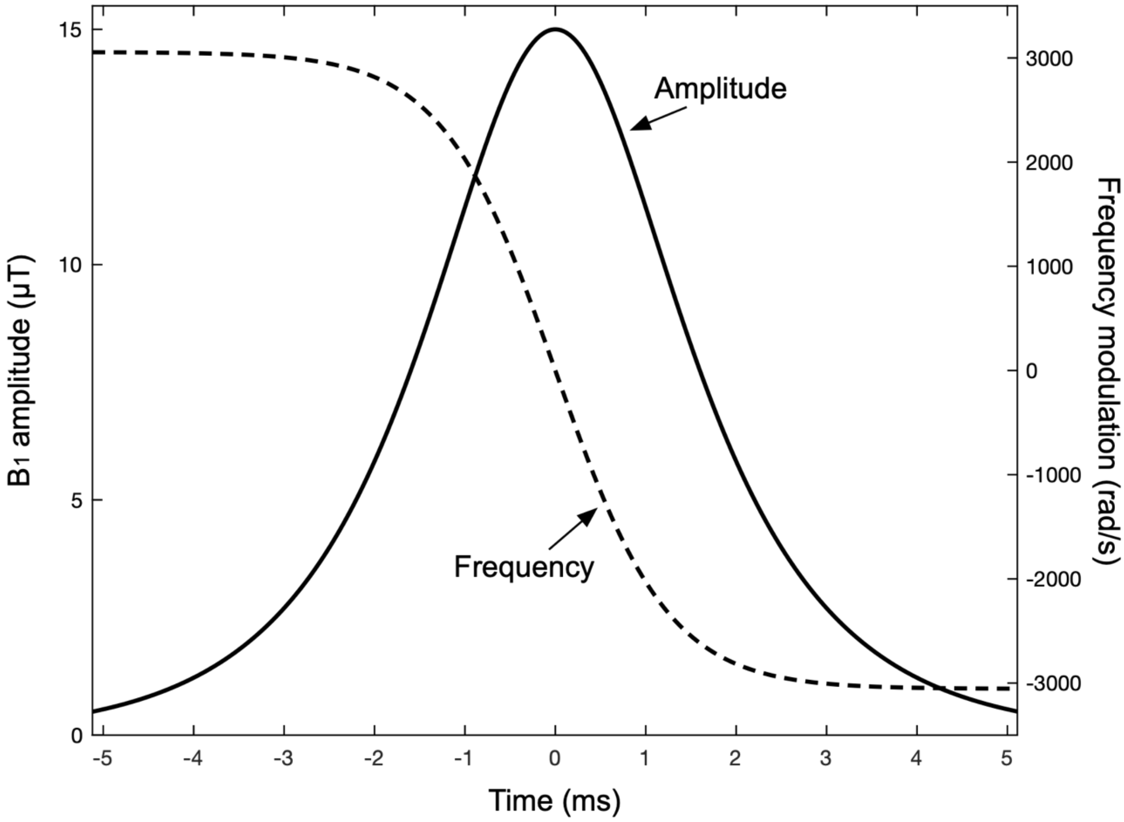

Figure 2.10: The B_1 amplitude and frequency modulation functions for a hyperbolic secant pulse with: T_p = 10.24 ms, A_0 = 15 μT, μ = 24, and β = 800 rad/s.

2.3.2 Frequency Offset Corrected Inversion (FOCI) Pulse

Another commonly used slice selective adiabatic inversion pulse is the ‘C-shape’ FOCI pulse.[29] This pulse type is based on the hyperbolic secant pulse, but modulates the B_1 amplitude, frequency modulation, and gradient amplitude functions by a shaping function:

$$A(t) = \begin{array}{cc} \cosh(\beta t) & \text{for } \cosh(\beta t) < A_{max} \\ A_{max} & otherwise \end{array}, \tag{2.24}$$

where A_{max} is the maximum FOCI modulation.

The FOCI transformed hyperbolic secant pulse uses the same peak B_1 amplitude as the corresponding unmodulated hyperbolic secant pulse and maintains the same trajectory of the effective field. Therefore, if the original hyperbolic secant pulse satisfied the adiabatic condition, the FOCI pulse will as well. However, the FOCI modulated pulse make use of the fact that spins at either edge of the selected slab are inverted at the

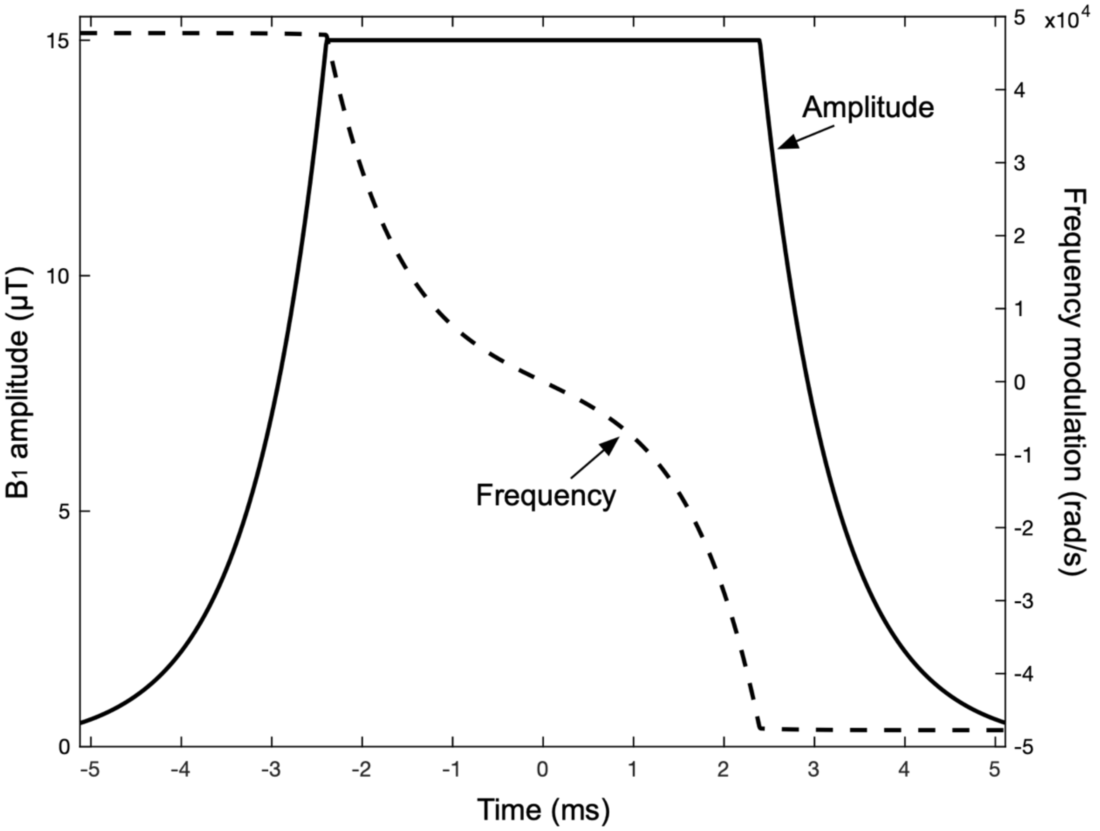

Figure 2.11: The B_1 amplitude and frequency modulation functions for a FOCI modulated hyperbolic secant pulse with: $A_{max} = 10$, $T_p = 10.24$ ms, $A_0 = 15\ \mu T$, $\mu = 24$, and $\beta = 1250$ rad/s.

beginning and end of the pulse. By modulating the B_1 amplitude, frequency modulation, and gradient amplitude functions simultaneously, the bandwidth at the edges of the pulse can be greatly increased, improving the effective slab profile.

A similar hyperbolic secant pulse to that shown in Figure 2.10 is shown in Figure 2.11 after being modulated by the FOCI shaping function in Equation 2.24 with $A_{max} = 10$. Note, the modulation angular frequency, β, has been increased to 1250 rad/s to maintain a similar level of pulse truncation to the original hyperbolic secant pulse.[30]

2.4 References

1 Haacke EM, Brown RW, Thompson MR, Venkatesan R. *Magnetic Resonance Imaging*. John Wiley & Sons Ltd: Chichester, UK, 1999 doi:10.1002/9781118633953.

2 Jezzard P, Matthews PM, Smith SM. *Functional Magnetic Resonance Imaging*. Oxford University Press, 2001 doi:10.1093/acprof:oso/9780192630711.001.0001.

3 Bernstein MA, King KF, Zhou XJ. *Handbook of MRI Pulse Sequences*. Elsevier: Amsterdam, 2004 doi:10.1016/B978-0-12-092861-3.X5000-6.

4 Bloch F. Nuclear Induction. *Phys Rev* 1946; **70**: 460–474.

5 Bloch F, Hansen WW, Packard M. The Nuclear Induction Experiment. *Phys Rev* 1946; **70**: 474–485.

6 Zur Y, Wood ML, Neuringer LJ. Spoiling of transverse magnetization in steady-state sequences. *Magn Reson Med* 1991; **21**: 251–263.

7 Hargreaves BA. Bloch Equation Simulator. http://mrsrl.stanford.edu/~brian/blochsim/ (accessed 1 Feb2013).

8 Haase A, Frahm J, Matthaei D, Hanicke W, Merboldt K-D. FLASH imaging. Rapid NMR imaging using low flip-angle pulses. *J Magn Reson* 1986; **67**: 258–266.

9 Lu H, Clingman C, Golay X, van Zijl PCM. Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn Reson Med* 2004; **52**: 679–682.

10 Lauterbur PC. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature* 1973; **242**: 190–191.

11 Jackson JI, Meyer CH, Nishimura DG, Macovski A. Selection of a convolution function for Fourier inversion using gridding (computerised tomography application). *IEEE Trans Med Imaging* 1991; **10**: 473–478.

12 Glover GH, Pauly JM. Projection Reconstruction Techniques for Reduction of Motion Effects in MRI. *Magn Reson Med* 1992; **28**: 275–289.

13 Mansfield P. Multi-planar image formation using NMR spin echoes. *J Phys C Solid State Phys* 1977; **10**: L55–L58.

14 Peters DC, Korosec FR, Grist TM, Block WF, Holden JE, Vigen KK *et al.* Undersampled projection reconstruction applied to MR angiography. *Magn Reson Med* 2000; **43**: 91–101.

15 Lai CM, Lauterbur PC. True three-dimensional image reconstruction by nuclear magnetic resonance zeugmatography. *Phys Med Biol* 1981; **26**: 851–856.

16 Mansfield P, Howseman AM, Ordidge RJ. Volumar imaging using NMR spin echoes: echo-volumar imaging (EVI) at 0.1 T. *J Phys E* 1989; **22**: 324–330.

17 Poser BA, Koopmans PJ, Witzel T, Wald LL, Barth M. Three dimensional echo-planar imaging at 7 Tesla. *Neuroimage* 2010; **51**: 261–266.

18 Feinberg DA, Oshio K. GRASE (gradient- and spin-echo) MR imaging: a new fast clinical imaging technique. *Radiology* 1991; **181**: 597–602.

19 Günther M, Oshio K, Feinberg DA. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn Reson Med* 2005; **54**: 491–498.

20 Meiboom S, Gill D. Modified Spin-Echo Method for Measuring Nuclear Relaxation Times. *Rev Sci Instrum* 1958; **29**: 688–691.

21 He X, Wengler K, Schweitzer ME. Diffusion sensitivity of 3D-GRASE in arterial spin labeling perfusion. *Magn Reson Med* 2018; **80**: 736–747.

22 Liang X, Connelly A, Tournier J-D, Calamante F. A variable flip angle-based method for reducing blurring in 3D GRASE ASL. *Phys Med Biol* 2014; **59**: 5559–5573.

23 Feinberg DA, Hale JD, Watts JC, Kaufman L, Mark A. Halving MR imaging time by conjugation: demonstration at 3.5 kG. *Radiology* 1986; **161**: 527–531.

24 Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med* 1999; **42**: 952–962.

25 Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J *et al.* Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002; **47**: 1202–1210.

26 Tannús A, Garwood M. Adiabatic pulses. *NMR Biomed* 1997; **10**: 423–434.

27 Norris DG. Adiabatic radiofrequency pulse forms in biomedical nuclear magnetic resonance. *Concepts Magn Reson* 2002; **14**: 89–101.

28 Silver M., Joseph R., Hoult D. Highly selective π/2 and π pulse generation. *J Magn Reson* 1984; **59**: 347–351.

29 Ordidge RJ, Wylezinska M, Hugg JW, Butterworth E, Franconi F. Frequency offset corrected inversion (FOCI) pulses for use in localized spectroscopy. *Magn Reson Med* 1996; **36**: 562–566.

30 Payne GS, Leach MO. Implementation and evaluation of frequency offset corrected inversion (FOCI) pulses on a clinical MR system. *Magn Reson Med* 1997; **38**: 828–833.

3

Current Arterial Spin Labelling Methods

3.1 Introduction

This chapter introduces key concepts in brain perfusion and angiographic imaging using arterial spin labelling. The topics covered include the most commonly used labelling techniques, dynamic sampling methods, image readouts, and background suppression (BGS) strategies. A frequently used PCASL perfusion model, which is used throughout this thesis, is also introduced. A more in-depth description of these concepts can be found in references [1–4] and the references therein.

3.2 Labelling Methods

Arterial spin labelling methods are numerous and varied, but one concept that they all share is that the ^{1}H in arterial blood supplying the brain is magnetically labelled and used as a freely diffusible endogenous tracer. This inflowing blood is typically labelled by either saturation,[5] or more commonly inversion.[6] It is then imaged while still in the large arteries (angiography) or once it exchanges with extravascular water in the tissue bed (perfusion). A second image, called the control image, in which the inflowing blood is not labelled, is also acquired. Subtraction of the labelled image from the control image cancels out static tissue signal (signal which is common in both images), leaving a signal relating only to the labelled blood, which for perfusion imaging is assumed proportional to the local rate of blood flow. This process is demonstrated in Figure 3.1.

The technique by which inflowing blood water is labelled varies and a few of the most common techniques are described here. The currently recommended labelling

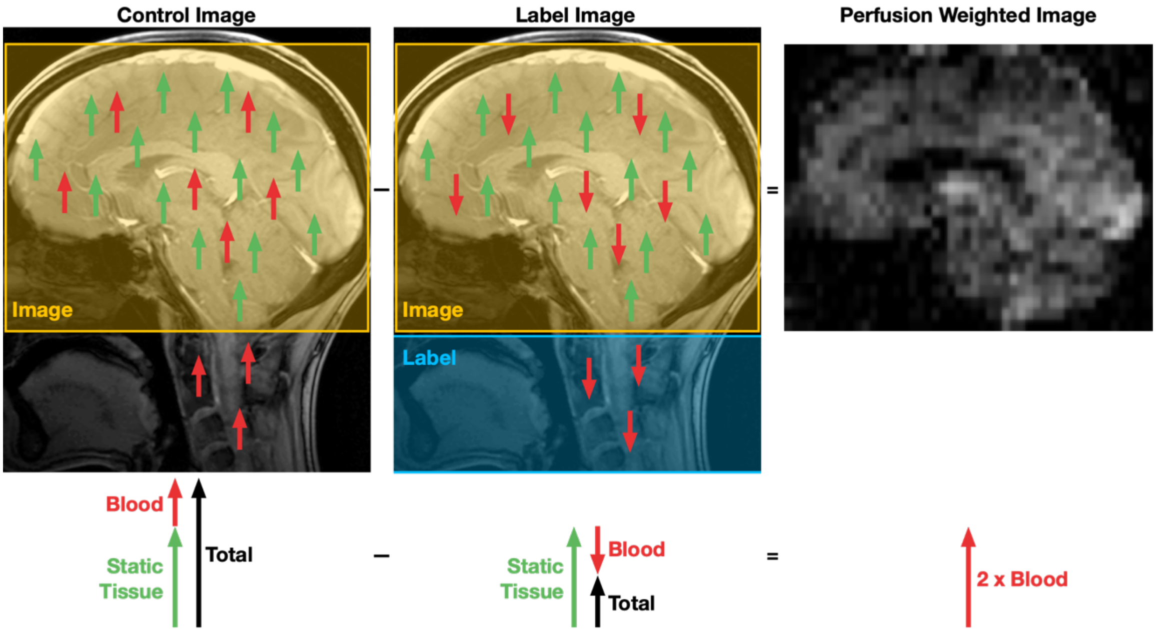

Figure 3.1: A visual depiction of the basic concepts behind ASL. Prior to acquiring the control image, the inflowing blood is left unperturbed, while in the label image the inflowing blood is inverted. In both cases, there is a delay before image acquisition to allow the bolus of inflowing blood to flow into the image volume. After this delay an image is acquired. When the label image is subtracted from the control image, the static tissue signal is removed, resulting in a signal proportional to the labelled blood signal. In this example, a long delay was used to allow the labelled blood to reach the tissue bed, resulting in a perfusion weighted image (PWI). T_1 *relaxation is ignored for simplicity.*

technique for standard brain perfusion imaging is pseudo-continuous ASL,[7,8] due to its superior SNR, ease of use, and widespread availability on commercial scanners.[1] PCASL labelling is used throughout this thesis for both perfusion imaging and angiography, though several other labelling techniques are also described for completeness.

3.2.1 Pulsed ASL (PASL)

Pulsed ASL methods use adiabatic fast passage pulses, similar to those described in Chapter 2, to invert spins across a large slab in a relatively short time of ~10 ms, the size of the slab being limited by the size and coverage of the transmit coil. As previously mentioned, adiabatic inversion pulses can be relatively insensitive to moderate variations in the B_1 amplitude. They can, therefore, achieve a relatively consistent high inversion efficiency of ~98% across a large region,[9] though the inversion efficiency in vivo is likely to be lower than this.[10] The largest size of the slab is typically limited to ~10 cm, which corresponds to a labelled bolus of blood with a temporal width of ~0.7 - 0.8 s.[9]

It is important to control for the effects of the labelling pulse in the imaging region, such as magnetisation transfer (MT),[11] otherwise the difference image may contain signal that is not related to perfusion. MT is a phenomenon related to the large range of resonant frequencies of bound water molecules. When the labelling pulse inverts blood distal to the imaging region, the ^1H in free water in the imaging region is off-resonance to the pulse and so will be negligibly affected. ^1H in bound water, such as those in macromolecules, may be on- or near-resonance, and so will be perturbed by the inversion pulse. Magnetisation between these two water pools exchanges and so the free water signal in the image region can be perturbed by the labelling pulse.

To control for this effect, most PASL methods use a similar RF pulse in the control condition to that used in the label image. How this is done varies across the many methods which have been developed. The multi-slice version of EPISTAR[12] uses two inversion pulses for the control image which are half the power of the pulse in the label condition, causing inflowing spins to be inverted then un-inverted, while approximately matching the MT effects in the label image. PICORE[13] uses the same labelling pulse as EPISTAR (proximal to the imaging region) but uses an off-resonance non-selective inversion pulse for the control condition. FAIR[14] uses a non-selective inversion pulse in the label condition and a slice-selective pulse covering just the imaging region in for the control condition.

3.2.2 Continuous ASL (CASL)

Continuous ASL also inverts blood proximal to the imaging region but does so using flow driven adiabatic inversion. Unlike the adiabatic pulses discussed previously, the frequency sweep in flow-driven adiabatic inversion comes from flowing spins moving through a gradient field, rather than the frequency of the pulse being swept.

A gradient field with constant amplitude is applied along the direction of flow at the same time as a continuous amplitude RF pulse. The carrier frequency of the RF pulse is set equal to the Larmor frequency along the gradient which corresponds to the desired location of the centre of inversion; this location is referred to as the labelling plane. As spins move along the gradient field and through the labelling plane, a process very similar to that described for adiabatic fast passage pulses in Chapter 2 takes place. In the rotating frame of reference of a spin flowing towards the labelling plane, the labelling pulse is far off-resonance when the spins are far away. As blood flow moves the spin closer to the labelling plane, the Larmor frequency of the spin is swept towards the carrier frequency of the labelling pulse; in the spin's frame of reference, the frequency of the pulse is swept towards the precession frequency of the spin. The result is that the effective field, $\vec{B}_{\text{eff}}$, that the flowing spin experiences as it flows past the labelling plane, is rotated from $+\hat{z}$ to $-\hat{z}$ (without loss of generality).

If the amplitude of the gradient is set appropriately for the velocity of the flowing blood, then the adiabatic condition is satisfied and spins flowing through the labelling plane will tend to precess around the slowly rotating effective field and become inverted. If the gradient amplitude is too large, then the effective field will rotate too quickly and spins will not follow it. If it is too small, spins will easily follow $\vec{B}_{eff}$, but will spend a long time in the transverse plane and experience a greater amount of T_2 relaxation. Fast and slow velocities have the same effect as large and small gradient amplitudes, respectively, resulting in the well-known velocity dependence of CASL.[15]

An advantage of CASL over PASL is that the temporal labelling duration can be much longer since it is not limited by the spatial coverage of the RF transmit coil. This means that even though the labelling efficiency is lower (≤90%),[15] the perfusion-weighted signal can be larger because of the greater accumulation of labelled blood in the tissue. A further advantage is that labelled blood only experiences T_1 decay from the position of the labelling plane. For PASL methods, the proximal end of the labelled bolus will have experienced much more T_1 relaxation by the time it reaches the tissue bed than the distal end, meaning CASL can have a higher SNR, even for matched label durations.[9]

The standard CASL sequence used a labelling plane placed above the head for the control condition to match MT effects (which are large for CASL due to the low gradient amplitude used). However, MT is only matched for a single slice equidistant to the label and control plane. A further development used a cosine modulated RF pulse for the control condition, creating two parallel labelling planes, both proximal to the brain. The 2nd labelling plane un-inverts the labelled blood while matching the MT effects of the labelling plane. However, this double inversion leads to a large reduction in the labelling efficiency.[16]

An alternative is to use a small dedicated labelling coil with coverage only over the neck.[17] The B_1 field strength is then negligible in the imaging region, removing MT effects and the need for an RF pulse in the control condition. However, dedicated hardware is necessary to achieve this.

Finally, the most prominent impediment to the use of CASL is that RF amplifiers on modern clinical scanners are unable to produce long duration continuous RF B_1 fields, having been optimised for short pulse durations.

3.2.3 Pseudo-Continuous ASL (PCASL)

The pseudo-continuous labelling pulse train[7,8] consists of a series of short (~0.5 ms), evenly spaced (~1 ms), RF pulses with an associated gradient field along the direction of flow that can be thought of as a discretised version of continuous labelling. The repeated pulses of PCASL leads to inversion in a somewhat similar manner to standard flow-driven adiabatic inversion, though there are important differences to CASL.

The PCASL inversion process is best viewed as a steady-state free-precession (SSFP) pulse sequence.[18] Ignoring relaxation effects, the steady-state longitudinal magnetisation of a series of repeated RF pulses is described by:[7]

$$M_z = \frac{\pm M_0 \sin(\alpha)\sin(\phi/2)}{\sqrt{(1-\cos(\alpha))^2 + \sin^2(\alpha)\sin^2(\phi/2)}}, \tag{3.1}$$

where α is the flip-angle and ϕ is the phase shift experienced between pulses. This behaviour is plotted in Figure 3.2. If we can get flowing magnetisation to traverse the red path in Figure 3.2 (the steady-state for $-180° < \phi < 180°$), it would give a similar effect to the inversion in PASL and CASL. The frequency of the RF pulse is again set to the desired location of the labelling plane.

In the gaps between RF pulses (~0.5 ms), gradient rewinders are used to keep the average gradient strength low (similar to that used in CASL) so that flowing spins experience a slow change in precession frequency, meeting the adiabatic condition. The

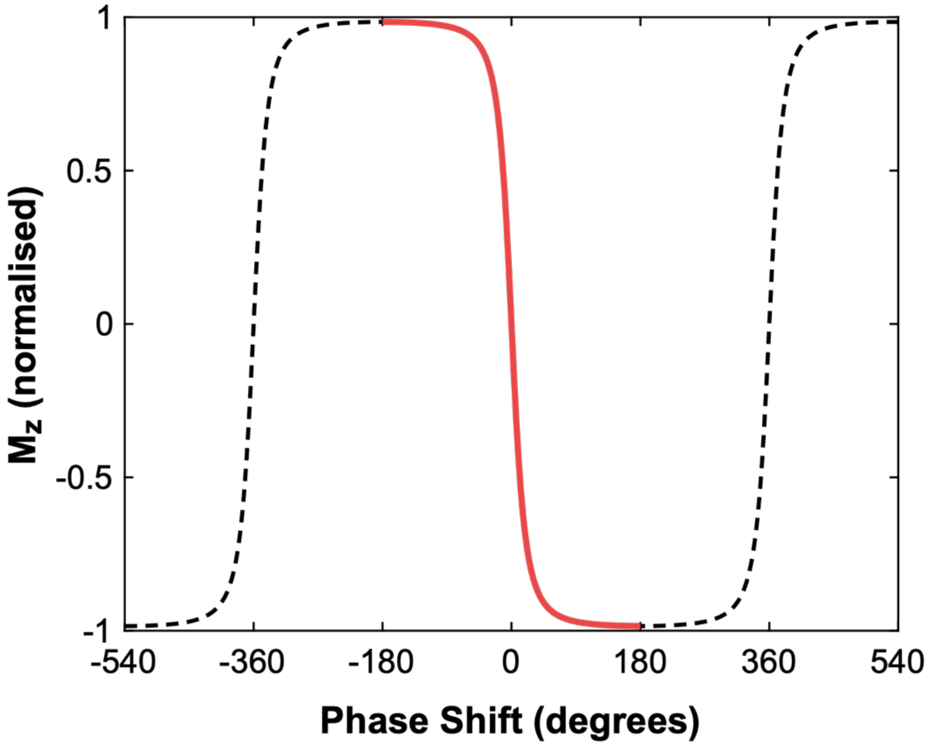

Figure 3.2: The steady-state behaviour of PCASL labelling shown across a range of inter-pulse phase shifts, ϕ, for a flip-angle, α, of 20°. The red section highlights the desired behaviour. Relaxation effects are ignored.

phase of each RF pulse is incremented based on the phase accumulation at the labelling plane due to the average gradient field. Since the phase accumulation is simply the time integral of the precession frequency, as spins move through the average gradient field and past the labelling plane, the phase shift between pulses moves from $\phi < 0$ to $\phi > 0$. If this is done slowly, spins flowing past the labelling plane will tend to follow the steady-state solution and become inverted.

The periodic nature of the PCASL labelling pulses leads to aliased labelling planes, which can be seen in the periodicity of the M_z steady-state solution across ϕ. The aliased labelling planes can be suppressed by using slice-selective pulses, through a combination of high gradient amplitudes during the RF pulses, noting that the distance between labelling planes is dependent on the mean gradient, and by limiting the pulse bandwidth, typically done by using Hanning or Gaussian shaped pulses.

Two varieties of the control condition have been proposed, referred to as unbalanced[7,8] and balanced[19] PCASL. In both cases, the control condition adds an RF phase increment of 180° to each RF pulse. This leads to spins to be flipped back and forth, resulting in minimal M_z perturbation, while matching the MT effects of the label condition. The difference between the two methods lies in the gradients used. In unbalanced PCASL, the mean gradient in the control condition is set to zero, reducing off-resonance sensitivity, but also resulting in different (unbalanced) gradients used for the label and control condition. This could lead to different eddy currents in each condition and result in artefactual signal in the difference images.[20] Balanced PCASL uses the same gradients in both label and control condition, better matching possible eddy currents, but also increasing the sensitivity to off-resonance. Balanced PCASL is necessary for vessel-encoded PCASL,[19] but unbalanced PCASL is recommended for standard use.[1] The implementation used in this thesis uses balanced PCASL. In-house testing of this sequence with a multi-phase protocol[21] has not shown significant signal loss due to off-resonance at the labelling plane in healthy controls at 3T.[22]

One issue with PCASL is an increased sensitivity to off-resonance fields. Like CASL, an inhomogeneous magnetic field will shift the labelling plane. However, unlike CASL, where this can have minimal effect on labelling efficiency, inversion efficiency can be lost with PCASL. One cause of this lost efficiency is the difference in phase between flowing blood and the RF pulses at the labelling plane. These phase offsets occur during the time between RF pulses and lead to a loss of inversion efficiency in the label condition and increased perturbation of the magnetisation in the control condition, ultimately decreasing the ASL difference signal. In an extreme case where spins accumulate an extra 180° phase shift between pulses, the label condition will become the control and the control condition will become the label. Proposed methods that can correct for this are: (1) improved shimming of the labelling plane,[8] (2) the use of multi-phase PCASL,[21] where the RF phase is purposefully incremented over several values and used to estimate the phase offset and (3) insertion of an extra RF phase term and transverse gradient blips between pulses calibrated using a prescan.[23–26]

Another way of thinking about this cause of lost inversion efficiency due to off-resonance is that the labelling plane will be shifted by a greater distance than the excitation slice profile of the RF pulses, because the low average gradient dictates the position of the labelling plane, while the high RF slice-select gradient dictates the centre of the RF excitation profile.[27] Shifting of these relative to each other will reduce the B_1 amplitude at the labelling plane and reduce inversion efficiency. Using this description, one method to improve robustness to this effect is to use unbalanced PCASL and broaden the slice-profile of the RF pulses, while being careful not to include the aliased labelling planes.[27]

Due to the much more efficient control condition for multi-slice imaging of PCASL compared to CASL, PCASL can achieve higher labelling efficiency.[20] It should be noted that the MT effects in PCASL are reduced compared to CASL, due to the larger gradient used during the RF pulses and near-perfect matching of the label and control conditions.

Like CASL, PCASL labelling efficiency has a dependency on the flow velocity[8,28] which varies across the cardiac cycle,[29] between arteries,[30] subjects,[30] disease states,[31] and during gas challenges.[28,32] Corrections for this can be performed by using phase contrast data as a normalisation factor[28] or by directly measuring the labelling efficiency using a prescan[33] or during the ASL scan itself.[34,35] Attempts to maximise the labelling efficiency across the cardiac cycle include the work of Luh et al.,[36] where the flow velocity was measured during the cardiac cycle and the PCASL B_1^+ is then dynamically varied accordingly.

3.3 Perfusion Imaging

3.3.1 Quantification

After inflowing blood water has been labelled (Section 3.2), the bolus of labelled blood travels down the arterial tree into the capillaries where it exchanges with water in the tissue. The tissue water compartment is much larger than the intravascular water compartment, meaning that most of the labelled water exchanges quickly. This larger compartment size also means that the label is slow to leave the tissue via venous outflow. It is this accumulation of labelled blood in the tissue that makes the longer possible label durations of CASL/PCASL advantageous.

3.3.1.1 Signal Model

A commonly used kinetic model for describing the accumulation of labelled blood in tissue was given by Buxton et al.[37] and is referred to as the standard model. The CASL/PCASL standard model is used exclusively throughout this thesis to describe the perfusion signal. The model is given as:

$$
\begin{aligned}
\Delta M(t) &= 0 && 0 < t < \Delta t \\
&= 2M_{0B} f T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}} \left(1 - e^{\frac{-(t-\Delta t)}{T_1'}}\right) && \Delta t < t < \tau + \Delta t \\
&= 2M_{0B} f T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}} e^{\frac{-(t-\tau-\Delta t)}{T_1'}} \left(1 - e^{\frac{-\tau}{T_1'}}\right) && \tau + \Delta t < t
\end{aligned}
\tag{3.2}
$$

where $1/T_1' = 1/T_{1t} + f/\lambda$, ΔM is the difference in longitudinal magnetisation between label and control images, t is the time from the start of labelling (s), M_{0B} is the equilibrium magnetisation of arterial blood, f is the CBF (s^{-1}), λ is the equilibrium brain-blood water partition coefficient (mL g^{-1}), α is the labelling efficiency, Δt is the arterial transit time (s), τ is the labelling duration (s), and T_{1b} and T_{1t} are the longitudinal relaxation time constants for arterial blood and tissue (s), respectively.

The assumptions of the model include the following. (1) The created labelled bolus of blood is a boxcar function. (2) As blood travels from the labelling plane to a voxel, it moves via uniform plug flow such that the time taken for this journey is well defined by a single transit time, Δt. This assumption results in the input function $c(t) = e^{-\Delta t/T_{1b}}$ for the time interval $\Delta t < t < \tau + \Delta t$ and zero elsewhere. (3) All blood arriving in a voxel instantaneously exchanges into the tissue. Before exchange, the label decays with T_{1b} but after exchange, decays with T_{1t}. The relaxation function is then $m(t) = e^{-t/T_{1t}}$. (4) Single-compartment kinetics is used to describe the water exchange between tissue and venous blood, such that the tissue compartment is well-mixed and the concentration ratio between tissue and venous blood is constant and equal to λ. The residue function, the fraction of labelled blood that previously arrived in the voxel and still remains at time *t*, is then $r(t) = e^{-ft/\lambda}$.

The model given in Equation 3.2 is a convolution of the input function with the relaxation function and the residue function, the result of which is scaled by $2\alpha M_{0B} f$:

$$\Delta M(t) = 2\alpha M_{0B} f\{c(t) \otimes [r(t) \cdot m(t)]\} \tag{3.3}$$

The analytical form of Equation 3.3 with the previously described assumptions is the formula given in Equation 3.2.

Other models have been proposed, such as two-compartment models that do not assume instantaneous exchange from the intra- to the extra-vascular space[38,39] and models that do not assume a well-mixed tissue compartment.[40] However, when only grey matter perfusion is considered, which is the case for this thesis, then the standard model is reasonably accurate.[41]

Further models which include a pass-through intravascular component[42–44] or incorporate dispersion of the labelled bolus[45–48] have also been proposed. However, these models require a larger number of parameters to be estimated from the data, leading to greater uncertainty in the perfusion estimates. Due to their increased complexity, they are not used for estimation of perfusion in this thesis.

3.3.1.2 Measuring CBF

As is evident from the standard model (Equation 3.2), if we wish to measure CBF there are many parameters which potentially must be either be estimated from the data, measured in a separate scan or assumed known. Many of these are signal scaling factors, so do not affect the shape of the signal curve, but do scale the absolute CBF estimate.

PCASL labelling creates a well-defined temporal bolus of labelled blood and so is assumed to be a known experimentally imposed parameter. The labelling efficiency, α, varies with blood velocity and can be either estimated from a prescan[28,33] or during the

PCASL scan,[34,35] but in practice it is often assumed to be 0.85.[8,28] The brain-blood water partition coefficient, λ, can be measured using a prescan,[49] but a brain-wide literature value of 0.9 mL/g is typically used.[50] T_{1b} and T_{1t}, the blood and tissue longitudinal relaxation times, can also be measured during a prescan,[51–55] but often literature values are used. In ASL analysis, T_{1b} is often assumed to be 1.65 s at 3T,[1,54] though it varies across subjects and with disease.[54,56,57] One reason for the variation reported in the literature is the linear dependence of T_{1b} on haematocrit fraction,[54] though this relationship may not hold in some diseases, such as sickle cell disease.[57] T_{1b} also varies linearly with temperature and has a small dependence on blood oxygenation.[54] There is also a large variation in literature T_{1t} values, with the value varying across the brain,[53,58,59] subjects,[60] and disease states.[60–62] In this thesis, where only healthy subjects were scanned, a single value of 1.445 s is used for grey matter at 3T.[55] T_{1t} has been shown to vary regionally by a large amount in the presence of tumours,[61,62] stroke,[62] and Parkinson disease.[60] This variation can have a large impact on CBF estimates when using ASL and the acquisition of a T_1 map should be considered when designing an ASL experiment.[63] A single value for M_{0B} can be estimated from a calibration image (an image with identical readout as the ASL data but without labelling or BGS) using a reference tissue,[64] typically CSF,[65,66] and literature values for relaxation times and λ. However, a simpler approach is to divide a smoothed calibration image by λ to use as a voxelwise estimate of M_{0B}.[1] This approach has the benefit of also approximately correcting for B_1^+ and B_1^- inhomogeneity, but leads to an increase in spatial noise due to the image division. The regionally dependent SNR caused by varying transmit and receive sensitivities can also not be corrected for.

The only two remaining unknown parameters are CBF and the arterial transit time (ATT). The signal model given in Equation 3.2 is shown in Figure 3.3 for a range of CBF and ATT values, demonstrating the variation in the kinetic model when these parameters vary. The signal evolution is composed of three sections: (1) the ATT period, where the leading edge of the labelled bolus has not yet reached the voxel; (2) the inflow period, where the label is accumulating in the tissue; (3) the signal decay period, where the signal is dominated by T_1 relaxation and a small amount of venous outflow.

It can be seen that CBF is mostly a scaling factor of the difference signal. The ATT shifts the signal curve, leading to large signal changes during the inflow period, while not greatly affecting the signal decay period where $t > \Delta t_{max} + \tau$. Therefore, if the label and control images are acquired after a post labelling delay (PLD) longer than or equal to the longest ATT in the image, the image intensity is mostly dependent on CBF, while being insensitive to ATT[67] (Note, the PLD is the time between the end of labelling and image acquisition).

CBF can be estimated by fitting the model to the data; if a single PLD is acquired, only CBF can be estimated while an assumed value must be used for ATT. If instead, multiple PLDs are acquired, both CBF and ATT can be estimated.

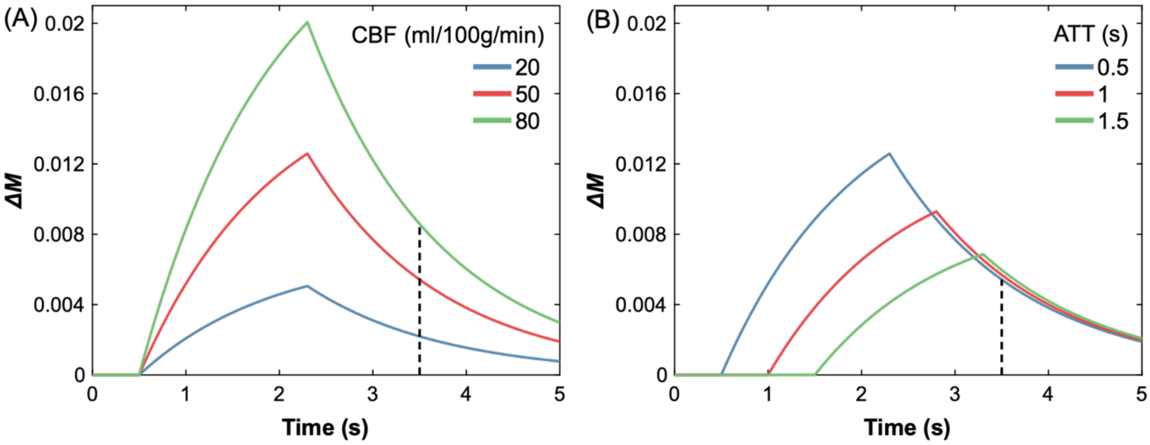

Figure 3.3: The PCASL difference signal evolution using the standard model (Equation 3.2) when varying CBF (A) and ATT (B). The common parameters used were: $\tau = 1.8$ s, $\alpha = 0.85$, $\lambda = 0.9$ mL/g, $T_{1b} = 1.65$ s, $T_{1t} = 1.445$ s, $M_{0B} = 1$. In (A), ATT = 0.5 s, while in (B), CBF = 50 mL/100g/min. The vertical dashed line highlights a region relatively insensitive to variation in the ATTs plotted. The signal is positive because the labelled image (with decreased longitudinal signal) is subtracted from the control image.

3.3.2 Sampling Strategies

This section describes a range of protocol designs for acquiring perfusion data for CBF quantification. The techniques described are used and compared in Chapters 4 and 5.

3.3.2.1 Single PLD

It is most common to use a single label duration (LD) and post-labelling delay (PLD) to measure CBF. This is the default technique recommended by the ISMRM Perfusion Study Group and the European ASL in Dementia consortium.[1] However, use of a single-PLD assumes that either the ATT is known or that the PLD is longer than the longest ATT in the imaging region. If ATT > PLD, large CBF underestimates can occur due to incomplete arrival of the label, while if the PLD is unnecessarily long, signal-to-noise (SNR) is lost, resulting in less precise measurements. For healthy subjects, it has been recommended to use a LD of 1.8 s and a PLD of 1.8 s.[1]

Multiple images (on the order of 30-50) are typically acquired and averaged across to improve the low SNR of ASL based CBF measurements, as well as to average over scanner instabilities and physiological variability.

3.3.2.2 Multiple PLDs

Alternatively, images can be acquired at multiple PLDs, sampling the signal curve at different points in its evolution. CBF and ATT can then be simultaneously estimated, reducing the potential ATT related error in CBF.[37] The ATT measurements can also be useful in and of themselves, such as being a possible indicator of disease severity for vertebrobasilar large artery disease[68] and an indicator of collateral flow.[69]

Multiple PLDs are typically sampled by acquiring one image at a given PLD per ASL preparation, with the PLD being sequentially varied for each acquisition. Multiple PLDs can also be acquired by using a Look-Locker type readout,[70] where flip-angles < 90° are used with multiple readout modules at different PLDs per ASL preparation. This has been shown to be more SNR-efficient than sequentially acquired multiple-PLDs, despite the reduction in flip-angle and the signal attenuation across Look-Locker frames,[42,71] but has limited slice coverage. Furthermore, the attenuation of ASL signal over the repeated readouts needs to be taken into account, potentially increasing the number parameters to be estimated.[42]

The acquisition of multiple PLDs, rather than a single-PLD, reduces the number of averages that can be acquired for each PLD in a given scan time. However, since all of the PLDs are combined during fitting, this is not necessarily detrimental to SNR, depending on the PLDs used. This concept is explored in Chapter 4.

3.3.2.3 Time-Encoded Multi-PLD

As an alternative to acquiring images at separate PLDs, the PLDs can be encoded into the PCASL preparation using a method called time-encoded PCASL.[72] This process is outlined in Figure 3.4. The labelling duration is divided into multiple blocks, or sub-boluses, where the PCASL condition switches between label and control. Each acquired image has a varying contribution of labelled signal accumulation from each sub-bolus. The signal from each individual sub-bolus can be extracted by addition and subtraction of the acquired images according to the encoding pattern.

A Hadamard encoding scheme, as used in Figure 3.4, is the most efficient encoding scheme because each row is orthogonal to the others, meaning that for a given row, no linear combination of the other rows is equal to it. This is more time-efficient than acquiring separate unencoded multiple PLDs because N difference images can be decoded from N+1 encoded images, rather than from $2 \cdot N$ images. There is also a factor of $\sqrt{(N+1)/2}$ reduction in the noise SD compared to conventional multi-PLD (assuming white Gaussian noise) from the averaging inherent in the decoding process.[73] However, the LDs used are typically shorter in time-encoded PCASL, reducing the accumulation of signal in the tissue and potentially reducing the accuracy of CBF estimates.[73,74]

Several variants of the standard time-encoded protocol design, where each sub-bolus has the same LD, have been suggested including the "free-lunch" method, where a long LD for the first block is used, with shorter equal LDs for the remaining blocks.[75] After decoding, the image corresponding to the first block provides similar information to the single-PLD protocol. However, because the conventional long PLD has been filled with short time-encoded blocks, additional temporal information is acquired without consequence, hence "free-lunch". Another variant is the T_1-adjusted design, where the LDs

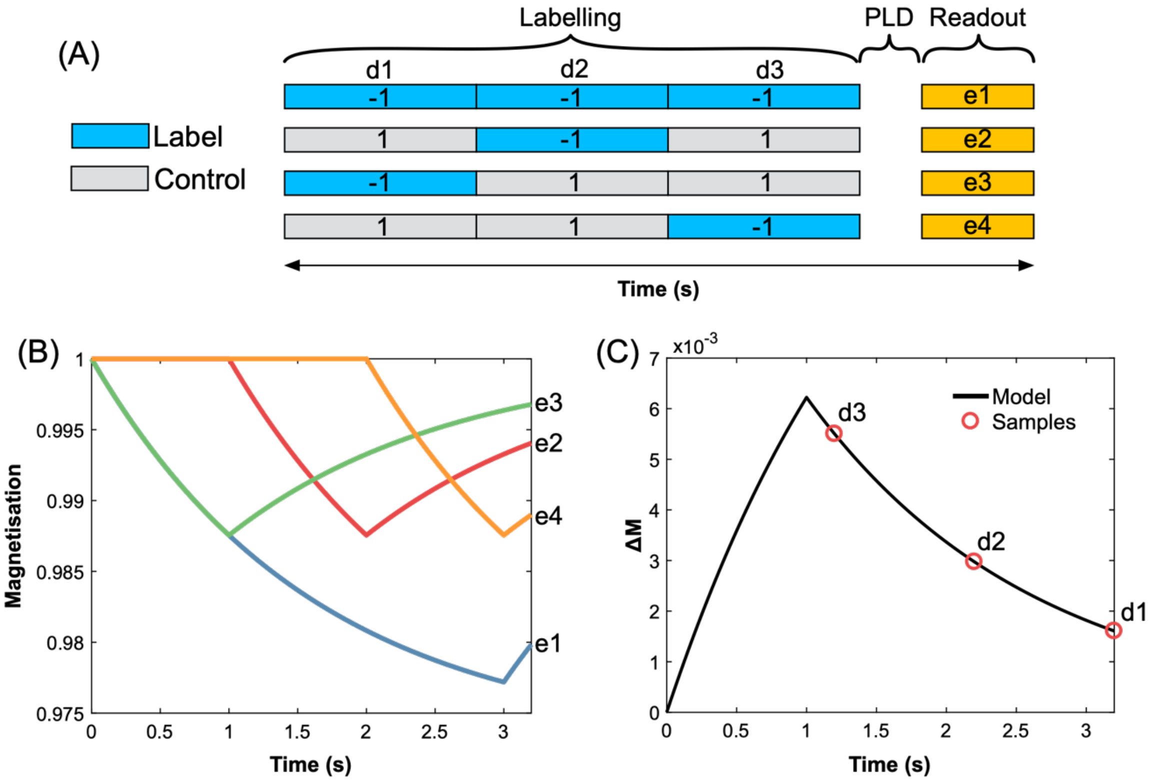

Figure 3.4: Graphical representation of time-encoded PCASL. (A) A simplified time-encoded protocol diagram. The label and control conditions are encoded into the PCASL labelling preparation using a 4×3 *Hadamard matrix, where each row corresponds to a separate ASL TR. The PCASL module is followed by a short PLD then the image readout. Four images are acquired with different encoded preparations (e1, e2, e3, e4) in order to decode the signal from the individual sub-boluses (d1, d2, d3). (B) The signal evolution over time for the encoded images. The equilibrium magnetisation is set to 1, and inflowing labelled blood decreases the signal. The images are acquired at 3.2 s. (C) The decoded signal corresponding to each block (red) with the full signal model. Settings: LD 1 s, PLD 0.2 s, CBF 50 mL/100g/min, ATT 0 s. Other settings as in Figure 3.3.*

are adjusted to account for T_1 relaxation, such that at the time of image acquisition, the summed signal across each sub-bolus is equal.[75] This results in progressively shorter LDs for the blocks temporally closer to the readout.

The signal decoding process assumes that contributions from each sub-bolus are constant over the entire acquisition, i.e. that physiology is stable, in much the same way that the standard subtraction of label from control assumes physiological stability. Averaging in standard PCASL single-PLD and multi-PLD has much the same effect as the effective averaging in the decoding process of time-encoded PCASL, where scanner and physiological variations are averaged across.

Time-encoded PCASL is more sensitive to motion than standard PCASL because if an image becomes corrupted it is no longer possible to decode any of the images for the entire encoded set, rather than just one label and control. A Walsh-ordered Hadamard matrix has been suggested to reduce this sensitivity, enabling initial images to be decoded while some of the encoded images are yet to be acquired, allowing a rescan to be performed if required.[76]

3.3.3 Readouts

The ISMRM/EU ASL perfusion consensus paper[1] recommends the use of fast segmented 3D readouts such as 3D GRASE[77,78] (as described in Chapter 2) or stack-of-spirals,[79,80] due to their superior SNR, uniform PLD across the entire volume, and compatibility with BGS techniques (described in Section 3.5). Segmenting the acquisitions (i.e. using multiple acquisitions to acquire all of *k*-space, also known as multi-shot imaging) helps to reduce T_2 blurring, but also increases sensitivity to subject motion.

Fast multi-slice 2D readouts such as EPI[81] and spirals[82] have been widely used with ASL, but lead to varying PLDs and poor BGS performance for most slices. Parallel imaging in the form of multi-band excitation[83,84] can be used with multi-slice readouts with

ASL,[85] decreasing the variability in PLDs and BGS across slices and accelerating data acquisition. Alternatively, the PLD variability can be utilised to benefit the acquisition of multi-PLD data by cycling the order of slice acquisition using the round-robin technique.[86]

Large voxel sizes (3-4 mm in-plane and 4-8 mm through-plane)[1] are typically used to improve SNR. This can cause significant partial volume effects of CBF estimates. Since grey matter (GM) typically has much higher blood flow than white matter (WM), CBF maps often spatially resemble grey matter partial volume estimate maps. It is potentially possible to reduce this bias during fitting,[87] where the different temporal dynamics of GM and WM can be used to aid separation.[88]

3.4 Angiography

In ASL angiography,[89–91] it is the contrast of labelled spins in the large arteries that is of interest. As such, images are acquired at much shorter PLDs. Due to the small size of intracranial arteries (typically 1.4 mm - 4.2 mm[92] at the circle of Willis, with downstream arteries typically even smaller), high spatial resolutions are necessary for vessel visualisation. Since images are acquired much sooner after labelling, minimal T_1 relaxation will have occurred and the signal will be concentrated in the arteries. SNR is, therefore, much higher than for perfusion imaging, facilitating higher spatial resolutions. Lower readout bandwidths are also typically used in angiography than perfusion imaging, further compensating resolution related SNR decreases.

3.4.1 Static Angiography

Static angiograms, which provide structural information about the arterial supply, can be acquired after filling the vasculature with labelled blood: this can be achieved with PCASL via the use of a long LD,[93] or with PASL labelling with a sufficiently long inversion time to allow inflow.[90] The duration of the PCASL label is not limited, so a sufficiently long

LD can be found to fill even very distal vessels. However, the limited physical size of the PASL label and the need to set an appropriate inversion time may lead to poor visibility of either proximal or distal vessels.

3.4.2 Dynamic Angiography

Dynamic angiograms are typically achieved by the use of a Look-Locker type readout[70,94] which acquire the same region of *k*-space at multiple time-points (frames) during a single flow-through of the label. The remaining regions of *k*-space are acquired over subsequent ASL preparations and readout periods, with the *k*-space data being combined before reconstruction. The result is that high temporal and spatial resolution images can be acquired. An alternative is to acquire data for a single image with much greater *k*-space coverage per ASL preparation and increment the PLD across different preparations,[95] similar to the sequential acquisition of PLDs for perfusion imaging (discussed in Section 3.3.2.2). The downside of this approach is that the temporal footprint of each image will be much larger, risking motion-related artefacts from the movement of labelled blood.

Dynamic angiograms provide additional haemodynamic information that may be useful in assessing disease, such as carotid artery disease.[96] They can also be fit with a kinetic signal model[97] which may provide further quantitative markers sensitive to disease. Synthetic angiograms can also be generated from the fitted data, which have been found to improve visualisation of small vessels compared to the raw images.[97]

Both PCASL and PASL can be used for dynamic angiography. However, with a PASL preparation, image acquisition can begin almost immediately, facilitating visualisation of the inflow of blood into the imaging region. Short PCASL LDs would need to be used to be able to image the inflow of blood. Nonetheless, similar information can be achieved with long PCASL LDs if inflow-subtraction is used.[98,99]

Time-encoded PCASL preparations have previously been combined with dynamic angiographic readouts to reduce the number of readout frames required, leading to improved depiction of distal vessels in later frames.[100]

3.4.3 Readouts

As mentioned above, dynamic ASL angiographic readouts typically acquire a small portion of *k*-space per frame per ASL preparation. This can be achieved with single-line readouts where a single line of *k*-space is acquired after each excitation. To achieve a desired temporal resolution, X, the number of *k*-space lines, L, that can be acquired during each ASL preparation is limited to X/TR. The number of ASL preparations required to fill *k*-space (referred to as the number of segments) is equal to N/L, where N is the total number of *k*-space lines. The scan time scales with the number of segments, all else being equal, since typically only a single average is acquired. It is therefore of great advantage to use acquisition acceleration, such as parallel imaging,[101,102] because the scan time is inversely proportional to the acceleration factor, R. This differs from perfusion imaging where the whole of *k*-space is often acquired after each ASL preparation, and additional measurements are used for averaging, so acceleration does not affect scan time as dramatically. Static angiograms are able to make use of longer readout times, but as mentioned previously, if the temporal footprint is too long, motion artefacts could appear.

Single-line readouts have been demonstrated with both SPGR[94,99,103,104] (as described in Chapter 2) and balanced steady-state free precession (bSSFP)[98,105,106] sequences. bSSFP readouts fully refocus transverse magnetisation using balanced gradients, preserving it for the next TR, making bSSFP readouts more SNR efficient than SPGR readouts, which spoil the transverse magnetisation. However, bSSFP readouts suffer from signal loss due to resonance offsets, which can be particularly problematic for large FOVs.[98] SPGR readouts don't suffer from this issue, but because the longitudinal magnetisation is attenuated by a factor of $\cos(\alpha)$ by each excitation pulse, the flip-angle, α, is limited to smaller values than for bSSFP readouts.

Both Cartesian[94,104,106] and radial trajectories[89,103,107,108] have been demonstrated with ASL angiography. Radial trajectories have the advantage of reduced sensitivity to motion[109] and diffuse angular undersampling artefacts.[110] The sparse nature of angiograms means that the artefacts from moderate angular acceleration are typically acceptable,[107] while the relatively incoherent undersampling artefacts also make radial trajectories more suitable for compressed sensing reconstructions.[111] Radial trajectories are also more robust to signal variations across the readout, which manifest as angular undersampling artefacts in radial trajectories,[112] but as a loss of resolution in Cartesian trajectories.[113]

For the SPGR readout, it is common to use a series of identical low flip-angle excitation pulses, referred to as constant flip-angles (CFAs),[99,103,104] but variable flip-angles (VFAs) across excitations can also be used and have been found to improve the depiction of small distal vessels.[114]

3.5 Background Signal Suppression

The aim of the image subtraction between label and control images in ASL is to subtract out static tissue signal which is present in both images. However, this subtraction will not be perfect due to subject motion and physiological variation such as pulsatility, resulting in increased ‘physiological noise’. This physiological noise scales with the static tissue signal and so BGS techniques have been developed to suppress this signal while leaving the ASL signal relatively unaffected.

In 1991, Dixon et al.[115] introduced a BGS technique for ASL angiography that is fundamentally similar to those used today for both angiography and perfusion imaging. This technique is demonstrated in Figure 3.5. The basis of this technique is as follows: (1)

at time $t = 0$, a presaturation pulse saturates the spins in the imaging region. Ideally, this sets $M_z = 0$, making the static tissue longitudinal magnetisation independent of the tissue signal at the end of the previous ASL TR. (2) Typically, immediately after the presaturation pulse, the spatially-selective spin labelling takes place, which labels spins proximal to the imaging region. (3) During the inflow of blood, two non-selective inversion pulses are played out which invert the labelled blood and the static tissue signal. These inversion pulses are timed relative to the presaturation pulse in order to null static tissue with two different T_1 values at the time of the image readout. The control-label difference is unaffected because it is simply inverted and reinverted. (4) The labelled image is acquired. This process is repeated in an identical manner for the control image, resulting in greatly reduced static tissue signal in both images, while the ASL signal is unaffected.

Further developments for ASL angiography by Mani et al.[116] included the use of a B_1 insensitive multiple pulse presaturation module with 4 B_1 and B_0 insensitive hyperbolic secant pulses to improve signal nulling across a wide range of T_1 values. Ye et al.[79] then demonstrated the effectiveness of BGS for suppressing background tissue related subtraction artefacts in ASL perfusion imaging using a presaturation module with two inversion pulses.

The BGS scheme used in this thesis for both ASL perfusion imaging and angiography is conceptually similar to the original method described by Dixon et al.[115] but with some key differences.

Rather than using a single 90° pulse for presaturation, a series of 4 pulses with optimised flip-angles are used, referred to as WET presaturation (**W**ater suppression **E**nhanced through ***T_1*** effects). The WET presaturation module was first used by Golay et al.,[117] having been adapted from the spectroscopic work of Ogg et al.[118] The flip-angles, in order, are: 88.9°, 98.7°, 82.5°, 159.0°. After each RF pulse, gradient spoilers, using a cycled scheme, are used to dephase the transverse signal. Assuming perfect spoiling after each pulse, this set of flip-angles can robustly saturate spins with T_1 values in the range of

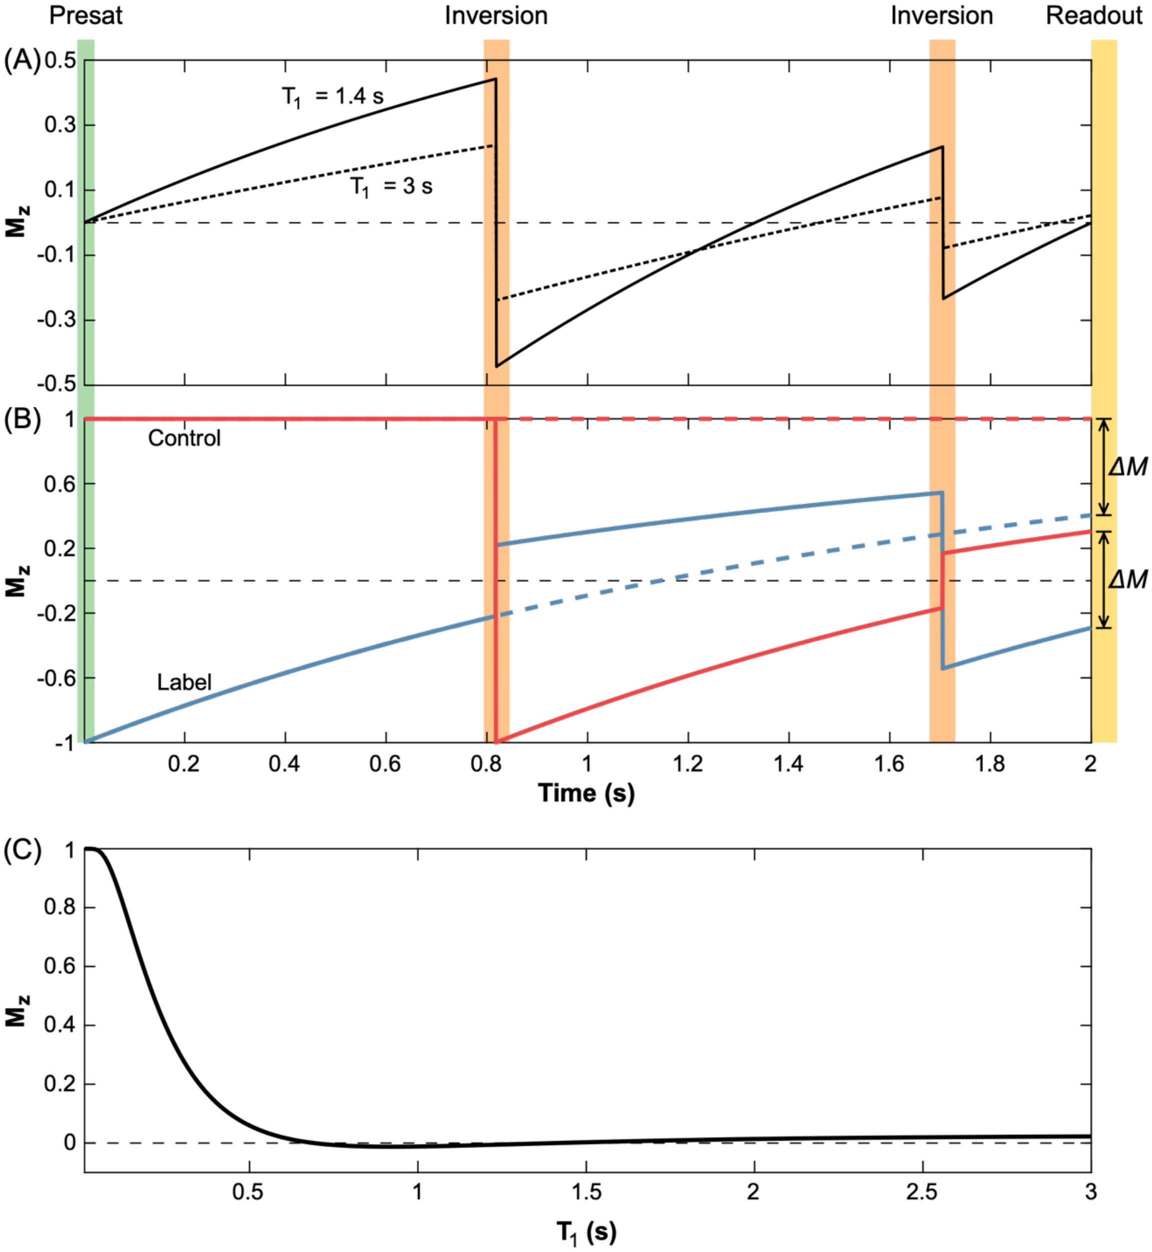

Figure 3.5: The effect of background suppression on static tissue and labelled/control blood. (A) The presaturation module saturates the static tissue which then starts to relax with T_1. The two inversion pulses are timed to null $T_1 = 0.7$ s and 1.4 s using Equation 3.4. At the readout, the static tissue is well supressed even for other T_1 values. (B) The blood is labelled/controlled instantaneously after the presaturation pulse in this example but is not affected by the presaturation which is applied only to the imaging region. The inversion pulses invert the labelled/controlled blood, but the difference in M_z between them remains the same. This is demonstrated by the dashed lines which show the relaxation of blood if inversion pulses were not used. $T_{1b} = 1.65$ s. (C) The final M_z of tissue for a range of T_1 values. Tissue with T_1 values larger than 0.7 are all well supressed for the timings used. Note the difference in scale between the graphs. Instantaneous RF pulses and perfect saturation and inversion were assumed.

400 - 4200 ms and with B_1 field variations of ±10%.

Two adiabatic inversion pulses are used for BGS in this thesis. The timings of these inversion pulses relative to the end of the presaturation module are calculated using the analytical solution given by Günther et al.[78] to null two T_1 values, $T_{1,opt}$ and $2 \cdot T_{1,opt}$, at $t = \mathrm{TI}$. The optimal inversion timings, τ_1 and τ_2, are:

$$\tau_{1,2}(\mathrm{TI}) = \mathrm{TI} + \frac{2}{R_{1,opt}} \cdot \ln\left(\left(\frac{1}{2} \pm \frac{1}{4}\right) + \left(\frac{1}{2} \mp \frac{1}{4}\right) \cdot e^{-\frac{1}{2}\mathrm{TI} \cdot R_{1,opt}}\right), \tag{3.4}$$

where TI is the time from the end of the presaturation module to the start of the image readout and $R_{1,opt} = 1/T_{1,opt}$. This formula was used to calculate the timings of the inversion pulses in Figure 3.5.

BGS schemes using a larger number of inversion pulses can achieve improved suppression across a wide range of T_1 values,[116,119] but analytical solutions to their optimal timings have not been given in the literature and may not exist. The timings must, therefore, be numerically optimised, which can make adjusting protocol timings and the associated BGS inversion timings a cumbersome process. Furthermore, the use of a larger number of inversion pulses leads to greater attenuation of the ASL contrast (see below) and so it is not necessarily given that improved BGS through the use of a larger number of inversion pulses leads to improved SNR.

The BGS simulation shown in Figure 3.5 assumed perfect inversion of tissue and blood M_z. The use of adiabatic inversion pulses improves robustness to variation in B_1, but 100% efficient inversion pulses do not exist. Even with very high B_1 amplitudes in simulation, some signal is always lost to T_2 relaxation during inversion. Garcia et al.[10] also demonstrated that the in vivo inversion efficiency of inversion pulses is much lower than in simulation or in phantoms, possibly due to slow MT effects in blood. For this reason, the ASL signal is attenuated by these inversion pulses. Each inversion pulse effectively reduces the ASL contrast, and so a correction should be used for quantitative perfusion imaging. This can be approximated by a simple scaling factor of the ASL difference signal.[1,120]

3.6 Conclusions

In this chapter, current methods for ASL perfusion imaging and angiography were described. Several widely used labelling schemes were discussed, with the conclusion being that PCASL leads to higher SNR due the longer LDs and reduced T_1 relaxation compared to PASL while having higher SNR and better matched MT effects than CASL for multi-slice imaging.

A commonly used perfusion model for CBF quantification was introduced, with its assumptions explained. The argument for using the standard model rather than more intricate models is due to its simplicity and widespread use. Three broad perfusion protocol designs from the literature were described along with their relative advantages and disadvantages. All three protocol designs are compared in later chapters and so are discussed further in the remainder of this thesis. The current recommendations for perfusion imaging readouts were summarised. In chapter 4, a multi-slice 2D EPI readout was used for demonstration of the protocol optimisation algorithm developed in that chapter. In chapter 5, a single-shot 3D GRASE readout was used due to its improved SNR and uniform PLDs and BGS across slices. To reduce the risk of motion-related artefacts, acquisition segmentation was not used.

Current ASL angiographic methods were introduced. An SPGR readout is used in Chapter 6 of this thesis due to its robustness to B_0 inhomogeneities and the intention of extending the methods developed in that chapter to whole brain 3D imaging. A radial trajectory is also used due to its robustness to motion and signal variations as well as its suitability for image acceleration.

The importance of BGS for ASL methods was explained with the development of techniques being summarised. The BGS scheme used throughout this thesis was described along with the trade-off between alternative schemes.

3.7 References

1 Alsop DC, Detre JA, Golay X, Günther M, Hendrikse J, Hernandez-Garcia L *et al.* Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn Reson Med* 2015; **73**: 102–116.

2 Wong EC. An introduction to ASL labeling techniques. *J Magn Reson Imaging* 2014; **40**: 1–10.

3 Jezzard P, Chappell MA, Okell TW. Arterial spin labeling for the measurement of cerebral perfusion and angiography. *J Cereb Blood Flow Metab* 2018; **38**: 603–626.

4 van Osch MJ, Teeuwisse WM, Chen Z, Suzuki Y, Helle M, Schmid S. Advances in arterial spin labelling MRI methods for measuring perfusion and collateral flow. *J Cereb Blood Flow Metab* 2018; **38**: 1461–1480.

5 Detre JA, Leigh JS, Williams DS, Koretsky AP. Perfusion imaging. *Magn Reson Med* 1992; **23**: 37–45.

6 Williams DS, Detre JA, Leigh JS, Koretsky AP. Magnetic resonance imaging of perfusion using spin inversion of arterial water. *Proc Natl Acad Sci U S A* 1992; **89**: 4220–4220.

7 Garcia DM, Bazelaire C, D. Alsop. Pseudo-continuous flow driven adiabatic inversion for arterial spin labeling. In: *Proceedings of the 13th Annual meeting of the ISMRM, Miami Beach, USA*. 2005, p 37.

8 Dai W, Garcia D, De Bazelaire C, Alsop DC. Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magn Reson Med* 2008; **60**: 1488–1497.

9 Wong EC, Buxton RB, Frank LR. A theoretical and experimental comparison of continuous and pulsed arterial spin labeling techniques for quantitative perfusion imaging. *Magn Reson Med* 1998; **40**: 348–55.

10 Garcia DM, Duhamel G, Alsop DC. Efficiency of inversion pulses for background suppressed arterial spin labeling. *Magn Reson Med* 2005; **54**: 366–372.

11 Wolff SD, Balaban RS. Magnetization transfer contrast (MTC) and tissue water proton relaxationin vivo. *Magn Reson Med* 1989; **10**: 135–144.

12 Edelman RR, Chen Q. EPISTAR MRI: multislice mapping of cerebral blood flow. *Magn Reson Med* 1998; **40**: 800–805.

13 Wong EC, Buxton RB, Frank LR. Implementation of quantitative perfusion imaging techniques for functional brain mapping using pulsed arterial spin labeling. *NMR Biomed* 1997; **10**: 237–249.

14 Kim S-G. Quantification of relative cerebral blood flow change by flow-sensitive alternating inversion recovery (FAIR) technique: Application to functional

mapping. *Magn Reson Med* 1995; **34**: 293–301.

15 Maccotta L, Detre JA, Alsop DC. The efficiency of adiabatic inversion for perfusion imaging by arterial spin labeling. *NMR Biomed* 1997; **10**: 216–221.

16 O'Gorman RL, Summers PE, Zelaya FO, Williams SCR, Alsop DC, Lythgoe DJ. In vivo estimation of the flow-driven adiabatic inversion efficiency for continuous arterial spin labeling: A method using phase contrast magnetic resonance angiography. *Magn Reson Med* 2006; **55**: 1291–1297.

17 Zaharchuk G, Ledden PJ, Kwong KK, Reese TG, Rosen BR, Wald LL. Multislice perfusion and perfusion territory imaging in humans with separate label and image coils. *Magn Reson Med* 1999; **41**: 1093–1098.

