Improving Acquisition Speed and Efficiency of Advanced Arterial Spin Labelling MRI

Sonja Sophie Schauman
St. Cross College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2020

Acknowledgements

Although I am the only one with my name on the front page of this thesis, many other people contributed hugely to it. I could not have written it without the support, encouragement, and love from the people around me.

First, I would like to thank my supervisors, Tom Okell, Mark Chiew, and Peter Jezzard, who guided me through the DPhil process. Tom and Mark tirelessly met with me every week, helped me when I was stuck, and gave detailed and valuable feedback on all abstracts, practice presentations, postdoc applications, funding proposals, manuscripts, and more. Peter was always there in the background with vast experience and sound advice when needed. Apart from the science they taught me, they showed me how to be a scientist. Tom and Mark showed me how being an academic does not mean you have to devote 100% of your life to the lab, but you can have a home life with a growing family whilst still being the best supervisor a DPhil candidate could ask for. My supervisors made time for me, but always put their families first. As it should be. One day, I hope to be like them: professional, knowledgeable, and most importantly, caring.

My colleagues and friends at FMRIB made the work over the last three years a joy, both before and after the onset of the COVID-19 pandemic that forced us to work from home and moved our conversations online. I loved the chats we had over lunch and tea breaks, I loved the overall supporting atmosphere at the Centre, and I loved that people always took the time to help each other out, as well as time to celebrate each other's achievements. Although everyone at FMRIB contributed to this atmosphere, I would like to thank a few people in particular that deserve a special mention: The genius ladies of WIN methods (GLOW), the best group of women supporting each other. Carinne Piekema for helping me get into public engagement as well as being a good friend. Jack Allen and Yan Tong for being the best desk neighbours. Joe Woods and Harry Mason for paving the way for me, working on similar problems to mine before me, saving me a lot of trouble along the way. Nicola Dinsdale, who came to my house with a thesis writing survival kit. Istvan Huszar, Amy Howard, and Rick Lange, who were in the same boat as me. The four of us came from the ONBI programme and joined FMRIB in 2017, we did this journey together, thank you for the ride.

I also had tremendous support from friends from outside the office. The people in my ONBI cohort that after the first year of studying together were scattered across the University in different departments as well as further away in Nottingham: Anna Booth, Aurelien Barbiton, Ellie Thompson, Alex Daniel, Julius Ossenberg-Engels, Helen Theissen, Jingwen Xiao, Liliana Barbieri, Lorenzo Venturini, and Nick Hall. The friends I made through rowing with Wolfson College Boat Club, especially the fierce women that got us to Head of the River in Summer Eights 2019: Jess Dunham, Jill Betts, Kate Erickson, Bea Faleri, Ellie Watts, Miriam Stricker, Rachel Anderson, and Claire Holubowskyj, as well as our coach Thomas Chapman. My boardgame gang, who kept going online even during the pandemic: Matt Gallagher, Jenny Gallagher, Sarah Fynes-Clinton, Rick Lange (again), Alex Smith, and Allison Scott. And also old friends in Finland, who, despite it having been many years since I lived there, keep in touch and meet up with me whenever I visit Helsinki. Special thanks go to Lotta Huikuri and Matti Fontell.

I also owe much to the one person, who thought I was worth moving to Oxford for, Edd Haigh. He left his life in London in order to join me in Oxford to do this degree, and I never thanked him enough for that. He is one of my biggest sources of inspiration, and has taught me so much about what success means and how the "standard path" is not the only path. He makes me better. He also contributed to the writing of this thesis by feeding me copious amounts of ice cream in the final weeks. Additionally, he finally convinced me to extend our little family, and in June we welcomed the puppy Sudo. Despite the stress of looking after a puppy during thesis writing, she has brought me much joy, and forced me to go outside and get some fresh air, which definitely kept me sane in the final weeks of writing this thesis.

Finally, the people I owe everything to is my family: Mamma, Pappa, and my brother Björn. They are the people who were around during my childhood. They are the people who shaped me into the person I am today. They are the ones who party the hardest when I succeed. This thesis would not exist if it wasn’t for my parents support and encouragement for me to constantly keep challenging myself. I am looking forward to many more celebrations with them.

Abstract

Magnetic resonance imaging (MRI) is a notoriously slow imaging method compared with other imaging modalities, for example computed tomography. It can, however, produce contrasts and attain biophysical information that is unattainable by other methods and is thus a valuable clinical tool.

In recent years, image acquisition times have been reduced through image reconstruction methods that require less data than traditional methods. In this thesis, these kinds of methods are used and extended upon for improved acquisition speed of an MRI method that is particularly slow due to requiring multiple encoded acquisitions to produce a single composite image: arterial spin labelling (ASL).

In conventional ASL, two encodings are used; one with "labelled" blood and one without "labelling", such that subtracting one from the other gives an image of just the blood signal. Angiography can be used to visualise blood flow through the arteries, and perfusion imaging to assess oxygen and nutrient supply to the tissue. Some advanced ASL methods require even more encodings that can be decoded to reveal more information about the cerebral haemodynamics. An example of such an advanced ASL method is vessel-encoded ASL, which allows for generation of separate images of blood originating from different arteries.

In this thesis, modern MRI sampling and reconstruction methods are optimised for vessel-encoded ASL and other ASL variants, with the aim of bringing these modalities towards clinically feasible scan times, which would eventually allow for more information-rich assessments of the cerebrovasculature and the perfusion state of the brain in, for example, patients suffering from stroke, dementia, and arteriovenous malformations.

By careful joint consideration of the multi-dimensional data, the data acquisition, and the reconstruction, very high acceleration factors can be achieved. This is demonstrated, first for vessel-encoded ASL angiography in 2D and 3D, then in similar advanced ASL methods (time-encoded ASL angiography, and combined angiography and perfusion imaging). A new radial sampling scheme is also presented and assessed on ASL angiographic data, that could have impact on imaging methods beyond ASL, in particular other dynamic MRI modalities.

Contents

List of Figures **vii**

List of Abbreviations **ix**

1 Introduction **1**

1.1 Motivation 1

1.2 Structure 3

2 Background **5**

2.1 Introduction 5

2.2 Magnetic Resonance Imaging 6

2.3 Data Sampling and Image Reconstruction 15

2.4 Quantifying Image Quality 31

2.5 Imaging Cerebral Haemodynamics 35

2.6 Summary and Conclusion 42

3 Feasibility of Accelerating Vessel-Encoded ASL Angiography **44**

3.1 Introduction 45

3.2 Methods 46

3.3 Results 57

3.4 Discussion 66

3.5 Conclusion 70

4 Bringing VE-ASL Angiography to Higher Dimensions **71**

4.1 Introduction 71

4.2 Improved Sampling Efficiency 73

4.3 Improved Sparsity 83

4.4 More Precise Regularisation Terms 90

4.5 Combining the Extensions to Acquire 4D VE-ASL Angiography . . 95

4.6 Discussion 102

4.7 Conclusion 106

5 Applications beyond VE-ASL angiography 107
5.1 Introduction 107
5.2 Time-encoded ASL angiography 109
5.3 Combined Angiography and Perfusion ASL 117
5.4 Discussion 122
5.5 Conclusion 125

6 Efficient Radial Sampling Beyond the Golden Angle Method 126
6.1 Introduction 126
6.2 Theory 128
6.3 Methods 135
6.4 Results 141
6.5 Discussion 145
6.6 Conclusion 150

7 Summary and Future Work 152
7.1 Summary 152
7.2 Directions of Future Work 156
7.3 Final Remarks 161

Appendices

A Toeplitz Embedding 164

B SNR Optimal Weighting Factors 166

References 168

List of Figures

2.1 An RF pulse can tip the magnetisation out of equilibrium along $\mathbf{B}_0$. 9
2.2 The simplest MR experiment. Free induction decay with no localisation 10
2.3 Without localisation only the total signal can be measured. 12
2.4 How a gradient field gives rise to a spatially varying phase pattern. 15
2.5 Sampling and windowing changes the resolution and field of view. . 20
2.6 The effect of multichannel acquisitios with different trajectories. . . 23
2.7 The iterative step of ISTA. 26
2.8 A koosh ball and a koosh ball trajectory. 27
2.9 Golden means method of sampling in 2D and 3D 29
2.10 SNR affects the ability to distinguish a features in a signal. 31
2.11 Flow driven adiabatic inversion. 40
2.12 Left:CASL, Right: pCASL. 40
2.13 VE-pCASL. 41

3.1 Simple VE numerical phantom 48
3.2 SNR comparison between in vivo and simulations. 50
3.3 pCASL angiography acquisition diagram 53
3.4 Vessel masks used in feasibility study 56
3.5 Comparison of non-VE and VE image quality in a numerical phantom 57
3.6 Comparison of image quality of VE and non-VE in a realistic simulation 58
3.7 Optimal regularisation overlap between subjects 59
3.8 The qualitatative effect of varying regularisation factors 60
3.9 Value of each regularisation term in feasibility study 62
3.10 Quality of non-VE and VE reconstructions 63
3.11 Temporal dynamics of non-VE and VE reconstructions - example 1 64
3.12 Temporal dynamics of non-VE and VE reconstructions - example 2 65
3.13 High resolution accelerated VE-ASL angiogram 65

4.1 Effect of sampling and aliasing patterns on VE-ASL signal decoding. 76
4.2 The three encoding and trajectory combinations that were trialled. 79
4.3 Comparison of the trialled reconstruction and sampling methods. . 81
4.4 Examples of the different approaches to samplign and reconstruction. 81

4.5 Example m-PSFs. 82
4.6 Effect on the m-PSF of acquiring the same spokes in different encodings. 83
4.7 Schematic for patch based sparsification method. 86
4.8 Sampling patterns used for simulated acqisitions. 88
4.9 Reconstruction time and quality for varying number of patches. . . 89
4.10 Image quality with different sampling and sparsification approaches. 90
4.11 Temporal basis function generation. 93
4.12 Image quality with and without the use of a spatial prior. 94
4.13 Image quality with and without the use of temporal regularisation . 95
4.14 Spoke ordering in simulations. 97
4.15 Ordering of spokes for the hybrid method in a multishot sequence. . 98
4.16 Image quality in 4D simulations. 101
4.17 Image quality in 4D in vivo experiments. 103
4.18 Different views of the 4D acquisition (5 minute acquisition). 104

5.1 Three different approaches to achieveing dynamic ASL imaging. . . 110
5.2 TEnc reconstructions at R = 12. 113
5.3 Line graph through two small vessels. 114
5.4 Frame-by-frame TEnc reconstructions. 115
5.5 TEnc reconstruction quality at varying acceleration factors. 116
5.6 Regularisation factor grid search for CAPRIA. 119
5.7 CAPRIA reconstructions for different reconstruction methods. . . . 120
5.8 Qualitative result of reconstructing with different amounts of data. 121
5.9 Quantitative result of reconstructing with different amounts of data. 121
5.10 Example frames from a 50 s (7 preps.) CAPRIA acquisition 123

6.1 Properties of set increment sampling 130
6.2 Monte Carlo simulations to measure SNR using different reconstruction methods. 134
6.3 Dynamic phantom used for SILVER experiments. 137
6.4 Spoke ordering used in SILVER in vivo experiments. 140
6.5 Vessel masks 141
6.6 Uniformity in different window sizes for SILVER and GR. 142
6.7 SILVER optimised for continuous ranges or discrete sets 143
6.8 SNR predictions and measurements in dynamic phantom. 144
6.9 SNR measurements across time. 145
6.10 Noise amplification in multi-coil SILVER experiments. 146
6.11 In vivo SNR measurements. 147
6.12 In vivo uniform, GR, and SILVER reconstructions. 148

List of Abbreviations

1D, 2D, ND	One-, two-, or N-dimensional.
ASL	Arterial spin labelling.
BA	Basilar artery.
CAPRIA	Combined angiography and perfusion using radial imaging and arterial spin labeling.
CASL	Continuous arterial spin labelling.
CS	Compressed sensing.
CT	Computed tomography.
DCE	Dynamic contrast enhanced (MRI).
DFT	Discrete Fourier transform.
DSA	Digital subtraction angiography.
FA	Flip angle.
FFT	Fast Fourier transform.
FISTA	Fast iterative soft thresholding algorithm.
FOV	Field of view.
GB	Gigabyte
GR	Golden ratio.
GRAPPA	Generalized autocalibrating partial parallel acquisition.
LICA	Left internal carotid artery.
MIP	Maximum intensity projection.
m-PSF	Multidimensional point spread function.
MR	Magnetic resonance.
MRI	Magnetic resonance imaging.
NMR	Nuclear magnetic resonance.
NRMSE	Normalised root mean square error.

non-VE-ASL	Non vessel encoded arterial spin labelling.
NUFFT	Non-uniform fast Fourier transform.
PASL	Pulsed arterial spin labelling.
pCASL	Pseudo-continuous arterial spin labelling.
PI	Parallel imaging.
PLD	Post labelling delay.
PSF	Point spread function.
RAM	Random Access Memory
RF	Radio frequency.
RICA	Right internal carotid artery.
RMSE	Root mean square error.
ROI	Region of interest.
SD	Standard deviation.
SENSE	Sensitivity encoding.
SILVER	Set increment with limited views encoding ratio.
SNR	Signal to noise ratio.
SSIM	Structural similarity index.
T1, T2, T2*	Tissue dependent time constants that determines magnetic resonance signal decay and recovery rates.
TE	Echo time.
TEnc-ASL	Time encoded arterial spin labelling.
TI	Inversion time.
TOF	Time of flight.
TR	Repetition time.
VE-ASL	Vessel encoded arterial spin labelling.

The length of exposure (one minute in sunlight) is still too long for the portrait. It was fifteen minutes when I first began my work. Progress may continue. Life is short and progress is slow.

— Gabriel Lippmann on Colour Photography 1908

1

Introduction

Contents

1.1 Motivation **1**
1.2 Structure **3**

1.1 Motivation

Just like photography during the 20th century evolved from a slow and cumbersome method of imaging to a fast, portable, and readily available method, magnetic resonance imaging (MRI) methods have been evolving rapidly during the 21st century. Some of the greatest leaps have been in acquisition speed thanks to revolutions such as parallel imaging (PI) [1, 2], and compressed sensing (CS) [3, 4]. These methods rely on inserting prior knowledge about the imaging system or imaged object into the reconstruction algorithm. In PI this additional information comes in the form of knowledge of the coil sensitivity profiles from pre-scans or calibration data. In CS, the information is instead regarding the object's sparsity in some domain. The insertion of external information allows for reconstruction with fewer sampled data points, and thus faster acquisition.

Similarly, like photography has improved going from black-and-white to colour, MRI methods are constantly improving to include more information about the object. Often, however, this extra information is gained at the cost of increased acquisition time as more encodings of the property of interest are needed. One such method, and the main focus of this thesis, is dynamic vessel-encoded arterial spin labelling (VE-ASL [5]). This imaging method provides detailed information about blood flow in the brain, separating blood from different feeding arteries, and thus allowing for more detailed cerebral haemodynamic investigations. Time-resolved and vessel-selective imaging can be useful, for example, in assessment of stroke [6] and complicated blood supply such as in arteriovenous malformations and other lesions [7]. Because VE-ASL is completely non-invasive, and requires no exogenous contrast agents or use of ionising radiation, it is safer than other methods for studying the vasculature and perfusion of the brain. VE-ASL is, however, not yet commonly used in clinical settings. One of the main drawbacks of VE-ASL is its long acquisition time, both in comparison to other angiographic and perfusion imaging methods and simpler non-vessel-encoded ASL (non-VE-ASL), the "black-and-white" version of this versatile imaging method.

One reason to hypothesise that VE-ASL can be accelerated, is that VE-ASL acquisition requires multiple encoded images that share a lot of information. By elimination of redundancies in the acquisition, and instead sharing of information across the encodings, each individual encoding could potentially be accelerated.

Angiographic methods are also known to be well suited to CS reconstruction because of their inherent sparsity in image space. Sparsity based methods could therefore be applied to VE-ASL angiography to bring it one step closer to clinical use. Although CS reconstruction has been applied in similar settings, for example to non-VE-ASL angiography [8] and time-of-flight (TOF) angiography [9], prior to the work presented in this thesis, CS had not been combined with VE-ASL.

The aim for this thesis is thus to fill this gap, to: i) investigate the technical hurdles and opportunities of applying non-linear reconstruction methods and efficient, information sharing sampling methods to VE-ASL angiography, ii) explore the potential benefits and limitations of these types of reconstruction and sampling methods, and iii) extend and generalise the methods developed for VE-ASL angiography to other MRI methods.

1.2 Structure

This thesis is organised into seven chapters. Following this short introductory chapter, the second chapter lays out the background information needed to put this thesis into context and highlights relevant previous literature. Chapters 3 to 6 are research chapters describing both theoretical, simulation, and experimental work that has been undertaken as part of this doctoral project.

The first research chapter, 3, describes a feasibility study on combining VE-ASL angiography with a CS reconstruction and highlights both challenges and opportunities related to combining the two. The content of chapter 3 has been published as a full paper in *Magnetic Resonance in Medicine* [10].

In chapter 4 the ideas developed in chapter 3 are extended upon and new ideas relating to encoding-sampling interactions are introduced. In this chapter, the acceleration factors of VE-ASL angiography are pushed to their limits to produce high quality 4D vessel-encoded angiograms in clinically feasible scan times. The contents in this chapter have been published as various abstracts presented at the *International Society for Magnetic Resonance in Medicine*'s (*ISMRM*) annual meetings [11–13].

In chapters 5 and 6, the developed framework is further generalised. In chapter 5 the reconstruction framework is applied to MRI methods beyond just VE-ASL angiography. Here both time-encoded ASL angiography and combined angiography and perfusion ASL are considered. The section on time-encoded ASL has been presented at the ISMRM annual meeting [14]. In chapter 6 a novel radial sampling scheme that improves upon golden angle sampling for dynamic MRI is presented and evaluated. The framework is called SILVER (the Set Increment with Limited Views Encoding Ratio) and improves sampling efficiency for dynamic acquisitions that require flexibility of reconstructing multiple different temporal resolutions. It is evaluated on non-VE-ASL angiography.

The thesis is wrapped up in chapter 7 with a summary and discussion about the ideas presented in the research chapters. Ideas for further work and research directions are also presented.

An image of an object may be defined as a graphical representation of the spatial distribution of one or more of its properties.

— Paul Lauterbur on Magnetic Resonance Imaging
1973

2
Background

Contents

2.1 Introduction **5**
2.2 Magnetic Resonance Imaging **6**
  2.2.1 Signal Generation 6
  2.2.2 Localisation 11
2.3 Data Sampling and Image Reconstruction **15**
  2.3.1 The Point Spread Function 18
  2.3.2 Parallel Imaging 21
  2.3.3 Compressed Sensing 22
  2.3.4 Sampling Trajectories 25
2.4 Quantifying Image Quality **31**
  2.4.1 Image SNR 31
  2.4.2 Normalised Root Mean Square Error 32
  2.4.3 Structural Similarity Index 33
  2.4.4 Correlation 34
2.5 Imaging Cerebral Haemodynamics **35**
  2.5.1 Non-MR Methods 35
  2.5.2 MR methods 36
2.6 Summary and Conclusion **42**

2.1 Introduction

This chapter outlines the underlying physics and mathematics of magnetic resonance imaging, the signal processing required to reconstruct MRI, ways of analysing image quality, and common methods for imaging cerebral haemodynamics. This background should be sufficient to put this thesis into context and guide the reader through the reasoning of the following chapters.

2.2 Magnetic Resonance Imaging

Magnetic resonance imaging (MRI) was first described as an extension to nuclear magnetic resonance (NMR), by which the NMR signal could be localised in space by the addition of magnetic field gradients [15]. But what is the NMR signal, and how is it generated? What is meant by addition of magnetic field gradients, and how does that localise the signal? The next two sections will aim to answer these questions briefly. For a more thorough description the interested reader may consult, for example Bernstein [16] or Nishimura [17], from which most of the material in this section is drawn.

2.2.1 Signal Generation

The origin of the MR signal is the quantum mechanical property spin. Any atomic nucleus that has an odd number of protons and/or neutrons has non-zero spin and can exhibit magnetic resonance. The most common atom to excite using MR is ^{1}H, because of its abundance in the human body, mainly as a component of H_2O. Spin gives rise to a magnetic dipole moment in charged elemental particles (neutrons have a magnetic dipole moment because they are made up of charged quarks). The magnetic moment follows the equation:

$$\boldsymbol{\mu} = \gamma \mathbf{S} \tag{2.1}$$

where $\boldsymbol{\mu}$ is the magnetic moment vector of a single nucleus, γ is a scaling factor called the gyromagnetic ratio that depends on the species of nucleus, and $\mathbf{S}$ is the nuclear spin angular momentum operator that is quantised. In the presence of a magnetic field, a hydrogen nucleus has two spin eigenstates with different energy levels. This phenomenon is known as Zeeman splitting. The spin state parallel to the field has slightly lower energy than the anti-parallel state. The difference in energy, ΔE, between the parallel and anti-parallel spin state depends on the external magnetic field strength, B, the gyromagnetic ratio of the species of nucleus, γ (267.5×10^6 rad s^{-1} T^{-1} for ^1H), and $\hbar$, which is the Planck constant divided by 2π ($1.054571817 \times 10^{-34}$ J s), one of the fundamental constants of physics.

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

This energy difference results in two populations forming based on the Boltzmann distribution

$$\frac{n_{parallel}}{n_{anti-parallel}} = e^{\Delta E/kT} \tag{2.3}$$

Where $n_{parallel}$ and $n_{anti-parallel}$ are the proportion of spins that would be found in the parallel and anti-parallel state respectively if measured one-by-one, k is the Boltzmann constant (1.38065×10^{-23} J K^{-1}), and T is the temperature in units of Kelvin. In a typical MR experiment, where T is body temperature, 37°C (310.15 K), the nucleus in question is ^1H, and the magnetic field strength is 3 Tesla, the fraction of spins in the parallel state ($n_{parallel}$) outnumber the anti-parallel by approximately 0.002%. This small difference in probability of finding one state over the other results in a measurable net magnetic moment, **M**, when a large number of nuclei (millions or more) are measured simultaneously. Most behaviour of the net magnetic moment in the presence of both static and time-varying magnetic fields can, however, be considered using classical physics by equating the net magnetisation with a magnetic dipole.

At equilibrium, **M** will be aligned with the applied external magnetic field, $\mathbf{B}_0$. By convention, the direction of $\mathbf{B}_0$ is chosen to be the z-direction in the coordinate system, and the component of vectors along this direction are called longitudinal components. Vector components perpendicular to this direction are called transverse components. If the system conditions are changed such that the magnetic moment vector points in a different direction than $\mathbf{B}_0$, according to classical mechanics, a torque is applied that changes the angular momentum, which is proportional to the magnetisation as per equation 2.1. The vector cross product between the magnetisation and the external magnetic field scaled by the gyromagnetic ratio gives the rate of change of the magnetisation vector:

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma \mathbf{B}_0 \tag{2.4}$$

By solving the differential equation, 2.4, it is found that $\mathbf{M}$ must precess around $\mathbf{B}_0$ at an angular frequency, ω_0:

$$\omega_0 = \gamma |\mathbf{B}_0| \tag{2.5}$$

The relationship expressed in Eq. 2.5 is called the Larmor equation and ω_0 is similarly called the Larmor frequency. This is the resonant frequency of the system, and photons with this frequency have the same energy as the energy difference of the two quantum eigenstates (equation 2.2) and can thus interact with the magnetised nuclei.

If an additional magnetic field, $\mathbf{B}_1$, is added to the system, the magnetisation $\mathbf{M}$ will follow equation 2.4 and start precessing around the new net magnetic field ($\mathbf{B} = \mathbf{B}_0 + \mathbf{B}_1$). If $\mathbf{B}_1$ is the magnetic component of an electromagnetic radiofrequency (RF) field applied such that $\mathbf{B}_1$ is perpendicular to $\mathbf{B}_0$ and circularly polarised such that it rotates around $\mathbf{B}_0$ at the Larmor frequency, resonance will occur, such that the magnetisation will simultaneously rotate around $\mathbf{B}_0$ and $\mathbf{B}_1$ and spiral out of the equilibrium direction along $\mathbf{B}_0$ as shown in figure 2.1.

The duration and intensity of the RF pulse determines the flip angle (FA), how far out of equilibrium the magnetisation is rotated from its equilibrium direction along $\mathbf{B}_0$. The MR signal that can be detected arises from the time-varying magnetic field caused by the precessing magnetisation, $\mathbf{M}$, that induces a voltage in a receive coil. The coil is only sensitive to the transverse component of $\mathbf{M}$ because it is tuned to time-varying magnetic fields changing at the Larmor frequency of the system. The transverse component of $\mathbf{M}$ can be expressed as either a vector $\mathbf{M}_{xy}$, or a complex number m:

$$M_{xy} e^{i\omega_0 t} = m \tag{2.6}$$

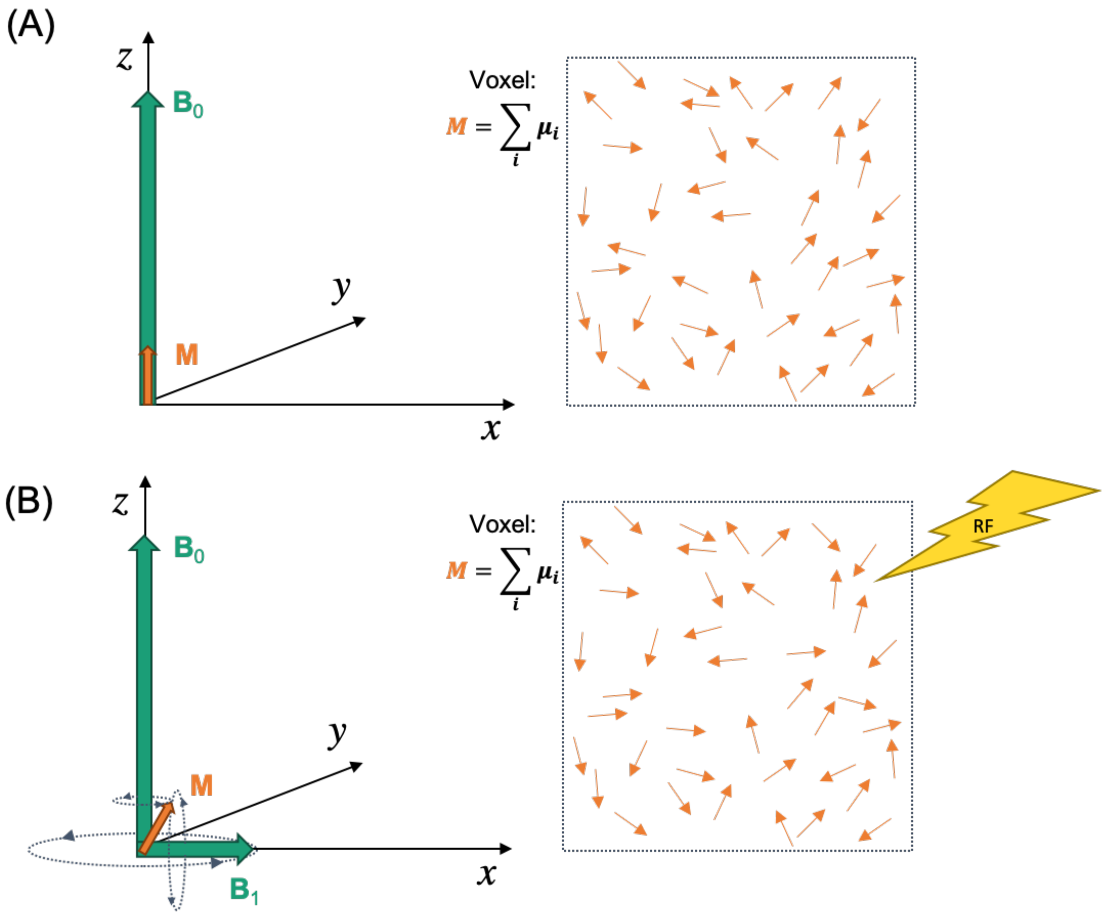

Figure 2.1: An RF pulse at the Larmor frequency can tip the magnetisation out of equilibrium (A) along $\mathbf{B}_0$ by rotating $\mathbf{M}$ around the net magnetic field $\mathbf{B}_0 + \mathbf{B}_1$ (B). On the level of individual spins, the RF pulse only changes the probability of finding them in the anti-parallel state by a very small amount.

If we consider the system in a frame of reference that is rotating at the Larmor frequency, ω_0, we can drop the $e^{i\omega_0 t}$ factor above. Unless the local magnetic field strength is changed the spins and the frame rotate at the same frequency. Similarly, the spiralling out of equilibrium when an RF pulse is applied, can in the rotating frame be viewed as a simple tipping of the magnetisation vector. From now on we shall consider the system in the rotating frame.

After excitation, the magnetisation will eventually return to an equilibrium state aligned with the main magnetic field through various processes of relaxation that will be discussed in more detail below. A simple MR experiment could look something like what is shown in figure 2.2, with a simple RF pulse causing $\mathbf{M}$ to be

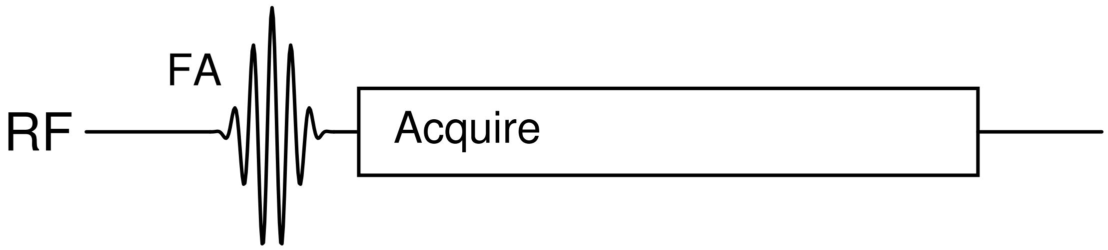

Figure 2.2: A simple MR experiment. An RF pulse at the Larmor frequency tips the magnetisation by some flip angle (FA) and the receive coil starts picking up a signal due to the resulting precessing magnetisation. Eventually the signal dies out as the magnetisation returns to equilibrium. No localisation of the signal is achieved.

tipped by some FA and the signal then being measured as it returns to equilibrium.