18 Carr HY. Steady-State Free Precession in Nuclear Magnetic Resonance. *Phys Rev* 1958; **112**: 1693–1701.

19 Wong EC. Vessel-encoded arterial spin-labeling using pseudocontinuous tagging. *Magn Reson Med* 2007; **58**: 1086–1091.

20 Wu WC, Fernández-Seara M, Detre JA, Wehrli FW, Wang J. A theoretical and experimental investigation of the tagging efficiency of pseudocontinuous arterial spin labeling. *Magn Reson Med* 2007; **58**: 1020–1027.

21 Jung Y, Wong EC, Liu TT. Multiphase pseudocontinuous arterial spin labeling (MP-PCASL) for robust quantification of cerebral blood flow. *Magn Reson Med* 2010; **64**: 799–810.

22 Berry ESK, Jezzard P, Okell TW. An Optimized Encoding Scheme for Planning Vessel-Encoded Pseudocontinuous Arterial Spin Labeling. *Magn Reson Med* 2015; **74**: 1248–1256.

23 Shin DD, Liu TT, Wong EC, Shankaranarayanan A, Jung Y. Pseudocontinuous arterial spin labeling with optimized tagging efficiency. *Magn Reson Med* 2012; **68**: 1135–1144.

24 Luh W-M, Talagala SL, Li T-Q, Bandettini PA. Pseudo-continuous arterial spin labeling at 7 T for human brain: Estimation and correction for off-resonance effects using a Prescan. *Magn Reson Med* 2013; **69**: 402–410.

25 Berry ESK, Jezzard P, Okell TW. Off-resonance correction for pseudo-continuous arterial spin labeling using the optimized encoding scheme. *Neuroimage* 2019; : 272.

26 Jahanian H, Noll DC, Hernandez-Garcia L. B0 field inhomogeneity considerations in pseudo-continuous arterial spin labeling (pCASL): effects on tagging efficiency and correction strategy. *NMR Biomed* 2011; **24**: 1202–1209.

27 Zhao L, Vidorreta M, Soman S, Detre JA, Alsop DC. Improving the robustness of pseudo-continuous arterial spin labeling to off-resonance and pulsatile flow velocity. *Magn Reson Med* 2017; **78**: 1342–1351.

28 Aslan S, Xu F, Wang PL, Uh J, Yezhuvath US, Van Osch M *et al.* Estimation of labeling efficiency in pseudocontinuous arterial spin labeling. *Magn Reson Med* 2010; **63**: 765–771.

29 Wu WC, Mazaheri Y, Wong EC. The effects of flow dispersion and cardiac pulsation in arterial spin labeling. *IEEE Trans Med Imaging* 2007; **26**: 84–92.

30 Albayrak R, Degırmencı B, Acar M, Haktanır A, Colbay M, Yaman M. Doppler sonography evaluation of flow velocity and volume of the extracranial internal carotid and vertebral arteries in healthy adults. *J Clin Ultrasound* 2007; **35**: 27–33.

31 Václavů L, Baldew ZA V., Gevers S, Mutsaerts HJMM, Fijnvandraat K, Cnossen MH *et al.* Intracranial 4D flow magnetic resonance imaging reveals altered haemodynamics in sickle cell disease. *Br J Haematol* 2018; **180**: 432–442.

32 Bulte DP, Kelly M, Germuska M, Xie J, Chappell MA, Okell TW *et al.* Quantitative

measurement of cerebral physiology using respiratory-calibrated MRI. *Neuroimage* 2012; **60**: 582–591.

33 Chen Z, Zhang X, Yuan C, Zhao X, van Osch MJP. Measuring the labeling efficiency of pseudocontinuous arterial spin labeling. *Magn Reson Med* 2017; **77**: 1841–1852.

34 Chen Z, Zhao X, Zhang X, Guo R, Teeuwisse WM, Zhang B *et al.* Simultaneous measurement of brain perfusion and labeling efficiency in a single pseudo-continuous arterial spin labeling scan. *Magn Reson Med* 2018; **79**: 1922–1930.

35 Günther M. Arterial Spin Labeled Input Function (ASLIF): signal acquisition during pseudo-continuous arterial spin labeling. In: *Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France*. 2018, p 305.

36 Luh WM, Wong EC, Talagala SL, Bandettini P a. Tagging Efficiency Improvement Using Velocity-matched Pseudo-continuous Arterial Spin Labeling and VERSE. In: *Proceedings of the 16th Annual Meeting of the ISMRM, Toronto, Canada*. 2008, p 3341.

37 Buxton RB, Frank LR, Wong EC, Siewert B, Warach S, Edelman RR. A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magn Reson Med* 1998; **40**: 383–396.

38 Parkes LM, Tofts PS. Improved accuracy of human cerebral blood perfusion measurements using arterial spin labeling: Accounting for capillary water permeability. *Magn Reson Med* 2002; **48**: 27–41.

39 Zhou J, Wilson DA, Ulatowski JA, Traystman RJ, van Zijl PCM. Two-Compartment Exchange Model for Perfusion Quantification Using Arterial Spin Tagging. *J Cereb Blood Flow Metab* 2001; **21**: 440–455.

40 St. Lawrence KS, Frank JA, McLaughlin AC. Effect of restricted water exchange on cerebral blood flow values calculated with arterial spin tagging: A theoretical investigation. *Magn Reson Med* 2000; **44**: 440–449.

41 Parkes LM. Quantification of cerebral perfusion using arterial spin labeling: Two-compartment models. *J Magn Reson Imaging* 2005; **22**: 732–736.

42 Francis ST, Bowtell R, Gowland PA. Modeling and optimization of look-locker spin labeling for measuring perfusion and transit time changes in activation studies taking into account arterial blood volume. *Magn Reson Med* 2008; **59**: 316–325.

43 Li K, Zhu X, Hylton N, Jahng G-H, Weiner MW, Schuff N. Four-phase single-capillary stepwise model for kinetics in arterial spin labeling MRI. *Magn Reson Med* 2005; **53**: 511–518.

44 Chappell M a., MacIntosh BJ, Donahue MJ, Günther M, Jezzard P, Woolrich MW. Separation of macrovascular signal in multi-inversion time arterial spin labelling MRI. *Magn Reson Med* 2010; **63**: 1357–1365.

45 Hrabe J, Lewis DP. Two analytical solutions for a model of pulsed arterial spin labeling with randomized blood arrival times. *J Magn Reson* 2004; **167**: 49–55.

46 Gallichan D, Jezzard P. Modeling the effects of dispersion and pulsatility of blood flow in pulsed arterial spin labeling. *Magn Reson Med* 2008; **60**: 53–63.

47 Chappell MA, Woolrich MW, Kazan S, Jezzard P, Payne SJ, MacIntosh BJ. Modeling dispersion in arterial spin labeling: Validation using dynamic angiographic measurements. *Magn Reson Med* 2013; **69**: 563–570.

48 Kazan SM, Chappell MA, Payne SJ. Modeling the Effects of Flow Dispersion in Arterial Spin Labeling. *IEEE Trans Biomed Eng* 2009; **56**: 1635–1643.

49 Roberts DA, Rizi R, Lenkinski RE, Leigh JS. Magnetic resonance imaging of the brain: Blood partition coefficient for water: Application to spin-tagging measurement of perfusion. *J Magn Reson Imaging* 1996; **6**: 363–366.

50 Herscovitch P, Raichle ME. What is the Correct Value for the Brain-Blood Partition Coefficient for Water? *J Cereb Blood Flow Metab* 1985; **5**: 65–69.

51 Ethofer T, Mader I, Seeger U, Helms G, Erb M, Grodd W *et al.* Comparison of longitudinal metabolite relaxation times in different regions of the human brain at 1.5 and 3 Tesla. *Magn Reson Med* 2003; **50**: 1296–1301.

52 Wansapura JP, Holland SK, Dunn RS, Ball WS. NMR relaxation times in the human brain at 3.0 tesla. *J Magn Reson Imaging* 1999; **9**: 531–8.

53 Stanisz GJ, Odrobina EE, Pun J, Escaravage M, Graham SJ, Bronskill MJ *et al.* T1, T2 relaxation and magnetization transfer in tissue at 3T. *Magn Reson Med* 2005; **54**: 507–512.

54 Lu H, Clingman C, Golay X, van Zijl PCM. Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn Reson Med* 2004; **52**: 679–682.

55 Lin C, Bernstein M, Huston J, Fain S. Measurements of T1 Relaxation times at 3.0T: Implications for clinical MRA. In: *Proceedings of the 9th Annual Meeting of the ISMRM, Glasgow, Scotland*. 2001, p 1391.

56 Wu WC, Jain V, Li C, Giannetta M, Hurt H, Wehrli FW *et al.* In vivo venous blood T1 measurement using inversion recovery true-FISP in children and adults. *Magn Reson Med* 2010; **64**: 1140–1147.

57 Václavů L, van der Land V, Heijtel DFR, van Osch MJP, Cnossen MH, Majoie CBLM *et al.* In Vivo T1 of Blood Measurements in Children with Sickle Cell Disease Improve Cerebral Blood Flow Quantification from Arterial Spin-Labeling MRI. *Am J Neuroradiol* 2016; **37**: 1727–1732.

58 Clare S, Jezzard P. Rapid T(1) mapping using multislice echo planar imaging. *Magn Reson Med* 2001; **45**: 630–634.

59 Liberman G, Louzoun Y, Ben Bashat D. T1 Mapping using variable flip angle SPGR data with flip angle correction. *J Magn Reson Imaging* 2014; **40**: 171–180.

60 Vymazal J, Righini A, Brooks RA, Canesi M, Mariani C, Leonardi M *et al.* T1 and T2 in the Brain of Healthy Subjects, Patients with Parkinson Disease, and Patients with Multiple System Atrophy: Relation to Iron Content. *Radiology* 1999; **211**: 489–495.

61 Bastin ME, Sinha S, Whittle IR, Wardlaw JM. Measurements of water diffusion and T1 values in peritumoural oedematous brain. *Neuroreport* 2002; **13**: 1335–1340.

62 Naruse S, Horikawa Y, Tanaka C, Hirakawa K, Nishikawa H, Yoshizaki K. Significance of proton relaxation time measurement in brain edema, cerebral infarction and brain tumors. *Magn Reson Imaging* 1986; **4**: 293–304.

63 Debacker CS, Daoust A, Köhler S, Voiron J, Warnking JM, Barbier EL. Impact of tissue T1 on perfusion measurement with arterial spin labeling. *Magn Reson Med* 2017; **77**: 1656–1664.

64 Wong EC, Buxton RB, Frank LR. Quantitative imaging of perfusion using a single subtraction (QUIPSS and QUIPSS II). *Magn Reson Med* 1998; **39**: 702–708.

65 Wu W-C, Wong EC. Intravascular effect in velocity-selective arterial spin labeling: The choice of inflow time and cutoff velocity. *Neuroimage* 2006; **32**: 122–128.

66 Macintosh BJ, Pattinson KTS, Gallichan D, Ahmad I, Miller KL, Feinberg DA *et al.* Measuring the Effects of Remifentanil on Cerebral Blood Flow and Arterial Arrival Time Using 3D Grase MRI with Pulsed Arterial Spin Labelling. *J Cereb Blood Flow Metab* 2008; **28**: 1514–1522.

67 Alsop DC, Detre JA. Reduced Transit-Time Sensitivity in Noninvasive Magnetic Resonance Imaging of Human Cerebral Blood Flow. *J Cereb Blood Flow Metab* 1996; **16**: 1236–1249.

68 MacIntosh BJ, Marquardt L, Schulz UG, Jezzard P, Rothwell PM. Hemodynamic

alterations in vertebrobasilar large artery disease assessed by arterial spin-labeling MR imaging. *Am J Neuroradiol* 2012; **33**: 1939–1944.

69 Lou X, Yu S, Scalzo F, Starkman S, Ali LK, Kim D *et al.* Multi-delay ASL can identify leptomeningeal collateral perfusion in endovascular therapy of ischemic stroke. *Oncotarget* 2017; **8**: 2437–2443.

70 Look DC, Locker DR. Time Saving in Measurement of NMR and EPR Relaxation Times. *Rev Sci Instrum* 1970; **41**: 250–251.

71 Günther M, Bock M, Schad LR. Arterial spin labeling in combination with a look-locker sampling strategy: Inflow turbo-sampling EPI-FAIR (ITS-FAIR). *Magn Reson Med* 2001; **46**: 974–984.

72 Günther M. Highly efficient accelerated acquisition of perfusion inflow series by Cycled Arterial Spin Labeling. In: *Proceedings of the 15th Annual Meeting of the ISMRM, Berlin, Germany*. 2007, p 380.

73 Dai W, Shankaranarayanan A, Alsop DC. Volumetric measurement of perfusion and arterial transit delay using hadamard encoded continuous arterial spin labeling. *Magn Reson Med* 2013; **69**: 1014–1022.

74 Guo J, Holdsworth SJ, Fan AP, Lebel MR, Zun Z, Shankaranarayanan A *et al.* Comparing accuracy and reproducibility of sequential and Hadamard-encoded multidelay pseudocontinuous arterial spin labeling for measuring cerebral blood flow and arterial transit time in healthy subjects: A simulation and in vivo study. *J Magn Reson Imaging* 2018; **47**: 1119–1132.

75 Teeuwisse WM, Schmid S, Ghariq E, Veer IM, van Osch MJP. Time-encoded pseudocontinuous arterial spin labeling: Basic properties and timing strategies for human applications. *Magn Reson Med* 2014; **72**: 1712–1722.

76 von Samson-Himmelstjerna F, Madai VI, Sobesky J, Guenther M. Walsh-ordered hadamard time-encoded pseudocontinuous ASL (WH pCASL). *Magn Reson Med* 2016; **76**: 1814–1824.

77 Feinberg DA, Oshio K. GRASE (gradient- and spin-echo) MR imaging: a new fast clinical imaging technique. *Radiology* 1991; **181**: 597–602.

78 Günther M, Oshio K, Feinberg DA. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn Reson Med* 2005; **54**: 491–498.

79 Ye FQ, Frank JA, Weinberger DR, McLaughlin AC. Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). *Magn Reson Med* 2000; **44**: 92–100.

80 Vidorreta M, Balteau E, Wang Z, De Vita E, Pastor MA, Thomas DL *et al.* Evaluation of segmented 3D acquisition schemes for whole-brain high-resolution arterial spin labeling at 3 T. *NMR Biomed* 2014; **27**: 1387–1396.

81 Mansfield P. Multi-planar image formation using NMR spin echoes. *J Phys C Solid State Phys* 1977; **10**: L55–L58.

82 Vidorreta M, Wang Z, Rodríguez I, Pastor MA, Detre JA, Fernández-Seara MA. Comparison of 2D and 3D single-shot ASL perfusion fMRI sequences. *Neuroimage* 2013; **66**: 662–671.

83 Moeller S, Yacoub E, Olman CA, Auerbach E, Strupp J, Harel N *et al.* Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magn Reson Med* 2010; **63**: 1144–1153.

84 Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn Reson Med* 2012; **67**: 1210–1224.

85 Suzuki Y, Okell TW, Chappell MA, van Osch MJP. A framework for motion correction of background suppressed arterial spin labeling perfusion images acquired with simultaneous multi-slice EPI. *Magn Reson Med* 2019; **81**: 1553–1565.

86 Lee W, Janik R, Scouten A, Stefanovic B, Sled JG. Efficient sampling of early signal arrival for estimation of perfusion and transit time in whole-brain arterial spin labeling. *Magn Reson Med* 2012; **68**: 179–187.

87 Liang X, Connelly A, Calamante F. Improved partial volume correction for single inversion time arterial spin labeling data. *Magn Reson Med* 2013; **69**: 531–537.

88 Chappell MA, Groves AR, MacIntosh BJ, Donahue MJ, Jezzard P, Woolrich MW. Partial volume correction of multiple inversion time arterial spin labeling MRI data. *Magn Reson Med* 2011; **65**: 1173–1183.

89 Dixon WT, Du LN, Faul DD, Gado M, Rossnick S. Projection angiograms of blood labeled by adiabatic fast passage. *Magn Reson Med* 1986; **3**: 454–462.

90 Nishimura DG, Macovski A, Pauly JM, Conolly SM. MR angiography by selective inversion recovery. *Magn Reson Med* 1987; **4**: 193–202.

91 Wang SJ, Nishimura DG, Macovski A. Multiple-readout selective inversion recovery angiography. *Magn Reson Med* 1991; **17**: 244–251.

92 Kamath S. Observations on the length and diameter of vessels forming the circle of Willis. *J Anat* 1981; **133**: 419–23.

93 Koktzoglou I, Gupta N, Edelman RR. Nonenhanced extracranial carotid MR angiography using arterial spin labeling: Improved performance with pseudocontinuous tagging. *J Magn Reson Imaging* 2011; **34**: 384–394.

94 Günther M, Warmuth C, Zimmer C. Sub-millimeter Dynamic Spin Labeling Cerebral 2D-Angiography with 40ms Temporal Resolution. In: *Proceedings of the 10th Annual Meeting of the ISMRM, Honolulu, USA*. 2002, p 1100.

95 Wu H, Block WF, Turski PA, Mistretta CA, Rusinak DJ, Wu Y *et al.* Noncontrast dynamic 3D intracranial MR angiography using pseudo-continuous arterial spin labeling (PCASL) and accelerated 3D radial acquisition. *J Magn Reson Imaging* 2014; **39**: 1320–1326.

96 MacIntosh BJ, Sideso E, Donahue MJ, Chappell MA, Günther M, Handa A *et al.* Intracranial Hemodynamics Is Altered by Carotid Artery Disease and After Endarterectomy. *Stroke* 2011; **42**: 979–984.

97 Okell TW, Chappell M a., Schulz UG, Jezzard P. A kinetic model for vessel-encoded dynamic angiography with arterial spin labeling. *Magn Reson Med* 2012; **68**: 969–979.

98 Okell TW, Schmitt P, Bi X, Chappell MA, Tijssen RHN, Sheerin F *et al.* Optimization of 4D vessel-selective arterial spin labeling angiography using balanced steady-state free precession and vessel-encoding. *NMR Biomed* 2016. doi:10.1002/nbm.3515.

99 Kopeinigg D, Bammer R. Time-resolved angiography using inflow subtraction (TRAILS). *Magn Reson Med* 2014; **72**: 669–678.

100 Suzuki Y, Teeuwisse WM, Schmid S, Helle M, Cauteren M Van, Osch MJP Van. Improving 4D pCASL angiography by combining Hadamard time-encoding with Look-Locker readout. In: *Proceedings of the 22nd Annual Meeting of the ISMRM, Milan, Italy*. 2014, p 6460.

101 Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med* 1999; **42**: 952–962.

102 Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J *et al.* Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson*

Med 2002; **47**: 1202–1210.

103 Wu H, Block WF, Turski PA, Mistretta CA, Johnson KM. Noncontrast-enhanced three-dimensional (3D) intracranial MR angiography using pseudocontinuous arterial spin labeling and accelerated 3D radial acquisition. *Magn Reson Med* 2013; **69**: 708–715.

104 Okell TW, Chappell M a., Woolrich MW, Günther M, Feinberg D a., Jezzard P. Vessel-encoded dynamic magnetic resonance angiography using arterial spin labeling. *Magn Reson Med* 2010; **64**: 698–706.

105 Koktzoglou I, Meyer JR, Ankenbrandt WJ, Giri S, Piccini D, Zenge MO *et al.* Nonenhanced arterial spin labeled carotid MR angiography using three-dimensional radial balanced steady-state free precession imaging. *J Magn Reson Imaging* 2015; **41**: 1150–1156.

106 Bi X, Weale P, Schmitt P, Zuehlsdorff S, Jerecic R. Non-contrast-enhanced four-dimensional (4D) intracranial MR angiography: A feasibility study. *Magn Reson Med* 2010; **63**: 835–841.

107 Peters DC, Korosec FR, Grist TM, Block WF, Holden JE, Vigen KK *et al.* Undersampled projection reconstruction applied to MR angiography. *Magn Reson Med* 2000; **43**: 91–101.

108 Okell TW. Combined angiography and perfusion using radial imaging and arterial spin labeling. *Magn Reson Med* 2018. doi:10.1002/mrm.27366.

109 Glover GH, Pauly JM. Projection Reconstruction Techniques for Reduction of Motion Effects in MRI. *Magn Reson Med* 1992; **28**: 275–289.

110 Scheffler K, Hennig J. Reduced circular field-of-view imaging. *Magn Reson Med* 1998; **40**: 474–480.

111 Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med* 2007; **58**: 1182–1195.

112 Theilmann RJ, Gmitro AF, Altbach MI, Trouard TP. View-ordering in radial fast spin-echo imaging. *Magn Reson Med* 2004; **51**: 768–774.

113 Tan ET, Huston J, Campeau NG, Riederer SJ. Fast inversion recovery magnetic resonance angiography of the intracranial arteries. *Magn Reson Med* 2010; **63**: 1648–1658.

114 Schmitt P, Speier P, Bi X, Weale P, Mueller E. Non-contrast-enhanced 4D intracranial MR angiography: Optimizations using a variable flip angle approach. In: *Proceedings of the 18th Annual Meeting of the ISMRM, Stockholm, Sweden*. 2010, p 402.

115 Dixon WT, Sardashti M, Castillo M, Stomp GP. Multiple inversion recovery reduces static tissue signal in angiograms. *Magn Reson Med* 1991; **18**: 257–268.

116 Mani S, Pauly J, Conolly S, Meyer C, Nishimura D. Background suppression with multiple inversion recovery nulling: Applications to projective angiography. *Magn Reson Med* 1997; **37**: 898–905.

117 Golay X, Petersen ET, Hui F. Pulsed Star Labeling of Arterial Regions (PULSAR): A robust regional perfusion technique for high field imaging. *Magn Reson Med* 2005; **53**: 15–21.

118 Ogg RJ, Kingsley PB, Taylor JS. WET, a T1- and B1-insensitive water-suppression method for in vivo localized 1H NMR spectroscopy. *J Magn Reson B* 1994; **104**: 1–10.

119 Maleki N, Dai W, Alsop DC. Optimization of background suppression for arterial spin labeling perfusion imaging. *Magn Reson Mater Physics, Biol Med* 2012; **25**: 127–133.

120 Shin DD, Liu HL, Wong EC, Liu TT. Effect of background suppression on CBF

quantitation in pseudo continuous arterial spin labeling. In: *Proceedings of the 19th Annual Meeting of the ISMRM, Montreal, Canada*. 2011, p 2101.

4

Optimal Experimental Design for Multi-PLD Perfusion Imaging

4.1 Introduction

In ASL perfusion imaging, it is common to use a single PLD,[1] which allows many averages to be acquired within a short scan time, thus increasing SNR. However, if the PLD is shorter than the ATT, CBF can be severely underestimated due to incomplete arrival of the labelled bolus. The use of longer PLDs reduces the risk of such errors but must be balanced against the loss of tracer signal through T_1 decay and increased noise due to fewer averages being achievable within a given scan time.

An alternative method is to use multiple PLDs and to fit a kinetic signal model to the resulting dynamic data.[2] In this way, both CBF and ATT can be estimated, reducing the bias caused by unknown ATT and also providing extra, potentially clinically useful, information.[3] However, multiple PLDs present a more complicated experimental design problem: when should the ASL signal be sampled to give the most accurate CBF and ATT measurements? This increased complexity, along with a more involved analysis process, has restricted the use of multi-PLD techniques.[1] Studies that do use multiple PLDs often use equally spaced PLDs over a range of times reflective of the expected ATTs and based on previous experience.[4,5]

The field of optimal experimental design provides a mathematical framework, the Cramér-Rao lower bound (CRLB),[6,7] with which to design experiments to minimise the variance of estimated parameters. Previous studies, which used the CRLB to optimise the inversion time (TI) in PASL experiments, have shown promise, but were limited by: 1) aggregating a series of locally optimal TIs,[8,9] 2) optimising across both ATT and CBF prior distributions[9,10] and 3) use of Gaussian prior distributions over the parameter values.[8–10] Kramme et al.[11] also proposed increasing the number of averages at longer PLDs to improve CBF accuracy, but this method does not directly consider the parameter uncertainty.

In this chapter, a flexible framework is presented for designing ASL protocols by maximising the accuracy of CBF or both CBF and ATT estimates. This work builds on previous studies,[8,10] to produce a simplified framework which combines the information obtained across all acquisitions, optimising the experiment as a whole, within a defined scan time.

The effectiveness of this framework is demonstrated by optimising the PLDs of a multi-PLD pseudo-continuous ASL[12] (PCASL) experiment, using a 2D multi-slice gradient-echo EPI readout across a uniform ATT distribution appropriate for grey matter (GM) in healthy volunteers. Two protocols were generated: one which minimises both CBF and ATT errors, and one which minimises only CBF errors while remaining insensitive to ATT variation. These protocols were compared against matched scan-time reference multi-PLD[13] and single-PLD protocols, using Monte Carlo simulations and in vivo experiments, and demonstrate that the use of a broad, population-specific, uniform ATT distribution can successfully reduce CBF and ATT errors.

The bulk of this chapter is formed from work that was presented at the 34th annual meeting of the European Society for Magnetic Resonance in Medicine and Biology[*] and was subsequently published in Magnetic Resonance in Medicine.[†]

---

[*] Woods JG, Chappell MA, Okell TW. Optimizing Post Labeling Delays in Multiple-Delay Arterial Spin Labeling MRI for Cerebral Perfusion Imaging. In: *Proceedings of the 34th Annual Meeting of the ESMRMB, Barcelona, Spain*. 2017, p 245.

[†] Woods JG, Chappell MA, Okell TW. A general framework for optimizing arterial spin labeling MRI experiments. *Magn Reson Med* 2019; **81**: 2474–2488.

4.2 Sequence Implementation

The multi-slice gradient-echo 2D-EPI PCASL sequence used for the work in this chapter was an existing custom-built sequence from our laboratory. It is identical to the PCASL sequence used in Okell et al.[13] and was implemented on a 3T Siemens Prisma scanner.

Each TR can be decomposed into 5 separate components, as shown in Figure 4.1, which are: a WET presaturation module, the PCASL label/control preparation, a PLD, two global hyperbolic secant (HS) inversion pulses, and a 2D-EPI multi-slice readout. Further information about each of these components are described in more detail in Chapter 2 and Chapter 3.

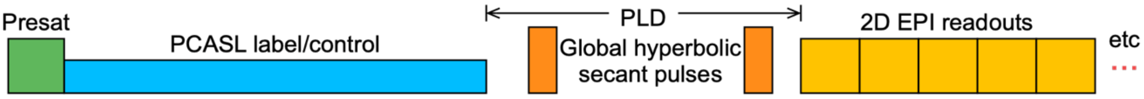

Figure 4.1: A simplified schematic of the multi-slice 2D-EPI PCASL sequence used in this work. The WET presaturation module ensures signal in the imaging region is 'reset' at the beginning of each TR and forms the first part of the BGS. The global HS pulses are timed relative to the presaturation module to null $T_{1,opt}$ and $2 \cdot T_{1,opt}$ at the start of the EPI readout. However, they are restricted to only play out during the PLD. The 2D-EPI readout is as described in Chapter 2.

Each TR begins with the WET presaturation module, to 'reset' the magnetisation. It forms an important part of the BGS by ensuring that each TR starts with the longitudinal component of the static tissue signal recovering from 0 (assuming perfect saturation). Furthermore, any remaining labelled blood signal from the proceeding TR is also saturated.

The WET presaturation module is immediately followed by the start of the PCASL pulse train, which continues for the duration of the specified label duration and ends with a gradient spoiler along the inferior-superior direction.

The PLD is the time between the end of the PCASL pulse train and the centre of the excitation pulse for the first acquired slice. Multiple PLDs can be defined in the sequence user interface and are run sequentially: the label and control conditions for each

PLD are run before the sequence moves onto the next PLD. If multiple averages are to be acquired, the sequence will first loop through all of the PLDs sequentially before repeating the loop to acquire the next average for each PLD. Other sequence components can run during the PLD, such as the two HS inversion pulses used for BGS and the spectrally-selective fat saturation pulse used immediately before each EPI excitation pulse.

The two BGS HS inversion pulses were timed using the formula in Equation 3.4 to null tissues with T_1 equal to $T_{1,opt}$ and $2 \cdot T_{1,opt}$ at the centre of the first EPI excitation pulse. The HS shaping parameters were $\mu = 24$ and $\beta = 800$ rad/s. The peak B_1^+ was set to the maximum available level.

The gradient echo 2D-EPI readout is as described in Chapter 2. A spectrally-selective fat saturation pulse and associated spoiler gradients are followed by a sinc excitation pulse and EPI readout. The slices are acquired in ascending order.

4.3 Theory

4.3.1 Cramér-Rao Lower Bound

The Cramér-Rao lower bound[6,7] provides a mathematical expression for the lower bound on the variance (uncertainty) of parameters estimated from a set of data. More specifically, it states that the inverse of the Fisher information matrix (FIM) is the lower bound of the covariance matrix for deterministic parameters when using a maximum likelihood estimator:

$$\mathbf{F}^{-1} \leq cov(\vec{\theta}) \tag{4.1}$$

where $\mathbf{F}^{-1}$ is the inverse of the FIM and $cov(\vec{\theta})$ is the covariance matrix for a column vector of model parameters, $\vec{\theta}$.

If a white Gaussian noise model is assumed, then the FIM for an ASL experiment can be expressed as:

$$\mathrm{F}\left(\vec{t};\vec{\theta}\right)_{jk}=\frac{A}{\sigma^2}\sum_{i=1}^{N}\left[\frac{\partial\Delta M\left(t_i;\vec{\theta};\rho\right)}{\partial\theta_j}\ \frac{\partial\Delta M\left(t_i;\vec{\theta};\rho\right)}{\partial\theta_k}\right] \quad (4.2)$$

where N is the number of acquisitions, A is the integer number of averages for each acquisition achievable in a given scan time, σ^2 is the normally distributed noise variance of the acquired images, ΔM is the ASL difference signal model, $\vec{t}$ is a vector containing the protocol timings, $\vec{\theta}$ is a vector of the model parameters to be inferred, ρ are the remaining model parameters which are assumed fixed, and $\frac{\partial\Delta M}{\partial\vec{\theta}}$ is the jacobian matrix of the signal model with respect to the unknown parameters, otherwise know as the sensitivity functions. Since it is common to assume known values for all model parameters except CBF and ATT, $\vec{\theta}=(f,\Delta t)^{\mathrm{T}}$, where f is the CBF in units of s^{-1}, Δt is the ATT in units of s, and $^{\mathrm{T}}$ is the transpose. Explicit reference to ρ is dropped for simplicity.

Note that identical noise variance across acquisitions is assumed here, unlike Xie et al.,[8] where the authors found an empirical noise model across acquisitions. This simplification means that the noise variance is inversely proportional to the number of averages, for the assumed case of white Gaussian noise. An optimal protocol can then be generated within a fixed scan time by calculating the number of averages possible for a given design within this time. A known value for the noise variance is not required for optimisation since it is only a scaling factor, though knowledge of this value will generate CRLB variances in physiological units. This also means that the optimisation will be equally valid for all voxels in images with spatially varying SNR, such as when a head array coil is used for signal detection.

4.3.2 Optimality Criteria

The most widely used optimality criterion is the determinant of the covariance matrix. Since the determinant of a matrix is equal to the product of its eigenvalues, it is proportional to the volume of the confidence ellipsoid of the estimated parameters. By minimising the determinant, the volume of this confidence ellipsoid is minimised. The D-optimality criterion can be defined as:

$$\underset{\vec{t}}{\operatorname{argmin}}\, \phi(\vec{t};\vec{\theta})_{\mathrm{D-optimal}} = \underset{\vec{t}}{\operatorname{argmin}} \det\left(cov(\vec{t};\vec{\theta})\right) = \underset{\vec{t}}{\operatorname{argmin}} \frac{1}{\det\left(\mathbf{F}(\vec{t};\vec{\theta})\right)} \quad (4.3)$$

because $1/\det(\mathbf{A}) = \det(\mathbf{B})$ when $\mathbf{A} = \mathbf{B}^{-1}$, where $\phi(\vec{t};\vec{\theta})$ is the cost function to be minimised for a set of parameters $\vec{\theta}$.

The second optimality criterion explored in this study is the minimisation of only the CBF variance, since CBF is often the main parameter of interest, with ATT being a confounding parameter. Conceptually, this objective function aims to minimise CBF variance, including minimisation of the impact of ATT on CBF measurements. This criterion, referred to as L-optimality, is given by:

$$\underset{\vec{t}}{\operatorname{argmin}}\, \phi(\vec{t};\vec{\theta})_{\mathrm{L-optimal}} = \underset{\vec{t}}{\operatorname{argmin}}\, \mathrm{trace}\left(\mathbf{W}\,\mathbf{F}(\vec{t};\vec{\theta})^{-1}\right) \quad (4.4)$$

where **W** is a symmetric non-negative definite matrix and has the same dimensions as **F**.[14] In this case, W_{11} is the only non-zero element, to select only the CBF variance.

4.3.3 (P)CASL Signal Model

The CASL general kinetic model from Buxton et al.[2] (Equation 3.2) is used in this work and it is repeated here for the benefit of the reader:

$$\begin{aligned}
\Delta M(t) &= 0 && 0 < t < \Delta t \\
&= 2M_{0B} f T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}}\left(1 - e^{\frac{-(t-\Delta t)}{T_1'}}\right) && \Delta t < t < \tau + \Delta t \\
&= 2M_{0B} f T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}} e^{\frac{-(t-\tau-\Delta t)}{T_1'}}\left(1 - e^{\frac{-\tau}{T_1'}}\right) && \tau + \Delta t < t
\end{aligned} \tag{4.5}$$

where $1/T_1' = 1/T_{1t} + f/\lambda$, ΔM is the ASL signal difference between label and control images, t is the time from the start of labelling (s), M_{0B} is the equilibrium magnetisation of arterial blood, f is the CBF (s^{-1}), λ is the equilibrium brain/blood water partition coefficient (mL g^{-1}), α is the labelling efficiency, Δt is the arterial transit time (s), τ is the labelling duration (s), and T_{1b} and T_{1t} are the longitudinal relaxation time constants for arterial blood and tissue (s), respectively. Note that any appropriate analytical signal model may be used instead.

4.3.4 Sensitivity Functions

To simplify the CBF sensitivity function, it is assumed that the apparent tissue relaxation time, T_1', is fixed with respect to f (fixed outflow), as proposed by Xie et al.[8] Using this assumption, with outflow set to 50 mL/100g/min, leads to a maximum error of ~1.78% in the simplified CBF sensitivity function, when the true value of f is in the range 0 – 100 mL/100g/min. The sensitivity functions are then:

$$\begin{aligned}
\frac{\partial \Delta M(t)}{\partial f} &= 0 && 0 < t < \Delta t \\
&= 2M_{0B} T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}}\left(1 - e^{\frac{-(t-\Delta t)}{T_1'}}\right) && \Delta t < t < \tau + \Delta t \\
&= 2M_{0B} T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}} e^{\frac{-(t-\tau-\Delta t)}{T_1'}}\left(1 - e^{\frac{-\tau}{T_1'}}\right) && \tau + \Delta t < t
\end{aligned} \tag{4.6}$$

$$
\begin{aligned}
\frac{\partial \Delta M(t)}{\partial \Delta t} &= 0 && 0 < t < \Delta t \\
&= 2M_{0B} f T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}} \left[-\frac{1}{T_{1b}} - e^{\frac{-(t-\Delta t)}{T_1'}} \left(\frac{1}{T_1'} - \frac{1}{T_{1b}} \right) \right] && \Delta t < t < \tau + \Delta t \qquad (4.7) \\
&= 2M_{0B} f T_1' \alpha e^{\frac{-\Delta t}{T_{1b}}} e^{\frac{-(t-\tau-\Delta t)}{T_1'}} \left(1 - e^{\frac{-\tau}{T_1'}} \right) \left(\frac{1}{T_1'} - \frac{1}{T_{1b}} \right) && \tau + \Delta t < t
\end{aligned}
$$

Shown in Figure 4.2, they demonstrate that the model is most sensitive to CBF when the signal is maximised ($t = \tau + \Delta t$) and most sensitive to ATT during inflow ($\Delta t < t < \tau + \Delta t$). The equivalent sensitivity functions for PASL may be found in reference [8]. The CBF sensitivity function where T_1' is not constant with respect to *f* is given in Appendix A.

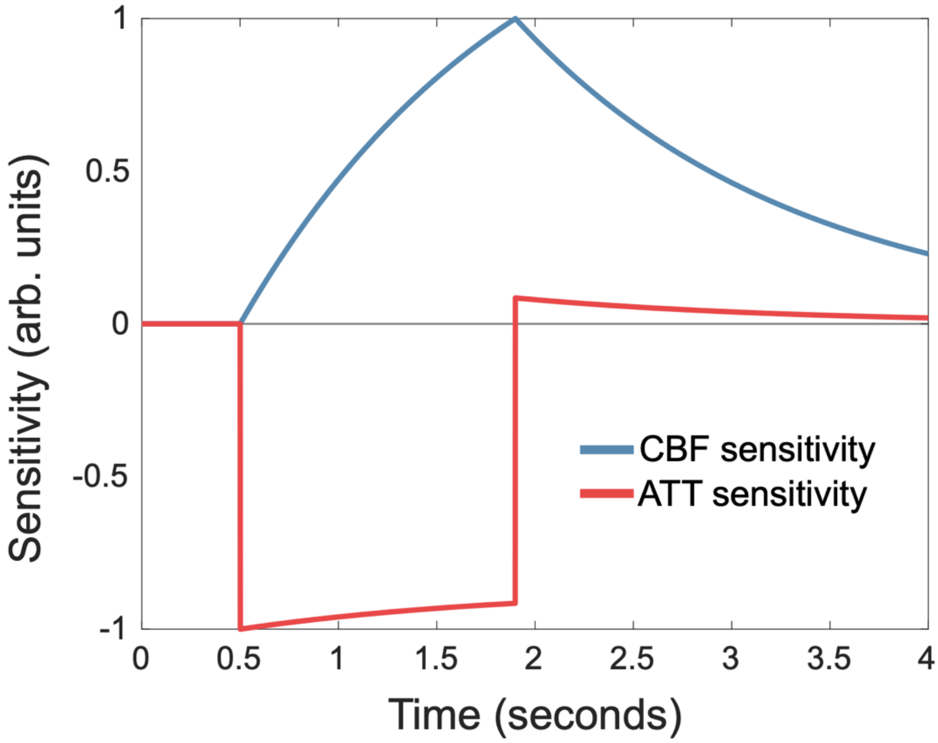

Figure 4.2: Example of the normalised sensitivity functions for CBF (blue) and ATT (red). Parameters used: labelling duration = 1.4 s; CBF = 50 mL/100g/min; ATT = 0.5 s; T_{1b} = 1.65 s; T_{1t} = 1.445 s; lambda = 0.9.

4.3.5 Prior Probability Distributions

The sensitivity functions depend on the unknown parameters being estimated: CBF and ATT. A protocol which minimises the CRLB at a single set of CBF and ATT values (locally optimal), will be a poor choice at values far from these.[14,15] *A priori* information of the likely range of parameter values is therefore required, in order to minimise the estimator variance over them.

The method proposed by Gilmour et al.,[14] referred to as average or pseudo-Bayesian optimal design, is used in this work. Here, the optimality criterion, $\phi(\vec{t};\vec{\theta})$, is averaged across a prior probability distribution, $p(\vec{\theta})$, and minimised. Therefore, the optimisation problem is:

$$\underset{\vec{t}}{\operatorname{argmin}}\,\Psi(\vec{t};\vec{\theta}) = \underset{\vec{t}}{\operatorname{argmin}}\,\mathrm{E}\left[\phi(\vec{t};\vec{\theta})\right] = \underset{\vec{t}}{\operatorname{argmin}}\int_{\boldsymbol{\theta}} \phi(\vec{t};\vec{\theta})p(\vec{\theta})d\vec{\theta} \qquad (4.8)$$

where $\Psi(\vec{t};\vec{\theta})$ is the cost function to be minimised for a range of ATTs and E is the expectation operator.

This integral is difficult to evaluate analytically, so it is approximated with a number of equally spaced samples from the prior distribution:

$$\widetilde{\Psi}(\vec{t}) = \frac{1}{r}\sum_{l=1}^{r} \phi(\vec{t};\vec{\theta}=\vec{\theta}_l)p(\vec{\theta}=\vec{\theta}_l) \simeq \Psi(\vec{t};\vec{\theta}) \qquad (4.9)$$

where $\vec{\theta}_l$ is a sample of parameters from the prior distribution, and r is the number of samples.

In previous work, both CBF and ATT prior probability distributions were used.[8–10] However, using Equations 4.6 and 4.7, it can be seen that the CBF sensitivity function is independent of CBF, while the ATT sensitivity function is linearly proportional to CBF. This means that $\phi_{\mathrm{L-optimal}}$, which in the current case is simply the CBF variance, does not depend on CBF and $\phi_{\mathrm{D-optimal}}$ is inversely proportional to the square of the CBF, due to the same dependence of the ATT variance on CBF. Therefore, the optimal protocol generated using each optimality criterion will be identical for any value of CBF. Consequently, the optimal design will only depend on the ATT distribution, which greatly reduces the required prior knowledge and the number of calculations involved in the optimisation, since a point prior may be used for CBF.

The relationships of the CBF and ATT variances with the absolute CBF value tells us that as CBF increases, the ATT estimation variance (the inverse of ATT accuracy)

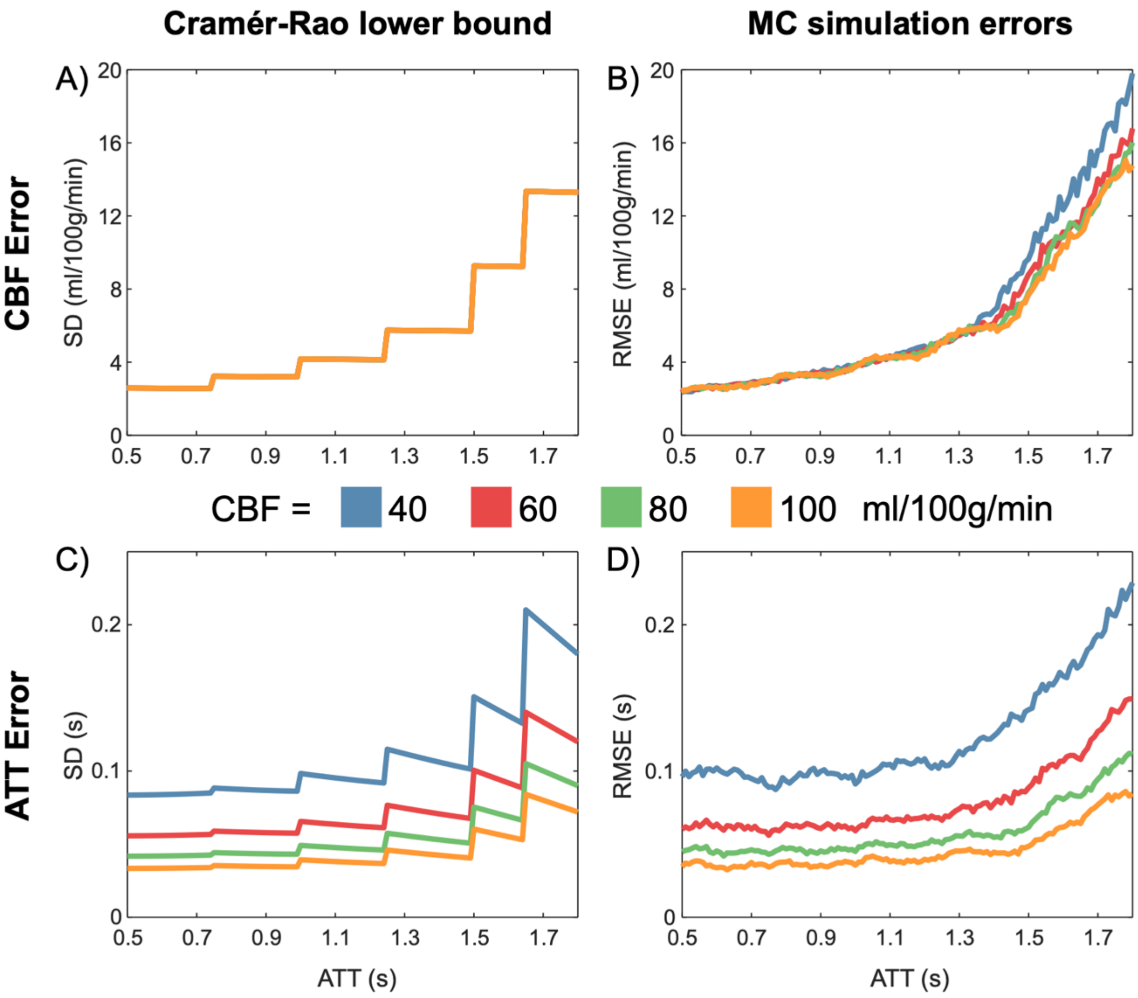

Figure 4.3: The effect of true CBF on the estimation errors of CBF (top row) and ATT (bottom row). Both the predicted Cramér-Rao lower bound (A,B) and RMSEs for Monte Carlo simulations (B,D) are shown. These results demonstrate that the CBF estimation errors do not vary greatly with CBF, while the ATT errors are inversely proportional to the square of the CBF. The Monte Carlo simulations broadly agree with the theoretical prediction and were performed as described in the Methods section using the reference multi-PLD protocol for one slice.

decreases with $1/f^2$, but the CBF variance remains unchanged. This is demonstrated in Figure 4.3.

4.4 Methods

All optimisations, simulations and analysis were performed using MATLAB (The MathWorks, Inc., Natick, MA, USA).

4.4.1 PLD Optimisation

The optimisation theory introduced above is general and can be used to optimise any type of ASL experiment for parameter inference. Here, specific methods are described that were employed for optimising a 2D multi-slice, multi-PLD, fixed label duration PCASL experiment, to find the optimal set of PLDs for CBF and ATT estimation.

Two protocols, referred to as CBF-ATT$_{opt}$ and CBF$_{opt}$, were generated using the optimality criteria in Equations 4.3 and 4.4 (D-optimal and L-optimal), respectively. The optimisation algorithm developed for this study uses an iterative exchange method similar to Xie et al.[8] and aims to find the optimal set of PLDs which minimises Equation 4.9. The pseudocode outline is shown in Figure 4.4. It takes as inputs the number of PLDs, ATT prior probability distribution, available scan duration, number of slices, slice duration and the remaining constants in Equations 4.6 and 4.7. The PLDs were initialised to be equally distributed in the interval $0.25 \leq \text{PLD} \leq 1.5$ s. The algorithm loops through the PLDs, selecting the optimal PLD at each step from a list of possible values. The PLD list for t_i (the ith PLD) was restricted to the interval $0.2 \text{ s} \leq t_{i-1} \leq t_i \leq t_{i+1} \leq 3$ s, using 25 ms increments. This reduces the size of the search without restricting the algorithm from finding the approximately optimal solution. The minimum PLD of 0.2 s was to provide enough time for the BGS inversion pulses used in vivo. The longest PLD chosen by the algorithm during testing was always at most equal to the longest ATT in the ATT prior

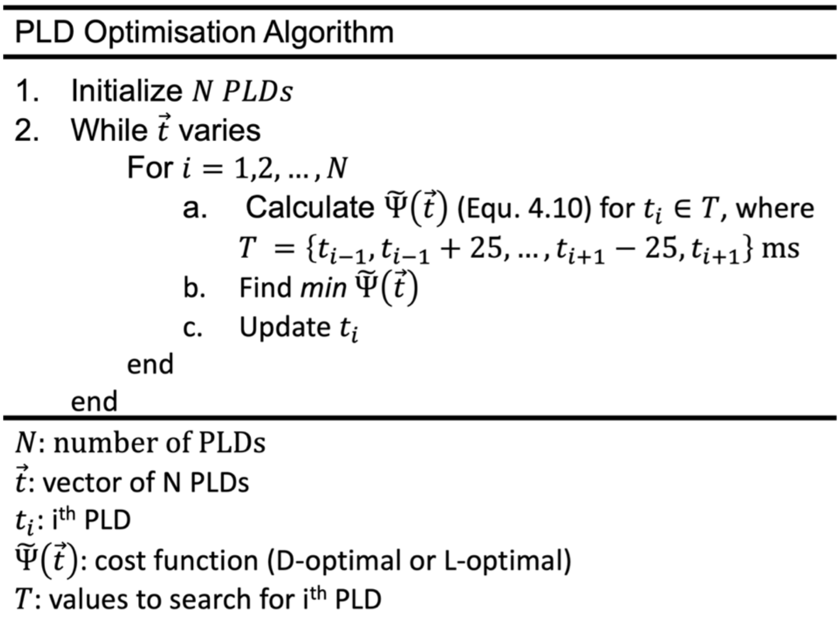

```
PLD Optimisation Algorithm
1. Initialize $N$ $PLDs$
2. While $\vec{t}$ varies
       For $i = 1,2,\dots,N$
           a. Calculate $\widetilde{\Psi}(\vec{t})$ (Equ. 4.10) for $t_i \in T$, where
              $T = \{t_{i-1}, t_{i-1}+25, \dots, t_{i+1}-25, t_{i+1}\}$ ms
           b. Find $min\ \widetilde{\Psi}(\vec{t})$
           c. Update $t_i$
       end
   end
$N$: number of PLDs
$\vec{t}$: vector of N PLDs
$t_i$: ith PLD
$\widetilde{\Psi}(\vec{t})$: cost function (D-optimal or L-optimal)
$T$: values to search for ith PLD
```

Figure 4.4: Pseudocode outlining the PLD optimisation algorithm used in this study. The optimal number of PLDs, N, can be found by running this algorithm for a range of N and finding the design which minimises Equation 4.9.

probability distribution. The PLD upper bound of 3 s will, therefore, not restrict the choice of PLDs in this study (see below), but will reduce the PLD search space.

An ATT range of $0.5 \leq \text{ATT} \leq 1.8$ s was assumed for GM in healthy volunteers.[1,16] A uniform prior probability distribution for this range was used to ensure equal weighting of all ATTs. To avoid edge effects, the ATT distribution was extended on either side by 0.3 s with linearly decreasing probability. The CBF point prior was set to 50 mL/100g/min. The readout duration was assumed to be 1.275 s to give realistic scan times for a whole-brain 24 slice readout, which is the number of slices previously used with this sequence with the reference multi-PLD protocol.[13] However, only 5 slices were optimised for, simulated, and acquired in vivo for this study (see below). The total allowable scan time was set to 5 minutes. Variable repetition time (TR) was assumed such that a short PLD is acquired with a correspondingly short TR to minimise dead time in the sequence.[17] The number of averages in Equation 4.2 was calculated as:

$$A = \left\lfloor \frac{S}{2\cdot\sum_{i=1}^{N}(\tau+t_i+R)} \right\rfloor, \tag{4.10}$$

where S is the total scan time available, R is the readout duration, the factor of 2 accounts for the doubling of the acquisitions for label and control images, and $\lfloor\ \rfloor$ is the floor function. The noise SD used in the optimisation and subsequent simulations was derived from preliminary in vivo data. The noise SD for each control or label image was 0.001545 (normalised by M_{0B}), so the noise SD for the difference images was $0.001545 \cdot \sqrt{2}$. All other model parameters are given in Table *4.1*.

Parameter	***Value***
General	
Label duration (τ)	1.4 s
T_1 of blood (T_{1b})	1.65 s[18]
T_1 of tissue (T_{1t})	1.445 s[19]
Labelling efficiency (α)	0.85[12]
Brain/blood water partition coefficient (λ)	0.9 mL/g[20]
Slice duration	53.125 ms
Number of slices	5
Optimisation	
Fixed CBF in apparent T_1 (T_1')	50 mL/100g/min
Readout duration	1.275 s
In vivo experiments	
RF labelling pulse duration	600 µs duration (Gaussian)
RF labelling pulse separation	1 ms
RF labelling flip-angle	20°
Mean labelling gradient	0.8 mT/m
Gradient during labelling pulses	6 mT/m
Nominal voxel size	3.4 x 3.4 x 5 mm
Matrix size	64 x 64
Bandwidth	2004 Hz/pixel
Partial Fourier factor	6/8
Echo time (TE)	21 ms
Velocity encoding (VENC)	4 cm/s[1]

Table 4.1*: Parameters used for optimisations, simulations and in vivo experiments*

For this proof of principle study, the protocols were optimised so that the mean of Equation 4.9 across 5 slices was minimised. The number of slices was kept small in order to minimise the range of PLDs and BGS performance across slices. For a given slice, if all effective PLDs exceed the ATT then the FIM becomes severely ill-conditioned due to a lack of ATT information. This is an inherent problem for CASL methods since the inflow of signal is more difficult to sample than for PASL. To avoid ill-conditioned matrices, the ATT probability distribution was truncated for each slice based on the shortest attainable PLD. The sum in Equation 4.9 was then additionally weighted by the number of contributing slices. Analytical inversion of the FIM was performed to improve algorithm speed.

The two optimised protocols were compared against a previously used evenly distributed multi-PLD protocol[13] (referred to as the reference multi-PLD protocol) and a single-PLD protocol using the recommended 1.8 s PLD.[1] All protocols used a 1.4 s labelling duration in order to match the labelling duration of the reference multi-PLD protocol. The timings for these protocols are given in Table 4.2 and shown in Figure 4.7. CRLB values were generated for all four protocols across the ATT range. For the single-PLD protocol, the CRLBs were estimated using $RMSE = \sqrt{variance^2 + bias^2}$, where the bias is the systematic error due to the assumed ATT.

4.4.2 Simulation Experiments

To validate the theoretical effects of the optimisation, Monte Carlo simulations were performed for all four ASL protocols. 2000 sets of data were generated for each ATT value in the range $0.5 \leq \text{ATT} \leq 1.8$ s, at 0.01 s intervals, using Equation 4.5 (the standard CASL model) with the parameters in Table *4.1*, for 5 slices. White Gaussian noise was added to label and control data before pairwise subtraction.

CBF and ATT were estimated from the data using Equation 4.5 and MATLAB's non-linear least squares (NLLS) function, 'fmincon'. Fitting was initialised using a coarse grid search and bounded by $0 \leq CBF \leq 200$ mL/100g/min and $0 \leq ATT \leq 2.5$ s. For the single-PLD protocol, CBF was also estimated using Equation 4.5, but with a fixed $ATT = 1.25$ s in the fitting process, since this resulted in the smallest theoretical error across the ATT range 0.5 - 1.8 s and uses the same assumptions as the multi-PLD data (see Figure 4.5). Using the recommended method for CBF estimation from single-PLD data[1] would result in a systematic underestimation of CBF relative to Equation 4.5 due to the assumption that $T_{1t} = T_1' = T_{1b}$. The RMSEs of the estimates, relative to the true values, were calculated for comparison. The RMSE represents a measure of accuracy and is a combination of both systematic bias and variance due to noise.

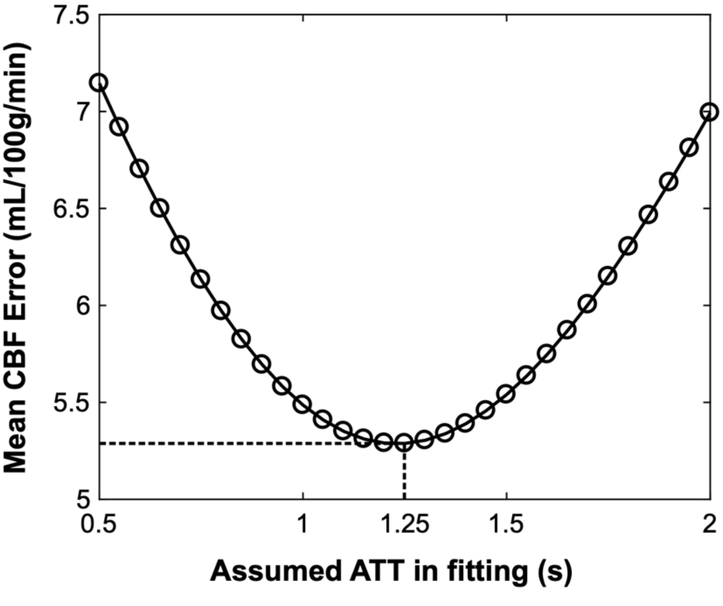

Figure 4.5: Finding the lowest CBF Error for the single-TI protocol for a range of assumed ATTs in the fitting process. The graph shows the theoretical root-mean-squared-error (circles are the evaluated assumed ATTs: 0.5 - 2 s, sampled every 50 ms). Data was generated for a grid of parameters: CBF 10 - 100 mL/100g/min sampled every 5 mL/100g/min, ATT 0.5 - 2 s sampled every 50 ms, and the parameters in Table 4.1.

4.4.3 In Vivo Experiments

4.4.3.1 Acquisition

In vivo data was acquired for the four protocols given in Table 4.2 to confirm that the theoretical benefits of the optimisation were realised in practice. Seven healthy volunteers (3 female, 23 to 27 years old) were scanned under a technical development protocol, agreed with local ethics and institutional committees. All in vivo data was acquired on a 3T Prisma system (Siemens Healthcare, Erlangen, Germany) with a 32-channel receive-only head coil. Additional scans performed were: a 3D single-slab time-of-flight angiography sequence (resolution $0.31 \times 0.31 \times 1.3$ mm^3), for placement of the PCASL labelling plane; a T_1 weighted structural image (resolution $1.5 \times 1.5 \times 1.5$ mm^3), for registration and tissue segmentation; and a B_0 field map (resolution $2 \times 2 \times 2$ mm^3, TE1 4.92 ms, TE2 7.38 ms), for distortion correction of the ASL data.

For the PCASL data, imaging parameters were: single-shot gradient-echo EPI readout, 90° flip-angle excitation, 5 transverse slices positioned to bisect the thalamus, ascending slice order, no acceleration, and spectrally-selective fat saturation (see Table *4.1* for the remaining imaging parameters). The labelling plane was placed in transverse orientation at the middle of the V3 section of the vertebral arteries, where the vertebral and internal carotid arteries are approximately parallel to each other and perpendicular to the transverse plane. PCASL labelling was achieved with: 600 μs duration Gaussian RF pulses, 1 ms spacing, 20° flip angle and 1.4 s labelling duration. To reduce the impact of macrovascular signal, flow signal crushing was applied in the inferior-superior direction with a velocity encoding cutoff of 4 cm/s.[1] BGS was achieved with a WET pre-saturation module[21] and two optimally timed global hyperbolic secant inversion pulses applied after the PCASL labelling, as in previous studies.[13,22] Further ASL scan parameters are given in

Table *4.1*. The spatial placement of the presaturation slab, PCASL labelling plane, and EPI slices are shown in Figure 4.6.

A fully relaxed M_0 image was acquired with identical acquisition parameters to the PCASL data but without BGS and PCASL labelling. This was used as the target for motion correction of the PCASL data and voxelwise calibration. A whole-brain M_0 image was also acquired, with matched resolution but with 24 slices, to aid registration between the T_1-weighted structural and partial brain ASL images.

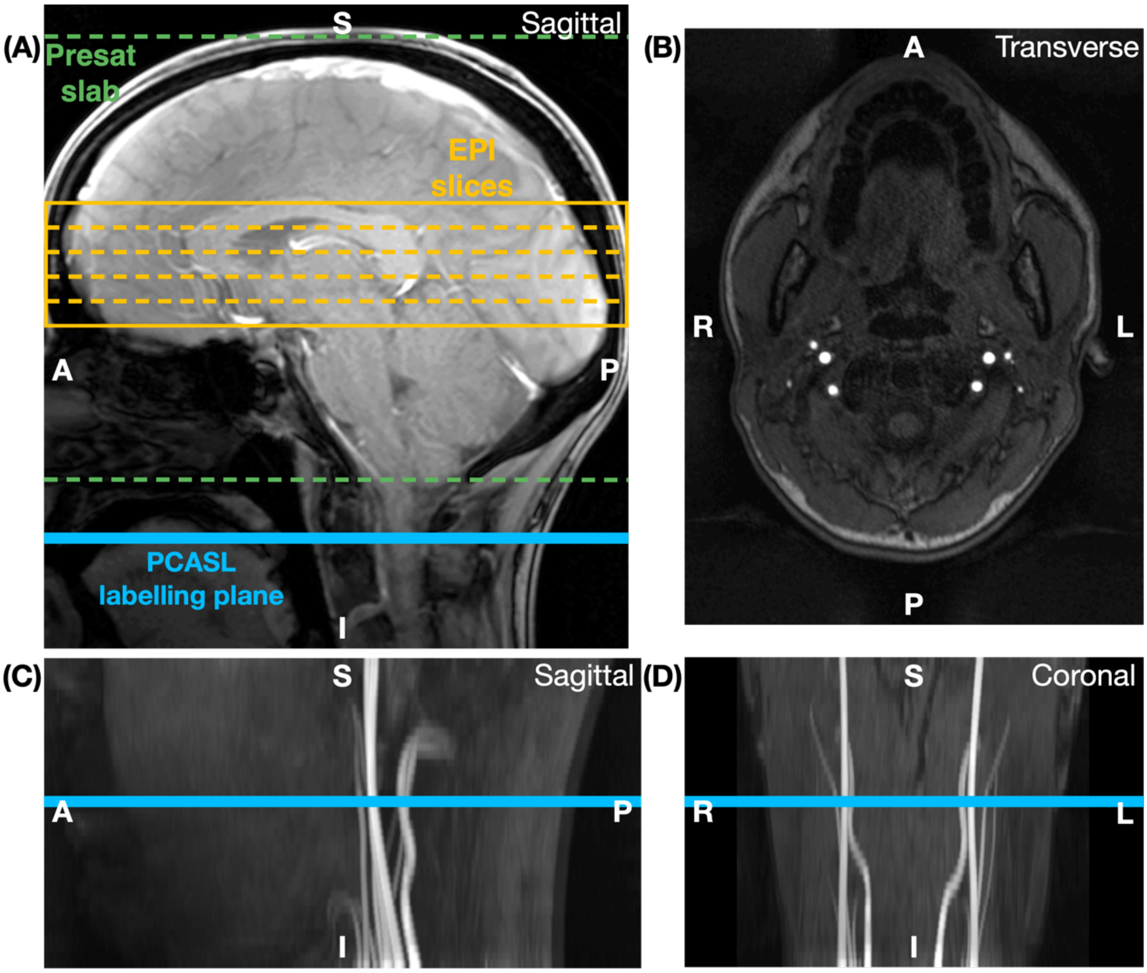

Figure 4.6: The in vivo experimental setup. (A) The approximate placement of the presaturation slab, PCASL labelling plane, and the EPI slices overlaid on a sagittal localiser image. The PCASL labelling plane was placed using a TOF image of the feeding arteries as described in the text. (B) A transverse slice of the TOF image showing the slice where the centre of the PCASL labelling plane was placed. (C) and (D) show sagittal and coronal maximum intensity projections (MIPs) of the TOF image demonstrating the approximate placement of the PCASL labelling plane.

To mitigate effects caused by changing physiology during the scan session, the PLDs from each of the four protocols were interleaved and distributed throughout the session. In preliminary phantom scans, we noticed that large variations in the BGS effectiveness, such as when a short PLD followed a long PLD or vice versa, caused the online EPI B_0 drift correction to create artificial sub-voxel movements in the phase-encode direction. In order to minimise this effect, the PLDs were not randomly ordered but were distributed across the scan session with the PLD duration gradually increasing and decreasing in 4 cycles. This resulted in a well-distributed coverage across the total ASL scan duration for each protocol, but also maintained gradual variations in the PLDs and resulting BGS performance thus minimising erroneous shifting in the image. Four extra PLDs were also acquired to ensure a maximum spacing of 75 ms across the acquired PLD range. This was used for the ground truth data (described below) to ensure a high temporal sampling of the data without any significant gaps.

Protocol	***Post-labelling delays (s)***	***# PLDs***	***# Averages***
Single-PLD	1.8	1	33
Reference multi-PLD	0.25, 0.5, 0.75, 1, 1.25, 1.5	6	7
CBF-ATT$_{opt}$	0.2, 0.2, 0.225, 0.3, 0.375, 0.45, 0.5, 0.55, 0.6, 0.6, 0.625, 0.625, 0.65, 0.65, 0.675, 0.675, 0.7, 0.7, 0.7, 0.7, 1.25, 1.275, 1.3, 1.35, 1.375, 1.4, 1.425, 1.425, 1.475, 1.5, 1.675, 1.75, 1.8, 1.825, 1.85, 1.875, 1.9, 1.925, 1.95, 1.975	40	1
CBF$_{opt}$	0.2, 0.7, 0.825, 1, 1.125, 1.25, 1.325, 1.4, 1.475, 1.55, 1.625, 1.675, 1.7, 1.725, 1.75, 1.775, 1.8, 1.825, 1.85, 1.85, 1.875, 1.875, 1.9, 1.925, 1.925, 1.95, 1.975, 1.975, 2, 2.025, 2.025, 2.05, 2.075, 2.075	34	1

Table 4.2: Protocol timings

4.4.3.2 Preprocessing

GM masks were generated from the T_1 structural image using FSL's FAST[23] tool. The ASL images were motion-corrected using FSL's FLIRT[24] with three degrees-of-freedom (two within-plane translations and one within-plane rotation). Distortion correction was performed using the B_0 field map and FSL's FUGUE tool. Further registrations were performed to transform the GM mask into the ASL native space, thresholding it at 50%. Prior to voxelwise calibration, the M_0 image was smoothed with a Gaussian kernel ($\sigma =$ 2.5 mm), as recommended in the recent consensus paper.[1]

4.4.3.3 Analysis

The kinetic model was fit to the in vivo data exactly as described for the simulated data above. In order to evaluate the error associated with the CBF and ATT estimates from each protocol, ground truth estimates were generated by combining and fitting all unaveraged ASL data across all the protocols, so as to equally weight the data from each protocol. The CBF and ATT errors from each protocol were then calculated relative to these ground truth estimates. This was done due to the absence of true ground truth values for the in vivo data. Since the PLDs from each protocol were interleaved and distributed across the entire scan session, the ground truth CBF and ATT estimates represent a precise estimate of the average CBF and ATT across the scan session.

To ensure only well-fit ground truth GM estimates were used in further analysis, the following restrictions were imposed. The CRLB (Equations 4.1 and 4.2) was calculated voxelwise and used as an approximation of the ground-truth CBF and ATT variance (uncertainty). The noise was estimated by the summed squared residuals of each fit, normalised by the statistical degrees of freedom.[25] Voxels were excluded if the estimated ground-truth CBF or ATT SD were greater than 5 mL/100g/min or 0.1 s, respectively. The analysis was further restricted to voxels with ground truth ATT estimates in the range of interest, $0.5 \leq \text{ATT} \leq 1.8$ s. The RMSE of the estimates from each protocol, relative to the ground truth values, were calculated for comparison. A two-tailed non-parametric test (Wilcoxon signed-rank test) was used to test for significant differences across subjects ($p < 0.05$).

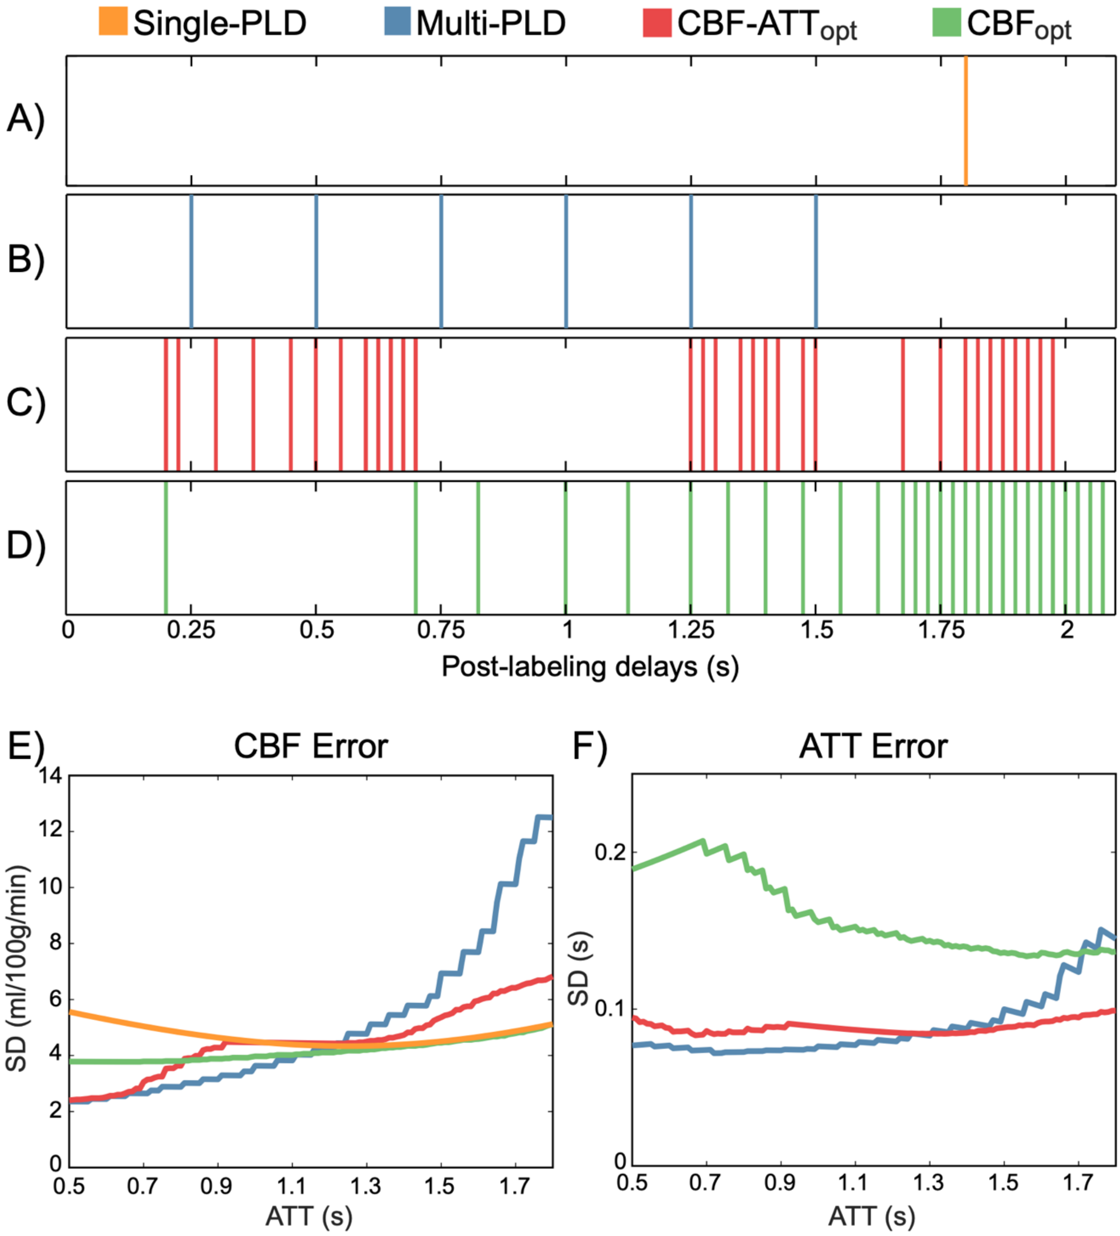

Figure 4.7: The PLDs (A,B,C,D) and the predicted CBF and ATT errors (Cramér-Rao lower bound (CRLB) SD) (E,F) for each of the protocols. The reference single-PLD protocol (A) has a fixed PLD at 1.8 s, the reference multi-PLD protocol (B) uses evenly distributed PLDs between 0.25 – 1.5 s, while the optimised protocols, CBF-ATT$_{opt}$ (C) and CBF$_{opt}$ (D), have more targeted PLDs. Repeated PLDs are not shown, but are listed in full in Table 4.2. The CRLB SDs for CBF and ATT demonstrate the impact that the choice of PLDs have on inference accuracy.

4.5 Results

4.5.1 PLD Optimisation

The PLDs for each protocol are given in Table 4.2 and shown in Figure 4.7. Both optimised schemes have very distinct sampling patterns with a high number of PLDs with few repeats. The CBF-ATT$_{opt}$ protocol has two main groups: one of short PLDs, sampling signal inflow, and a group of longer PLDs, sampling the signal peak and decay for most of the ATT distribution. The CBF$_{opt}$ protocol covers a similar range of PLDs, but has very few short PLDs and increased density of longer PLDs, leading to more sampling of the signal curve after complete bolus arrival.

The theoretical CBF and ATT errors (CRLBs) are also shown in Figure 4.7 and demonstrate the predicted improvement in accuracy of the optimised protocols across the ATT prior range. Both optimised protocols have less variable CRLB profiles than the

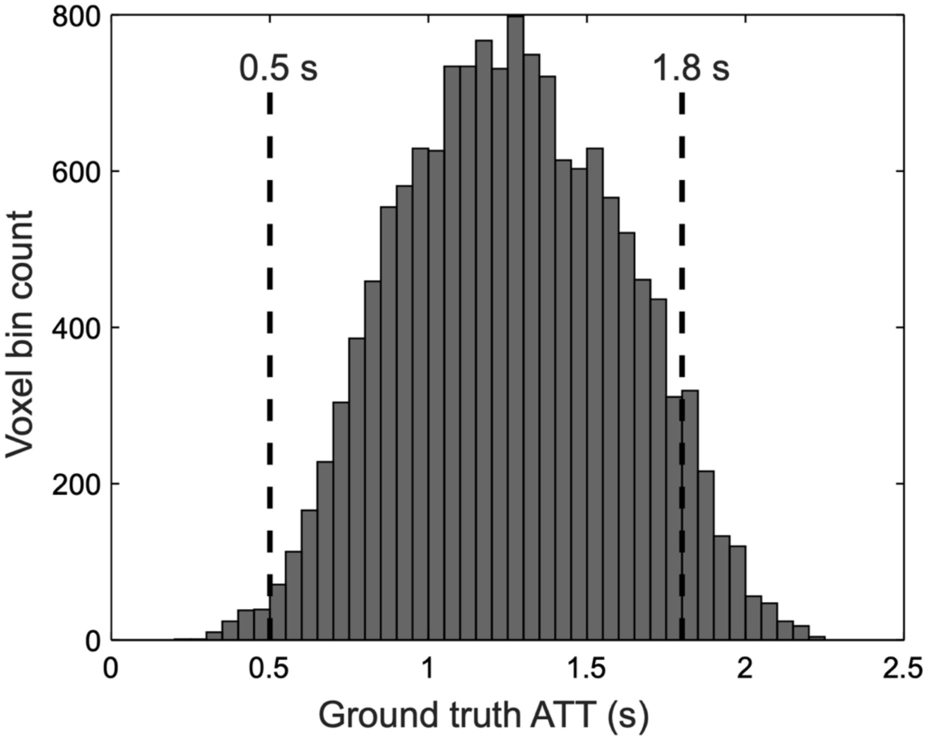

Figure 4.8: ***Histogram of the ground truth ATT estimates which had an estimated CBF and ATT maximum likelihood distribution SD less than 5 mL/100g/min and 0.1 s, respectively. The range of ATTs included in further analysis are shown by vertical dashed lines.***

reference multi-PLD protocol, while the $\mathrm{CBF_{opt}}$ protocol demonstrates the lowest average CBF error. The increase in the ATT CRLB of the $\mathrm{CBF_{opt}}$ protocol is not reflected in its CBF errors due to its designed insensitivity to ATT accuracy. The single-PLD protocol achieves its minimal error at $ATT = 1.25$ s but is less accurate than the $\mathrm{CBF_{opt}}$ protocol.

4.5.2 In Vivo Data Selection

Before presenting the MC simulation and in vivo results together, the results of the in vivo data selection are described. Across the 7 subjects, the GM masks contained 19732 voxels in total. Of these, 31.6% were excluded because they did not meet one or more of the restrictions imposed on the ground truth data: 26.3% did not meet the uncertainty restrictions and a further 5.3% were not in the specified ATT range. A histogram of the

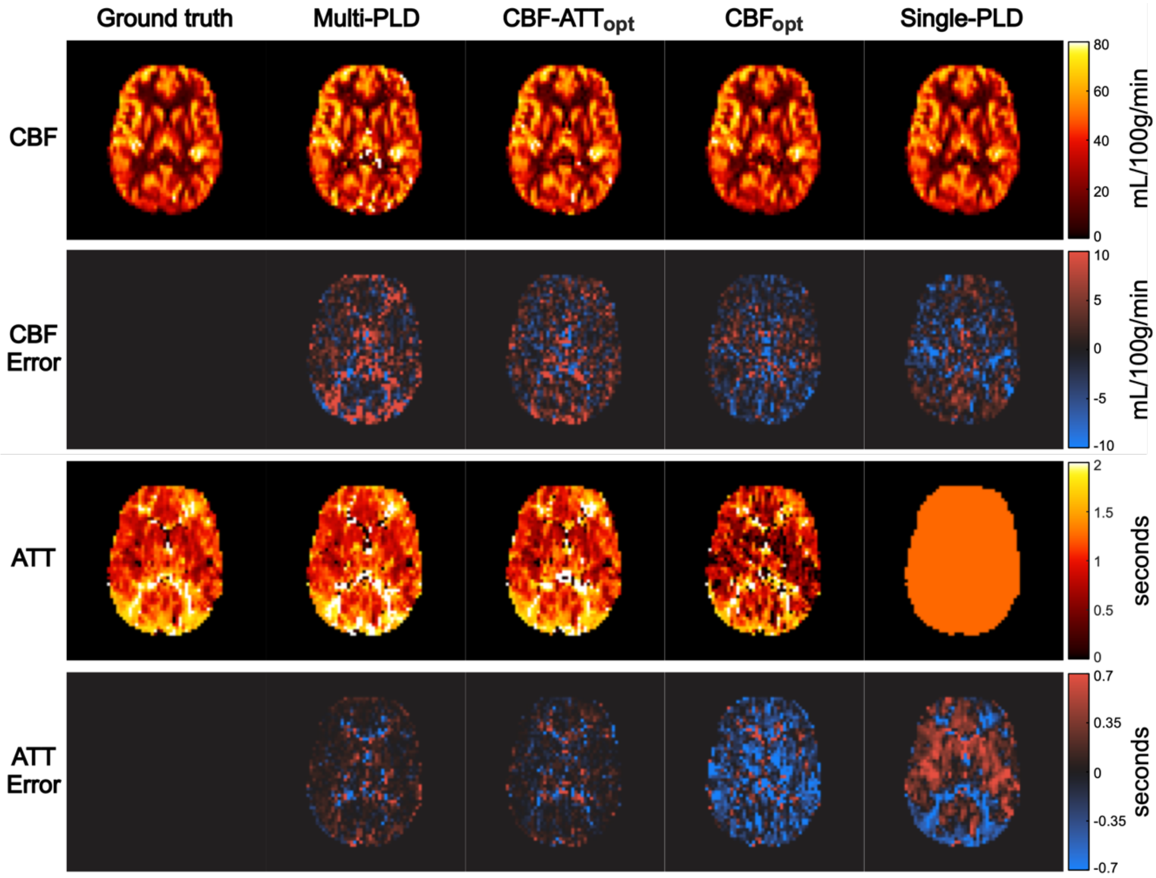

Figure 4.9: Representative CBF and ATT maps and for the ground truth estimates and the four tested protocols. The associated error maps relative to the ground truth data are also shown. The maps show an axial slice from a single subject. Note, there is no ATT map for the single-PLD protocol.

ATTs for the well-fit ground truth voxels is shown in Figure 4.8, demonstrating that the majority of well-fit voxels were within the targeted ATT range (94.7%). ATTs outside the range $0.5 \leq ATT \leq 1.8$ s were not used in further analysis.

4.5.3 Protocol Comparison

Representative CBF and ATT maps and error maps for one subject are shown in Figure 4.9. There is a reasonable qualitative agreement between the CBF maps. However, the reference multi-PLD protocol map exhibits noticeably more noise. Good agreement between the ATT maps for ground truth, reference multi-PLD and CBF-ATT$_{opt}$ protocols can be seen, while the CBF$_{opt}$ ATT map is clearly less accurate.

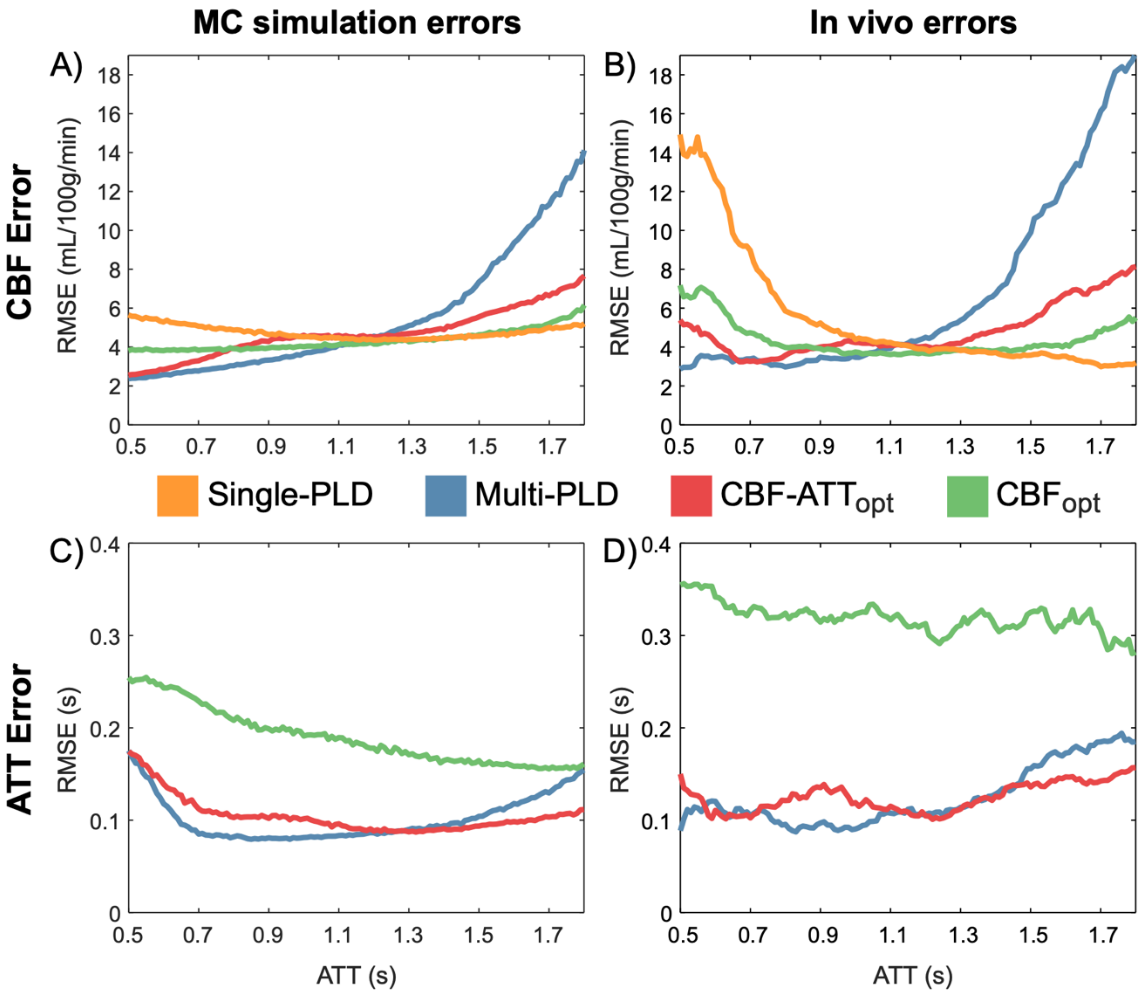

Figure 4.10: CBF (top row) and ATT (bottom row) RMSEs for the Monte Carlo simulations (A,C) and in vivo experiments (B,D). The in vivo data are the combined data across all 7 subjects, which has been smoothed using a sliding window mean (window width = 100 ms; increment = 10 ms).

Figure 4.10 shows the RMSE in CBF and ATT estimation for each protocol from the MC simulations and the in vivo data across the ATT range. For the in vivo data, all 7 subjects' data were combined using a sliding window to aid interpretability. There is a strong correspondence between the MC simulations and in vivo data, which both agree extremely well with the trends seen in the predicted errors (Figure 4.7), demonstrating the expected improvement in CBF and ATT estimation with the optimised protocols. As expected, the CBF_{opt} protocol resulted in high CBF accuracy while having less accurate ATT estimates.

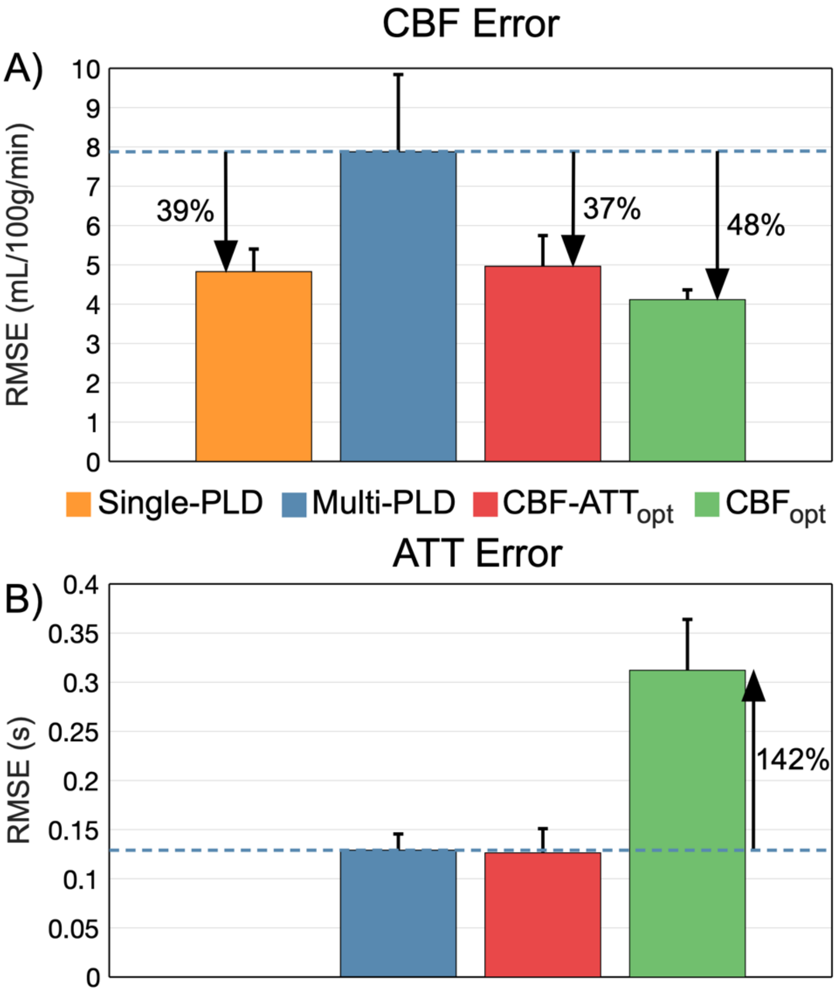

Figure 4.11: In vivo CBF (A) and ATT (B) RMSEs across subjects. The height of each bar graph is the mean RMSE across subjects, while the error bar shows the SD across subjects. The errors were checked for significant differences using a non-parametric paired test (Wilcoxon signed rank test), $p < 0.05$. All differences are significant, except between the single-PLD and CBF-ATT$_{opt}$ CBF errors and reference multi-PLD and CBF-ATT$_{opt}$ ATT errors.

The average in vivo RMSEs across subjects are shown in Figure 4.11. The CBF RMSEs were (mean ± SD mL/100g/min): 4.11 ± 0.25 (CBF$_{opt}$), 4.83 ± 0.57 (single-PLD), 4.97 ± 0.78 (CBF-ATT$_{opt}$), and 7.88 ± 1.97 (reference multi-PLD). Note that both optimal protocols had significantly reduced CBF errors compared to the reference multi-PLD protocol, while CBF$_{opt}$ also had significantly reduced CBF errors compared to the single-PLD protocol. The CBF$_{opt}$ ATT RMSE (0.31 ± 0.05 s) was significantly higher than the reference multi-PLD (0.13 ± 0.02 s) and CBF-ATT$_{opt}$ (0.13 ± 0.02 s) ATT RMSE. These results are very consistent with the relative performance of each protocol across subjects, as shown in Figure 4.12, demonstrating their robustness.

The RMSE is a useful metric for comparison since it combines both the bias and variance in the estimates. However, these measures can also be separately examined. This

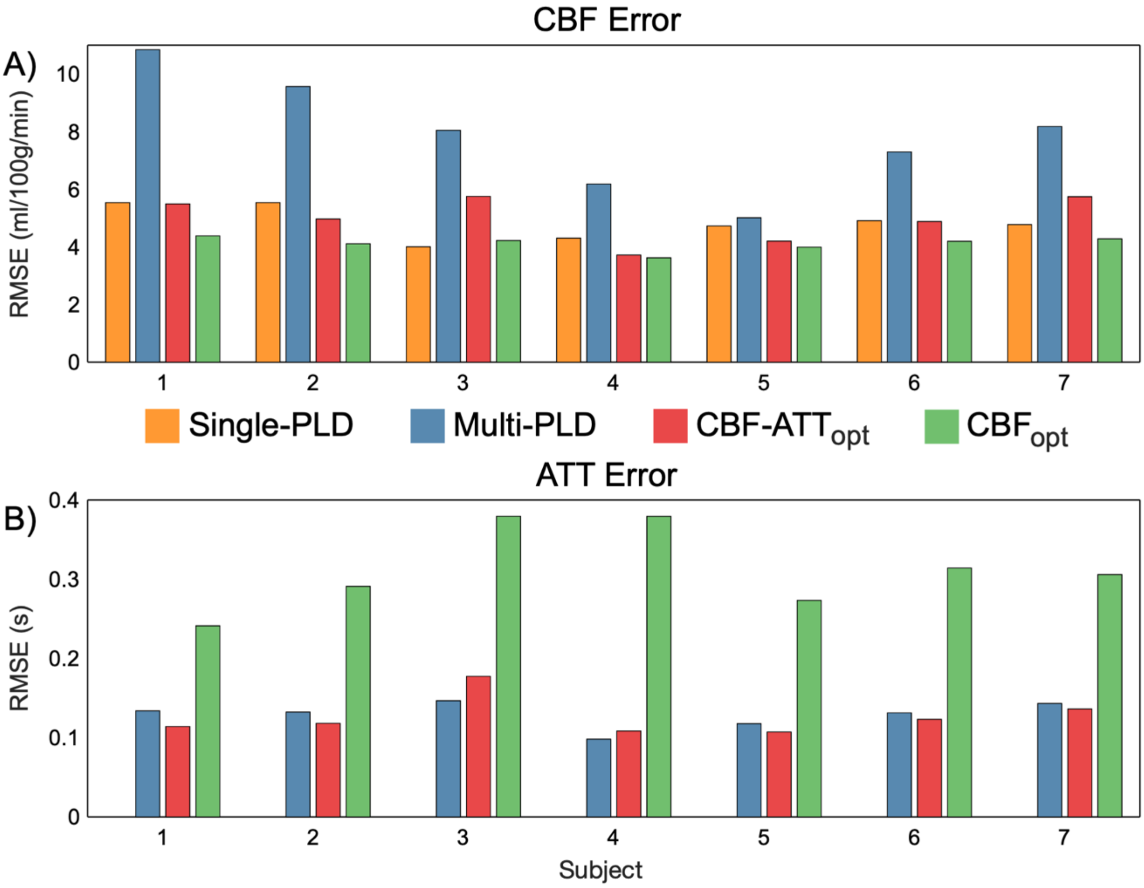

Figure 4.12: In vivo CBF (A) and ATT (B) RMSEs for each of the 7 subjects. The trends are extremely similar across the subjects, demonstrating the robustness of the optimisation.

breakdown is shown for the in vivo data in Figure 4.13. For the CBF estimates, the $\mathrm{CBF_{opt}}$ protocol had a significantly smaller bias and SD than both the reference single-PLD and multi-PLD protocols, while the $\mathrm{CBF\text{-}ATT_{opt}}$ protocol had a significantly lower SD than the reference multi-PLD protocol, but not the single-PLD protocol. For the ATT estimates, there were no significant differences in the biases (though a trend of underestimation can be seen for $\mathrm{CBF_{opt}}$) while $\mathrm{CBF_{opt}}$ had significantly higher SD than the reference multi-PLD and $\mathrm{CBF\text{-}ATT_{opt}}$ protocols.

4.6 Discussion

In this work, a general framework for optimising ASL experiments was presented. This framework can be used with any labelling scheme, for which there is an analytical model,

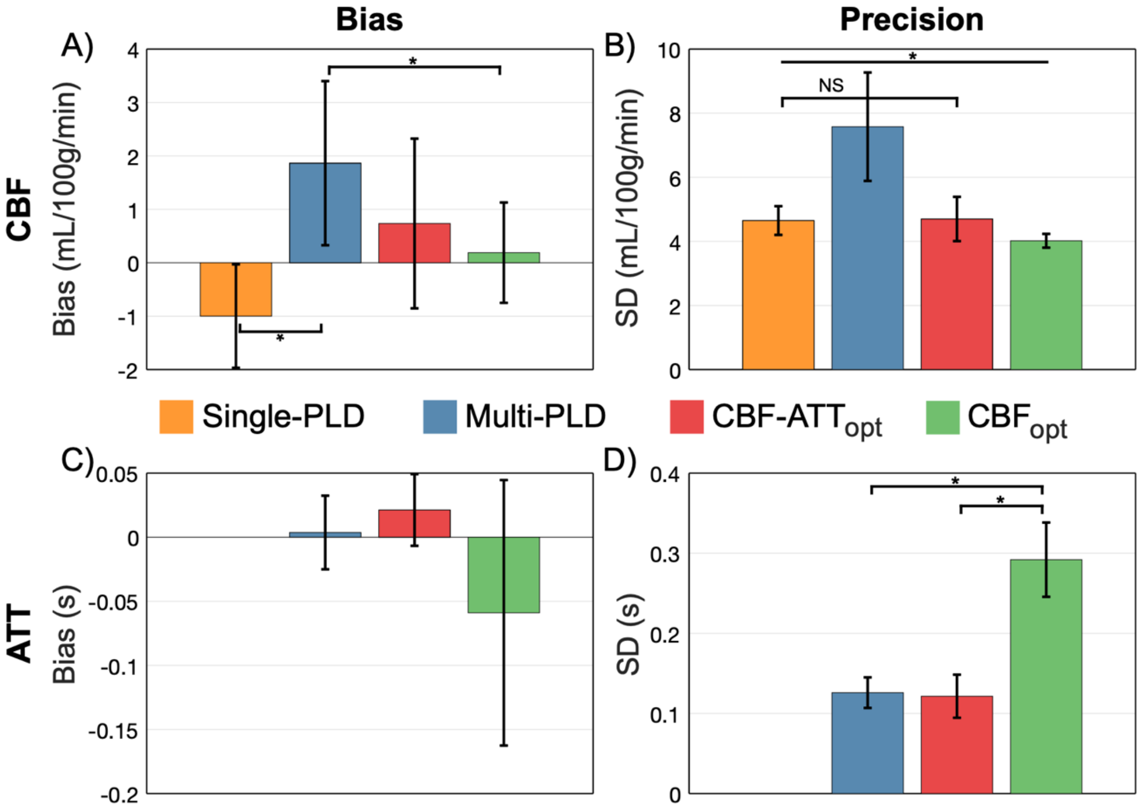

Figure 4.13: In vivo CBF (top row) and ATT (bottom row) bias (A,C) and precision (B,D) across subjects. The height of each bar graph shows the mean across subjects, while the error bar shows the SD across subjects. An asterisk (*) signifies significant differences using a non-parametric paired test (Wilcoxon signed rank test), $p < 0.05$. All differences in (B) are significant except Single-PLD vs $\mathrm{CBF\text{-}ATT_{opt}}$.

and any readout. The practical benefits of using this framework in the specific case of fixed label duration, multi-PLD PCASL experiments with a 5 slice EPI readout were demonstrated. Two protocols were designed which either aimed to improve both CBF and ATT accuracy or just CBF accuracy. These protocols were shown to achieve their respective aims using Monte Carlo simulations and in vivo experiments, including a strong agreement with the predicted performance. The CBF optimised protocol resulted in a 48% reduction of in vivo CBF errors relative to the reference multi-PLD protocol while optimising for both CBF and ATT resulted in a 37% reduction of in vivo CBF errors, without a loss in ATT accuracy. It was also shown that a well optimised multi-PLD protocol can produce more accurate CBF estimates (15% reduction in error) than a single-PLD protocol in the same scan time, as well as providing potentially physiologically useful ATT maps.[3]

4.6.1 Optimisation

Simplifications that were made to the FIM (see Theory 4.3.5) show that a CBF distribution is not required for optimisation because the optimal protocol is insensitive to the value chosen. These simplifications improve ease of use compared to previous PASL work,[8–10] since only an ATT prior probability distribution needs to be defined, and also greatly speeds up the optimisation process due to the reduced dimensionality of the problem. Furthermore, by including the average term in Equation 4.2, the scan time can also be pre-specified, meaning that strict clinical scan times can be adhered to and the optimal protocol for the time available may be found.

Equation 4.9, as used in the algorithm outlined in the pseudocode in Figure 4.4, combines the CRLB information across the ATT distribution before finding the optimal timing for each PLD. This approach results in PLD timings that maximise the information obtained across the entire ATT prior simultaneously. This is a distinct advantage over previous work where the PASL inversion times (TIs) were chosen based on a histogram of locally optimal TIs for individual ATT and CBF parameter pairs drawn from the priors, which was divided between a fixed number of TIs.[8,9] Further PASL work used an average approach, employing an adaptive quadrature technique to reduce the number of samples required from the ATT and CBF priors, thus improving the speed of optimisation.[10] However, after removing the need for a CBF prior distribution, the number of calculations required is drastically reduced, limiting the benefit of using such a technique. The direct sampling method used here is also more easily implemented, assisting the adoption of CRLB optimisation for ASL studies.

The use of a uniform ATT prior probability distribution was proposed, rather than a normal distribution as used in previous studies.[8–10] ATT has been shown to be regionally dependent,[26] so even if a pre-defined normal distribution does exactly match the distribution of ATTs found in a given subject, the brain regions corresponding to the edges of the distribution will be less accurately estimated than regions corresponding to the centre of the distribution. This may be desirable if certain brain regions are of more interest to the experimenter, but we opted to equally weight the entire brain in this study. However, the general framework described here would allow any shape or range of ATT distribution to be used.

The ATT range used in this work was appropriate for healthy volunteers.[1,16] This is the main a priori knowledge necessary for optimisation using this method but may not always be known. When scanning specific patient groups where the typical ATT range is known, this should be used to inform the ATT prior probability distribution. However, when no specific a priori information is available, a conservative large range of ATTs could be used.

The algorithm was found to be reasonably insensitive to initialisation, with PLDs in the final protocol only varying by one or two increments of 25 ms if different initial conditions are chosen. This suggests that solutions close to the global minimum were achieved. It should be noted that care must be taken to initialise the PLDs to cover most of the ATT range. If not, the PLDs may fail to update due to all possible PLD choices being dominated by ill-conditioned FIMs within the ATT distribution.

4.6.2 Optimal Protocols

The optimised multi-PLD protocols found in this work differ greatly from commonly used evenly spaced multi-PLD protocols. The designs reflect the underlying information content, demonstrated by the sensitivity functions (Figure 4.2). For a single ATT value, the CBF-ATT$_{\text{opt}}$ PLDs were found at $\text{PLD} = [\text{ATT},\ \text{ATT} + \tau]$, as found by Xie et al.[8] However, the CBF-ATT$_{\text{opt}}$ protocol did not simply reflect this relationship and the shape of the ATT prior. Instead, a starkly different pattern is seen in Figure 4.7, which is a result of the optimisation accounting for the combined information obtained across all PLDs simultaneously, rather than considering the local optimisation of each PLD.

The CBF$_{\text{opt}}$ protocol was also introduced, which minimises the CBF variance while being insensitive to ATT variation. Xie et al.[8] also used a CBF specific optimisation but did not include the ATT sensitivity function in the FIM. If the ATT sensitivity function is not included in the FIM, it is implicitly assumed that the ATT is known, which results in much larger errors (see Figure 4.14).

An alternative optimality criterion not investigated in this work is the non-zero weighted trace of the CRLB matrix, in Equation 4.4. This could be used to choose any trade-off between CBF and ATT accuracy, resulting in a much more flexible tool for optimising ASL experiments. However, the relative weightings for the desired trade-off would need to be empirically deduced; a potentially burdensome process. Furthermore, other model parameters, such as tissue T_1, could be included in the FIM with the aim of minimising sensitivity to them, similar to the CBF$_{\text{opt}}$ approach to ATT.

For both protocols, the optimal number of PLDs were found to be large because a large number of PLDs provides the most flexibility for minimising the predicted variance. This may seem counterintuitive since each data point is relatively noisy. However, the information across PLDs is combined during fitting, giving similar results to averaging, but with greater information content. For ease of implementation, the proposed framework can be restricted to choose a smaller number of PLDs with more than one average at each.

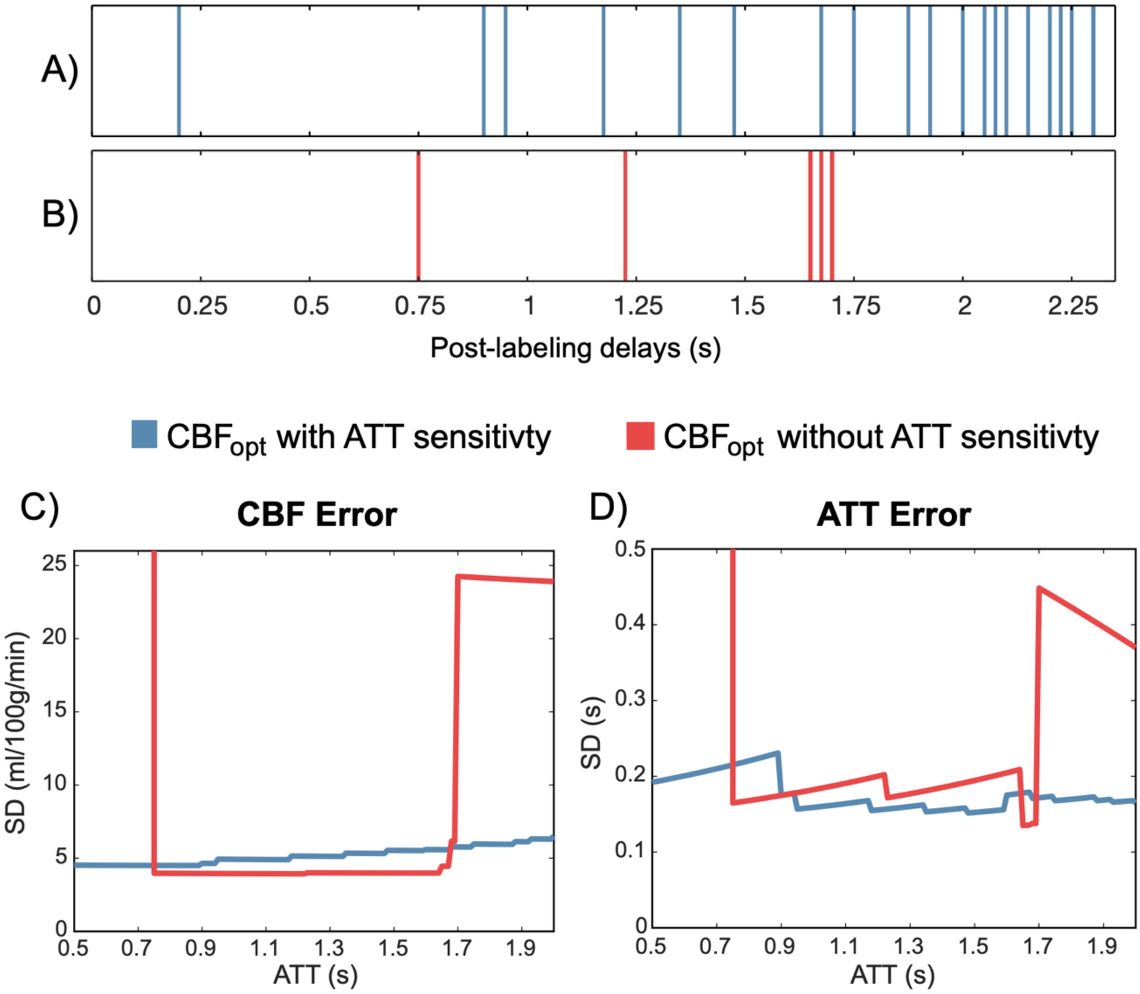

Figure 4.14: The PLDs (A,B) and the predicted CBF and ATT errors (Cramér-Rao lower bound (CRLB) SD) (C,D) for CBF$_{\text{opt}}$ *with and without the ATT sensitivity function in the FIM. Assumptions: 0.5 s readout duration for a 3D acquisition, 28 PLDs and an ATT range of 0.5 – 2 s (similar to Figure 4.18). Repeated PLDs are not shown. CBF and ATT CRLBs at ATTs shorter than 0.75 s are extremely large for the protocol without the ATT sensitivity function included and are out of view for clarity.*

For example, restricting the number of CBF-ATT$_{opt}$ PLDs to 10 or fewer resulted in a 10 PLD protocol which only increased the predicted mean CBF and ATT errors by ~0.5% and ~1%, respectively (see Figure 4.15). This relatively small increase suggests that a smaller number of PLDs can still be used effectively and should facilitate cases where a high degree of readout segmentation is required, such as with whole-brain 3D readouts, where each PLD takes several TRs to acquire.

We did not test whether the optimised protocols were more or less affected by loss of data, so the effect of the optimisation framework on motion-robustness is unknown. This could be evaluated by fitting the simulation or in vivo data many times while randomly exluding a given proportion of the data. Motion-robustness could also be

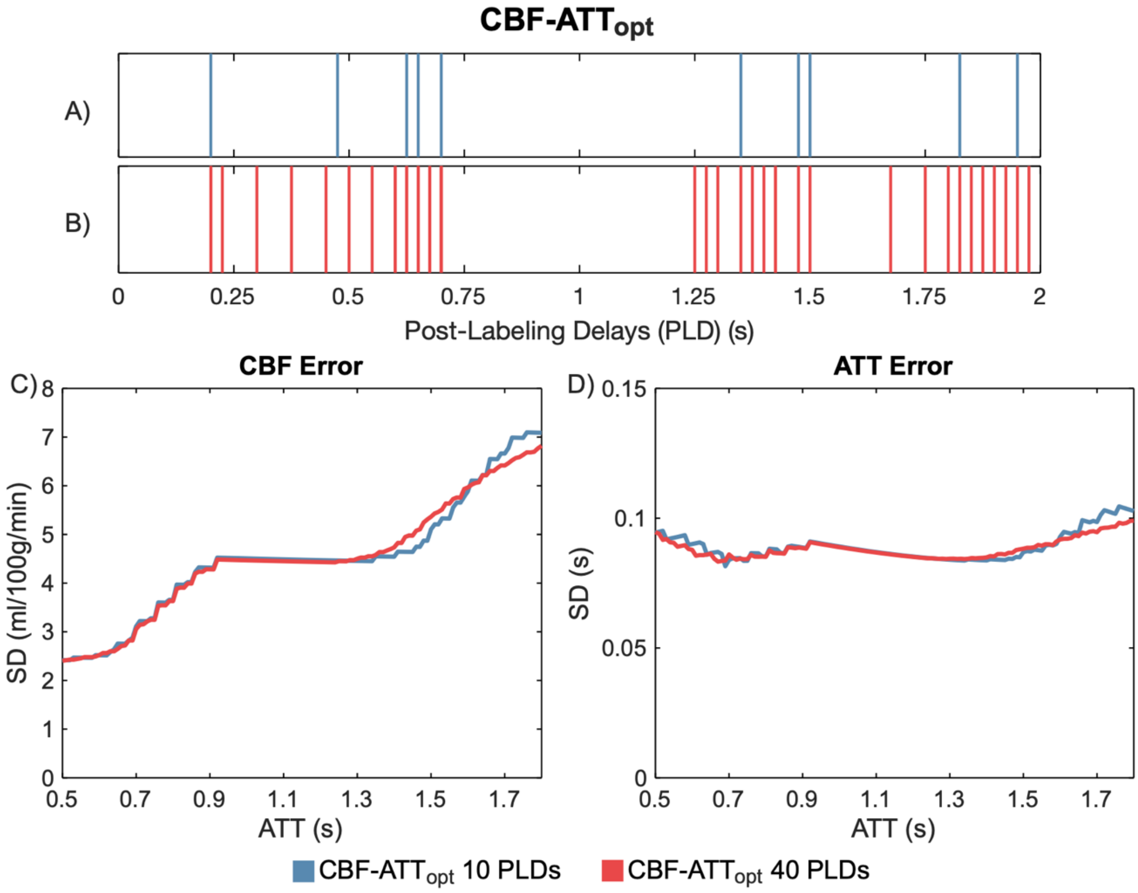

Figure 4.15: The PLDs (A,B) and the predicted CBF and ATT errors (Cramér-Rao lower bound (CRLB) SD) (C,D) for the CBF-ATT$_{opt}$ protocol with 10 or 40 PLDs. The 10 PLD protocol uses 4 averages of the 10 PLDs, while the 40 PLD protocol only has one average for each PLD. Using 10 PLDs rather than 40 PLDs only resulted in a ~0.5% and ~1% increase in the CBF and ATT CRLBs, respectively.

incorporated into the optimisation framework by finding the protocol timings which result in the highest CBF and ATT accuracy on average after excluding a random selection of data points.

In this study, the PCASL labelling duration was fixed at that used in the reference protocol[13] for simplicity. However, using this framework, a single label duration for the whole experiment or different values for each PLD could be optimised, yielding greater flexibility in the design and potentially leading to further reductions in CBF and ATT estimation errors. Furthermore, the sub-bolus timings in time-encoded PCASL,[27] a method which efficiently produces multi-delay ASL data, could be optimised, potentially yielding greater improvements in estimation accuracy. The potential of these more advanced techniques is explored in Chapter 5 of this thesis.

4.6.3 Simulation and In Vivo Experiments

Figure 4.7 and Figure 4.10 demonstrate the excellent agreement between the theoretical, simulated and in vivo CBF and ATT errors, validating the optimisation framework outlined here. However, there was a poorer agreement between the simulation and in vivo CBF results at ATTs less than 0.8 s, where the CBF errors for all protocols were higher than expected, with the single-PLD protocol being most affected. This could be due to the reduced data available at these ATT values, making the results more susceptible to noisy estimates, and/or the presence of residual macrovascular signal causing CBF overestimates in the ground truth data. A further explanation, which is explored in Chapter 5, is that short ATTs occur near the large arteries which are typically located closer to the centre of the brain where SNR is lower with the 32-channel receive array coil used here. Another noticeable difference was the increase in the simulation ATT RMSE at short transit times, compared to the CRLBs. This is due to poor ATT estimation in the top few slices which was not possible to resolve due to the longer effective PLDs at these slices.

4.6.3.1 Fitting Method

The CRLB is based on maximum-likelihood estimation, as is NLLS fitting. However, the data is not restricted to being fit with an NLLS method, though we chose to here for consistency. Figure 4.16 demonstrates the differences between fitting the in vivo data with the NLLS approach and a variational Bayesian approach, as implemented in the BASIL toolbox in FSL,[28] which is commonly applied to in vivo data. The general trends of the protocols are very similar for both fitting methods, and therefore this choice does not influence the conclusions of this work. BASIL reduced CBF errors, particularly in regions with very large errors in the NLLS fitting. Mean ATT errors were also reduced but were larger at short ATTs. This suggests that BASIL produces better CBF estimates from noisy

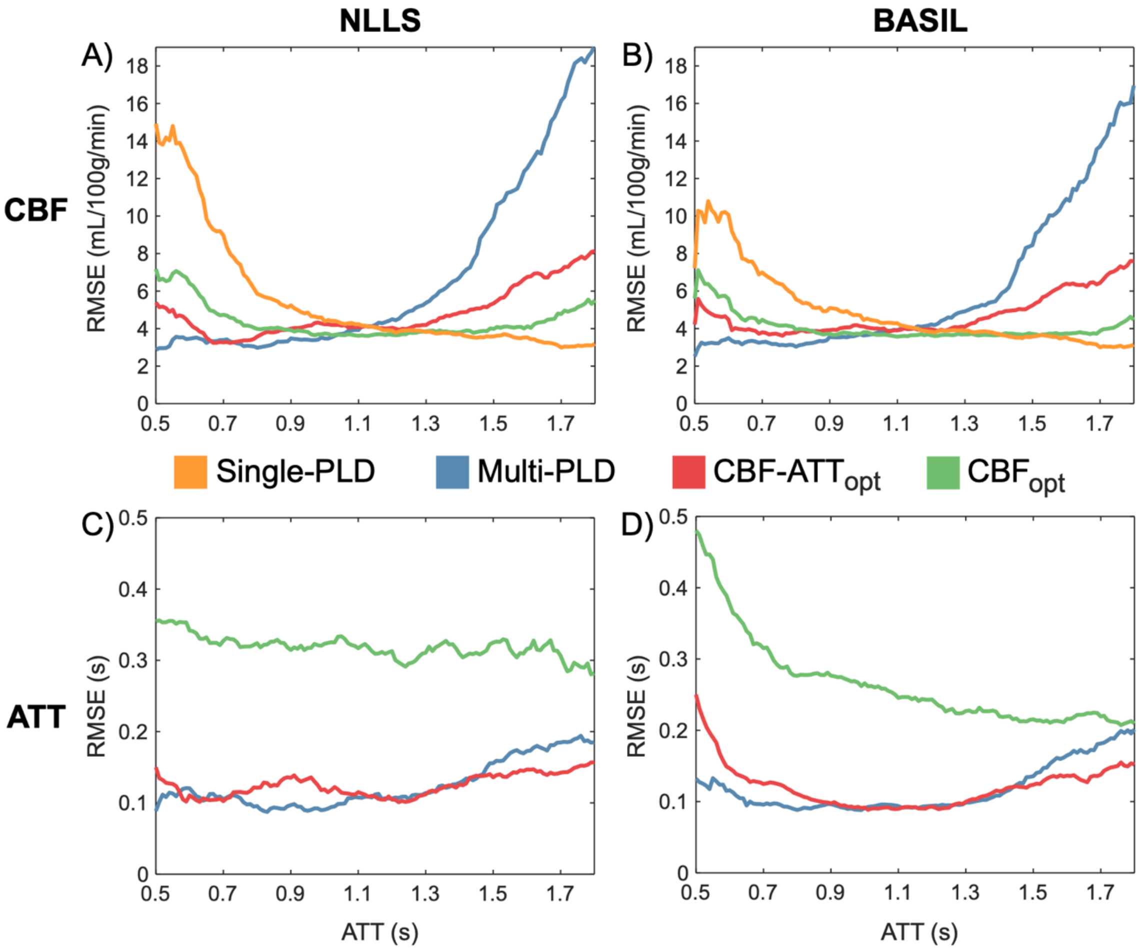

Figure 4.16: In vivo RMSEs of CBF (top row) and ATT (bottom row) for data fitted with the NLLS method (A,C), as in Figure 4.10(B,D), and with BASIL (B,D), a variational Bayesian algorithm.

data than a naïve NLLS algorithm, but there remain significant benefits from appropriately optimising the PLDs in both cases. The fitting priors used for the BASIL fitting were (mean ± SD) 0 ± 10^6 mL/100g/min and 1.25 ± 1 s, for CBF and ATT respectively. The same exclusion criteria as before were applied to the BASIL fitted data, but the marginal posterior distribution SDs output by BASIL were used instead of the CRLB SDs.

4.6.3.2 Exclusion of Data

Although the majority of well-fit voxels in the healthy volunteers were within the ATT range optimised for, 4.8% had ATTs longer than 1.8 s. If the optimisation ATT range was

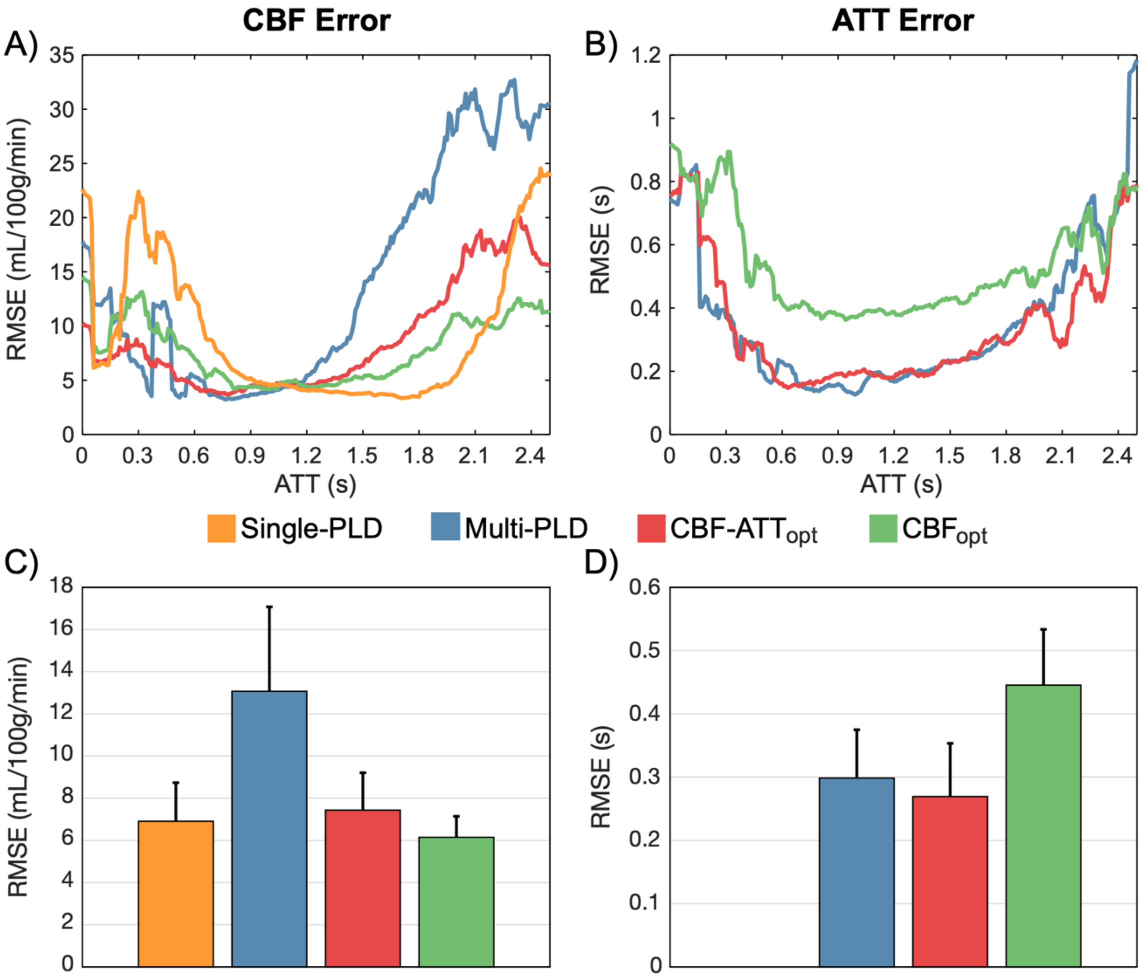

Figure 4.17: In vivo RMSEs of CBF (left column) and ATT (right column) estimates for all GM data for the 7 subjects with no other data exclusion criteria imposed. The top row (A,B) shows the RMSE trends across ATTs, while the bottom row (C,D) shows the mean and SD of the RMSEs across subjects. The general trends remain unchanged, compared to Figure 4.10 and Figure 4.11.

extended, the number of well-fit voxels with longer ATTs would probably increase. For this reason, future work in healthy volunteers should consider using a wider ATT prior.

To demonstrate the robustness of the results to the ground truth data exclusion criteria, all of the GM data were re-analysed without the restrictions on fitting accuracy or ATTs placed on the ground truth data; the results are shown in Figure 4.17. The general trends and relative performance seen in Figure 4.10 and Figure 4.11 are largely unchanged. As expected, there is a general increase in RMSEs due to the removal of the ground truth data exclusion criteria. Figure 4.17 (A) and (B) agree well with Figure 4.10 in the ATT prior range of $0.5 \leq \text{ATT} \leq 1.8$. The trends in (C) and (D) also agree well with Figure 4.11, though the CBF_{opt} and single-PLD CBF RMSEs are no longer significantly different. This is to be expected since there is greater noise in the ground truth estimates and data from outside the optimised range has been included.

4.6.3.3 Use of Ground-Truth Data

The in vivo CBF and ATT estimates were compared with ground truth estimates, derived from the same, but combined, data from the individual protocols. This has the potential to benefit the accuracy of some protocols over others if the estimates are biased in some way. However, when this analysis process was simulated, the relative performance of each protocol was largely unaffected. Specifically, using estimated ground truth values lead to a relatively even decrease in the CBF and ATT RMSE across all of the protocols: 0.49 ± 0.08 mL/100g/min and 0.011 ± 0.004 s, respectively.

4.6.3.4 Noise

Magnitude data has a Rician noise distribution, but a Gaussian noise distribution was used in the simulations and data fitting. It has been shown, however, that when the SNR of data is greater than 3, a Gaussian distribution is a good approximation of the true Rician distribution.[29] To test if this condition was met with the in vivo ASL data used in this study, the temporal SNR (tSNR) of the reference multi-PLD and single-PLD control GM magnitude data was calculated using the formula: $tSNR = S \ / \ \sigma$, where S and σ are the mean and standard deviation of the signal over the repeats for each PLD. Only 0.54% of the GM voxels had a tSNR less than 3 with the median tSNR being 31.4. However, if extremely efficient BGS was used or complex subtraction performed before magnitude reconstruction of the perfusion-weighted data, then a Gaussian distribution would no longer be an appropriate approximation.

In the optimisation and fitting, it was assumed that there was an identical noise magnitude across PLDs. However, the BGS strategy used in this work resulted in variable static tissue signal suppression and therefore variable noise levels across the PLDs. Appendix B details the measurement of this variable noise level across PLDs and its use in further Monte Carlo simulations. When comparing the results of these variable noise simulations with the results presented in Figure 4.10, minor differences were found to the CBF and ATT errors, but the broad trends are unchanged.

The variable noise model could be incorporated into the optimisation, but it is highly specific to the sequence used. Its effect on the estimates could also be reduced by weighting the data appropriately during fitting.[28] Alternatively, a more flexible BGS scheme which allows interleaving of the inversion pulses with labelling, such as that used by Dai et al.,[5,30] would result in more comparable BGS across PLDs, better meeting the assumption of equal noise at all PLDs. A more even level of BGS across PLDs will also further reduce the small image shift artefacts that occurred when the BGS varied.

4.6.3.5 Image Readout

For this proof-of-concept study, a multi-slice 2D readout was used with a small number of slices to minimise the variation in timings and BGS. However, the use of a single-excitation 3D readout would ensure identical BGS and PLDs across all slices, leading to a simpler optimisation problem and improved static tissue attenuation across the imaged volume. This would enable the full benefits of this optimisation framework to be realised for whole-brain perfusion measurements.

Although this framework was only validated in a limited cross-section of the brain, it is expected that this method will benefit perfusion measurements across the whole brain because the range of ATTs in these data (Figure 4.8) is comparable to previously reported ranges across the whole brain[5,16,31] and inferior-superior ATT variations have previously been shown to be comparable to within-slice variations.[22]

Protocol	***Post-labelling delays (s)***	***# PLDs***	***# Averages***
$\mathbf{0.5 \leq ATT \leq 2\ s}$			
Even	0.5, 0.725, 0.925, 1.15, 1.35, 1.575, 1.775, 2	8	6
CBF-ATT$_{opt}$	0.2, 0.9, 0.9, 1.15, 2, 2.2	6	8
CBF$_{opt}$	0.2, 0.975, 1.425, 1.85, 2.025, 2.15, 2.25, 2.3, 2.3, 2.3	10	4
$\mathbf{1 \leq ATT \leq 3\ s}$			
Even	1, 1.25, 1.5, 1.75, 2, 2.25, 2.5, 2.75, 3	9	4
CBF-ATT$_{opt}$	0.7, 1.875, 1.9, 1.9, 1.925, 2.75, 2.95, 3.1, 3.3	9	4
CBF$_{opt}$	0.7, 1.8, 2.5, 2.875, 3.075, 3.3, 3.3, 3.325	8	4

Table 4.3: Optimal PLDs for 3D acquisitions

Due to the increasing use of 3D readouts with ASL, such as 3D GRASE and stack-of-spirals,[1] 3D (or single-slice) specific CBF-ATT$_{opt}$ and CBF$_{opt}$ PLD timings have been included for completeness. These were generated for two ATT ranges: a standard ATT range of $0.5 \leq ATT \leq 2$ s and a prolonged ATT range of $1 \leq ATT \leq 3$ s, in a similar way to as described in Section 4.4.1 except that the readout duration was changed to 500 ms and only a single slice was used. In each case, the number of PLDs was constrained to ≤10 in order to facilitate readout segmentation, necessary to reduce T_2/T_2^* blurring or off-resonance distortions in whole-brain 3D readouts with long echo-train lengths. The

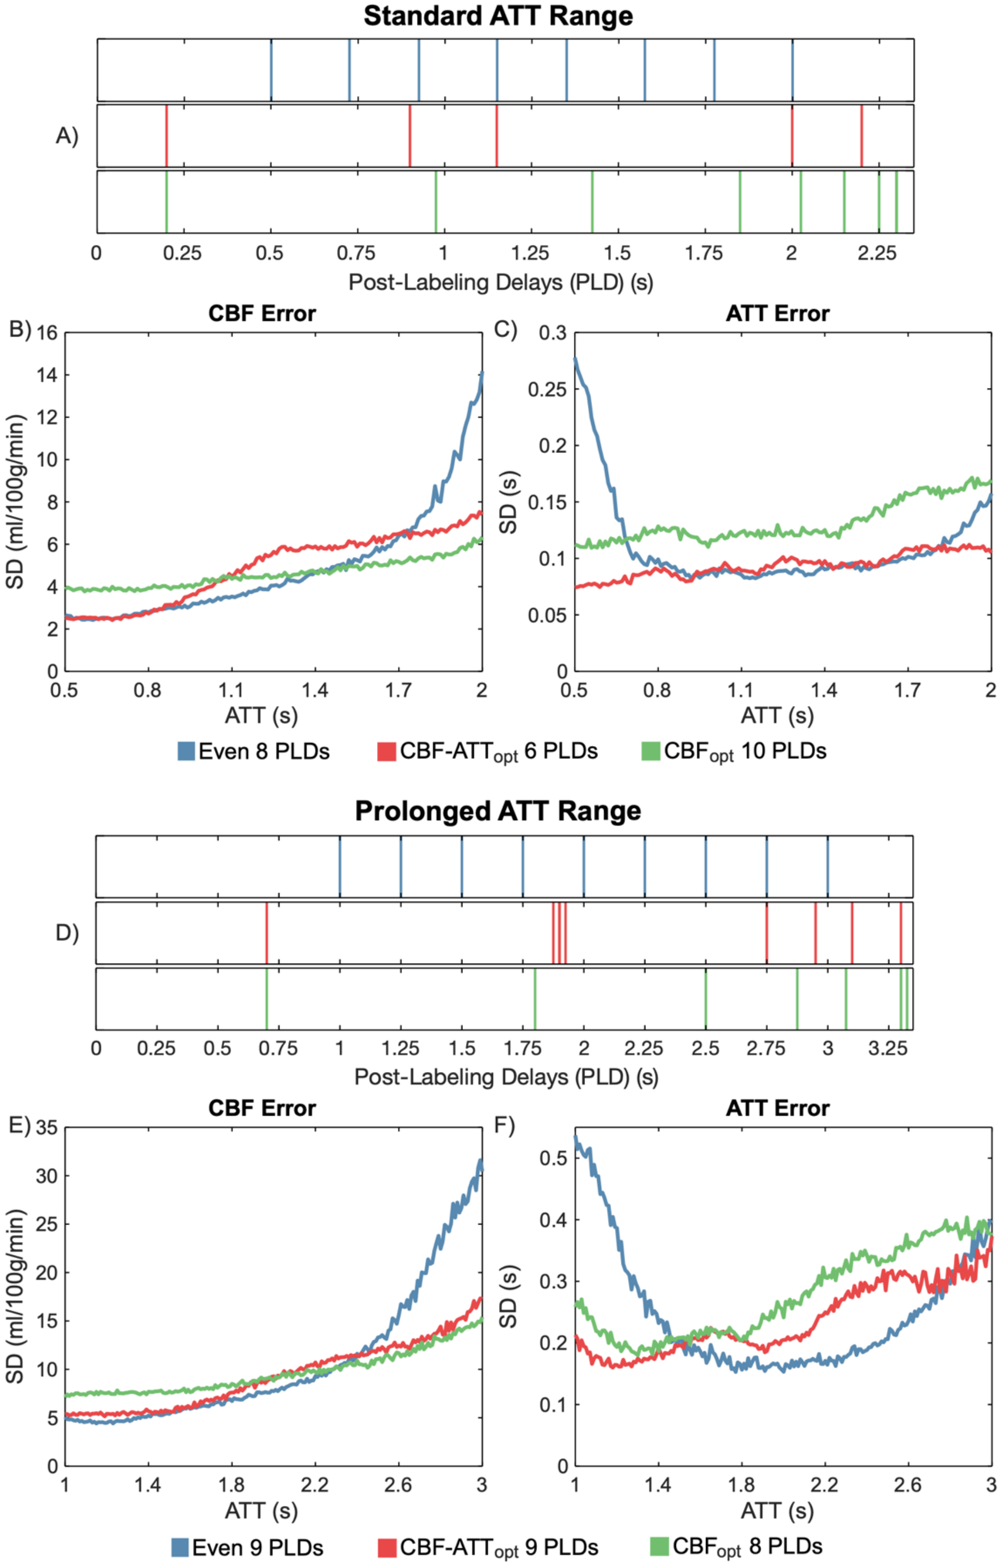

Figure 4.18: The PLDs (A,D) and the MC simulation RMSEs (B,C,E,F) for evenly distributed, CBF-ATT$_{opt}$ and CBF$_{opt}$ protocols over a standard healthy range of ATTs ($0.5 \le ATT \le 2$) s and a prolonged ATT range ($1 \le ATT \le 3$). The timings shown are for a 3D/single slice acquisition. Repeated PLDs are not shown.

optimised protocols were compared to evenly distributed protocols in each case.

The timings can be found in Table 4.2 (where the number of averages assume a single-shot readout) and the MC simulation RMSEs are shown in Figure 4.18. For both ATT ranges, the optimised protocols reduced the average errors across the ATT distributions relative to the evenly distributed protocol. However, for late ATTs in the prolonged ATT range, the CBF and ATT errors were large for all protocols, demonstrating the difficulty in estimating CBF in the presence of delayed ATT and short scan durations. Further questions, therefore, remain as to the best approach to improve CBF and ATT estimation accuracy in cases of delayed ATT. An alternative method to PCASL that holds promise for CBF estimation in the presence of delayed ATT is velocity-selective ASL,[32] which effectively eliminates the ATT and so experiences minimal tracer T_1 decay compared to whole-brain PCASL.

For both 2D and 3D readouts, it is assumed that the image is static after the acquisition excitation and during the readout. Motion during this period can result in ghosting or blurring depending on the type of readout used. The movement of ASL signal which is still intravascular during the readout can, therefore, lead to artefacts in the acquired images if flow compensation is not used. This effect is likely to lead to larger artefacts in 3D acquisitions due to the longer echo-trains which are typically used.

4.6.3.6 White Matter Perfusion

The work in this chapter has focused on optimising protocols for GM perfusion. However, the framework could also be applied to measuring white matter (WM) perfusion, which typically has longer ATTs, shorter T_1 and lower CBF than GM, making it a much more challenging application, thus increasing the importance of appropriately optimising protocols. The low SNR of WM could also be aided by moving to higher field strengths than 3T due to the longer T_1 relaxation times.[33]

4.6.3.7 Reference Protocol

It should be noted that the reference multi-PLD protocol[13] was originally intended for a 24 slice whole-brain scan, meaning that the average effective PLDs across the brain are on the order of 0.5 s longer than those given in Table 4.2. This would likely result in more accurate CBF estimates at longer ATTs in higher slices, but reduced accuracy at shorter ATTs in those same slices, than is presented here. The protocols were not optimised or acquired across 24 slices due to the wide-range of PLDs that result from such a large number of slices. The round-robin approach suggested by Lee et al.,[34] which permutes the order in which the slices are acquired in order to achieve multi-PLD data, could be used to improve the estimation accuracy of whole-brain multi-slice 2D readouts across the slices. This approach could also make the optimisation of such a protocol more compatible with the optimisation framework introduced in this chapter.

4.7 Conclusions

In conclusion, a general framework for optimising ASL experiments has been developed. It was validated in the specific case of multi-PLD PCASL, showing significant improvements over reference single-PLD and multi-PLD protocols. The clinical use of this framework will lie in the development and testing of standardised, patient population specific protocols.

4.8 References

1 Alsop DC, Detre JA, Golay X, Günther M, Hendrikse J, Hernandez-Garcia L *et al.* Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn Reson Med* 2015; **73**: 102–116.

2 Buxton RB, Frank LR, Wong EC, Siewert B, Warach S, Edelman RR. A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magn Reson Med* 1998; **40**: 383–396.

3 MacIntosh BJ, Marquardt L, Schulz UG, Jezzard P, Rothwell PM. Hemodynamic

alterations in vertebrobasilar large artery disease assessed by arterial spin-labeling MR imaging. *Am J Neuroradiol* 2012; **33**: 1939–1944.

4 Gonzalez-At JB, Alsop DC, Detre JA. Cerebral perfusion and arterial transit time changes during task activation determined with continuous arterial spin labeling. *Magn Reson Med* 2000; **43**: 739–746.

5 Dai W, Robson PM, Shankaranarayanan A, Alsop DC. Reduced resolution transit delay prescan for quantitative continuous arterial spin labeling perfusion imaging. *Magn Reson Med* 2012; **67**: 1252–1265.

6 Rao CR. Information and the Accuracy Attainable in the Estimation of Statistical Parameters. In: Kotz S, Johnson NL (eds). *Breakthroughs in Statistics. Springer Series in Statistics (Perspectives in Statistics)*. Springer New York: New York, NY, 1992, pp 235–247.

7 Cramer H. *Mathematical Methods of Statistics.* 9th ed. Princeton University Press: Princeton, NJ, 1946https://www.jstor.org/stable/2280199?origin=crossref.

8 Xie J, Gallichan D, Gunn RN, Jezzard P. Optimal design of pulsed arterial spin labeling MRI experiments. *Magn Reson Med* 2008; **59**: 826–834.

9 Sanches J, Sousa I, Figueiredo P. Bayesian fisher information criterion for sampling optimization in ASL-MRI. In: *2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Rotterdam, Netherlands*. 2010, pp 880–883.

10 Owen D, Melbourne A, Thomas D, De Vita E, Rohrer J, Ourselin S. Optimisation of Arterial Spin Labelling Using Bayesian Experimental Design. In: Ourselin S, Joskowicz L, Sabuncu M, Unal G, Wells W (eds). *Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2016. Lecture Notes in Computer Science*. Springer, Cham, 2016, pp 511–518.

11 Kramme J, Gregori J, Diehl V, Madai VI, von Samson-Himmelstjerna FC, Lentschig M *et al.* Improving perfusion quantification in arterial spin labeling for delayed arrival times by using optimized acquisition schemes. *Z Med Phys* 2015; **25**: 221–229.

12 Dai W, Garcia D, De Bazelaire C, Alsop DC. Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magn Reson Med* 2008; **60**: 1488–1497.

13 Okell TW, Chappell MA, Kelly ME, Jezzard P. Cerebral blood flow quantification using vessel-encoded arterial spin labeling. *J Cereb Blood Flow Metab* 2013; **33**: 1716–1724.

14 Gilmour SG, Trinca LA. Bayesian L-optimal exact design of experiments for biological kinetic models. *J R Stat Soc C* 2012; **61**: 237–251.

15 Jones JA, Hodgkinson P, Barker AL, Hore PJ. Optimal Sampling Strategies for the Measurement of Spin–Spin Relaxation Times. *J Magn Reson Ser B* 1996; **113**: 25–34.

16 Guo J, Holdsworth SJ, Fan AP, Lebel MR, Zun Z, Shankaranarayanan A *et al.* Comparing accuracy and reproducibility of sequential and Hadamard-encoded multidelay pseudocontinuous arterial spin labeling for measuring cerebral blood flow and arterial transit time in healthy subjects: A simulation and in vivo study. *J Magn Reson Imaging* 2018; **47**: 1119–1132.

17 Johnston ME, Lu K, Maldjian JA, Jung Y. Multi-TI arterial spin labeling MRI with variable TR and bolus duration for cerebral blood flow and arterial transit time mapping. *IEEE Trans Med Imaging* 2015; **34**: 1392–1402.

18 Lu H, Clingman C, Golay X, van Zijl PCM. Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn Reson Med* 2004; **52**: 679–682.

19 Lin C, Bernstein M, Huston J, Fain S. Measurements of T1 Relaxation times at 3.0T:

Implications for clinical MRA. In: *Proceedings of the 9th Annual Meeting of the ISMRM, Glasgow, Scotland*. 2001, p 1391.

20 Herscovitch P, Raichle ME. What is the Correct Value for the Brain-Blood Partition Coefficient for Water? *J Cereb Blood Flow Metab* 1985; **5**: 65–69.