The process of relaxation and return to equilibrium is explained by two separate processes, spin-lattice, and spin-spin-relaxation. Spin-lattice relaxation, also called T1 relaxation, governs the regrowth of M_z, the longitudinal component of $\mathbf{M}$. Spin-spin relaxation, also called T2 relaxation, on the other hand, governs the loss of the transverse component, M_{xy}, as the spins in the ensemble lose rotational coherence. Rotational coherence is also lost due to magnetic field inhomogeneities in the sample as some spins experience a different local magnetic field, and thus precess with a different angular frequency from the pure Larmor frequency. This causes the signal to decay faster than it would with pure T2 relaxation. This type of signal decay that combines the effects of spin-spin relaxation and local field inhomogeneities is called T2* relaxation.

The measured complex signal intensity thus depends on the volume and density of excited nuclei as well as time since excitation, and intrinsic T1 and T2(*) values of the tissue, along with the flip angle of the excitation pulse. T1 and T2(*) vary between different tissue types and in imaging, the sequence timings are often designed to create contrast between tissues with different relaxation properties (e.g. white and grey matter in the brain). Not all MR contrast is, however, generated through these relaxation time constants. For example, in arterial spin labelling that is presented further in section 2.5.2, the blood signal is modified prior to image acquisition to generate contrast between blood and static tissue. For this type of imaging it is worth sampling the signal as soon as possible after excitation to get as high signal as possible before the spins relax and the magnetisation returns to equilibrium. Sampling the signal later causes lower signal-to-noise ratios (SNR), which is a fundamental concept in MRI, and will be examined more thoroughly in section 2.4.1.

2.2.2 Localisation

But what is the value of achieving contrast between different tissues if you cannot separate them in space? Without localisation of signal there are no images. The signal, s, that is measured by the scanner is simply the integral of all transverse components of all spins in space weighted by the receive coil's sensitivity profile. The signal recorded at a specific time is then:

$$s(t) \propto \int_x \int_y \int_z m(x, y, z, t) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}z \tag{2.7}$$

Where $m(x, y, z, t)$ is the transverse component of the magnetic moment vector $\mathbf{M}$ as presented in equation 2.6, and it is a complex quantity. The constant of proportionality, omitted above, can be derived from Faraday's Law of induction [18]. Here, without the addition of spatially or temporally varying magnetic fields, $m(x, y, z, t)$ changes as a function of time only due to the relaxation processes described above. Figure 2.3(A) shows the spatial distribution of magnetisation at some time t, and figure 2.3(B) shows the signal recorded by the scanner. Without any localisation measures the magnetisation in each voxel can only be assumed to be the average of the object's total magnetisation.

There are two different ways of localising the signal, spatially selective excitation, and phase encoding. Both methods rely on the addition of spatially varying magnetic fields, most commonly linear gradient fields, that can be turned on and off. Linear gradient fields causes the local field strength to vary linearly along one dimension such that the precession frequency varies linearly across that dimension.

$$\omega(x, y, z) = \gamma(B_0 + G_x x + G_y y + G_z z) \tag{2.8}$$

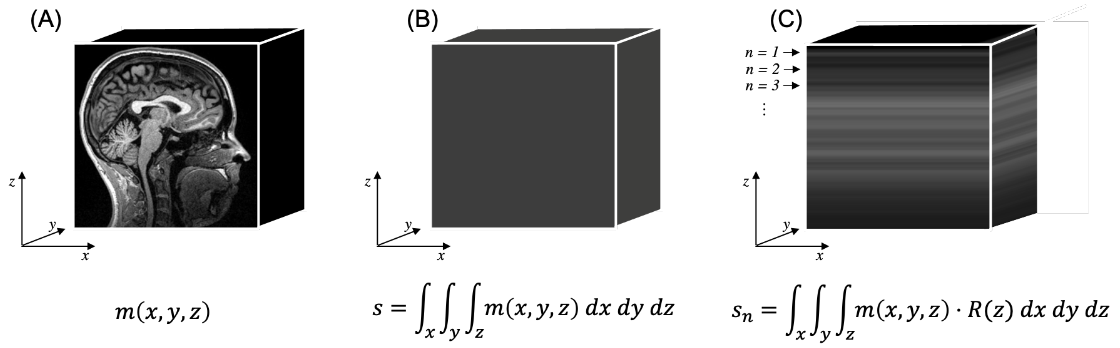

Figure 2.3: Without localisation of the signal, (A), only the total signal within a volume can be measured, (B). Slice selection, (C), can tell the average signal in a plane, but more is needed to encode the signal in 3D.

Here, $G_{x,y,z}$ is the gradient field strength along the x, y, or z dimension respectively. $G_{x,y,z}$ is measured in mT/m and most systems have gradient coils that can produce gradient fields with amplitudes on the order of 50 mT/m [17]. Note that the direction of the magnetic field vector of the gradient fields is always parallel or anti-parallel to the main magnetic field. In the rotating frame the B_0 term in equation 2.8 can be omitted.

Spatially Selective Excitation

The first method for localising the signal is to simply only excite signal in a limited portion of space, then one can be certain that the measured signal came from there. The simplest, and most common form of this type of localisation is called slice-selection, but more controlled spatial excitation is possible using parallel transmission [19] and/or specially designed RF pulses [20]. However, a review of these methods is beyond the scope of this thesis.

Slice selection is done by application of a gradient during excitation with a band-limited RF pulse. Because different locations along the gradient direction have different resonant frequencies, only spins at locations that correspond to the frequency bandwidth of the RF pulse will experience resonance and get excited.

This corresponds to a slice (or slab) of tissue perpendicular to the applied gradient direction. The ideal slice selection pulse would excite a slice with equal flip angle everywhere and a sharp cut-off to locations outside the slice (a box function). To do this, however, an infinitely long RF pulse would be required. In the small flip angle regime, the frequency content of an RF pulse is determined by the Fourier transform of its amplitude envelope, and the Fourier transform of a box function is a sinc function. Instead of an infinitely long sinc function, an apodised sinc function with only a small number of side lobes is commonly used.

If slice selection is used along the z-axis, the measured signal equation will instead be:

$$s(t) \propto \int_x \int_y \int_z m(x, y, z, t) \cdot R(z) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}z \tag{2.9}$$

where $R(z)$ is a function describing the shape of the slice profile along z. By acquiring data with different slice locations with different central frequencies of the excitation RF pulses, a 1D representation of the object can be built up slice-by-slice (figure 2.3(C)). Without additional encoding within the slices, this is not enough to adequately represent a three dimensional object.

Phase Encoding

The other method of localisation that can be used either together with slice-selection or without, is phase encoding. If phase encoding is used in combination with slice selection of thin slices, it is often sufficient to do a 2D encoding and assume that the signal across the slice select direction is constant within a slice. Phase encoding can also be performed as a full 3D encoding with or without spatially selective excitation. In phase encoding, multiple measurements are made with different 2D or 3D phase patterns imposed on the precessing magnetisation and the measurements are combined to separate signal coming from different spatial locations.

As above, by application of linear gradient fields the resonant frequency will vary linearly along the direction of the gradient. However, in phase encoding, the gradient is applied after excitation of the signal and thus the spatially varying resonant frequencies will result in different amounts of phase accrual, ϕ, of the transverse component of $\mathbf{M}$.

$$\phi(x,y,z,t) = \int_0^t \omega(x,y,z)\,\mathrm{d}\tau \tag{2.10}$$

which combined with equation 2.8, this becomes:

$$\phi(x,y,z,t) = \gamma \int_0^t (B_0 + G_x x + G_y y + G_z z)\,\mathrm{d}\tau \tag{2.11}$$

again, in the rotating frame the B_0 term can be omitted.

Over time sinusoidal phase patterns with increasing spatial frequencies emerge along the gradient direction (figure 2.4). When the gradient field is turned off, the accrued spatially varying phase remains although the frequency returns to the Larmor frequency, ω_0. The signal measured by the system can now be written:

$$s(t) \propto \int_x \int_y \int_z m(x,y,z,t) \cdot e^{-\gamma i \int_0^t (G_x(\tau)x + G_y(\tau)y + G_z(\tau)z)\,\mathrm{d}\tau}\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z \tag{2.12}$$

where $G_{x,y,z}(\tau)$ is the time course of the gradient fields. The time integral of $G_{x,y,z}$ governs the spatial frequency of the accrued sinusoidal phase pattern and can also be interpreted as a location in a 3D space [21, 22], normally referred to as k-space. With the k-space formalism, equation 2.12 can be rewritten as:

$$s(t) \propto \int_x \int_y \int_z m(x,y,z,t) \cdot e^{-i2\pi[k_x(t)x + k_y(t)y + k_z(t)z]}\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z \tag{2.13}$$

with $k_x = \frac{\gamma}{2\pi} \int_0^t G_x(\tau)\,\mathrm{d}\tau$, referring to k-space k_x coordinate. Similarly k_y and k_z are the other k-space coordinates. The value of k is the number of cycles per unit length the phase pattern has in each dimension.

The form of equation 2.13 is that of a 3D Fourier transform. If slice selection is used, we can ignore one of the dimensions (usually the z-dimension) and only consider the 2D problem.

Sampling a point in k-space is thus sampling of one Fourier component of the spatial distribution of m. If the signal is sampled continuously whilst a gradient

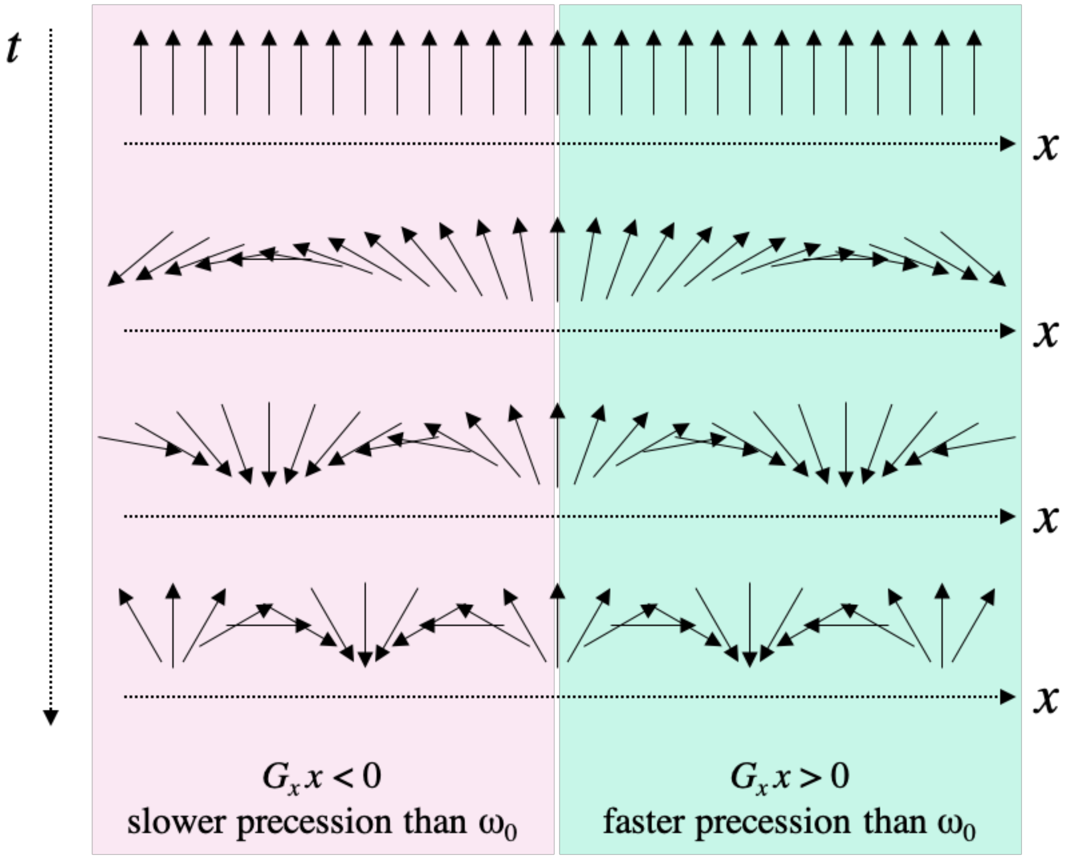

Figure 2.4: How a gradient field gives rise to a spatially varying phase pattern. The longer the gradient is on for, or the stronger it is, higher frequency patterns will emerge.

field is turned on, a trajectory through k-space can be traced. If enough samples are collected the Fourier transform can be inverted and the image reconstructed. Further discussion on this point can be found in the following section.

2.3 Data Sampling and Image Reconstruction

In the previous section we saw how the addition of linear gradient fields can turn the MRI scanner into a physical Fourier transform operation and alluded to how that allows us to localise the signal in space. In this section we generalise the concept of localisation and show how localisation of MRI signal, and thus image reconstruction, can be considered an inverse problem. This allows us to use more efficient tools from the field of signal processing to reconstruct the signal into images. We also move from a continuous description of the signal to a discrete one since

MRI scanners sample the signal at discrete time points whilst moving along the k-space trajectory. The object $m(x, y, z, t)$ is a continuous property and so is $s(t)$, but we can only record a sampled version of $s(t)$. From now on we shall call the sampled signal $\mathbf{s}$, with each element of the vector representing a time point (or equivalently a k-space location). We will also assume that the object, m, is finite and discrete, such that it can be represented by a vector $\mathbf{m}$.

To formalise image reconstruction as a discrete inverse problem, where the aim is to reconstruct a number voxels that map out a property of interest (most commonly the spatio-temporal MR signal distribution), the forward problem is formulated:

$$\mathbf{s} = \mathbf{E}\mathbf{m} + \mathbf{n} \tag{2.14}$$

Where $\mathbf{E}$ is a linear operator that models the effect the imaging system has on the property of interest. In general, this operator includes the sampling of the magnetisation's Fourier components (k-space trajectory), and spatial modulation by multiple receive coil sensitivities. The coil sensitivities are complex maps of space that define the receive coil's weighting of magnetisation in each point in space. Physically, they are governed by the Biot-Savart law [23], but in practice they are mostly experimentally measured. $\mathbf{n}$ is a vector containing the system noise added to the measurements (in k-space). We will initially consider $\mathbf{n}$ to be complex white noise, but in reality, with crosstalk between the receive coils, and noise originating in randomly occurring current loops in the imaged sample, this noise can also be partially correlated between channels. Noise correlation affects, for example, parallel imaging and SNR. For simplicity, this initial discussion will exclude consideration of the multiple receive coils and assume a single coil with uniform sensitivity.

If the acquisition operator, $\mathbf{E}$, is invertible, for example when it is a 2D or 3D discrete Fourier transform (DFT), we can simply multiply the acquired data vector, $\mathbf{s}$, with its inverse, $\mathbf{E}^{-1}$ to produce a reconstruction of the image, $\mathbf{m}$. This does not take into account the effect of noise from $\mathbf{n}$, however, reconstruction with an orthogonal transform, such as the inverse DFT, results in Gaussian noise remaining Gaussian when transformed. In order for the acquisition operator to be a DFT, k-space samples have to be placed in a regular rectangular or cubic pattern on a Cartesian grid.

The linear system can also be over or under determined ($\mathbf{E}$ has unequal numbers of rows and columns). In these cases, the pseudo-inverse can be used instead of the inverse to create an image with a least-squares fit: $\mathbf{m} = (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{s}$. Here, the superscript H denotes the conjugate transpose of the matrix.

The form of the pseudo-inverse can be derived by starting from the forward model (equation 2.14) and left multiplying both sides by $\mathbf{E}^{\mathrm{H}}$, which gives

$$\mathbf{E}^{\mathrm{H}}(\mathbf{E}\mathbf{m} + \mathbf{n}) = \mathbf{E}^{\mathrm{H}}\mathbf{s} \tag{2.15}$$

Even though $\mathbf{E}$ is not a square matrix, $\mathbf{E}^{\mathrm{H}}\mathbf{E}$ is a square matrix, which can be inverted to isolate $\mathbf{m}$:

$$\mathbf{m} = (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}(\mathbf{s} - \mathbf{n}) \tag{2.16}$$

Again, assuming that the noise is Gaussian and the transform is orthogonal, the noise term remains Gaussian and an unbiased estimate of $\mathbf{m}$ can be reconstructed without knowledge of $\mathbf{n}$.

Solving the pseudo-inverse analytically can be computationally intensive as it includes multiple matrix multiplications and an inversion of a potentially large matrix. More efficient algorithms have been developed. The most important one in this context is the fast Fourier transform (FFT [24]), that can perform the operation of a DFT operator in a time proportional to $N \log N$, where N is the number of sampled points, rather than N^2 which a standard matrix multiplication would take. Iterative methods are also often used to avoid inverting $\mathbf{E}^{\mathrm{H}}\mathbf{E}$ directly.

The following sections will consider what happens when $\mathbf{E}$ is underdetermined (more columns than rows) and/or non-orthogonal. To do this, we will introduce the concept of a point spread function (PSF), which is a central concept in image reconstruction, and present some common linear and non-linear approaches to image reconstruction of underdetermined systems. In the final section some sampling trajectories central to this thesis are presented.

2.3.1 The Point Spread Function

The PSF, also called the system's impulse response, is a common way to characterise an imaging system. In MRI, it determines the intrinsic spatial resolution (ability to separate two point sources in space) of the imaging system, as well as the field of view, which is the maximum size the imaged object can be. In this section, we will limit the discussion to the effects of the sampling trajectory on the PSF, but this concept can be generalised to include more general models of the acquisition operator, and is further elaborated upon in chapter 4.

Let us first consider MR imaging sampled on a regular (Cartesian) grid. For simplicity, we will consider 1D sampling. The extension into 2D and 3D is trivial, as higher dimensional Fourier sampling can be achieved by sequential 1D operations along the different dimensions of the data. The fully sampled acquisition operator, $\mathbf{E}_{full}$ can thus be defined as the unitary DFT, $\mathbf{F}$:

$$\mathbf{F} = \frac{1}{\sqrt{N}} \begin{bmatrix} e^{-i2\pi\cdot(0\cdot0)/N} & e^{-i2\pi\cdot(0\cdot1)/N} & \dots & e^{-i2\pi\cdot(0\cdot(N-1))/N} \\ e^{-i2\pi\cdot(1\cdot0)/N} & e^{-i2\pi\cdot(1\cdot1)/N} & \dots & e^{-i2\pi\cdot(1\cdot(N-1))/N} \\ \vdots & \vdots & \ddots & \vdots \\ e^{-i2\pi\cdot((N-1)\cdot0)/N} & e^{-i2\pi\cdot((N-1)\cdot1)/N} & \dots & e^{-i2\pi\cdot((N-1)\cdot(N-1))/N} \end{bmatrix} \tag{2.17}$$

or equivalently

$$F_{i,j} = \frac{1}{\sqrt{N}} e^{-i2\pi\cdot(k_i\cdot x_j)/N} \tag{2.18}$$

with k_i and $x_j \in [0, 1, 2, ..., (N-1)]$. Each row in $\mathbf{F}$ maps out a phase pattern with increasing spatial frequency, and the matrix has N columns, corresponding to the length of the object vector, $\mathbf{m}$, that is modelled to be made up of N pixels. In reality, the object is a continuous function. If more pixels of the object are modelled (a longer vector) to better capture the continuous spatial variation of $\mathbf{m}$, the DFT's rows can be elongated to allow for the matrix multiplication, but the problem becomes under-determined as the number of columns now exceeds the number of rows unless the k_i values are extended too. If the problem is undersampled, the object might contain spatial frequencies that have not been sampled, and the information about their existence is fundamentally lost in the signal generation ($\mathbf{s}$

is a lower dimensional vector than $\mathbf{m}$). The PSF, $\mathbf{E}^{\mathrm{H}}\mathbf{E}$, can be used to visualise just what information has been lost. The effect of the PSF can be studied by taking $\mathbf{E}^{\mathrm{H}}\mathbf{E}\mathbf{m}_{\delta}$, with $\mathbf{m}_{\delta}$ being a vector of all zeros and a single element set to one. In the simple case of $\mathbf{E} = \mathbf{F}$, the PSF is shift invariant it does not matter which element is non-zero. In more elaborate models of the acquisition operator (e.g. including coil sensitivity maps or using non-linear gradients [25]) the PSF is no longer necessarily shift invariant.

In the Cartesian case, the extent of sampling (how high frequencies are included in the DFT) directly determines the resolution, and the spacing between samples directly determines the alias-free field-of-view, FOV, as demonstrated in figure 2.5. When the signal is sampled on a regular grid with even spacing the continuous Fourier transformed object can be thought of as being multiplied with a comb-function (a function that is 1 at the sample locations separated by Δk and 0 everywhere else). The Fourier transformed object can contain frequencies tending towards infinity, whereas sampling is always done up to some maximum frequency. This "windowing" can be modelled as a multiplication with a box-function that is 1 between $-k_{max}$ and $+k_{max}$ (here k is defined in physical units rather than an index as in equations 2.17 and 2.18) and zero elsewhere. Because of the Fourier convolution theorem we can see that the object will be convolved with a comb function with each delta function separated by $\frac{1}{\Delta k}$ as well as a sinc function (the Fourier transform of a box-function). Not sampling high enough frequencies thus leads to blurring as each point in image space is spread out with a sinc function and this sets the fundamental spatial resolution. Not sampling with small enough Δk steps instead leads to copies of the object being reconstructed on top of the object itself. The maximum FOV is thus defined by the distance between the comb function's non-zero points in image space.

The way imaging is sped up is by reducing the number of rows in the DFT (points in k-space). This can be done for Cartesian trajectories either by 'skipping lines' in k-space, which leads to larger gaps in k-space (along one dimension) and thus smaller a FOV, or acquiring data at a lower resolution, which leads to blurring. Other methods,

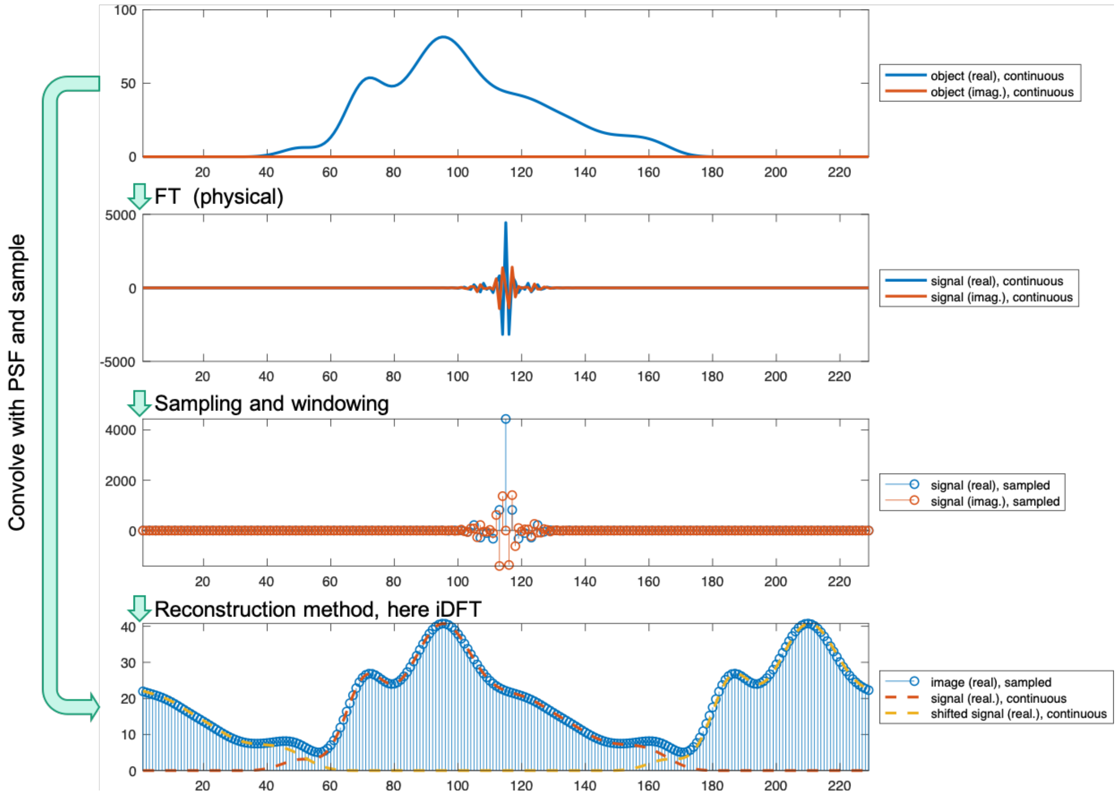

Figure 2.5: In MRI a signal is first transformed into its Fourier transform by physically imposing phase patterns across it. That Fourier transform is then sampled. Here, the sampling is done with slightly too large gaps between the samples, which causes a shifted copy of the signal to appear slightly overlapping with the original signal when reconstructed). The windowing (not sampling infinitely many Fourier components) limits the spatial resolution, but the effect cannot easily be seen in this example that is a very smoothly varying signal. The system can also be analysed directly by convolving the signal with the system's PSF and sampling.

such as partial Fourier sampling [26] have more subtle effects on the PSF that also lead to blurring of the images, unless special reconstruction approaches are utilised.

When samples are acquired off the Cartesian grid the acquisition operator changes such that instead of the regularly sampled DFT, **E** has elements:

$$E_{i,j} = \frac{1}{\sqrt{N}} e^{-i2\pi\cdot(k_i\cdot x_j)/N} \tag{2.19}$$

k_i and $x_j \in \mathbb{R}$ (unlike for the DFT where the elements were regularly spaced integers). The non-Cartesian samples can be 're-gridded' onto a Cartesian grid to allow for reconstruction with the FFT, the PSF can then also be analysed as if Cartesian sampling had been performed. Alternatively the acquisition and reconstruction can be modelled with $\mathbf{E}$ directly and the PSF can still be studied using the operator $\mathbf{E}^{\mathrm{H}}\mathbf{E}$.

2.3.2 Parallel Imaging

Multi-coil acquisitions allow for undersampling of k-space when the sensitivity maps provide additional spatial information to the system. This is referred to as parallel imaging (PI) and can be done in two main ways: One is to interpolate in k-space and synthesize the missing samples based on some calibration data (e.g. generalised autocalibrating partially parallel acquisitions, GRAPPA [2]). The other is to reconstruct the corrupted images for each channel separately and combine to recover the underlying signal using explicit knowledge of the coil sensitivity distributions in image space (sensitivity encoding, SENSE [1]). These are both linear reconstruction methods and are mathematically equivalent [27].

In this section, the focus will be on the second method, explicit use of coil sensitivity maps in image space, as that is what have been used throughout this thesis.

The purpose of SENSE-like reconstructions is to untangle aliased signal from multiple spatial sources that end up stacked on top of each other when k-space is undersampled. The coil sensitivity profiles weight the object with different spatial weights for each receive channel (figure 2.6(A)). The value in one pixel in one coil channel is a superposition of aliased signals as determined by the system's PSF and the coil's sensitivity profile. In Cartesian imaging, the result of regular undersampling is shifted copies of the object (figure 2.6(B)), which makes the 'untangling' problem somewhat easier as the number of mixed source locations is limited to R, the acceleration factor. On the other hand, in non-Cartesian imaging, the aliasing patterns are more complex, and energy from every pixel is likely to end up on top of every other pixel (figure 2.6(C)). Including the coil sensitivity maps into the forward model can make the problem well-posed and allow for alias-free reconstruction. Let's first consider the effect on the PSF, if data from each channel were reconstructed separately.

$$\mathrm{PSF}_c = \mathbf{E}_c^{\mathrm{H}}\mathbf{E}_c\mathbf{m}_\delta \tag{2.20}$$

with $\mathbf{E}_c = \mathbf{F}\mathbf{S}_c$, where $\mathbf{F}$ is the sampled Fourier operator, and $\mathbf{S}_c$ is the diagonal matrix that describes the coil sensitivity for the c^{th} coil. Here, $\mathbf{E}$ is of size $N_k \times N_v$ (same as $\mathbf{F}$, with N_k equal to the number of k-space samples, N_v equal to the number of voxels, and $N_k < N_v$). If, instead, each subproblem is stacked on top of each other ($\mathbf{E} = \begin{bmatrix} \mathbf{E}_1 \\ \mathbf{E}_2 \\ \vdots \\ \mathbf{E}_{N_c} \end{bmatrix}$), $\mathbf{E}$ changes shape to $N_k N_c \times N_v$, which makes the problem well-posed as long as $N_k N_c \geq N_v$, and all columns are linearly independent. The coils are likely to be linearly independent, but less likely to be orthogonal.

Non-orthogonality leads to poorer conditioning of the matrix and noise amplification. The amount of noise amplification is determined by the so-called geometry factor (g-factor) and is spatially varying. Usually, the g-factor is highest in the centre of the field of view where the coil sensitivity maps are the most similar. Noise correlation between coils also leads to noise amplification. If the system's noise correlation matrix is known (e.g. from an independent noise scan), the noise matrix can be included in the acquisition model and used to improve a SENSE reconstruction. Details on this can be found in the original SENSE paper [1].

To solve a SENSE type reconstruction problem, iterative solvers are often used as inversion of $\mathbf{E}^{\text{H}}\mathbf{E}$ is too computationally expensive for large matrices [28]. However, in the Cartesian case, where the number of overlapping voxels is limited to R, $\mathbf{E}^{\text{H}}\mathbf{E}$ will be extremely sparse and possible to break up into a much smaller unfolding matrices for each pixel that can be easily inverted [1].