21 Golay X, Petersen ET, Hui F. Pulsed Star Labeling of Arterial Regions (PULSAR): A robust regional perfusion technique for high field imaging. *Magn Reson Med* 2005; **53**: 15–21.

22 Günther M, Oshio K, Feinberg DA. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn Reson Med* 2005; **54**: 491–498.

23 Zhang Y, Brady M, Smith S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging* 2001; **20**: 45–57.

24 Jenkinson M. Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *Neuroimage* 2002; **17**: 825–841.

25 Seber GAF, Wild CJ. *Nonlinear Regression*. John Wiley & Sons, Inc: New York, NY, 1989 doi:10.1002/0471725315.

26 MacIntosh BJ, Filippini N, Chappell MA, Woolrich MW, Mackay CE, Jezzard P. Assessment of arterial arrival times derived from multiple inversion time pulsed arterial spin labeling MRI. *Magn Reson Med* 2010; **63**: 641–647.

27 Günther M. Highly efficient accelerated acquisition of perfusion inflow series by Cycled Arterial Spin Labeling. In: *Proceedings of the 15th Annual Meeting of the ISMRM, Berlin, Germany*. 2007, p 380.

28 Chappell MA, Groves AR, Whitcher B, Woolrich MW. Variational Bayesian Inference for a Nonlinear Forward Model. *IEEE Trans Signal Process* 2009; **57**: 223–236.

29 Gudbjartsson H, Patz S. The rician distribution of noisy MRI data. *Magn Reson Med* 1995; **34**: 910–914.

30 Dai W, Robson PM, Shankaranarayanan A, Alsop DC. Benefits of Interleaved Continuous Labeling and Background Suppression. In: *Proceedings of the 18th Annual Meeting of the ISMRM, Stockholm, Sweden*. 2010, p 1748.

31 Chen Y, Wang DJJ, Detre JA. Comparison of arterial transit times estimated using arterial spin labeling. *MAGMA* 2012; **25**: 135–144.

32 Wong EC, Cronin M, Wu W-C, Inglis B, Frank LR, Liu TT. Velocity-selective arterial spin labeling. *Magn Reson Med* 2006; **55**: 1334–1341.

33 Gardener AG, Jezzard P. Investigating white matter perfusion using optimal sampling strategy arterial spin labeling at 7 Tesla. *Magn Reson Med* 2015; **73**: 2243–2248.

34 Lee W, Janik R, Scouten A, Stefanovic B, Sled JG. Efficient sampling of early signal arrival for estimation of perfusion and transit time in whole-brain arterial spin labeling. *Magn Reson Med* 2012; **68**: 179–187.

5

Comparison of Existing and Novel PCASL Perfusion Protocols

5.1 Introduction

In Chapter 4, it was demonstrated that a Cramér-Rao lower bound based optimisation framework was effective at objectively designing multi-PLD PCASL protocols to either maximise CBF accuracy in the presence of unknown ATTs or to maximise both CBF and ATT accuracy.

In this chapter, the aim was to investigate the relative CBF accuracy of a range of advanced PCASL protocol designs, including time-encoded strategies. An example range of the protocols compared are shown in Figure 5.1 and are described in Section 5.2. The optimisation framework presented in Chapter 4 was used to objectively optimise each protocol so that it achieves close to its maximum possible CBF accuracy across a healthy range of ATTs. The objective was then to study the inherent accuracy of each protocol design, using simulations and in vivo data, to assess which protocol design provides the highest CBF accuracy for this ATT range. As well as optimising and comparing several existing protocol designs, a novel design, which is a combination of time-encoding and sequential PLDs, was also explored.

Several studies have previously compared the CBF and ATT accuracy of different PCASL protocols. However, these studies only compared a small subset of protocols and did not explicitly optimise them for CBF or ATT accuracy.[1–4]

Since the benefit of varying both label durations and PLDs is explored in this chapter, the term 'multi-TI' (inflow-time) is used rather than simply 'multi-PLD', to refer to ASL protocols with multiple LDs and/or PLDs.

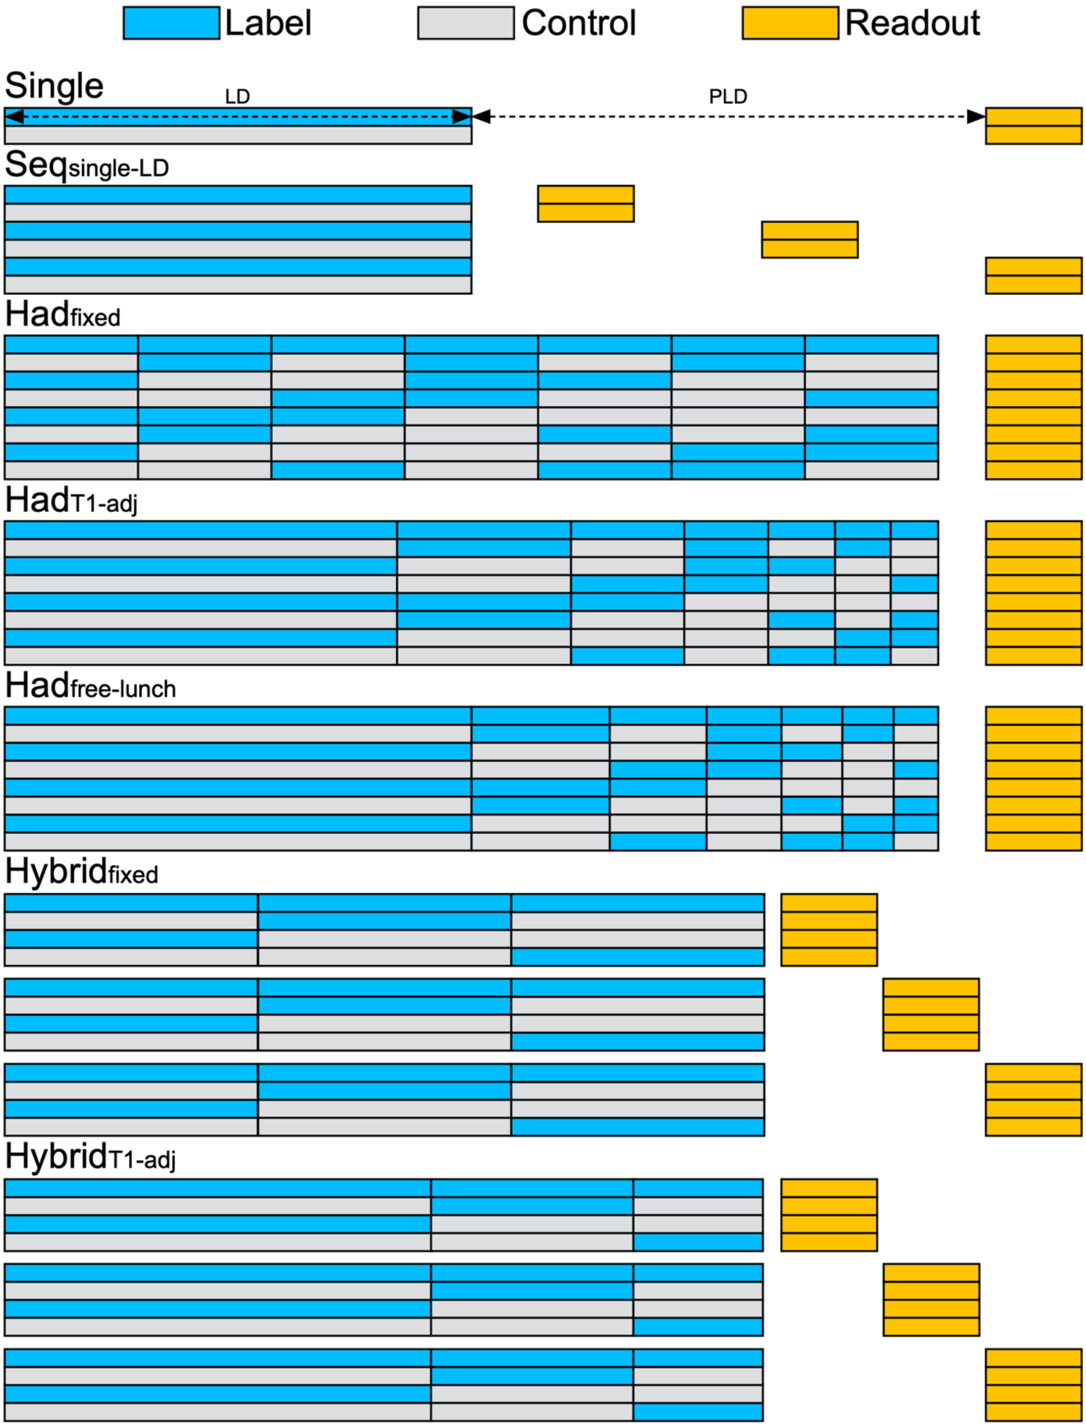

Figure 5.1 Overview of the protocols explored in this work. The first 5 protocols have been previously proposed in the literature. The Had$_{free\text{-}lunch}$ protocol is shown with T_1-adjusted final LDs, but fixed-duration LDs were also explored for the final LDs. The hybrid protocols are Hadamard-encoded schemes with multiple PLDs to increase the timing flexibility.

Some of the work described in this chapter was presented at the 26th annual meeting of the International Society for Magnetic Resonance in Medicine.[*][†]

5.2 Protocol Designs

5.2.1 Single-TI

As described in Chapter 4, ASL perfusion imaging is often performed using a single LD and PLD. The PLD should be set equal to the longest ATT within the imaging region to avoid large underestimates of CBF when PLD < ATT, and to avoid unnecessary SNR loss when PLD > ATT_{max}.

For the LD, Zun et al.[5] demonstrated that it is more SNR efficient to use long LDs of 3 - 4 s rather than shorter LDs with more averages. This suggests that the increase in ASL signal accumulation in the tissue outweighs the loss of noise averaging from the increase in TR. Lebel et al.[6] further demonstrated that longer LDs are beneficial in terms of temporal signal stability (possibly due to an increased spread of ASL label generation over the cardiac and respiratory cycles) and reduced sensitivity to longer than expected ATTs. However, since a LD of 1.8 s is recommended for routine use,[7] this was used as the standard LD for single-TI protocols in this chapter, with the use of much longer LDs explored later in the chapter.

---

[*] Woods JG, Chappell MA, Okell TW. Comparison of optimized single-PLD, sequential multi-PLD and time-encoded PCASL for cerebral blood flow measurements. In: *Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France*, 2018. p. 2162.

[†] Woods JG, Chappell MA, Okell TW. A novel hybrid of time-encoded and sequential multi-PLD PCASL for improved cerebral blood flow estimation. In: *Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France*, 2018. p. 2161.

5.2.2 Sequential Multi-TI

If multiple TIs are sampled (multiple LDs and PLDs), the ATT can also be estimated, reducing this possible source of large errors. Another advantage of acquiring multiple TIs is that the signal can be sampled near the peak of the kinetic signal curve across a wide range of ATTs. Assuming the noise magnitude is equal at all ATTs and TIs, multi-TI protocols can, therefore, achieve higher average SNR across a wide range of ATTs than sampling at a single-TI.

The LDs can also be varied for multi-TI protocols, rather than just the PLDs, and it is likely that LDs longer than 1.4 s, as used in Chapter 4, will also be beneficial for estimating CBF with multi-TI protocols. This cannot be evaluated for multi-TI protocols as easily as for single-TI protocols, where the SNR efficiency can simply be used, because the CBF variance is not simply proportional to the noise variance. Instead, the Cramér-Rao lower bound can be used to calculate the CBF and ATT variance for any given protocol.

It is common to use a fixed LD while varying the PLD, but it is also possible to either vary just the LD[8] or both the LD and PLD.[3,9] Therefore, the two types of sequential multi-TI protocol designs explored here were: (1) a fixed (but also optimised) LD with varying PLDs ($\text{Seq}_{\text{single-LD}}$), and (2) an independently varying LD with each varying PLD ($\text{Seq}_{\text{multi-LD}}$). These two protocol designs encapsulate all of the previously proposed sequential multi-TI designs in the literature.

5.2.3 Time-Encoded Multi-TI

Time-encoded multi-TI methods switch between the PCASL label and control conditions during the ASL preparation to encode the LDs and PLDs.[1,2,10] The most time-efficient encoding is a Hadamard-encoding scheme, where each encoded preparation is orthogonal to the others. In this way, N perfusion-weighted images with different PLDs (and possibly different LDs as well) can be decoded from N+1 encoded images, rather than from $2{\cdot}N$ images for the sequential multi-TI method.

The decoding process also results in signal averaging, because all N+1 images are used to decode each of the N TIs, giving a $\sqrt{(N+1)/2}$ reduction in noise SD compared to a single conventional label and control per TI. Therefore, for matched protocol timings, a time-encoded protocol can acquire the same data as a sequential protocol, but in $(N+1)/(2\cdot N)$ of the time with a $\sqrt{(N+1)/2}$ times smaller noise SD.[2]

However, the PLDs, and therefore the temporal sampling, are dictated by the LDs of each sub-bolus and the final PLD. This results in a trade-off between optimal temporal sampling (PLDs) and tracer accumulation (LDs), meaning the previously stated theory may not predict time-efficiency in practice when parameter estimation accuracy is the aim.

The time-encoded protocol designs explored in this work were: (1) each sub-bolus has the same fixed LD ($\text{Had}_{\text{fixed}}$);[10] (2) T_1-adjusted LDs ($\text{Had}_{T1\text{-adj}}$), where the LDs are chosen such that at the readout excitation time the total signal from each block is equal (adjusting for T_1 decay);[11] and (3) the free-lunch method ($\text{Had}_{\text{free-lunch}}$), where the first block has the same LD as the single-TI method, with the other blocks filling the remaining preparation time.[11] For the remaining $\text{Had}_{\text{free-lunch}}$ blocks, the use of both fixed-duration and T_1-adjusted LDs were explored.

5.2.4 Hybrid Multi-TI

A novel multi-TI protocol design, which is a hybrid of the time-encoded and sequential methods, was also investigated. Rather than using a fixed final PLD after the time-encoded PCASL preparation and acquiring multiple averages, the final PLD is allowed to sequentially vary for each full encoding matrix. This gives $N{\cdot}M$ TIs, where M is the number of LDs encoded into the PCASL preparation and N is the number of final PLDs. This design is able to trade-off the superior noise averaging of the time-encoding methods (with larger encoding matrices resulting in more signal averaging) and the increased signal accumulation from longer LDs (which are typically only possible with smaller encoding matrices). A further advantage of smaller encoding matrices is the increased flexibility of the decoded PLDs, because, with multiple different final PLDs, the effective PLDs will not only be determined by the LDs of each block.

For example, with this approach a lower dimension Hadamard-encoding matrix could be used, such as a 4 × 3 scheme, while still maintaining high temporal resolution by using several different final PLDs. This would allow longer LDs than is typically achieved with time-encoded preparations, as well as increasing the flexibility of the timings that can be acquired. In this work, hybrid fixed-duration schemes (Hybrid$_{\text{fixed}}$) and hybrid T_1-adjusted schemes (Hybrid$_{T_1\text{-adj}}$) were explored, where the same encoding matrix is acquired with N independent final PLDs.

5.3 Sequence Implementation

In order to evaluate the in vivo performance of the protocol designs outlined above, several changes were made to the 2D EPI PCASL pulse sequence described in Chapter 4. A simplified schematic of the final pulse sequence is shown and briefly described in Figure 5.2.

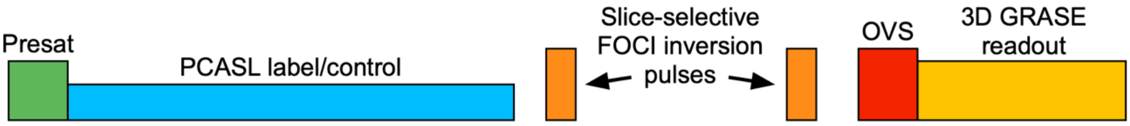

Figure 5.2: A simplified schematic of the PCASL sequence with a 3D GRASE readout. The WET presaturation module is unchanged from Chapter 4, except RF spoiling is now used. The slice-selective FOCI pulses are used for background suppression in place of the global hyperbolic-secant pulses so that they can be interleaved with the PCASL pulse train. The outer-volume suppression (OVS) module is used to improve the 3D GRASE slab profile by saturating the magnetisation just outside of the imaging volume. The OVS pulses are run during the PLD. The 3D GRASE readout is largely as described in Chapter 2, though minor modifications are described below.

The sequence was updated so that multiple LDs could be used. This enabled the sequence to run protocols that either used a single LD with multiple PLDs, multiple LDs with a single PLD, or multiple LDs which are each paired with a specific PLD.

A time-encoded preparation was incorporated which enabled the encoding of label and control conditions into the PCASL preparation. The size of the Hadamard-encoding matrix is flexible and is set by the number of LDs defined in the protocol, while a single PLD is used for all encoded images. The encoded images, corresponding to each row of the Hadamard matrix, were looped over first, before acquiring any averages. This is similar to the sequential case, where each LD and PLD are looped through before acquiring the averages.

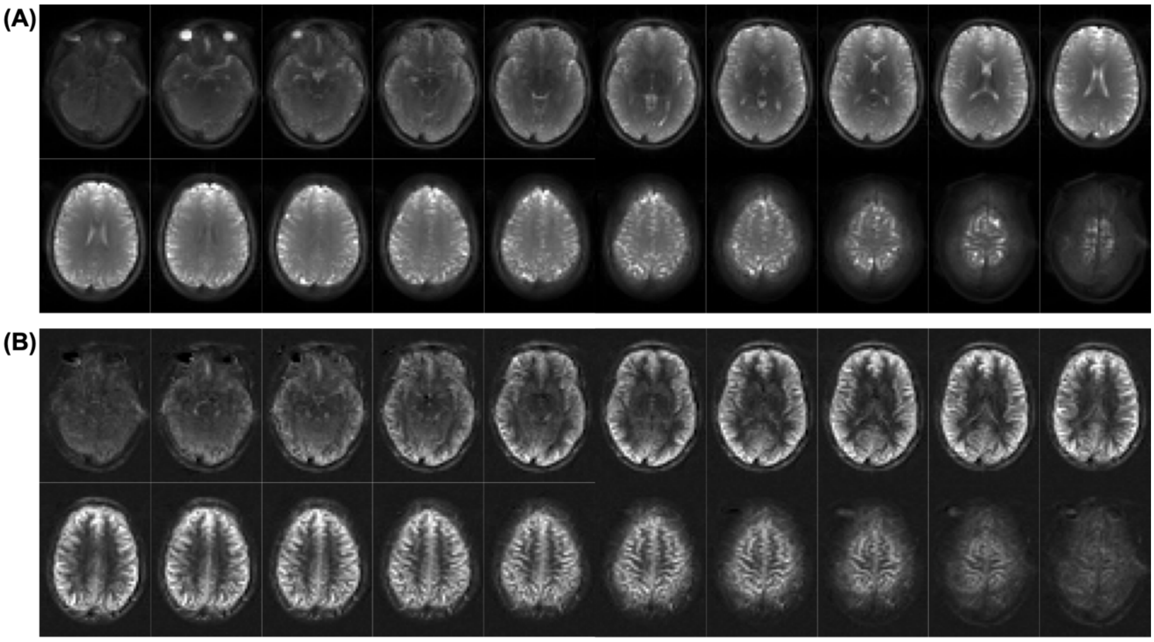

Figure 5.3: Example in vivo images from one subject for all 20 slices using the final 3D GRASE sequence settings described in Appendix C and the methods section of this chapter. (A) The M_0 acquisition. (B) The mean perfusion-weighted image from the single-TI protocol described in Table 5.2. The images have been corrected for motion and distortion but not for the receive-coil sensitivity profile.

5.3.1 3D GRASE Readout

Due to the higher SNR of 3D readouts, as well as the identical PLDs and BGS across slices, a 3D GRASE readout[12,13] was incorporated into the sequence in place of the 2D EPI readout using an existing software library.

Several experiments were performed to improve the default slice profile of the existing 3D GRASE profile, including the use of outer-volume suppression.[13] These experiments are described in Appendix C.

Example M_0 and perfusion-weighted images are shown in Figure 5.3 to demonstrate the final image quality. Reasonable SNR in the perfusion-weighted images has been achieved. However, due to the long echo-train of the 3D readout, blurring is evident in the top slices of the images. This could be further reduced with the use of acquisition segmentation and acceleration, though these options are not explored in this thesis.

5.3.2 Background Suppression

The BGS inversion pulses used in Chapter 4 were restricted to only play out during the PLD, referred to here as $\text{BGS}_{\text{Constrained}}$. As shown in Appendix B Figure B.1, this resulted in variable static tissue signal across PLDs, which led to a variable noise level.

An alternative timing strategy is to allow the inversion pulses to play out during both the PLD and the PCASL labelling, referred to as $\text{BGS}_{\text{Flexible}}$, thus increasing the flexibility of the inversion timings.[14] This is a particularly attractive strategy for time-encoded PCASL protocols, where the PLD is typically very short and has previously been demonstrated to result in a negligible loss of label.[15]

Figure 5.4 demonstrates the effect of the two inversion pulse timing strategies on the static tissue signal for a range of T_1 values. Note, $T_{1,opt}$ was set to 700 ms to ensure good suppression of approximate WM and GM T_1 values at 3T (700 - 1100 ms[16–21] and 1200 - 1800 ms[18–22], respectively). The timings are also set to null signal 100 ms before excitation for both cases, which was found to be robust for ensuring positive signal in all cases. This is important because magnitude subtraction is used in this chapter and Chapter 4. Static tissue signal magnitude is typically much larger than the labelled blood signal, so if the static tissue signal were inverted, the ASL contrast would also be inverted. Since T_1 varies regionally, and the static tissue signal may not be inverted for all T_1 values, it is difficult to correct for this. If the static tissue signal magnitude was smaller than the ASL signal, the ASL contrast would be reduced, making quantification difficult.

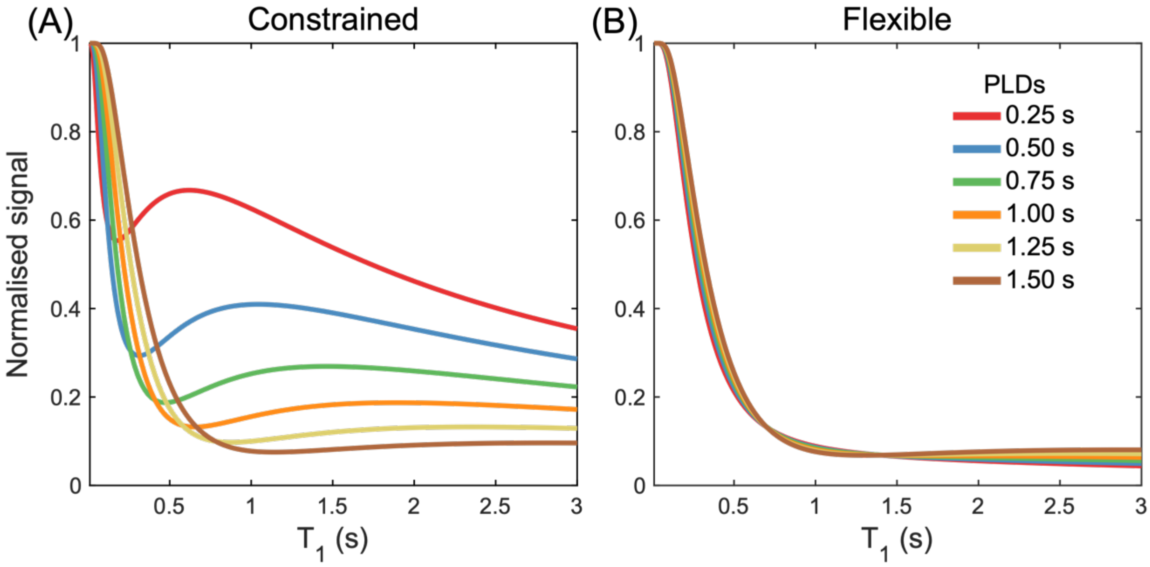

Figure 5.4: The theoretical residual static tissue longitudinal magnetisation at the readout time after the presaturation module and two inversion pulses across a range of T_1 values for $T_{1,opt}$ = 700 ms. (A) The residual magnetisation for the constrained inversion pulse timing strategy ($BGS_{Constrained}$); (B) the residual magnetisation for the flexible timing strategy ($BGS_{Flexible}$). For both (A) and (B), the residual longitudinal magnetisation is shown for 6 different PLDs (0.25 - 1.5 s) and a LD of 1.4 s, which demonstrates that the $BGS_{Flexible}$ strategy results in a much more comparable BGS performance across PLDs. Instantaneous pulse durations, perfect spoiling, and perfect inversion are assumed. The null-time has been set to 100 ms before the excitation, to ensure positive signal in all cases.

It is evident that the $BGS_{Flexible}$ timing strategy results in a more comparable BGS performance across PLDs than the $BGS_{Constrained}$ strategy, which better meets the protocol optimisation assumption in Chapter 4 of a uniform noise magnitude across PLDs.

When interleaved inversion pulses are used, they must be slice-selective to ensure that only previously labelled/controlled blood is inverted, while not perturbing blood that is yet to flow past the labelling plane.[14,15] After the first inversion pulse, the PCASL condition must be switched from label/control to control/label until the 2nd inversion pulse (where the condition must be switched back again), to match the previously created, but now inverted, ASL contrast. This process is demonstrated in Figure 5.5 for a single-TI and time-encoded protocol. For consistency between protocols, slice-selective pulses were used even when an inversion pulse was not interleaved.

To ensure that as little ASL contrast was lost to the imperfect slice profile of the inversion pulses as possible, experiments were performed to optimise the pulse slice profile, including the use of FOCI[23] modulation of the hyperbolic secant pulses. These experiments are described in Appendix D.

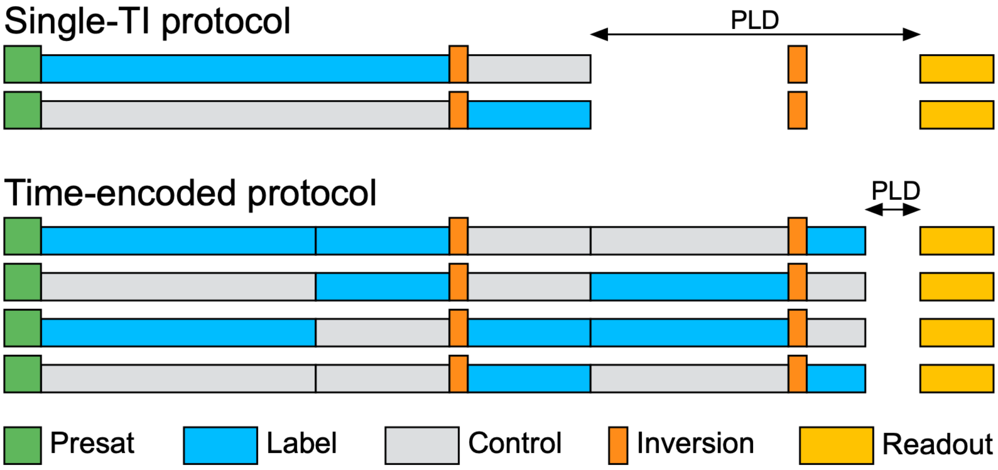

Figure 5.5: A protocol schematic demonstrating the switching of PCASL label and control conditions for a single label and control pair of a sequential protocol and a 4x3 time-encoded protocol. The PCASL preparation switches from label/control to control/label depending on what the condition of the current preparation or encoded block is. The OVS module is not shown for simplicity.

5.4 Methods

5.4.1 Protocol Optimisation

The protocols in this chapter were optimised for CBF accuracy only, with ATT treated as a confounding variable, using the L-optimal cost function in Equation 4.4. This comparison was restricted to CBF optimised protocols for two reasons: 1) CBF is often the main parameter of interest with knowledge of ATT predominantly being used to correct ATT related biases in the CBF estimates, and 2) optimising for only one parameter makes interpretation of the final protocols and their relative parameter estimation accuracy simpler. However, there is nothing to prevent the optimisation framework being used to also optimise ATT accuracy.

The original optimisation algorithm, described in Chapter 4, iterated through each of the N PLDs of the multi-PLD protocol, and for each, searched for the value which minimised the average CRLB across ATTs. For a sequential multi-TI protocol, the LD can also be included in the optimisation by either finding the optimal LD and PLD pair for each of the N sequential TIs ($\mathrm{Seq}_{\text{multi-LD}}$) or by searching for a single optimal LD for the whole scan at each PLD iteration ($\mathrm{Seq}_{\text{single-LD}}$). In both cases, the i^th^ PLD was updated from the list $\mathrm{PLD}_{\min} \leq \mathrm{PLD}_{i-1} \leq \mathrm{PLD}_{i} \leq \mathrm{PLD}_{i+1} \leq \mathrm{PLD}_{\max}$, while the i^{th} LD was only restricted to $\mathrm{LD}_{\min} \leq \mathrm{LD}_{i} \leq \mathrm{LD}_{\max}$. Monotonically increasing LDs were not enforced because it was assumed that it may not necessarily be optimal for a long PLD to be paired with a long LD and a short PLD with a short LD.

For the $\mathrm{Had}_{\text{fixed}}$ and $\mathrm{Had}_{T1\text{-adj}}$ protocol designs, only the first LD and final PLD, for any encoding size, $(N+1) \times N$, must be searched over, making it possible to carry out a global grid search for all possible timing combinations for each N. The $\mathrm{Had}_{\text{free-lunch}}$ protocol differs only in that the first block LD is fixed equal to the single-TI protocol LD, with the other blocks being optimised in the same way as $\mathrm{Had}_{\mathrm{fixed}}$ and $\mathrm{Had}_{T1\text{-adj}}$.

The optimisation of the hybrid protocols was done by iterating through each of the final PLDs and updating the time-encoded LDs on each iteration. The same time-encoded preparation is used with all of the final PLDs, so updating the LDs for one, updates it for the entire protocol.

For each multi-TI protocol design, the number of TIs were optimised for by running the algorithm for a range of TI numbers and finding the protocol which minimised the CRLB. However, they were restricted to $N \leq 15$: sequential protocols $N = 2,\ldots,15$; Hadamard protocols $N = 3,7,11,15$; hybrid protocols $N = 3,6,7,9,11,12,14,15$ (corresponding to: 4×3 encoding with 1-5 PLDs, 8×7 encoding with 1-2 PLDs, 12×11 and 16×15 with 1 PLD).

The optimisation used a uniform ATT prior probability distribution within the range 0.5 - 2 s, which was found to be representative of the healthy subjects scanned in Chapter 4. This ATT distribution was sampled at 1 ms intervals and was extended on either side by 0.3 s with a linearly decreasing weighting to reduce edge effects. A CBF distribution is not required for optimisation (see Chapter 4), so a point prior of 50 mL/100g/min was used. The LD grid search was restricted to 0.1 - 1.8 s, ensuring the minimum LD was greater than 100 ms, as suggested by Teeuwisse et al.,[11] with the longest LD matching the recommended single-TI LD.[7] It is likely that longer LDs than 1.8 s would be beneficial for multi-TI protocols, as they are for single-TI protocols,[5,6] but the maximum LD was restricted to 1.8 s to maintain a similar level of signal dependence on tissue T_1 across protocols. The PLD search list was 0.075 - 2.3 s. Both lists were sampled at 25 ms intervals. Variable TR and a maximum scan time of 5 minutes was assumed, as before.

The CRLB-based cost function used the simplified sensitivity functions derived in Chapter 4 (Equations 4.6 and 4.7) which used the standard CASL model from Buxton et al.,[24] the model parameters given in Table 5.1, and assumed white Gaussian noise. The noise magnitude was calculated from preliminary in vivo data (noise SD of label and control image = 1.3×10^{-3} relative to M_0).

Further specific details for each protocol optimisation are as follows. Single-TI: LD 1.8 s and PLD 2 s (ensuring PLD ≥ ATT). $\text{Seq}_{\text{single-LD}}$: initialised with LD 1.8 s and N PLDs spaced evenly between 0.075 - 2.3 s. $\text{Seq}_{\text{multi-LD}}$: N LD and PLD pairs randomly initialised between 0.8 - 1.8 s and 0.075 - 2.3 s, respectively. The sequential protocol optimisations iterated through each TI in a randomly permuted manner and were run 20 times each for robustness. $\text{Had}_{\text{free-lunch}}$: the first LD was fixed at 1.8 s (matching Single-TI) with the remaining LDs being either fixed-duration or T_1-adjusted; the final PLD was also optimised for, therefore not guaranteeing that the first LD and PLD will match the Single-TI protocol. The $\text{Had}_{\text{fixed}}$ and $\text{Had}_{T1\text{-adj}}$ protocols did not require initialisation because the entire timing parameter space could be evaluated in one step. $\text{Hybrid}_{\text{fixed}}$ and $\text{Hybrid}_{T1\text{-adj}}$: all N final PLDs initialised at 0.075 s - the LDs did not require initialisation because they are all globally optimised at each step, similar to the time-encoded protocols.

5.4.2 Monte Carlo Simulations

Monte Carlo simulations were performed to investigate the performance of the optimised protocols under ideal conditions where the ground truth is known. Simulated data were generated for each protocol using the standard CASL model[24] with the parameters in Table 5.1 for ATTs between 0.5 - 2 s, sampled every 0.01 s. Gaussian white noise was added to 2000 replicas of the label and control (or encoded) data at each ATT sample, using the same noise magnitude as the protocol optimisations above. The noisy perfusion data at each PLD was then decoded according to the encoding scheme for each protocol. The data were then fit, and the estimates compared, as described in Sections 5.4.3.3 and 5.4.3.4 below.

5.4.3 In Vivo Experiments

5.4.3.1 Acquisition

To investigate the relative performance of each protocol in vivo, 10 healthy volunteers (5 female, mean age 30.7, range 24 - 40) were recruited and scanned under a technical development protocol, agreed with local ethics and institutional committees, on a 3T Prisma system (Siemens Healthcare, Erlangen, Germany) with a 32-channel receive-only head coil.

All scanning occurred during a single session for each subject (total scan duration ~50 minutes). Volunteers were asked to lie still and stay awake throughout the scan. A nature documentary was shown to help maintain a similar level of alertness throughout the session.

Each session began with a 3-plane localiser followed by a 3D single-slab TOF angiography sequence used to place the PCASL labelling plane, as described in Chapter 4. A 3D T1-weighted MPRAGE sequence (resolution $1.5 \times 1.5 \times 1.5$ mm^3) was then acquired for generating the brain mask and grey matter (GM) mask used in the analysis. Following this, 4 calibration images were acquired with the same readout module as the PCASL data (see below). Two of these had right-left (R-L) in-plane phase encoding and two had L-R, used to correct off-resonance distortions. Finally, the ASL scans were acquired in a pseudo-randomly permuted order for each subject to reduce the impact of physiological-drift from biasing the comparison.

The imaging parameters for the ASL scans were: single-shot 3D gradient and spin-echo (GRASE) readout,[12,13] TE 28.5 ms, excitation flip-angle 90°, refocussing flip-angle

120°,[25,26] FOV $230 \times 168 \times 100$ mm^3, matrix $64 \times 46 \times 20$, spatial resolution $3.59 \times 3.65 \times 5$ mm^3 (an odd number of phase-encode lines are acquired, but an even number of columns are reconstructed), 20 acquired partitions, no acceleration, no slice-oversampling, centric partition ordering, bandwidth 2298 Hz/px, total readout duration 583 ms, spectrally-selective fat saturation. A rectangular FOV with R-L in-plane phase encoding was used to reduce the echo train length (747 ms readout duration with square FOV). The imaging slab was placed in the transverse plane with the superior edge flush with the top of the brain. Outer-volume suppression (OVS), using a cosine-modulated water suppression enhanced through T_1 effects (WET) module,[27,28] was used to reduce slab profile errors aliasing in the through-plane direction (see Appendix C), similar to Günther et al.[13] Readout, phase-encode and 3D encode directions were anterior-posterior, right-left, and feet-head, respectively. Variable TR was used with all PCASL protocols to maximise scan-time efficiency.

PCASL labelling used 500 μs duration Gaussian RF pulses, 1 ms spacing, 20° flip-angle. Background suppression was performed to reduce motion-induced subtraction errors[29] and other sources of physiological noise using a WET presaturation module[27,28] and two optimised C-shaped FOCI inversion pulses **(**Appendix D**).**[23,30] The timing of the inversion pulses used Equation 3.4 to null $T_{1,opt}$ and 2 x $T_{1,opt}$ ($T_{1,opt}$ = 700 ms) 100 ms before excitation. The inversion pulses were able to interleave with the PCASL labelling, similar to Dai et al.[14,15] The transverse presaturation and inversion slabs were manually placed so that the inferior edge intersected the centre of the labelling plane and the superior edge was above the skull.

The calibration images were acquired using the WET presaturation module followed by a 10 s delay to allow controlled and near-complete magnetisation recovery before the 3D-GRASE readout.

5.4.3.2 Preprocessing

The preprocessing of the in vivo data made use of a range of tools in the FSL (FMRIB software library) toolbox.[31]

The individual ASL scans were first motion-corrected using rigid-body registration with MCFLIRT.[32] The temporal-mean of each was then registered to the first calibration image using FLIRT.[32,33] These transformations were applied in a single step using sinc interpolation to reduce interpolation smoothing.

The susceptibility-induced off-resonance field was estimated with TOPUP,[34] using the 4 calibration images with opposite in-plane phase-encode directions. The 2 R-L phase-encoded images were then motion-corrected using MCFLIRT, distortion corrected using TOPUP, and averaged. Distortion correction was also applied to the motion-corrected ASL data.

The T_1-weighted MPRAGE structural image was brain-extracted using BET,[35] generating a brain mask. Tissue segmentation was then performed using FAST[36] to create GM partial volume estimates (PVE). The undistorted calibration image was registered to the structural image using the rigid-body boundary-based registration in FLIRT, providing a structural to calibration transformation. The brain mask and GM PVE were then registered to the calibration image using this transformation and had thresholds applied (brain mask 90%, GM mask 50%). This brain mask was applied to the calibration image and ASL data.

The edges of the masked calibration image were eroded before being expanded using a mean filter and brain masked again, to remove low-intensity voxels at the edge of the brain which can lead to erroneously high CBF values during the voxelwise calibration step. It was then smoothed (Gaussian kernel, $\sigma = 2.5$ mm) to improve SNR, as recommended.[7]

Perfusion-weighted images (PWI) were then generated by subtracting/decoding the undistorted, motion-corrected, ASL images. These PWI images were finally calibrated prior to fitting to account for global scaling factors by multiplying by the blood–brain partition coefficient ($\lambda = 0.9$ mL/g)[37], dividing by the labelling efficiency ($\alpha = 0.85$)[38] and voxelwise dividing by the smoothed calibration image.

Parameter	**Value**
Model	
T_1 of blood (T_{1b})	1.65 s[39]
T_1 of tissue (T_{1t})	1.445 s[16]
Labeling efficiency (α)	0.85[38]
Brain/blood water partition coefficient (λ)	0.9 mL/g[37]
Sequence	
RF labeling pulse duration	500 µs duration (Gaussian)
RF labeling pulse separation	1 ms
RF labeling flip angle	20°
Mean labeling gradient	0.8 mT/m
Gradient during labeling pulses	6 mT/m
Echo time (TE)	28.5 ms
Readout duration	583 ms
Analysis	
CBF prior	0 ± 10^6 mL/100g/min
ATT prior	1.3 ± 10^6 s

Table 5.1: Model and sequence parameters used in the optimisations, Monte Carlo simulations and in vivo experiments.

5.4.3.3 Model fitting

CBF and ATT were estimated in the same way for the multi-TI simulated data and the in vivo perfusion data using a variational Bayesian inference algorithm as implemented in

FSL's BASIL.[40] In each voxel, this approach not only provides estimates of CBF and ATT but also uncertainty estimates in the form of the standard deviation of the marginal posterior distributions. The standard CASL model[24] was used with the parameters in Table 5.1. Fitting was initialised with a coarse grid search to avoid local minima bounded by $0 \leq \text{CBF} \leq 200$ mL/100g/min and $0 \leq \text{ATT} \leq 2.5$ s and sampled every 1 mL/100g/min and 0.01 s. The BASIL fitting priors were noninformative (CBF 0 ± 10^6 mL/100g/min and ATT 1.3 ± 10^6 s) to minimise bias in the resulting parameter estimates. Negative CBF and ATT estimates were set equal to zero.

CBF was estimated from the single-TI data using BASIL with the same parameters and model as the multi-timepoint data but with ATT fixed at 1.3 s. This ATT value was chosen because it minimised the theoretical CBF bias caused by an assumed ATT over the ATT range 0.5 - 2 s for a single PLD at 2 s, similar to Figure 4.5 but using a different LD, PLD, and ATT range.

5.4.3.4 Statistical Analysis

The CBF estimates for each protocol were compared in three different ways to more thoroughly examine the data than in Chapter 4: (1) the posterior probability distribution SDs output by BASIL were used as a measure of uncertainty (or confidence) in the CBF estimates,[40] which is sensitive to the amount of CBF information, noise, and other non-perfusion related effects, such as macrovascular signal and imaging artefacts in the data; (2) the root-mean-squared-error (RMSE) relative to ground truth estimates (see below) were used as a measure of accuracy, incorporating both systematic bias and random noise contributions (similar to Chapter 4); and (3) the test-retest RMSE for each scan was calculated by splitting the data into two 2.5 minute data sets and separately fitting each half, giving a measure of reproducibility, which is independent of any estimated ground truth or uncertainties from the fitting process.

Note, for (3) the $\mathrm{Had}_{T1\text{-adj}}$ data was split into the first 4 and last 5 averages while the $\mathrm{Hybrid}_{\mathrm{fixed}}$ errors were not calculated because there were only 3 averages (see Table 5.2).

5.4.3.4.1 In Vivo Specific Analysis

Each of the 3 different comparison metrics were calculated on a voxelwise and subject-wise basis. Standard errors for the voxelwise metrics were calculated by bootstrap sampling across the 10 subjects using 1000 samples, where the relevant statistical measure (mean SD, RMSE, or test-retest RMSE) is performed on each sample. Each sample is a selection of 10 randomly chosen subjects, selected with replacement, which means that a given sample could contain multiple copies of the same subject's data. The SDs generated from these sampling distributions reflect the variability that the sampled subjects contribute to the voxelwise metrics.

Ground truth CBF and ATT estimates were generated for the in vivo data in a similar manner to Chapter 4 by combining all of the data from each protocol and fitting for CBF and ATT as described above. To account for the different noise levels in the non-Hadamard-encoded, Hadamard-encoded and hybrid data (caused by the averaging effect of Hadamard-decoding), BASIL was given 3 noise magnitudes to estimate, which is similar to weighted NLLS fitting (see Chappell et al.[40]). To investigate whether these ground truth estimates were biased towards particular protocols and whether modelling the three different noise magnitudes is beneficial, ground truth estimates for the MC simulation data were generated in the same way.

Only voxels within the GM masks were used in the analysis. To eliminate poorly fit ground truth data from the analysis, Tukey's outer-fence method[41] was used, where voxels with ground truth posterior SDs of more than 3 times the inter-quartile range above the 75th percentile for either CBF or ATT were excluded. This resulted in upper bounds on the ground truth posterior SDs of 2.9 mL/100g/min and 0.061 s, respectively. Furthermore, voxels were also excluded if the posterior SDs for any individual scan were > 500 mL/100g/min and 50 s, which would suggest an extremely poor fit, perhaps arising from motion or other artefacts, and could bias the resulting comparison.

All significance testing used a two-sided paired Wilcoxon signed-rank test with Bonferroni correction for 6 comparisons (α = 0.05).

5.5 Results

5.5.1 Optimised protocols

Protocol	*Label durations (ms)*	*Post-label delays (ms)*	*Averages*	*Total acquisitions*	*Scan duration*
Single-TI	1800	2000	34	64	5:02 min
***Seq**single-LD*	1800	175, 1050, 1425, 1725, 2075, 2200, 2300, 2300, 2300	4	72	5:00 min
***Seq**multi-LD*	1800, 1500, 1800, 1800, 1800, 1800, 1800, 1800, 1800, 1800, 1800, 1800	200, 825, 1350, 1475, 1800, 1850, 2100, 2225, 2300, 2300, 2300, 2300	3	72	5:00 min
***Had**fixed*	550, 550, 550, 550, 550, 550, 550	100	8	64	4:54 min
***Had**T1-adj*	1150, 675, 475, 375, 300, 250, 225	75	9	72	5:00 min
***Had**free-lunch*	1800, 625, 450, 350, 300, 250, 225	125	8	64	5:05 min
***Hybrid**fixed*	1275, 1275, 1275	75, 150, 600, 850, 1000	3	60	5:00 min
***Hybrid**T1-adj*	1800, 850, 550	200, 650, 900, 900	4	64	4:49 min

Table 5.2: Optimised protocol timings for the protocol designs compared in this work

The optimised timings for each protocol are shown in Table 5.2 with the predicted CBF errors (CRLBs) as a function of ATT shown in Figure 5.10. All of the multi-TI

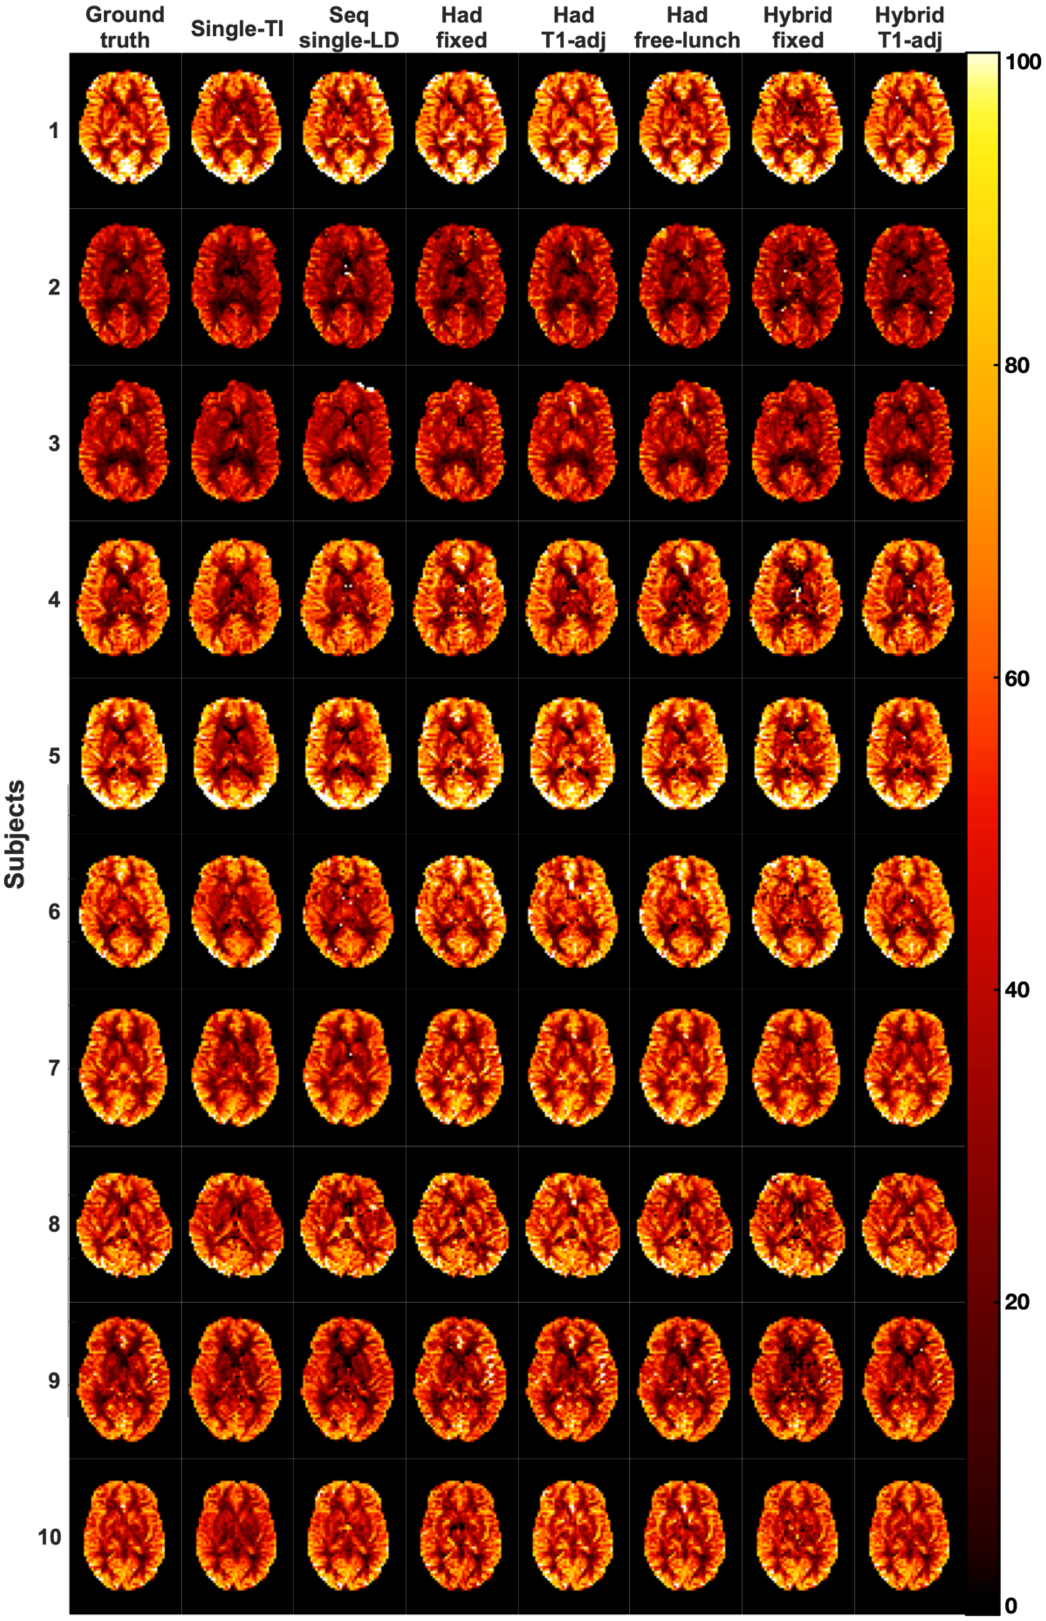

Figure 5.6: A single slice of the CBF maps for each subject for each of the protocols and the ground truth estimates. The colourmap is in units of mL/100g/min

protocols have effective PLDs which span at least the range 0.2 to 2.3 s, allowing them to sample near the peak signal of the shortest and longest ATTs in the range of interest.

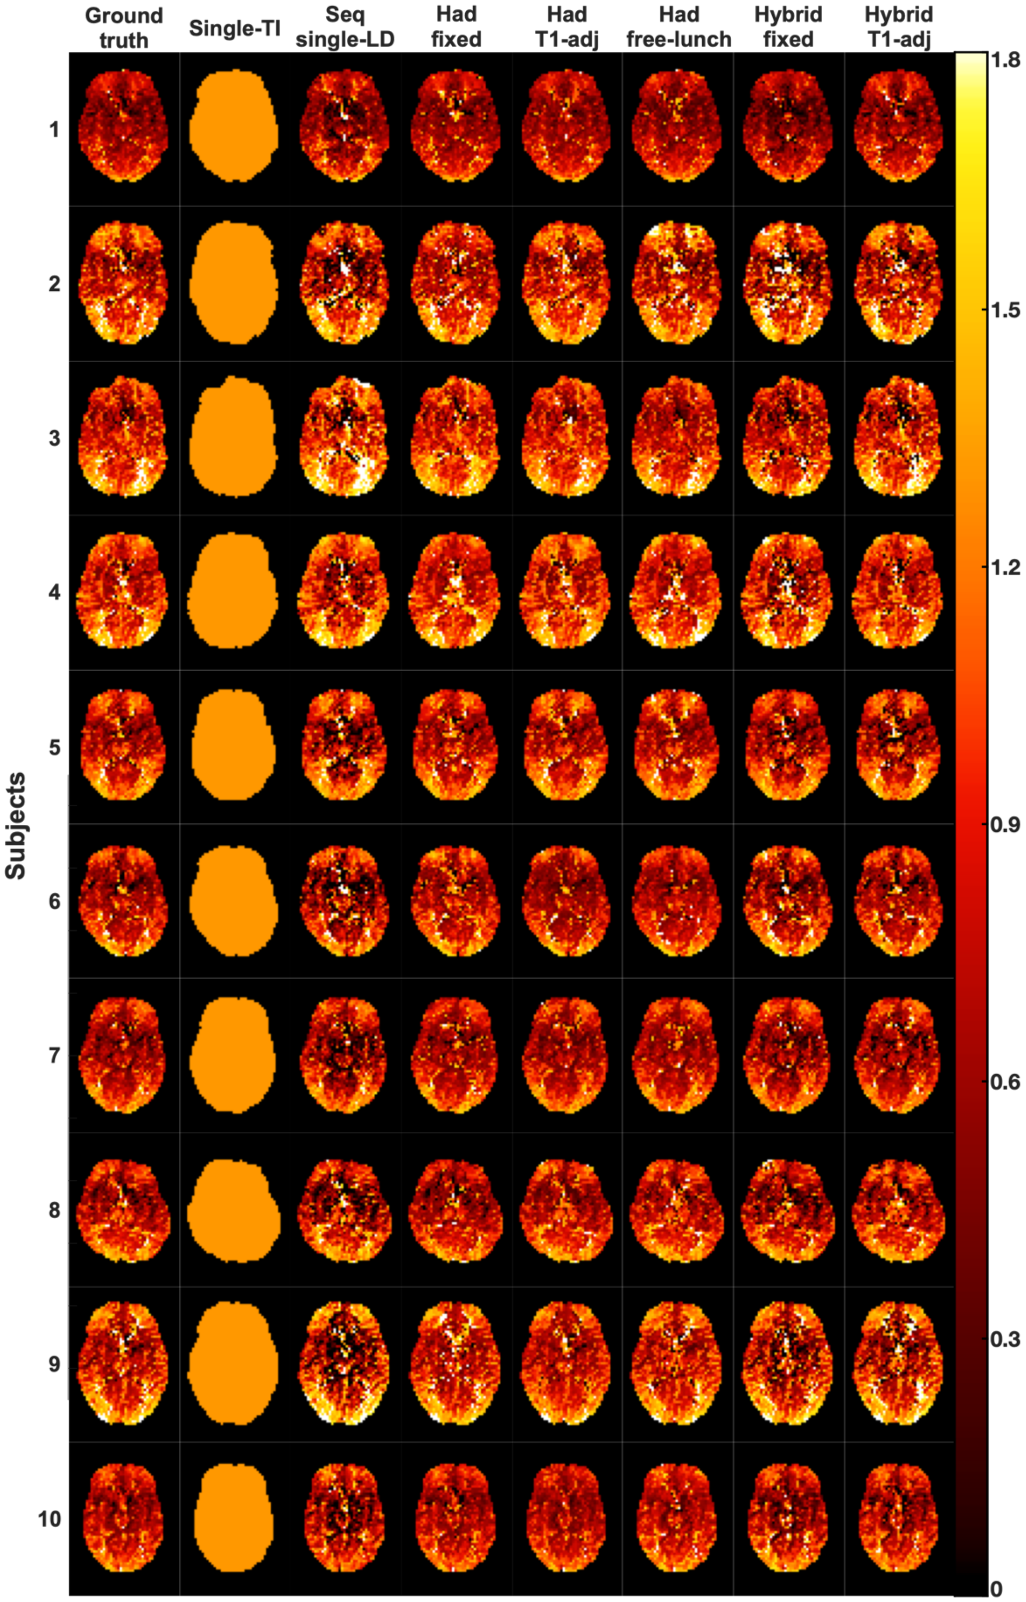

Figure 5.7: A single slice of the ATT maps for each subject for each of the protocols and the ground truth estimates. The colourmap is in units of seconds.

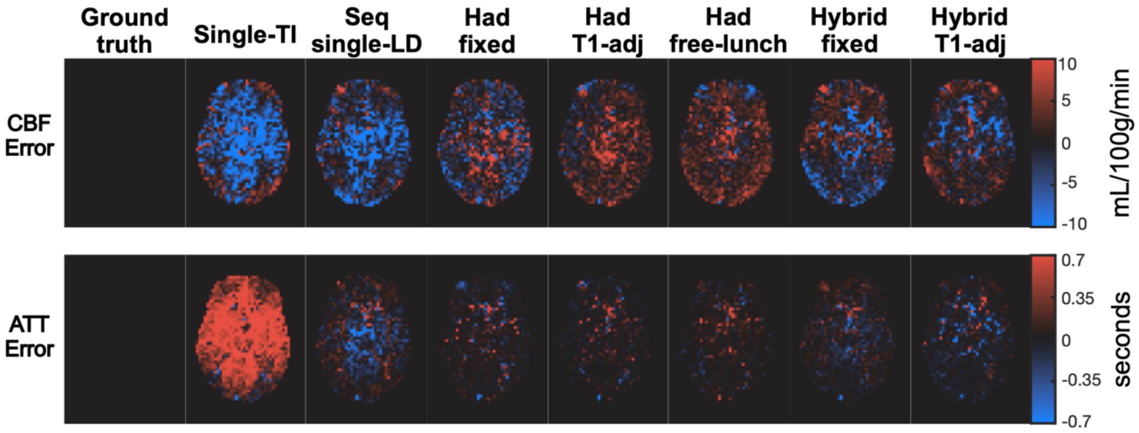

Figure 5.8: Example CBF and ATT error maps relative to the ground truth estimates. The error maps are shown for subject 7 for the same slice shown in Figure 5.6 and Figure 5.7.

The Seq$_{\text{single-LD}}$ optimisation chose a LD of 1.8 s with 9 PLDs of increasing density at long PLDs. The pattern of the PLDs is similar to the CBF optimised protocol in Chapter 4 but differs slightly because here because a 3D readout is used. The Seq$_{\text{multi-LD}}$ protocol chose a similar set of PLDs to Seq$_{\text{single-LD}}$ and all but one of the 12 LDs was 1.8 s. Due to the similarity with the Seq$_{\text{single-LD}}$ protocol and the marginal improvement in predicted CBF errors, Seq$_{\text{multi-LD}}$ was not used in further comparisons.

For the Had$_{\text{fixed}}$ protocol, a 4 × 3 encoding with LDs of 1.125 s had 15% lower average simulated CBF uncertainty than the 8 × 7 encoding in simulations. However, due to the unknown performance of the 4 × 3 encoding in vivo and the more common use of the 8 × 7 encoding, the optimised 8 × 7 timings were used in further comparisons.

A 4 × 3 encoding was also optimal for the Had$_{\text{free-lunch}}$ protocol, with the last 2 LDs using fixed LDs (1100 ms). When only the more common 8 × 7 encodings or larger were considered, an 8 × 7 encoding with T_1-adjusted LDs for the final 6 LDs (shown in Figure 5.1) was optimal. The 8 × 7 protocol was used in further comparisons. An 8 × 7 encoding was optimal for the Had$_{T_1\text{-adj}}$ protocol.

For the hybrid protocols, the Hybrid$_{\text{fixed}}$ optimisation chose a 4 × 3 encoding with an LD of 1.275 s and 5 PLDs, resulting in 15 TIs. The Hybrid$_{TI\text{-adj}}$ optimisation also chose a 4 × 3 encoding scheme with a first LD of 1.8 s and 4 PLDs, giving 9 unique TIs.

5.5.2 Comparisons

A single slice from the CBF and ATT maps from all of the subjects are shown in Figure 5.6 and Figure 5.7. The CBF maps for all protocols are consistent and of good quality, but show some subtle differences. The ATT maps share similar major features such as longer transit times in the occipital lobe, but there are also large variations in transit times near the expected position of the large arteries. Large differences in the CBF and ATT maps can be seen across subjects. This is likely due to the range in age and sex of the subjects. For example, subjects 2 and 3, who both have lower mean CBF than the other subjects, are both male and are near the upper end of the age range while subject 1 is female and at the lower end of the age range. When partial volume correction is not performed, as in this work, global perfusion has been found to be negatively correlated with age.[42,43] Mean CBF has also been found to be higher in women,[44] which may be due to average differences between the sexes in haematocrit fraction.[45] Other possible causes of the differences in CBF seen across subjects may include differing blood flow velocities, which could affect the PCASL labelling efficiency, and the true subject-dependent T_{1t} values. Mean ATT has also been found to be lower for women[46] and longer in older subjects.[47]

Figure 5.8 shows the CBF and ATT error maps for subject 7 relative to the ground truth estimates, demonstrating the spatial patterns of the errors for each protocol. Underestimation of the CBF estimates for the single-TI and Seq$_{\text{single-LD}}$ protocols can be seen while the errors in the single-TI ATT map reflect the incorrect assumption of a single brain-wide ATT value.

5.5.2.1 In Vivo Data Selection

There were a total of 79211 voxels in the GM masks across all 10 subjects. Of these, 6.2% were excluded due to poor ground truth CBF and ATT fits (posterior SDs > 2.9 mL/100g/min or > 0.061 s) and a further 4.1% were excluded because there were extremely poor fits in one or more of the individual scans (posterior SDs > 500 mL/100g/min or > 50 s). Of the included voxels, 90% of the ground truth ATTs lay between 0.5 - 1.51 s (5th - 95th percentiles, median = 0.97 s). Figure 5.9 shows that the excluded voxels are consistently located where one would expect the large arteries to be, suggesting that the exclusion criteria have mostly excluded data contaminated by macrovascular signal.

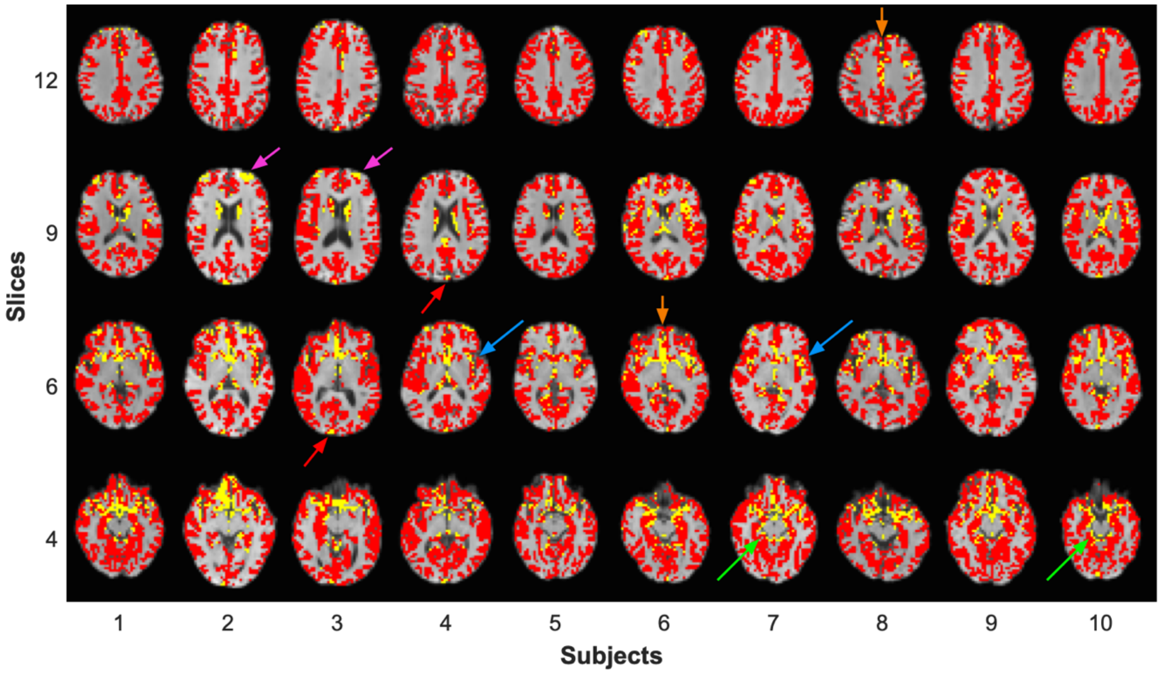

Figure 5.9: Maps of the voxels excluded from the in vivo analysis (yellow) overlaid on the GM masks (red) with the T_1-weighted structural images (grey) behind for use as an anatomical reference. Four slices are shown for each of the 10 subjects. The excluded voxels are consistently located near the circle of Willis (green arrows), middle cerebral arteries (blue arrows) and anterior cerebral arteries (orange arrows), for each subject. The superior sagittal sinus (red arrows) and blurred signal from the eyes (pink arrows) can also be seen in some subjects.

5.5.2.2 Trends across ATTs

The predicted CBF errors (CRLBs) for the existing and novel hybrid protocols are shown in Figure 5.10(A,D) as a function of ATT. The single-TI error was flat across the ATT range because it is only dependent on the noise magnitude, which is assumed to be constant across all ATTs - the RMSEs plotted in Chapter 4 also included CBF bias which varies across ATTs and so were not flat. The multi-TI protocols all had increasing CBF errors as the ATT increased. The sharp changes in error across ATTs for the multi-TI protocols are where ATT = PLD or ATT = LD + PLD for one of the LD/PLD pairs; at these points either the inflow of the tracer begins to be sampled, which leads to an increase in sensitivity to ATT, causing an increase in CBF variance overall, or the inflow of the tracer is not sampled and so this data point is not sensitive to CBF at all.

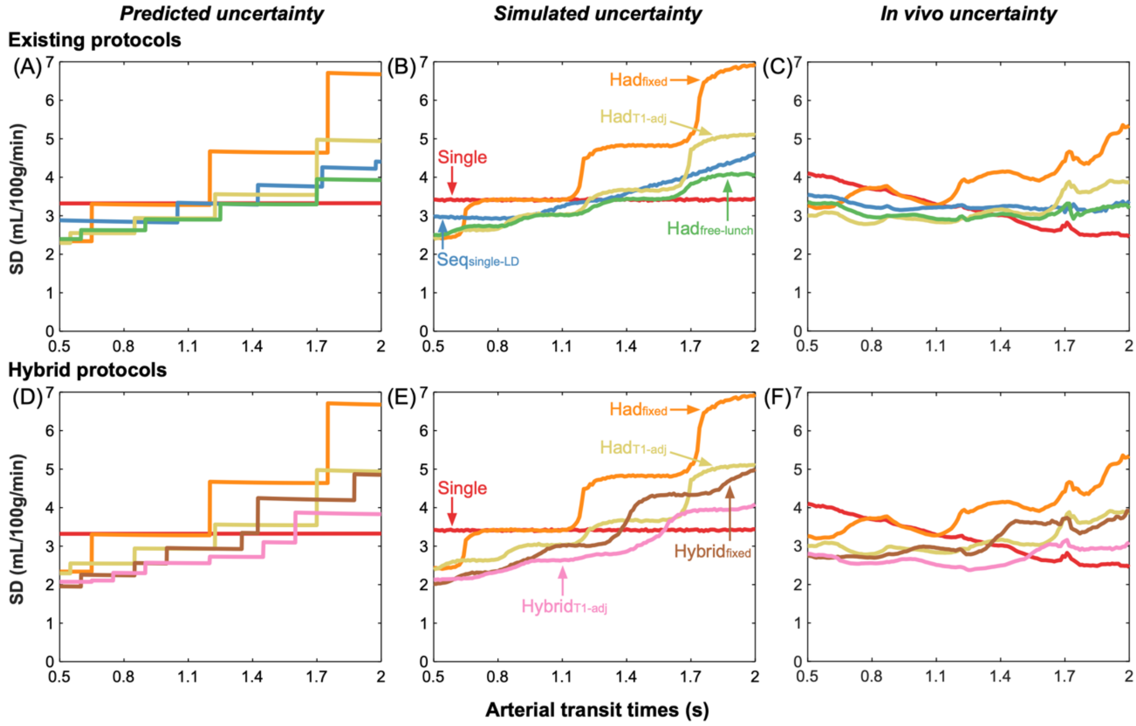

Figure 5.10: The predicted (Cramér-Rao lower bound SDs) (A,D), simulated (Monte Carlo simulation posterior SDs) (B,E), and in vivo (posterior SDs) (C,F) CBF uncertainty measures for the existing protocols (top) and the proposed hybrid protocols (bottom). The hybrid protocols are shown alongside their non-hybrid analogues and the single-TI protocol to demonstrate the improvement in CBF estimation certainty. For the simulation and in vivo results, the median SD at each ATT is plotted. A sliding window was used to plot the in vivo data, as done in Chapter 4: window size 0.1 s, step size 0.01 s.

Of the existing protocols, $\text{Had}_{\text{fixed}}$ had the highest predicted CBF errors across most of the ATT range. $\text{Had}_{T1\text{-adj}}$ and $\text{Had}_{\text{free-lunch}}$ had similar errors across ATTs, but $\text{Had}_{\text{free-lunch}}$ maintained lower errors at long ATTs. $\text{Seq}_{\text{single-LD}}$ had slightly higher errors than $\text{Had}_{T1\text{-adj}}$, except at long ATTs. Most of the multi-TI protocols had reduced CBF errors at short ATTs compared to the single-TI protocol but were worse at longer ATTs.

Both of the hybrid protocols achieved lower predicted CBF errors at almost all ATTs relative to their non-hybrid analogues, but the $\text{Hybrid}_{T1\text{-adj}}$ protocol also maintained a lower CBF error than the other multi-TI protocols at almost all ATTs and has lower errors than the single-TI protocol for the majority of the distribution.

The median CBF errors from the Monte Carlo simulations, shown as posterior probability SDs from the Bayesian fitting (Figure 5.10(B,E)), follow the trends of the predicted errors extremely closely, validating the expected performance of each protocol under ideal conditions. The jumps in CBF error are visible but are more gradual due to the blurring effect of noise.

The in vivo median CBF errors (plotted using a sliding window, similarly to in Chapter 4), also shown as posterior SD trends, (Figure 5.10(C,F)) exhibit similar relative performance for each protocol, but there is a general decrease in CBF errors at longer ATTs for all protocols compared to the predicted and simulated CBF errors. Similar jumps in SDs where the ATT exceeds one of the chosen PLDs or where ATT = LD + PLD can be seen, especially for the $\text{Had}_{\text{fixed}}$ protocol.

5.5.2.3 Mean In Vivo Uncertainty

The in vivo voxelwise and subjectwise CBF posterior SDs for each protocol, representing the average uncertainty in the CBF estimates across the GM masks from the fitting process, are shown in Figure 5.11. All of the voxel-wise differences between protocols were significant. Of the existing protocols, $\text{Had}_{\text{fixed}}$ had the highest SDs, followed by the single-

TI protocol. This gives the result that some of the multi-TI methods outperform the single-TI method, despite the larger number of averages which can be used in the latter case: Seq$_{\text{single-LD}}$, Had$_{T1\text{-adj}}$ and Had$_{\text{free-lunch}}$ all produced lower CBF SDs than the single-TI protocol, with Had$_{T1\text{-adj}}$ giving the lowest errors of the conventional protocols. Both novel hybrid approaches outperformed their conventional analogues with Hybrid$_{T1\text{-adj}}$ producing the smallest CBF errors of all the protocols considered.

These results are reflected in the subject-wise comparisons, though Single, Seq$_{\text{single-LD}}$, and Had$_{\text{fixed}}$ SDs were not significantly different to each other and neither were Had$_{T1\text{-adj}}$, Had$_{\text{free-lunch}}$, and Hybrid$_{\text{fixed}}$. Hybrid$_{\text{T1-adj}}$ still had significantly smaller SDs than all other protocols at the subject level.

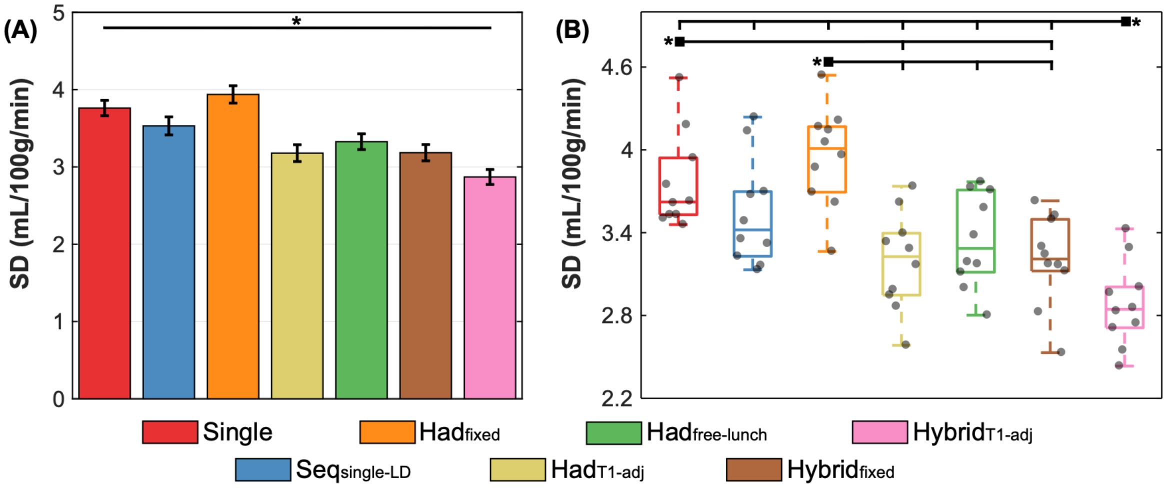

Figure 5.11: The in vivo posterior SDs output by BASIL plotted across all voxels (A) and across subjects (B). In (A), the mean and standard error (see methods) of the posterior SDs across voxels are shown. In (B), the boxplots show the quartiles and range of the subject SDs. All protocols in (A) had significantly different mean posterior SDs to each other. In (B), the protocols with significantly different subject mean SDs are explicitly highlighted.

5.5.2.4 Errors Relative to Ground Truth

The voxelwise and subjectwise CBF RMSEs relative to the ground truth estimates for each protocol are shown in Figure 5.12, demonstrating the accuracy of each protocol, including any systematic bias.

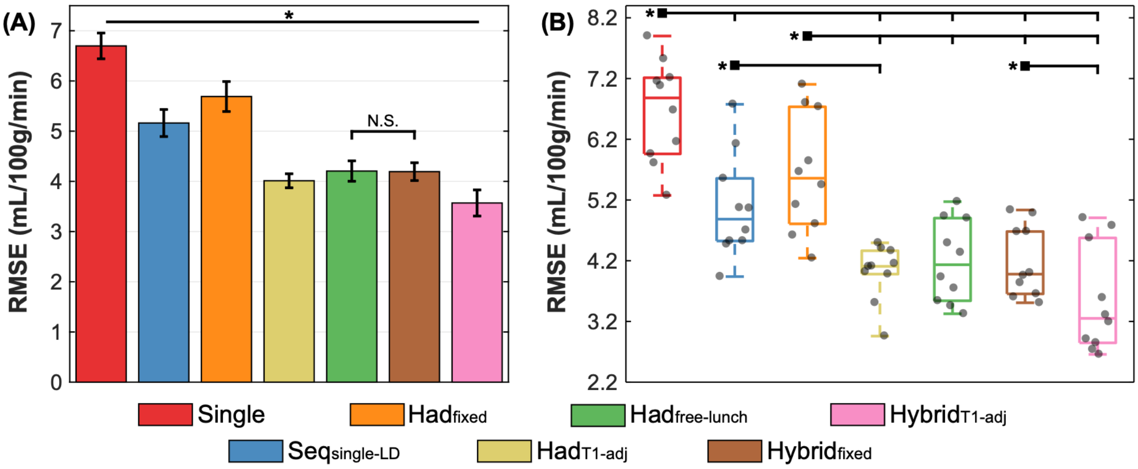

Figure 5.12: The in vivo RMSEs relative to the estimated ground truth values across all voxels (A) and across subjects (B). In (A), the mean and standard error (see methods) of the RMSEs across voxels are shown. In (B), the boxplots show the quartiles and range of the subject RMSEs. All protocols in (A) had significantly different RMSEs to each other except Had$_{free\text{-}lunch}$ and Hybrid$_{fixed}$. In (B), the protocols with significantly different subject RMSEs are highlighted.

The voxelwise results are similar to the relative performances shown for the posterior SDs, but here the single-TI protocol has the highest errors due to the greater bias with this approach. There are also relatively larger errors for Seq$_{\text{single-LD}}$ and Had$_{\text{fixed}}$, compared to the other time-encoded and hybrid protocols. As before, Had$_{T1\text{-adj}}$ had the smallest errors of the conventional protocols, but Hybrid$_{T1\text{-adj}}$ performed the best overall. These results are again reflected in the subjectwise comparisons, although some differences did not reach significance due to the reduced statistical power of this approach.

Figure 5.13 demonstrates that using estimated ground truth values (fitting the combined data from all protocols) in the MC simulations showed no obvious bias towards any particular protocol, justifying the use of these ground truth estimates. However, this analysis also shows that if the different noise magnitudes across the different acquisition schemes were not taken into account during the ground truth fitting, then large biases could be introduced.

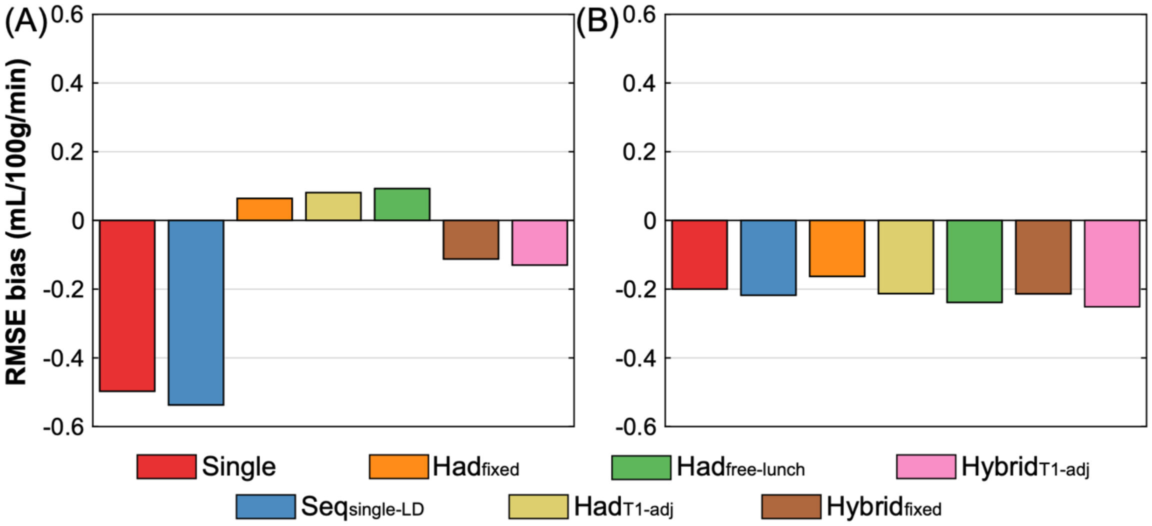

Figure 5.13: Bias in the ground truth Monte Carlo simulation CBF RMSEs when using 2 different noise models: (A) ground truth values fitted using 1 noise magnitude for all of the data and (B) ground truth values fitted using 3 noise magnitudes (1 each for: non-Hadamard protocols, Hadamard protocols, and the hybrid protocols). The noise magnitudes were estimated from the ground truth data in both cases (see methods). When 1 noise magnitude is used in the fitting, there is a large variation in the bias across protocols, but when 3 noise magnitude are used the RMSEs are much more similarly underestimated for all of the protocols by -0.21 ± 0.03 mL/100g/min.

5.5.2.5 Test-retest errors

The voxelwise and subjectwise CBF test-retest RMSEs for each protocol are shown in Figure 5.14, demonstrating the reproducibility of each protocol. A further 2.4% of the GM voxels were excluded from this comparison because one or more of the 2.5 minute scans had CBF or ATT posterior SDs > 500 mL/100g/min or > 50 s, suggesting very poor fits.

The voxelwise results from this reproducibility metric reflect the posterior SDs and errors relative to the ground truth estimates described above. However, in this case, $\text{Had}_{\text{fixed}}$ had the least reproducible estimates, with the single-TI protocol also having higher RMSEs than the other protocols. This again demonstrates that more robust CBF estimates can be obtained with multi-TI protocols than a single-TI protocol, in this case using a metric which is not reliant on uncertainty estimates from the fitting algorithm nor any estimated ground truth. $\text{Seq}_{\text{single-LD}}$ had the lowest test-retest RMSEs of all of the protocols (7.00 ±

0.24 mL/100g/min), marginally lower than Hybrid$_{T1\text{-adj}}$ (7.13 ± 0.82 mL/100g/min). All of the voxelwise test-retest RMSEs differences were significant.

For the subjectwise comparisons, Had$_{\text{fixed}}$ had significantly higher RMSEs than all protocols except the single-TI protocol. No other differences were significant, but the trends reflect the results in the voxelwise comparisons. The test-retest RMSEs were quite varied across subjects, which may in part be due to changing physiology. For example, the highlighted subject in Figure 5.14B for the Hybrid$_{T1\text{-adj}}$ scan had an average increase in CBF of 10 mL/100g/min between the two halves of the scan, perhaps reflecting a change in the alertness of the subject. The effect of this large variability can be seen in the large SDs of the voxelwise comparisons.

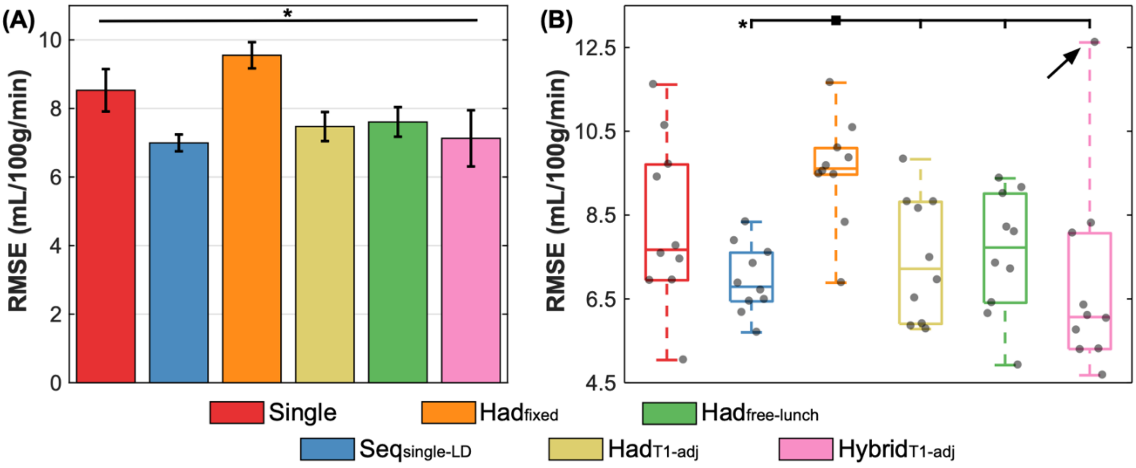

Figure 5.14: The in vivo test-retest RMSEs between the two 2.5 minute sub-datasets for each protocol across all voxels (A) and across subjects (B). In (A), the mean and standard error (see methods) of the RMSEs across voxels are shown. In (B), the boxplots show the quartiles and range of the subject RMSEs. All protocols in (A) had significantly different RMSEs to each other except Seq$_{single\text{-}LD}$ and Hybrid$_{T1\text{-}adj}$. In (B), the protocols with significantly different subject RMSEs are highlighted. The arrow in (B) highlights a single subject that had particularly large test-retest errors for the Hybrid$_{T1\text{-}adj}$ protocol.

5.5.3 Arterial Transit Time

Although the protocols were not optimised for ATT accuracy, the results of the ATT comparisons are summarised here. The voxelwise measures of uncertainty, accuracy and reproducibility are shown in Figure 5.15 and demonstrate that the time-encoded and hybrid protocols all have more confident, accurate, and reproducible ATT estimates than Seq$_{\text{single-LD}}$. The Had$_{T1\text{-adj}}$ protocol had the most confident ATT estimates overall, but compared to Hybrid$_{T1\text{-adj}}$ its accuracy relative to the ground truth was not significantly different (0.096 ± 0.006 s versus 0.097 ± 0.004 s) and was only marginally more repeatable (0.163 ± 0.012 s versus 0.165 ± 0.010 s). The relative performance of each protocol is also reflected in the subjectwise comparisons (not shown).

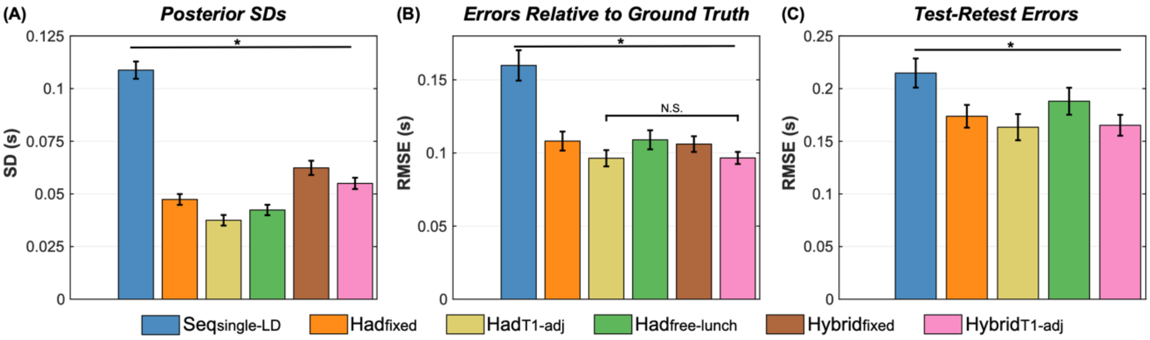

Figure 5.15: The in vivo voxelwise ATT measures of (A) uncertainty (posterior distribution SDs), (B) accuracy (RMSEs relative to the ground truth estimates) and (C) repeatability (test-retest RMSEs). The mean and standard error (see methods) of the RMSEs across voxels are shown. All protocols had significantly different measures, unless highlighted as not-significant (NS).

5.6 Discussion

In this study, a wide range of PCASL protocol designs were optimised for CBF accuracy and their CBF estimates compared using the theoretical Cramér-Rao lower bounds, Monte Carlo simulations and in vivo experiments, which were in good agreement. For the Monte Carlo simulations and in vivo data, the CBF estimates were compared: (1) using the standard deviation of the posterior distributions from the fitting algorithm as a measure of uncertainty or confidence; (2) using the RMSEs of the estimates relative to the ground truth estimates as a measure of accuracy, which includes both random variability and systematic biases; and (3) using the RMSEs of the test-retest estimates as a measure of reproducibility.

The single-TI CBF estimates were found to be less confident, less accurate, and less reproducible than all of the scan-time matched multi-TI protocols tested, except for the $\text{Had}_{\text{fixed}}$ protocol, which had less confident and reproducible estimates. These results were demonstrated using Monte Carlo simulations and confirmed in vivo.

The multi-TI methods have lower SNR at each TI than the single-TI protocol due to fewer averages, but the noise is in effect averaged across TIs during the fitting process. This results in similar noise averaging to the single-TI data, but with TIs that more effectively sample the signal curve for the whole ATT distribution on average. Multiple-TI methods also have the additional advantage of providing estimates of ATT, a potentially clinically useful physiological parameter in its own right,[48] and a useful indicator of delayed blood arrival.

Of the existing protocols tested from the literature, $\text{Had}_{TI\text{-adj}}$ had the highest CBF confidence and accuracy, while $\text{Seq}_{\text{single-LD}}$ was the most reproducible. Finally, it was demonstrated that the hybrid methods were more confident, accurate, and reproducible than their standard Hadamard-encoded analogues and that $\text{Hybrid}_{TI\text{-adj}}$ produced the most confident and accurate CBF estimates of all of the tested protocols and had a similar reproducibility to $\text{Seq}_{\text{single-LD}}$.

5.6.1 Optimal Protocol Designs

Except for $\text{Had}_{\text{fixed}}$ and $\text{Hybrid}_{\text{fixed}}$, the longest PLD of the multi-TI protocols was roughly 2.3 s. A PLD of 2.3 s will sample the peak of the signal for the longest ATT in the prior distribution (including the 0.3 s taper used to minimise edge effects). $\text{Had}_{\text{fixed}}$ had a longest effective PLD of 3.4 s, where very little signal from its 550 ms LD will remain, demonstrating the difficulty of balancing the choice of LDs and PLDs with this protocol design. The Hybrid$_{\text{fixed}}$ protocol also had 5 effective PLDs which were longer than 2.3 s, but this protocol benefits from longer LDs.

The single-TI PLD of 2 s was adequate for the expected ATT range of 0.5 - 2 s. If the PLD was shorter, CBF would be underestimated in voxels with ATT > PLD due to incomplete label arrival. With a longer PLD, signal would be lost unnecessarily across the whole ATT range. If an ATT pre-scan were performed with the single-TI acquisition, as in Dai et al.,[14] then a shorter PLD could be used to improve the average SNR across the ATT range. However, this approach requires sacrificing some scan time for the pre-scan, reducing the SNR of the single-TI data. It is possible, though, that due to the lower-resolution of the ATT pre-scan (which assumes that ATTs vary smoothly across the brain), there could be an overall improvement in CBF accuracy compared to the multi-TI methods. Nevertheless, this approach will not be appropriate for situations where it is preferable to estimate CBF and ATT simultaneously, such as in situations where a stimulus or drug has the potential to modify both parameters during the course of an experiment.

The variable LDs of the Had$_{T1\text{-adj}}$ and Had$_{\text{free-lunch}}$ protocols are beneficial to CBF estimation at long ATTs (Figure 5.10) because of the longer LDs and more appropriate PLDs of the first few blocks, compared to Had$_{\text{fixed}}$. The sensitivity of the CRLBs to the protocol design is particularly clear when comparing Had$_{T1\text{-adj}}$ and Had$_{\text{free-lunch}}$ because they have similar timings except for the first LD. The long first LD of Had$_{\text{free-lunch}}$ leads to a reduction in CBF uncertainty for ATTs between 1.7 s - 2 s, compared to Had$_{T1\text{-adj}}$. However, Had$_{\text{free-lunch}}$ has one less average than Had$_{T1\text{-adj}}$, which leads to a minor increase in CBF uncertainty at short ATTs. In simulations, the average CBF errors across the ATT distribution were slightly lower for Had$_{\text{free-lunch}}$ than for Had$_{T1\text{-adj}}$. However, when averaged across all voxels in vivo, $\mathrm{Had}_{T1\text{-adj}}$ had a slightly lower overall error, probably because of the shorter than expected ATTs encountered in this group of subjects (95% of voxels had ATTs less than 1.51 s).

It should be noted that the $\mathrm{Had}_{\text{free-lunch}}$ protocol had a different PLD for the first block (LD 1.8 s, PLD 2.3 s) than the single-TI protocol (LD 1.8 s, PLD 2 s), so the CBF estimates from the single-TI protocol are not directly comparable with the estimates calculated from just the first decoded block of the $\mathrm{Had}_{\text{free-lunch}}$ protocol, as in Teeuwisse et al.[11] This longer PLD was necessary to reduce large uncertainty in the ATT and CBF estimates near the upper edge of the ATT prior distribution which occurs for the multi-TI protocols.

For both $\mathrm{Had}_{\text{fixed}}$ and $\mathrm{Had}_{\text{free-lunch}}$, a 4 × 3 Hadamard-encoding matrix was predicted to have improved average accuracy compared to an 8 × 7 encoding, however, the next best 8 × 7 encoding was used in each case in the comparison. This was done because this larger encoding size is more widely used in the literature and it was not known how using only 3 TIs under in vivo conditions would affect the CBF accuracy. Example

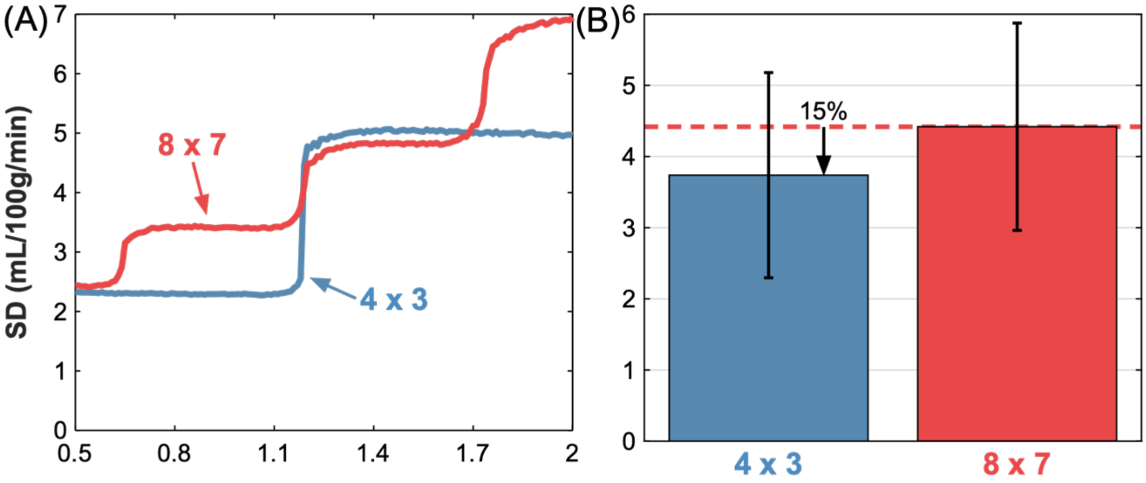

Figure 5.16: The simulated CBF uncertainty (Monte Carlo simulation posterior SDs) for the optimised $\mathrm{Had}_{\mathit{fixed}}$ *protocol when using a* 4 × 3 *or* 8 × 7 *Hadamard matrix. (A) the median uncertainty across ATTs, (B) the mean and SD uncertainty for each encoding size across the whole ATT range. Due to the* 4 × 3 *Hadamard protocol only having 3 PLDs, there is a large jump in uncertainty in the middle of the ATT distribution where the protocol becomes more sensitive to longer ATTs. However, overall, the* 4 × 3 *protocol has a lower average uncertainty than the* 8 × 7 *protocol.*

Monte Carlo results are shown for the $\text{Had}_{\text{fixed}}$ protocol in Figure 5.16, which demonstrate the large improvement in CBF accuracy at short and long ATTs. The benefit of the smaller encoding matrix likely comes from the long LDs used (1125 ms compared to 550 ms). Future work could further investigate the advantages and disadvantages of standard 4×3 Hadamard encodings in vivo. However, due to the improved performance of the hybrid protocol designs, it is likely that future protocol design improvements will come from moving away from the standard Hadamard designs which use a single final PLD.

The Hadamard-encoded protocols are relatively simple to optimise because the design restrictions reduce the dimensionality of the parameter space to be explored. All of the possible timing combinations, therefore, can be exhaustively searched rather than iteratively updated, as is the case for the sequential protocol designs. As a result of the larger parameter space of the sequential protocols, repeated random timing initialisations were required to avoid the many local minima. This was not found to be a problem when only optimising the PLDs for the sequential protocols in Chapter 4, but it becomes necessary for the $\text{Seq}_{\text{multi-LD}}$ protocol.

Hadamard-encodings schemes were used for the time-encoded protocols because these provide the most efficient encodings. However, they can only be of size $(\text{rows} \times \text{columns})\ 2k \times (2k - 1)$, for k = 1, 2, 4, 6, 8, 10, …. Less efficient encodings may provide more flexibility in the LDs and PLDs and could be explored with the same optimisation framework used in this work.

Each of the multi-TI protocols were optimised for CBF accuracy, with ATT treated as a confounding parameter. If both parameters are of interest, it is possible to optimise the protocols so that any trade-off in CBF and ATT accuracy can be achieved by minimising the weighted sum of the CRLB variances or by minimising the determinant of the CRLB matrix. Alternatively, one could optimise a protocol specifically for ATT estimation. This could either be done by ignoring CBF in the optimisation and fitting for just ATT, or by optimising for ATT, as done in this work with CBF, and treating CBF as a confounding parameter. Since the time-encoded methods give improved ATT measurements compared to $\text{Seq}_{\text{single-LD}}$ (due to their denser sampling of inflowing blood and improved averaging) it is likely these protocols have the greatest potential for any trade-off where ATT accuracy is important.

5.6.2 Simulation and In Vivo Comparisons

The results show that the $\text{Had}_{\text{fixed}}$ protocol produced significantly less confident, less accurate, and less reproducible CBF estimates than the other multi-TI protocols. However, its ATT estimates were more confident, more accurate, and more reproducible than those of $\text{Seq}_{\text{single-LD}}$, due to its denser sampling of the signal inflow. This is in agreement with a previous study that compared the $\text{Had}_{\text{fixed}}$ and $\text{Seq}_{\text{single-LD}}$ protocols, though with different timings, and found that $\text{Seq}_{\text{single-LD}}$ had more reproducible CBF estimates while $\text{Had}_{\text{fixed}}$ had more reproducible ATT estimates.[4]

The single-TI CBF estimates were also less confident, less accurate, and less reproducible than the other multi-TI protocols, suggesting that simply averaging at one conservative TI for a given ATT distribution does not produce more accurate CBF estimates than an optimised multi-TI protocol. This is despite also estimating ATT with the multi-TI protocols, which unavoidably increases uncertainty in the CBF estimates due to the overlapping sensitivities of each parameter.

The multi-TI protocols produced more confident and reproducible CBF estimates than the single-TI protocol because they sample near to the peak ASL signal more often across the whole ATT distribution, increasing the average SNR efficiency. Nonetheless, at ATTs > 1.6 s, the single-TI protocol had more confident CBF estimates than $\text{Hybrid}_{TI\text{-adj}}$, the best performing multi-TI protocol (Figure 5.10). This is largely unavoidable because the 34 averages of the single PLD at 2 s will always estimate CBF more accurately when the true ATT is close to 2 s than is possible with multi-TI protocols (optimised for a range of ATTs) with matched scan time. The advantage of the multi-TI protocols is that they can have lower average CBF uncertainty across a given ATT range.

The $\mathrm{Seq}_{\text{single-LD}}$ protocol had a similar CBF test-retest reproducibility to $\mathrm{Hybrid}_{T1\text{-adj}}$ despite having larger posterior SDs than all of the other multi-TI protocols except $\mathrm{Had}_{\text{fixed}}$. This may be in part due to the highly variable test-retest RMSEs seen across subjects, possibly due to the lower SNR in the CBF estimates of the 2.5 minute data sets as well as fluctuating subject physiology. For instance, if the highlighted outlier subject in Figure 5.14(B) was removed from this comparison, $\mathrm{Hybrid}_{T1\text{-adj}}$ had significantly lower RMSEs (6.33 ± 0.41 mL/100g/min), and so better reproducibility, than $\mathrm{Seq}_{\text{single-LD}}$ (7.01 ± 0.27 mL/100g/min).

5.6.2.1 Model Fitting

The single-TI data were fit with the same model as the multi-TI data in order to minimise systematic differences between the CBF estimates. The assumed ATT of 1.3 s represents the best guess for the ATT prior distribution used because it minimised the systematic errors.

Fitting the single-TI data with the approach recommended in the ASL white paper[7] ($T_{1t} = T_{1b}$, $f/\lambda = 0$ and ATT $\leq$ PLD) led to the mean GM CBF being significantly underestimated by (mean ± SD) 8.0 ± 2.6 mL/100g/min compared to the multi-TI data, on the subject level. None of the other protocols had significantly different mean GM CBF estimates from each other on the subject level, including the single-TI data fit with the standard CASL model.

5.6.2.2 Choice of ATT Prior

A uniform ATT prior distribution of 0.5 - 2 s was chosen based on the ATT estimates in Chapter 4, which used a similar labelling plane placement. However, the ATTs encountered in vivo here were generally shorter than expected, with 95% of the ground truth ATTs ≤ 1.51 s. This may be partly explained by the use of a visual stimulus to maintain alertness in the volunteers during the scan session: there is evidence that a visual stimulus can lead to a reduction of ATTs[49] (and an increase in CBF)[50] in the visual cortex, which typically has longer ATTs than other brain regions.[47] Furthermore, 5.2% of the ground truth voxels had ATTs < 0.5 s, which was outside of the optimised ATT range. Nevertheless, if these voxels were excluded from the analysis, the overall results are similar and so the conclusions remain unchanged (results not shown).

Flow crushing gradients were not used, which have been shown to significantly lengthen measured ATTs across the brain,[47] though the spoiler gradients either side of the 3D GRASE refocussing pulses will have caused some amount of flow crushing.[13] Since resting perfusion ASL scans do not usually use a visual stimulus, and it is recommended not to use vascular crushing,[7] it is likely that an ATT prior distribution range of 0.5 - 1.8 s is sufficient for protocol optimisation for young healthy volunteers. However, a range of 0.5 - 2 s, as seen in Chapter 4, may be more appropriate if vascular crushing is used.

5.6.2.3 Motion sensitivity

Similarly, to other subtractive imaging techniques, motion during or between the label and control images can corrupt the perfusion estimates. For the single-TI and sequential protocols, corrupted images can be identified and excluded from the fitting process, affecting only those images which are corrupted. However, because the time-encoded and hybrid protocols use a signal encoding process, an image corrupted by motion will propagate artefacts to all of the TIs decoded using that image. Motion, therefore, affects the time-encoded and hybrid protocols to a greater extent than the single-TI and sequential protocols, with protocols using larger encoding matrices being more affected. This is not necessarily a problem with the compliant healthy volunteers scanned in this work, but it could make the time-encoded methods difficult to use in non-compliant patient groups, such as in acute stroke imaging.

Samson-Himmelstjerna et al.[25] suggested reordering the rows of the Hadamard-encoding matrix so that low temporal resolution images can be decoded before the entire encoding matrix has been acquired, with no time penalty. This could then be used by the scan operator to identify and reacquire certain volumes while the patient is still in the scanner, or to simply be able to salvage certain decoded timepoints from an otherwise corrupted set of encoded images.

There are several methods which could be employed to reduce the motion sensitivity of ASL methods in general, which are briefly summarised here. In this work, retrospective image-based motion correction was used to re-align the label and control images to correct for inter-volume motion. Other methods which can correct for inter-volume motion include fast navigator acquisitions which can be inserted into the sequence at regular intervals to either perform prospective motion correction by updating the 3D FOV,[51] by reacquiring corrupted data,[52] or retrospective correction of *k*-space data.[53] Self-navigating readouts, such as GRASE PROPELLER,[54] can also be used. Intra-volume motion is more difficult to correct, so real-time image quality monitoring and reacquisition techniques are likely easier to implement. Sensitivity to intra-volume motion can be reduced by using fast readouts, achieved either by *k*-space acquisition segmentation or by *k*-space undersampling, using parallel imaging,[55,56] compressed sensing,[57,58] or a combination of these for reconstruction of undersampled data.

5.6.2.4 Single-Shot 3D GRASE

A single-shot 3D GRASE readout was used for the in vivo ASL comparison. In order to reduce the through-plane blurring, a rectangular FOV was used to reduce the echo-train length and a refocusing flip-angle of 120° was used[25,26] to preserve more signal for later echoes. Although there remained a considerable amount of blurring in the ASL images, this should not affect the conclusions of this study since each protocol used an identical readout and so is directly comparable. In addition, the number of PLDs were restricted to ≤ 15 to ensure that more than one average for each protocol could be achieved. Therefore, the optimised protocols may also be used with segmented 3D readouts in a similar scan time by dividing the number of averages used here by the number of desired segments. The relative performance of each protocol should remain constant in this regime because the predicted CBF and ATT variance simply scales with the number of averages.

5.6.2.5 Variable Noise Across ATTs

There was a gradual increase in the in vivo posterior distribution SDs at shorter ATTs compared to the Monte Carlo simulations, which assumed equal noise across all ATTs. A possible explanation is that shorter ATTs were generally located closer to the middle of the brain and so further from the head-coil receive elements than longer ATTs. This could result in an SNR level that was negatively related to the ATT.

To test this hypothesis, the voxelwise temporal noise, σ, was calculated from the calibrated in vivo single-TI control images (which should have very little perfusion weighting) by taking the SD across averages. A linear model, $\sigma = a \cdot ATT + b$, was fitted to these data using the ground-truth ATT estimates and a linear-least-squares fit. This resulted in the fitted parameters $a = -4.28 \times 10^{-4}$ and $b = 20.29 \times 10^{-4}$ $(R^2 = 0.58)$, showing that there was indeed a strong correlation between noise magnitude and ATT.

This model was then rescaled so that the predicted noise SD at ATT = 1.25 s (the middle of the ATT distribution) was equal to the noise SD used in the original simulations.

This variable noise level was used in further Monte Carlo simulations (similar to those described in the methods), the results of which are shown in Figure 5.17. The simulation CBF posterior SD trends across ATTs now much more closely match the in vivo trends, suggesting that the relationship between coil sensitivity and ATTs does largely

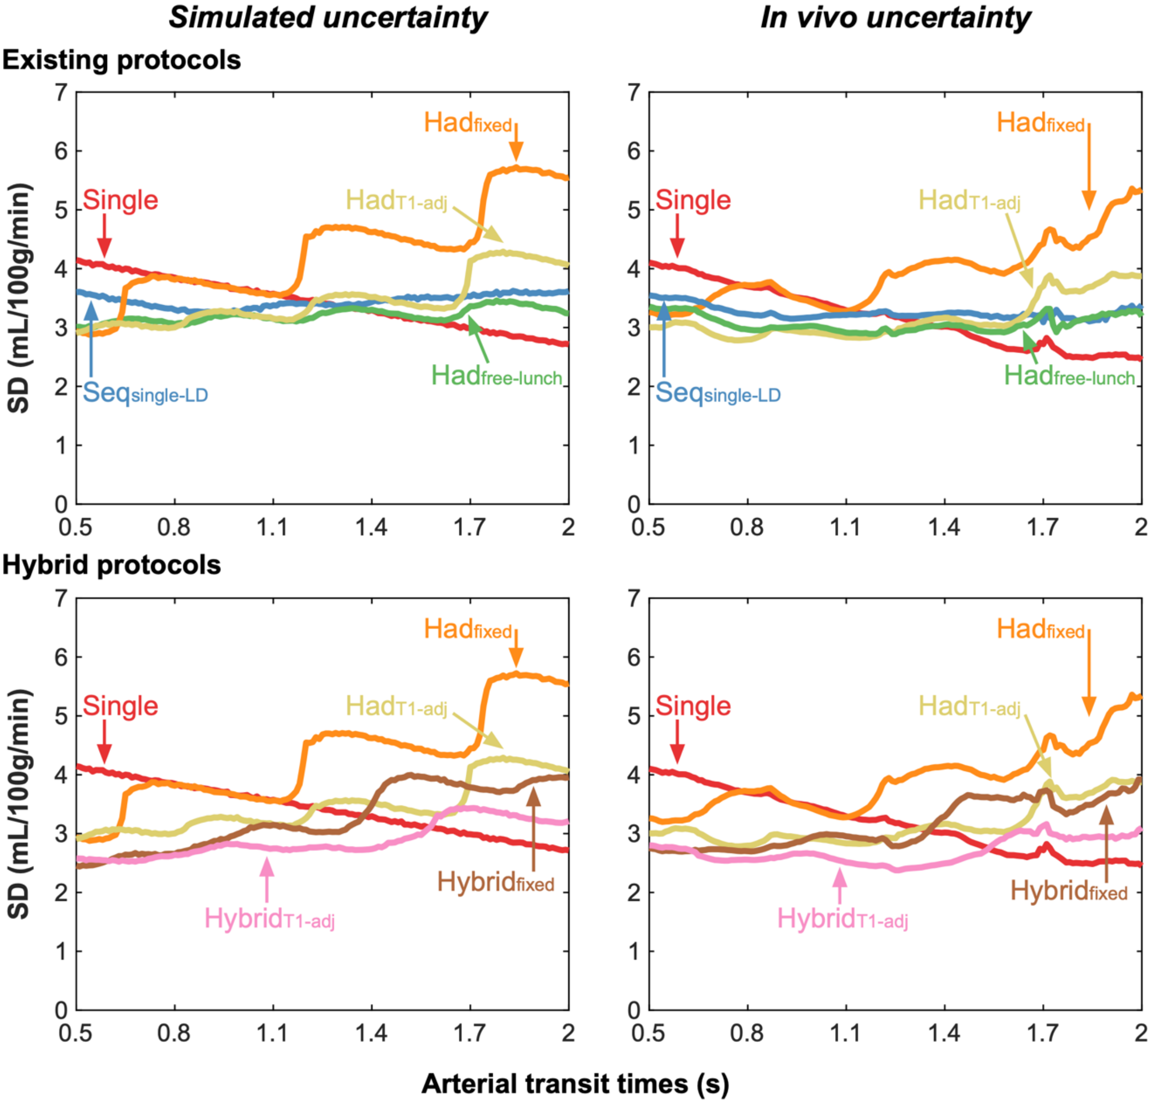

Figure 5.17: The simulated (Monte Carlo simulation posterior SDs) and in vivo (posterior SDs) median CBF uncertainty measures for the existing protocols (top) and the proposed hybrid protocols (bottom). These MC simulations use the predicted variable noise levels across ATTs calculated from the in vivo data and scaled to the noise SD used in the original MC simulations. The MC posterior SD trends across ATTs now better match the in vivo trends, suggesting that the relationship between the ATTs and the variable noise level of the head-coil largely explains the varying posterior SDs we see in vivo.

explain the average trends seen in vivo. However, this ATT dependent noise model is unlikely to be useful for protocol optimisation because it will vary across subjects, subject placement, and head coil design.

5.6.3 Extended Investigations

5.6.3.1 Hadamard and Hybrid Protocols with Variable Label Durations

The effect of individually optimising each of the LDs in the Hadamard-encoded and hybrid protocols was also investigated, referred to as $\text{Had}_{\text{variable}}$ and $\text{Hybrid}_{\text{variable}}$, respectively. Allowing complete freedom for the LDs, as opposed to the fixed-duration or T_1-adjusted design restrictions, provides the greatest flexibility to the standard Hadamard design but also greatly increases the size of the protocol timing parameter space and the chance of finding a local minimum.

Early testing of optimising $\text{Had}_{\text{variable}}$ with randomised initial timings suggested that the best protocols had LDs which were longest at the start of the PCASL preparation and successively decreased in duration. Therefore, the i^{th} LD was restricted by $\text{LD}_{i\text{-}1} \leq \text{LD}_i \leq \text{LD}_{i+1}$, helping to reduce the number of candidate LDs at each update iteration. The same minimum and maximum LDs of 0.1 s and 1.8 s were used, as previously. Both protocols were optimised in a similar manner to those described previously, but with 50 random timing initialisations and with each time-encoded block being updated in a randomly permuted order. For the $\text{Hybrid}_{\text{variable}}$ protocol, the encoded LDs were allowed to vary for each separate PLD, but the Hadamard-encoding size was kept constant within each protocol.

Protocol	*Label durations (ms)*	*Post-label delays (ms)*	*Averages*	*Total acquisitions*	*Scan duration*
Had$_{\text{variable}}$	1725, 750, 650, 375, 150, 150, 125	100	8	64	4:58 min
Hybrid$_{\text{variable}}$	1800, 1050, 775, 1800, 1225, 550, 1800, 850, 750, 1800, 800, 800	100, 525, 575, 700	4	64	4:55 min

Table 5.3: Optimised protocol timings for Had$_{variable}$ and Hybrid$_{variable}$. The Hybrid$_{variable}$ PLDs correspond in order to each of the 4 sets of encoded LDs.

The optimal timings for Had$_{\text{Variable}}$ and Hybrid$_{\text{variable}}$ are given in Table 5.3. The Had$_{\text{variable}}$ LDs and PLD are quite similar to the Had$_{\text{free-lunch}}$ protocol (with T_1-adjusted final LDs), but the 2nd - 4th LDs are longer while the 5th - 7th LDs are shorter than the Had$_{\text{free-lunch}}$ LDs.

For the Hybrid$_{\text{variable}}$ protocol, 4 different PLDs, with associated different encoded LDs, were chosen, effectively giving 4 different 4 × 3 Hadamard-encoded sub-protocols, each with 4 averages, which are combined to give optimal measurements for CBF estimation.

Figure 5.18 shows the CBF estimation performance of these 2 variable LD protocols relative to a subset of previously compared protocols. Across the ATT distribution, although Had$_{\text{variable}}$ and Hybrid$_{\text{variable}}$ have lower average uncertainty than the Hadamard and hybrid protocols, respectively, there is only a very small improvement. Compared to Had$_{\text{free-lunch}}$, Had$_{\text{variable}}$ has lower uncertainty at shorter ATTs, but at the cost of higher uncertainty at longer ATTs. The Hybrid$_{\text{variable}}$ protocol reduces uncertainty at short and long ATTs compared to Hybrid$_{T_1\text{-adj}}$ but has higher uncertainty at ATTs in the centre of the distribution.

Due to the increased complexity and computational cost of optimising the variable LD protocols, and the very small gains in CBF accuracy, it seems that the more rigid Hybrid$_{T1\text{-adj}}$ design has a better trade-off between CBF accuracy and protocol optimisation complexity.

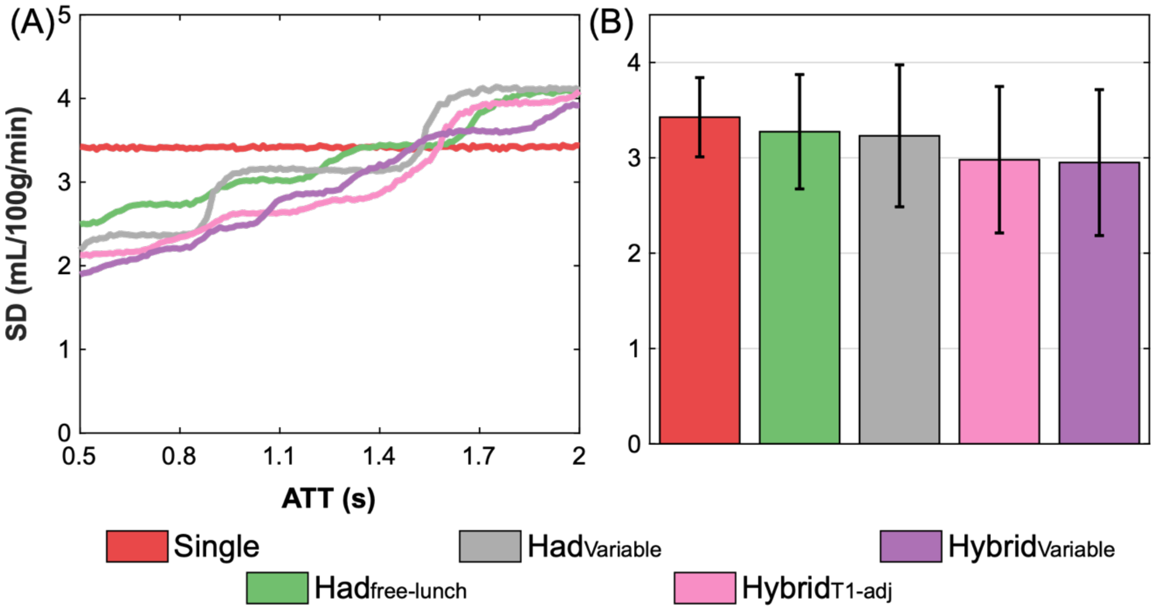

Figure 5.18: The simulated CBF uncertainty (Monte Carlo simulation posterior SDs) for Had$_{variable}$, Hybrid$_{variable}$, and a selection of previously compared protocols. (A) the median uncertainty for each protocol across ATTs, (B) the mean and SD uncertainty for each protocol across the whole ATT range.

5.6.3.2 Long Label Durations (> 1.8 s)

The use of long label durations was also investigated as a method for improving CBF accuracy further by increasing the potential for greater signal accumulation.

A subset of the optimisations described previously were re-run with a maximum LD of 5 s. In order to fairly compare the single-TI protocol with the optimised multi-TI protocols, the LD of the single-TI protocol was optimised by maximising the tSNR efficiency, similar to Zun et al.,[5] but with a fixed scan duration of 5 minutes and a PLD of 2 s (see Figure 5.19). The optimised protocol timings are given in Table 5.4.

The Had$_{TI\text{-adj}}$ protocol had the same timings as when the LD was constrained to be ≤ 1.8 s due to the trade-off between the LDs and PLDs, and so the timings are not repeated here. All of the other protocols that were optimised used much longer LDs, suggesting that longer LDs with fewer averages was beneficial for CBF accuracy with these protocols.

Figure 5.20 shows the simulated CBF uncertainty for each of the long LD optimised protocols. All of the protocols, except Had$_{TI\text{-adj}}$, improved CBF accuracy compared to their short LD versions (by (mean ± SD) 6 ± 1% on average), demonstrating that the advantage of increasing the LDs outweighs the cost of reducing the number of averages.

It can be seen that in the long LD regime, the Seq$_{\text{single-LD}}$ and Seq$_{\text{multi-LD}}$ protocols have improved CBF accuracy compared to Had$_{TI\text{-adj}}$. This suggests that the increase in signal accumulation from the longer LDs of the two sequential protocols is more beneficial to CBF estimation than the ideal signal averaging of the Hadamard-decoding.

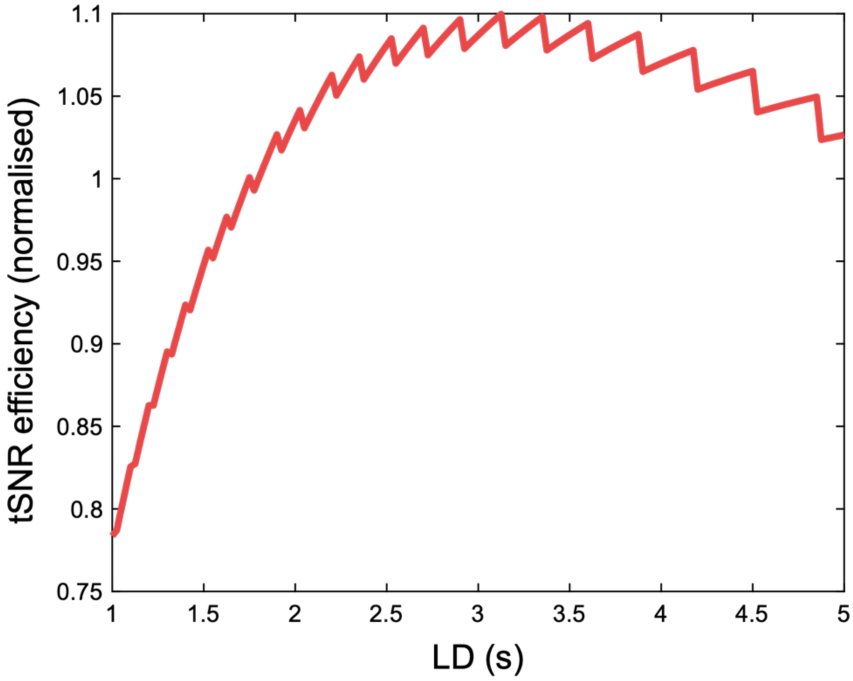

Figure 5.19: Optimising the LD for the single-TI protocol. The tSNR efficiency is calculated as $tSNR_{efficiency} = signal \cdot \sqrt{averages}$*, where the signal is the mean of the theoretical signal across the ATT distribution and the number of averages is rounded down to the nearest integer. For white Gaussian noise, the noise variance is expected to scale with the reciprocal of the number of averages. The tSNR efficiency has been normalised relative to the tSNR efficiency at LD = 1.8 s. The optimal LD of 3.125 s increases the tSNR efficiency by 10% compared to LD = 1.8 s. The PLD was fixed at 2 s.*

Protocol	*Label durations (ms)*	*Post-label delays (ms)*	*Averages*	*Total acquisitions*	*Scan duration*
Single-TI	3125	2000	26	52	4:58 min
Seq$_{\text{single-LD}}$	2800	75, 975, 1600, 2100, 2300, 2300	5	60	5:00 min
Seq$_{\text{multi-LD}}$	2025, 1400, 3425, 4225, 3850	200, 925, 1975, 1475, 2300	6	60	5:00 min
Hybrid$_{T1\text{-adj}}$	3550, 1050, 625	75, 275, 550, 625	3	48	5:00 min
Hybrid$_{\text{variable}}$	3975, 1125, 300, 3325, 1425, 550, 3300, 1225, 500, 3575, 725, 625	200, 325, 575, 675	3	48	5:00 min

Table 5.4: Optimised protocol timings for the listed protocols with a maximum LD of 5 s. The Hybrid$_{variable}$ PLDs correspond in order to each of the 4 sets of encoded LDs.

In the short LD regime, Seq$_{\text{single-LD}}$ and Seq$_{\text{multi-LD}}$ performed similarly to each other in terms of CBF and ATT accuracy (not shown), but with the increased flexibility of the long LD regime, the improvement in CBF accuracy of Seq$_{\text{multi-LD}}$ is more noticeable (4% improvement relative to Seq$_{\text{single-LD}}$). However, the clearer advantage of the Seq$_{\text{multi-LD}}$ design is its improvement in ATT accuracy compared to Seq$_{\text{single-LD}}$ (20% reduction in ATT simulation posterior SDs), implying that by varying both the LDs and PLDs in a sequential protocol design, ATT estimation accuracy can be greatly improved while not diminishing CBF accuracy. However, the time-encoded protocols still maintained superior ATT estimation (34 - 48% lower simulation ATT uncertainty than Seq$_{\text{multi-LD}}$).

The Hybrid$_{T1\text{-adj}}$ and Hybrid$_{\text{variable}}$ protocols perform very similarly to each other, indicating again that the increased flexibility of the Hybrid$_{\text{variable}}$ protocol does not lead to an obvious improvement in CBF accuracy. This may, of course, be due to the large parameter space of the Hybrid$_{\text{variable}}$ optimisation and the globally optimal protocol timings have just not been found. Both hybrid protocols had the lowest overall CBF uncertainty, again demonstrating the superior design characteristics of the hybrid protocols: balancing ASL signal accumulation (long LDs) with efficient noise averaging (Hadamard decoding) and ideal dynamic signal sampling (PLDs).

Compared to the single-TI protocol, the Seq$_{\text{single-LD}}$, Seq$_{\text{multi-LD}}$, and Had$_{T1\text{-adj}}$ protocols improved CBF accuracy at short ATTs, but have worse accuracy at long ATTs, similar to the short LD case. However, the Hybrid$_{T1\text{-adj}}$ protocol only had worse CBF accuracy than the single-TI protocol at ATTs longer than 1.75 s, but with much improved CBF accuracy at all other ATTs.

In vivo, longer LDs may also have the advantage of averaging the generation of the ASL tracer over multiple heartbeats, reducing sensitivity to pulsatility.[6] However, the impact of increased sensitivity to tissue T_1, due to the greater amount of time the labelled blood has to accumulate in the tissue, would need to be evaluated. Even with shorter LDs, though, if accurate absolute CBF estimates are desired, it is likely that tissue T_1 (and blood T_1) should be estimated where possible, reducing the concern of tissue T_1 dependence.[59]

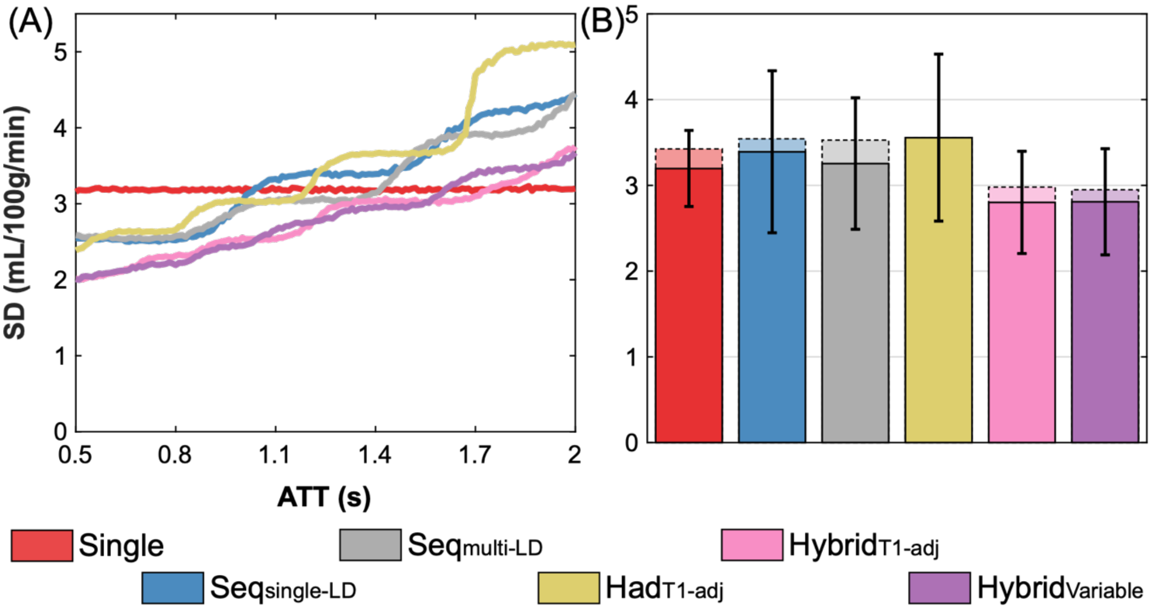

Figure 5.20: The simulated CBF uncertainty (Monte Carlo simulation posterior SDs) for the protocols optimised under the long LD condition. (A) the median uncertainty for each protocol across ATTs, (B) the mean and SD uncertainty for each protocol across the whole ATT range. In (B), the mean CBF uncertainty for the corresponding short LD protocols are shown in the background to demonstrate the improvement gained from using long LDs.

This is because tissue and blood T_1 relaxation varies across subjects and with disease, as discussed in Chapter 3.[17–19,39,60–64] If not accounted for, this variation in T_1 can contribute to inter-subject CBF measurement variability or even create false group CBF differences or obscure them. Possible deviations from the assumed model parameter values should, therefore, be carefully considered in each experiment.

5.7 Conclusions

It has been demonstrated that appropriately optimised multi-TI protocols can provide more confident, more accurate, and more reproducible CBF estimates than a single-TI protocol.

When using a maximum labelling duration of 1.8 s, a T_1-adjusted time-encoded protocol results in the most accurate CBF estimates of the conventional protocols from the literature. However, when long labelling durations are used, simulations suggest that a sequential protocol can outperform the conventional time-encoded protocols. Nevertheless, time-encoded protocols were found to provide more accurate ATT estimates in both cases.

The Hybrid protocol designs are able to find a more optimal balance between the LDs, PLDs and Hadamard-encoding size, compared to the existing protocols, giving the most robust CBF estimates of all the protocols.

The large increase in flexibility of the $\text{Hybrid}_{\text{variable}}$ protocol increased the computational burden of the optimisation while not providing much improvement to CBF accuracy. It seems, therefore, preferable to use the more constrained T_1-adjusted technique for choosing the LDs with the hybrid protocol design.

5.8 References

1 Wells JA, Lythgoe MF, Gadian DG, Ordidge RJ, Thomas DL. In vivo hadamard encoded continuous arterial spin labeling (H-CASL). *Magn Reson Med* 2010; **63**:

1111–1118.

2 Dai W, Shankaranarayanan A, Alsop DC. Volumetric measurement of perfusion and arterial transit delay using hadamard encoded continuous arterial spin labeling. *Magn Reson Med* 2013; **69**: 1014–1022.

3 Johnston ME, Lu K, Maldjian JA, Jung Y. Multi-TI arterial spin labeling MRI with variable TR and bolus duration for cerebral blood flow and arterial transit time mapping. *IEEE Trans Med Imaging* 2015; **34**: 1392–1402.

4 Guo J, Holdsworth SJ, Fan AP, Lebel MR, Zun Z, Shankaranarayanan A *et al.* Comparing accuracy and reproducibility of sequential and Hadamard-encoded multidelay pseudocontinuous arterial spin labeling for measuring cerebral blood flow and arterial transit time in healthy subjects: A simulation and in vivo study. *J Magn Reson Imaging* 2018; **47**: 1119–1132.

5 Zun Z, Lebel RM, Shankaranarayanan A, Zaharchuk G. What Is the Ideal Labeling Duration for Pseudocontinuous Arterial Spin Labeling? In: *Proceedings of the 22nd Annual Meeting of the ISMRM, Milan, Italy*. 2014, p 4566.

6 Lebel RM, Shankaranarayanan A, Smith EE, McCreary C, Frayne R, Dai W *et al.* The many advantages of arterial spin labeling with long label duration. In: *Proceedings of the 23rd Annual Meeting of the ISMRM, Toronto, Canada*. 2015, p 2333.

7 Alsop DC, Detre JA, Golay X, Günther M, Hendrikse J, Hernandez-Garcia L *et al.* Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn Reson Med* 2015; **73**: 102–116.

8 Borogovac A, Habeck C, Small SA, Asllani I. Mapping Brain Function Using a 30-Day Interval between Baseline and Activation: A Novel Arterial Spin Labeling fMRI Approach. *J Cereb Blood Flow Metab* 2010; **30**: 1721–1733.

9 Guo J, Wong EC. Increased SNR efficiency in velocity selective arterial spin labeling using multiple velocity selective saturation modules (mm-VSASL). *Magn Reson Med* 2014; **705**: n/a-n/a.

10 Günther M. Highly efficient accelerated acquisition of perfusion inflow series by Cycled Arterial Spin Labeling. In: *Proceedings of the 15th Annual Meeting of the ISMRM, Berlin, Germany*. 2007, p 380.

11 Teeuwisse WM, Schmid S, Ghariq E, Veer IM, van Osch MJP. Time-encoded pseudocontinuous arterial spin labeling: Basic properties and timing strategies for human applications. *Magn Reson Med* 2014; **72**: 1712–1722.

12 Feinberg DA, Oshio K. GRASE (gradient- and spin-echo) MR imaging: a new fast clinical imaging technique. *Radiology* 1991; **181**: 597–602.

13 Günther M, Oshio K, Feinberg DA. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn Reson Med* 2005; **54**: 491–498.

14 Dai W, Robson PM, Shankaranarayanan A, Alsop DC. Reduced resolution transit delay prescan for quantitative continuous arterial spin labeling perfusion imaging. *Magn Reson Med* 2012; **67**: 1252–1265.

15 Dai W, Robson PM, Shankaranarayanan A, Alsop DC. Benefits of Interleaved Continuous Labeling and Background Suppression. In: *Proceedings of the 18th Annual Meeting of the ISMRM, Stockholm, Sweden*. 2010, p 1748.

16 Lin C, Bernstein M, Huston J, Fain S. Measurements of T1 Relaxation times at 3.0T: Implications for clinical MRA. In: *Proceedings of the 9th Annual Meeting of the ISMRM, Glasgow, Scotland*. 2001, p 1391.

17 Clare S, Jezzard P. Rapid T(1) mapping using multislice echo planar imaging. *Magn*

Reson Med 2001; **45**: 630–634.

18 Stanisz GJ, Odrobina EE, Pun J, Escaravage M, Graham SJ, Bronskill MJ *et al.* T1, T2 relaxation and magnetization transfer in tissue at 3T. *Magn Reson Med* 2005; **54**: 507–512.

19 Liberman G, Louzoun Y, Ben Bashat D. T1 Mapping using variable flip angle SPGR data with flip angle correction. *J Magn Reson Imaging* 2014; **40**: 171–180.

20 Wansapura JP, Holland SK, Dunn RS, Ball WS. NMR relaxation times in the human brain at 3.0 tesla. *J Magn Reson Imaging* 1999; **9**: 531–8.

21 Ethofer T, Mader I, Seeger U, Helms G, Erb M, Grodd W *et al.* Comparison of longitudinal metabolite relaxation times in different regions of the human brain at 1.5 and 3 Tesla. *Magn Reson Med* 2003; **50**: 1296–1301.

22 Wright PJ, Mougin OE, Totman JJ, Peters a. M, Brookes MJ, Coxon R *et al.* Water proton T 1 measurements in brain tissue at 7, 3, and 1.5T using IR-EPI, IR-TSE, and MPRAGE: results and optimization. *Magn Reson Mater Physics, Biol Med* 2008; **21**: 121–130.

23 Ordidge RJ, Wylezinska M, Hugg JW, Butterworth E, Franconi F. Frequency offset corrected inversion (FOCI) pulses for use in localized spectroscopy. *Magn Reson Med* 1996; **36**: 562–566.

24 Buxton RB, Frank LR, Wong EC, Siewert B, Warach S, Edelman RR. A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magn Reson Med* 1998; **40**: 383–396.

25 von Samson-Himmelstjerna F, Madai VI, Sobesky J, Guenther M. Walsh-ordered hadamard time-encoded pseudocontinuous ASL (WH pCASL). *Magn Reson Med* 2016; **76**: 1814–1824.

26 He X, Wengler K, Schweitzer ME. Diffusion sensitivity of 3D-GRASE in arterial spin labeling perfusion. *Magn Reson Med* 2018; **80**: 736–747.

27 Ogg RJ, Kingsley PB, Taylor JS. WET, a T1- and B1-insensitive water-suppression method for in vivo localized 1H NMR spectroscopy. *J Magn Reson B* 1994; **104**: 1–10.

28 Golay X, Petersen ET, Hui F. Pulsed Star Labeling of Arterial Regions (PULSAR): A robust regional perfusion technique for high field imaging. *Magn Reson Med* 2005; **53**: 15–21.

29 Ye FQ, Frank JA, Weinberger DR, McLaughlin AC. Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). *Magn Reson Med* 2000; **44**: 92–100.

30 Payne GS, Leach MO. Implementation and evaluation of frequency offset corrected inversion (FOCI) pulses on a clinical MR system. *Magn Reson Med* 1997; **38**: 828–833.

31 Jenkinson M, Beckmann CF, Behrens TEJ, Woolrich MW, Smith SM. FSL. *Neuroimage* 2012; **62**: 782–790.

32 Jenkinson M. Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *Neuroimage* 2002; **17**: 825–841.

33 Jenkinson M, Smith S. A global optimisation method for robust affine registration of brain images. *Med Image Anal* 2001; **5**: 143–156.

34 Andersson JLR, Skare S, Ashburner J. How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *Neuroimage* 2003; **20**: 870–888.

35 Smith SM. Fast robust automated brain extraction. *Hum Brain Mapp* 2002; **17**: 143–155.

36 Zhang Y, Brady M, Smith S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging* 2001; **20**: 45–57.

37 Herscovitch P, Raichle ME. What is the Correct Value for the Brain-Blood Partition Coefficient for Water? *J Cereb Blood Flow Metab* 1985; **5**: 65–69.

38 Dai W, Garcia D, De Bazelaire C, Alsop DC. Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magn Reson Med* 2008; **60**: 1488–1497.

39 Lu H, Clingman C, Golay X, van Zijl PCM. Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn Reson Med* 2004; **52**: 679–682.

40 Chappell MA, Groves AR, Whitcher B, Woolrich MW. Variational Bayesian Inference for a Nonlinear Forward Model. *IEEE Trans Signal Process* 2009; **57**: 223–236.

41 Tukey JW. *Exploratory Data Analysis*. Addison-Wesley: London, 1977.

42 Chen JJ, Rosas HD, Salat DH. Age-associated reductions in cerebral blood flow are independent from regional atrophy. *Neuroimage* 2011; **55**: 468–478.

43 Parkes LM, Rashid W, Chard DT, Tofts PS. Normal cerebral perfusion measurements using arterial spin labeling: Reproducibility, stability, and age and gender effects. *Magn Reson Med* 2004; **51**: 736–743.

44 Vernooij MW, van der Lugt A, Ikram MA, Wielopolski PA, Vrooman HA, Hofman A *et al.* Total Cerebral Blood Flow and Total Brain Perfusion in the General Population: The Rotterdam Scan Study. *J Cereb Blood Flow Metab* 2008; **28**: 412–419.

45 Henriksen OM, Kruuse C, Olesen J, Jensen LT, Larsson HBW, Birk S *et al.* Sources of Variability of Resting Cerebral Blood Flow in Healthy Subjects: A Study Using 133 Xe SPECT Measurements. *J Cereb Blood Flow Metab* 2013; **33**: 787–792.

46 MacIntosh BJ, Filippini N, Chappell MA, Woolrich MW, Mackay CE, Jezzard P. Assessment of arterial arrival times derived from multiple inversion time pulsed arterial spin labeling MRI. *Magn Reson Med* 2010; **63**: 641–647.

47 Dai W, Fong T, Jones RN, Marcantonio E, Schmitt E, Inouye SK *et al.* Effects of arterial transit delay on cerebral blood flow quantification using arterial spin labeling in an elderly cohort. *J Magn Reson Imaging* 2016. doi:10.1002/jmri.25367.

48 MacIntosh BJ, Marquardt L, Schulz UG, Jezzard P, Rothwell PM. Hemodynamic alterations in vertebrobasilar large artery disease assessed by arterial spin-labeling MR imaging. *Am J Neuroradiol* 2012; **33**: 1939–1944.

49 Qiu M, Paul Maguire R, Arora J, Planeta-Wilson B, Weinzimmer D, Wang J *et al.* Arterial transit time effects in pulsed arterial spin labeling CBF mapping: Insight from a PET and MR study in normal human subjects. *Magn Reson Med* 2010; **63**: 374–384.

50 Kawasaki T, Kiyosawa M, Ishii K, Senda M. Regional cerebral blood flow response to visual stimulation measured quantitatively with PET. *Neuro-Ophthalmology* 1998; **20**: 79–89.

51 Frost R, Hess AT, Okell TW, Chappell MA, Tisdall MD, van der Kouwe AJW *et al.* Prospective motion correction and selective reacquisition using volumetric navigators for vessel-encoded arterial spin labeling dynamic angiography. *Magn Reson Med* 2016; **76**: 1420–1430.

52 Zun Z, Shankaranarayanan A, Zaharchuk G. Pseudocontinuous arterial spin labeling with prospective motion correction (PCASL-PROMO). *Magn Reson Med* 2014; **72**: 1049–1056.