2.3.3 Compressed Sensing

Compressed sensing is a method that can be used for reconstructing images from undersampled data, which relies on a-priori knowledge of an image's sparsity in some transform domain. This is possible when the data is undersampled in a incoherent pattern that causes the aliasing to have noise-like characteristics and thus making it possible to separate aliasing artifacts from the original signal. An image can be said to have a sparse representation if it can be represented by mainly zeros in some domain (and thus is compressible, hence the name compressed sensing). To find an

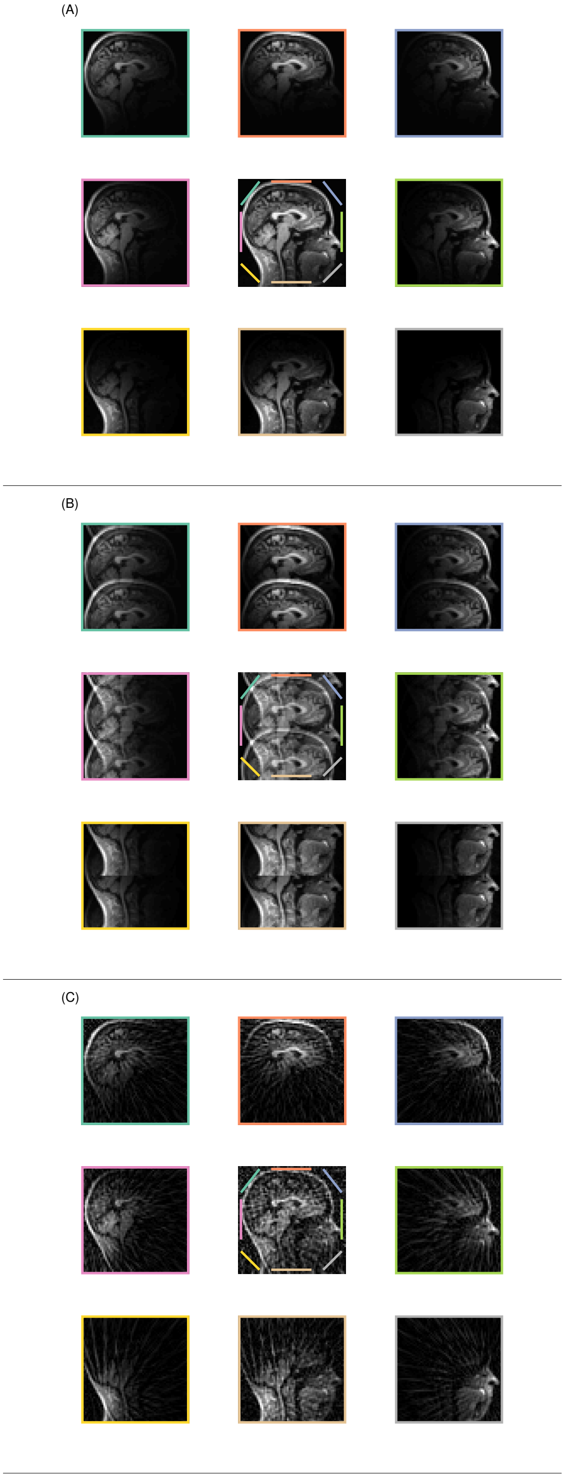

Figure 2.6: The image in the middle of each group shows what the image would look like reconstructed with a single uniform coil using a fully sampled acquisition, (A), an undersampled Cartesian acquisition, (B), or an undersampled radial acquisition, (C). The coloured lines on the middle images represent coil locations and the images around the central image shows what eight different coil sensitivities do to the image.

image that is both consistent with the data and sparse in some transform domain, a cost function, $f(\mathbf{x})$, is defined with the aim of minimising it to find the optimal fit.

$$f(\mathbf{x}) = \frac{1}{2}||\mathbf{Em} - \mathbf{s}||_2^2 + \lambda||\mathbf{Sm}||_1 \tag{2.21}$$

where $\mathbf{E}$ is the acquisition matrix, which describes the undersampled acquisition of the data in k-space, $\mathbf{m}$ is the image that we are trying to solve for, $\mathbf{S}$ is an operator that transforms $\mathbf{x}$ into a sparse domain, $\mathbf{s}$ is the acquired k-space data. $||\cdot||_p$ is the p-norm. In the case of a vector $\mathbf{v}$ it is calculated $||\mathbf{v}||_p = (\sum_i |v_i|^p)^{1/p}$. Sparse images have a low 1-norm and are thus prioritised when this cost function is minimised. The reason for using the 1-norm rather than the more intuitive "0-norm" (which is not a true norm, but counts the number of non-zero elements in a vector) as a regularisation factor is that solving a minimisation problem of the 0-norm is computationally hard to solve, but it has been shown that the 1-norm works as a substitute [3]. The factor λ is a number weighting the first term (data consistency) against the second term (sparsity of the image). If $\lambda = 0$, minimising the cost function finds a solution that is consistent with the data, however, in the noiseless case, there are infinitely many solutions, and in the noisy case the solution with the shortest distance between $\mathbf{Ex}$ and $\mathbf{s}$ and thus the least-squares-fit to the data is found. With $\lambda \neq 0$ a solution which is both consistent with the data and sparse in some domain is chosen.

The cost function, as described above in equation 2.21, is a convex function, but due to the second term, not fully differentiable. This means that a standard gradient descent algorithms alone does not work on this type of problem. A common algorithm for problems of this form is a fast iterative shrinkage-thresholding algorithm (FISTA [29]) that combines gradient descent of the first smooth term of the cost function with shrinkage or soft thresholding that pushes the solution towards the minimum of the second term of the cost function.

FISTA is a modified version of an iterative shrinkage-thresholding algorithm (ISTA) that is faster in terms of iterations required for convergence. The iterative step of ISTA is described in figure 2.7. Mathematically it can be written as:

$$\mathbf{m}_{n+1} = T_t(\mathbf{m}_n - \alpha \mathbf{E}^{\mathrm{H}}(\mathbf{E}\mathbf{m}_n - \mathbf{s})) \tag{2.22}$$

here α is a step size parameter that can be anything smaller than the inverse of the largest eigenvalue of $\mathbf{E}^{\mathrm{H}}\mathbf{E}$. $T_t(\mathbf{z})$ is a function that performs the non-linear soft thresholding operation on each element in $\mathbf{z}$ with a thresholding level $t \geq 0$, and is defined as:

$$T_t(\mathbf{z})_i = \begin{cases} z_i - t & \text{if } z_i \geq t \\ 0 & \text{if } -t \geq z_i < t \\ z_i + t & \text{if } z_i \leq -t \end{cases} \tag{2.23}$$

in the case of complex valued z, such as in MRI, the thresholding is applied to the magnitude of each element maintaining phase information.

FISTA has been shown to converge more quickly than ISTA because it utilises a specific linear combination of the two previous steps instead of just the previous step in its iterative update:

$$\begin{aligned} \mathbf{m}_n &= T_t(\mathbf{w}_n - \alpha \mathbf{E}^{\mathrm{H}}(\mathbf{E}\mathbf{w}_n - \mathbf{s})) \\ t_{n+1} &= \frac{1 + \sqrt{1 + 4t_n^2}}{2} \\ \mathbf{w}_{n+1} &= \mathbf{m}_n + \left(\frac{t_n - 1}{t_{n+1}}\right)(\mathbf{m}_n - \mathbf{m}_{n-1}) \end{aligned} \tag{2.24}$$

2.3.4 Sampling Trajectories

The most common way of sampling k-space is on a Cartesian grid. When sampling on a regular rectangular grid, the properties of the PSF are easily explainable with Fourier relationships as shown above. It is also the sampling method that produces the highest SNR as each point in k-space has equal weighting on the final image, minimising the noise. For k-space interpolation methods (GRAPPA), Cartesian imaging is also helpful as the same interpolation kernel can be used for every missing k-space point. In compressed sensing type reconstructions Cartesian imaging is non-optimal because of how the aliased signal generally only spreads in one dimension

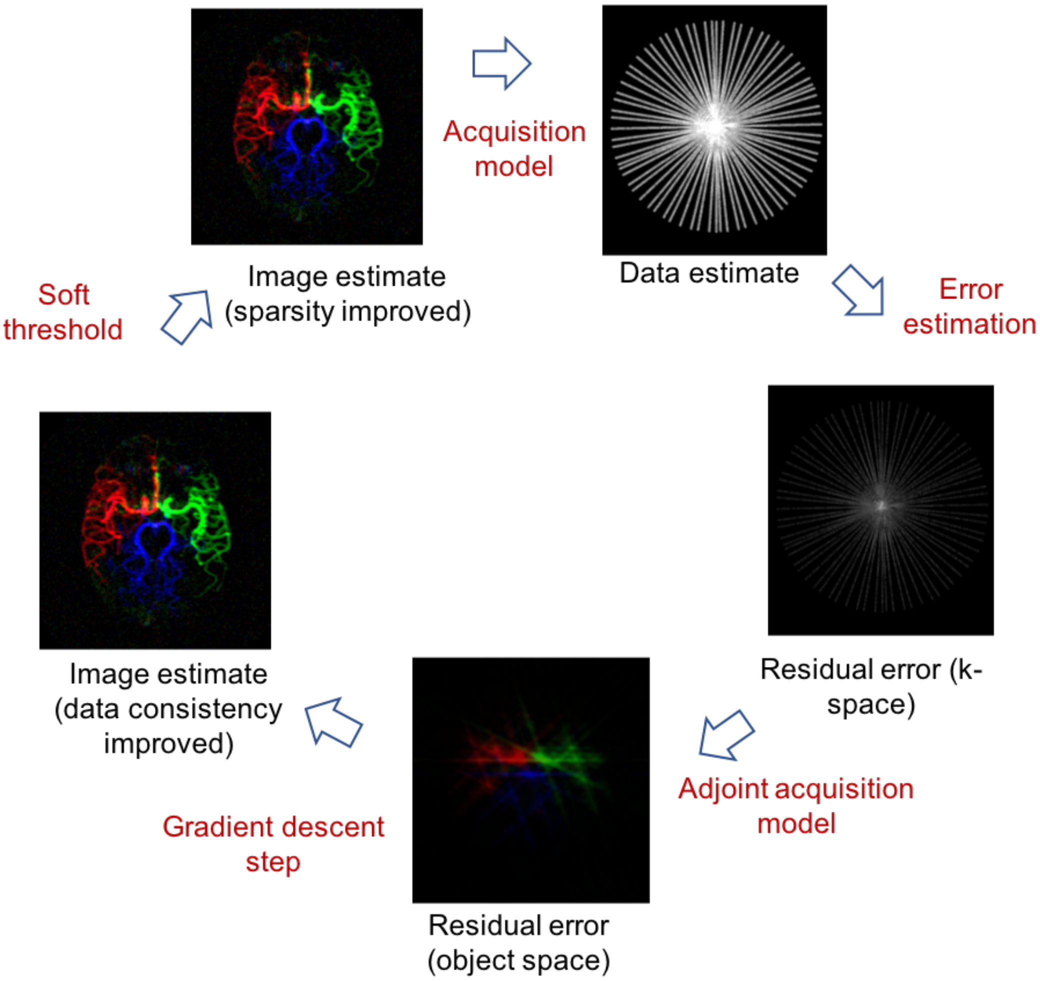

Figure 2.7: The iterative step of ISTA. Starting with an image estimate (top left) it is possible to model what the output k-space data would be (top right). That estimate is then compared to the collected data and the difference is found (bottom right). Using the adjoint acquisition transform the difference is brought back to object space (bottom middle), this is the local gradient of the cost function. The image estimate is updated by moving the along the gradient to a solution that is more consistent with the data (bottom left). Finally, the soft thresholding operation is applied to the resulting image enforcing sparsity, and the process is repeated.

for 2D imaging, and two dimensions for 3D imaging, which generally makes it less noiselike and harder to remove with CS, which is inherently a denoising method.

Other methods are better suited for CS, due to their incoherent aliasing patterns. Radial sampling, which is the main focus in this thesis, is commonly used in dynamic

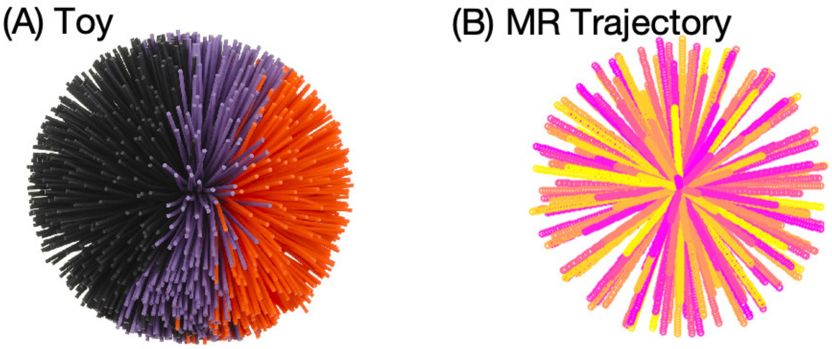

Figure 2.8: A koosh ball (A) and a koosh ball trajectory (B), which is a 3D radial k-space sampling trajectory.

imaging that needs high undersampling factors. The benefits of radial sampling include reduced sensitivity to motion [30], and spatially incoherent aliasing when sampled below the Nyquist limit [31].

Commonly, 2D radial sampling is performed with uniform angular gaps between the spokes such that each window (spokes reconstructed together) is reconstructed with N spokes $\frac{1}{N} \times 180°$ apart. In 3D, radial trajectories are often referred to as koosh-ball trajectories because of their resemblance to a children's toy with that name (figure 2.8). Multiple methods for determining the placement and ordering of spokes in 3D radial sampling have been proposed [32, 33] but we will limit our discussion to 3D radial sampling with the multidimensional golden means method [34] described below.

In the same way as the separation of k-space samples determines the alias free FOV side length in Cartesian sampling, the largest gap between k-space samples in a radial trajectory determines the alias free FOV diameter. The number of radial spokes needed to achieve the same resolution and FOV as a Cartesian image with N lines is $N \times \frac{\pi}{2}$ [16]. This extends to 3D sampling too. Therefore, $N \times \frac{\pi}{2}$ is the number of spokes we will use as the $R = 1$ reference point when we define the acceleration factor. If the spokes are not sampled uniformly, e.g. with the golden ratio sampling method as explained below, we shall still assume this number of spokes for $R = 1$, although the true largest k-space gap might be slightly larger, especially if the number of spokes is low.

Golden Ratio Based Sampling Methods

A 2D golden ratio based sampling trajectory [35] is used a lot in the thesis because of its flexibility for reconstructing images with varying temporal resolution and/or undersampling factors. Moving to 3D requires an extension of the golden ratio method. Such an extension was presented by Chan et al. [34], and was coined the multidimensional golden means method. In this section, the properties of the golden ratio method and the extension into the multidimensional golden means method is reviewed.

The idea behind golden ratio sampling is that each new spoke acquired with a constant increment intersects the largest gap in k-space by the golden ratio, ϕ. In 1D, the golden ratio is a way of partitioning a line into two segments, such that the ratio of the segments to each other is the same as the ratio between the two segments combined and the longer segment (figure 2.9(A)). When the segment in question is curved around the arc of a semi-circle, a golden angle can be derived from the golden ratio (figure 2.9(B)). This relationship can also be expressed mathematically like this:

$$\phi = \frac{1}{1+\phi} \tag{2.25}$$

Equation 2.25 extends into a quadratic equation, $\phi^2 + \phi - 1 = 0$ with solutions $\phi = \frac{1 \pm \sqrt{5}}{2}$ and the corresponding golden angles are $180° \times \phi$. Normally, the smaller solution of ϕ is used, ($\frac{1-\sqrt{5}}{2} \approx -0.61803$). The minus sign can be ignored as the ordering of the two segments are irrelevant, only the ratio of segment lengths is of interest. By placing the tip of each radial spoke on a unit circle in steps of $\phi \times 180° \approx 111.246°$, golden ratio sampling in 2D is achieved.

To extend this method to 3D sampling, the concept of a golden ratio has to be extended to 2D. Rather than picking points on a line segment and wrapping that around a circle for spoke directions, points are chosen on a square and wrapped

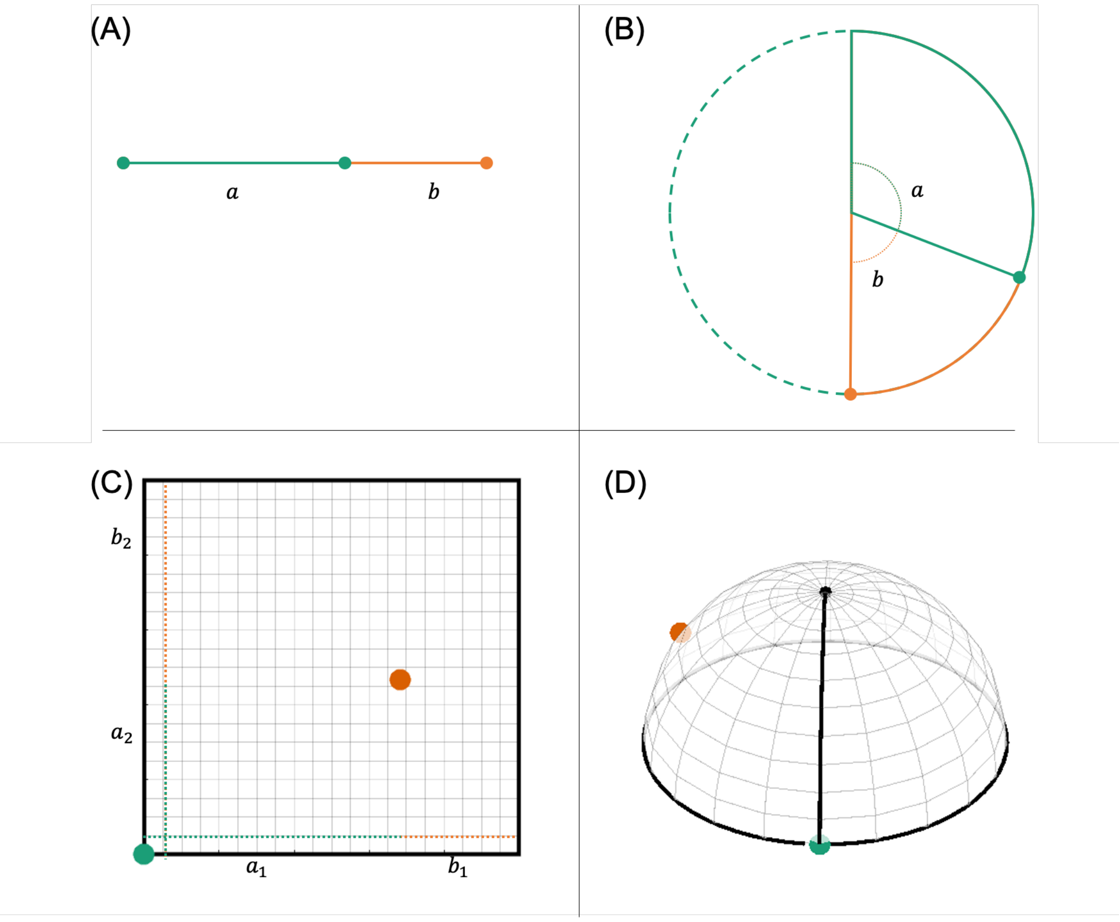

Figure 2.9: The 1D golden ratio (A) allows for 2D sampling (B) by wrapping the line around a semi circle. The 2D golden mean (C), allows for 3D sampling (D) by wrapping the plane around a hemisphere.

around a hemisphere with the same surface area and used as sampling spoke directions (figure 2.9(C) & (D)).

The extension to 2D is easier to explain if we derive the golden ratio in a different way. The golden ratio is the limit of the ratio between consecutive elements of the Fibonacci sequence. In the Fibonacci sequence each element is the sum of the two previous elements and can be written as a recursive relationship:

$$\begin{bmatrix} F_{i-1} \\ F_i \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} F_{i-2} \\ F_{i-1} \end{bmatrix} \tag{2.26}$$

From this relationship the same golden ratios as above can be found as the eigenvalues of the Fibonacci matrix. To extend to higher dimensions an alternative sequence, called Narayana's cow sequence or the supergolden sequence, can be used instead of Fibonacci. In the supergolden sequence the previous number of the sequence is added to the one two places before that. The recurrence relationship is thus:

$$\begin{bmatrix} F_{i-2} \\ F_{i-1} \\ F_i \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} F_{i-3} \\ F_{i-2} \\ F_{i-1} \end{bmatrix} \tag{2.27}$$

and the eigenvalues of this matrix, when normalised such that the largest one is equal to 1, can be used to sample a unit square with the same efficiency as a unit line segment is sampled with the standard golden ratio. For more details see [34, 36].

To transform from 2D sampling of a square to radial 3D sampling in a sphere, Chan proposes the following transform that maintains area around each point from the square to the surface of the sphere:

$$\theta = x \times 2\pi \tag{2.28}$$

$$\psi = \sin^{-1} y \tag{2.29}$$

Where θ is the azimuthal angle (direction along the equator of the circle), and ψ is the elevation angle. x and y are the Cartesian coordinates on the unit square.

The 2D golden means method does, however, not provide all of the same properties as the simple 1D golden ratio method because of this transform from a 2D plane to a 3D sphere. The sampling uniformity suffers from similar distortions like flat maps of the Earth do. The transform from the unit square to the surface of a sphere, as described by Chan et al., maintains the area coverage of each point, but near the poles the areas get elongated, meaning that points are sampled more tightly in one direction than another. Because of the non-linear nature of the plane-to-sphere transform, unlike in 2D sampling with the 1D golden ratio method, the spoke distributions of subsequent frames are not just rotations of the previous spoke frame. Subsequent frames therefore have somewhat different sampling efficiencies. However, as Chan et al. showed in their work, the multidimensional golden means method still has good temporal stability, i.e. subsequent frames of arbitrary numbers of

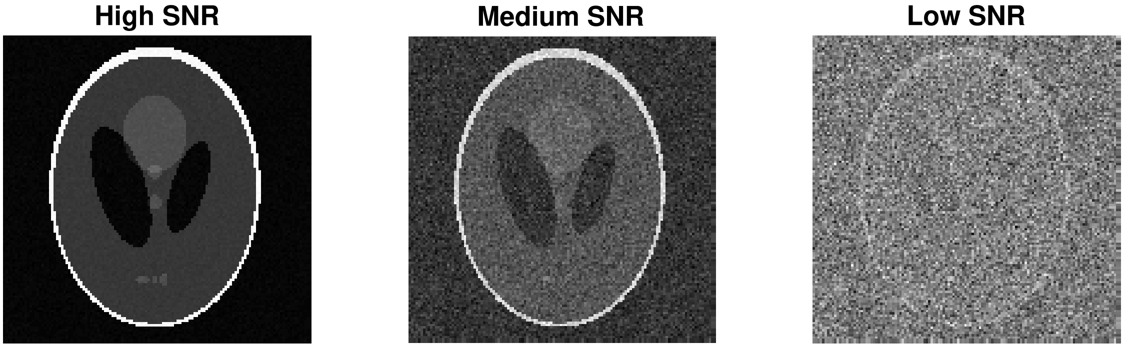

Figure 2.10: SNR affects the ability to distinguish a features in a signal. At high SNR, (A), all features can be distinguished, at medium SNR, (B), some lower contrast features are lost, and for very low SNR, (C), the object might be completely unrecognisable.

spokes have similar sampling efficiencies, compared with linear and bit-reversed [37] radial sampling methods that work best when the number of spokes to use in each frame is known a-priori.

2.4 Quantifying Image Quality

To quantify how well any one reconstruction method is working, some type of metric that describes image quality is needed. Preferably the metric would correlate with the qualitative assessment of a human viewer. It has been shown [38] that simple metrics such as the absolute error between a reconstructed image and the ground truth, or the peak signal-to-noise-ratio, which are commonly used metrics, do not correlate well with human perception of image quality. So, what does? This is still an unsolved problem in image processing and analysis, and human perception cannot be easily digested into a single quality metric, so the choice of metric must be based on the use case, with knowledge of its inherent limitations. This section lists some metrics commonly used to assess image quality, and outlines their suitability in different types of analyses.

2.4.1 Image SNR

One of the most common metrics in image analysis is signal-to-noise ratio (SNR). In theory the metric is very simple: SNR is, as the name suggests, the ratio of the signal to the noise level. A simple 2D example is shown in figure 2.10.

SNR can be measured by averaging signal in a region of interest (ROI) and assuming that the signal is constant such that the SNR numerator is the mean signal in that area, and the denominator is the standard deviation of voxel values in the same area. This method only really works in simple simulations as there is almost always real signal variation within any ROI in vivo. In that case, the standard deviation can instead be measured in another ROI in the background. Alternatively, SNR can be measured on a pixel-by-pixel level by taking multiple measurements with independent noise samples, e.g. across time (assuming the signal stays constant), or, in simulations, running the simulation multiple times with many instances of added noise. Any artifacts or non-Gaussian noise (e.g. physiological noise) will hinder accurate SNR measurements. Often SNR of an image is reported as mean SNR in some mask over the object or as peak SNR (pSNR).

SNR can be a useful metric to compare when no ground truth is available. However, accurate measurement of SNR is often hard to do because of the issue with artifacts mentioned above, and assumption of even signal in the ROI where noise is measured (if the background is used in a magnitude image, Gaussian noise becomes Rician and requires a correction factor to be comparable with SNR measured on complex images). Comparisons between different subjects might not be accurate, but comparisons between different reconstructions of the same data can be indicative of image quality.

When non-linear reconstruction is used, SNR is not a very useful metric, as it can be very biased towards, for example, over-regularised reconstructions, where the noise variance is negligible, and the signal is highly corrupted but still measurable. In the case where the noise variance tends to zero, SNR tends to infinity, as long as there is some signal. This is an especially large problem if pSNR is reported.

2.4.2 Normalised Root Mean Square Error

Another common metric is the root-mean-squared error (RMSE), which has arbitrary units, or a version normalised by the signal in the ground truth image (normalised-root-mean-square error, NRMSE), which normally is reported as a percentage error.

RMSE and NRMSE are calculated using the following equations:

$$RMSE(x, x_0) = \sqrt{\frac{\sum_i^N (x_{0,i} - x_i)^2}{N}} \tag{2.30}$$

$$NRMSE(x, x_0) = \frac{RMSE(x, x_0)}{\sum_i^N x_{0,i}/N} \tag{2.31}$$

with x being the reconstruction and x_0 being the ground truth.

RMSE is a metric similar to SNR but measures both bias and variance, instead of variance alone. However, this requires ground truth, whereas SNR does not.

Although the metric is commonly reported as evidence of improved image reconstruction and in training of deep learning reconstruction methods [39], it does not always correlate well with perceptual quality [38].

2.4.3 Structural Similarity Index

Structural Similarity Index is a metric originally developed to compare digital natural images with their compressed versions. It was specifically designed to correspond with visual perception of image quality. It is calculated locally, but often averaged across the whole image to give one single score for an image. It works both for 2D and 3D data. The metric is calculated using the following formula:

$$SSIM(x, x_0) = \frac{(2\bar{x}\bar{x_0} + c_1)(2\sigma_{x,x_0} + c_2)}{(\bar{x}^2 + \bar{x_0}^2 + c_1)(\sigma_x^2 + \sigma_{x_0}^2 + c_2)} \tag{2.32}$$

where $\bar{x}$ signifies the mean, σ_x the variance, and σ_{x,x_0} the covariance between the two images. c_1 and c_2 are stabilising constants that have values that have been experimentally determined along with the dynamic range, L. $c_1 = (0.01L)^2$, $c_1 = (0.03L)^2$.

The formula takes into account the luminance, contrast, and structure of the images, however, the formula for how these are related is not very intuitive. The dynamic range parameter, L, contained within the constants c_1 and c_2, can change the results drastically, so depending on whether the dynamic range in the data, or the maximum dynamic range available based on the numerical data type, is used the results will vary. Most studies do not report how this parameter is chosen. If the images, as in medical imaging, especially angiography, contains many background pixels, the results can easily become biased too, i.e. the mean SSIM is dominated by the values in the noisy background, which contains no useful information and is easy to ignore visually. If images that are reconstructed have intrinsically different image intensities, it can be hard to compare their results because of the strong dependence on the dynamic range parameter. Another parameter that is hidden within this formula, is the local neighbourhood that SSIM is assessed over. In MATLAB the default setting of the kernel is an isotropic Gaussian with a standard deviation of 1.5 pixels. However, if images with different resolution are compared, again, this will bias the result if not corrected for.

If these parameters can be set up properly, and the data can be normalised for a fair comparison, SSIM can be a useful metric, especially if focused in on the relevant pixels (e.g. a vessel mask for angiography) [38].

2.4.4 Correlation

A metric that is not as commonly used is the correlation coefficient [40] between a ground truth image and a reconstruction, however, it is used extensively throughout this thesis as it was found to be a robust metric that generally agreed with visual inspection of the images. It is invariant to absolute image scaling, and easier to understand and interpret than for example SSIM explained above. It does not take into account spatial features, just voxel values, but by masking the images that are compared, it can be focused in on the relevant pixels. The correlation coefficient between two images (or sets of voxels) is:

$$r(x, x_0) = \frac{1}{N-1}\sum_{i=1}^{N}\left(\frac{x_i - \bar{x}}{\sigma_x}\right)\left(\frac{x_{0,i} - \bar{x_0}}{\sigma_{x_0}}\right) \tag{2.33}$$

with $\bar{x}$ signifying the average within the set, and σ being the standard deviation within the set.