53 Gallichan D, Marques JP, Gruetter R. Retrospective correction of involuntary

microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magn Reson Med* 2016; **75**: 1030–1039.

54 Tan H, Hoge WS, Hamilton CA, Günther M, Kraft RA. 3D GRASE PROPELLER: Improved image acquisition technique for arterial spin labeling perfusion imaging. *Magn Reson Med* 2011; **66**: 168–173.

55 Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med* 1999; **42**: 952–962.

56 Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J *et al.* Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002; **47**: 1202–1210.

57 Lustig M, Donoho DL, Santos JM, Pauly JM. Compressed Sensing MRI. *IEEE Signal Process Mag* 2008; **25**: 72–82.

58 Zhao L, Fielden SW, Feng X, Wintermark M, Mugler JP, Meyer CH. Rapid 3D dynamic arterial spin labeling with a sparse model-based image reconstruction. *Neuroimage* 2015; **121**: 205–216.

59 Debacker CS, Daoust A, Köhler S, Voiron J, Warnking JM, Barbier EL. Impact of tissue T1 on perfusion measurement with arterial spin labeling. *Magn Reson Med* 2017; **77**: 1656–1664.

60 Wu WC, Jain V, Li C, Giannetta M, Hurt H, Wehrli FW *et al.* In vivo venous blood T1 measurement using inversion recovery true-FISP in children and adults. *Magn Reson Med* 2010; **64**: 1140–1147.

61 Václavů L, van der Land V, Heijtel DFR, van Osch MJP, Cnossen MH, Majoie CBLM *et al.* In Vivo T1 of Blood Measurements in Children with Sickle Cell Disease Improve Cerebral Blood Flow Quantification from Arterial Spin-Labeling MRI. *Am J Neuroradiol* 2016; **37**: 1727–1732.

62 Vymazal J, Righini A, Brooks RA, Canesi M, Mariani C, Leonardi M *et al.* T1 and T2 in the Brain of Healthy Subjects, Patients with Parkinson Disease, and Patients with Multiple System Atrophy: Relation to Iron Content. *Radiology* 1999; **211**: 489–495.

63 Bastin ME, Sinha S, Whittle IR, Wardlaw JM. Measurements of water diffusion and T1 values in peritumoural oedematous brain. *Neuroreport* 2002; **13**: 1335–1340.

64 Naruse S, Horikawa Y, Tanaka C, Hirakawa K, Nishikawa H, Yoshizaki K. Significance of proton relaxation time measurement in brain edema, cerebral infarction and brain tumors. *Magn Reson Imaging* 1986; **4**: 293–304.

6
Time-Encoded Dynamic Angiography

6.1 Introduction

Dynamic angiograms[1,2] can be obtained to visualise the pass-through of the ASL tracer by acquiring multiple sequential (highly segmented) images after each ASL preparation, providing much richer haemodynamic information than a static angiogram. However, due to the large number of excitations, low flip-angles are necessary, as described in Chapter 2, which limits SNR. By encoding some, or all, of the desired temporal resolution into the PCASL pulse train using a time-encoded preparation,[3,4] higher flip-angles can be used during the readout, because fewer excitations are required than in a matched sequential protocol.

In Chapter 5, it was demonstrated that under certain regimes, such as long LDs, traditional time-encoded preparations can have worse CBF accuracy than a sequential protocol, because the benefit of improved noise averaging is outweighed by the reduction in signal magnitude due to the use of short LDs. In ASL angiography,[5,6] it is the pass-through of labelled blood in the arterial tree that is visualised, not the accumulation of signal in tissue, and so the short label durations used in time-encoded PCASL may not be so detrimental.

In this chapter, the time-encoded PCASL preparation is combined with a temporally resolved spoiled gradient echo (SPGR) readout using a radial trajectory (described in Chapter 2). Using simulations and preliminary in vivo data, it is demonstrated that the higher flip-angles available when using a time-encoded preparation can greatly increase the SNR of the dynamic angiograms. However, if a constant flip-angle (CFA) scheme is used (described in Chapter 2) as is the case for previous time-encoded PCASL angiography work,[3,4] signal discontinuities appear between images decoded from different encoding blocks (see Figure 6.7). This is because each RF pulse in the SPGR readout attenuates the ASL M_z by a factor of $\cos(\alpha)$, where α is the flip-angle, with larger flip-angles causing greater attenuation. A variable flip-angle (VFA) scheme is derived and implemented which maintains a constant signal across all excitations and removes these discontinuities. This VFA scheme is also shown to further improve SNR and that a time-encoded protocol with the VFA scheme can double the SNR compared to a scan time matched sequential CFA protocol, greatly enhancing the depiction of small distal vessels.

The use of time-encoding extends the scan time compared to standard sequential protocols which typically only acquire a single average. Therefore, initial results on the effects of reducing the scan time of the time-encoded protocol by angularly undersampling the radial acquisition are presented. Angular undersampling leads to an increase in noise, due to reduced *k*-space sampling, but also an increase the aliased energy. For radial acquisitions, undersampling artefacts manifest as streaks in the image which appear noise-like due to their dispersed nature, and so can be tolerable in angiography.[7]

The BGS implementation used in Chapter 5, which nulls static tissue signal before the start of the readout, was found to lead to a varying tissue signal level across the SPGR readout which led to a varying noise level. Therefore, improvements to the BGS to improve tissue nulling across the entire readout are described and preliminary data is presented.

Some of the work described in this chapter was presented at the 27th annual meeting of the International Society for Magnetic Resonance in Medicine.[*]

---

[*] Woods JG, Schauman SS, Chiew M, Chappell MA, Okell TW. Optimization of time-encoded pseudo-continuous ASL angiography with a variable flip-angle scheme. In: *Proceedings of the 27th Annual Meeting of the ISMRM, Montreal, Canada*. 2019, p 844.

6.2 Conventional PCASL Dynamic Angiography

As described in Chapter 3, PCASL angiography differs from standard perfusion imaging in that the ASL tracer is imaged while still in the large arteries, rather than the tissue, requiring much shorter PLDs than perfusion imaging. Also, since the arteries at the Circle of Willis typically have outer diameters between 1.4 mm - 4.2 mm,[8] with downstream arteries being even narrower, higher spatial resolutions are required for angiography than for perfusion imaging. However, since the tracer signal is concentrated in the large arteries and has undergone less T_1 decay, SNR is higher and so higher spatial resolutions are achievable.

In order to achieve high spatial and temporal resolutions, only a small amount of *k*-space is sampled during each readout frame, with the acquisition being segmented across multiple ASL preparations.[9] For instance, if a 192×192 matrix is sampled with a Cartesian trajectory and 12 *k*-space lines are acquired during each frame, then 16 ASL preparations (segments) will be required to fill all 192 *k*-space lines. The scan time is, therefore, directly proportional to the number of segments, all else being equal.

6.3 Time-Encoded PCASL Dynamic Angiography

Time-encoded PCASL dynamic angiography differs from conventional PCASL angiography in that some of the temporal information is encoded into the PCASL preparation itself, rather than achieving all of the temporal resolution with the temporally resolved readout. Similarly, to the perfusion sequences in Chapter 5, a Hadamard-encoding matrix can be used to efficiently encode the label and control conditions into the PCASL preparation for angiography, decoding the temporal resolution during image decoding.

When a time-encoded preparation is used in conjunction with multiple readout frames, the number of temporally resolved images is equal to $N \cdot M$, where N+1 is the number of Hadamard encoding cycles and M is the number of readout frames. The resulting temporal information differs from time-encoded data with only one PLD (used for the standard time-encoded perfusion protocols in Chapter 5), but is similar to the hybrid methods, introduced in Chapter 5, which have varying PLDs for the same encoding matrix. The difference being here that an SPGR readout is used, giving multiple images per ASL

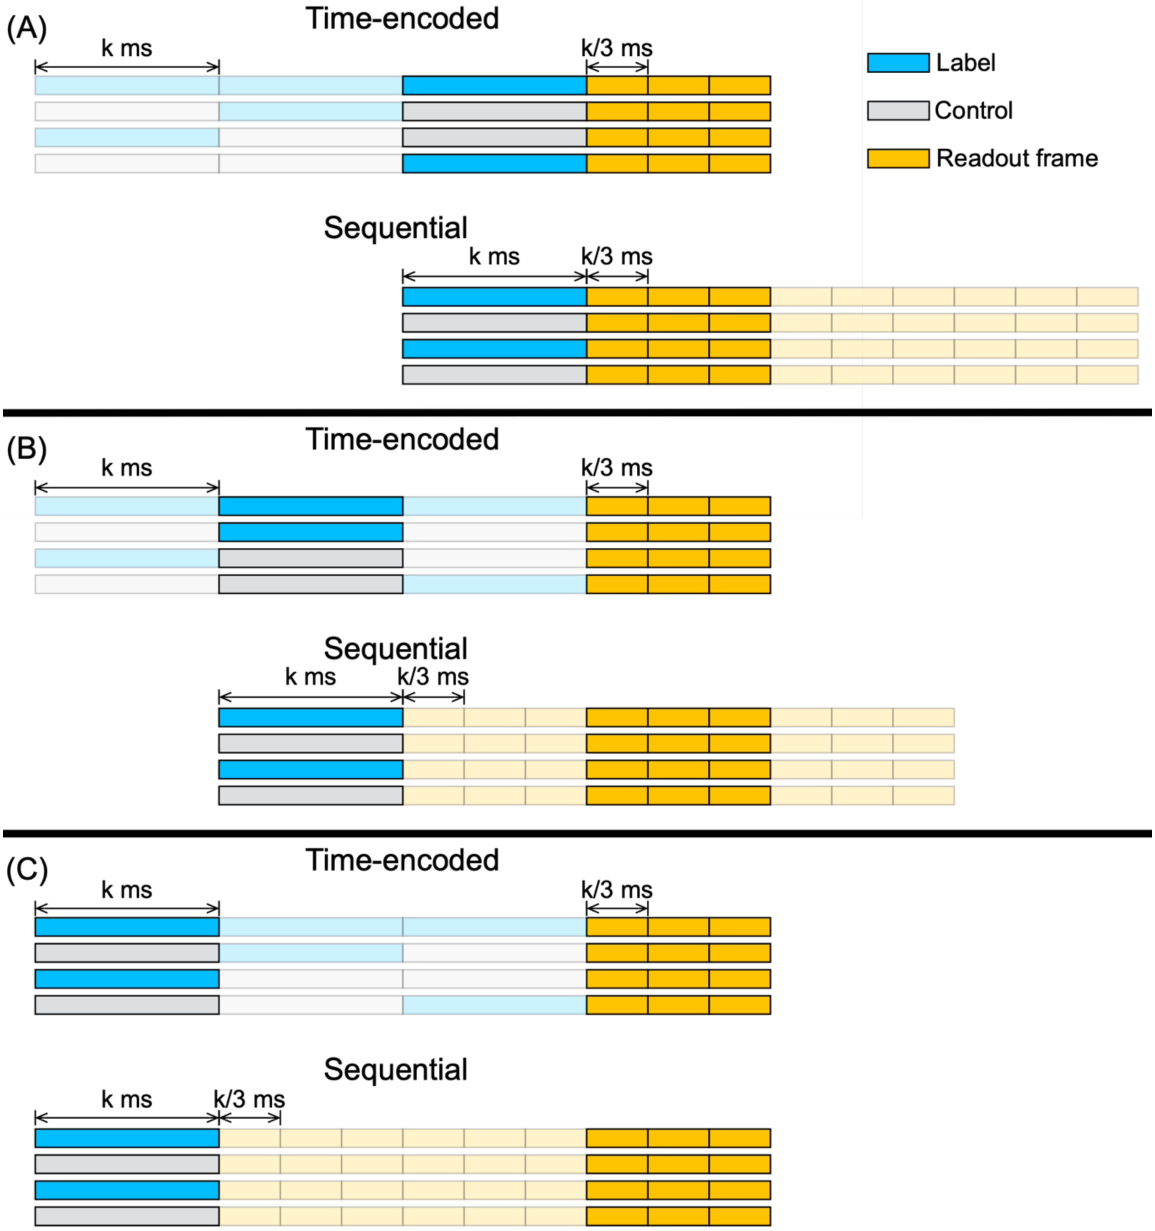

Figure 6.1: Demonstration of how matched timing sequential and time-encoded angiography protocols can achieve the same temporal information. (A), (B), and (C) demonstrate the symmetry between the protocols when decoding each time-encoded block for the 3 readout frames. In this example, the LD is set to be 3 times larger than the duration of each temporal frame in the readout so that the decoded images have the same LDs and PLDs for each protocol. However, this is not essential. The sequential protocol has two averages to match the scan time of the time-encoded protocol.

preparation. The SPGR acquisition is segmented over multiple preparations to decrease the temporal footprint of each image and increase the temporal resolution.

An example of how a time-encoded PCASL angiography protocol can acquire the same images as a standard sequential protocol is shown in Figure 6.1. In this example, the LD and readout frame duration are linked by the size of the encoding matrix (LD = 3·(frame duration)) so that the protocols give temporally matched images - for the time-encoded protocol, the 3 readout frames uniformly span each encoded LD block. The LD and frame duration can be chosen to be any value, but in this work, this design was chosen to give matched uniformly spaced images making it simpler to compare sequential and time-encoded protocols.

Since all of the temporal information in the sequential protocol is acquired using a time-resolved readout, Hadamard-encoding does not have the scan time reduction benefits it had for the perfusion acquisitions. Additionally, there is no noise averaging benefit either: the time-encoded protocol in Figure 6.1 acquires 4 images, giving a $\sqrt{2}$ reduction in noise, but the scan time of the time-encoded protocol is equivalent to 2 averages of the sequential protocol, which also results in a $\sqrt{2}$ reduction in noise. The benefit of a time-encoded preparation for dynamic PCASL angiography comes from a reduction in the number readout frames that are required to acquire the same temporal information, therefore allowing larger flip-angles to be used (see Figure 6.3).

6.4 Sequence Implementation

There was an existing sequence in our group that incorporated PCASL labelling with the SPGR readout described in Chapter 3. The sequence outline is shown in Figure 6.2(A). It uses a WET presaturation module[10,11] for BGS at the start of each ASL TR (where ASL TR refers specifically to the repetition time for the entire label or control preparation and readout), followed immediately by the PCASL labelling module,[12] as described for the perfusion sequences in Chapters 4 and 5. The key difference between the perfusion sequences and this angiographic sequence is that the readout is a segmented temporally-resolved SPGR readout.[9] The readout excitation train is split into separate ‘frames’, where the same lines of *k*-space are acquired during each frame for each ASL preparation. The remaining *k*-space lines are acquired during subsequent ASL TRs with identical BGS and PCASL preparation. The segmented *k*-space acquisitions for each frame are then combined before reconstruction.

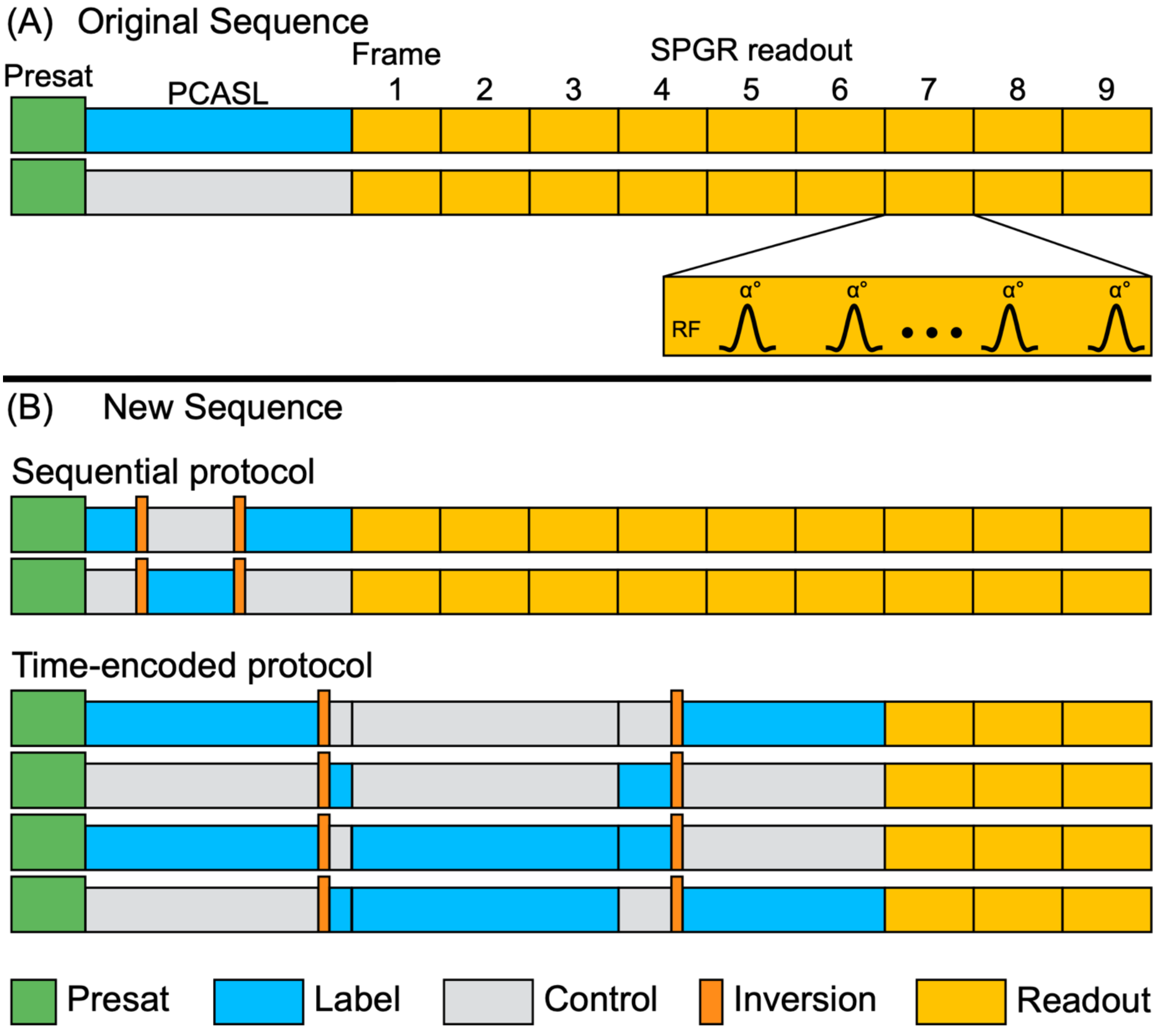

Figure 6.2: Schematic of the original PCASL dynamic angiography sequence (A) and the new sequence (B) which is capable of performing both sequential and time-encoded experiments. The new sequence also incorporates the interleaved FOCI pulses introduced in Chapter 5, leading to improved BGS.

To enable the use of time-encoded PCASL labelling for angiography and inversion pulses interleaved with the labelling period, this sequence was combined with the time-encoded PCASL preparation module and interleaved FOCI pulses described in Chapter 5.

These sequence changes can be viewed in the simplified schematics shown in Figure 6.2(B), where a short LD sequential and time-encoded protocol are shown with BGS (presaturation and 2 interleaved inversion pulses). These sequences have matched LD and PLD timings for all of the decoded temporally resolved images, as in Figure 6.1.

6.5 Flip Angle Optimisation

During the SPGR readout, each RF pulse attenuates the ASL signal, reducing the available longitudinal magnetisation of the tracer for the next excitation pulse. The choice of flip-angle is, therefore, a trade-off between SNR at early time frames and later frames (Figure 6.3). When using PCASL labelling, at a given position inside an artery there is a continual supply of labelled blood (for the duration of the labelled bolus) which has experienced the same amount tracer loss due to T_1 relaxation.[13] Therefore, when choosing the excitation flip-angles, T_1 relaxation does not need to be accounted for, only the effect of the previous flip-angles. Throughout this work, it is assumed that all of the ASL bolus is equally attenuated by each excitation pulse.

Each of the flip-angle formulas described below were optimised using a grid search across a range of flip-angles 1° to 60°, sampled every 0.1°. The optimisations used the formulas: $M_z(n) = M_z(n-1) \cdot \cos\big(\alpha(n)\big)$ and $M_{xy}(n) = M_z(n-1) \cdot \sin\big(\alpha(n)\big)$, where $M_z(n)$ and $M_{xy}(n)$ are the longitudinal and transverse magnetisation after the nth excitation pulse (n = 1,..,N), respectively, $\alpha(n)$ is the nth flip-angle, N is the total number of excitations after each ASL preparation, $M_{xy}(0) = 0$, and $M_z(0) = 1$. These simulations assume a constant supply of labelled blood throughout the entire readout, instantaneous RF pulses and perfect spoiling of residual transverse signal before the next excitation.

There are several signal characteristics which could be optimised, such as maximising the mean signal or maximising the minimum signal. The tracer bolus will have travelled furthest down the vascular tree by the end of the readout, and so will be in the smaller distal vessels where the labelled signal is less concentrated and has undergone a greater degree of RF attenuation. It is, therefore, important to ensure there is sufficient signal by the end of the readout period to visualise these downstream arteries. The visibility of distal vessels has previously been used as a sequence performance measure.[14–16] A simple method for preserving signal is to maximise the minimum signal acquired across the entire readout, ensuring that the signal acquired at any given excitation is not much lower than for any other excitation.

6.5.1 Constant Flip-Angle Scheme

For a constant flip-angle (CFA) scheme, where each excitation flip-angle across the whole readout is identical, the transverse signal at the n^{th} excitation will always be lower than at the $(n-1)^{\text{st}}$ excitation. In this case, maximising the signal at the last excitation is identical to maximising the minimum signal. Figure 6.3 shows the acquired signal for a range of different CFAs for 108 excitations (the number of excitations used for the sequential protocol described in Section 6.5.3). A flip-angle of 5.5° maximises the signal at the last excitation, larger flip-angles use up the available longitudinal magnetisation earlier during the readout, while smaller flip-angles do not excite enough of the magnetisation to produce a strong signal.

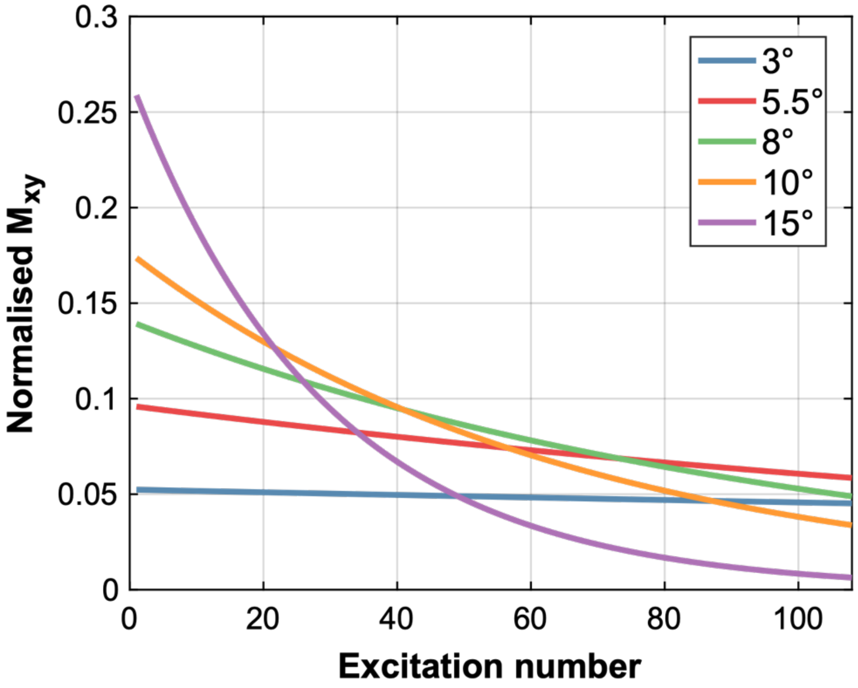

Figure 6.3: The simulated transverse signal acquired with a CFA scheme for a range of flip-angles when using an SPGR readout with 108 excitations. The simulations assume a single voxel with continually inflowing labelled blood and so do not take T_1 relaxation into account.

6.5.2 Variable Flip-Angle Schemes

6.5.2.1 Quadratic Formula

As an alternative to the CFA scheme, the flip-angles can be varied across excitations. Schmitt et al.[2] suggested using a quadratic formula to vary the flip-angles (referred to as $\mathrm{VFA}_{\mathrm{quadratic}}$), enabling higher signal at late excitations than can be acquired with CFAs:

$$\alpha(n) = \mathrm{FA}_1 + (\mathrm{FA}_2 - \mathrm{FA}_1) \cdot ((n-1)/(N-1))^2, \tag{6.1}$$

where FA_1 and FA_2 are the first and last flip-angle, respectively. This formula can ramp up the flip-angles over the course of the readout and so it is possible to approximately compensate for the attenuating effect of the previous RF pulses. For this scheme, it is necessary to simulate the signal across the entire readout in order to maximise the minimum signal because it is not guaranteed to occur at the last excitation.

Figure 6.4 shows the flip-angles and the acquired signal for a range of different values for FA_1 and FA_2 for 108 excitations. The flip-angles which maximise the minimum signal are $\mathrm{FA}_1 = 4.9°$ and $\mathrm{FA}_2 = 11.4°$. When FA_1 is lowered to 3° but FA_2 is fixed at 11.4°, the transverse signal at early excitations is lower, leaving more longitudinal magnetisation

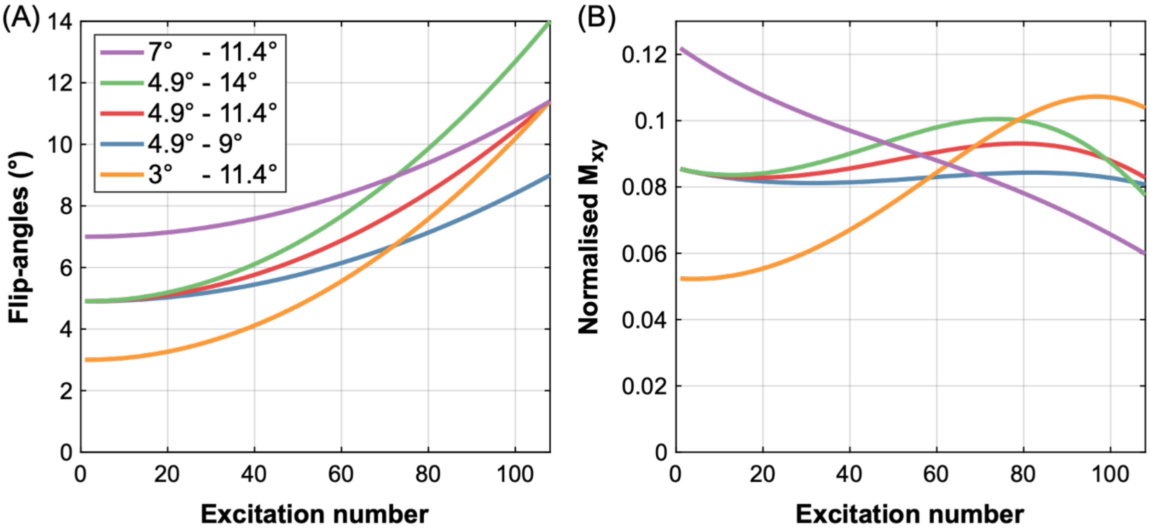

Figure 6.4: The flip-angles (A) and the simulated transverse magnetisation acquired (B) with the $\mathrm{VFA}_{\mathrm{quadratic}}$ formula (Equation 6.1) for a range of FA_1 and FA_2 when using an SPGR readout with 108 excitations.

available when the flip-angle ramps up, leading to higher signal at the end of the readout. The opposite is true when FA_1 is set higher than the optimal but FA_2 is still fixed at 11.4°. When FA_1 is fixed at 4.9° but FA_2 is lowered to 9°, the flip-angle ramp leads to lower acquired signal across the entire readout. If FA_2 is instead increased to 14°, then the flip-angles ramp up more quickly, using up more of the longitudinal magnetisation during the middle of the readout, leading to lower signal at the end of the readout.

6.5.2.2 Maintain Formula

Another VFA scheme which can be used is one which ramps up the flip-angles in such a way that each RF pulse directly accounts for the attenuating effect of the previous pulse. This formula is derived below in a similar way to Wang et al.,[1] except that T_1 relaxation is not taken into account, as explained previously.

To maintain a constant transverse signal across excitations, the following condition must be met:

$$M_{xy}(n+1) = M_{xy}(n), \tag{6.2}$$

where $M_{xy}(n)$ is the transverse magnetisation after the n^{th} excitation. This is equivalent to

$$M_{xy}(n+1) = M_z(n) \cdot \sin\big(\alpha(n+1)\big) = M_z(n-1) \cdot \sin\big(\alpha(n)\big) = M_{xy}(n), \tag{6.3}$$

where $M_z(n)$ is the longitudinal magnetisation just after the n^{th} excitation. Since

$$M_z(n) = M_z(n-1) \cdot \cos\big(\alpha(n)\big), \tag{6.4}$$

we have that

$$M_z(n-1) \cdot \cos\big(\alpha(n)\big) \cdot \sin\big(\alpha(n+1)\big) = M_z(n-1) \cdot \sin\big(\alpha(n)\big). \tag{6.5}$$

Solving for $\alpha(n)$, to generate a backwards recursive formula, gives

$$\alpha(n) = \tan^{-1}\Big(\sin\big(\alpha(n+1)\big)\Big). \tag{6.6}$$

This formula (referred to as VFA_{maintain}) is independent of the initial longitudinal magnetisation and is maximised when $\alpha(N) = 90°$, where all of the available magnetisation will have been used during the readout. Zhao et al.[17] demonstrated that when $\alpha(N) = 90°$, this formula can be expressed as

$$\alpha(n) = \tan^{-1}\left(\frac{1}{\sqrt{N-n}}\right). \tag{6.7}$$

However, Equation 6.6 is more general, and so is used in this work.

Equation 6.7 inherently maximises the minimum acquired signal by ensuring that the signal is maximally constant across all excitations. There are, however, several potential problems associated with using a final flip-angle of 90°. The first is that the SAR can be high when using this formula with minimum ASL TR protocols, as is done here. The second is that care needs to be taken to avoid voltage clipping. It was necessary to reduce the time-bandwidth-product (TBWP) of the excitation pulses due to the short pulse durations. Thirdly, the static tissue signal recovers over the course of the readout. Since higher static tissue signal leads to an increase in subtraction errors related to physiological noise and aliased signal, the apparent noise in the final angiograms will increase during the readout for all of the flip-angle schemes given above, assuming the tissue signal is nulled before the start of the readout. However, because of the large ramp in flip-angles at the end of the readout for the $\text{VFA}_{\text{maintain}}$ formula, the acquired static tissue signal will be relatively larger (see Section 6.7 for more details on the BGS). To mitigate all three of these problems, $\alpha(N)$ was set equal to 30°. For a readout with 36 excitations (used for the time-encoded protocol described in Section 6.5.3), the final excitation will only have 3.9% lower ASL signal compared to $\alpha(N) = 90°$, but 7.5% lower static tissue signal (assuming T_{1t} = 0.791 s for WM at 3T[18], TR = 10 ms (the excitation TR during the SPGR readout), and the static tissue is nulled 100 ms before the first excitation).

Figure 6.5 shows the $\text{VFA}_{\text{maintain}}$ flip-angles and the signal generated for 108 excitations. The first flip-angle is ~5.4° and the last is 30°. The flip-angles ramp up in such a way so as to perfectly balance the effect of the previous excitation pulse. This means that the acquired signal maintains at a constant magnitude across the entire readout, where the signal magnitude depends only on the number of excitations during the readout. As a comparison, the VFA$_{\text{quadratic}}$ scheme is plotted with FA$_1$ and FA$_2$ equal to the first (5.4°) and last (30°) flip-angles of the VFA$_{\text{maintain}}$ scheme. The VFA$_{\text{quadratic}}$ formula initially ramps up the flip-angles much more quickly, leading to very high signal in the middle of the readout but very low signal at the end.

6.5.3 Simulation Comparison

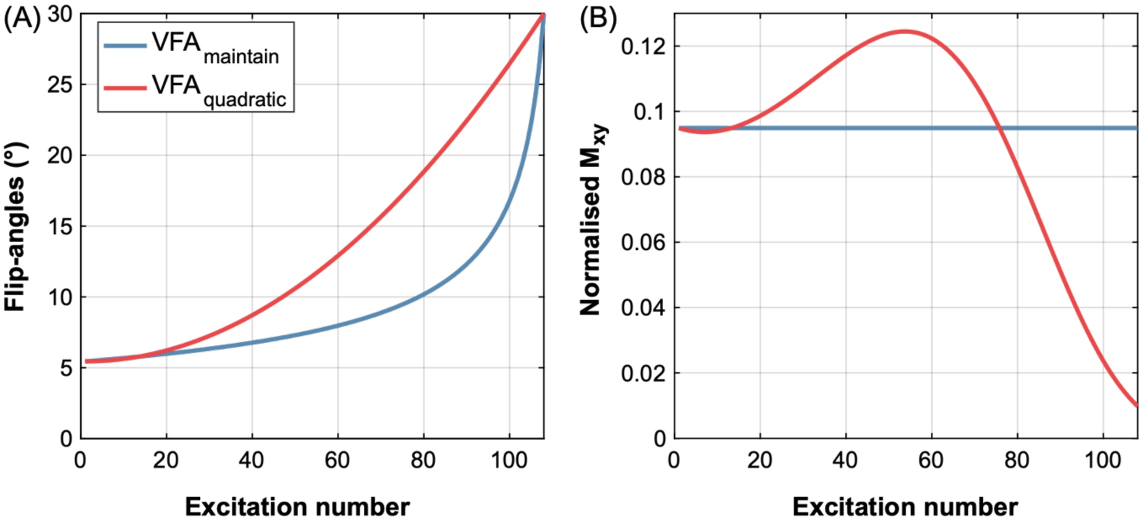

Figure 6.5: The flip-angles (A) and the simulated transverse signal acquired (B) with the VFA$_{\text{maintain}}$ *formula (Equation 6.6) when using an SPGR readout with 108 excitations. The VFA*$_{\text{quadratic}}$ *scheme with the same first and final flip-angles is shown as a comparison.*

6.5.3.1 Simple Signal Model

The above 3 flip-angle formulas were optimised to maximise the minimum signal, for timing-matched sequential and time-encoded protocols. For the sequential protocol, 9 readout frames with 12 excitations in each frame were chosen, resulting in 108 excitations overall. This was found to be a reasonable trade-off in terms of the single-line readout TR (10 ms), the temporal coverage of the total readout (1080 ms), the temporal resolution, and the number of segments required to cover the full *k*-space (16 segments for a matrix size of 192 × 192 and a Cartesian acquisition) - see Section 6.6 for full details of the sequence settings. The time-encoded protocol used a 4 × 3 Hadamard encoding matrix and 3 readout frames, with other settings being identical to the sequential protocol. The LD of both protocols was 360 ms (3 times longer than the duration of each readout frame), in order to match the LDs and PLDs between the protocols and to enable visualisation of the inflow of labelled blood.

The effective PLDs (defined as the time from the end of labelling to the centre of each readout frame) for both protocols were: 62 ms, 182 ms, 302 ms, 422 ms, 542 ms, 662 ms, 782 ms, 902 ms, and 1022 ms (including the 2 ms spoiler at the end of the PCASL labelling).

Figure 6.6 shows the optimised flip-angles and resulting signal for each protocol using each flip-angle formula. The flip-angles used are given in Table 6.1 and the protocol timings in Table 6.2. For the time-encoded protocol, only the acquired 36 excitations are shown for clarity. When all 108 effective samples are decoded, the following 72 excitations are simply repeats of the data shown for this simplistic simulation.

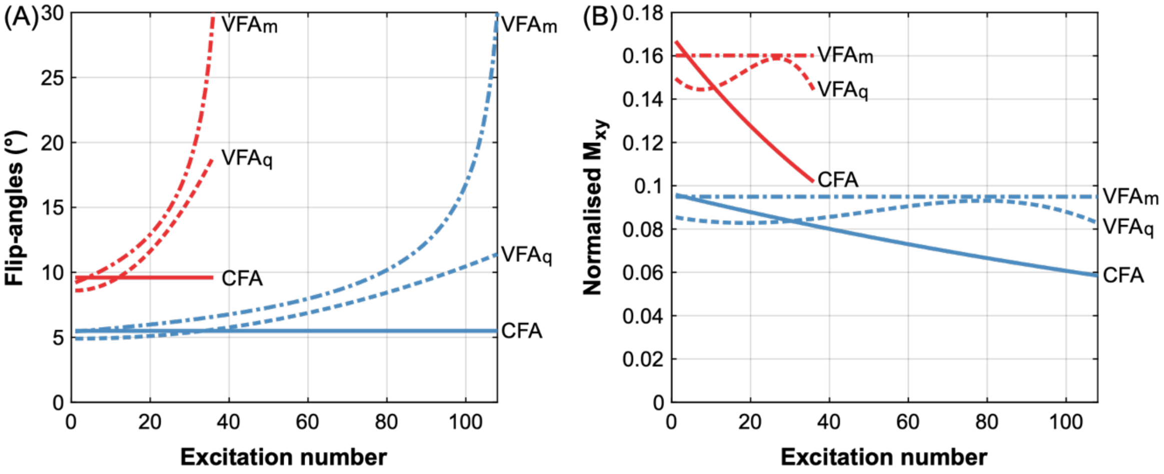

Figure 6.6: The flip-angles (A) and the simulated transverse signal (B) for the sequential (blue) and time-encoded (red) protocols with the three different flip-angle formulas: CFA, $VFA_{quadratic}$ (VFA_q), and $VFA_{maintain}$ (VFA_m). The flip-angles were optimised to maximise the minimum signal for each protocol using each formula. The time-encoded protocol has 36 excitations across 3 readout frames, while the sequential protocol has 108 excitations across 9 frames.

There are three key conclusions from Figure 6.6. (1) It is clear that whichever flip-angle formula is used, the time-encoded protocol achieves higher signal across the entire readout due to the much larger flip-angles which can be used. (2) Allowing the flip-angles to ramp up during the readout results in higher signal, particularly at the end of the readout. (3) For both the sequential and time-encoded protocols, the $\text{VFA}_{\text{maintain}}$ formula results in higher signal at almost all excitations in the readout. Due to the superior performance of $\text{VFA}_{\text{maintain}}$ over $\text{VFA}_{\text{quadratic}}$, the $\text{VFA}_{\text{quadratic}}$ scheme was not used in any further comparisons.

Flip-angle formula	*Flip-angles*
Sequential	
CFA	5.5°
$\text{VFA}_{\text{quadratic}}$	$\text{FA}_1 = 4.9°$, $\text{FA}_2 = 11.4°$
$\text{VFA}_{\text{maintain}}$	Max FA = 30°
Time-encoded	
CFA	9.6°
$\text{VFA}_{\text{quadratic}}$	$\text{FA}_1 = 8.6°$, $\text{FA}_2 = 18.9°$
$\text{VFA}_{\text{maintain}}$	Max FA = 30°

Table 6.1: The flip-angles used for each protocol with each flip-angle formula.

6.5.3.2 Realistic Signal Model

The above comparison does not take into account the duration or shape of the labelled bolus travelling through a vessel. For the following comparison, the angiography model described by Okell et al.[13] was used, which models a labelled bolus of blood passing

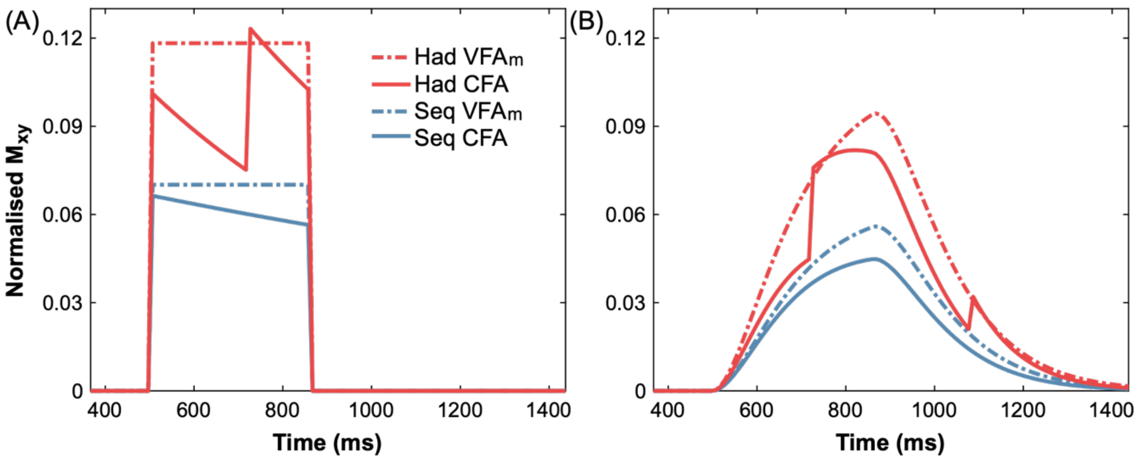

Figure 6.7: The acquired signal for the sequential (Seq) and time-encoded (Had) protocols with CFA and VFA$_{maintain}$ (VFA$_m$) flip-angle schemes assuming two different levels of dispersion with the angiography model: (A) no dispersion ($p = 0$ ms, $s = 1000$ s^{-1}) and (B) with dispersion ($p = 100$ ms, $s = 10$ s^{-1}). Other model parameters were: $A = 1$, $T_1 = 1.65$ s, $\Delta t = 0.5$ s. The signal was evaluated at each excitation. The time is shown relative to the start of labelling.

through a single voxel and incorporates RF attenuation, T_1 relaxation, and dispersion (characterised by a gamma dispersion kernel). However, here it is assumed that all of the labelled blood is attenuated by every excitation pulse, which is equivalent to setting $\delta_{t,min} = 0$ in the Okell et al.[13] model, where $\delta_{t,min}$ is the time at which the labelled blood first arrives in the imaging region.

Figure 6.7 shows the acquired signal when using the CFA and VFA$_{maintain}$ flip-angle schemes for each protocol with the angiography model. Two different cases are shown: (A) (approximately) no dispersion and (B) with moderate dispersion. The difference between the maximum signal in the no dispersion case (Figure 6.7(A)) and those in Figure 6.6(B) are due to T_1 relaxation of the tracer, which is simulated here. The signal discontinuities in the time-encoded CFA protocol are evident in these simulations. These discontinuities are a result of the signal progressively reducing across the readout. Each decoded bolus uses the same readout, so when the signal from all of the boluses are plotted in chronological order, jumps in the signal occur as if the signal had been reset several times during the ‘effective’ readout. Since the $VFA_{maintain}$ formula ensures a constant signal level across the entire readout, no discontinuities are evident and the signal time-course is smooth, similar to the sequential time-course. The conclusions from Figure 6.6 also still hold: that the time-encoded protocol leads to higher signal for either of the flip-angle schemes and that the $VFA_{maintain}$ scheme provides greater signal than the CFA scheme.

6.6 Experimental Comparisons

In order to validate the above simulation comparison, in vivo comparisons were performed for the sequential and time-encoded protocols described in Section 6.5.3 using the CFA and $VFA_{maintain}$ schemes, giving 4 protocols in the comparison: (1) sequential CFA, (2) sequential $VFA_{maintain}$, (3) time-encoded CFA, and (4) time-encoded $VFA_{maintain}$. The protocol timings are given in Table 6.2 and the flip-angles in Table 6.1.

Protocols	*LDs (ms)*	*Effective PLDs (ms)*	*Frames*	*Averages*	*Scan Time*
Sequential	360	62, 182, 302, 422, 542, 662, 782, 902, 1022	9	2	2:30 min (low res) 4:35 min (high res)
Time-encoded	360, 360, 360	62, 182, 302, 422, 542, 662, 782, 902, 1022	3	1	2:30 min (low res) 4:35 min (high res)

Table 6.2: Sequence timings. Low res and high res refer to the in-plane resolution of the acquisitions which is described in Section 6.6.1.1

6.6.1 Methods

6.6.1.1 Acquisition

Four volunteers (3 male, mean age 33 ± 12, range 22 - 48) were recruited and scanned under a technical development protocol, agreed with local ethics and institutional committees, on a 3T Verio system (Siemens Healthcare, Erlangen, Germany) with a 32-channel receive-only phased-array head coil. All scanning occurred during a single scan session for each volunteer. The volunteers were asked to lie still during the scan but were not required to stay awake.

Each scan started with a 3-plane localiser followed by a single-slab 3D TOF sequence (resolution $0.31 \times 0.31 \times 1.3$ mm^3) used to place the PCASL labelling plane, as described in Chapter 4.

The ASL angiography protocols were acquired using a 2D SPGR readout to facilitate rapid scans so that all of the protocols could be acquired and compared within a single scan session. Two different spatial resolutions were used: (1) a lower in-plane resolution of 1.15×1.15 mm^2 to facilitate short scans times enabling all 4 protocols to be compared, and (2) a higher in-plane resolution of 0.63×0.63 mm^2 to compare the ability of the $\mathrm{VFA}_{\mathrm{maintain}}$ protocols to depict the very small vessels not visible at lower resolutions.

The common imaging parameters for the angiography scans using the SPGR readout were: FOV 220×220 mm, slice thickness 70 mm, radial trajectory, TR 10 ms, 12 excitations per frame, temporal resolution 120 ms, RF spoiling[19] using a quadratic RF phase increment of 50°. No fat suppression was used. The 2D single-slice was positioned so that it covered the circle of Willis and some distal arteries, as shown in Figure 6.8.

PCASL labelling was achieved using 500 μs duration Gaussian RF pulses, 1 ms pulse interval, and 20° flip-angle, with the labelling plane positioned in the same way as described in Chapter 4. Background suppression was achieved using a WET presaturation

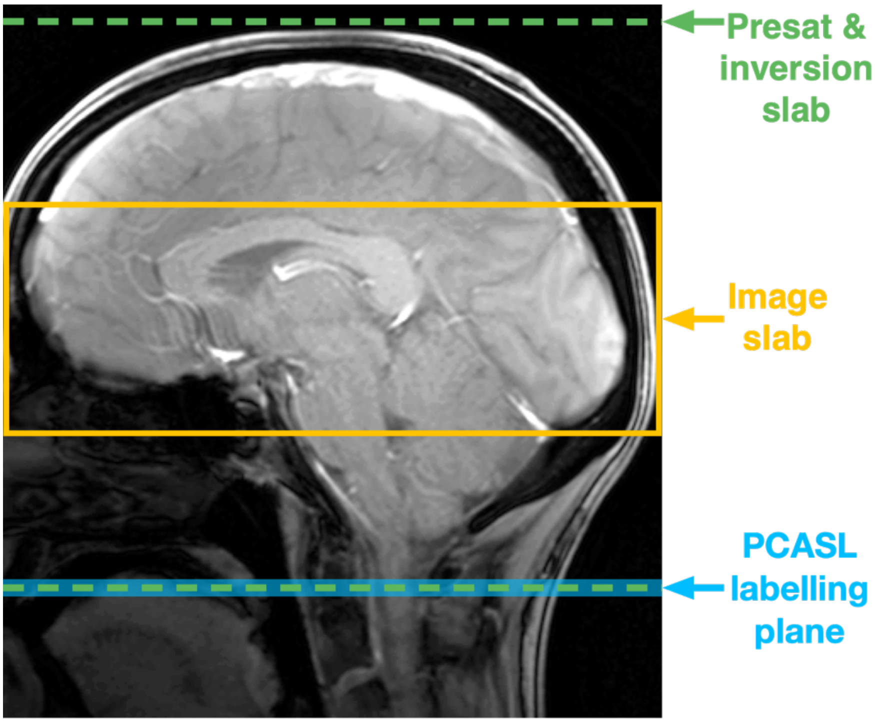

Figure 6.8: The approximate placement of the image slab, PCASL labelling plane, WET presaturation slab, and the FOCI inversion slab. Note, the presaturation and FOCI slab are placed identically with the superior edge above the brain and the inferior edge flush with the centre of the labelling plane. Covering the entire brain with the presaturation module ensures that all of the labelled blood from the previous label/control TR is saturated to avoid contaminating the signal in the current TR. The FOCI inversion slab is flush with the labelling plane in order to efficiently invert the blood that is already in the label/control condition, while not perturbing inflowing blood that is yet to reach the labelling plane.

module[10,11] and two FOCI pulses[20,21] timed using Equation 3.4 to null tissue 100 ms before the first excitation in the readout, $T_{1,opt}$ 700 ms, and positioned as shown in Figure 6.8 and as described in Chapter 5. The peak B_1^+ of the FOCI pulses was set to ~10 µT due to SAR restrictions and the shaping parameters were: A_{max} 20, µ 1.5, and ß = 1349.17 rad/s (see appendix D).

The 1.15 × 1.15 mm^2 in-plane resolution specific imaging parameters were: matrix 192 × 192, bandwidth 131 Hz/pixel, TE 5.08 ms, 300 radial spokes, 25 segments. The sequential protocols used 2 averages to match the scan time of the time-encoded protocols. This data was acquired for 3 of the volunteers.

The 0.63×0.63 mm^2 in-plane resolution specific imaging parameters were: matrix 352×352, bandwidth 138 Hz/px, TE 5.27 ms, 552 radial spokes, 46 segments. For this resolution, the sequential VFA$_{\text{maintain}}$ protocol used 1104 radial views (92 segments) in order to match the scan time of the time-encoded VFA$_{\text{maintain}}$ protocol. This factor of 2 oversampling will result in similar noise averaging to the 2 average case and so represents an alternative method for collecting twice as much data for the sequential protocol. This data was acquired for 3 of the volunteers.

In practice, ASL angiograms are often acquired with only a single average. In order to match the scan time of the time-encoded protocol to that of a single average of the sequential protocol, the time-encoded acquisition can simply be angularly undersampled by a factor of 2, halving the number of segments. To test the effect of angularly undersampling the time-encoded VFA$_{\text{maintain}}$ data, prospectively undersampled data in 3 of the volunteers was acquired to compare with the 1st average of the sequential VFA$_{\text{maintain}}$ data. The protocol settings were identical to the previously described settings for the 1.15×1.15 mm^2 in-plane resolution scans, except that only 156 radial spokes were acquired, requiring only 13 segments rather than 25, giving an angular acceleration factor of $\mathrm{R}_\theta \sim 1.93$ compared to a fully sampled acquisition. The scan time was 1:18 min, compared to 1:15 min for one average of the sequential VFA$_{\text{maintain}}$ protocol.

6.6.1.2 Reconstruction

Reconstruction of the data was performed offline in MATLAB for two reasons: (1) the way the sequence looping method discussed in Section 6.6.1.3 is implemented means that the data is not combined appropriately by the online reconstruction software and (2) offline reconstruction allows more advanced reconstruction techniques and corrections to be used without implementing them on the scanner. The offline reconstruction used tools from the

Michigan Image Reconstruction Toolbox,[22] including the non-uniform FFT (NUFFT),[23] and code written by Mark Chiew, Thomas Okell, and Sophie Schauman.

A common method for reconstructing images from non-uniformly sampled *k*-space data is to perform a gridding procedure where *k*-space samples are interpolated and re-sampled onto a uniform grid which can then be used with the IFFT algorithm. The NUFFT is a general form of gridding, where the optimal interpolation kernel is precomputed for each local *k*-space neighbourhood. The implementation used here finds the interpolation kernel that minimises the worst-case approximation to the true signal by using min-max optimisation.[23]

The reconstruction process was implemented as follows:

1. Generate the radial *k*-space sampling trajectory.
2. Estimate the coil sensitivities using the adaptive combine algorithm.[24]
3. Generate the NUFFT operator[23] using the *k*-space samples and coil sensitivities.
4. Generate the density compensation weights using an iterative procedure.[25]
5. Align the phase of each radial spoke to that of the corresponding spoke in the first image.
6. Apply the density compensation weights to the data.
7. Apply the adjoint of the NUFFT operator.
8. Subtract/decode the reconstructed complex images.
9. Take the magnitude of the complex difference images.

For the coil sensitivities, *k*-space is first averaged across PCASL conditions before images are reconstructed for each coil using a naïve (knowledge of the trajectory only) NUFFT reconstruction and then averaged across readout frames. The adaptive combine algorithm[24] provides a near-optimal method for combining multiple coil data and can be used to estimate the coil sensitivities; a kernel size of approximately 10 mm was used for both resolutions. Incorporating the coil sensitivities into the NUFFT operator is equivalent to performing Roemer coil combination.[26] Parallel imaging was not used. The density compensation weights are necessary to pre-weight the *k*-space data due to the variable density with which *k*-space is sampled with a radial trajectory; the density being inversely proportional to the distance from the centre. This down-weights the oversampled centre of *k*-space data and upweights the less densely sampled outer *k*-space data.

The phase alignment step assumes that the static tissue signal is the dominant source of signal in each image. Since the complex static tissue signal in each image should be identical, any difference in phase is assumed to be due to effects such as B_0 drift. By aligning the phase for each *k*-space spoke, this drift can be minimised. If this is not done, any artefactual difference in phase will cause a reduction in the apparent ASL signal. For this step, the mean difference in the phase of each radial spoke to that of the same spoke from the first image is subtracted from the phase of each spoke. A weighted mean is taken across coils and the samples along each radial spoke, weighted by the magnitude of the signal. This is done because estimating the phase from low magnitude data is inherently less accurate.

The difference images are subtracted/decoded and averaged after reconstruction of the complex label/control/encoded images. In the case of undersampled radial data, artefacts that are caused by convolution of the PSF with the static tissue signal will be subtracted out in an ideal subtraction/decoding process, although artefacts from the ASL difference signal will remain.

6.6.1.3 Sequence looping

To further reduce phase differences due to B_0 drift, the sequence first loops over the ASL label and control conditions and then over the *k*-space segments. For example, the first 12 radial lines of the acquisition are acquired in the label condition for each readout frame and then, in the following ASL TR, those same lines of *k*-space are acquired in the control condition. This minimises the time, and therefore B_0 drift, between matched lines of *k*-space. For the time-encoded data, each encoding is looped over in the same way. Any averages are also acquired within the same inner loop.

6.6.1.4 Quantitative SNR Analysis

To quantitatively compare the protocols in vivo, a vessel mask was created for each subject using the combined temporal mean across all of the data included in each comparison to improve the SNR of the angiograms and to reduce the risk of biasing the results towards one protocol. For each protocol, the mean signal within the vessel mask was divided by the root-mean-squared (RMS) of the noise within a background ROI to find the SNR at each temporal frame. The RMS was used as a measure of the background noise because in magnitude data the noise converges to a Raleigh distribution as the magnitude goes to zero. Assuming that the complex background noise is normally distributed with zero mean and with equal SD, σ, for the real and imaginary parts, the RMS of the magnitude noise is equal to $\sigma \cdot \sqrt{2}$. The factor of $\sqrt{2}$ was not corrected for because it will simply scale the calculated SNRs in an identical manner for all of the protocols. Due to the use of coil sensitivities in the reconstruction, the noise magnitude will vary spatially across the image. However, this will occur in a similar manner for all of the protocols within each comparison meaning that it is still possible to compare their SNR values.

The method used to create the vessel masks, and a depiction of the background ROI used, is shown in Figure *6.9*. The pipeline begins with the temporal mean of the combined data. 1) The background signal is extracted to identify non-vessel specific signal (using a ~10 mm disk kernel and MATLAB's 'strel' and 'imopen' functions, which perform a morphological opening operation which removes fine details but preserves broad signal variations). 2) The background signal is subtracted to improve the contrast between the background and vessel signal. 3) The contrast is enhanced using MATLAB's 'imadjust' function to saturate the top 1% of pixels. 4) A subject specific brain mask (generated using FSL's BET on the mean unsubtracted image) is applied to the image to remove high-contrast artefacts and external arteries. 5) A global threshold is applied to the image to isolate the vessels (threshold was 70% of the threshold calculated with Otsu's method[27]), then the image is binarised. 6) Clusters of less than 10 contiguous pixels are removed using

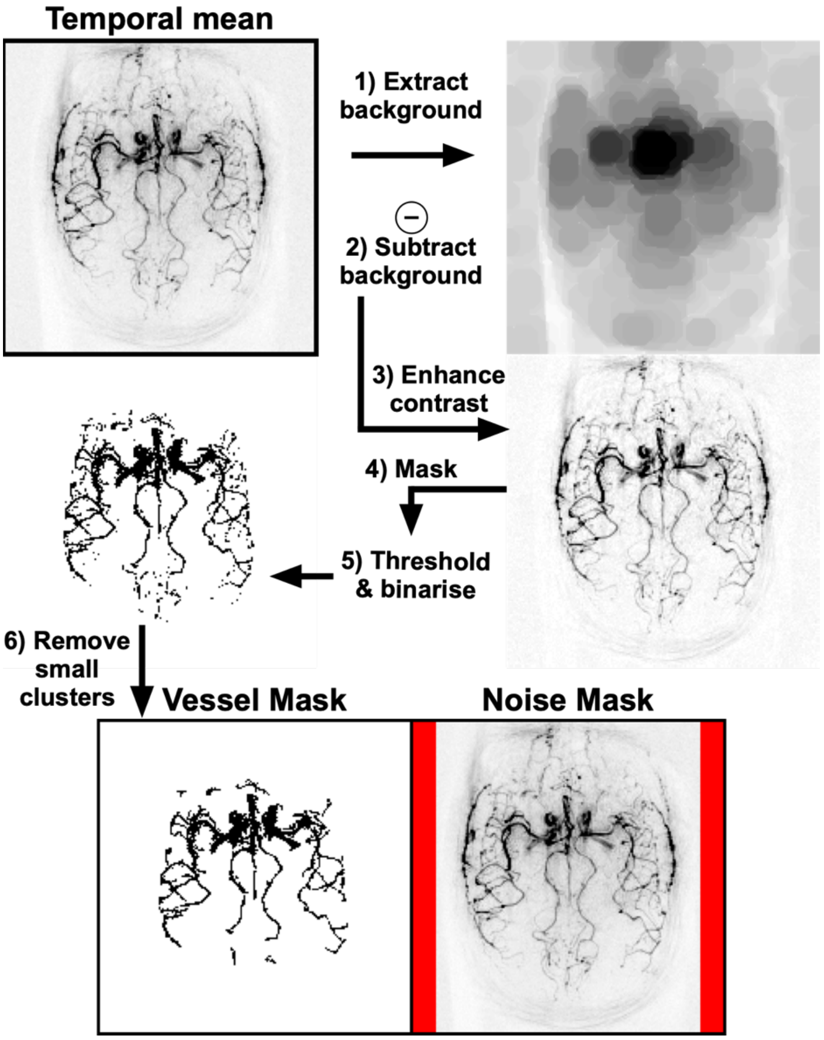

Figure 6.9: Processing pipeline for creation of the vessel masks. The contrast of each of the images shown here was enhanced to clearly depict each step.

MATLAB's 'bwareaopen' function. The noise mask was two rectangular ROIs on either side of the images that did not overlap with the head of any of the subjects.

6.6.2 Results

6.6.2.1 Phase correction

The reconstruction was performed with and without the phase-correction step for the sequential $\text{VFA}_{\text{maintain}}$ data for 1 subject to investigate the effect this has on the reconstructed angiograms. The phase correction step was found to improve the robustness of the reconstruction to signal drift, as demonstrated in Figure 6.10. When the phase-correction step is used, static tissue subtraction artefacts are reduced compared to the no phase-correction reconstruction. This is particularly evident around the edge of the brain.

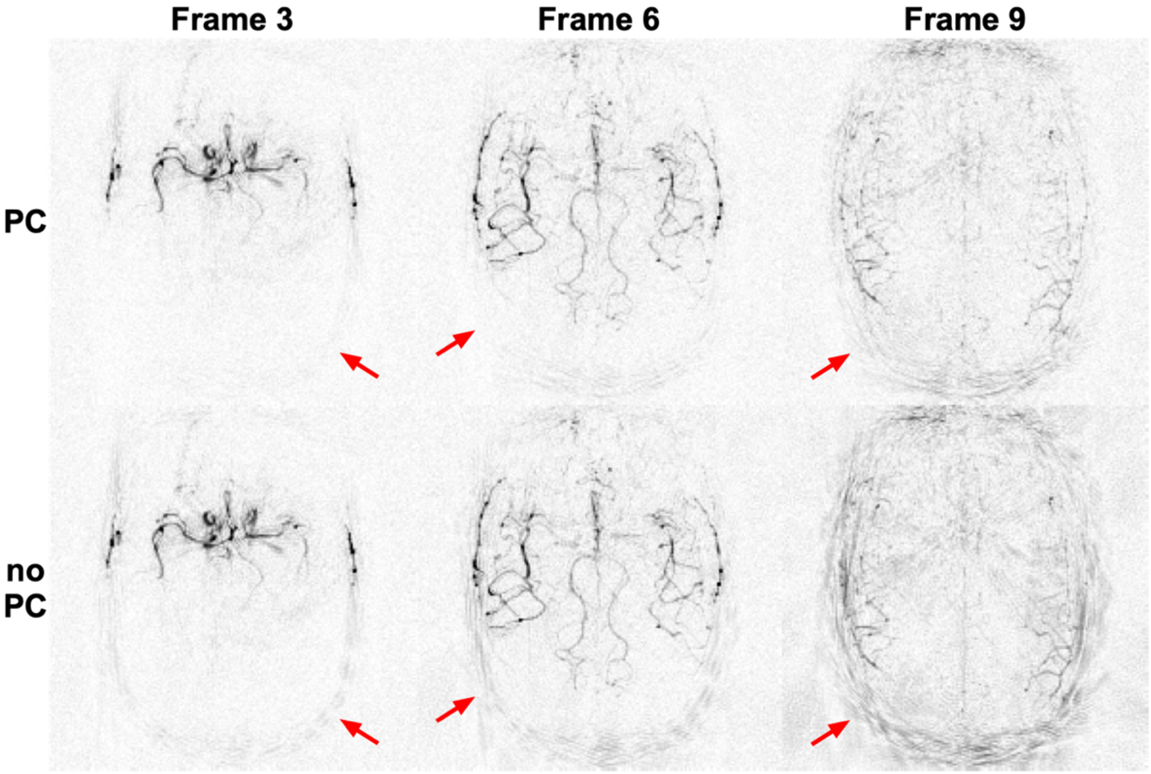

Figure 6.10: A comparison of the sequential $\text{VFA}_{\text{maintain}}$ *reconstructed data with (top row) and without (bottom row) the phase correction step in the reconstruction for 1 subject. The data shown are frames 3, 6, and 9, for a single subject. Use of the phase-correction step greatly reduces static tissue subtraction artefacts which are particularly visible around the edge of the brain (highlighted with arrows) and improves the depiction of the small distal vessels in the final frame.*

The last frame shows the most obvious static tissue artefacts and so the clearest improvement (seen across the whole brain, but particularly clearly at the edge of the brain) due to the more relaxed static tissue signal and the larger flip-angles used for this image. Phase correction was therefore used for the remainder of this chapter.

6.6.2.2 Lower Resolution: 1.15×1.15 mm^2

Example frames from one subject are shown in Figure 6.11. It is clear that the time-encoded data has higher arterial signal than the sequential data for both flip-angle formulas

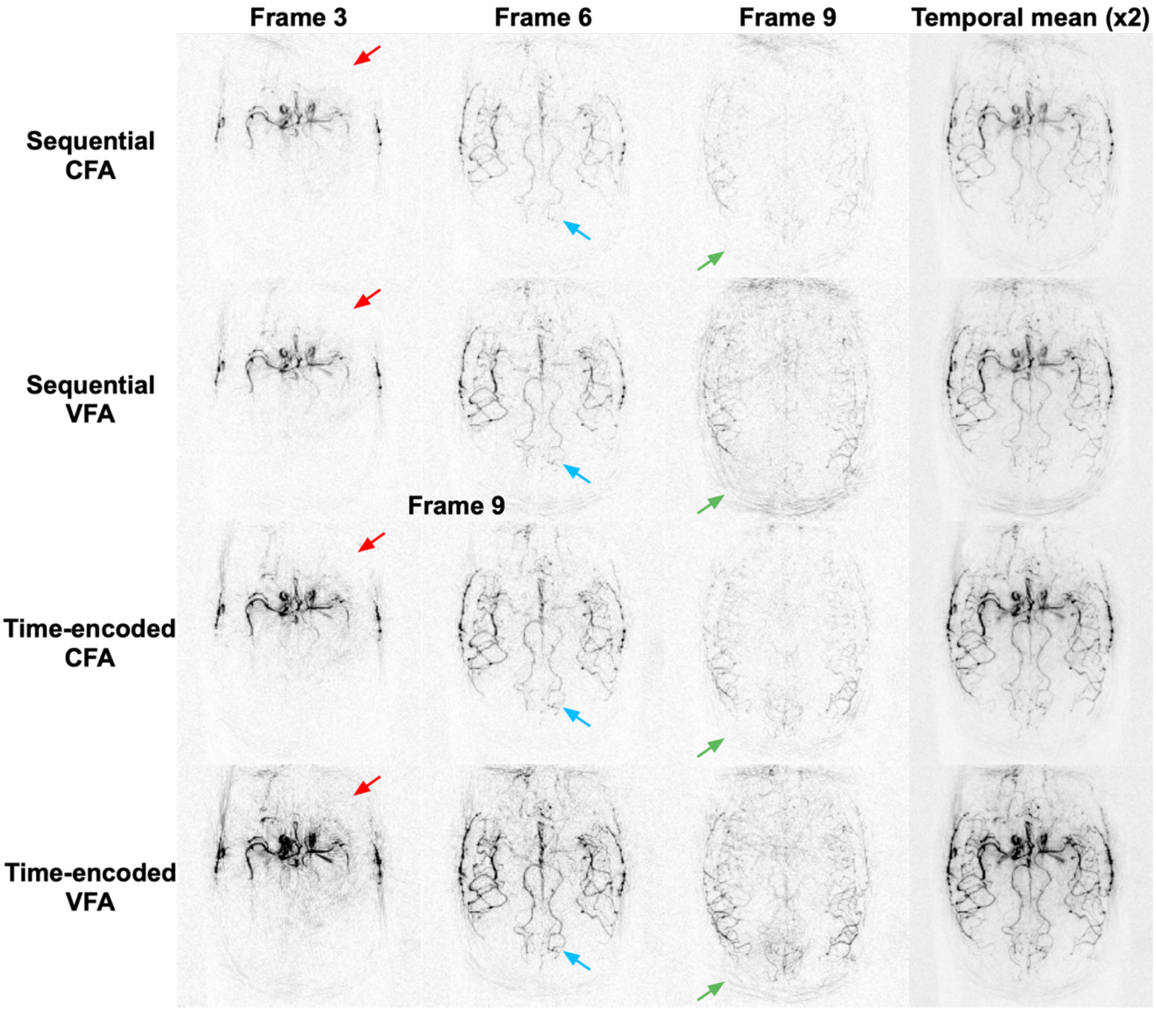

Figure 6.11: Selected frames (3rd, 6th, and 9th) and the temporal means for a single subject for the resolution 1.15×1.15 *mm*^{2}*. The red arrows highlight the increased static tissue signal in the time-encoded* $VFA_{maintain}$ *data. The blue arrows highlight the increased signal in the distal vessels for the time-encoded* $VFA_{maintain}$ *data. The green arrows highlight the subtraction errors at the back of the head which is lower in both CFA images, but is also lower in the time-encoded* $VFA_{maintain}$ *data compared to the sequential* $VFA_{maintain}$ *data. All of the images are shown on the same intensity scale except the temporal means which have been magnified by a factor of 2.*

and the VFA formula increases the signal for both sequential and time-encoded protocols, as expected from the simulations. The data from the 2 other volunteers show qualitatively similar trends.