2.5 Imaging Cerebral Haemodynamics

To understand the health and physiology of the brain, studying blood flow on multiple scales is important. Angiograms can shed light on abnormal vessel anatomy such as intracranial aneurysms [41], vascular malformations [42], and narrowings such as carotid stenosis, and thromboembolisms [43]. It can also be useful in studying blood supply to tumours [44] and help plan cerebrovascular surgery [45]. Perfusion imaging, on the other hand, can be used to study the function and physiology of the brain as there are strong correlations between, for example, perfusion, oxygenation, and neural activity [46]. Abnormal neurovasculature does also not always lead to abnormal tissue perfusion, which is often what is most crucial in terms of the brain functioning normally, so combinations of both angiographic and perfusion information are useful in a range of diseases [47].

Historically, imaging of cerebral haemodynamics have been performed with methods requiring ionising radiation that carry an inherent risk. Modern non-ionising methods of imaging have emerged as alternatives utilising the flexibility of MRI.

Some common methods for imaging the cerebrovasculature and perfusion are presented briefly below. Many other imaging methods and modalities also exist and for example [48, 49] provide extensive overviews of both non-MR and MR methods.

2.5.1 Non-MR Methods

Digital Subtraction Angiography

The gold standard way of visualising the cerebral arteries is using subtraction angiography, a method that relies on subtraction of two x-ray images acquired before and after administration of a, typically iodine based, radio-opaque contrast agent. Because the only difference before and after administration of contrast agent is the appearance of the vessels, subtraction of the two images results in an image of the contrast agent in the vessels only. This method is often referred to as DSA, Digital Subtraction Angiography, as the subtraction is done computationally nowadays, but the method was first described as a manual method with two x-ray films, one with inverted contrast, taped to each other [50].

On top of the risk associated with ionising radiation, DSA is an invasive method that requires arterial access, use of contrast agents, and is often performed under general anesthaesia to keep the patient immobile for the subtraction to work. It has been shown to carry a risk of causing stroke or bleeding [51].

Computed Tomography Angiography and Perfusion Imaging

As an alternative to DSA, less invasive methods, such as computed tomography angiography, CTA, can provide comparable diagnostic value without the need for arterial access [52]. In CTA, the contrast is given intravenously, and a 3D image is reconstructed to be able to see the vessels despite being surrounded by the skull.

CT can also be used for perfusion imaging, and is commonly used in acute stroke imaging. Although MRI methods (often diffusion weighted imaging) of assessing the core (dead tissue) and penumbra (tissue at risk) are more sensitive and specific [53], CT is a fast and common alternative.

Although less invasive than DSA, CTA and CT perfusion still carry the inherent risk involved with ionising radiation and the use of contrast agents.

2.5.2 MR methods

Dynamic Contrast Enhanced MRI

Dynamic contrast enhanced (DCE) MRI is a method often used for quantitative or semi-quantitative perfusion mapping [54]. It is an imaging method where a gadolinium based contrast agent is administered intravenously, which shortens the T1 of the tissue due to being paramagnetic. The passage of contrast agent through the brain vasculature can then be analysed using pharmacokinetic modelling for perfusion quantification.

Alternatively, DCE can also be used for angiography. In this modality images are acquired rapidly after injection of contrast agent, such that the first pass of contrast agent can be captured. For subsequent passes the bolus will be dispersed, as well as residing in the venous system, with the effect of worsened arterial contrast.

Because the need to capture this first pass of bolus with high spatial and temporal resolution, very high acceleration is required [55].

Gadolinium based contrast agents are generally considered safe for patients with normal kidney function [56]. However, emerging evidence of gadolinium retention in the brain a long time after injection has raised concerns regarding the safe use of these exogenous contrast agents [57].

Time of Flight Angiography

There are also MR methods that do not require exogenous contrast agents. One method for angiography is time-of-flight (TOF)[58]. In TOF imaging contrast is generated by repeated saturation of a slab of tissue whilst non-saturated blood with high signal intensity arrives into the slab producing angiographic contrast. This method is a static method, as there is no way of following a bolus through the arteries. It is also only sensitive to blood vessels arriving from outside the imaging region, as any blood that starts off within the slab (e.g. in a vessel parallel to the imaging plane) will experience the saturation pulses and thus not be enhanced in the way fresh blood arriving from outside the imaging slab is.

Phase Contrast Angiography

Another contrast agent free MR method is phase contrast angiography [59]. It relies on a pair of bipolar gradients that causes a shift in phase for moving spins, but leaves stationary spins unaffected as the effect of the second gradient block undoes the effect of the first if the spin is in the same location as when the first was applied. The resulting acquired phase is directly related to the velocity the spins were moving at, assuming constant velocity and bulk motion of blood within voxels. Phase contrast angiography is thus a quantitative method of imaging moving blood. Normally, the velocity encoding is only applied in one dimension, e.g. inferior-superior, but multiple encodings can be performed to get a multidimensional vector of velocity. Especially for 3D imaging, the scan time required compared with TOF is long. Even 1D phase contrast angiography has generally longer scan times than TOF because the increased TE and TR required to include the encoding gradients in the sequence, and because phase contrast imaging requires a reference scan to remove background phase variation (either with no phase encoding, or phase encoding with the opposite gradients) [16]. Phase contrast angiography can also suffer from partial volume effects if a voxel includes both stationary and moving spins and thus require very high spatial resolution, which also forces the scan times to be long.

Arterial Spin Labelling

The main topic of this thesis, is arterial spin labelling (ASL [60–62]), which is a completely non-invasive imaging method, requiring no exogenous contrast agent, that can be used both for angiography and perfusion imaging, much like DCE. The origin of contrast is blood that has been "labelled" with an inversion pulse before entering the imaging region. The total magnetisation in a voxel that will give rise to a resonant signal when imaging is $\mathbf{M}_{brain} + \mathbf{M}_{blood}$, where $\mathbf{M}_{brain}$ is the magnetisation coming from the static brain tissue, and $\mathbf{M}_{blood}$ is the magnetisation from the blood (ignoring T1 decay between tagging and imaging). The magnitude of $\mathbf{M}_{blood}$ is generally one to two orders of magnitude smaller than the magnitude of $\mathbf{M}_{brain}$. However, by taking two images, one with $\mathbf{M}_{blood}$ inverted (= a tag image) and one with it not inverted (= a control image), the signal $2|\mathbf{M}_{blood}|$ can be retrieved by subtracting the two images from each other. This relationship can be shown more formally by the equation

$$\begin{gathered}\mathbf{s}_{voxel} = \mathbf{A}\mathbf{m}_{voxel} \\ \mathbf{s}_{voxel} = \begin{bmatrix} s_1 \\ s_2 \end{bmatrix}, \mathbf{A} = \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}, \mathbf{m}_{voxel} = \begin{bmatrix} B \\ S \end{bmatrix}\end{gathered} \tag{2.34}$$

where s_1 and s_2 represent the tagged and untagged voxel values respectively, B is the signal coming from the blood, and S is the signal coming from static tissue. $\mathbf{A}$ is the encoding matrix that describes the tagging and controlling of blood. Note that the blood water signal can come either from blood in blood vessels or from perfused tissue depending on the timing between tagging and imaging. Assuming that the tag and control images have been reconstructed well, for each voxel one can then take the inverse (or pseudo-inverse if $\mathbf{A}$ is non-square) of $\mathbf{A}$ to get the expected result that $B = \frac{1}{2}(s_2 - s_1)$.

$$\begin{aligned} \mathbf{m}_{voxel} &= \mathbf{A}^{+}\mathbf{s}_{voxel} \\ \mathbf{A}^{+} = (\mathbf{A}^{\mathrm{H}}\mathbf{A})^{-1}\mathbf{A}^{\mathrm{H}} &= 0.5 \times \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix} \end{aligned} \tag{2.35}$$

There are different methods for labelling and controlling the blood signal: pulsed, continuous, and pseudo-continuous ASL (PASL, CASL, and pCASL). In PASL the spins in a large volume (in the neck) are all inverted by a single RF pulse, then the blood is allowed to travel to the imaging region (in the brain) and imaged there. CASL, on the other hand, is based on a continuous flow-driven adiabatic inversion pulse that labels blood in a labelling plane as the blood travels through that plane before imaging it in the brain. CASL gives rise to stronger blood inversion signal than PASL, but is hard to implement due to RF amplifier limitations. CASL also has magnetisation transfer problems, which makes it harder to get a well matched control condition. Therefore pCASL, a method that mimics CASL without the need for a continuous RF pulse has been developed [63] and is the most common ASL labelling method. CASL uses so-called flow driven adiabatic inversion pulses to invert the spins flowing through the labelling plane. Flow driven adiabatic inversion works by having spins experience a slowly varying net magnetic field produced by a combination of a gradient and the $\mathbf{B}_1$ field (figure 2.11). The main difference between CASL and pCASL is that in pCASL the RF pulse is broken up into shorter pulses (and changed from rectangular pulses to Hanning pulses to give a smooth variation in the effective flip angle experienced by the magnetisation as it travels through the labelling plane). The splitting of the continuous RF to a pulsed version leads to aliased labelling planes, whose separation is determined by the average gradient applied during labelling. To make sure that aliased labelling planes are outside the region that are affected by the RF, such that flowing spins only experience one inversion, the gradient is strong during RF excitation which leads to narrow RF pulses, and a gradient in the opposite direction is applied between RF pulses, such that the average gradient strength still approximately

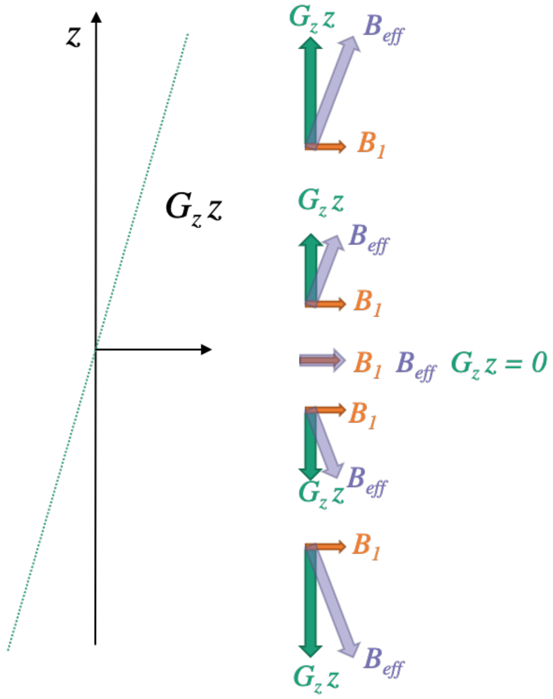

Figure 2.11: Flow driven adiabatic inversion. In the rotating frame spins will precess around the net magnetic field produced by the gradient and $\mathbf{B}_1$ field. If the net field is slowly changed as the spin moves in the z-direction the magnetisation will follow the net field and thus be inverted after having passed through the labelling plane.

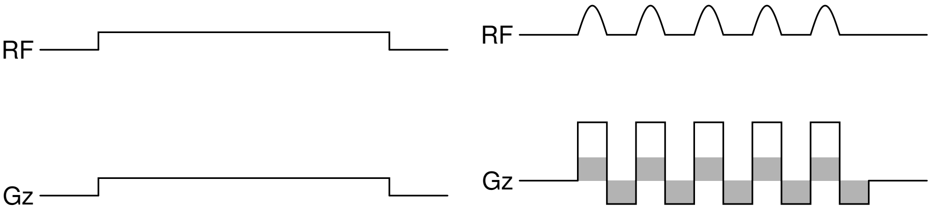

Figure 2.12: Left: CASL, Right: pCASL. In both cases the "tag" condition is showed. For "control", CASL can locate the labelling plane opposite the imaging region, whereas pCASL can achieve a control condition by alternating the polarity of the RF pulses. In the pCASL diagram the grey shading represents equal areas. The unshaded area in the pCASL sequence is equal to the total area under the CASL gradient.

matches that of a CASL pulse. Pulse sequence diagrams for both CASL and pCASL labelling are shown in figure 2.12.

As an extension to pCASL, it is possible to modulate the inversion efficiency across the labelling plane by applying transverse gradients during the tagging [5] (figure 2.13). This method is called vessel encoded ASL (VE-ASL). By taking a number of images, where blood from different arteries are either labelled, not

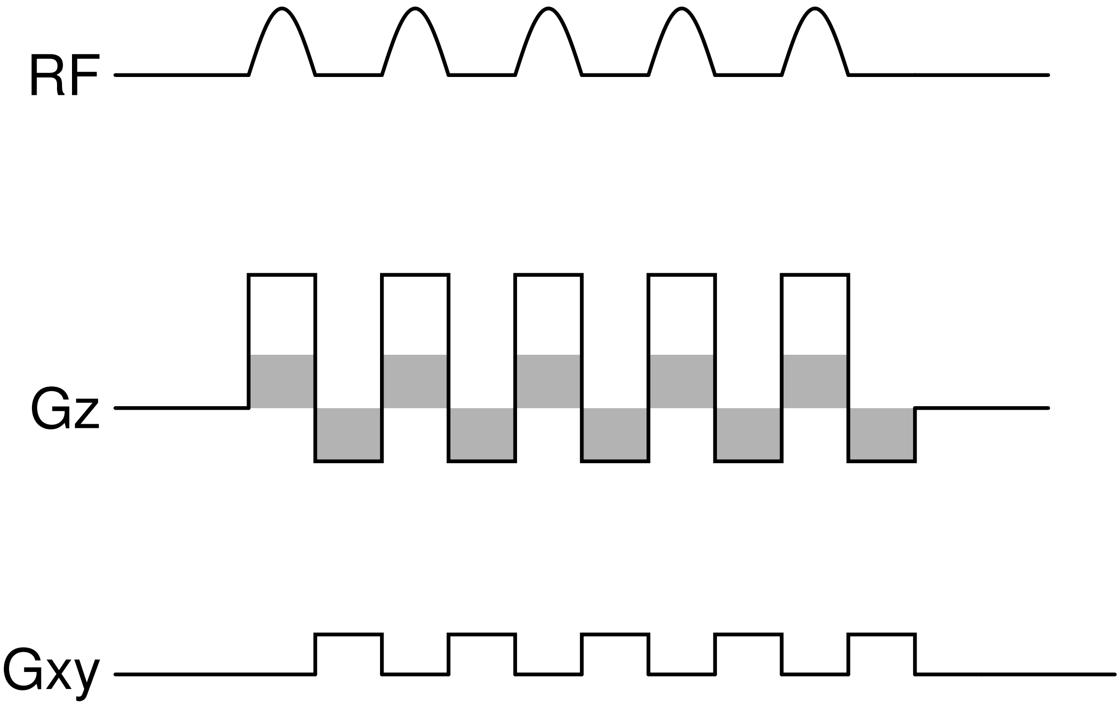

Figure 2.13: VE-pCASL. Between each pulse a transverse gradient is applied such that depending on the spins' location within the labelling plane, the following pulse either undoes the effect of the previous pulse or continues the inversion process. This results in spatially varying labelling efficiency.

labelled, or partially labelled, one can theoretically reconstruct the blood signal from each supplying artery separately by solving a similar problem to equation 2.34, with the encoding matrix modified to reflect the tagging of individual arteries for example like this:

$$\mathbf{y}_{voxel} = \mathbf{A}\mathbf{m}_{voxel}$$
$$\mathbf{y}_{voxel} = \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ s_4 \end{bmatrix}, \mathbf{A} = \begin{bmatrix} -1 & 1 & -1 & 1 \\ 1 & -1 & -1 & 1 \\ -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \mathbf{m_{voxel}} = \begin{bmatrix} R \\ L \\ B \\ S \end{bmatrix} \tag{2.36}$$

where $\mathbf{y}_{voxel}$ is the resulting selectively tagged images, and $\mathbf{m}_{voxel}$ is a matrix with number of rows equal to number of components, and number of columns equal to the number of voxels in the image. The components in this example are R for the right carotid blood, L for the left carotid blood and B for blood supplied by the the basilar arteries. S still stands for the static tissue signal. To find $\mathbf{m}_{voxel}$ you can simply solve it using the pseudoinverse of the encoding matrix in the same way as shown in (2.35). Probabilistic methods for calculating the blood flow from different arteries also exist [64].

A should be designed such that it is full rank and each component is labelled or controlled an equal number of times for optimal signal to noise ratio. The ideal choice is a Hadamard matrix like the one shown in (2.36). Wong [5] showed that vessel-encoding using a Hadamard matrix results in the same SNR as a standard pCASL acquition acquired with the same scan time. This can be understood by how Hadamard encoding weights each component equally, and thus the combined acquisition operator remains orthogonal and does not cause noise amplification.

The encoding operator can be included in the full signal acquisition operator used in reconstruction by stacking the different vessel components (m_1, m_2, ...) and the different vessel encodings (s_1, s_2, ...) in the vectors **m** and **s** as described in equation 2.14 above. The acquisition operator, **E** then becomes **FHS**, with **F** describing the Fourier sampling, **H** being the vessel encoding (same as **A** but with blocks of ± the identity operator instead of ± ones), and **S** the coil sensitivities The full size of **E** is then $N_{k-space-samples}N_{coils} \times N_{voxels}N_{vessel-components}$.

Although the static tissue is theoretically removed from the vessel image(s) by the decoding process, background suppression methods are often used to lower the intensity of the static tissue and thus the physiological noise related to it.

2.6 Summary and Conclusion

In this thesis, only a small corner of the vast field of magnetic resonance research is explored. MRI is a versatile method of imaging with intricate localisation and contrast mechanisms. It can, among many other things, provide clinical information about the vasculature and perfusion state of the brain, with or without the use of contrast agents. In this thesis, the focus is on a specific MRI method called arterial spin labelling.

Arterial spin labelling comes in many flavours, and again, this thesis is mainly about a very specific type of arterial spin labelling, namely vessel encoded ASL. Vessel-selectivity is not generally attainable with other methods of imaging, and thus provides unique information that could be of clinical or other research value. However, as vessel encoding increases the number of encoded images compared to standard ASL, which is a fairly slow method of imaging already, acceleration of the acquisition is of high importance.

Linear and non-linear approaches to acceleration through reconstruction of undersampled MR acquisitions were briefly presented in this chapter. These methods rely on both extension of the acquisition operator to make the problem better posed, and inclusion of priors on the reconstruction to choose the most correct reconstruction among all possible solutions that are compatible with the data.

In the next chapter a first attempt of combining vessel-encoded ASL with non-linear reconstruction is presented.

It is truly an amazing feeling when you know that you have built something that no one else ever has – and it actually works.

— Donna Strickland in her Nobel Banquet Speech
2018

3

Feasibility of Accelerating Vessel-Encoded ASL Angiography

Contents

3.1 Introduction 45
3.2 Methods 46
  3.2.1 Forward Model 46
  3.2.2 Simulations 48
  3.2.3 In Vivo Acquisition 49
  3.2.4 Pre-processing 53
  3.2.5 Reconstruction 54
  3.2.6 Analysis 55
3.3 Results 57
  3.3.1 Effect of Increased Relative Sparsity 57
  3.3.2 Simulations on Real Data 58
  3.3.3 Regularisation Optimisation 59
  3.3.4 In Vivo Comparison of VE and non-VE 61
  3.3.5 High Resolution Imaging 64
3.4 Discussion 66
  3.4.1 Relative Sparsity as Driving Factor 66
  3.4.2 Temporal Regularisation 67
  3.4.3 Lambda Optimisation 67
  3.4.4 Feasibility of Accelerating VE and non-VE-ASL 69
3.5 Conclusion 70

3.1 Introduction

In this chapter, a proof-of-principle study into combining radially sampled VE-ASL angiography with compressed sensing and parallel imaging is introduced. This work was published in *Magnetic Resonance in Medicine* [10].

As outlined in the background chapter, section 2.5.2, VE-ASL angiography is a versatile method of imaging the cerebral vasculature. The main drawback of VE-ASL is its longer acquisition times compared to other angiographic methods, and in particular to non-vessel-encoded ASL (non-VE-ASL) because $N + 1$ images are required to separate blood coming from N arteries compared with only a tag and control image for non-VE-ASL. To achieve equal scan time, and thus equal SNR, the VE images have to be acquired with higher undersampling factors than non-VE. VE-ASL angiography is, however, well suited for acceleration.

The two main reasons why VE-ASL angiograms might be particularly well-suited to undersampled reconstruction are related to intrinsic properties of angiographic data:

First of all, angiograms are spatially sparse. This can be exploited in a compressed sensing (CS [3, 4]) acquisition and reconstruction framework. Compared with non-VE-ASL angiograms, VE-ASL images have higher relative sparsity because approximately the same number of non-zero voxels are distributed across multiple vessel-selective images. If there is no mixing of blood from different arteries, the number of non-zero voxels is exactly the same for VE and non-VE. However, if there is some mixing, the number of non-zero voxels is slightly higher for VE, but the overall relative sparsity would still be higher for VE unless the blood signal from all separated feeding arteries was mixed completely. Because relative sparsity, along with image dimensionality, and signal-to-noise ratio (SNR), contributes to the performance of a CS reconstruction [65] VE-ASL angiograms should perform better than non-VE angiograms in a sparsity-constrained reconstruction.

Secondly, at sufficiently high temporal resolution the signal varies smoothly in time [66] as the bolus of labelled blood passes through the arterial tree. This temporal smoothness can be exploited to further regularise the underdetermined image reconstruction problem. While exploiting redundancy in the temporal domain is possible in the dynamic acquisitions provided by both VE and non-VE-ASL, non-dynamic methods like time-of-flight angiography cannot benefit from this extra dimension of information.

In this study, an accelerated acquisition and reconstruction method for dynamic VE-ASL angiography is presented. The reconstruction approach is based on the enhanced spatial sparsity of vessel-specific angiograms and the smoothness of their temporal evolution. The proposed method produces VE-ASL images of comparable quality to non-VE-ASL at matched scan duration at acceleration factors varying from R = 2 to R = 34, providing vessel-specific information at no additional scan time or image fidelity cost.

3.2 Methods

3.2.1 Forward Model

The imaging system was modelled with a linear equation as outlined in the background chapter, section 2.3:

$$\mathbf{s} = \mathbf{E}\mathbf{m} + \mathbf{n} \tag{3.1}$$

where $\mathbf{s}$ is a vector containing the complex signal measured by all the receive coils, with each entry representing one point in k-space in one coil. Noise, $\mathbf{n}$, is a vector of complex noise, whose covariance depends on the coil noise covariance matrix. For simplicity, in simulations, the added noise will be Gaussian and uncorrelated between channels. The imaged object, $\mathbf{m}$, is a vector containing the complex magnetisation of blood from each vessel component as well as the static tissue for every position in physical space. Its length is therefore the number of voxels by the number of time points by the number of vessel components, i.e. a three-vessel VE image would have four components (three vessels and static tissue) and a non-VE image only two components (vessels and static tissue), thus making the vector $\mathbf{m}$ twice as long in the VE case. In this work, a three-vessel VE encoding was used, separating blood originating from three main arteries: the right and left internal carotid arteries (RICA, LICA), and the basilar artery (BA). **E** is the linear operator that models the encoding and acquisition of the signal. It contains three parts: i) the linear combination of signal from blood and static tissue components depending on the applied vessel-encoding scheme, ii) the spatial coil sensitivity profiles for each of the receive coils in the system, and iii) the Fourier sampling operator.

Simulations of the imaging system and subsequent reconstruction of both simulated and in vivo data was performed using MATLAB (Release 2017a, The MathWorks, Inc., Natick, Massachusetts, United States). The encoding operator (**E** in equation 3.1) was implemented as a composition of functions instead of one single linear operator because of the large size of the problem. The vessel-encoding part of **E** was implemented directly as a pixel-wise matrix multiplication of either a 2 × 2 Hadamard matrix (tag and control) for the non-VE case or a 4 × 4 Hadamard matrix for the VE case, as follows:

$$\mathrm{non-VE}: \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} RICA + LICA + BA \\ S \end{bmatrix} \tag{3.2}$$

$$\mathrm{VE}: \begin{bmatrix} -1 & -1 & -1 & 1 \\ -1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & 1 \end{bmatrix} \begin{bmatrix} RICA \\ LICA \\ BA \\ S \end{bmatrix} \tag{3.3}$$

RICA, LICA, and BA represents signal from the blood coming from the respective arteries, and S represents the static tissue signal intensity.

The coil sensitivity operator and its conjugate transpose were applied as pointwise multiplication on the image and weighted combination of coils for the forward and adjoint transform respectively. The transforms between non-uniform k-space samples and image space were implemented using the non-uniform fast Fourier transform (NUFFT [67]) in the Michigan Image Reconstruction Toolbox [68]. The k-space sampling trajectory was a 2D golden angle radial trajectory [35].

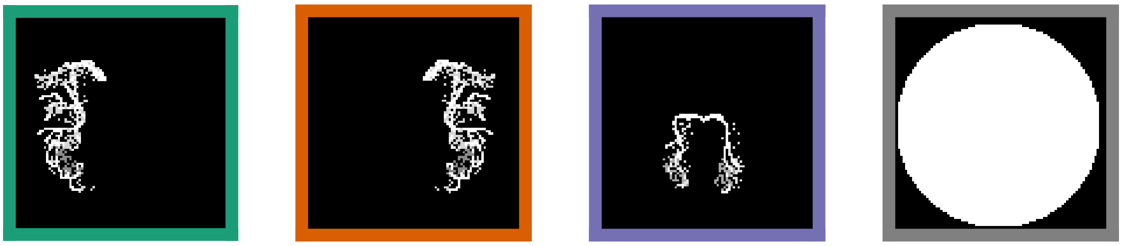

Figure 3.1: The simple numerical phantom used for sparsity experiments. The image with the green border mimics blood flow from the RICA, orange mimics the LICA, and purple mimics the BA. The "static tissue component" (grey border) has approximately 100 times higher intensity than the "vessel components". This colour scheme will be used in all subsequent figures to depict the RICA, LICA and BA.

3.2.2 Simulations

Two data sets were used as ground truths for numerical simulations. One was a simple numerical phantom and the other a fully sampled dynamic VE-ASL angiogram from a previous study [69].

The numerical phantom consisted of a single frame of a hand-drawn “vessel-like” image on a 96 × 96 pixel grid. It consisted of three "vessel components" and one "static tissue" component. The first "vessel component" occupied primarily the left part of the field-of-view; a mirror image of it, the second "vessel component", occupied primarily the right part of the field-of-view; and the third component occupied the lower space between the two. The "vessel components" overlapped in only nine pixels. The vessel component images had pixel intensities ranging from 0.54 to 1.00, and the "static tissue component" was a simple filled circle covering all of the “vessels,” with uniform intensity of 100.00. The phantom is shown in figure 3.1.

This simple phantom was used in the initial experiments to confirm that the hypothesis regarding increased relative sparsity improving the reconstruction quality holds in a simplified system. This phantom was used to generate both VE and non-VE images (in non-VE the three "vessel components" were added together to create a single "vessel component"). In terms of sparsity, the non-VE data had 14% non-zero values in the "vessel component" whereas the VE data contained only 5% non-zero values distributed across the three "vessel components". The levels of sparsity in the phantom are realistic for 2D angiographic imaging at, or just above, the circle of Willis, as confirmed by the in vivo data presented in section 3.2.3. The in vivo images had sparsity levels of 5.3% ± 0.7% for VE and 15.2% ± 2.2% for non-VE. No noise was added in the simplified system, and only one receive coil with uniform spatial sensitivity was modelled so that just the effect of differences in sparsity between VE and non-VE angiography could be studied. The images were transformed into k-space data using the forward model described in section 2.3 and reconstructed with 100%, 50%, 25%, 12.5%, 6.25%, and 3.125% of the number of samples needed to reach the Nyquist limit.

The second simulated acquisition and reconstruction was done on a high SNR, and fully sampled dynamic angiogram. It was used to mimic the in vivo system as closely as possible, but with a well-defined ground truth and controlled noise conditions. Coil sensitivity maps measured using a phantom in a 32-channel head coil were included to generate multichannel data in the simulated acquisition, however, for reconstruction the coil sensitivity maps were estimated directly from the undersampled data, as explained further in section 3.2.3. Complex Gaussian noise was added in k-space. k-space SNR was defined as:

$$SNR_k = \frac{rms(I)}{\sigma_{noise}} \tag{3.4}$$

where $rms(I)$ represents the root-mean-square intensity of the noiseless k-space measurements, and σ is the standard deviation of the added noise signal. Three noise conditions were simulated on the second phantom: (1) no noise, (2) moderate SNR ($SNR_k = 185.8$), and (3) low SNR ($SNR_k = 92.9$). These simulated data sets were subsequently undersampled and reconstructed in exactly the same manner as the in vivo data (see section 3.2.3). The high noise condition produced image SNR comparable to the in vivo image (figure 3.2) when reconstructed with a linear least-squares reconstruction on fully sampled data.

3.2.3 In Vivo Acquisition

Six healthy volunteers (all male, age range: 25 to 34, mean age: 29) were scanned on a 3T Magnetom Verio (Siemens Healthineers, Erlangen, Germany) MRI scanner

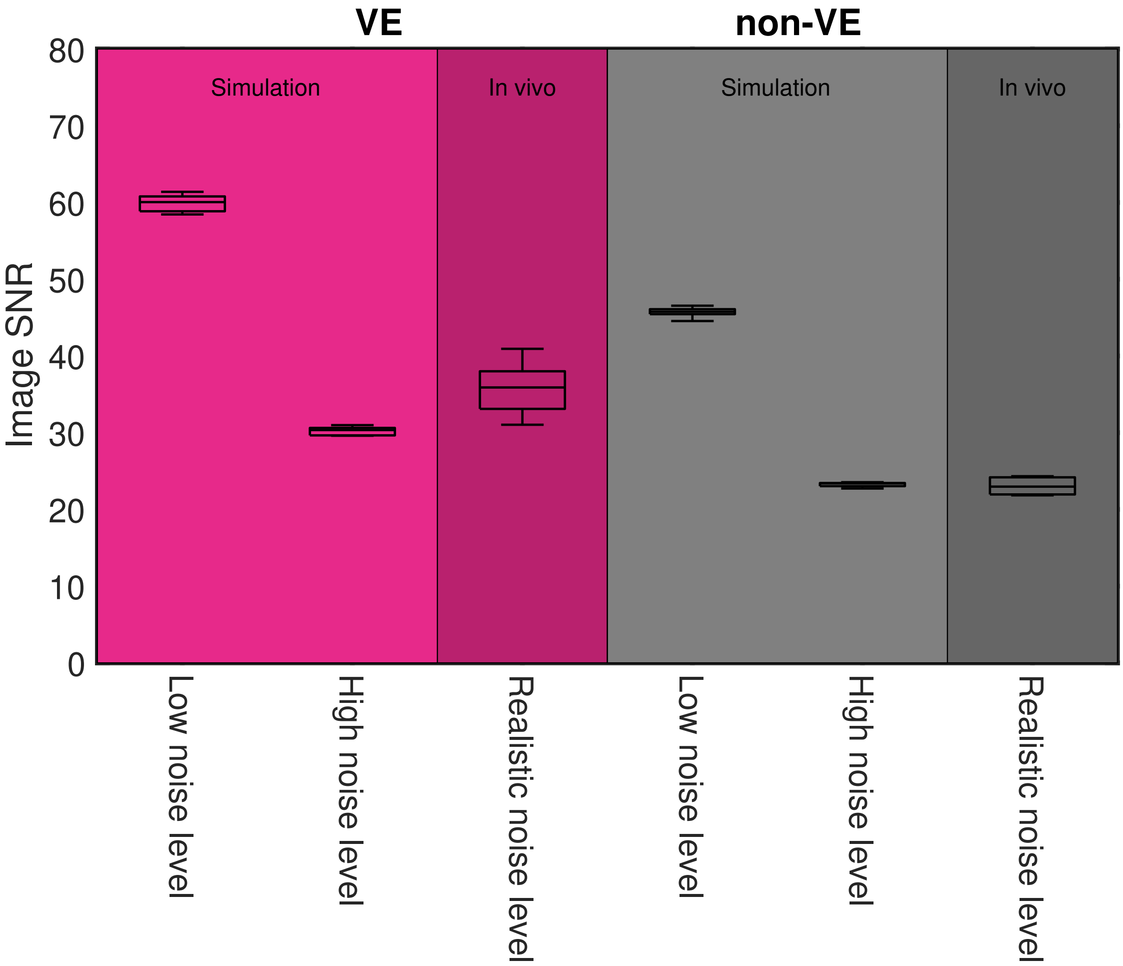

Figure 3.2: Image SNR comparison between the in vivo acquisitions and the simulations in a fully sampled non-regularised reconstruction. Because both the VE and non-VE acquisitions are fully sampled, the VE images have the expected factor of $\sqrt{2}$ higher SNR than non-VE due to doubled acquisition time.

using a 32-channel head coil, with a dynamic 2D golden-angle radial VE-ASL and non-VE-ASL sequence, similar to the implementation previously described by Okell [70]. Data from five subjects were used for the main comparison of VE and non-VE at matched scan times, and 1 subject was used to study the generalisability of the method at higher spatial resolution. All in vivo data were acquired under a technical development protocol approved by the local ethics committee.

Vessel-selective labelling was performed with pseudo-continuous ASL using transverse gradients (G_{xy}) of alternating polarity applied between the RF pulses to modulate the inversion efficiency across the labelling plane [5]. The labelling plane was set just below the confluence of left and right vertebral arteries. For this study, the vertebral arteries were treated as a single artery to allow a 4-cycle Hadamard encoding scheme to be performed. Hadamard encoding ensures SNR optimality, but is limited to encoding matrices of size 1, 2, or multiples of 4 [71]. This means that to encode the two vertebral arteries separately with optimal encoding would require a Hadamard encoding of order 8, which again would double the number of encodings needed compared to 3-vessel encoding. With more than three vessels, it can also become challenging to fully tag or control all vessels based on the vessel geometry. The labelling was performed inferior to where the vertebral arteries merge to form the BA to ensure artifacts associated with the labelling plane did not overlap with the imaging region. The location of the labelling plane, as well as transverse modulation of inversion efficiency for VE, was set on subject-by-subject basis based on a quick time-of-flight angiography pre-scan.

A spoiled gradient echo 2D golden angle radial readout (TR = 11.73 ms, TE = 5.95 ms, FA = 7°) was initiated immediately after the pseudo-continuous ASL labeling pulse train (labeling duration: 1000 ms) (figure 3.3). Each preparation was preceded by a presaturation module for background suppression, and 108 radial spokes were read out during the 1266.8-ms-long imaging period. These 108 spokes were split into 12 frames of 9 spokes in the reconstruction. The same set of 108 spokes were read out for each encoding before moving on to the next set of 108 spokes that were ordered, such that for each frame, each shot (providing 9 new spokes per frame), carried on the golden-angle ordering from the previous shot that had finished, as described in [70], and summarised in table 3.1.

The images for the five subjects used to compare VE against non-VE at matched scan time were reconstructed with 1.1 × 1.1 × 50.0 mm spatial resolution (matrix size 192 × 192) and 105.57-ms temporal resolution. The high-resolution data set was acquired with 0.68-mm^2 isotropic in-plane resolution and a matrix size of 320 × 320. A total of 34 ASL preparations were needed for each encoding to fully sample the 1.1-mm^2 resolution images, and 56 for the 0.68-mm^2 resolution images. The ASL preparation module was repeated every 2400 ms, so for 34 shots and 4 encodings (VE) the total scan time to reach the Nyquist limit (306 spokes per

Spoke:	Data acquisition along a straight line through the centre of k-space at some angle. If two spokes have the same angle, they are referred to as the same spoke.
Frame:	A number of subsequently acquired spokes. The number of subsequent spokes in a frame along with TR sets the temporal resolution.
Shot:	Multiple frames acquired after an ASL preparation module. Multiple shots are acquired so that frames with a set temporal resolution can contain more spokes than limited by the sequence TR. Each shot carries on the ordering of spokes such that each frame has a set of spokes with a golden ratio distribution.
Encoding:	Different ASL preparation of the magnetisation. Each shot is repeated for every encoding.

Table 3.1: Glossary for radial golden ratio ASL data acquisition method

frame) was 5 minutes 26 seconds, and for 2 encodings (non-VE) the total scan time was 2 minutes 43 seconds. For the high-resolution VE data (56 shots, 4 encodings), the total scan time for R = 1 was 8 minutes 58 seconds.

An oversampled 1.1 mm^2 resolution data set (acquisition time 10 min 53 s) for both the non-VE (R = 0.25) and VE (R = 0.5) cases were acquired to be used as ground truth. Then, independently acquired test data sets for both VE and non-VE (total acquisition time 5 min 26 s each) were used to assess the reconstruction method. Before reconstructing, the test data were split into multiple subsets by grouping sequentially acquired spokes, such that the number of spokes in a group corresponded to a specific acceleration factor. For example, subset 1 at R = 2 would include the first 153 spokes in each frame (306 needed for R=1), and subset 2 would include the following 153 spokes. The images were reconstructed at three different levels of acceleration with matched scan time between non-VE and VE: i) High undersampling - R = 34 for VE (maximum acceleration as this only used one single ASL preparation for each encoding) and R = 17 for non-VE, scan time 10 s, ii) Medium undersampling - R = 8.5 for VE and 4.25 for non-VE, scan time 38 s, iii) Low undersampling - R = 2 for VE and R = 1 (no undersampling) for non-VE, scan time 2 min 43 s.

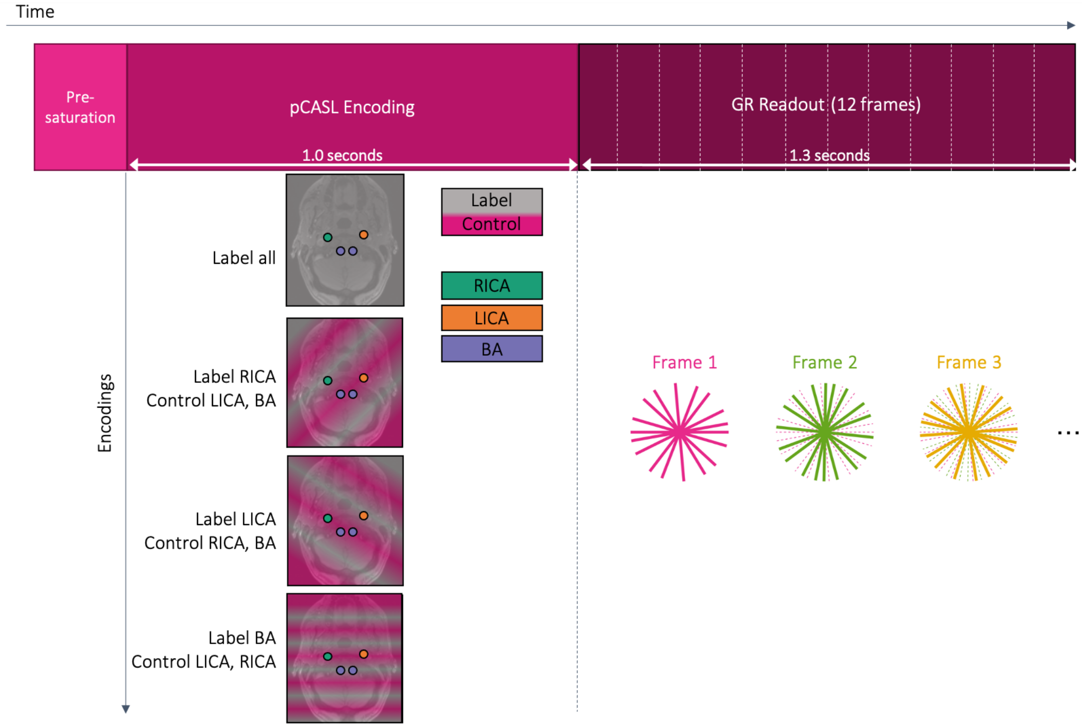

Figure 3.3: The imaging sequence consisted of a pre-saturation module for background suppression, pCASL labelling, and a spoiled gradient-echo readout in a radial golden ratio (GR) trajectory. The continuous readout was separated into 12 frames in reconstruction such that each frame had the same but interleaved trajectory. Four different encodings were used to encode the RICA, the LICA, and the BA that was tagged in the vertebral arteries that were close together and therefore encoded as a single vessel. The encodings were generated with transverse gradients within the pCASL pulse train that generates spatial modulations of inversion efficiency across the labelling plane.

3.2.4 Pre-processing

The images were reconstructed and decoded in the complex domain throughout. Therefore, the data were sensitive to phase errors. Because of this, the in vivo data was pre-processed with a phase correction step that was applied to account for B_0 drift during the scan. The aim of the phase correction is to minimise the phase difference between the same spokes ($\mathbf{k}_n$) in different encodings ($n = \{0, 1, 2, 3\}$) with a scalar phase correction factor ($e^{i\theta_n}$). This problem can be written as:

$$\min_{\theta_n} |\mathbf{k}_0 - \mathbf{k}_n e^{i\theta_n}|_2^2 \tag{3.5}$$

and it is minimised when $\theta_n = \arg(\mathbf{k}_0^H \mathbf{k}_n)$. This cost is sensitive to phase changes of the static tissue since the overwhelming majority of signal originates in the static tissue, and not in the blood signal. By phase aligning the data from the different encodings, complete static tissue cancellation is enforced, which cleans the background and also improves the blood signal.

Coil sensitivity calibration images were reconstructed by combining k-space data across temporal frames to give a fully sampled or near fully sampled temporal average image per coil. These images were then used to estimate the relative coil sensitivity profiles for every point in space using the adaptive combine method [72]. These estimated coil sensitivity profiles were used for generating the encoding operator, **E**, for each subject/simulation. To improve speed and reduce the memory burden of the reconstruction, the 32 coil channels were compressed to just 8 channels using singular value decomposition [73].

3.2.5 Reconstruction

Both the simulated and in vivo data were reconstructed using the same method. Image reconstruction was achieved by the optimisation of a non-linear cost function:

$$c(\mathbf{m}) = \frac{1}{2} \left|\mathbf{E}\mathbf{m} - \mathbf{s}\right|_2^2 + \lambda_1 \left|\mathbf{m}\right|_1 + \lambda_2 \left|\nabla_t \mathbf{m}\right|_2^2 \tag{3.6}$$

In the cost function **m** is the unknown image (with the vessel components and the static tissue at all time points concatenated), **E** is the image acquisition operator, ∇_t is a finite difference operator in the temporal domain, and **s** is the raw k-space signal. Here, the first term imposes data consistency on the reconstruction, the second term imposes image domain sparsity, and the third term enforces temporal smoothness. λ_1 and λ_2 weigh the importance of the regularising terms against the data consistency term. This cost function was minimised using the fast iterative shrinkage thresholding algorithm (FISTA [29]), using a Toeplitz embedding formulation to replace NUFFTs with FFTs for reduced computation time [74] (see appendix A for further details).

The regularisation factors in equation 3.6, were determined experimentally by a grid search across a range of potential values. The (λ_1, λ_2) search space was chosen to be wide enough to ensure it fully characterised the target optima. The trialled λ_1's ranged from 0 to 10^{-5} in steps of 10^{-6} for all acceleration factors both in vivo and for the simulated data. For the highest undersampling factors (R = 17 and 34), λ_2 was varied from 0 to 6 in steps of 0.2 for the in vivo search, and 0 to 2 in simulations. For medium undersampling (R = 4.25 and 8.5), λ_2 was varied 0 to 10 in steps of 1 for both the in vivo and simulation case. For low undersampling (R = 1 and 2) it was varied in steps of 2 from 0 to 20.

For the simulations, the combination of λ_1 and λ_2 with the highest correlation with the ground truth (as explained in section 3.2.6 below) for each acceleration factor and noise level was used. In vivo, the average performance across all subjects and eight subsets of data (except for the VE R = 2 and non-VE R = 1 case where only two subsets of data were acquired) were calculated and the regularisation factors that produced the highest correlation score on average were chosen and used for all further analysis. Subject specific optimal ranges of regularisation factors, that resulted in less than 1% quality reduction from the subject specific optimum were also calculated. This was done in order to confirm that the group optimum was reasonable for all subjects. The overlap of these subject-specific optimal regularisation ranges was also inspected to see how subject dependent the optimal regularisation factors were.

3.2.6 Analysis

All reconstructions were compared against ‘ground truth’ images. For the simulations, the input image was used directly for comparison. For the in vivo data the oversampled acquisition was reconstructed with minimal regularisation applied for denoising ($\lambda_1 = 10^{-6}$, $\lambda_2 = 0$) and used as ground truth.

For quantitative assessment of image quality, non-overlapping vessel-specific masks (figure 3.4) were applied to both the reconstruction and the ground truth as it was important that the quantitative assessment focused on the relevant pixels in

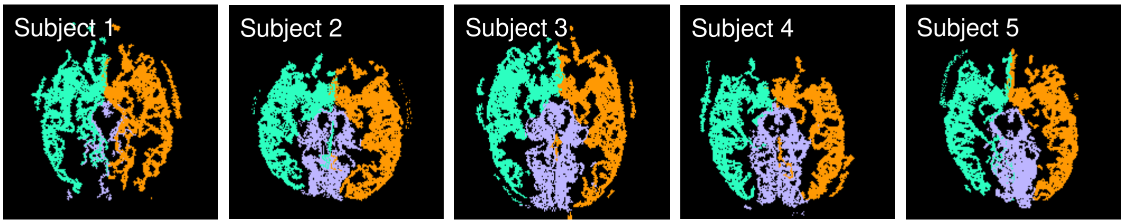

Figure 3.4: Subject-specific masks used in the assessment of image reconstruction. In pixels where blood supply was mixed, the most intense vessel component in the ground-truth image was chosen.

the sparse images to avoid bias due to irrelevant artifacts far away from the vessels (e.g. from eye motion). Focusing on the vessels does thus not capture the presence of artifacts far from the vessels, which could be a weakness of using a mask for assessment. However, the main artifacts that degrade image quality in radial imaging with CS reconstruction are streaking and loss of faint vessels. Streaking would appear near sources of signal (vessels), and loss of vessels are also local artifacts. Masking, does therefore not substantially limit the image quality assessment.

The masks were generated by thresholding the fully sampled reconstruction and then dilating the mask with a 3×3 kernel, $\begin{bmatrix} 0 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix}$, to include background pixels surrounding the blood vessels. The masks were then applied to each frame of each vessel component (or the total vessel component for non-VE). The Pearson correlation coefficient (r) between the ground truth and reconstructed pixel values across all time points within the masks were then calculated. This metric captures signal variation, which allows for a joint assessment of faint vessels, arteries with strong signal, and the signal-less background voxels within the mask. Before settling on the correlation coefficient as the image quality metric, structural similarity index [75] (section 2.4.3) and normalised-root-mean-squared error (section 2.4.2) were trialled, but the results were inconsistent with visual assessment of image quality.

When comparing non-VE against VE, the correlation coefficients, r, for each vessel mask were Fisher transformed to a z-score ($z = \text{arctanh}(r)$) to make the distributions of correlation coefficients more Gaussian. This then allowed Student's t-tests to be performed to determine statistical significance at a 98.3% confidence

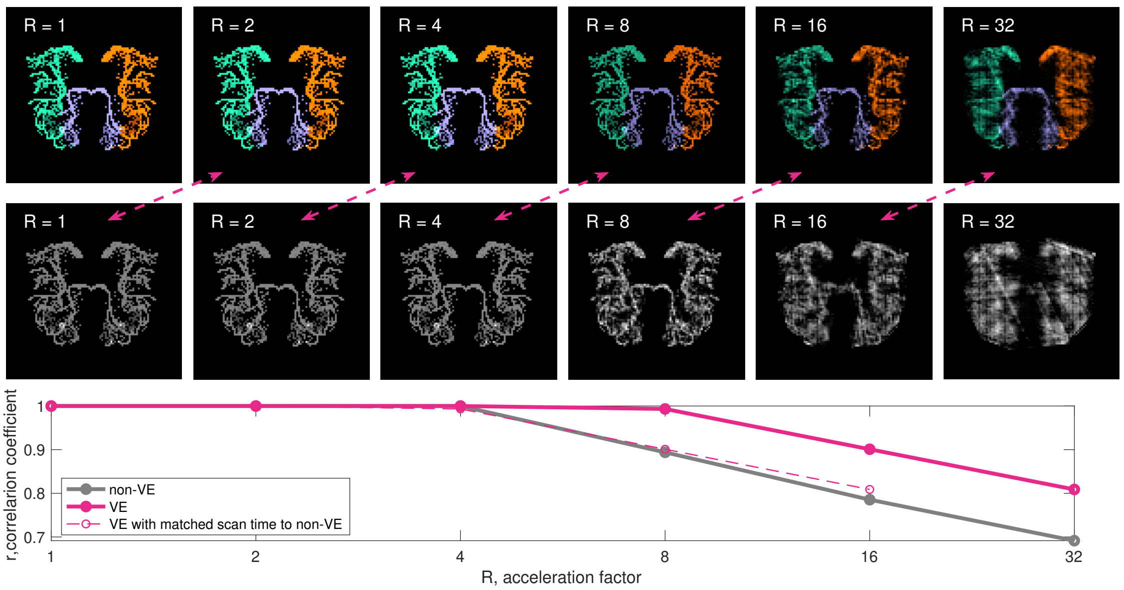

Figure 3.5: Numerical phantom simulations comparing VE and non-VE with different amounts of undersampling in a simple noiseless phantom. For R > 4, the quality of the non-VE reconstruction starts to decrease rapidly. For VE, this only occurs at approximately twice the undersampling factor, R > 8. The dashed line on the bottom graph shows a shifted copy of the VE line, to illustrate the reconstruction quality of a 3-vessel VE angiogram at the same acquisition time as the non-VE angiogram. The dashed arrows connect the VE and non-VE angiograms that would have the same acquisition time.

interval (95% with added Bonferroni correction for multiple comparisons of the three vessel components).

3.3 Results

3.3.1 Effect of Increased Relative Sparsity

In the simple numerical phantom, where the only difference between non-VE and VE was the ratio of non-zero to zero voxels, the VE data was reconstructed more robustly at higher undersampling factors than non-VE (figure 3.5). VE and non-VE data could be reconstructed essentially perfectly at low undersampling factors, but for R > 4 (less than 25% sampling), the reconstruction of the sparser VE simulation outperformed the non-VE version, achieving approximately matched performance for twice the undersampling factor, negating the factor of two time penalty that would be needed to perform three-vessel VE instead of non-VE angiography. Qualitatively, increased blurring and streaking artifacts were observed in the non-VE case.

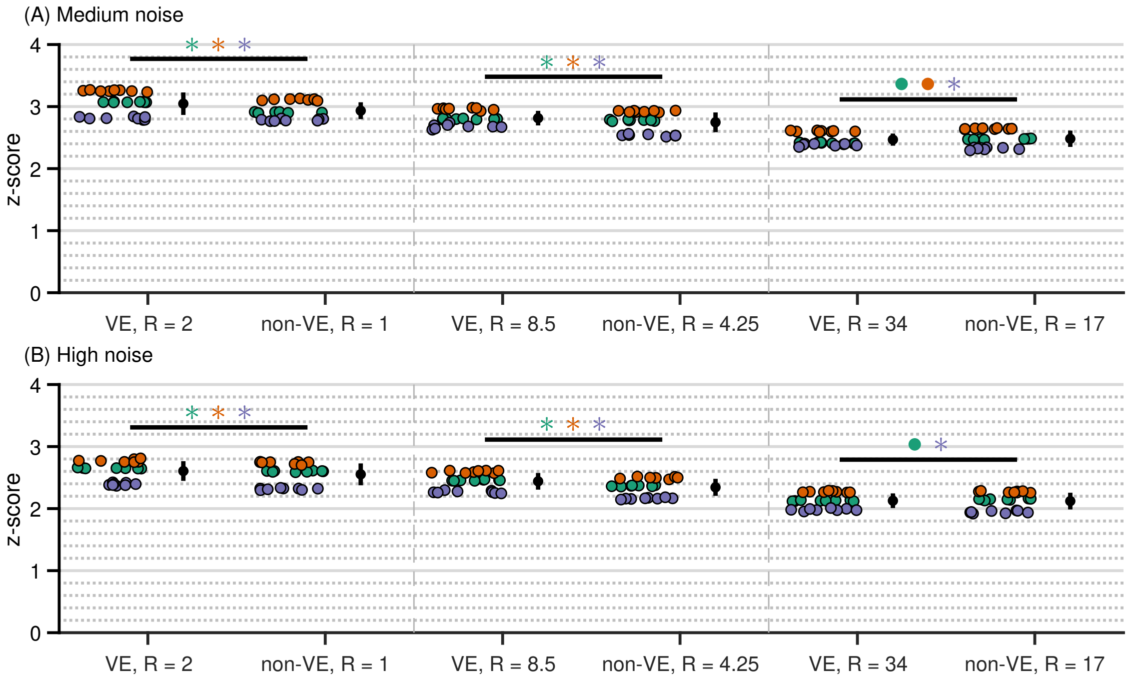

Figure 3.6: Reconstruction quality in simulations with medium noise (A) and high noise (B). Each scatter point represents the Fisher transformed correlation coefficient calculated in a mask (RICA, green; LICA, orange; BA, purple) for one reconstruction. Statistical significance between the time-matched non-VE and VE groups is represented by an asterisk ($*$) if VE had a higher correlation coefficient and a dot ($\bullet$) if non-VE did. Each pair of datasets, separated by dashed vertical lines, would have matched scan time.

3.3.2 Simulations on Real Data

Similar results to the numerical phantom were observed in the real data simulations (figure 3.6). However, with no added noise the VE and non-VE reconstruction quality was high ($r > 0.99$) at all acceleration factors, and no difference was found between VE and non-VE. The realistic experiment included the use of multiple coils and temporal regularisation that was not present in the simple experiment, which could explain this discrepancy. With added noise and simulated matched scan times (equal SNR, but twice the undersampling for VE), VE was reconstructed marginally, but significantly, better ($p < 0.01$) than non-VE for low and medium acceleration factors in both medium and high noise conditions. At high acceleration factors, the results were more varied and VE performed better than non-VE in some vessels but worse in others and for some there was no statistically significant difference.

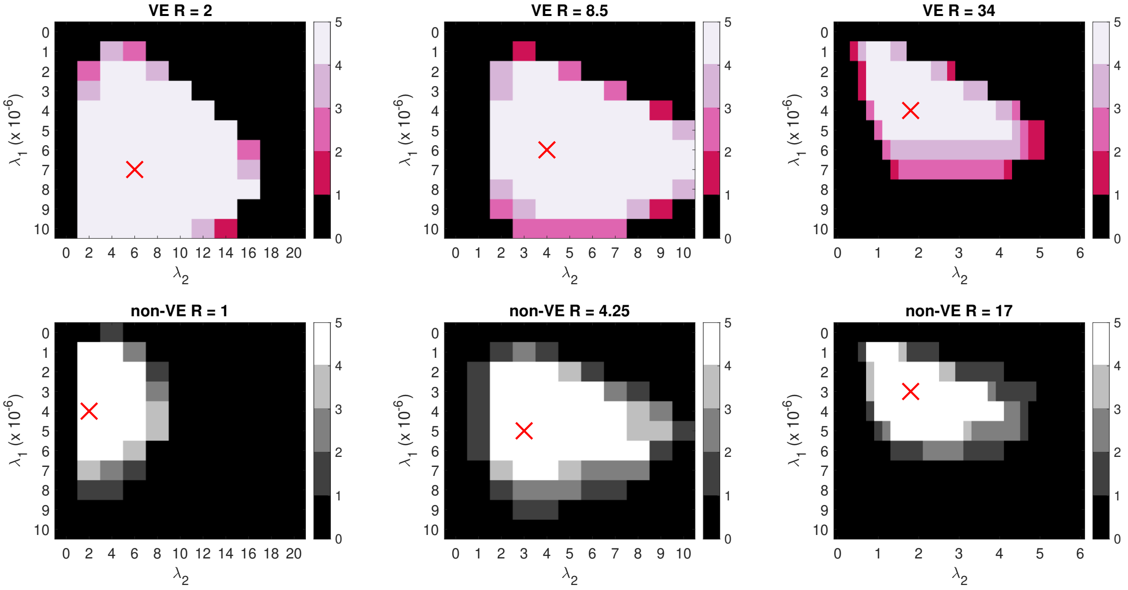

Figure 3.7: Optimal regularisation factors (marked with red “x”) were within the optimal area (within 1% of optimum) for every subject. The colour scale represents how many subjects had optimal reconstruction at each combination of regularisation factors

3.3.3 Regularisation Optimisation

The optimal regularisation factors for the in vivo reconstructions did not vary considerably between different subjects and their optimal ranges (within 1% from optimum) had considerable overlap at all acceleration factors. The group optimum was within the subject specific optimal ranges for both VE and non-VE (figure 3.7).

Varying the regularisation factor within the overall optimal range resulted in a sensitivity and specificity tradeoff, i.e. reduced noise against the cost of losing visibility of small vessels. Some examples of this tradeoff are shown in figure 3.8. For the sake of comparing VE with non-VE reconstructions in an unbiased way, the group optimum regularisation factors for each acquisition method were used for all further analyses. The optimal regularisation factors for each acceleration factor and subject or noise level are summarised in table 3.2.

Using both regularisation terms improved the overall reconstruction quality in all cases. Average correlation coefficients for the in vivo data improved 4.6-55.2% by having a non-zero λ_2 (difference between figure 3.9(C) and (D)), 2.3-8.9% by having a non-zero λ_1 (difference between figure 3.9(B) and (D)), and 15.3-98.7% by having both regularisation factors non-zero(difference between figure 3.9(A) and

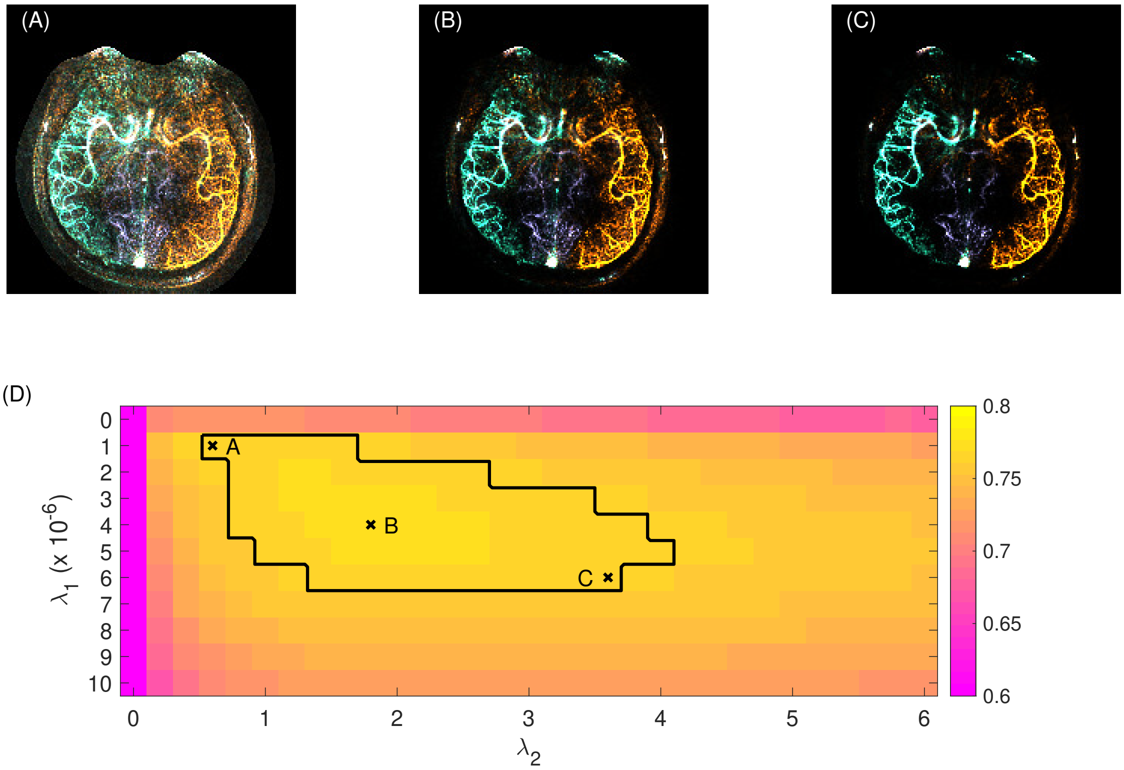

Figure 3.8: The same raw data set (one of the in vivo vessel-encoded R = 34 subsets) reconstructed with 3 different combinations of regularisation factors: minimal regularization, and therefore more noise (A); optimal regularisation based on the average correlation coefficient (B); and maximal regularisation, resulting in a heavily denoised reconstruction (C). (D) shows the average correlation coefficient across all subjects at each combination of regularization factors. The black border in (D) represents the area where the reconstruction was within 1% of optimum, and the black "x"s denote the regularization factors used for (A), (B), and (C) respectively.