The frames chosen for Figure 6.11 correspond to the highest level of tissue signal in the encoded images for the time-encoded protocols, due to recovery of the magnetisation during the readout and the larger flip angles towards the end of the VFA schedule, and therefore have the largest amount of static tissue signal artefacts. Comparatively, the static

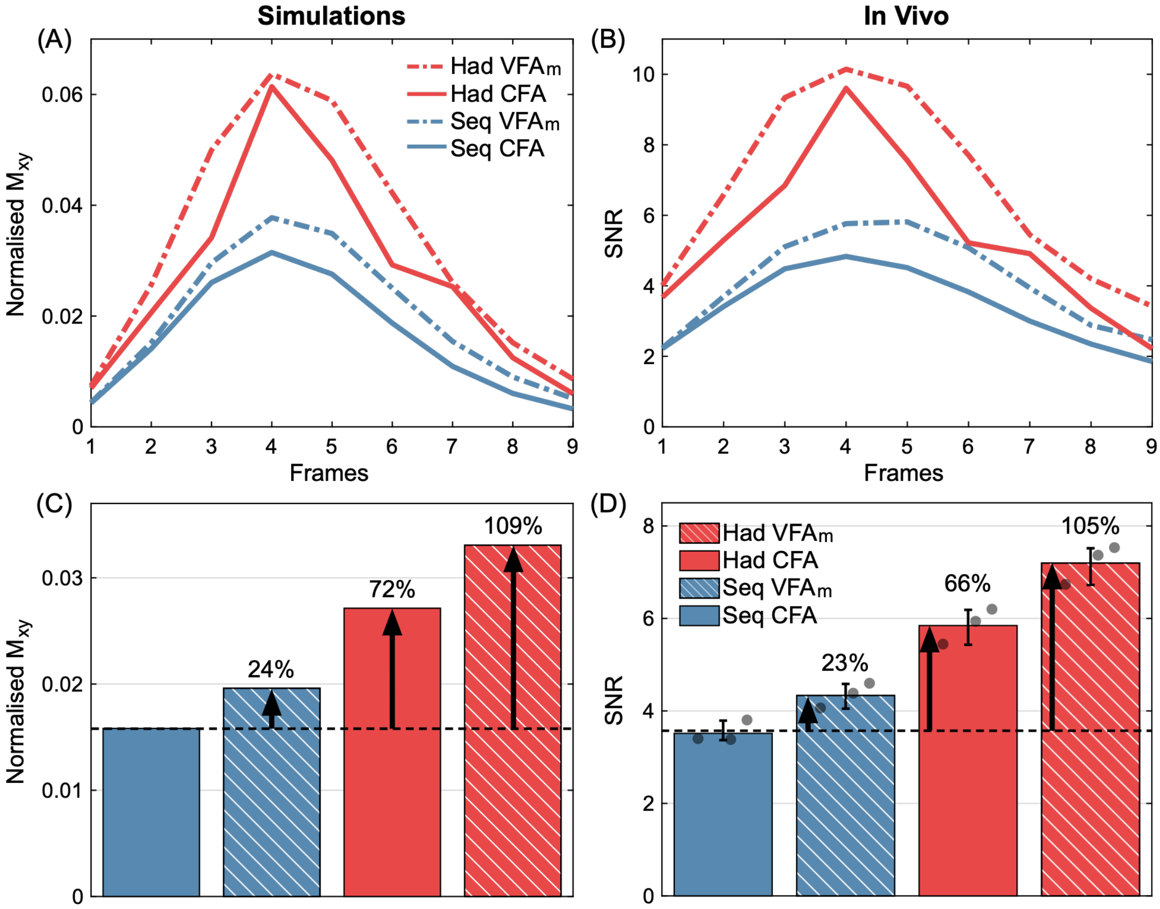

Figure 6.12: SNR comparisons of the sequential (Seq) and time-encoded (Had) protocols using the CFA and VFA$_{maintain}$ (VFA$_m$) flip-angle schemes. Simulations of the acquired signal using the angiography model described previously (A,C) are compared to the SNR measurements from a single subject (B) and the average SNR across 3 subjects (D) using the vessel and noise masks described in Figure 6.9. *The top row shows the simulated signal and in vivo SNR for each frame, while the bottom row shows the average across all frames. The simulations show the mean of a range of ATTs to better reflect the in vivo data:* Δt = *0.3 s - 0.6 s, sampled every 25 ms. Other parameters were:* p = *100 ms,* $s = 5\ s^{-1}$, $A = 1$, $T_1 = 1.65$ *s. In (D), each subject's mean SNR is shown as a dot, while the bar graphs show the mean and range across subjects. The relative improvement in SNR of each method compared to the sequential CFA protocol is displayed.*

tissue signal in the sequential data can be seen to increase steadily across all frames, with frame 9 appearing to have a larger level of static tissue signal artefacts at the back of the head than the time-encoded data. This is due to the reduced T_1 relaxation over the course of the shorter readout for the time-encoded protocol.

The SNR measurements using the vessel and noise masks for 3 subjects are shown in Figure 6.12. It is clear that the use of the $VFA_{maintain}$ flip-angle scheme improved the SNR for both the sequential and time-encoded protocols compared to the CFA scheme. These results also confirm that the time-encoded protocol yields higher SNR than the sequential protocol, irrespective of which of these two flip-angle formulas are used. The signal discontinuities in the time-encoded CFA data can be seen in vivo, but these discontinuities are removed by the $VFA_{maintain}$ formula.

The relative performance of the in vivo SNRs are similar to the simulations, but with some minor differences. This is probably due to the differences between the comparisons. The simulations assume a single set of dispersion parameters and a uniform ATT distribution between 300 - 600 ms, whereas the in vivo data likely has a different distribution of ATTs and the level of dispersion likely correlates with the ATT (with dispersion increasing further down the vascular tree). Furthermore, it is the SNR that is plotted for the in vivo data, which accounts for the varying level of background noise, whereas the simulations only show the acquired ASL signal, so differences in the noise level across frames are not considered. It can then be expected that some of the predicted improvement in signal due to the use of higher flip-angles (due to time-encoding or variable flip-angles) is accompanied by an increase in the physiological noise level due to a higher static tissue signal.

6.6.2.3 High Resolution: $0.63 \times 0.63\ \mathrm{mm}^2$

The temporal mean of the high in-plane resolution data set, acquired to evaluate differences in small vessel visibility, for the sequential and time-encoded $\mathrm{VFA}_{\mathrm{maintain}}$ protocols is shown in Figure 6.13 for one subject. The temporal mean is shown for clarity, but the individual frames show similar trends to the low-resolution comparison. The data from the

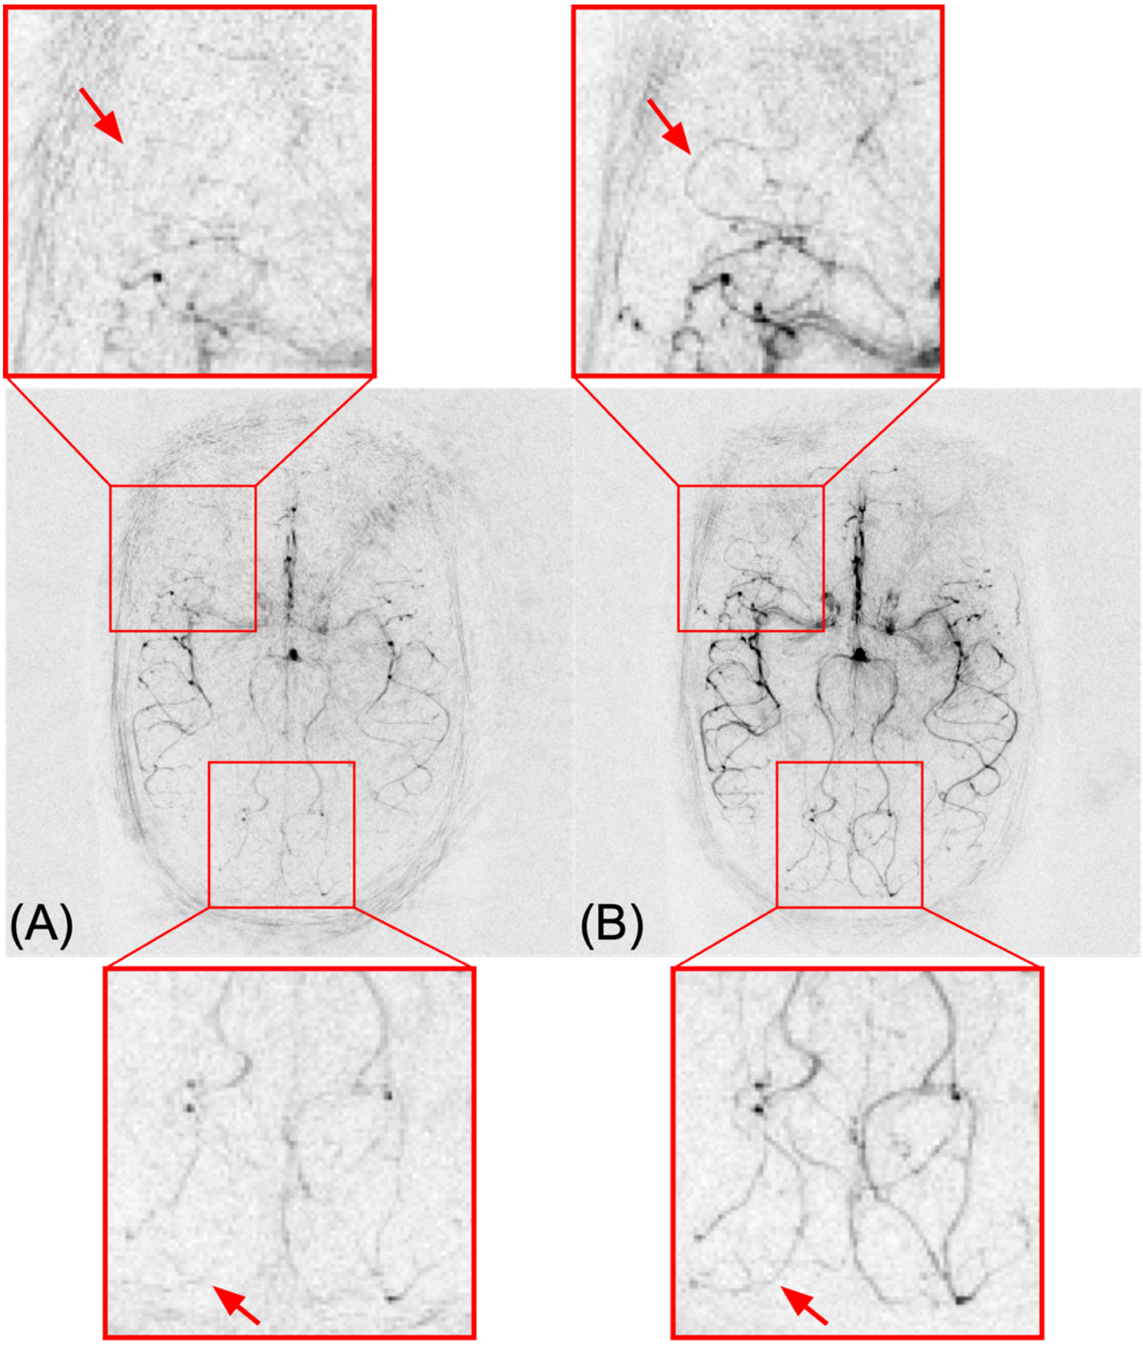

Figure 6.13: The temporal mean of the sequential (A) and time-encoded (B) $\mathit{VFA}_{\mathit{maintain}}$ high resolution data. Two regions of interest are shown which demonstrate the visual improvement of the very small distal vessels due to the higher SNR of the time-encoded $\mathit{VFA}_{\mathit{maintain}}$ data.

2 other subjects also qualitatively agree with this comparison. The SNR is theoretically approximately 2.5 times lower at this resolution than for the 1.15×1.15 mm^2 resolution data due to the smaller voxel size, larger bandwidth, and the slightly longer TE. However, the larger number of radial spokes acquired somewhat reduces the loss in SNR.

It is clear that there is higher overall SNR in the time-encoded data than the sequential data, with many vessels appearing very faint relative to the background noise in the sequential data. The magnified regions highlight two distal vessels that are very difficult to see in the sequential data, due to the low SNR but are much more clearly depicted in the time-encoded data. These results demonstrate that the improvement in SNR of the time-encoded $\mathrm{VFA}_{\mathrm{maintain}}$ protocol results in an improved depiction of small distal vessels.

6.6.2.4 Undersampled Comparison

The results of the comparison between the undersampled time-encoded $\mathrm{VFA}_{\mathrm{maintain}}$ protocol and 1-average of the sequential $\mathrm{VFA}_{\mathrm{maintain}}$ protocol are shown in Figure 6.14 where the temporal mean for both protocols and their fully-sampled (2-averages for the sequential protocol) counterparts are displayed for each subject.

Despite being angularly undersampled by almost a factor of 2, the undersampled time-encoded $\mathrm{VFA}_{\mathrm{maintain}}$ data looks remarkably similar to the fully sampled-data, with most of the vessels still clearly depicted, albeit with an increase in the estimated noise. The general increase in noise is due to the reduced number of *k*-space samples: a reduction by a factor of $\mathrm{R}_\theta = 300/156 = 1.92$, resulting in a theoretical factor of 1.39 increase in the noise SD. However, there is also an increase in the aliased energy, which for radial acquisitions is fairly incoherent and appears noise-like due to the dispersed nature of the aliasing.[7,28]

The quantitative SNR measurements for the undersampled time-encoded and 1-average sequential VFA$_{maintain}$ data are shown in Figure 6.15 demonstrating that despite under-sampling, the time-encoded protocol still outperforms the sequential protocol. The time-encoded SNR is likely to be conservative, and probably an underestimate, because

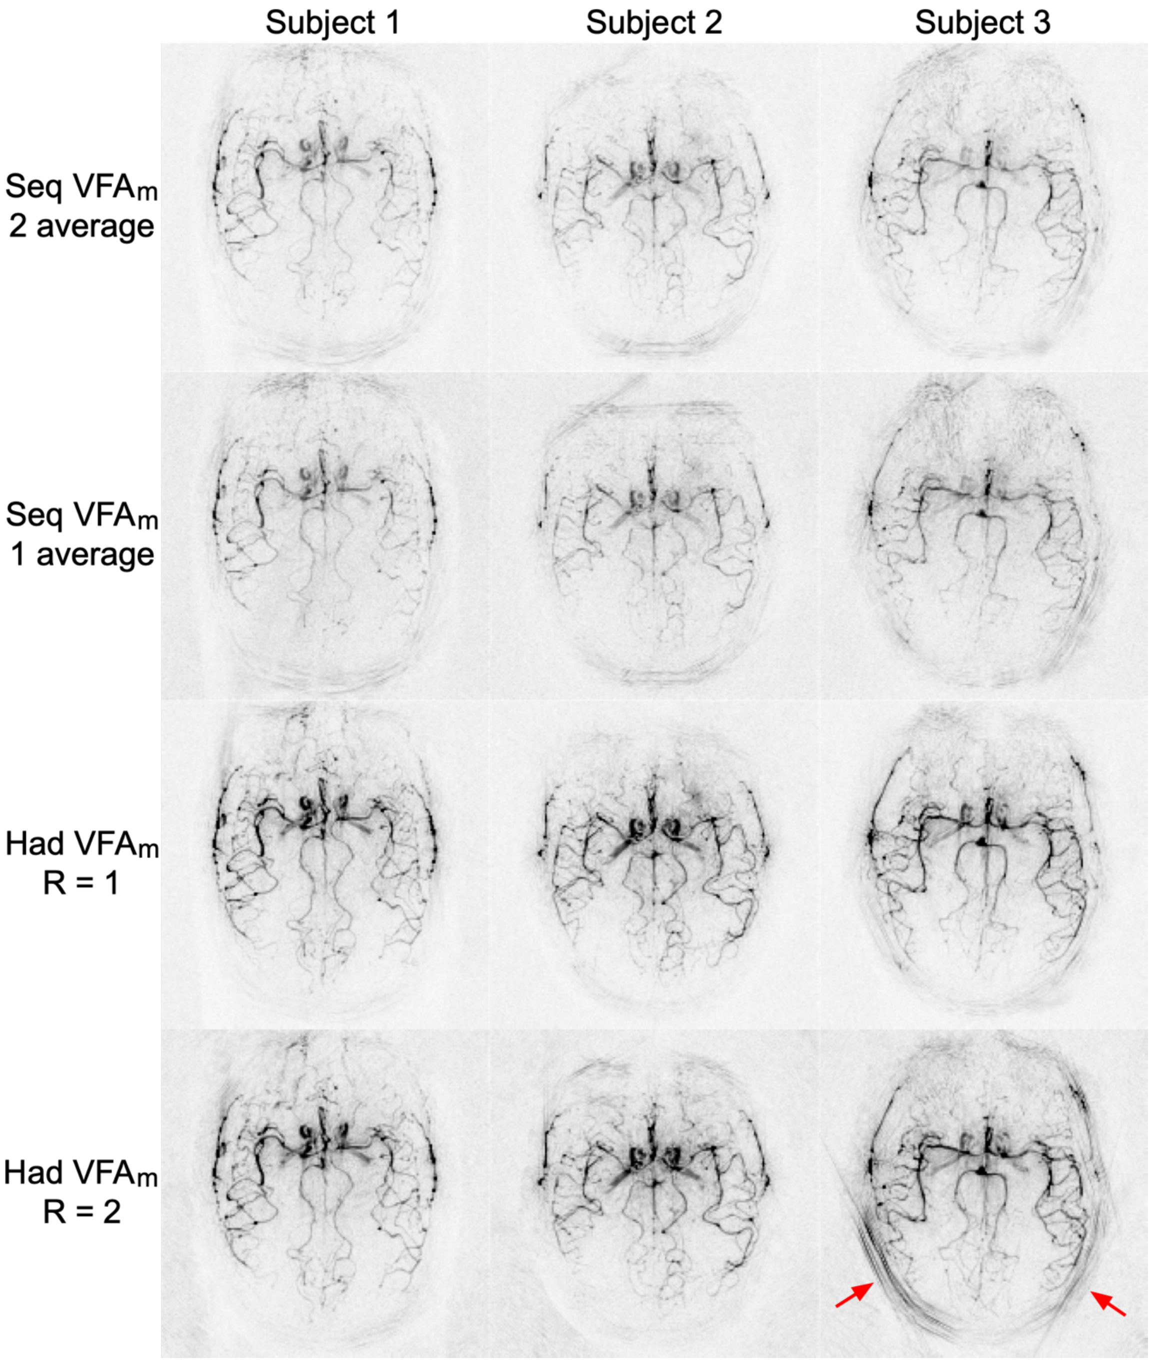

Figure 6.14: The temporal mean of the sequential (Seq) VFA$_{maintain}$ (VFA$_m$) data using 2 or 1 average and the fully-sampled (R=1) and undersampled (R=2) time-encoded (Had) VFA$_{maintain}$ data for 3 subjects. The arrows highlight static tissue artefacts possibly caused by motion.

the noise mask (two ROIs at the edge of the images) likely includes a larger amount of aliased energy than the voxels immediately surrounding the arteries, due to the reduced FOV effect.[28] The effect of the aliasing artefacts on the SNR can be seen in that the time-encoded data has an average 44% higher SNR than the sequential data, but for the fully-sampled protocol, the time-encoded data had an average 66% higher SNR than the 2-average sequential protocol. If just the noise estimates are considered, the 1-average sequential protocol had an average factor of 1.40 greater noise SD than the 2-average protocol, in agreement with the $\sqrt{2} = 1.41$ expected increase, while the undersampled time-encoded protocol had a factor of 1.48 greater noise, larger than the expected 1.39 increase.

The undersampled data for subject 3 has several streak artefacts at the back of the head which appear to be due to motion during a small portion of the acquisition. Despite this, the majority of the arteries are still clearly visible due to the position of the artefacts.

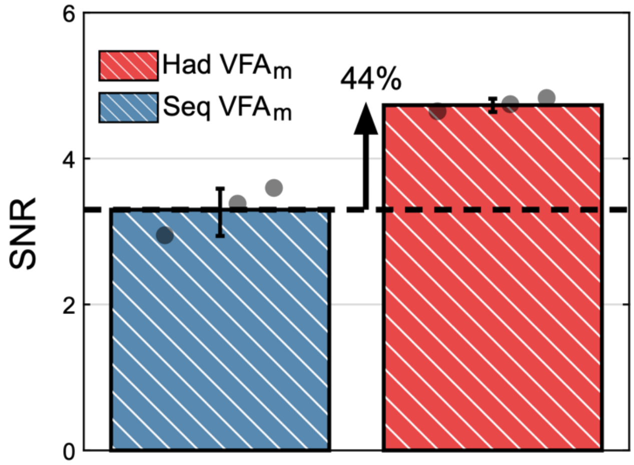

Figure 6.15: The average SNR measurements for the single-average sequential (Seq) VFA$_{maintain}$ (VFA$_m$) data and the undersampled time-encoded (Had) VFA$_{maintain}$ data for 3 subjects, using the same noise mask as described previously, but using a vessel mask generated from these short scan time data. Each subject's mean SNR is shown as a dot, while the bar graphs show the mean and range across subjects. The relative improvement in SNR of the time-encoded protocol compared to the sequential protocol is displayed.

6.7 Background Suppression Optimisation

As shown in Section 6.6.2, artefacts from the increasing static tissue signal across the readout frames contribute to an increasing level of physiological 'noise' in the image. This may hinder image interpretation, particularly for the time-encoded data where the background tissue signal level changes rapidly across frames reconstructed from different time-encoded blocks, rather than gradually increasing over the course of the readout as it does for the sequential protocols.

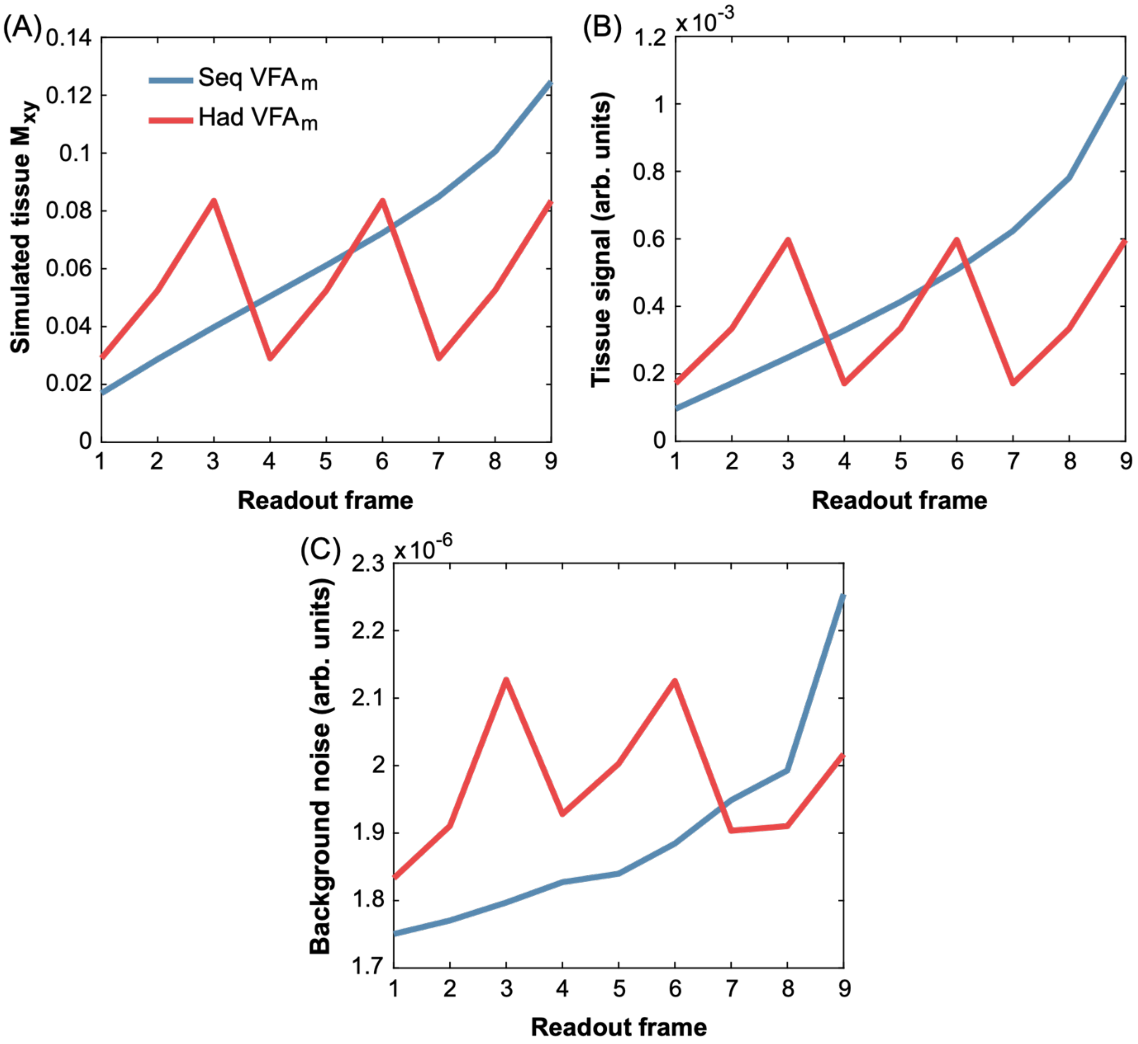

Figure 6.16: The static tissue signal and background noise for the fully-sampled low-resolution sequential VFA$_{maintain}$ (Seq VFA$_m$) and time-encoded VFA$_{maintain}$ (Had VFA$_m$) protocols. (A) The simulated tissue signal, (B) the mean in vivo tissue signal, and (C) the in vivo background noise. The simulated tissue signal was simulated for each excitation then averaged across each frame. The in vivo data is calculated from 1 subject.

Figure 6.16 demonstrates how the simulated and fully-sampled low-resolution in vivo tissue signal varies across readout frames for the sequential and time-encoded VFA$_{\text{maintain}}$ protocols and how this affects the background noise level. The simulations assumed $T_{1t} = 0.791$ s[18] and instantaneous pulses, as before, and accounts for the excitation flip-angles. The in vivo static tissue signal was calculated as the mean signal within the brain mask across all unsubtracted images for each frame (across label and control images or encodings). The RMS of the background noise was calculated within the previously used background ROIs from the difference images. The in vivo tissue signal can be seen to vary in a similar manner to the simulations. The background noise also varies in a similar manner to the tissue signal, though the noise in the time-encoded images is higher than expected relative to the sequential noise. This confirms that the magnitude of the tissue signal has an effect on the apparent noise in the angiograms.

In the in vivo comparison in Section 6.6.2, the timing of the BGS inversion pulses were chosen to null static tissue 100 ms before the first excitation pulse of the readout. This was implemented for the benefit of the perfusion sequences in Chapter 5 which used magnitude subtraction and so required a positive static tissue M_z to avoid loss of ASL signal. Since complex subtraction was used for the angiography data in this chapter, it is not required that the static tissue M_z is positive, since phase differences are accounted for in the subtraction. An alternative for the angiography sequence, therefore, is to move the null-time closer to the centre of the readout. The static tissue would then have negative M_z at the beginning of the readout, be nulled near the centre, and have positive M_z at the end of the readout. The absolute static signal would then be smaller on average across the readout.

6.7.1 Grid-Search Optimisation

A grid-search was performed for each of the CFA and $\mathrm{VFA}_{\mathrm{maintain}}$ protocols for the null-time which minimised the maximum simulated absolute static tissue signal. First, the range of null-times that satisfy the following conditions for each protocol were found: 1) the inversion pulses must occur within the labelling period, 2) the inversion pulses must not overlap with each other or the readout. The static tissue signal for each possible null-time was then simulated (assuming instantaneous pulses and accounting for the excitation flip-angles) and the null-time found which minimised the maximum absolute tissue signal. The grid-search was performed across the null-time range of -1000 ms to 1000 ms at 1 ms intervals, where a null-time of 0 corresponds to the centre of the first excitation pulse and a negative null-time corresponds to during the readout. The relevant settings used were: TR 10 ms, T_{1t} 0.791 s[18], $T_{1,opt}$ 700 ms.

Protocol	***Optimal Null-Time***
Sequential CFA	-87 ms
Sequential $\mathrm{VFA}_{\mathrm{maintain}}$	-87 ms
Time-encoded CFA	-149 ms
Time-encoded $\mathrm{VFA}_{\mathrm{maintain}}$	-217 ms

Table 6.3: The optimal null-times for each protocol which minimises the mean absolute static tissue signal across the readout. A null-time of 0 corresponds to the centre of the first excitation pulse and a negative null-time corresponds to the static tissue being nulled during the readout.

The optimal null-times are given in Table 6.3. These null-times were robust to the value of T_{1t} used, with only the time-encoded CFA null-time changing in the T_{1t} range of 0.5 - 3 s (null-time = -137 ms and -166 ms for each end of this range). For all of the protocols, negative null-times were optimal, ensuring that the longitudinal magnetisation of the static tissue passed through zero during the readout.

For both flip-angle formulas, the time-encoded protocol was able to use longer (negative) null-times due to the longer PCASL preparation period, but the $VFA_{maintain}$ protocol did not reach its ideal null-time of -237 ms due to timing constraints. Due to the short preparation period of the sequential protocol, a null-time of only -87 ms could be used for both flip-angle formulas (the longest negative null-time possible with this protocol), which is far from the ideal null-times of -338 ms and -613 ms for the CFA and $VFA_{maintain}$ protocols, respectively, when no restrictions are placed on the inversion pulse timings.

The ideal null-times were not at the centre of the readouts (-540 ms and -180 ms for the sequential and time-encoded protocols, respectively) due to the contributory effects of the recovering tissue signal and the flip-angles used.

The simulated static tissue signal for the sequential $VFA_{maintain}$ and time-encoded $VFA_{maintain}$ protocols with the original null-time of 100 ms and the optimal null-times are

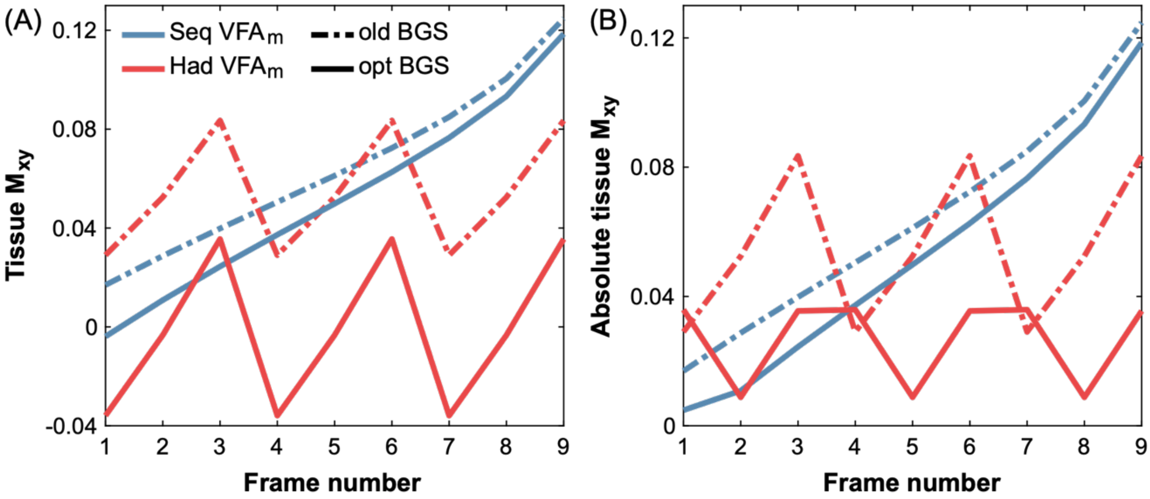

Figure 6.17: The simulated static tissue signal ($T_{1t} = 0.791$ *s) for the sequential* $VFA_{maintain}$ *(Seq* VFA_m*) and time-encoded* $VFA_{maintain}$ *(Had* VFA_m*) protocols with two different BGS null-times: (1) the original 100 ms (dashed lines) and (2) the optimised null-times (solid lines). (A) shows the unmodulated static tissue signal while (B) shows the absolute static tissue signal to demonstrate the evolution of the static tissue signal and its magnitude. In both cases, the mean signal across each frame is shown.*

shown in Figure 6.17. The optimised null-times improved the theoretical BGS performance for all of the protocols, but the largest improvement was for the time-encoded $\text{VFA}_{\text{maintain}}$ protocol which had a mean decrease in the absolute tissue signal of 55%, compared to a 19% decrease for the sequential $\text{VFA}_{\text{maintain}}$ protocol. For the time-encoded protocol, the 2nd, 5th, and 8th frames experienced the largest decrease (85%) followed by the 3rd, 6th, and 9th frames which had a 60% decrease in absolute static tissue signal.

6.7.2 Methods

Two of the volunteers included in Section 6.6 were scanned with the low-resolution sequential and time-encoded $\text{VFA}_{\text{maintain}}$ protocols during the same scan session as the previously presented data using the optimised null-times given in Table 6.3. Data for the CFA protocols with optimised null-times were not acquired because we wanted to focus on the $\text{VFA}_{\text{maintain}}$ approach with its superior SNR. All other acquisition parameters, reconstruction pipeline, and analysis methods were the same as previously described in Section 6.6.1.

6.7.3 Results

Substantial motion was seen in both the sequential and time-encoded optimised null-time data for 1 volunteer. This meant that the arteries in each scan were misaligned and severe artefacts were present which made any qualitative or quantitative comparison difficult. Therefore, only a single subject was included in this preliminary comparison.

Figure 6.18 shows the in vivo absolute tissue signal from the unsubtracted images and the background noise from the angiograms across frames for a single subject, calculated in the same way as for Figure 6.16. The tissue signal varied in a very similar manner to that seen in the simulations in Figure 6.17(B), with the optimised null-times for each protocol leading to an overall reduction in tissue signal. The background noise did

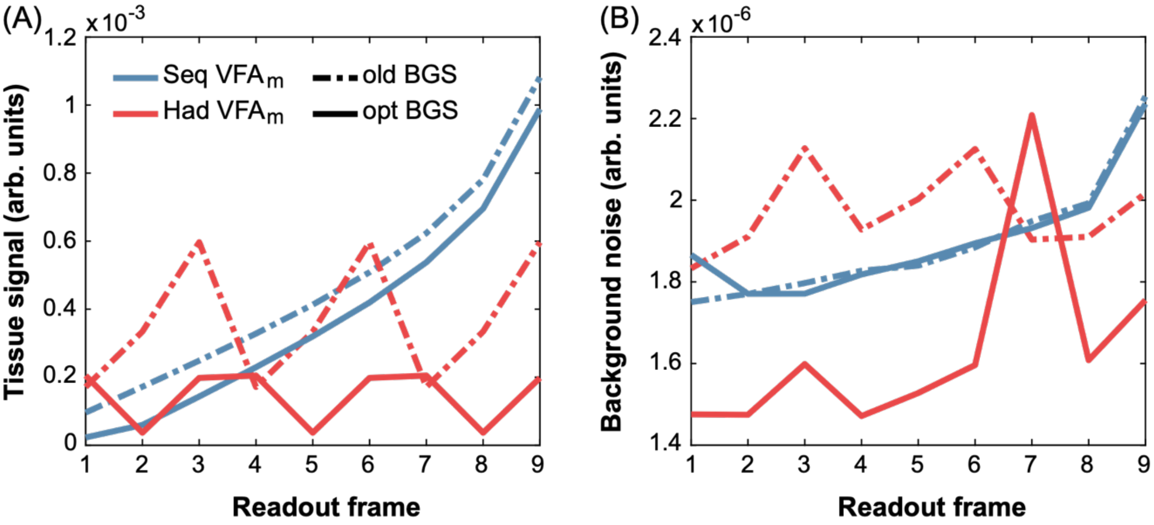

Figure 6.18: The in vivo tissue signal and background noise with (solid) and without (dashed) the optimised BGS null-time. (A) The mean tissue signal from the unsubtracted images with the brain mask and (B) the background noise from the difference images shown across frames for the sequential (Seq) and time-encoded (Had) VFA$_{maintain}$ (VFA$_m$) protocols. The original null-time was 100 ms (dashed) and the optimised null-times are given in Table 6.3.

not decrease for the sequential VFA$_{\text{maintain}}$ protocol when using the optimised null-time, but there was a large decrease in the noise in the time-encoded VFA$_{\text{maintain}}$ data, except for the 7th frame.

Figure 6.19 shows all 9 frames of the time-encoded data for the original and optimised BGS null-times. There was a good general agreement between both data sets, but some differences are highlighted. It can be seen that signal from around the skull and near the eyes was better suppressed in frames 2, 3, 5, 6, 8, and 9 of the optimised BGS protocol, which correspond to the last two readout frames which have improved BGS. The sequential data is not shown but both BGS cases were qualitatively similar.

Reconstructions with and without the phase-correction step resulted in a much smaller difference for the optimised null-time time-encoded data than for original null-time data, suggesting that the reduction in tissue signal reduced B_0 drift related subtraction errors (results not shown).

In frames 4, 5, and 6 of Figure 6.19, there was a loss of signal in the distal arteries

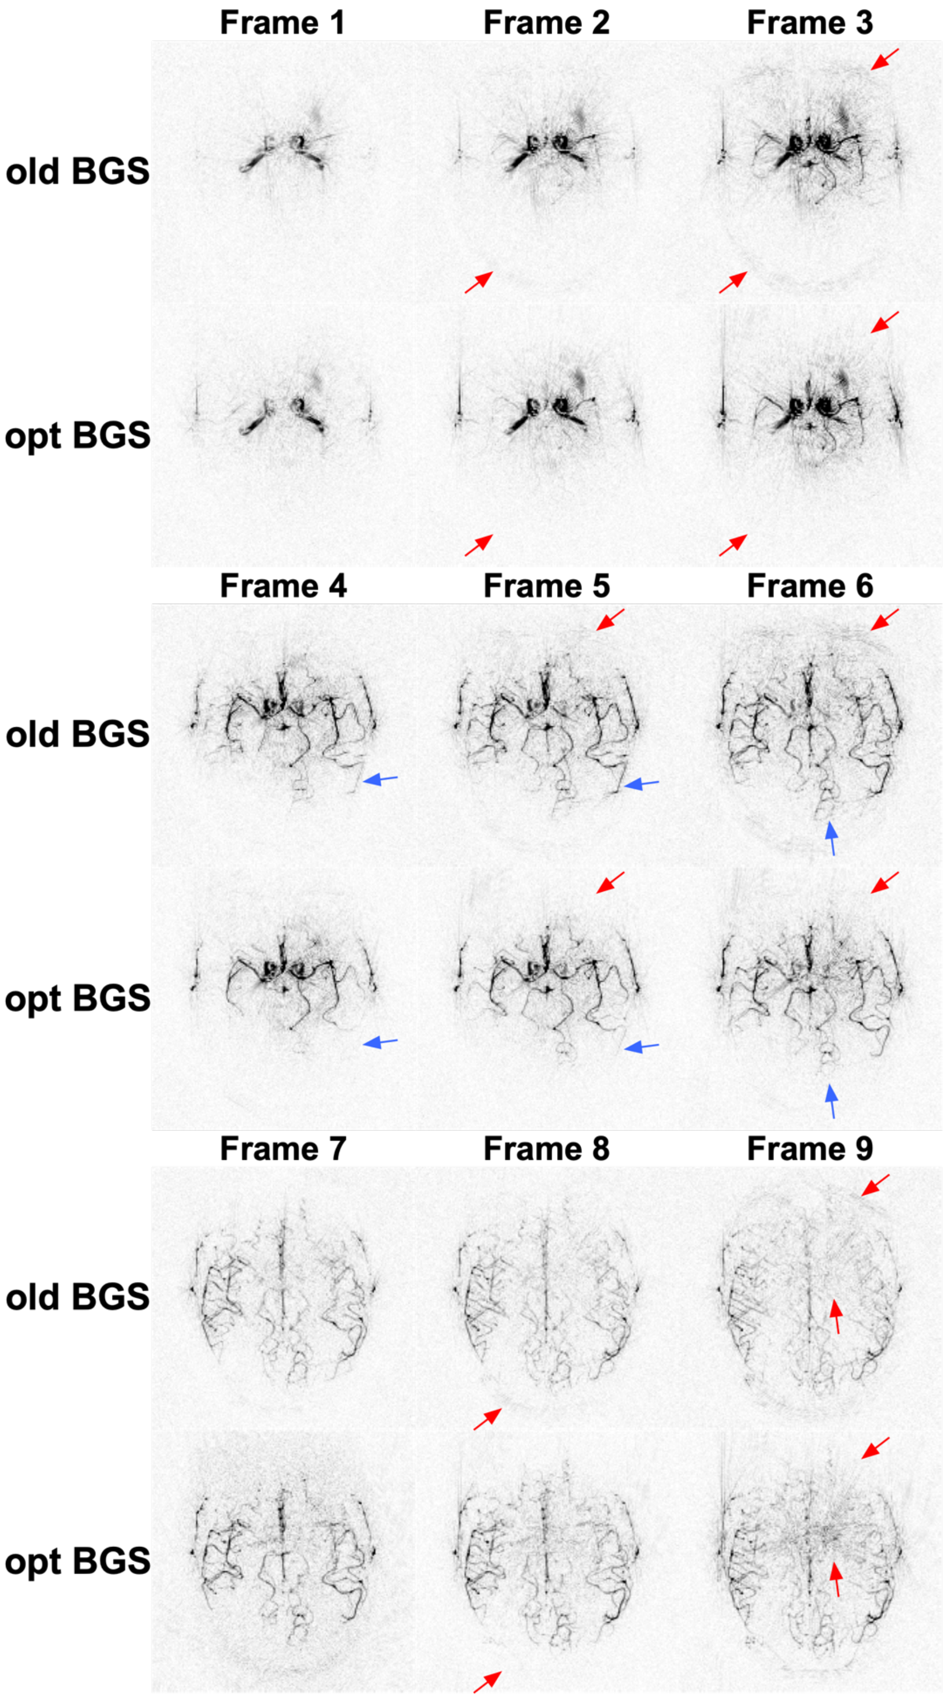

Figure 6.19: The individual frames for the time-encoded VFA$_{maintain}$ *protocol with the old BGS null-time and the optimised BGS null-time. Variations in tissue signal related 'noise' are highlighted in red while ASL signal differences are highlighted in blue.*

in the optimised BGS protocol. This may be due to physiological changes across the two scans, but it could also be due to the different inversion pulse timings. In the original protocol, the inversion pulses were centred at 317 ms and 793 ms after the start of labelling, which are near the end of the 1st encoded block and 73 ms after the start of the 3rd encoded block (see Figure 6.2). For the optimised null-time protocol, the inversion pulses were centred at 453 ms and 1070 ms after the start of labelling, which is 93 ms after the start of the 2nd encoded block and right at the end of the last encoded block. Therefore, it is possible that the 2nd encoded block in the optimised null-time data, which corresponds to frames 4, 5, and 6, had a loss of labelled signal at the front of the bolus, compared to the original protocol, due to saturation of existing labelled blood by the imperfect inversion slice profile as well as a loss of label generation time. However, a corresponding loss of signal in the 3rd encoded LD of the original protocol (frames 7, 8, and 9) is not as clearly seen.

The increase in the measured noise in the 7th frame of the optimised null-time protocol appears to be a global increase in noise, rather than due to motion artefacts.

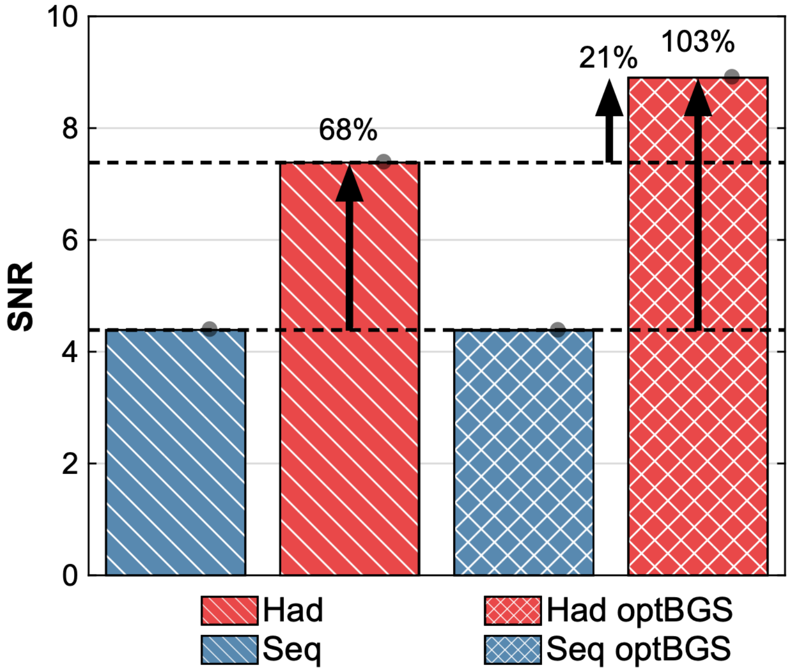

Figure 6.20: The mean SNR for the sequential (Seq) and time-encoded (Had) VFA$_{maintain}$ protocols with the original BGS null-time of 100 ms and with the optimised null-times. The relative improvement of the protocols is shown above each bar graph.

Additionally, an increase in noise in the centre of the brain in the 9th frame of the optimised null-time protocol is highlighted. The source of both occurrences is unclear.

The quantitative SNR comparison is shown in Figure 6.20. The ASL signal and noise were calculated as before with a vessel mask containing only the protocols included in this comparison. There was no clear difference in the SNR between the sequential protocols. Similar to the original BGS null-time comparison with 3 subjects, the time-encoded protocol improved the SNR by 68% in this subject. However, the optimised BGS time-encoded protocol increased the SNR by 103% compared to both sequential protocols and by 21% compared to the original null-time time-encoded protocol.

6.8 Discussion

In this chapter, it was demonstrated that the use of a time-encoded PCASL preparation could greatly increase the SNR in dynamic PCASL angiography. The source of this SNR improvement came from the use of larger flip-angles enabled by the reduction in the number of excitations during the readout. A VFA scheme which maintains a constant signal across the readout was developed and shown to eliminate the signal time-course discontinuities created by the decoding process of time-encoded data as well as maximise the ASL signal across the readout.

By angularly under-sampling the time-encoded data by almost a factor of 2, the scan time was reduced to a similar duration to that of a single average of the sequential protocol. Despite this large reduction in *k*-space sampling, most of the arteries were still clearly visible and the majority of the SNR benefit over conventional PCASL angiography was maintained.

Finally, preliminary results suggest that optimisation of the BGS null-time can lead to improved static tissue suppression which can reduce the physiological noise present in the angiograms. This is of particular benefit to protocols with short readout durations, which can be achieved with a time-encoded preparation because it is possible to for the null-time to occur relatively further into the readout time.

6.8.1 Label Duration

A short LD of 360 ms was chosen for both sequential and time-encoded protocols in order to decode temporally evenly spaced images from the time-encoded protocol which were matched with the timings of the sequential protocol. This was done so that the haemodynamic information in each protocol was similar, making it simpler to evaluate the benefits of using a time-encoded preparation with larger flip-angles during the readout.

A further benefit of using a short LD with a PCASL preparation is that the inflow of blood can be visualised. If a long LD were used, only the outflow could be visualised. However, it has previously been shown that "inflow-subtraction" can be used to visualise the inflow of blood, where each frame is subtracted from the first frame.[14] This subtraction results in an increase in the noise in the images, though, and assumes that all of the vessels are filled with labelled blood by the start of the readout.

Another downside of using longer LDs is that the scan time would also be increased: for example, a LD of 1 s would increase the low-resolution sequential protocol scan time from 2:30 min to 3:34 min and the high-resolution scan time from 4:35 min to 6:33 min. However, use of a long LD could improve the depiction of distal vessels, because the labelled bolus may have filled these vessels before the ASL tracer starts to be attenuated by the readout pulses. Furthermore, with moderate dispersion, the peak tracer signal can be reduced when using a short LD due to the temporal-smearing of the bolus (see Figure 6.7), but a long LD would reach a maximum plateau, improving SNR.[29]

An alternative strategy would be to use PASL labelling for angiography.[30] Since the labelling period is extremely short (on the order of 10 ms) imaging can begin almost immediately and visualise the entire inflow of blood. Another benefit of PASL methods is that the labelling efficiency is typically higher than for PCASL (~98%[31] rather than ~85%[12], respectively). However, because all of the label is created at the same time with PASL methods, by the time the distal end of the label arrives in the imaging volume the mean effective efficiency across the entire bolus is actually likely to be ~78% (assuming LD = 800 ms[32] and T_{1b} = 1.65 s[33]), although depiction of distal vessels will still benefit from the higher labelling efficiency at the proximal end of the labelled bolus. Nonetheless, as for the sequential PCASL protocol, all of the image frames would need to be acquired from the readout, meaning that low flip-angles would be necessary, which ultimately limits SNR.

6.8.2 Flip-angle scheme

The $\text{VFA}_{\text{maintain}}$ protocol was found to increase the ASL signal compared to the CFA scheme in vivo (and the $\text{VFA}_{\text{quadratic}}$ scheme through simulations) by making more effective use of the available magnetisation. Several downsides of this flip-angle scheme include higher SAR and increased tissue subtraction artefacts (discussed below).

SAR was reduced by lowering the TBWP of the excitation pulses and reducing the maximum flip-angle from 90° to 30°. Since ASL is a subtractive technique, any extra static tissue signal excited above or below the nominal slab by the low TBWP pulses should be subtracted out, and because a 2D readout was used, extra ASL signal will simply appear in the summed-intensity-projection. If a 3D radial readout were used, the standard 2 × oversampling along each spoke would reduce aliasing of signal into the FOV, though some artefacts would remain if angular undersampling were used. Alternatively, a large isotropic FOV could be used in conjunction with a head-only receive coil such that selective excitation becomes unnecessary and hard excitation pulses could be used, greatly reducing SAR.

For each flip-angle scheme, the minimum signal across the readout was maximised. Because the ASL tracer will have travelled further downstream by the end of the readout, an alternative cost function could be used to up-weight the signal towards the end of the readout, thus producing an improved depiction of distal vessels (for instance, the $VFA_{quadratic}$ scheme in Figure 6.4 with $FA_1 = 3°$ and $FA_2 = 11.4°$). Since the labelled blood is concentrated in the large vessels during the early frames, and so SNR is high, using reduced flip-angles for these frames may not be detrimental for vessel visualisation. However, this would result in a gradual brightening of the tracer within a given voxel across frames for the sequential protocol and would cause the tracer signal to oscillate for the time-encoded protocol.

The flip-angle optimisations assumed that all of the ASL bolus experienced each RF excitation pulse during the readout. However, this is unlikely to have happened with the experimental setup used in this work. For example, for one of the subjects, there was a 53 mm gap between the bottom of the image volume and the centre of the labelling plane. Assuming a constant non-pulsatile blood velocity of 20 cm/s and straight vessels which are perpendicular to the labelling plane and imaging volume, it would take 265 ms for the labelled blood to reach the image volume, much longer than the 2 ms gap between the end of the PCASL pulse train and the start of the readout in this work. This could introduce further signal variations to the signal time-course. This effect could be taken into account in the flip-angle optimisations or corrected for in post-processing or model-fitting.[13] However, in other experimental setups, this assumption would be better met: for example, if the labelling plane were closer to the imaging volume, or if a larger imaging volume was used (such as might be used with a whole-brain 3D acquisition) which would reduce the gap.

6.8.3 SNR Quantification

The vessel masks were robustly created using an automated thresholding method. However, the small distal vessels were often not included in the masks due to their lower signal and small size. An alternative SNR quantification method was tested where small hand-drawn ROIs were placed at different points along the vascular tree. Similar results to the comparison presented in Section 6.6.2.2 were seen even in the distal vessels (not shown).

The background noise mask consisted of two ROIs on either side of the images. The noise calculated with these ROIs will incorporate both thermal noise as well as signal originating from the subject which appears noise-like in nature. This is because signal will be spread out from each point due to: the PSF, trajectory errors during the acquisition, the gridding process in the reconstruction, and flow artefacts. When the radial acquisition is angularly undersampled, these background noise ROIs will also include increased aliased energy resulting from undersampling. However, this aliased signal may not greatly affect visualisation of the arteries because a reduced circular region with a diameter of FOV/2 around each signal source will not contain aliased signal.[28] If most of the arteries with significant signal energy are within this distance of each other (e.g. the proximal arteries around the circle of Willis), then the undersampling artefacts will not have a big impact on any of the other vessels. The SNR calculation for the undersampled time-encoded data, therefore, will likely be an underestimate of the effective SNR in the local neighbourhood of the vessels.

6.8.4 Undersampling

The image resolution for radial readouts is approximately independent of the number of radial spokes but is instead dependent on the extent of each spoke in *k*-space (i.e. k_{max}).[7] The image resolution, therefore, does not suffer significantly when the number of projections are reduced. In contrast, the image resolution is reduced in Cartesian imaging when the number of phase encode steps are reduced (assuming the same FOV).

The undersampling artefacts are also diffuse and incoherent in nature, rather than discrete ghosts as for Cartesian imaging. Angiograms are inherently sparse, which makes undersampled radial trajectories well-suited for angiography due to the reduced potential for signal aliasing. The natural sparseness of angiograms in image space also makes them suitable for compressed sensing reconstructions,[34] which can improve image quality considerably under high levels of acceleration. Extension to a 3D readout will allow even greater rates of acceleration due to the increased natural sparseness of the angiograms.

6.8.5 BGS Null-Time Optimisation

The null-time for each protocol was optimised to minimise the maximum absolute static tissue signal acquired during the readout. An in vivo demonstration of the benefits of this optimisation was shown in only one subject, but reductions in the background noise could be qualitatively and quantitatively appreciated. This approach can somewhat mitigate the increased static tissue signal resulting from using larger flip-angles and the $\text{VFA}_{\text{maintain}}$ scheme.

There appeared to be a loss of distal tracer signal in the middle 3 frames in the optimised null-time time-encoded data compared to the original null-time data. The cause of this was hypothesised to be either of physiological origin, or a loss of tracer magnetisation due to saturation from the imperfect slice profile of the FOCI pulses as well as the loss of time spent labelling (a total loss of 20.48 ms for the 10.24 ms pulse and padding time to maintain a 50% RF duty cycle). The latter case is a general downside of the interleaved inversion pulses. However, scans on a greater number of subjects are required to confirm whether this is a systematic issue with the optimised BGS approach.

For long LDs, such as those typically used in perfusion imaging, a LD loss of 20 ms will proportionally have a small effect on the total label generated, a result found by Dai et al.[35] For the short LD used in this chapter, however, the loss of tracer will have a greater impact, especially for the sequential protocol in which both inversion pulses occur during the single LD. To reduce this loss, the slice profile of the pulses could be further improved and the duration of the pulses could be reduced. Another option for the time-encoded protocol is for the inversion pulses to only occur during the switches between encoded boluses, reducing the loss of label through fewer switches between label and control conditions; though this would likely result in non-optimal timings and negatively impact the BGS.

The sources of the oscillations in noise in several frames of the in vivo comparison in Figure 6.19 are not clear. They could be caused by subtleties in the timings and implementation of the BGS inversion pulses, but it could also be caused by unpredictable events such as instabilities in the scanner hardware. A larger number of subjects are therefore required to determine if the artefacts are consistent.

A more advanced method for minimising the static tissue signal during the readout would be to place one or more inversion pulses between readout frames, leading to the tissue M_z to pass through 0 more than once. This would be particularly beneficial for long readout durations where the tissue signal varies considerably and the null-time cannot reach its optimal value. However, this would also lead to varying levels of ASL signal loss across the readout frames due to imperfect inversion and different frames being acquired after a different number of inversion pulses.

6.8.6 Time-Encoding

A time-encoded PCASL preparation is an attractive alternative to standard PCASL labelling for dynamic angiography due to the reduced reliance on the readout for generating temporally resolved images, which enables large flip-angles to be used. Alternatively, the number of readout frames and flip-angles can be kept the same while the temporal resolution is increased.

Rather than using a 4 × 3 Hadamard-encoding matrix with 3 acquisition frames, an 8 × 7 encoding matrix could be used with only 1 readout frame, generating all of the temporal information from the time-encoded preparation, enabling the use of larger flip-angles during the readout. This method was tested in vivo using a constant LD of 160 ms, a single frame with a duration of 120 ms and the $\mathrm{VFA}_{\mathrm{maintain}}$ scheme. This gave effective PLDs of 62 ms, 222 ms, 382 ms, 542 ms, 702 ms, 862 ms, and 1022 ms (matching the longest PLD of the previously used protocols) and flip-angles that ramped up across the range 15° - 30°. The remaining sequence settings and reconstruction process were identical to those described for the low-resolution acquisitions in Section 6.6.1, except that 8 encodings were acquired, resulting in a scan time of 4:19 min. The 7 decoded frames from a single subject are shown in Figure 6.21. Although there is good visibility of the ASL tracer traversing the vessels, there is a large artefact emanating from the edge of the head.

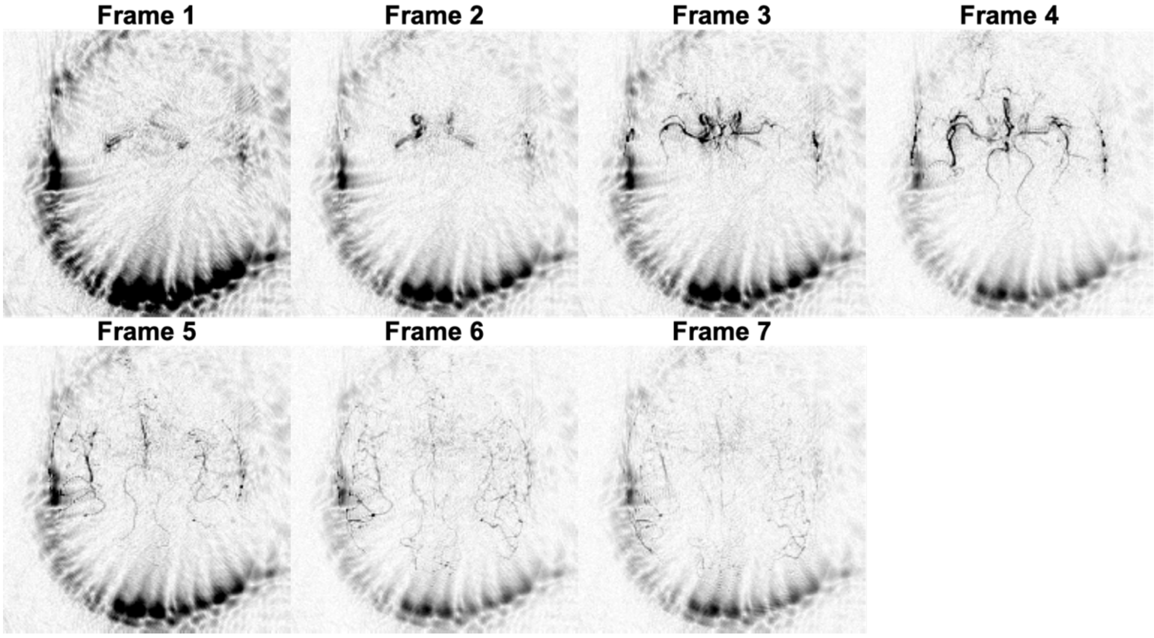

Figure 6.21: The 7 frames from the in vivo 8×7 *time-encoded acquisition in a single subject.*

This artefact is not visibly noticeable in the unsubtracted images due to its relatively small signal magnitude. Extensive sequence testing did not isolate the source of this artefact, but did suggest that it may be related to the use of large, and quickly varying, flip-angles during the readout. Epstein et al.[36] demonstrated that for an SPGR readout using a quadratic RF-spoiling scheme, constant flip-angles of 20° and above led to visible artefacts due to large signal oscillations during the approach to steady-state. These signal oscillations are likely made much worse by the quickly varying VFAs used with the time-encoded protocol in this work. Further investigation using extended phase graph (EPG)[37] simulations are required to confirm this as the source of the artefact. Options for mitigating this artefact include optimising the RF phase increment,[36] using an optimised set of VFAs to control the signal evolution during the approach to steady-state,[38] and use of more advanced gradient spoiling schemes.[39]

It is not a requirement to use Hadamard-encodings, but they are the most efficient in terms of decoding the largest number of images, N, from the fewest encoded images, N+1, with an (N+1)/2 signal averaging benefit relative to conventional PCASL. Another example of an encoding matrix that could be used is:

$$A = \begin{bmatrix} -1 & -1 \\ -1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix}, \tag{6.8}$$

generating 2 decoded timepoints. This encoding matrix could be of benefit due to the longer LDs that could be used, but it is not as efficient in terms of increasing the temporal resolution.

In this work, the encoded LDs were all of equal duration, but each block can use a different label duration. The first LD can be made longer without affecting the PLDs of the following blocks, better filling the downstream arteries for the longest PLDs, but if the

LDs of the following blocks are variably set, then the temporal resolution will no longer be uniformly spaced.

As mentioned in Chapter 5, a downside of using time-encoded preparations is the increased sensitivity to motion. If an image frame is corrupted by motion, the resulting artefacts will propagate to all of the decoded images from that frame. However, methods to improve motion robustness could mitigate this issue. Possible navigator-based motion correction methods include prospective motion correction[40–42] and retrospective correction of data before reconstruction.[43] As an alternative to using navigators, which typically extend scan time, the radial readout itself could be used to retrospectively perform motion correction, in a process called ‘self-navigation’. If the radial angle is incremented by the ‘golden-ratio angle’, 111.25°,[44] instead of by $180°/N$ (where N is the total number of radial spokes), then images can be reconstructed at arbitrary temporal resolution while maintaining a relatively even coverage of k-space.[45] This property can then be used to reconstruct high-temporal but low-spatial resolution images from the relatively densely sampled centre of k-space to use for motion correction of the k-space data before the final reconstruction.[46]

6.8.7 Readout

In this work, a 2D SPGR readout was used to keep the scan duration of each protocol short so that all of the data for each subject could be acquired within a single scan session. Future work will include the extension of time-encoded angiography to a whole-brain 3D SPGR radial readout, where the benefits of a time-encoded preparation are expected to be maintained.

Balanced SSFP (bSSFP) readouts[14,47] provide greater SNR than SPGR readouts because residual transverse magnetisation is not spoiled after each excitation, but is instead refocussed. However, bSSFP is sensitive to off-resonance fields, which are more common in whole-brain acquisitions because it is difficult to accurately shim the entire image volume at 3T. The presence of off-resonant spins leads to low-signal banding artefacts which can be particularly deleterious for ASL acquisitions because downstream ASL signal in later image frames can become partially saturated when a banding artefact crosses an upstream artery.[14] Due to these complications, and the desire to extend this work to use a whole-brain 3D acquisition, an SPGR readout was used.

6.8.8 Protocol Optimisation

As previously described, the LD, temporal resolution, number of readout frames, and level of acquisition segmentation were chosen subjectively in order to compare matched sequential and time-encoded protocols. As well as their use for qualitative assessment of arterial supply, dynamic angiograms can also be fit with a kinetic model to provide quantitative measurements of blood volume, transit times, and dispersion parameters.[13,48]

It would, therefore, be possible to extend the optimisation framework developed for ASL perfusion imaging protocols in Chapter 4 to optimise the labelling and acquisition parameters for ASL angiography protocols to more accurately estimate these physiological parameters. In a similar way to Chapter 5, sequential and time-encoded protocols could then be individually optimised with an objective framework and compared based on the accuracy with which they estimate these parameters. This may lead to a fairer comparison of each angiographic protocol, which was the aim of Chapter 5 for the perfusion protocols.

6.9 Conclusions

It has been demonstrated that a time-encoded preparation can be beneficial for dynamic PCASL angiography by allowing some of the temporal information to be generated from the PCASL preparation. This enabled the use of larger flip-angles during a shorter SPGR readout which led to an SNR improvement of 66%. Furthermore, a VFA formula,

$\text{VFA}_{\text{maintain}}$, was derived which led to a high constant ASL signal across the readout and removed the signal discontinuities present in the time-encoded data when a CFA scheme was used. The $\text{VFA}_{\text{maintain}}$ scheme led to an SNR increase of 105% relative to the sequential CFA protocol.

Preliminary data showed the potential of tailoring the BGS to maintain a low background tissue signal across readouts in order to reduce physiological noise in the angiograms as well as the potential of the time-encoded protocol to being angularly undersampled.

6.10 References

1 Wang SJ, Nishimura DG, Macovski A. Multiple-readout selective inversion recovery angiography. *Magn Reson Med* 1991; **17**: 244–251.

2 Schmitt P, Speier P, Bi X, Weale P, Mueller E. Non-contrast-enhanced 4D intracranial MR angiography: Optimizations using a variable flip angle approach. In: *Proceedings of the 18th Annual Meeting of the ISMRM, Stockholm, Sweden*. 2010, p 402.

3 Suzuki Y, Teeuwisse WM, Schmid S, Helle M, Cauteren M Van, Osch MJP Van. Improving 4D pCASL angiography by combining Hadamard time-encoding with Look-Locker readout. In: *Proceedings of the 22nd Annual Meeting of the ISMRM, Milan, Italy*. 2014, p 6460.

4 van der Plas MCE, Schmid S, Versluis M, van Osch MJP. Time-encoded golden angle radial ASL. In: *Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France*. 2018, p 306.

5 Dixon WT, Du LN, Faul DD, Gado M, Rossnick S. Projection angiograms of blood labeled by adiabatic fast passage. *Magn Reson Med* 1986; **3**: 454–462.

6 Nishimura DG, Macovski A, Pauly JM, Conolly SM. MR angiography by selective inversion recovery. *Magn Reson Med* 1987; **4**: 193–202.

7 Peters DC, Korosec FR, Grist TM, Block WF, Holden JE, Vigen KK *et al.* Undersampled projection reconstruction applied to MR angiography. *Magn Reson Med* 2000; **43**: 91–101.

8 Kamath S. Observations on the length and diameter of vessels forming the circle of Willis. *J Anat* 1981; **133**: 419–23.

9 Günther M, Warmuth C, Zimmer C. Sub-millimeter Dynamic Spin Labeling Cerebral 2D-Angiography with 40ms Temporal Resolution. In: *Proceedings of the 10th Annual Meeting of the ISMRM, Honolulu, USA*. 2002, p 1100.

10 Ogg RJ, Kingsley PB, Taylor JS. WET, a T1- and B1-insensitive water-suppression method for in vivo localized 1H NMR spectroscopy. *J Magn Reson B* 1994; **104**: 1–10.

11 Golay X, Petersen ET, Hui F. Pulsed Star Labeling of Arterial Regions (PULSAR): A robust regional perfusion technique for high field imaging. *Magn Reson Med*

2005; **53**: 15–21.

12 Dai W, Garcia D, De Bazelaire C, Alsop DC. Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magn Reson Med* 2008; **60**: 1488–1497.

13 Okell TW, Chappell M a., Schulz UG, Jezzard P. A kinetic model for vessel-encoded dynamic angiography with arterial spin labeling. *Magn Reson Med* 2012; **68**: 969–979.

14 Okell TW, Schmitt P, Bi X, Chappell MA, Tijssen RHN, Sheerin F *et al.* Optimization of 4D vessel-selective arterial spin labeling angiography using balanced steady-state free precession and vessel-encoding. *NMR Biomed* 2016. doi:10.1002/nbm.3515.

15 Li W, Xu F, Schär M, Liu J, Shin T, Zhao Y *et al.* Whole-brain arteriography and venography: Using improved velocity-selective saturation pulse trains. *Magn Reson Med* 2018; **79**: 2014–2023.

16 Tan ET, Huston J, Campeau NG, Riederer SJ. Fast inversion recovery magnetic resonance angiography of the intracranial arteries. *Magn Reson Med* 2010; **63**: 1648–1658.

17 Zhao L, Mulkern R, Tseng CH, Williamson D, Patz S, Kraft R *et al.* Gradient-echo imaging considerations for hyperpolarized 129Xe MR. *J Magn Reson B* 1996; **113**: 179–83.

18 Lin C, Bernstein M, Huston J, Fain S. Measurements of T1 Relaxation times at 3.0T: Implications for clinical MRA. In: *Proceedings of the 9th Annual Meeting of the ISMRM, Glasgow, Scotland.* 2001, p 1391.

19 Zur Y, Wood ML, Neuringer LJ. Spoiling of transverse magnetization in steady-state sequences. *Magn Reson Med* 1991; **21**: 251–263.

20 Ordidge RJ, Wylezinska M, Hugg JW, Butterworth E, Franconi F. Frequency offset corrected inversion (FOCI) pulses for use in localized spectroscopy. *Magn Reson Med* 1996; **36**: 562–566.

21 Payne GS, Leach MO. Implementation and evaluation of frequency offset corrected inversion (FOCI) pulses on a clinical MR system. *Magn Reson Med* 1997; **38**: 828–833.

22 Fessler JA. Michigan Image Reconstruction Toolbox. https://web.eecs.umich.edu/~fessler/code/ (accessed 26 Feb2018).

23 Fessler JA, Sutton BP. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Trans Signal Process* 2003; **51**: 560–574.

24 Walsh DO, Gmitro AF, Marcellin MW. Adaptive reconstruction of phased array MR imagery. *Magn Reson Med* 2000; **43**: 682–690.

25 Pipe JG, Menon P. Sampling density compensation in MRI: Rationale and an iterative numerical solution. *Magn Reson Med* 1999; **41**: 179–186.

26 Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. *Magn Reson Med* 1990; **16**: 192–225.

27 Otsu N. A Threshold Selection Method from Gray-Level Histograms. *IEEE Trans Syst Man Cybern* 1979; **9**: 62–66.

28 Scheffler K, Hennig J. Reduced circular field-of-view imaging. *Magn Reson Med* 1998; **40**: 474–480.

29 Okell TW, Chappell M a., Woolrich MW, Günther M, Feinberg D a., Jezzard P. Vessel-encoded dynamic magnetic resonance angiography using arterial spin labeling. *Magn Reson Med* 2010; **64**: 698–706.

30 Suzuki Y, Fujima N, Ogino T, Meakin JA, Suwa A, Sugimori H *et al.* Acceleration of ASL-based time-resolved MR angiography by acquisition of control and labeled

images in the same shot (ACTRESS). *Magn Reson Med* 2018; **79**: 224–233.