In vivo	Subj 1		Subj 2		Subj 3		Subj 4		Subj 5		Overall	
	λ_1	λ_2	λ_1	λ_2	λ_1	λ_2	λ_1	λ_2	λ_1	λ_2	λ_1	λ_2
non-VE R = 1	4.00E-06	4	4.00E-06	4	4.00E-06	2	4.00E-06	2	5.00E-06	2	4.00E-06	2
VE R = 2	7.00E-06	6	7.00E-06	6	7.00E-06	6	7.00E-06	6	7.00E-06	4	7.00E-06	6
non-VE R = 4.25	4.00E-06	4	4.00E-06	4	5.00E-06	4	4.00E-04	3	6.00E-06	3	5.00E-06	3
VE R = 8.5	6.00E-06	5	6.00E-06	4	6.00E-06	4	6.00E-06	5	7.00E-06	4	6.00E-06	4
non-VE R = 17	3.00E-06	1.8	3.00E-06	2.2	4.00E-06	2.2	3.00E-06	1.6	4.00E-06	1.8	3.00E-06	1.8
VE R = 34	1.00E-06	0.8	4.00E-06	2	4.00E-06	2	4.00E-06	2	5.00E-06	2	4.00E-06	1.8

Simulations	$SNR_k = \infty$		$SNR_k = 185.8$		$SNR_k = 92.8$	
	λ_1	λ_2	λ_1	λ_2	λ_1	λ_2
non-VE R = 1	0	0	1.00E-06	2	4.00E-06	2
VE R = 2	0	0	1.00E-06	2	7.00E-06	4
non-VE R = 4.25	1.00E-06	0	3.00E-06	1	5.00E-06	3
VE R = 8.5	1.00E-06	0	3.00E-06	2	7.00E-06	3
non-VE R = 17	1.00E-06	0.2	2.00E-06	0.6	4.00E-06	1.4
VE R = 34	1.00E-06	0.2	2.00E-06	0.6	5.00E-06	1.4

Table 3.2: Optimal regularisation factors for in vivo and simulation

(D)). This suggests that the effect of a non-zero λ_2 is somewhat greater for image quality than a non-zero λ_1, but that using both instead of just one has an even greater effect on image quality (as can also be seen as the sharp difference between the λ_1 or $\lambda_2 = 0$ row/column to the non-zero combinations in figure 3.8(D).

Similarly, in simulations with non-zero noise a 1.0-13.0% improvement was observed for non-zero λ_2, 0.1-10.0% improvement for non-zero λ_1, and 2.2-88.3% improvement by having both regularisation factors non-zero. Visually, the value of having non-zero values for both regularising terms is clearly shown in an example reconstruction in figure 3.9, with non-zero λ_2 causing better delineation of the vessels, and non-zero λ_1 reducing noise and noise-like artifacts. The blurring caused by having λ_2 set to zero seems to reduce the correlation coefficient more than background noise caused by $\lambda_1 = 0$. The larger effect of a non-zero λ_2 in vivo than in simulation could be due to the presence of temporally varying physiological noise in vivo, which benefits from temporal regularisation more than the simulations that had no temporally varying noise component.

3.3.4 In Vivo Comparison of VE and non-VE

Generally, no statistically significant difference in image quality was found between VE and non-VE images acquired at the same scan time despite VE requiring twice the undersampling factor. A single exception was the high acceleration RICA where the non-VE correlation coefficients were marginally higher ($p < 0.01$). The in vivo results are summarised in figure 3.10.

Qualitatively, the image quality of time-matched VE and non-VE images were similar as suggested by the quantitative results. Example reconstructions of all five subjects are also shown in figure 3.10. At high R, a loss of faint features compared with the ground truth was apparent for both non-VE and VE and some subsets of the full data sets included artifacts potentially caused by motion.

The temporal dynamics were also generally well conserved across acceleration factors. The signal was well preserved in the late frames at moderate acceleration factors, but at the highest acceleration factors some residual aliasing remained in

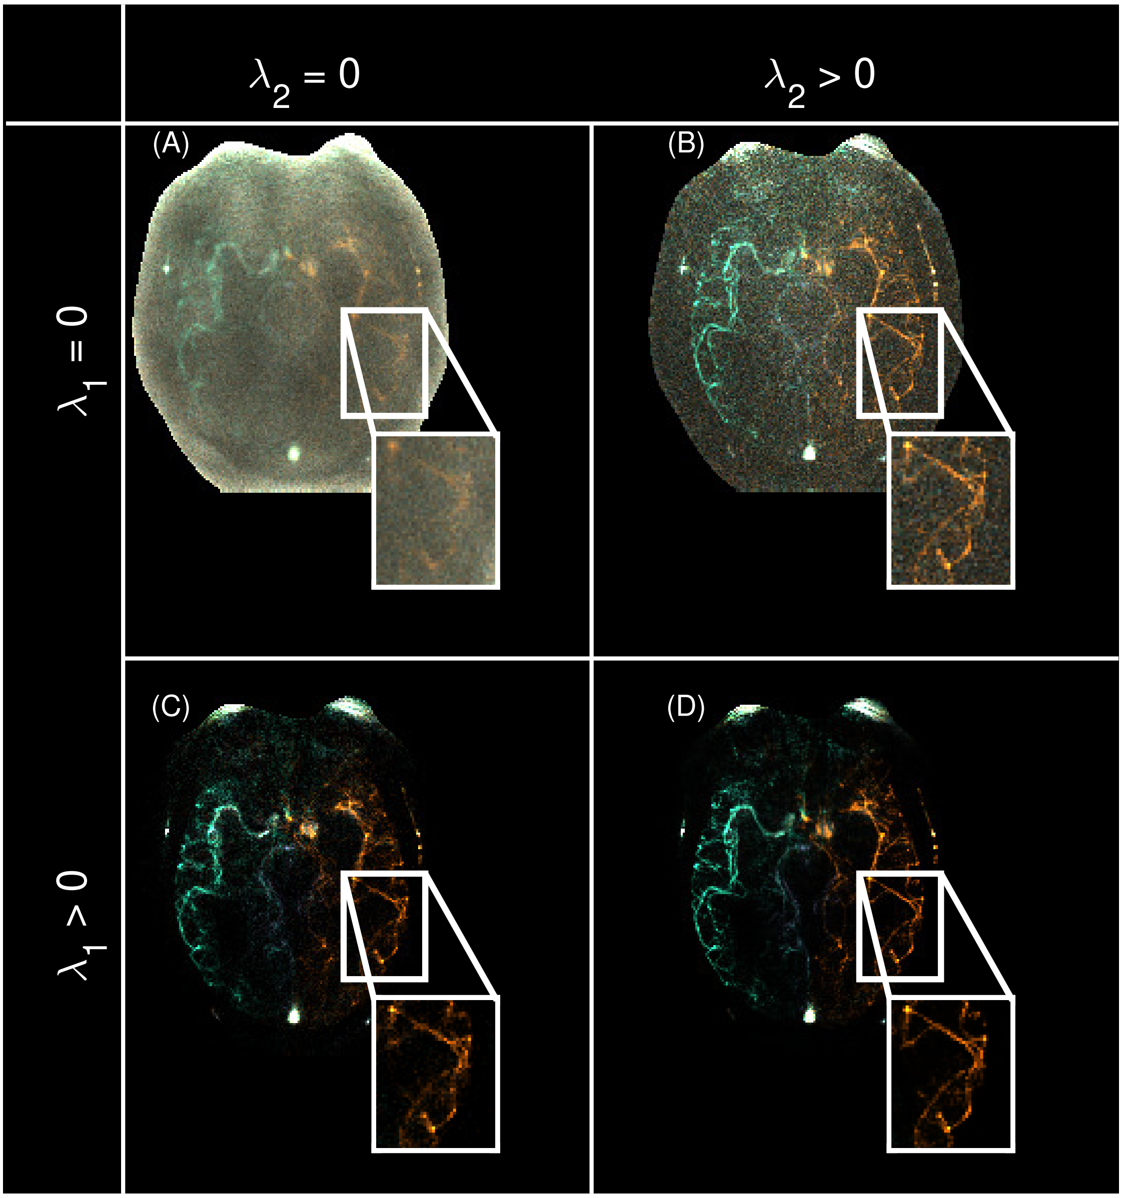

Figure 3.9: An example of the R = 34 VE-ASL reconstruction (temporal average) with no regularisation (A), only the L2 temporal smoothness constraint (B), only the L1 sparsity constraint (C), and both regularizing terms included (D). The L2 term sharpens the vessels, and the L1 term denoises the background. Both regularisation terms improve the overall reconstruction quality.

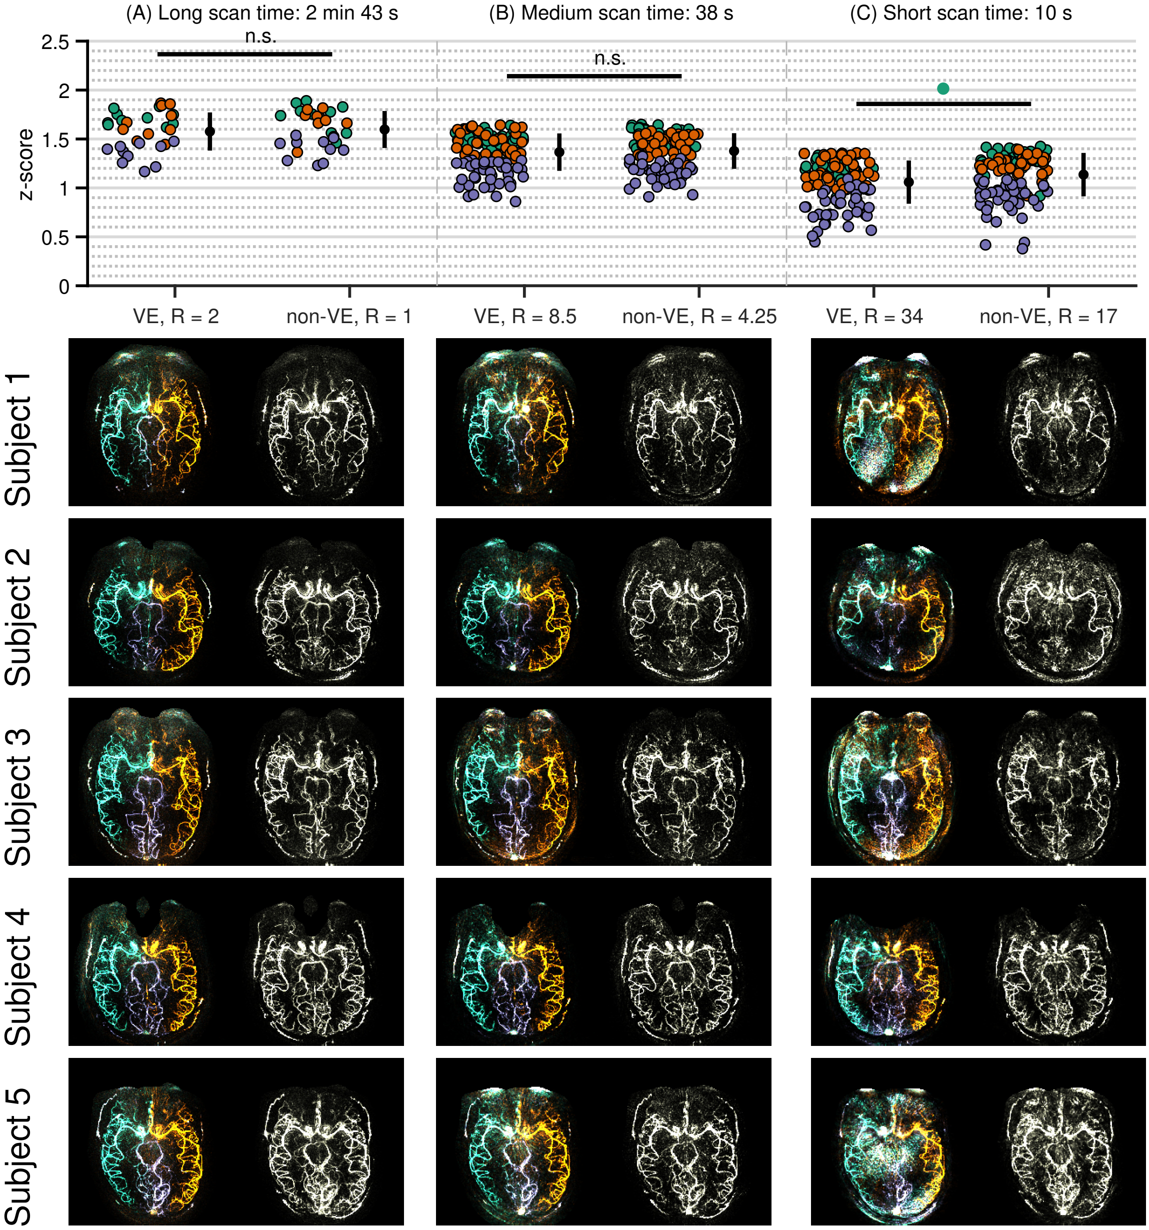

Figure 3.10: In vivo reconstructions: fully sampled non-VE versus R = 2 VE (A); moderate acceleration (R = 4.25 non-VE versus R = 8.5 VE) (B); and high acceleration (R = 17 non-VE versus R = 34 VE) (C). Below the quantitative graph examples of the reconstruction quality for one subset in all subjects for time-averaged VE and non-VE.

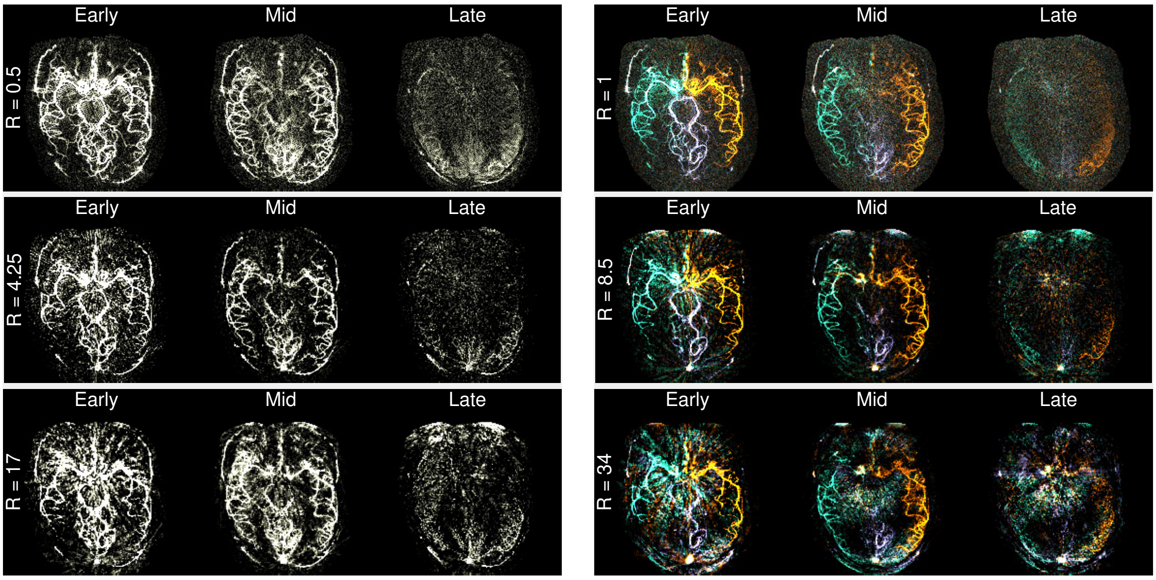

Figure 3.11: Temporal dynamics in an example subject at varying acceleration factors for a non-VE acquisition (left) and a VE acquisition (right). The early time point is frame 1, the mid-timepoint is frame 6, and the late time point is frame 12.

the later frames (figure 3.11) for both the non-VE and VE images. Figure 3.12 shows the temporal profile of the signal averaged in two 3 × 3 voxel regions in proximal and distal vessels in an example subject. In the distal vessel the SNR is lower and the temporal signal is noisy even in the ground truth case. The temporal regularisation smooths the signal and preserves overall shape.

3.3.5 High Resolution Imaging

In under a minute, (R = 9.33, scan time 58 s) very high-quality images could be acquired at high resolution as can be seen in figure 3.13. The regularisation factors used to create this image were on the lower end of the 1% optimal region ($\lambda_1 = 2 \times 10^{-6}$, $\lambda_2 = 2.0$) as higher regularisation factors removed many of the fainter vessels. This highlights the option of choosing regularisation factors based on preferred tradeoff between sensitivity and specificity.

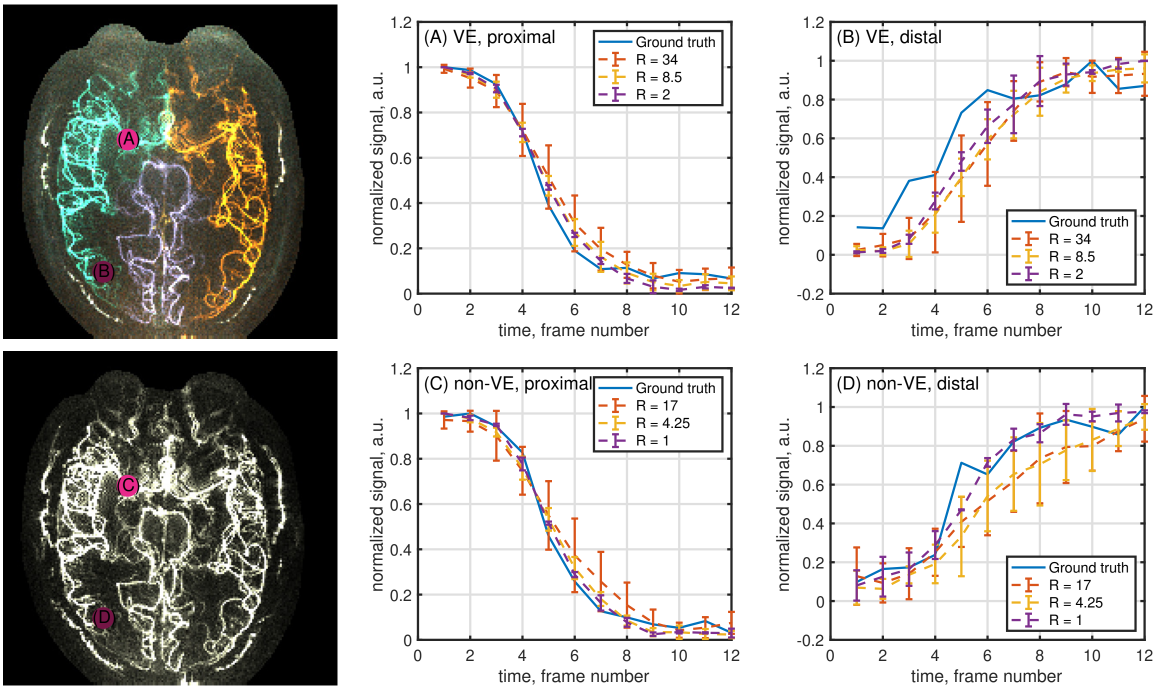

Figure 3.12: Temporal profile in two regions of interest, in a proximal vessel (A,C) and in a distal vessel (B,D), with blood supply from the RICA in one example subject. The error bars indicate the SD of the signal measured from reconstructions of different subsets of the raw data at each acceleration factor. (A), and (B) show points in the VE-ASL angiogram, (C) and (D) show the same points in a non-VE-ASL angiogram.

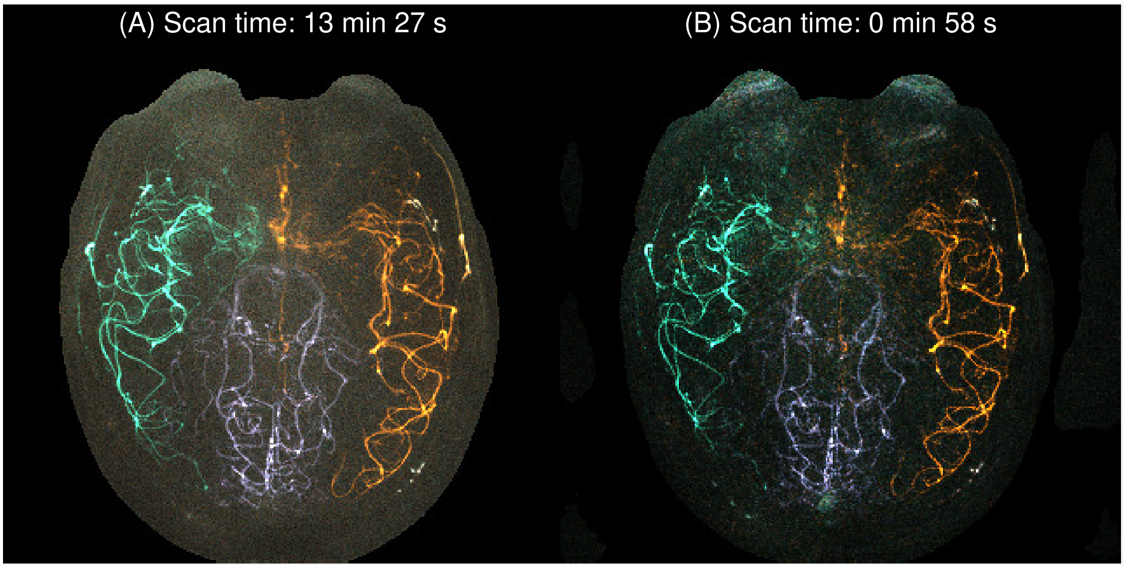

Figure 3.13: (A) High-resolution image from oversampled VE scan (R = 0.67). (B) Highly accelerated scan (R = 9.33)

3.4 Discussion

3.4.1 Relative Sparsity as Driving Factor

As hypothesised, the simple simulation experiment showed that relative sparsity (proportion of non-zero voxels to total number of voxels reconstructed) can drive a L1-regularised reconstruction. This agrees with underlying theory of compressed sensing [65]. How much relative sparsity drives the reconstruction quality compared to other factors such as SNR, and the spatial distribution of non-zero voxels, are topics worth further consideration. The theoretical and practical limitations of this method to get extra information "for free" in a system where relative sparsity increases should be further explored beyond this initial feasibility study.

One potential extension that was not explored in this thesis is applying this method to VE-ASL with labeling above the Circle of Willis, which would require more encodings because there are more vessel branches, but each decoded image would be sparser, allowing for potentially higher acceleration. However, practical issues such as with achieving an ideal Hadamard encoding for more complicated vessel geometries, might limit the improvements that the increased sparsity alone buys. A strategy for optimising encodings for complex geometries has, however, been proposed previously [76].

To further aid reconstruction, one could think that the non-sparse nature of the static tissue needs to be considered. However, as will be shown in the next chapter, the static tissue reconstruction does not interact with the reconstruction of the vessels when the same sampling trajectory is used for every encoding, as was the case in this chapter. If different sampling trajectories are used, the non-sparse nature of the static tissue does require consideration. Although good images of the static tissue signal are not necessary from a clinical perspective, the static tissue component needs to be reconstructed well in order to correctly decode the blood signal. Further improvements could be achieved by applying a sparsifying transform such as the wavelet transform, or another prior. In similar work on non-VE pulsed ASL (PASL) angiography wavelet regularisation on the control image was reported to improve vessel delineation [8]. By transforming the static tissue to a sparser domain, it should allow for better reconstruction in a compressed sensing framework.

3.4.2 Temporal Regularisation

In the same vein as increasing the spatial dimensionality by increasing the number of vessels encoded, the dynamic acquisition also allows you to take advantage of structure and redundancy in the temporal domain. In this work, the L2 norm of temporal finite differences was used to regularise the reconstruction. It greatly improved the reconstruction results beyond the L1 spatial sparsity alone. We observed that the introduction of the temporal regularisation mostly improved spatial delineation, which could be explained by the sharing of high spatial frequency k-space information across time frames. Over-regularised temporal finite difference constraints cause temporal blurring of the signal, but that was not observed in this work. The blood signal was changing smoothly enough, as predicted by the physiological kinetic model of ASL angiography [66]. In previous works, temporal constraints have been enforced by using sliding-window acquisitions [77] or with compressed sensing using L1 constraints in temporal total variation (TV) frameworks [78], or in the temporal frequency domain [79]. These approaches can, however, have unwanted effects such as increased blurring for the sliding-window method, temporal stair-casing artifacts for TV [80], or artificially introduced periodic behavior in the temporal frequency domain. Model based non-linear reconstruction is another option that has been explored in perfusion ASL [81], and whether any of these approaches could improve reconstruction for VE-ASL angiography could be studied further.

3.4.3 Lambda Optimisation

Among the five volunteers scanned for the 1.1 mm^2 resolution data, the difference in optimal regularisation factors was marginal and there was considerable overlap among their optimal regions. This suggests that once the reconstruction has been optimised for an acceleration factor and imaging protocol it should be robust for new subjects. To test this rigorously, a cross validation approach must be used on a larger number of subjects once this method is developed for clinical use. Because the optimal regularisation factors depend on the properties of the data itself (inherent sparsity and temporal smoothness), one has to be careful with extending this conclusion to patients and other populations. There was some variability in head size, angle of imaging slab, and two of the volunteers exhibited considerable mixing of RICA and BA blood due to an asymmetrical Circle of Willis configuration, but patients with, for example, arteriovenous malformations could have considerably lower image domain sparsity due to the presence of additional abnormal vessels that can affect optimal regularisation factors. Further studies in appropriate patient groups are therefore also required to determine how generalisable this result is.

In this study the Pearson's correlation coefficient, r, was used to objectively optimise the regularisation parameters and define the quality of the reconstruction. We found the Pearson's correlation coefficient to be a robust metric that corresponded well with visual perception of image quality, which was not the case for other metrics that were tested in preliminary work. As discussed in section 2.4, normalised-root-mean-square-error (NRMSE) is a straightforward metric but it does not correspond well with perceptual quality, as also reported in[38]. Structural similarity index (SSIM) [75] was developed specifically to correspond with visual perception worked well within a data set to e.g. tune the parameters, but was not suitable to compare different datasets as it was sensitive to the absolute scaling of the signal. Often, the signal-to-noise ratio (SNR) is reported as a summary metric for image quality, but it was also found to be unsuitable for this type of reconstruction, as the non-linear nature of the reconstruction method can cause high SNR when image fidelity is poor, by both attenuating the signal (high bias) and removing all noise (low variance). The masking improved robustness as otherwise the result was mainly driven by large areas with no signal and regions containing artifacts such as eye motion. In future work, tuning of the regularisation factors could be done manually by experts to balance specificity (i.e. noise removal) and sensitivity (preservation of faint signals) for optimal clinical utility.

3.4.4 Feasibility of Accelerating VE and non-VE-ASL

These results demonstrate the feasibility of acquiring vessel-encoded images without increasing scan time compared to fully sampled non-vessel-encoded. Similar image quality was obtained with R = 2 VE and R = 1 non-VE imaging in vivo. Further reductions in scan time have also been shown to be possible, both for non-VE and VE-ASL, with matched image quality. At the highest acceleration factor (R = 17 for non-VE and 34 for VE) the main features were still visible with scan times as short as 10 s, although some loss of faint features and artifacts were observed. The required image quality will depend on the clinical application of the technique. For example, if the scan is acquired to add information about mixing of blood from different sources to other angiographic images a highly accelerated scan of lower quality might be sufficient, whereas if it is to be used diagnostically on its own, moderate acceleration factors might be more appropriate (as shown in the high resolution data set).

In simulation the correlation coefficient, r, was consistently higher for VE at low and medium R, indicating not only equivalent but also slightly improved performance of VE over non-VE at matched scan time. However, although statistically significant, this difference was small compared with the effect of increasing or decreasing scan time or varying the SNR. At high R the performance results varied between being in favor of VE and non-VE, and differences in the qualitative assessment of the image quality were small. Therefore, this does not contradict the overall conclusion of this study that VE and non-VE images of similar quality can be achieved at matched scan times even at the highest acceleration factor.

Similarly, in vivo, the non-significant results at low and medium R indicate similar performance level. One potential reason why non-VE achieved significantly higher correlation coefficients in the RICA at high R is that two of the five subjects exhibited mixing of blood supply on the right side (figure 3.10). Mixing provides higher SNR for the non-VE reconstruction as the signal from multiple origins are added together rather than split into two vessel components.

3.5 Conclusion

This study has shown that the additional sparsity provided by vessel-encoding allows us to generate vessel-selective dynamic angiograms of equal quality as conventional non-vessel-selective ASL angiography acquired in the same scan time. This suggests that the vessel-selective information is accessible with no cost of either scan time or data quality, providing extra information about the cerebral vasculature and haemodynamics for free.

This study also demonstrated a method to produce high resolution VE angiograms from less than a minute of scanning data, reducing scan time by approximately an order of magnitude (R = 9.3).

For many clinical applications it would be desirable to extend this technique to 3D. The main reason 2D imaging of a single slab was chosen for this feasibility study, was to be able to acquire ground truth images in reasonable scan times. Extending to 3D will increase the relative sparsity further and should therefore allow for higher acceleration. The extension to dynamic 3D, or 4D, will be considered in the next chapter along with some modifications to the acquisition and reconstruction scheme designed to make the method even more sampling-time efficient.

These creatures you call mice, you see, they are not quite as they appear. They are merely the protrusion into our dimension of vast hyperintelligent pandimensional beings.

— Douglas Adams in *The Hitchhiker's Guide to the Galaxy* 1979

4

Bringing VE-ASL Angiography to Higher Dimensions

Contents

4.1 Introduction 71
4.2 Improved Sampling Efficiency 73
  4.2.1 Theory 74
  4.2.2 Methods 77
  4.2.3 Results 80
4.3 Improved Sparsity 83
  4.3.1 Methods 84
  4.3.2 Results 87
4.4 More Precise Regularisation Terms 90
  4.4.1 Methods 91
  4.4.2 Results 93
4.5 Combining the Extensions to Acquire 4D VE-ASL Angiography 95
  4.5.1 Methods 96
  4.5.2 Results 100
4.6 Discussion 102
4.7 Conclusion 106

4.1 Introduction

In the previous chapter a feasibility study into the use of compressed sensing based non-linear reconstruction methods to accelerate VE-ASL angiography was presented.

The conclusion was that VE-ASL angiography is well suited to a compressed sensing type reconstruction, and that its increased relative sparsity compared to non-VE-ASL allows it to be accelerated more whilst retaining the same image quality using the proposed reconstruction approach.

In this chapter, the framework is extended into 4D (dynamic 3D imaging). The main benefit of acquiring data in 4D, from a clinical point of view, is the ability to visualise the data from multiple angles, and thus get a better understanding of the anatomy and any pathology to use for example in surgical planning [82]. 4D images could also be achieved by acquiring multiple dynamic slices in 2D, however, using a dynamic 3D acquisition also brings technical benefits, such as improved SNR by acquiring data from a larger volume, isotropic resolution (i.e. not limited by slice profiles), and increased relative sparsity of the vessels.

However, much higher acceleration factors are needed to acquire data in 4D with feasible scan times. Dynamic angiography requires both high spatial and temporal resolution to be useful, so fully sampled images could easily take many hours to acquire if required to reach the Nyquist limit. This chapter focuses on improvements to both the acquisition and reconstruction approach from the previous chapter to make 4D imaging possible in clinically reasonable scan times ($<$ 10 min).

Three areas of potential improvement were identified:

1. Improved sampling efficiency.

2. Finding a sparser representation of the object.

3. More tailored reconstruction.

Each of these has the potential to improve the image quality and allow for higher undersampling factors, thus making 4D imaging possible, but they all come with challenges and limitations too. The following sections introduce each one in turn and present simulation experiments regarding their respective usefulness. The most promising improvements are then combined to attempt 4D imaging in vivo.

4.2 Improved Sampling Efficiency

To improve the incoherence of the acquisition operator, and to better leverage mutual information between the encoded images, the 4D image reconstruction problem can be extended to a 5D problem by considering the vessel components as their own dimension. The dimensions of the data are then: space (three dimensions), time (one dimension), and vessel components (one dimension). By considering the vessel components as its own dimension it is possible to begin considering how undersampling artifacts can spread in space, time, and across components. The more the aliased signal is spread out, the easier it is to remove through de-noising based reconstruction methods such as compressed sensing [4]. In this section we consider how to spread artifacts across the component dimension of VE-ASL.

The spreading of artifacts across the vessel components is controlled by jointly varying the vessel encoding scheme and k-space sampling. Because each encoded image is a combination of all the components, they contain complementary information, and sampling each encoded image with a different trajectory in k-space should better leverage this shared information than using the same trajectory for each one, similar to how high frequency information was shared between frames in the previous chapter because each frame used a different sampling trajectory. Using different trajectories for each encoding causes the aliased signal from one component to end up in the other components, as we shall see in section 4.2.1. However, when signals from different components mix together, challenges arise if one component, for example the static tissue component in the case of ASL, has much higher energy than the others, and therefore can swamp the other components signal with its aliased energy. For example, if one component has voxel values 10 times stronger than another, even if there is only 1% residual aliasing from the strong component into the weak component, the image of the weak component will be severely affected. In this section a proposed solution to this problem that allows for aliasing to spread out between the vessel components but not the static tissue component is presented.

4.2.1 Theory

In MRI, the physical property of interest that is imaged is often not measured directly, but is instead inferred from multiple measurements or encoded images; one such example is ASL. In ASL, the blood signal is inferred from a linear combination of encoded images. In non-VE-ASL there are only two encodings, the tag and the control, and in VE-ASL the vessel encoding scheme determines the number of encoded images to acquire. Many modern sampling and reconstruction methods consider the sampling only in k-space, or k-space and time [79, 83–85], instead of considering the sampling across the different encodings.

As shown in the previous chapter, the acquisition operator, $\mathbf{E}$, for VE-ASL angiography can be modelled as the matrix multiplication of three components:

1. The vessel encoding operator, $\mathbf{H}$, that describes how the blood signal originating in different arteries have been modulated for each acquired dataset.
2. The coil sensitivity modulations transform, $\mathbf{S}$, that weights each image according to the spatial modulation of each receive coil.
3. The Fourier transform, $\mathbf{F}$, taking the object between image space and k-space and incorporating the effect of sampling.

Combining all three gives us this relationship:

$$\mathbf{E} = \mathbf{FSH} \tag{4.1}$$

When treating reconstruction and decoding separately, the two simpler inverse problems are solved using $\mathbf{FS}$ for reconstruction and $\mathbf{H}$ for decoding. $\mathbf{S}$ is determined by the hardware design, so cannot be easily modified. Therefore, it will be excluded in this proof-of-principle project, and a single receive coil with spatially homogeneous sensitivity will be assumed. This is equivalent to setting $\mathbf{S}$ equal to the identity operator, $\mathbf{I}$. $\mathbf{F}$ is determined by the sampling trajectory, and this is normally the only part of the acquisition operator that is optimised for the reconstruction through use of a trajectory that produces incoherent artifacts. For example radial, spiral, or randomly undersampled Cartesian trajectories are often used for this purpose. **H** is normally chosen to be a Hadamard encoding to optimise the SNR [5] (because there is an equal contribution from all encoded images, and associated noise averaging, to each final decoded image).

Since **F** is limited to be a Fourier sampling matrix when using linear gradients for data acquisition, and k-space is sampled in continuous paths because it is time inefficient to sample discrete points, there is only so much freedom in how it can be optimised on its own. By combining it with **H** the incoherence of **E** can be better controlled, and additional degrees of freedom in shaping the point-spread-function (PSF) [86] are gained by choosing the trajectory and encoding scheme jointly.

The PSF is defined as the transform $\mathbf{E}^{\mathrm{H}}\mathbf{E}$, and is generally studied by applying the transform to a point source (a vector of all elements equal to zero except for one). When only **F** is included in **E**, the PSF is shift invariant in space and between the encoded components, meaning that the effect of the PSF on a point source does not depend on the position of the non-zero element in the vector. Including **H** into the acquisition operator can make the PSF component dependent, meaning that the effect of the PSF on a point source in one component can depend on which component the non-zero element is in (the PSF remains shift invariant in space). To distinguish the PSF that includes both **F** and **H** from the more common definition of PSF that only includes **F**, and to visualise the effect of the PSF on a delta function with the non-zero element in different components, the term multidimensional point spread function (m-PSF) will be used in the rest of this chapter. Visualisation of the m-PSF will show how the PSF applied to a point source in one component affects itself and all other components.

In the previous chapter, **F** and **H** were both included into **E**, but the m-PSF was still component invariant due to each encoded image being sampled with the same trajectory. This can be explained by considering that artifacts from undersampling behaved just like signal and cancelled out from the different encoded images when decoded and thus aliasing was constrained to the same component as the signal.

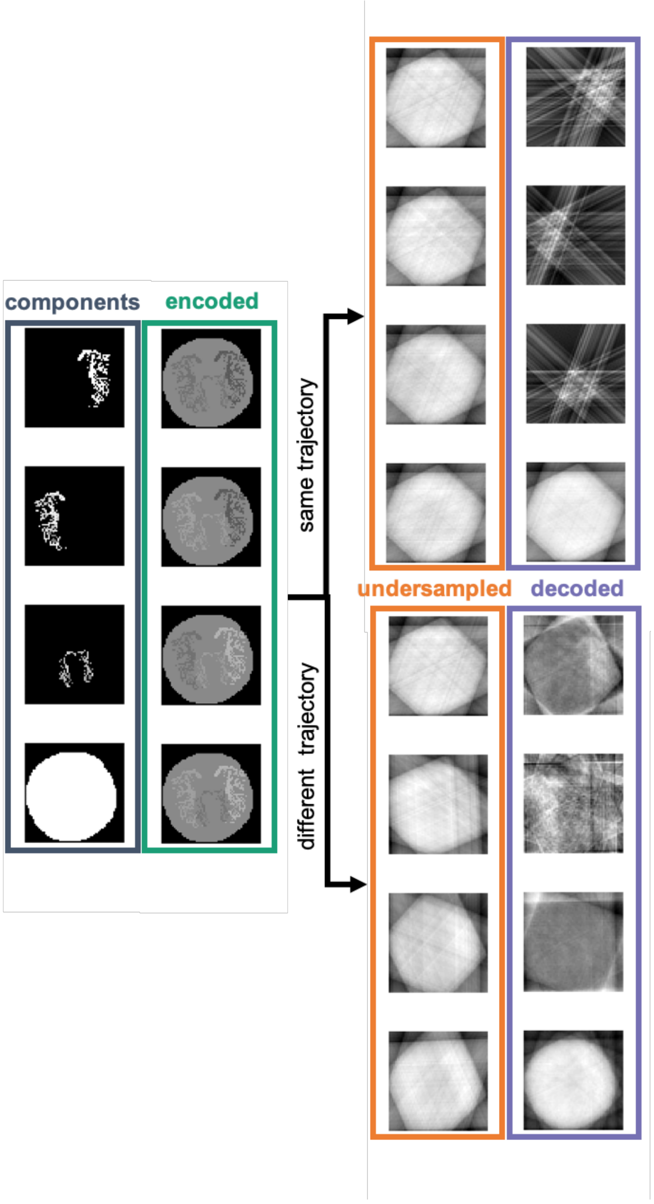

Figure 4.1: The effect of undersampling a VE-ASL signal with the same spokes every encoding (top) causes aliasing artifacts to be decoded in the same way as signal is decoded and thus artifacts do not 'spread' across components. If different trajectories are used to sample each encoding (bottom) the components mix, and especially the strong static tissue component causes severe artifacts in the decoded images.

The aliased signal starts to spread across encodings if different trajectories that create different aliasing patterns in each encoded image are used (figure 4.1).

Observations on the properties of the m-PSF are presented in the results section below.

4.2.2 Methods

Before going to 4D (section 4.5) the m-PSF was studied in a 2D numerical phantom. The effect of varying the k-space trajectory between encodings in VE-ASL was examined in the same phantom as used in chapter 3. It was resized to 64 x 64 pixels and its vessel components were randomly augmented with small changes in location, orientation, and intensity. The translation of each vessel component in each direction was drawn from Gaussian distributions (mean = 0, SD = 1), and then rounded to the nearest whole pixel. The angle of rotation of each vessel was also drawn from a Gaussian distribution (mean = 0, SD = 5°). Finally, the intensity of each vessel component was modulated by multiplying the original intensities (which were between 0 and 255) with a number drawn from a Gaussian distribution (mean = 1, SD = 0.5). Ten different versions of the phantom was used, and each version had ten instances of complex Gaussian noise added to them in k-space (SD = 10).

First, three different decoding and reconstruction approaches were used on the simple phantom to test the hypothesis that decoding and reconstructing jointly gives the best results as it allows aliasing to spread in the component dimension. The three approaches were:

1. Decode first then reconstruct – this can only be done when each encoding has the same k-space sampling and **H** and **F** commute.
2. Reconstruct encoded images, then decode – can be done with any sampling scheme but the images that are reconstructed are not sparse in image space. This is the most conventional way of reconstructing and decoding.
3. Jointly reconstruct and decode – include the encoding (**H**) in the forward model as per section 4.2.1. Here, the sampling operator (**F**) and encoding operator (**H**) can interact.

For reconstruction of the simple phantom, a simple CS reconstruction with sparsity enforced in decoded image space was used. The cost function was identical to the approach used in chapter 3 except it was without temporal regularisation.

Optimisation was done with FISTA with 300 iterations, and a step size of 0.01. The regularisation factor was set to 5 for all reconstructions.

Instead of trialling fully general sampling and encoding interactions, this study focused on three test cases that were all based on variations of 2D golden angle radial trajectories. For the simple phantom, either a total of 16 or 64 spokes spread across the encodings were used to reconstruct the images. The three approaches are described in figure 4.2, and were:

1. The same spokes method: Golden angle sampling with fixed sampling across all encodings with a 4 x 4 Hadamard encoding (4 or 16 spokes per encoding).
2. The varying spokes method: Golden angle with the trajectory rotated 45° for each encoding, still with a 4 x 4 Hadamard encoding (4 or 16 spokes per encoding).
3. The hybrid method: Golden angle with the trajectory rotated every other encoding with a paired 8 x 4 Hadamard encoding (2 or 8 spokes per encoding).

The paired encoding scheme, described in figure 4.2(C), is a hybrid between using the same spokes for each encoding and varying the spokes. With this approach, each pair of encoded images has the same trajectory and opposite encoding of the vessels, such that if they were subtracted from each other the static tissue background would be removed and a standard 4 x 4 Hadamard encoding would be retained for the vessel components only. With the 8x4 Hadamard encoding scheme, every pair of encoded images (e.g. rows 1 & 2) has a different trajectory from the other pairs (e.g. 3 & 4, 5 & 6, or 7 & 8), but requires double the number of encodings (8 instead of 4). Thus, half the number of sampling spokes can be acquired for each encoding to maintain scan time. This scheme decouples the static tissue component from the vessels, whilst still allowing for more unique k-space samples to be acquired. More details on why this works will be explained through examination of the m-PSF in the results and discussion sections.

To separate the effect of static tissue contamination from results that are due to the sampling efficiency only, simple phantom trials with both the static tissue

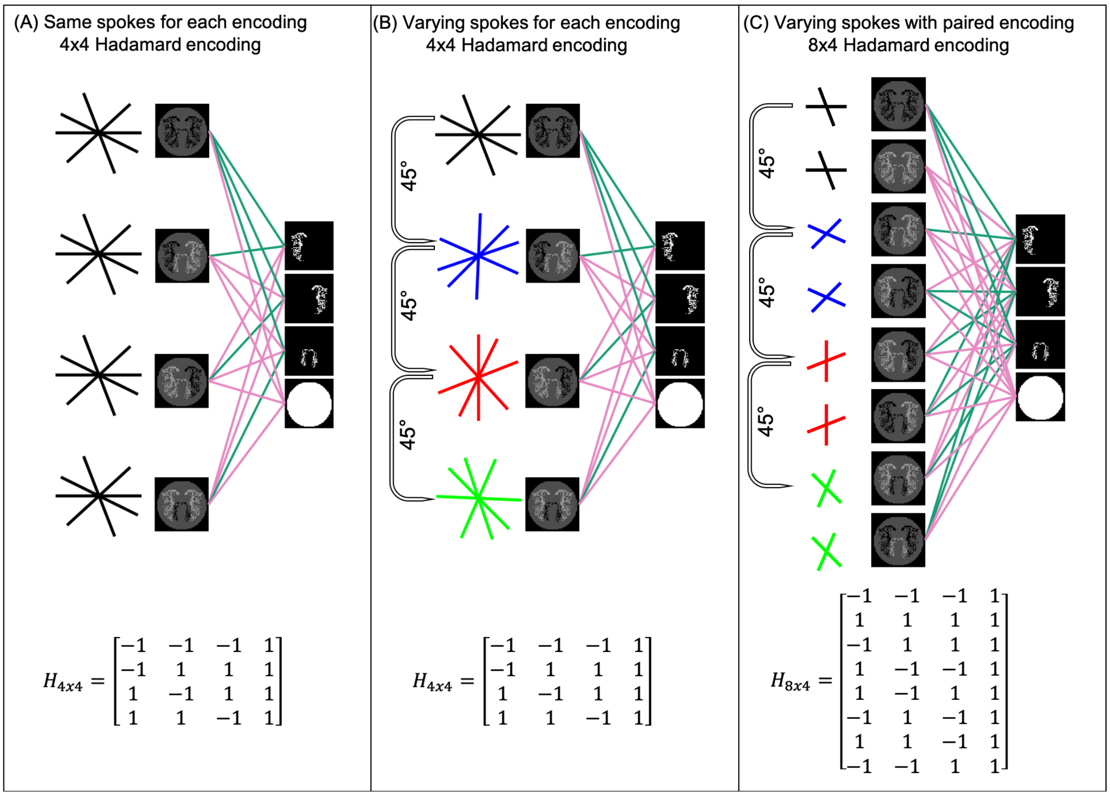

Figure 4.2: The proposed encoding and sampling methods. The four base components of the numerical phantom to the right are added (pink line) and subtracted (green line) according to a Hadamard encoding scheme (a simplified encoding matrix is shown at the bottom of each subfigure. In reality each ± 1 entry in $\mathbf{H}$ is an identity matrix of the size $N_{voxels} \times N_{voxels}$). The radial spokes on the left show the k-space trajectory for each encoding. In (A), the same spokes are acquired four times. In (B), four sets of different spokes are acquired. In (C), each matched sampling pair have the opposite tag-control condition.

component set to zero, and trials with high static tissue intensity (approximately 10 times the vessel signal), were run.

The joint decoding and reconstruction approach was further examined in the simple phantom by studying the effect of permuting the spokes acquired for each encoding on image quality and on the m-PSF. Three different spoke orderings were used, the same as used before, one where the spokes were as aligned as possible in each encoding, and one random permutation.

Image quality was again assessed using the voxelwise correlation coefficient between the ground truth and the reconstruction.

4.2.3 Results

As hypothesised, joint reconstruction and decoding was found to have the best image reconstruction quality for both 64 and 16 spokes, and both with or without strong background signal (figure 4.3). It was however, only slightly, but statistically significantly better ($p \ll 0.001$) than first decoding and then reconstructing when the same spokes were used for each encoding. Reconstructing the encoded images first and then decoding had the worst results. It is, however, worth keeping in mind that the reconstruction method was not optimised for non-sparse images such as non-decoded images.

Acquiring more unique spokes in k-space resulted in clear improvements compared to acquiring the same spokes every encoding in the joint reconstruction and decoding framework. The scheme with varying spokes every encoding performed the best in all reconstruction methods as long as there was no contamination by static tissue (figure 4.3(A,B)), but had the worst performance when realistic static tissue was included in the data (figure 4.3(C,D)). When static tissue was included, the hybrid method outperformed both the same spokes method (p < 0.01) and the fully varying spokes method ($p \ll 0.001$) with 16 spokes, for 64 spokes no significant difference between same spokes and the hybrid method was found. A clear improvement in reconstruction of fine details was observed with the hybrid method for 16 spokes, at 64 spokes all sampling approaches resulted in good reconstruction with the joint reconstruction method, see figure 4.4.

The m-PSF of each trialled sampling and encoding combination for the 16 spoke case is shown in figure 4.5. The m-PSF is symmetrical in all cases, meaning that component A aliases onto component B in exactly the same way as component B aliases into component A. In the case with same spokes every encoding, each component only aliases into itself as predicted. With different trajectories all components alias onto all other components, and with the hybrid approach, all but component 4 alias onto each other. This visualises why the hybrid method worked best when static tissue signal was present. The aliased energy of the static tissue swamps the other components with the varying spokes approach, but is

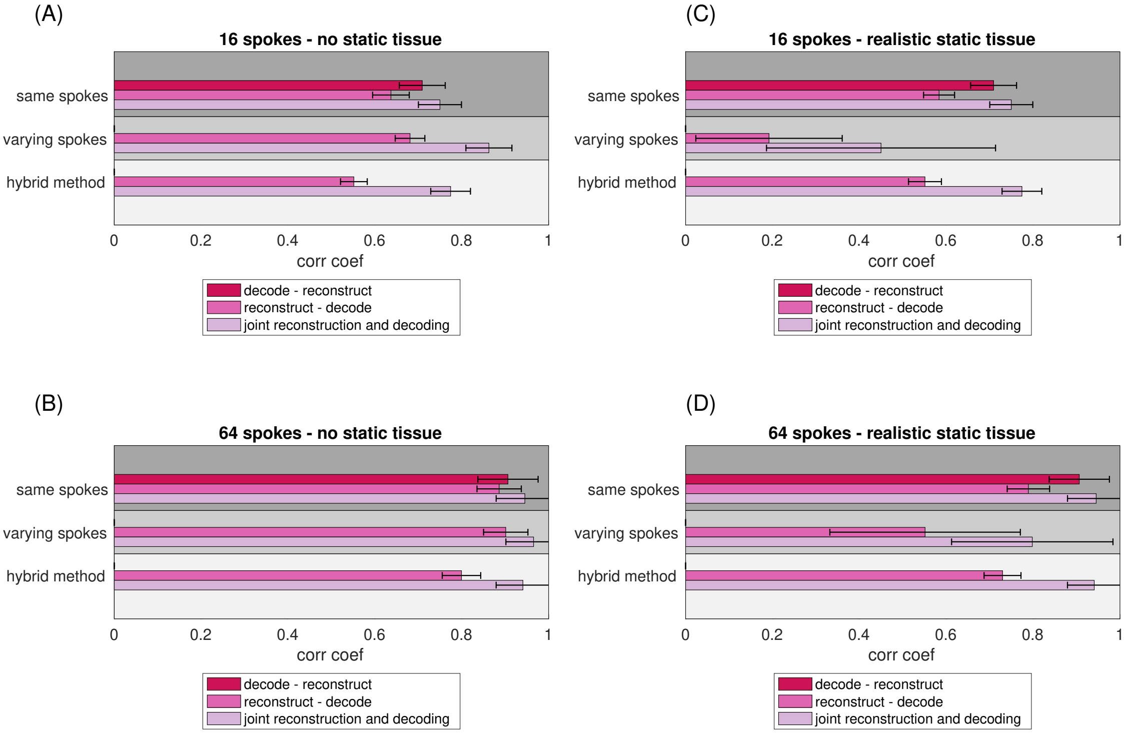

Figure 4.3: Comparison of the trialled reconstruction and sampling methods without (A,B) and with (C,D) strong static tissue background.

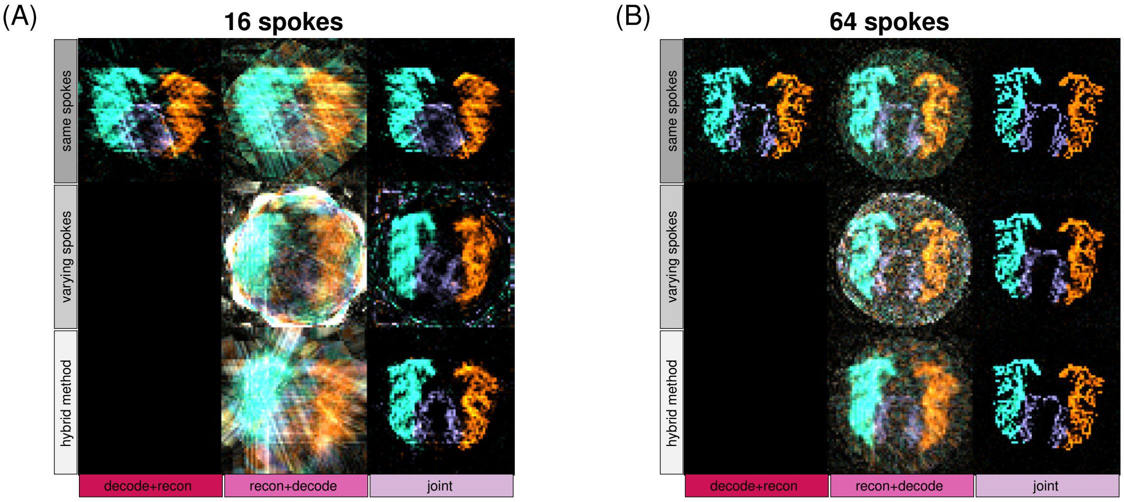

Figure 4.4: Examples of the different approaches to sampling and reconstruction. (A) shows the 16 spoke results with severe artifacts if not reconstructed and decoded jointly. The varying spokes method shows residual aliasing from the static tissue, and the same spokes approach is very streaky due to only 4 unique spokes being sampled. The hybrid method produced the cleanest results with the best vessel definition. In (B) the images are close to perfectly recovered using the joint and decode+recon methods with a total of 64 spokes sampled across all encodings.

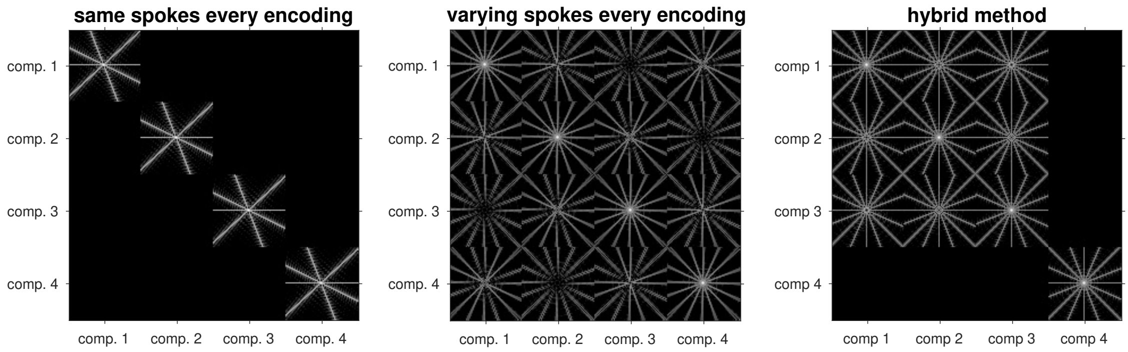

Figure 4.5: Example m-PSFs for the 16 spoke case. Acquiring the same spokes contains aliasing within components (off diagonal elements are zero). Varying spokes spreads the aliasing over all components and have more spread within the components due to more unique spokes being sampled. The hybrid method nulls cross spreading from component four thus mitigating the effect of static tissue contamination in VE-ASL but still allows for spreading of aliased energy between components one to three.

nulled in the paired approach. This property can be can be represented in the m-PSF framework by noting that component mixing is determined by energy of off-diagonal blocks in the m-PSF.

By permuting the sampled spokes between the different encodings, it was observed that the diagonal blocks are unaffected by which spokes are acquired in which encoding, but that the off-diagonals change when different sampling schemes are used. Permuting the spokes acquired across the paired encodings slightly changed the m-PSF off-diagonal elements, which resulted in minor differences in image quality (figure 4.6). The sampling with more closely aligned spokes in each encoding achieved better image quality than the randomly permuted sampling ($p < 0.05$), none of the other differences reached significance.

Although the regularisation factor, λ, was not optimised for each reconstruction method and ground truth permutation, it was confirmed that the effect of varying λ was smaller than the differences in reconstruction quality due to the reconstruction approach or sampling trajectory used. This was done by manually tuning the regularisation factor for one ground truth permutation using all reconstruction and sampling-encoding methods.

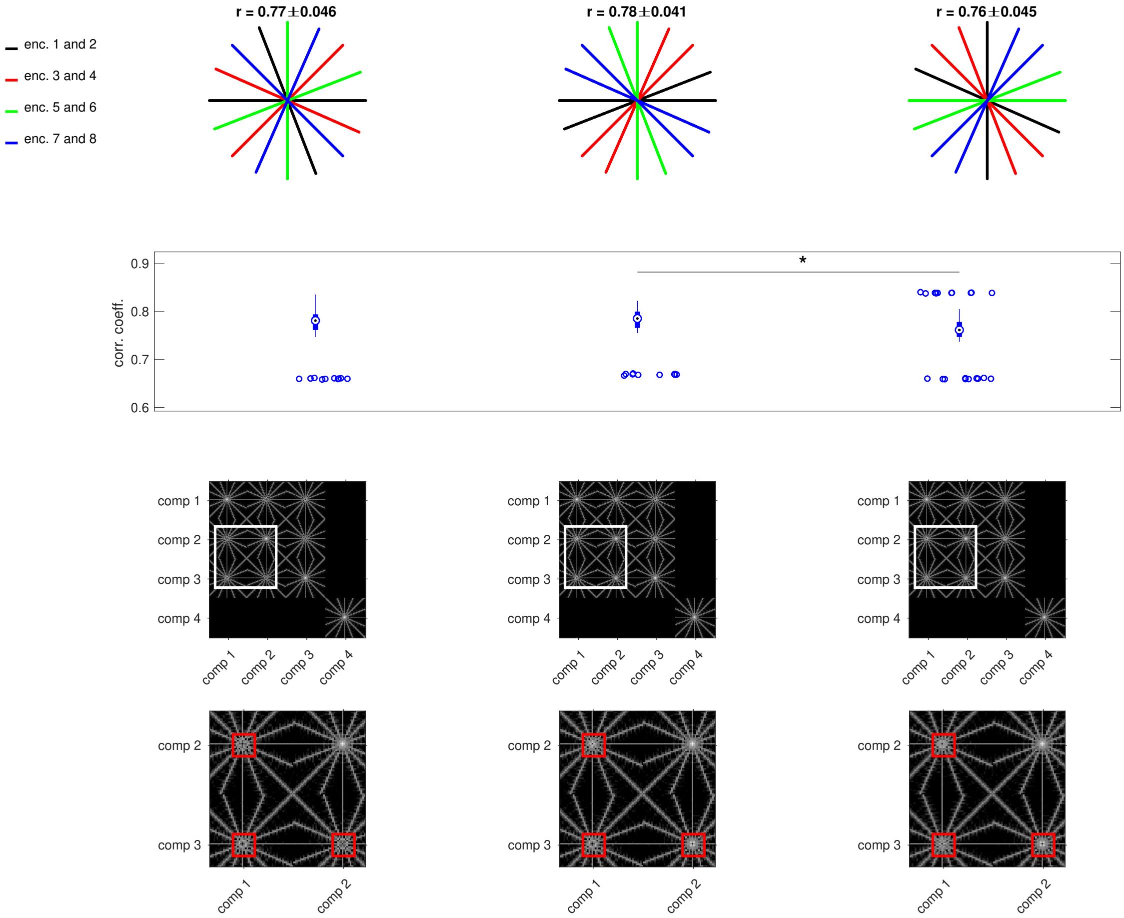

Figure 4.6: Effect on the m-PSF of acquiring the same spokes in different encodings. the left-most column shows the ordering that was used for the earlier experiments, the other two columns show sampling that have spokes acquired in a different order and small differences in their m-PSFs. The bottom row shows zoomed in versions of the m-PSF (zoomed in on the white box), differences are highlighted with red boxes.