31 Wong EC, Buxton RB, Frank LR. A theoretical and experimental comparison of continuous and pulsed arterial spin labeling techniques for quantitative perfusion imaging. *Magn Reson Med* 1998; **40**: 348–55.

32 Alsop DC, Detre JA, Golay X, Günther M, Hendrikse J, Hernandez-Garcia L *et al.* Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn Reson Med* 2015; **73**: 102–116.

33 Lu H, Clingman C, Golay X, van Zijl PCM. Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn Reson Med* 2004; **52**: 679–682.

34 Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med* 2007; **58**: 1182–1195.

35 Dai W, Robson PM, Shankaranarayanan A, Alsop DC. Benefits of Interleaved Continuous Labeling and Background Suppression. In: *Proceedings of the 18th Annual Meeting of the ISMRM, Stockholm, Sweden*. 2010, p 1748.

36 Epstein FH, Mugler JP, Brookeman JR. Spoiling of transverse magnetization in gradient-echo (GRE) imaging during the approach to steady state. *Magn Reson Med* 1996; **35**: 237–245.

37 Weigel M. Extended phase graphs: Dephasing, RF pulses, and echoes - pure and simple. *J Magn Reson Imaging* 2015; **41**: 266–295.

38 Mugler JP, Epstein FH, Brookeman JR. Shaping the Signal Response during the Approach to Steady State in Three-Dimensional Magnetization-Prepared Rapid Gradient-Echo Imaging Using Variable Flip Angles. *Magn Reson Med* 1992; **28**: 165–185.

39 Hess AT, Robson MD. Hexagonal gradient scheme with RF spoiling improves spoiling performance for high-flip-angle fast gradient echo imaging. *Magn Reson Med* 2017; **77**: 1231–1237.

40 Zun Z, Shankaranarayanan A, Zaharchuk G. Pseudocontinuous arterial spin labeling with prospective motion correction (PCASL-PROMO). *Magn Reson Med* 2014; **72**: 1049–1056.

41 Helle M, Koken P, Sénégas J. Improving Motion Robustness of Pseudo-Continuous Arterial Spin Labeling by using real-time Motion Correction. In: *Proceedings of the 23rd Annual Meeting of the ISMRM, Toronto, Canada*. 2015, p 270.

42 Frost R, Hess AT, Okell TW, Chappell MA, Tisdall MD, van der Kouwe AJW *et al.* Prospective motion correction and selective reacquisition using volumetric navigators for vessel-encoded arterial spin labeling dynamic angiography. *Magn Reson Med* 2016; **76**: 1420–1430.

43 Gallichan D, Marques JP, Gruetter R. Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magn Reson Med* 2016; **75**: 1030–1039.

44 Winkelmann S, Schaeffter T, Koehler T, Eggers H, Doessel O. An optimal radial profile order based on the golden ratio for time-resolved MRI. *IEEE Trans Med Imaging* 2007; **26**: 68–76.

45 Okell TW. Combined angiography and perfusion using radial imaging and arterial spin labeling. *Magn Reson Med* 2018. doi:10.1002/mrm.27366.

46 Graedel NN, McNab JA, Chiew M, Miller KL. Motion correction for functional MRI with three-dimensional hybrid radial-Cartesian EPI. *Magn Reson Med* 2017; **78**: 527–540.

47 Carr HY. Steady-State Free Precession in Nuclear Magnetic Resonance. *Phys Rev* 1958; **112**: 1693–1701.

48 Okell TW, Chappell MA, Jezzard P. A theoretical framework for quantifying blood volume flow rate from dynamic angiographic data and application to vessel-encoded arterial spin labeling MRI. *Med Image Anal* 2013; **17**: 1025–1036.

7 Summary and Future Work

In this final chapter, a summary of the work performed for this thesis is given along with a conclusion relating to the major findings. Following this, potential future research avenues, which build on the work in this thesis, are discussed.

7.1 Summary

The aim of this thesis was primarily to improve the time-efficiency of ASL based brain perfusion and angiography scans, due to the method's inherent low SNR. The approach taken within this thesis to address this problem was to improve the accuracy of fitted parameters, or to increase the image SNR, within a given scan time. The methods developed through this approach can equally be applied to the problem of maintaining a similar level of accuracy, or SNR, to an unoptimised protocol, but within a shorter scan time.

In Chapter 4, a flexible algorithm was developed for optimising the timings of multi-timepoint ASL perfusion protocols for a given scan-time in order to maximise the accuracy of either CBF and ATT or just CBF while being insensitive to variations in ATT. This algorithm used the Fisher information matrix to calculate the lower bound of the covariance matrix, in order to minimise the parameter variances. The effectiveness of this algorithm was tested on sequential multi-PLD protocols and demonstrated that there was good agreement between the predicted CBF and ATT variances (derived from the Cramér-Rao lower bound), the simulated errors (from Monte Carlo simulations), and the in vivo errors (with the use of ground truth parameter estimates). When compared to a scan time and label duration matched single-PLD protocol, the CBF optimised multi-PLD protocol produced more accurate CBF estimates, despite having fewer averages at each PLD.

In Chapter 5, this protocol optimisation algorithm was then applied to a wide range of advanced multi-TI protocol designs from the literature and several novel designs. This was done in order to fairly compare the range of protocols after each one had been objectively optimised to estimate CBF as accurately as possible, given their design constraints. Each protocol was then compared using Monte Carlo simulations and in vivo data. Of the existing designs, the T_1-adjusted and free-lunch time-encoded protocols produced the most accurate CBF estimates, while the $\text{Hybrid}_{\text{T1-adj}}$ protocol had the most accurate CBF estimates of all of the tested protocols. Similarly, to Chapter 4, all of the multi-TI protocols had more accurate CBF estimates than the single-TI protocol, when compared to the ground truth estimates. However, it was also found that all of the protocols, except the fixed label duration time-encoded protocol, had more confident CBF estimates (based on the posterior distribution SD) and more repeatable CBF estimates (based on test-retest RMSEs) than the single-TI protocol. This is despite the multi-TI protocols also being used to estimate ATT simultaneously. This suggests that there is more useful information in the multi-TI protocols for estimating CBF, resulting in detectable test-retest differences. Further improvements in CBF accuracy were demonstrated to be possible by increasing the label duration of the protocols.

In Chapter 6, a time-encoded PCASL preparation was used for dynamic angiography with the aim of reducing the number of readout frames. This enabled the use of larger flip-angles, increasing in vivo SNR by an average of 66% - 68% (for the CFA and $\text{VFA}_{\text{maintain}}$ schemes). This increased SNR improved visualisation of the arterial tree, which was particularly evident for small distal vessels, some of which were near the noise floor in the sequential protocol images. The scan time was approximately halved when the time-encoded data was angularly undersampled by a factor of ~2. However, the resulting angiograms still showed good visual delineation of the arterial tree. The noise level and aliasing energy increased when undersampled, which would likely negatively affect quantitative measures if a model were fit to the data, but a compressed sensing reconstruction may yield improved results. Preliminary results demonstrated that when the BGS timings were optimised to minimise static tissue signal across the readout period, physiological noise could be reduced, leading to an increased measure of SNR.

7.2 Future Research

7.2.1 Perfusion Protocol Optimisation

The perfusion protocol optimisation algorithm was demonstrated to achieve in vivo reductions in CBF and ATT errors. It was also shown that it can be used to optimise a range of different protocol designs and accurately predict which protocol would result in the most accurate CBF estimates. However, only CBF and ATT were estimated from the data, with the other parameters being fixed to literature values. Also, more physiologically correct models were not considered in this work. There are, therefore, several avenues which could be pursued, using the Cramér-Rao lower bound framework developed in this thesis, to further improve the precision and physiological accuracy of CBF measurements.

7.2.1.1 Including an Intravascular Compartment

It can be seen in the angiography data in Chapter 6 that the ASL signal persisted in the large arteries up to the end of the readout, at about ~1 s. Intravascular signal was ignored in the tissue perfusion model used for parameter estimation in Chapters 4 and 5, and the influence of this signal on the results of those chapters were reduced by the use of flow crushing gradients and data exclusion criteria.

It has previously been demonstrated that intravascular blood can be successfully modelled in multi-TI PASL data, by inclusion of a model component that describes labelled blood flowing through a voxel.[1] This intravascular component could be included into the standard perfusion model used for protocol optimisation, with the aim of reducing sensitivity to intravascular signal. The effect on CBF and ATT accuracy of additionally estimating intravascular signal from the data could then be analysed using the Cramér-Rao lower bound estimates and Monte Carlo simulations. However, choosing a realistic distribution of arterial blood volume fraction for protocol optimisation and simulations could prove problematic, since this will vary across the brain with vessel size and vessel density within the voxel.

A recent study has suggested that intravascular signal can be more accurately estimated from perfusion data when high temporal resolutions are used.[2] This makes sense due to the fast dynamics of intravascular signal, but it remains to be seen what effect this has on the accuracy and repeatability of CBF estimates.

7.2.1.2 Inclusion of Other Parameters

Tissue T_1 could be included in the optimisation process in order to investigate how well it can be estimated from ASL data. Alternatively, tissue T_1, as well as other parameters such as blood T_1, could also be included into the optimisation process, though not with the aim of estimating them. Instead, a prior probability distribution could be defined for each parameter in order to find the protocol that best estimates CBF and ATT or just CBF on average across these distributions. This would not reduce the protocol's sensitivity to variations in these parameters in the same way as ATT, for which sensitivity is reduced via estimation, but the resulting protocol timings would reflect the best case protocol for CBF and ATT estimation in the presence of an expected variation in these parameters.

7.2.1.3 Comparison to Look-Locker Readouts

The use of a look-locker (LL) readout for generating multi-timepoint data for perfusion imaging was not considered in this thesis. Previous studies have demonstrated this approach can be more SNR efficient than incrementally adjusting the TI in a PASL protocol for a single-slice.[3,4] It would therefore be of interest to compare the 'single readout' protocol designs used and developed in this thesis with the previously developed LL readout protocols. However, the need to account for signal attenuation with LL readouts would need to be taken into account during any comparison.

A recent study demonstrated the use of a time-encoded PCASL preparation with a LL readout was able to achieve temporal sampling resolutions of up to 75 ms.[2] Similar to the benefits found in this thesis of using a time-encoded preparation with a LL readout for angiography, it may be beneficial for perfusion parameter estimation to achieve some of the temporal sampling with a time-encoded preparation with the remainder being achieved using a LL readout.

7.2.1.4 Incorporation of *k*-space Sampling

Recent work has demonstrated that parametric maps can be directly estimated from undersampled noisy *k*-space data with the use of compressed sensing methods.[5,6] Further work has attempted to directly quantify the effect of non-uniform noise-amplification of undersampled data with multiple coil arrays on parametric maps.[7] It may, therefore, be possible to incorporate acceleration factors, or even *k*-space sampling trajectories, directly into the protocol optimisation.

By way of a simple example, it could then be possible to directly evaluate whether it is beneficial to use an acceleration factor of R = 2 for a single-slice EPI readout, taking into account the reduced TE and the scan time saved (with the resulting increased number of averages), or whether fully sampled data, with fewer averages and longer TE, should be acquired.

This could then extend the optimisation performed in this thesis to not only take into account the effects of the label durations, PLDs, and the number of averages on parameter accuracy, but could also include knowledge of coil sensitivities, gradient slew rates and maximum amplitudes. In this way, a protocol with non-uniform trajectory, variable label durations and PLDs, could be designed which optimally makes use of the scan time available to best estimate CBF and ATT.

7.2.2 Quantitative Angiography Protocol Optimisation

As mentioned at the end of Chapter 6, the perfusion optimisation algorithm from Chapter 4 could be used to optimise angiography protocols to accurately estimate blood volume, transit time, and dispersion parameters using the model described by Okell et al.[8,9] The sequential and time-encoded protocols could then be fairly compared by allowing their protocol timings and flip-angles to vary according to their separate design constraints, with the performance metric being parameter estimation accuracy, similar to Chapters 4 and 5.

Vessel delineation could also be used as a performance metric by using the fitted parameters to generate synthetic angiograms.[8] For visual evaluation of these synthetic images, a long ASL angiography or time-of-flight scan could be performed and used as a ground-truth image.

7.2.3 Extending Angiography to 3D

A simple extension to work in Chapter 6 is the use of a 3D SPGR radial readout. The benefits of a time-encoded readout, with the associated higher flip-angles, are expected to be preserved when using a 3D readout. However, the opportunity to angularly undersample the radial acquisition is increased due to the increased sparseness of the angiograms.

7.2.4 Time-Encoded Angiography with Compressed Sensing

In Chapter 6, angular undersampling of R = 2 was used with the time-encoded protocol to reduce the scan time to approximately the same length as for the sequential protocol with a single average. Higher levels of acceleration could be tested to further reduce the scan time of both protocols making use of compressed sensing reconstructions to reduce aliasing artefacts. It is expected that greater levels of acceleration will be possible with the time-encoded protocol, due to its higher SNR.

7.2.5 Improved BGS Optimisation

Preliminary results in Chapter 6 regarding optimisation of the BGS null time, demonstrated promising results in improved suppression of physiological noise in the dynamic angiograms. However, a larger number of subjects are required in order to evaluate the benefits of this method.

Potential evidence was also seen that may suggest the interleaved BGS inversion pulses visibly reduced the amount of label generated in the middle time-encoded block for the optimised null-time data (Figure 6.20). Testing in further subjects should demonstrate if this is a repeatable effect. Simulations of the signal time-course could also aid in understanding the extent of this effect and in generating alternative BGS options, such adjusting the inversion timings and the label durations such that the inversion pulses play out in-between the time-encoded blocks.

7.2.6 Applications

The improvements made to CBF estimation accuracy through the use of protocol optimisation and the hybrid protocol design will be useful for any application utilising ASL for quantitative CBF mapping. These applications include neuroscience research using ASL-based FMRI[10,11] and calibrated BOLD,[12,13] stroke,[14–16] arteriovenous malformations,[17] characterisation of neurodegenerative diseases such as Alzheimer's[18] and vascular dementia,[19] depression,[20] and neuro-oncology.[21] However, certain applications require different experimental designs to the resting CBF measurements described in chapters 4 and 5. For example, ASL-FMRI applications often require CBF estimates to be generated with a temporal resolution on the order of seconds, not minutes. This design constraint may result in an optimal protocol design quite different to those generated in this thesis. Calibrated BOLD techniques often use administered gas for calibration;[22] the physiological changes that these gases induce need to be taken into account for accurate calibration. For example, hypercapnic calibration, where the CO_2 fraction of the inspired gas is increased to elicit a change in both the BOLD signal and CBF, also affects the arterial blood flow velocity.[13] Since PCASL labelling efficiency is sensitive to the arterial blood velocity, ignoring this effect could lead to inaccurate calibration and erroneous CBF and OEF measurements. Therefore, further model parameters must also be estimated which should also be taken into account when designing the ASL protocol.

It would be particularly interesting to test the effectiveness of the protocols developed in this thesis in clinical cohorts which may not be as compliant as healthy volunteers. As mentioned in Chapters 5 and 6, time-encoded PCASL is potentially more sensitive to motion than standard PCASL. This increased sensitivity to motion could outweigh the improvements to CBF accuracy and angiography SNR that have been demonstrated in this work. Therefore, these protocols should be evaluated in clinical settings. It would also be interesting to evaluate their performance in the presence of delayed blood arrival, such as in older populations or the presence of collateral flow.

7.3 References

1. Chappell M a., MacIntosh BJ, Donahue MJ, Günther M, Jezzard P, Woolrich MW. Separation of macrovascular signal in multi-inversion time arterial spin labelling MRI.

Magn Reson Med 2010;63:1357–1365.

2. van der Plas MCE, Teeuwisse WM, Schmid S, Chappell M, van Osch MJP. High temporal resolution arterial spin labeling MRI with whole-brain coverage by combining time-encoding with Look-Locker and simultaneous multi-slice imaging. Magn Reson Med 2019;81:3734–3744.
3. Günther M, Bock M, Schad LR. Arterial spin labeling in combination with a look-locker sampling strategy: Inflow turbo-sampling EPI-FAIR (ITS-FAIR). Magn Reson Med 2001;46:974–984.
4. Francis ST, Bowtell R, Gowland PA. Modeling and optimization of look-locker spin labeling for measuring perfusion and transit time changes in activation studies taking into account arterial blood volume. Magn Reson Med 2008;59:316–325.
5. Zhao B, Lam F, Liang Z-P. Model-Based MR Parameter Mapping With Sparsity Constraints: Parameter Estimation and Performance Bounds. IEEE Trans Med Imaging 2014;33:1832–1844.
6. Zhao L, Fielden SW, Feng X, Wintermark M, Mugler JP, Meyer CH. Rapid 3D dynamic arterial spin labeling with a sparse model-based image reconstruction. Neuroimage 2015;121:205–216.
7. Hu C, Peters DC. SUPER: A blockwise curve-fitting method for accelerating MR parametric mapping with fast reconstruction. Magn Reson Med 2019:1–15.
8. Okell TW, Chappell M a., Schulz UG, Jezzard P. A kinetic model for vessel-encoded dynamic angiography with arterial spin labeling. Magn Reson Med 2012;68:969–979.
9. Okell TW, Chappell MA, Jezzard P. A theoretical framework for quantifying blood volume flow rate from dynamic angiographic data and application to vessel-encoded arterial spin labeling MRI. Med Image Anal 2013;17:1025–1036.
10. Aguirre GK, Detre JA, Zarahn E, Alsop DC. Experimental Design and the Relative Sensitivity of BOLD and Perfusion fMRI. Neuroimage 2002;15:488–500.
11. Segerdahl AR, Mezue M, Okell TW, Farrar JT, Tracey I. The dorsal posterior insula subserves a fundamental role in human pain. Nat Neurosci 2015;18:499–500.
12. Davis TL, Kwong KK, Weisskoff RM, Rosen BR. Calibrated functional MRI: Mapping the dynamics of oxidative metabolism. Proc Natl Acad Sci 1998;95:1834–1839.
13. Bulte DP, Kelly M, Germuska M, Xie J, Chappell MA, Okell TW, Bright MG, Jezzard P. Quantitative measurement of cerebral physiology using respiratory-calibrated MRI. Neuroimage 2012;60:582–591.
14. Wang DJJ, Alger JR, Qiao JX, Gunther M, Pope WB, Saver JL, Salamon N, Liebeskind DS. Multi-delay multi-parametric arterial spin-labeled perfusion MRI in acute ischemic stroke - Comparison with dynamic susceptibility contrast enhanced perfusion imaging. NeuroImage Clin 2013;3:1–7.
15. Harston GWJ, Okell TW, Sheerin F, et al. Quantification of Serial Cerebral Blood Flow in Acute Stroke Using Arterial Spin Labeling. Stroke 2017;48:123–130.
16. Chalela JA, Alsop DC, Gonzalez-Atavales JB, Maldjian JA, Kasner SE, Detre JA. Magnetic Resonance Perfusion Imaging in Acute Ischemic Stroke Using Continuous Arterial Spin Labeling. Stroke 2000;31:680–687.
17. Wolf RL, Wang J, Detre JA, Zager EL, Hurst RW. Arteriovenous Shunt Visualization in Arteriovenous Malformations with Arterial Spin-Labeling MR Imaging. Am J Neuroradiol 2008;29:681–687.
18. Wierenga CE, Hays CC, Zlatar ZZ. Cerebral blood flow measured by arterial spin labeling MRI as a preclinical marker of alzheimer's disease. J Alzheimer's Dis 2014;42:S411–S419.
19. Gao Y-Z, Zhang J-J, Liu H, Wu G-Y, Xiong L, Shu M. Regional Cerebral Blood Flow and Cerebrovascular Reactivity in Alzheimer's Disease and Vascular Dementia Assessed

by Arterial Spinlabeling Magnetic Resonance Imaging. Curr Neurovasc Res 2012;10:49–53.
20. Duhameau B, Ferré J-C, Jannin P, Gauvrit J-Y, Vérin M, Millet B, Drapier D. Chronic and treatment-resistant depression: A study using arterial spin labeling perfusion MRI at 3Tesla. Psychiatry Res Neuroimaging 2010;182:111–116.
21. Wang P, Li J, Diao Q, et al. Assessment of glioma response to radiotherapy using 3D pulsed-continuous arterial spin labeling and 3D segmented volume. Eur J Radiol 2016;85:1987–1992.
22. Blockley NP, Griffeth VEM, Simon AB, Buxton RB. A review of calibrated blood oxygenation level-dependent (BOLD) methods for the measurement of task-induced changes in brain oxygen metabolism. NMR Biomed 2013;26:987–1003.

Appendices

A Complete CBF Sensitivity Function

The CBF sensitivity function where T_1' is not constant with respect to f is:

$$
\begin{aligned}
\frac{\partial \Delta M(t)}{\partial f} &= 0 && 0 < t < \Delta t \\
&= 2M_{0B}\alpha e^{\frac{-\Delta t}{T_{1b}}}\Bigg[T_1'\left(1 - e^{\frac{-(t-\Delta t)}{T_1'}}\right) \\
&\quad + \left(\frac{fT_1'(t-\Delta t)}{\lambda} - f\frac{\partial T_1'}{\partial f}\right)e^{\frac{-(t-\Delta t)}{T_1'}} && \Delta t < t < \tau + \Delta t \\
&\qquad + f\frac{\partial T_1'}{\partial f}\Bigg] \\
&= 2M_{0B}\alpha e^{\frac{-\Delta t}{T_{1b}}}\Bigg[T_1' e^{\frac{-(t-\tau-\Delta t)}{T_1'}}\left(1 - e^{\frac{-\tau}{T_1'}}\right) \\
&\quad + \left(\frac{fT_1'(t-\Delta t)}{\lambda} - f\frac{\partial T_1'}{\partial f}\right)e^{\frac{-(t-\Delta t)}{T_1'}} && \tau + \Delta t < t \\
&\quad + \left(f\frac{\partial T_1'}{\partial f} - \frac{fT_1'(t-\tau-\Delta t)}{\lambda}\right)e^{\frac{-(t-\tau-\Delta t)}{T_1'}}\Bigg]
\end{aligned}
\tag{A.1}
$$

where $\frac{\partial T_1'}{\partial f} = -\frac{\lambda T_{1t}^2}{(\lambda + fT_{1t})^2}$.

B Variable Noise Measurements and Monte Carlo Experiments

This appendix describes the measurement of and the Monte Carlo simulations using the variable noise across the PLDs from the in vivo data in Chapter 4.

The tSNR, mean signal, and SD was calculated from the combined reference multi-PLD and single-PLD control images from all 7 subjects, as described in Chapter 4. This data is plotted across PLDs in Figure B.1, where it is clear that the noise magnitude is not constant across PLDs. The reference multi-PLD and single-PLD control data was used here because multiple averages at each PLD can be used to calculate the mean signal and SD voxelwise at each PLD. The noise magnitude appears to vary in a similar manner to the residual tissue signal, decreasing as the PLD approaches 1.25 s then increasing for longer PLDs.

The noise varies in a non-simple manner across PLDs, so in order to generate a noise model the tSNR was fit with an empirical exponential decay model and the background suppressed tissue signal was simulated using a series of simple saturation and inversion recovery calculations. The PLD dependent noise model could then be described by division of the tissue signal model by the fitted tSNR model.

The background suppressed tissue signal assumed perfect and instantaneous presaturation and inversion pulses and used the scanner inversion timings. The median tSNR was fit with an exponential decay model of the form $tSNR(t) = a \cdot e^{-b \cdot t} + c$, where t is the label duration plus the PLD, and $a,\ b,$ and c are constants to be estimated. The noise SD can then be described as $\sigma(t) = signal(t)/tSNR(t)$. This noise SD was scaled so that the average SD was equal to that used in Chapter 4. Monte Carlo simulations were performed in a similar manner to those in Chapter 4, except that the variable noise model was used to add noise to the data. Note, the variable noise level is not accounted for in the NLLS fitting process.

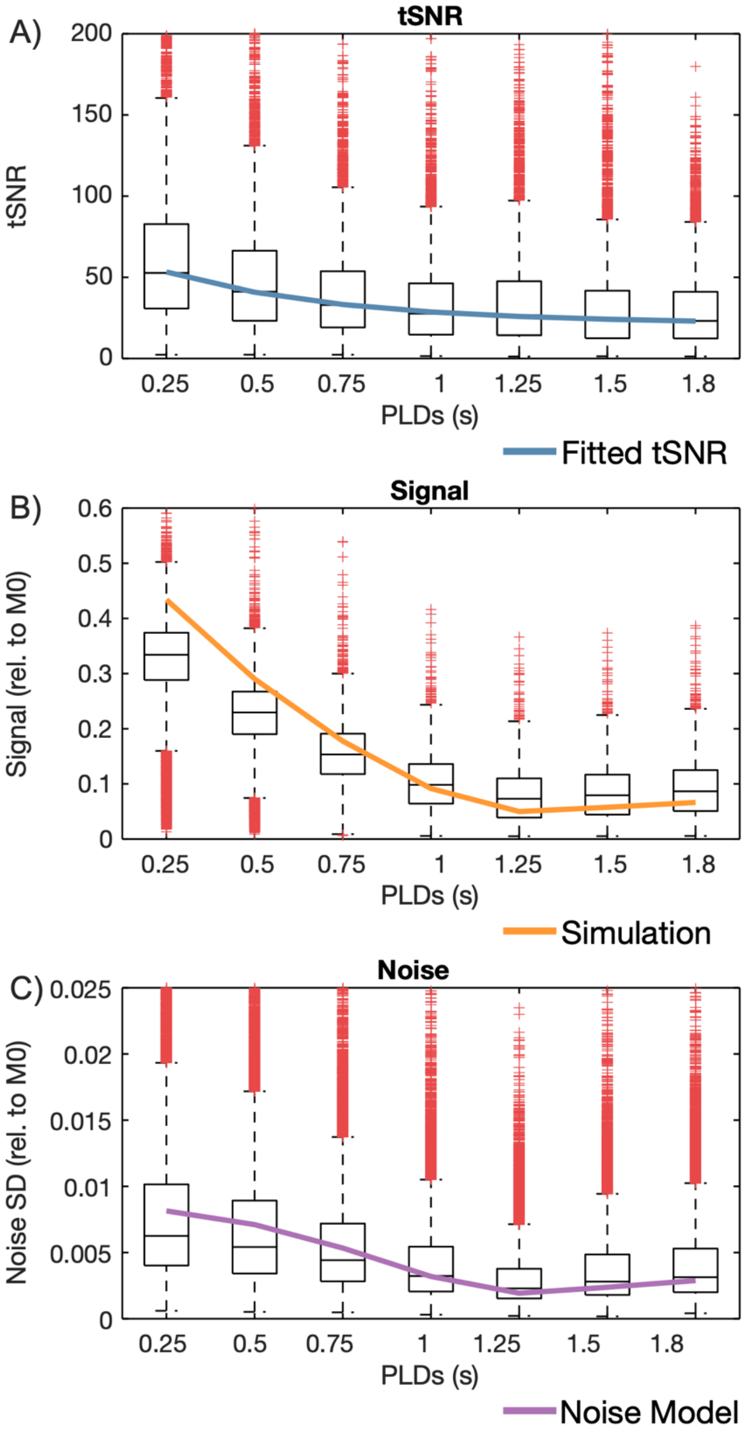

Figure B.1: Boxplots of the tSNR (A), signal (B), and noise (C) of the GM voxels from the control ASL data of the reference multi-PLD and single-PLD data for all 7 subjects. The median tSNR has been fit using an exponential decay model, while the tissue signal (normalised by M_{0B}) was simulated using a series of saturation and inversion recovery models. The noise can then be modelled as the simulated signal divided by the fitted tSNR.

The results of the tissue signal simulations, tSNR fitting, and the resulting (unscaled) noise model are shown overlaid on the in vivo data in Figure B.1. The RMSEs of the CBF and ATT estimates for each protocol are shown in Figure B.2 alongside the RMSEs of the uniform noise case. While the reference multi-PLD protocol CBF RMSE increased by an average of 0.80 mL/100g/min across the ATT range, the RMSE of the other protocols decreased by 0.62 mL/100g/min (CBF-ATT$_{opt}$), 1.02 mL/100g/min (CBF$_{opt}$), and 1.47 mL/100g/min (single-PLD), reflecting the reduced noise due to the PLDs of these protocol. All of the protocols had increased ATT RMSE, by: 0.074 s

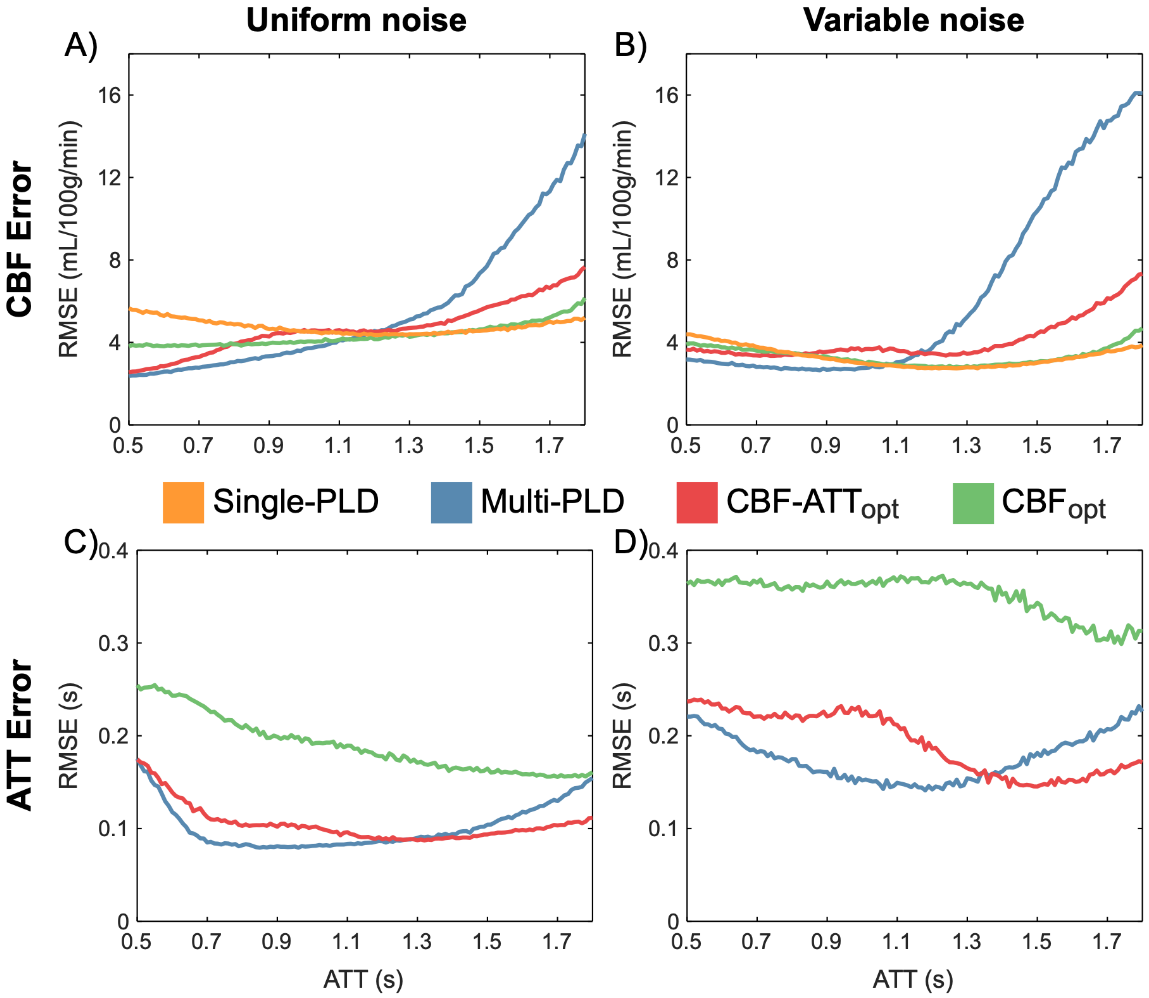

Figure B.2: The effect on CBF (A,B) and ATT (C,D) RMSEs when using uniform (A,C) or variable (B,D) noise across PLDs. The broad trends are consistent. However, all protocols except the reference multi-PLD protocol have reduced CBF errors at late ATTs and the ATT errors for all protocols increased across the entire ATT range. The single-PLD, CBF-ATT$_{opt}$ and CBF$_{opt}$ CBF estimates benefit from having many PLDs at times of reduced noise with the variable noise model.

(reference multi-PLD), 0.088 s (CBF-$\mathrm{ATT_{opt}}$), and 0.162 s ($\mathrm{CBF_{opt}}$), presumably because of the increased uncertainty due to equally weighting the data despite there being varying noise levels.

C 3D GRASE Readout

The ASL white paper[1] recommends using either a 3D multi-echo stack-of-spiral readout[2,3] or a 3D GRASE readout[4,5] for ASL perfusion measurements to ensure high SNR efficiency and identical whole brain BGS and PLDs. A 3D GRASE readout was used in Chapter 5 and this appendix details several experiments that were performed to improve the default slab profile. There were two aims: (1) to avoid the need for slice-oversampling which extends the readout time and (2) to increase the SNR in the edge slices due to signal roll-off.

C.1 3D GRASE Slab Profile

During the 3D GRASE readout, both the excitation and refocussing pulses are slab selective. In our default 3D GRASE readout, the excitation pulse has a relatively high time-bandwidth product (TBWP) compared to a low TBWP for the refocussing pulses. This was likely done to reduce the chance of voltage clipping of the higher flip-angle refocussing pulses and to reduce SAR. However, the low TBWP results in a poorer slab profile for the refocussing pulses compared to the excitation pulses. This is not necessarily a problem, since the refocussing pulses use identical crusher gradients immediately before and after the RF pulse in order to dephase newly created transverse signal, the FID, while preserving the desired transverse signal created by the 90° excitation pulse, in the form of spin- and stimulated-echoes. Therefore, if the refocussing pulses generate new transverse signal beyond the original excitation slab, this will be repeatedly spoiled by the crusher gradients. Nonetheless, a problem can arise if the transition band of the refocussing slab encroaches into the excited slab. If this happens, transverse signal at the edge of the slab created by the excitation pulse will experience a lower effective refocussing flip-angle than intended, reduced the signal of the corresponding edge slices.

A series of experiments were performed in a phantom to test whether the refocussing pulses were adversely affecting the effective 3D GRASE slab profile. This was done by rotating the image readout by 90° relative to the excitation slab so that the 3D-encode and in-plane phase-encode directions are effectively switched, similar to Günther et al.[5]

General experiment details were: large spherical oil phantom (diameter approximately 250 mm), FOV 250 x 250 mm^2, base matrix 63 x 64, nominal resolution 3.91 x 3.91 mm^2, EPI factor 63, nominal slice thickness 5.3 mm (maximum allowed), 30 partitions, 50% partition over-sampling (to reduce aliasing), turbo factor 15 (3 segments in 3D encode direction, to reduce through-plane blurring), bandwidth 2298 Hz/Px, TE 36.5 ms, TR 20 s, excitation flip-angle 90°, refocussing flip-angle 120°.[6,7] Readout, phase-encode, and 3D-encode directions were right-left, superior-inferior, and anterior-posterior, respectively, while the excitation slab was slice-selective along the superior-inferior axis.

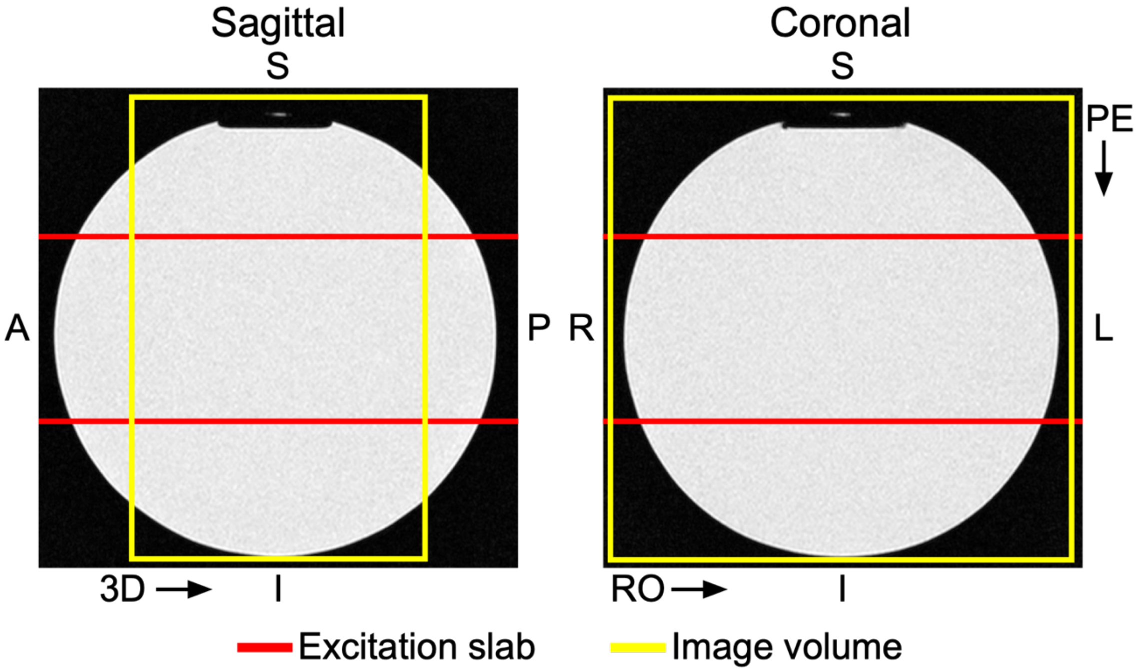

Figure C.1: The experimental setup for the 3D GRASE slab profile experiments. The excitation slab and image volume are shown in red and yellow, overlaid on the oil phantom. The physical orientations are displayed around the images (S: superior, I: inferior, A: anterior, P: posterior, R: right, L: left). The k-space encoding directions are displayed with arrows around the images (RO: readout, PE: phase-encode, 3D: 3D-encode).

The excitation slab thickness was 100 mm wide, which was found to be sufficient to cover most of the brain (the cortex and part of the cerebellum) in most volunteers and matches the slab width used by Günther et al.[5]

Figure C.1 shows the experimental set-up. For each acquisition, 2 averages were acquired with approximately 19 s pause between each acquisition. As default, the sequence was set up to excite a slab 10% wider than requested, in order to avoid signal drop off at

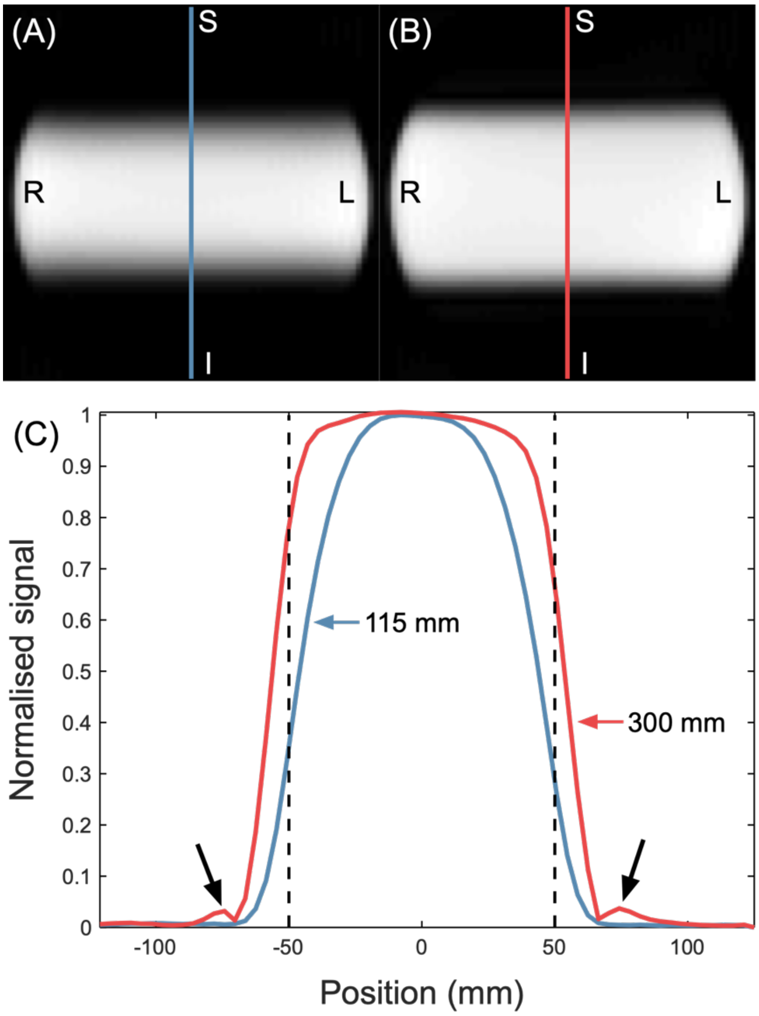

Figure C.2: A cross-section of the slab profile of the 3D GRASE readout. (A) The slab profile with the default settings (110 mm excitation width, 115 mm refocusing width); (B) the slab profile with an increased refocusing width of 300 mm; (C) the cross-section of (A) and (B) with the nominal width depicted with dashed black lines. The location of the cross-sections in (A) and (B) are shown with blue and red lines, respectively. The physical orientations are displayed around the images in (A). The location of the sidelobes from the sinc excitation in (C) are highlighted with black arrows. The data uses the mean of the middle 11 slices (A-P direction) and 2 averages.

the edge of the image, while the refocussing pulse widths were 15% wider than requested, due to the lower TBWP. Both excitation and refocussing pulses use windowed-sinc pulses.

Figure C.2, which uses the mean of the 11 middle slices and 2 averages, shows the difference between the default settings (excitation width 110 mm, refocussing width 115 mm) and setting the refocussing width to 300 mm. It is clear that the flip-angle roll-off of the default refocussing pulses reduced the signal at the edge of the slab created by the excitation pulse. When the refocussing width is larger, the slice profile better reflects the profile of the excitation pulse, including the characteristic sidelobes generated by the windowed-sinc excitation pulse. There is also a much more uniform signal magnitude inside the nominal imaging region, increasing the average signal within the spatial range of -50 mm to 50 mm by 14%. However, there is also a large increase in the signal outside the nominal region, which will alias into the image unless slice-oversampling is used. An alternative method for reducing this signal, without extending scan time by slice-oversampling, is to use outer-volume suppression (OVS) pulses immediately prior to the excitation pulse. Since the OVS pulses can play out during the PLD, they do not add any additional scan time but do constrain the shortest PLD that can be used.

C.2 Outer-Volume Suppression

Outer-volume suppression pulses aim to saturate signal from outside the imaging region in order to prevent it from being excited with the magnetisation in the imaging region. Transverse magnetisation outside the spatially encoded region will be aliased into the image if OVS or oversampling is not used. Günther et al.[5] used a technique similar to Q2TIPS (which suppresses signal proximal to the image volume by application of repeated suppression pulses) to suppress signal both proximal and distal to the image volume by modulating the Q2TIPS pulse with a cosine function to create two parallel, spatially separated, suppression pulses (a specific case of multi-banded pulses).[8] For the FAIR labelling scheme used in reference [5], this has the advantage of suppressing inflowing signal from both arteries and veins which may be proximal or distal to the image volume, respectively, but also has the added advantage of improving the 3D GRASE slab profile. This technique was used here solely for the purpose of improving the 3D GRASE slab profile in order to reduce the need for slice-oversampling.

The slice profile of a large slab, which will have a correspondingly large transition band, can be improved after application of OVS because magnetisation in the transition zone and side-lobes will already be saturated. All else being equal, the size of the transition band is proportional to the width of the slice, and since the OVS pulse width can be chosen to be smaller than the width of the 3D slab, it is possible to achieve a more well defined slice profile than for the 3D slab excitation. Therefore, the signal from just outside the nominal slab can be saturated with higher specificity, preventing the side-lobes and part of the transition width from being excited by the readout excitation.

There was existing sequence code available for sequentially applying an arbitrary number of 40 mm wide, cosine-modulated, 90° windowed-sinc excitation pulses immediately before the 3D GRASE excitation pulse. The pulses were centred on the middle of the image volume and separated from each other by the image volume width + 40 mm. The RF phase of each excitation pulse was varied using a quadratic scheme[9] to improve signal spoiling and each pulse was followed by crusher gradients with cycled amplitudes to prevent stimulated echoes.

A saturation module using the optimised WET presaturation flip-angles[10] are able to provide improved saturation across a wide range of T_1 values with reasonable B_1^+ insensitivity (±10%) compared to a series of 90° pulses. The same OVS scheme as above was therefore implemented, but for 4 pulses which used the WET flip-angles from Golay et al.[10] in order to compare this with the default OVS settings.

For the WET OVS implementation, the pulse durations and gradient cycling scheme were altered to match those in the WET presaturation module. The TBWP was also reduced by a small amount to prevent the final WET pulse, with flip-angle of 159°, being clipped due to insufficient amplifier voltage.

Figure C.3 demonstrates the different slab profiles that result from using no OVS, 4 default OVS pulses, and 4 WET OVS pulses. For an excitation width of 110 mm and a refocussing width of 300 mm, there is a clear reduction of signal outside the nominal region when using either OVS module compared to no OVS. However, when the WET OVS module is used, the signal outside is reduced by 12% more than when using the default OVS pulses. This is made more visually clear in Figure C.3(B) where the 3D GRASE excitation width was increased to 130 mm. One downside of using either OVS module is the slight reduction of signal inside the nominal region (in Figure C.3(A): default 1.2%, WET 1.5%), suggesting that the OVS modules are disturbing a small proportion of spins

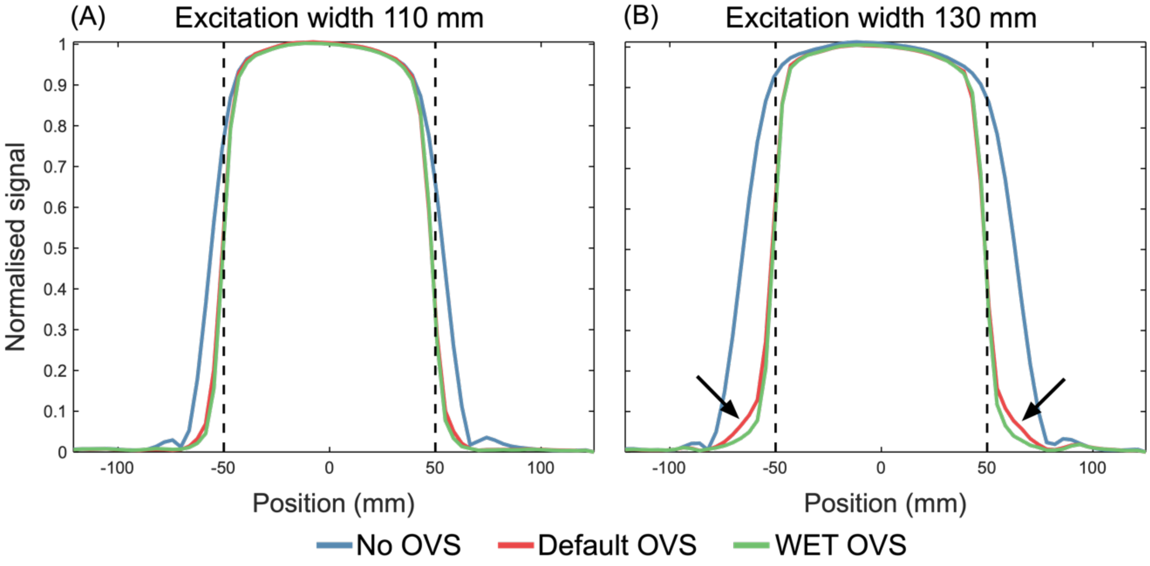

Figure C.3: Comparing the slab profile of the 3D GRASE readout with no OVS, 4 default OVS pulses, and with 4 WET OVS pulses. (A) The slab profile with an excitation width of 110 mm and (B) the slab profile with a 130 mm excitation width. In both cases, the refocusing width was 300 mm. The nominal width of 100 mm is depicted with dashed black lines. The arrows in (B) highlight the more obvious improvement of the WET OVS pulses when the excitation width is larger. The data uses the mean of the middle 11 slices (A-P direction) and 2 averages.

within the image volume. Nonetheless, this effect is small compared to the reduction in unwanted signal outside the nominal region.

Due to the improved performance of the WET OVS pulses, they were used in preference to the default OVS pulses in all further experiments.

C.3 Refocussing Pulse Width

During an in vivo experiment, each TR would ideally start immediatley after the end of the previous 3D GRASE readout to maximise time-efficiency. Since each TR begins with the WET presaturation module followed immediately by the PCASL labelling there is only a short time between the last refocussing pulse in the readout and the start of the PCASL labelling. It is, therefore, important that the refocussing pulses do not interfere with inflowing blood that will be labelled in the next TR. Although the 300 mm refocussing width used above provides an improved slab profile compared to the default 115 mm width, it extends by a nominal distance of 100 mm below the inferior edge of the image

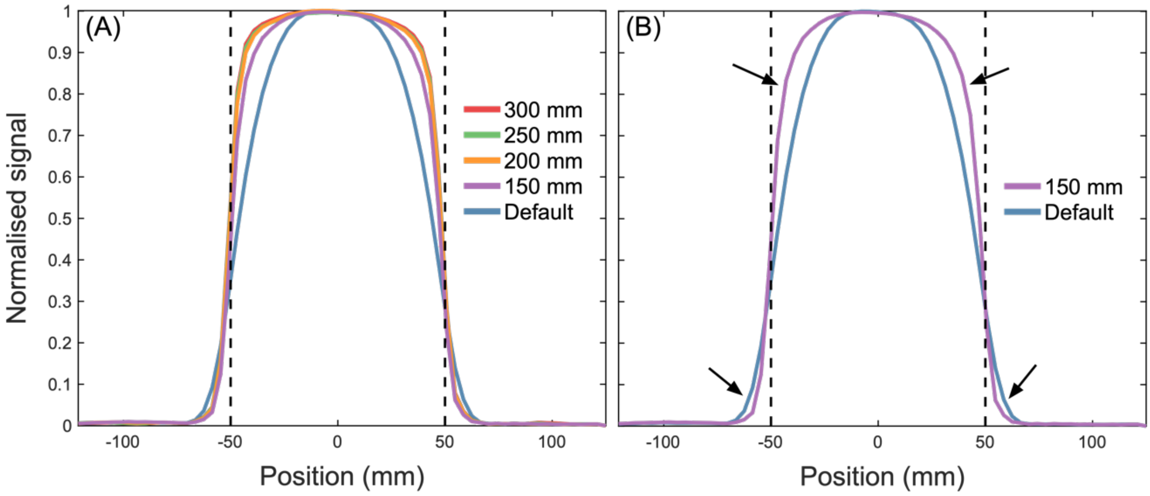

Figure C.4: (A) Comparing the 3D GRASE slab profile with the default settings (110 mm excitation width, 115 mm refocusing width, no OVS) and a variety of refocusing widths (150 - 300 mm) with WET OVS. (B) Showing only the default slab profile alongside the improved slab profile with a refocusing width of 150 mm and WET OVS. The new slab profile has a more uniform signal within the nominal region, and a reduced signal outside. The nominal width of 100 mm is depicted with dashed black lines. The data uses the mean of the middle 11 slices (A-P direction) and 2 averages.

volume, and likely influences spins beyond this due to the poorly defined profile of the pulse.

In a preliminary in vivo scan, where a 100 mm thick image volume was placed with its superior edge flush with the top of the brain and the PCASL labelling plane positioned as in Chapter 4, the bottom of the image volume was ~46 mm away from the centre of the labelling plane, well within the nominal 100 mm reach beyond the edge of the image volume of the refocussing pulses from the previous TR. Even when the blood flow rate is taken into account, it is still likely that inflowing blood which should be labelled in the nth TR may be peturbed by the final refocussing pulse in the (n-1)st TR (assuming a blood flow rate of 20 cm/s, plug flow, and a gap of 100 ms between the last refocussing pulse and the start of labelling only adds an extra 20 mm of effective distance between labelling plane and readout). If a refocussing width of 150 mm were used instead, the nominal inferior edge of the refocussing pulse would be 21 mm above the centre of the labelling plane, even without accounting for the blood flow rate.

Figure C.4(A) shows the 3D GRASE slab profile for a range of different refocussing widths with the WET OVS module and the default slab with no OVS. For the refoccusing width of 150 mm, there is a reduction of signal both inside and outside the nominal region of 3.1% and 17%, respectively, compared to the refocussing pulse width of 300 mm. Although there is a slight reduction in SNR from within the target region when using a refocussing width of 150 mm, there is also less signal from outside the nominal image volume.

Compared to the default slab profile without OVS, the 150 mm refocussing width with WET OVS has an average 8.9% higher signal inside the nominal region and 18% lower signal outside the target region. The refocussing width of 150 mm, along with the WET OVS, still improves the overall shape of the slab profile, as can be seen in Figure

C.4(B), and reduces signal which would normally be aliased into the image volume. Therefore, the final chosen settings for the 3D GRASE readout were: WET OVS, a refocussing width of 150 mm, and no slice-oversampling.

Further improvements to the OVS, which are beyond the scope of this thesis, could include using maximum phase pulses[8] to achieve improved signal spoiling, temporally offsetting the multi-banded pulses to reduce peak B_1^+ (allowing for a higher TBWP), use of SLR optimised pulse waveforms to improve the slice profile for each individual flip angle (the 159° pulse especially will have a degraded slice profile due to non-linearities in the Bloch equations[8]), and WET flip-angle optimisation to ensure that the saturation null-time occurs at the centre of the 3D GRASE excitation pulse.

D Background Suppression Inversion Pulses

This appendix details the simulations and experiments carried out to optimise the pulse profile of the slice-selective interleaved FOCI pulses used in Chapters 5 and 6.

The adiabatic inversion pulses used for BGS in Chapter 4 were hyperbolic-secant (HS) pulses with settings: $\mu = 24$, $\beta = 800$ rad/s, pulse duration (T_p) 10.24 ms, pulse truncation 3.33% and with the peak voltage set to the maximum available amplifier amplitude. Note, μ is the side-to-width ratio and so determines the sharpness of the slab profile, while β is the modulation angular frequency.

For the interleaved slice-selective (SS) inversion pulses used in Chapters 5 and 6, the inferior edge of the inversion plane must be placed such that it is flush with the centre of the PCASL labelling plane. Ideally, this inversion pulse would have a perfectly rectangular profile, so as to perfectly invert all spins right up to the labelling plane. However, since this is not possible in practice, we aimed to evaluate the inversion profile of the default SS HS pulse and to improve it by using FOCI modulation,[11] as described in Chapter 2.

A grid-search was performed across a range of FOCI pulse parameters, A_{max} (the maximum FOCI modulation factor) and μ, to find the set of parameters that achieved the best inversion efficiency inside a defined inversion region while minimising the perturbation of magnetisation outside, using Bloch simulations. In this way, the parameters that gave the best overall inversion efficiency and slab profile can be found.

Similar to Payne and Leach,[12] β was set to achieve a FOCI pulse truncation of 4%, using the formula $\beta = 2 \cdot \cosh^{-1}(A_{max}/0.04)/T_p$. Note, this formula differs from that given in reference [12] in that A_{max} is explicitly included to keep the truncation level at 4% no matter the value of A_{max}. T_p was fixed at 10.24 ms, and inversion width set to 200 mm.

The grid-search of the parameters evaluated were: A_{max} = 1 - 20, and μ = 0.5 - 24, both evaluated every 0.5. Bloch simulations used a spatial grid across the range -200 to 200 mm with samples every 1 mm. The simulation peak B_1^+ was set to 10 μT to ensure that the FOCI pulse performed well even at low power, rather than risk the pulse being underpowered in the centre. The maximum B_1^+ at isocentre on the Siemens Prisma scanner in the data collected for Chapter 4 was ~25 μT, so 10 μT represents a 60% drop in B_1^+. The inversion slab was centred at 0 mm. T_1 and T_2 of blood were assumed to be 1650 ms[1] and 150 ms,[13] respectively.

The optimisation aims to solve the following optimisation problem:

$$\underset{A_{max},\mu}{\operatorname{argmin}} \left(mean\left(M_z^{inside}\right) \cdot mean\left(M_z^{outside}\right) \right), \tag{D.1}$$

which has a minimal value of -1 when $mean\left(M_z^{inside}\right) = -1$ and $mean\left(M_z^{outside}\right) = 1$. In this way, the inversion efficiency inside versus perturbation of magnetisation outside a defined region can be balanced and the transition band of the inversion region minimised.

The optimisation returned the following values as having minimised Equation 5.1: A_{max} = 20, μ = 1.5, and ß = 1349.17 rad/s. The top row of Figure D.1 shows the simulated inversion profile for a range of peak B_1^+, for the default and optimised pulses immediately after the end of a single pulse. It is clear that the optimised FOCI pulse has a much improved inversion efficiency across the defined slab at all except very high B_1^+, where the default HS performs similarly to the optimised FOCI. The bottom row shows a zoomed in section of the left-hand side of the inversion slab for 3 different resonance conditions. Figure D.1(D) demonstrates the reduced transition width of the optimised FOCI pulse on-resonance, while Figure D.1(C) and (E) demonstrate that the larger gradient used at the edges of the optimised FOCI pulse results in a smaller off-resonance slice-shift than for the default HS.

Phantom experiments, and matched Bloch simulations, were also carried out to check the performance of each pulse using a range of B_1^+ values. The experiments and simulations both used a presaturation followed by two inversion pulses, optimally timed to null signal 100 ms before the readout using a TI of 250 ms and $T_{1,opt}$ of 700 ms (the centre of the inversion pulses were 39 ms and 114 ms after the presaturation), matching the layout of the standard BGS of the ASL data, but with a short TI to account for the short T_1 and T_2 of the phantom (not measured but assumed to be 500 ms and 50 ms, respectively). The simulations assumed perfect spoiling of magnetisation at t = 0 ms.

The experimental settings were as follows: 3D GRASE readout, matrix size 96 x 96, FOV 192 x 192 mm (giving 2 mm isotropic in-plane resolution), 8 partitions, slice thickness 2.7 mm, 4 segments along phase-encode direction to reduce TE, 2 segments along 3D encode direction to reduce blurring, bandwidth 1002 Hz/px, TE 31.8 ms,

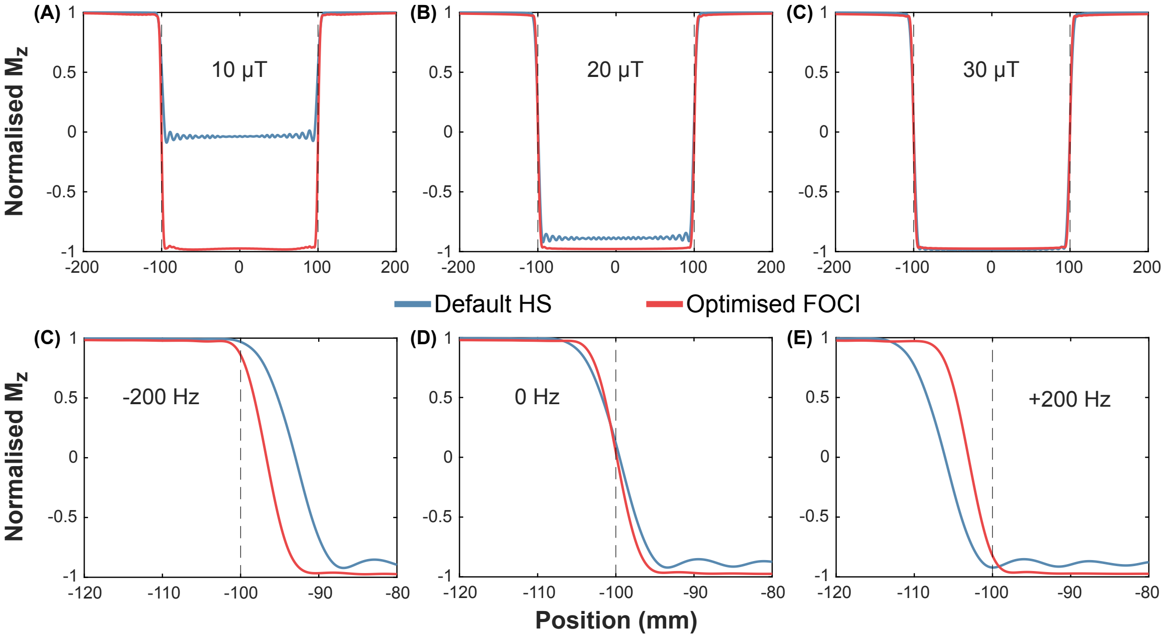

Figure D.1: Bloch simulations of the default hyperbolic-secant pulse (default HS) and the optimised c-shaped FOCI modulated hyperbolic-secant pulse (optimised FOCI). The top row shows the simulated normalised longitudinal magnetisation after each pulse has been applied with a maximum B_1^+ of (A) 10 µT, (B) 20 µT and (C) 30 µT. The bottom row shows the transition zone to the left of the inversion slab for 3 different off-resonance conditions with peak B_1^+ = 20 µT: (C) -200 Hz, (D) on-resonance and (E) +200 Hz.

transverse imaging slab. A pause of 3 s after each acquisition was used to allow recovery of the longitudinal magnetisation.

The WET presaturation slab was transverse and covered the entire image volume. The slice-selective inversion pulses were placed sagittal with the left-hand edge at the centre of the image volume. Global (non-selective: NS) HS inversion pulses with the default settings were also used as a comparison. The peak voltage was set to the maximum value, then halved, then quartered to give data at three different peak B_1^+ values. The B_1^+ values were calculated from the default reference voltage measured during scanner calibration. Data using the NS default HS pulse was only acquired with the maximum voltage and was used to normalise all other data, because the magnetisation is expected to

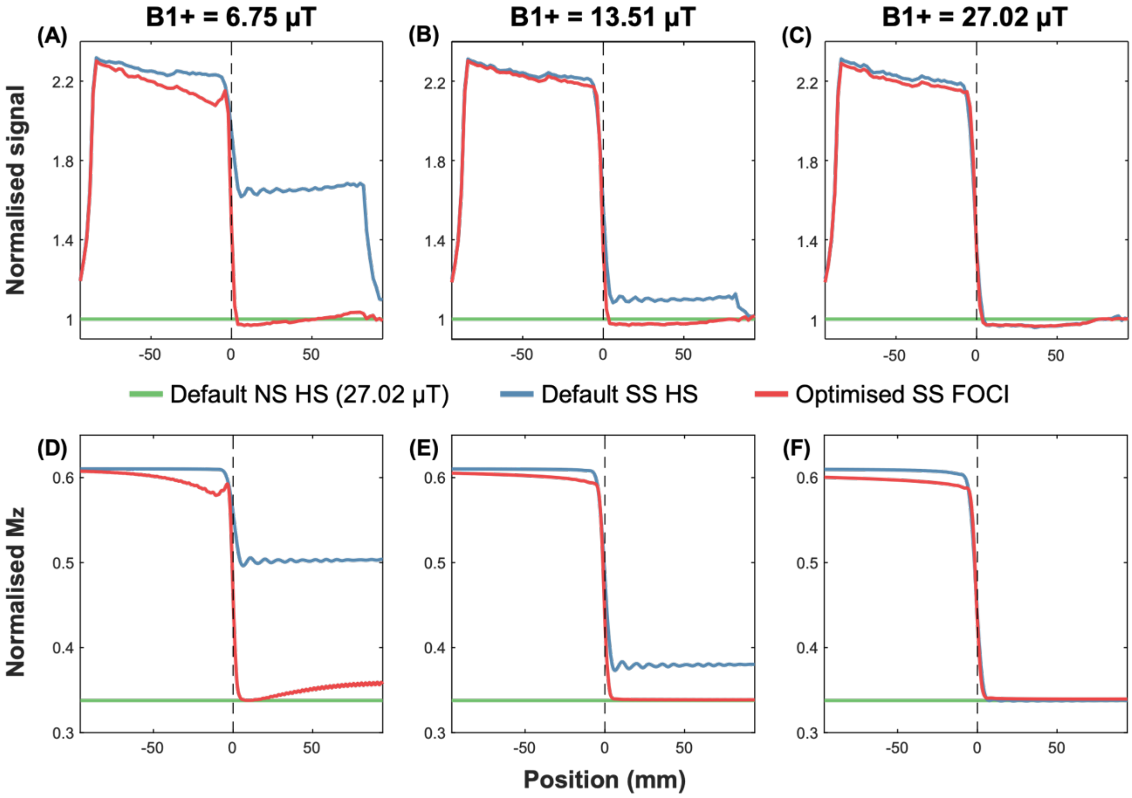

Figure D.2: Phantom data (top) and matched Bloch simulations (bottom) for the non-selective default HS pulse, the slice-selective default HS pulse and the optimised slice-selective FOCI pulse. For the phantom scans, the peak voltage was set to the maximum, half of maximum and quarter of maximum giving an estimated B_1^+ of 6.75 μT (A, D), 13.51 μT (B, E) and 27.02 μT (C, F), respectively. NS HS phantom data was only collected with B_1^+ = 27.02 μT, but is shown in (A) and (B) as a reference and was used to normalise all of the data. The reduced signal for the FOCI pulse simulation in (F) for spins at positions < 0, corresponds to a ~1.7 % reduction.

be ideally suppressed in this case. Simulations were run after the phantom experiments with matched B_1^+, slab positioning, and spatial resolution.

The top row of Figure D.2 shows the results of the phantom experiments and demonstrates that the optimised FOCI pulse does indeed maintain greater inversion efficiency than the default SS HS for a wide range of B_1^+ levels, except at the maximum B_1^+, when the two SS pulses perform similarly. Good agreement can be seen between the phantom and simulations for the same B_1^+ levels, including the oscillations in the default SS HS pulse profile, and the perturbation of magnetisation outside the inversion slab for the optimised SS FOCI pulse at $B_1^+ = 6.75$ µT. At $B_1^+ = 27.02$ µT, when both SS pulses achieved good inversion efficiency, the FOCI pulse can be seen to undesirably perturb magnetisation outside the inversion slab. However, the reduced signal of spins outside the inversion slab for the FOCI pulse only corresponded to a ~1.7% reduction in Mz compared to the SS HS pulse in simulations.

Due to the improved slice profile and robustness to B_0 and B_1^+ variation of the optimised FOCI pulse, these setting were used for all of the experiments using the interleaved BGS inversion scheme in Chapters 5 and 6.

References

1 Alsop DC, Detre JA, Golay X, Günther M, Hendrikse J, Hernandez-Garcia L *et al.* Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn Reson Med* 2015; **73**: 102–116.

2 Irarrazabal P, Nishimura DG. Fast Three Dimensional Magnetic Resonance Imaging. *Magn Reson Med* 1995; **33**: 656–662.

3 Yang Y, Glover GH, van Gelderen P, Mattay VS, Santha AKS, Sexton RH *et al.* Fast 3D functional magnetic resonance imaging at 1.5 T with spiral acquisition. *Magn Reson Med* 1996; **36**: 620–626.

4 Feinberg DA, Oshio K. GRASE (gradient- and spin-echo) MR imaging: a new fast clinical imaging technique. *Radiology* 1991; **181**: 597–602.

5 Günther M, Oshio K, Feinberg DA. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn Reson Med* 2005; **54**: 491–498.

6 von Samson-Himmelstjerna F, Madai VI, Sobesky J, Guenther M. Walsh-ordered hadamard time-encoded pseudocontinuous ASL (WH pCASL). *Magn Reson Med* 2016; **76**: 1814–1824.

7 He X, Wengler K, Schweitzer ME. Diffusion sensitivity of 3D-GRASE in arterial spin labeling perfusion. *Magn Reson Med* 2018; **80**: 736–747.

8 Bernstein MA, King KF, Zhou XJ. *Handbook of MRI Pulse Sequences*. Elsevier: Amsterdam, 2004 doi:10.1016/B978-0-12-092861-3.X5000-6.

9 Zur Y, Wood ML, Neuringer LJ. Spoiling of transverse magnetization in steady-state sequences. *Magn Reson Med* 1991; **21**: 251–263.

10 Golay X, Petersen ET, Hui F. Pulsed Star Labeling of Arterial Regions (PULSAR): A robust regional perfusion technique for high field imaging. *Magn Reson Med* 2005; **53**: 15–21.

11 Ordidge RJ, Wylezinska M, Hugg JW, Butterworth E, Franconi F. Frequency offset corrected inversion (FOCI) pulses for use in localized spectroscopy. *Magn Reson Med* 1996; **36**: 562–566.

12 Payne GS, Leach MO. Implementation and evaluation of frequency offset corrected inversion (FOCI) pulses on a clinical MR system. *Magn Reson Med* 1997; **38**: 828–833.

13 Lu H, Xu F, Grgac K, Liu P, Qin Q, van Zijl P. Calibration and validation of TRUST MRI for the estimation of cerebral blood oxygenation. *Magn Reson Med* 2012; **67**: 42–49.