4.3 Improved Sparsity

As shown in the previous chapter, improved relative sparsity can push sparsity based reconstructions to higher acceleration factors. Although relative sparsity is improved immediately by going to a full 3D acquisition, it is also somewhat reduced by the introduction of the labelling plane into the imaging region when using a non-selective excitation that causes a large non-sparse artifact. The artifact is due to modulation of static tissue at the labelling plane, that varies for the different VE preparations. This prompts consideration of other methods of improving sparsity of the image.

The vessel components are already very sparse in image space so let us focus our attention onto the static tissue component. Although it is not of interest clinically, in a joint reconstruction the quality of the static tissue can affect the quality of the vessel components. Previously, the static tissue had been left in the reconstruction un-modified although it is non-sparse in image space. In this chapter the static tissue component is considered more carefully to ensure that it does not interfere with the reconstruction of the vessel components. Both the coupling between the vessel images and the static tissue image and simple sparsifying transforms to move the static tissue into a sparse domain are explored.

4.3.1 Methods

The m-PSF analysis showed that using the same spokes each encoding or using the hybrid method decouples the static tissue from the vessel components. When the static tissue and the vessel components are decoupled, improving sparsity of the static tissue does not affect the vessel images. However, when varying spokes are used, the static tissue causes problems if not dealt with properly. If the static tissue could be fully reconstructed, its aliasing artifacts in the other components can be suppressed and it might be possible to use the most efficient method of sampling k-space and never repeat the same k-space locations. To get good reconstruction of the static tissue in a compressed sensing based reconstruction, it needs to be reconstructed in a sparse domain.

Two methods for sparsifying the static tissue were trialled and the effect measured on image quality (correlation coefficient between ground truth and reconstruction) of a realistic 2D simulation (the oversampled acquisition of subject 1 from chapter 3). Again, the theory was trialled in a lower dimensional reconstruction problem before moving to 4D (section 4.5) to be able to run many experiments with reasonable reconstruction times on non-specialised hardware.

The first method for sparsifying the static tissue was the very common method for sparsifying dense MR images, the wavelet transform. The second method was a data driven method that relied on an average static tissue image (across frames and encoded images) that could be reconstructed at a much lower undersampling factor when different spokes are acquired for each frame and encoding. The average image was then used to create a sparse basis for the static tissue and a linear sparsifying transform to be included in the compressed sensing framework. The sparsifying transform was generated by taking the data from all frames and encodings and reconstructing an average image using a simple sampling density weighted NUFFT-based re-gridding. The non-zero voxels of the image were then divided into a set number of patches. The assumption is then that the static tissue in each time frame, $\mathbf{s}_t$, is simply:

$$\mathbf{s}_t = \sum_{i}^{N} c_{i,t}\mathbf{p}_i \tag{4.2}$$

where $\mathbf{p}$ is the vector defining each patch of the image and N is the number of patches, which has to be small in comparison to the number of non-zero voxels to define a sparse basis. $c_{i,t}$ is a scaling factor that can be individually set for each frame. The square sparsifying operator is then built up by the combining the patch column vectors, $\mathbf{p}_i$, with their null-space to form an orthonormal basis transform. This is not a conventional approach to patch-based sparsification, which more commonly uses learned over-complete dictionaries [87], rather than a basis determined by a temporal mean, but this approach was taken as an initial exploration into static tissue sparsification. This process is summarised in figure 4.7. Because the patches were large and not optimised to contain only one tissue type per patch, this method will struggle to capture tissue specific intensity changes due to T1 decay, and thus mainly capture global signal attenuation between frames. This is a potential weakness of this method, but was accepted as a simple proof of principle approach to trial the effect of sparsification of the signal.

The cost function that was used for reconstructing VE-ASL angiograms was:

$$c(\mathbf{m}) = \frac{1}{2}\left|\mathbf{E}\mathbf{m} - \mathbf{s}\right|_2^2 + \lambda_1\left|\mathbf{P}\mathbf{m}\right|_1 + \lambda_2\left|\nabla_t\mathbf{m}\right|_2^2 \tag{4.3}$$

The only difference between this cost function and equation 3.6 that was used in the previous chapter, is the introduction of $\mathbf{P}$, a sparsifying transform in the

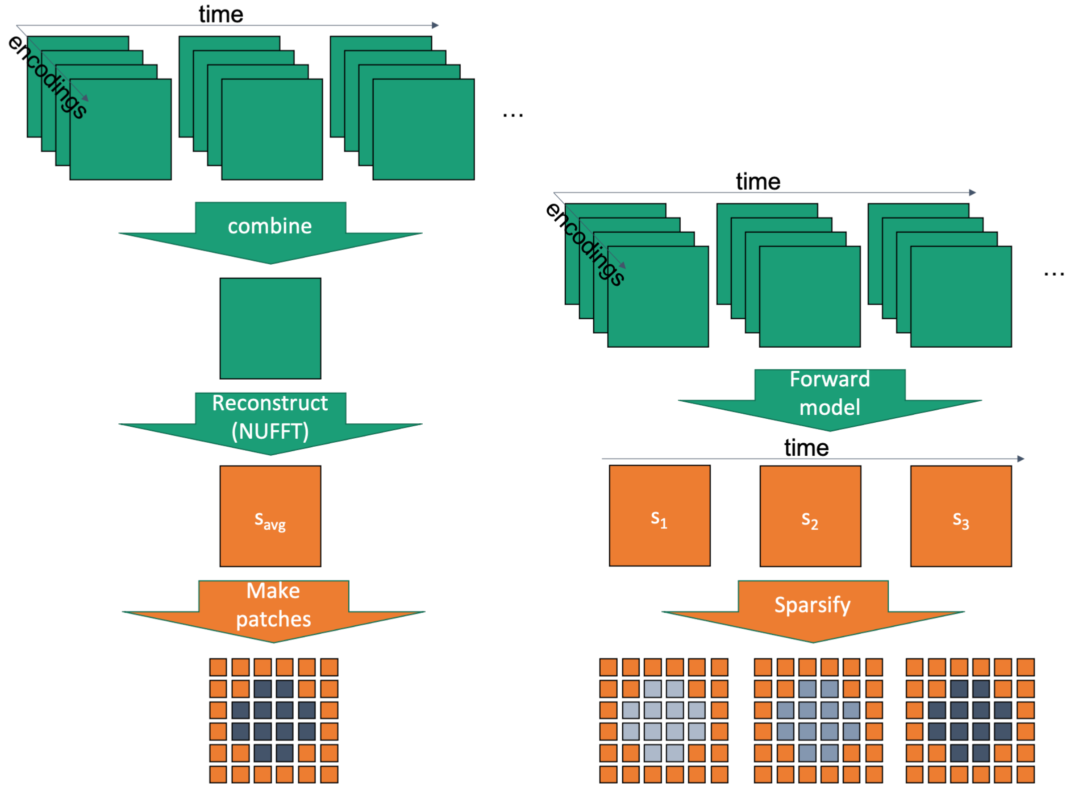

Figure 4.7: Schematic for patch based sparsification method. Green squares represent k-space acquisitions, orange squares reconstructed static tissue images.

L1-term. **P** was set either to a level 4 Daubechies db2 wavelet transform applied only to the static tissue component (an identity transform was applied to vessel components), or the data driven patch based basis transform.

The data was reconstructed with FISTA, 300 iterations, step size 0.01, $\lambda_1 = 10^{-6}$. λ_2 was set to 1 and to zero, to separate the effect of sparsity alone and interaction with the temporal regularisation. These values were chosen to be of similar magnitude as those optimised for in chapter 3.

The acquisition was simulated using the forward model, and white complex noise (SD = 0.000003) was added in k-space to achieve a $SNR_k = 107.63$ for the realistic data, which is similar to the noise levels for a VE-ASL acquisition, as measured in the previous chapter (section 3.2.2). k-space was undersampled with an acceleration factor of $R = 34$ (9 spokes per frame per encoding for same spokes and varying spokes) or $R = 61$ (5 spokes per frame per encoding for the hybrid method) using a golden angle trajectory across time. For the varying spokes reconstruction the trajectory was rotated nine times the golden angle (111.24...°) between every encoding, and for the hybrid method five times the golden angle. With this trajectory golden angle encoding across both time and encodings is achieved, which is slightly different than the method used in section 4.2 where the sampling was done using golden angle within each frame and uniform sampling across encodings. The sampling patterns used in this section are summarised in figure 4.8).

The effect of varying the number of patches in the data driven sparsifying operator was examined by running the reconstruction with 32, 128, and 1024 patches. Higher number of patches allows for more non-zero components in reconstruction, and thus lower image sparsity, which is bad for a compressed sensing reconstruction. However, more patches (and thus smaller patches) makes the sparsifying operator **P** sparser which is beneficial for reconstruction speed when using MATLAB's inbuilt sparse matrix data type.

In order to assess whether improved background suppression methods could help in future work, image quality (of the vessel components) was assessed at two different levels of background intensity; strong background suppression (the maximum intensity in the static tissue set to the same as the maximum intensity in the vessel components), and realistic static tissue (the acquired data used as it was in chapter 3, the maximum value in the static tissue was approximately 7.6 times higher than the maximum value in the vessels).

4.3.2 Results

Image reconstruction quality for 32 and 128 patches was similar, but at 1024 patches image quality was reduced. Time for the reconstruction of 1024 and 128 patches were similar but the time was increased for 32 patches. For this reason 128 patches were used in all further experiments (figure 4.9).

In the presence of weak static tissue all methods performed well, but the varying spokes method with static tissue sparsification outperformed the same spokes and hybrid method slightly. At the acquired background tissue intensity, decoupling the static tissue from the vessels either using same spokes or the hybrid method,

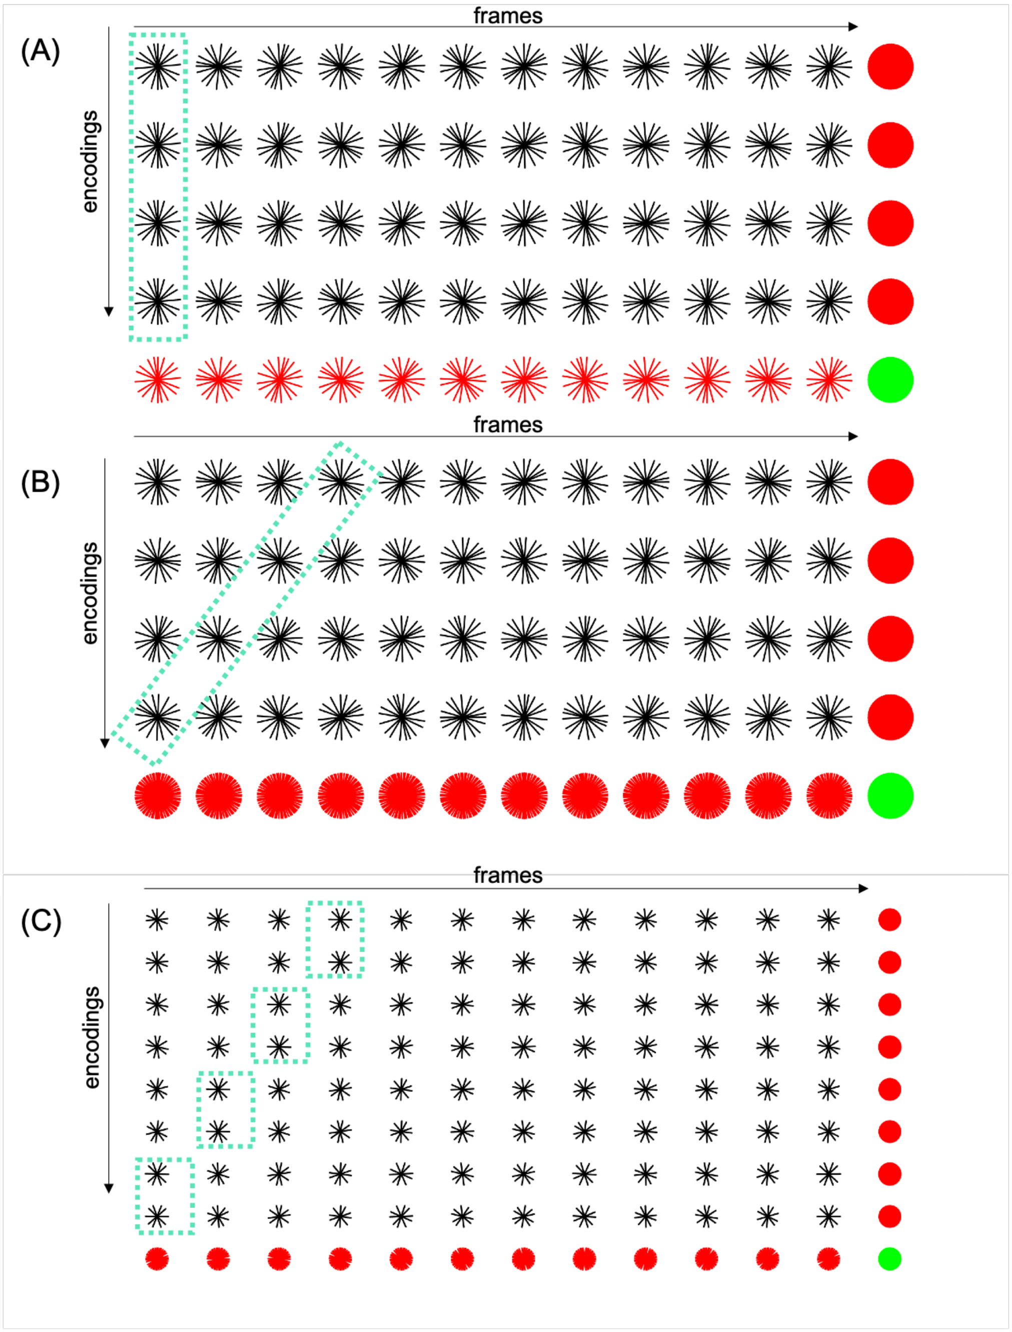

Figure 4.8: Sampling patterns used for simulated acquisitions. In black are the sampling patterns for each frame and each encoding, in red is the combination of spokes across frames (on the right) or across encodings (bottom). The bottom right green spokes shows all spokes acquired in the experiment. Examples of encodings/frames with the same sampling pattern are annotated with the green dashed box. (A) shows the "same spokes" method where each frame has the same sampling pattern for each encoding. (B) shows the "varying spokes method" where the sampling pattern is rotated across frames and time. (C) shows the "hybrid method" where the same pattern is repeated in paired encodings.

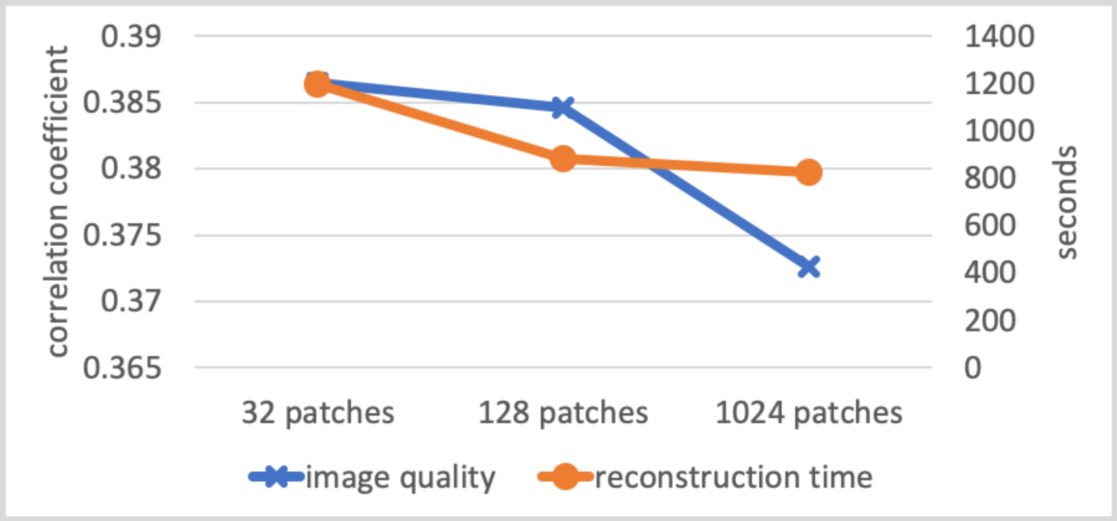

Figure 4.9: Reconstruction time and image quality for different number of patches used to sparsify image.

improved the reconstruction quality significantly. Interestingly, the same spokes method slightly outperformed the hybrid method when the L2 temporal smoothness constraint was included. This could be explained by the fact that neighbouring temporal frames in the hybrid method share some of the same spokes and thus have more smoothly varying artifacts across time. The results with temporal smoothness constraints are shown in figure 4.10(A).

If the temporal regularisation term was set to zero, the static tissue sparsification through the patch based basis transform improved image quality when the static tissue intensity was low. The wavelet transform did, however, provide no improvement to image quality for weak static tissue, and made the reconstruction quality worse for realistic static tissue. With the realistic static tissue intensity levels, the patch based method still improved reconstruction compared to doing nothing with the static tissue acquired with varying spokes. Again, decoupling the static tissue from the vessels by using the same spokes or hybrid methods of sampling and encoding performed the same in both cases, as they were not affected by the static tissue strength. Without L2-smoothness the hybrid method performed slightly better than using the same spokes, which is consistent with the results from section 4.2. Figure 4.10(B) shows the results of reconstructing without temporal smoothness constraints.

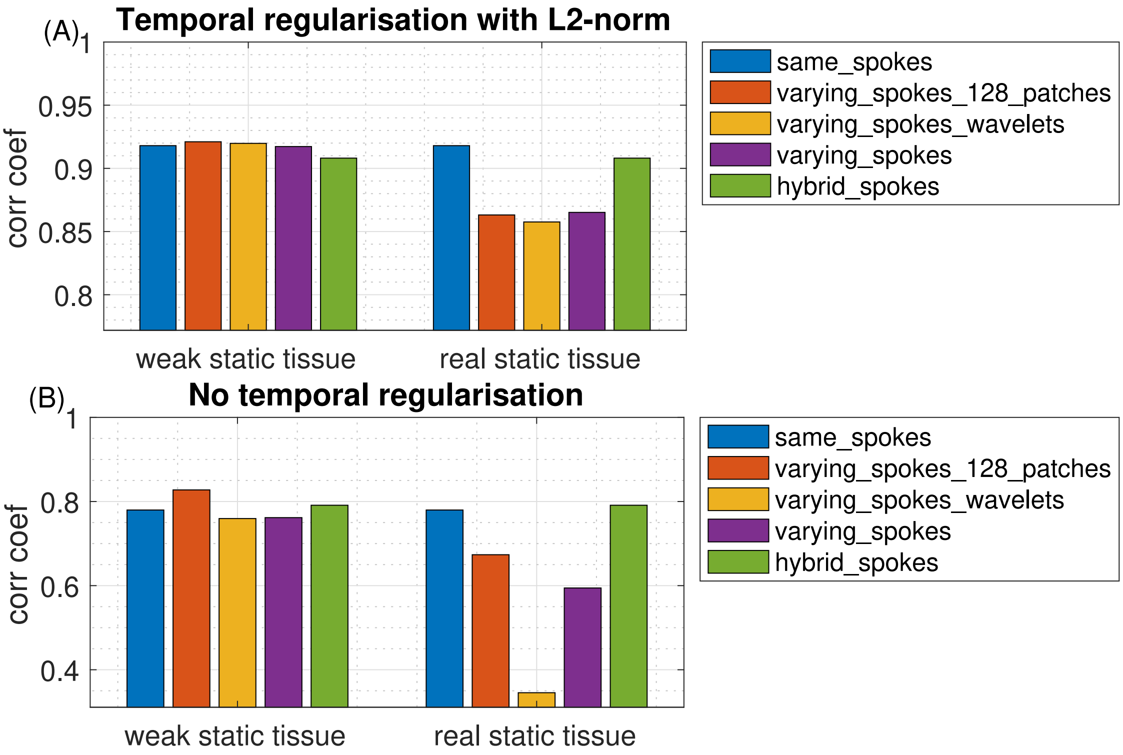

Figure 4.10: Image reconstruction quality with different sampling and sparsification approaches. With temporal smoothness constraints (A), and without (B).

In summary, both static tissue sparsification approaches seem to have limited benefits, especially if better background suppression cannot be performed. The conclusion from this section is therefore that decoupling static tissue is a more promising avenue to pursue than sparsifying the static signal.

4.4 More Precise Regularisation Terms

The third and final avenue of extensions that was explored as part of this project was the use of more tailored regularisation terms. In the previous chapter and sections, the regularisation terms have relied on very general properties of the data; spatial sparsity and temporal smoothness. Instead of these universal properties, new regularisation terms that can constrain VE-ASL angiograms specifically were considered.

VE-ASL has a strong spatial prior with blood from the RICA likely to end up in the right cerebral hemisphere, blood from the LICA in the left hemisphere, and blood from the BA in the posterior circulation. Blood from the three feeding arteries are also generally not mixed in the majority of the vessels in the brain. This information can be built into the reconstruction as a term that penalises signal that ends up in the "wrong" place, as it is likely to be aliasing rather than actual signal.

But how is the "wrong" place defined? One could imagine splitting up space into three boxes based on a localiser and forcing the signal from each vessel to stay in some non-overlapping anatomical regions. However, with pathology and non-standard configurations of the circle-of-Willis these types of hard constraints on blood ending up on the right/left/posterior might not hold as well, so a softer constraint is required. In this project, a low-resolution image was used to create regional priors for each vessel component. The regions were also allowed to overlap.

Similarly, the images can be more constrained temporally, as the temporal evolution of the VE-ASL signal follows a smooth dynamic model as described by Okell et al. [66]. The model includes both physiological and MR parameters. The physiological parameters include blood volume, bolus dispersion time to peak, dispersion sharpness, and transit time from labelling plane to the imaged voxel, whereas MR parameters include signal loss due to RF attenuation by early imaging pulses and relaxation due to T1 decay.

By assuming that the signal will follow this model and simulating time courses within physiologically plausible parameters and then performing a principal component analysis the possible time courses can be reduced to a smaller subspace. The temporal signal in each voxel can then be explained by only a few temporal components, rather than a value in each time frame, and thus sparsifying the signal in time. For a sparse signal, an L1-constraint can be used instead of L2-smoothness as in chapter 3.

4.4.1 Methods

The two new regularisation methods were applied to the same dynamic 2D data as was used in the static tissue sparsification experiments above with both same spokes for each encoding and the hybrid method.

Spatial Regularisation

The spatial regularisation was done by taking the data, and applying Gaussian radial weighting (SD = 50) to each k-space line (length: 384), then combining data across the 12 temporal frames, reconstructing the three vessel components, and thresholding the result at each vessel's 80th percentile to produce low resolution masks for where the vessels were expected to be. The areas where the vessels were expected to be were given no weighting (i.e. no additional cost or penalty), and areas were the vessels were not expected to be were given an extra cost set by λ_3 in the combined cost function using a diagonal matrix, $\mathbf{W}$, which had entries 1 corresponding to voxels outside the masks and 0 inside. The cost function was then:

$$c(\mathbf{m}) = \frac{1}{2} |\mathbf{Em} - \mathbf{s}|_2^2 + \lambda_1 |\mathbf{m}|_1 + \frac{1}{2}\lambda_2 |\nabla_\mathbf{t} \mathbf{m}|_2^2 + \frac{1}{2}\lambda_3 |\mathbf{Wm}|_2^2 \tag{4.4}$$

λ_1 was set to 10^{-6}, λ_2 and λ_3 were set to 1. FISTA with a step size of 0.01 was run for 300 iterations.

Temporal Regularisation

Plausible time courses, based on the forward model derived in [66], with fixed scan parameters and varying physiological parameters were simulated in order to find a sparse basis to use in reconstruction. The temporal evolution was simulated for each spoke (every 12 ms), but the signal time course was then averaged for all spokes in the frame. The simulated parameters are summarised in table 4.1. A total of 4800 different time courses were simulated. A visualisation of the time courses are shown in figure 4.11(A,B). The framewise time courses were then compressed using singular value decomposition. The first four components (out of 12) are shown in figure 4.11(C). They explained $> 99\%$ of all variance (figure 4.11(D)).

The temporal model was used instead of temporal smoothness and without spatial priors. The cost function was thus:

$$c(\mathbf{m}) = \frac{1}{2} |\mathbf{Em} - \mathbf{s}|_2^2 + \lambda_1 |\mathbf{Sm}|_1 \tag{4.5}$$

Scan parameters	Labelling duration (ms)	1000
	Repetition time (ms)	12
	Echo time (ms)	6
	Flip angle (degrees)	7
	Post labelling delay	0
	Readout duration (ms)	1296
	Spokes per frame	9
Physiological parameters	T1 of blood (ms)	1664
	Time from labelling plane to imaging region (ms)	0
	Time from labelling plane to voxel (ms)	0 to 1500 (steps of 100)
	Dispersion kernel sharpness (a.u., low sharpness implies high dispersion)	0.001 to 0.015 (steps of 0.001)
	Dispersion kernel time to peak (ms)	1 to 20 (steps of 1)

Table 4.1: Parameters used in simulation of plausible ASL angiography signal time courses.

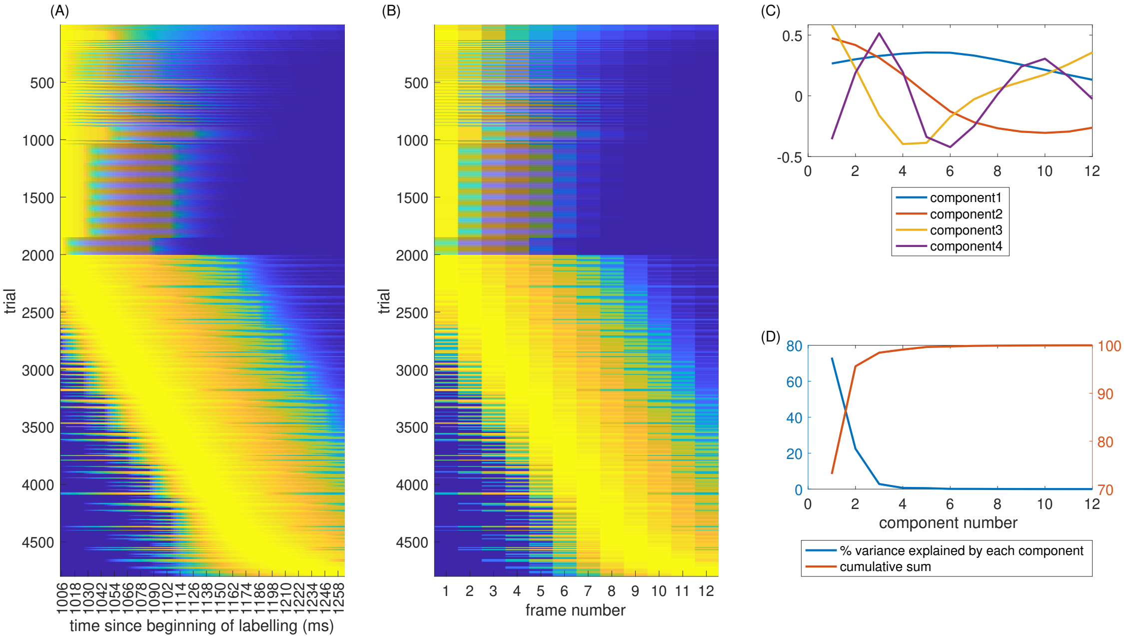

Figure 4.11: Temporal basis function generation. (A) shows all the simulated time courses with a resolution of one time point per spoke (TR = 12 ms). In (B) the signal from (A) is averaged into 12 frames. (C) shows the first four components from the principal component analysis, and (D) shows the amount of variance explained by each component.

With **S** being the basis transform to the simulated temporal bases. λ_1 was set to 10^{-6}. The reconstruction parameters were otherwise the same as above.

4.4.2 Results

The spatial regularisation method did not improve results for either the hybrid or same spokes method. Indeed, it made the reconstructions slightly worse (figure 4.12). The lack of improvement can be explained by the aliasing from an undersampled radial trajectory appearing as radial spokes, with high energy deposition closer to the point itself than further away, so penalising signal far away from their sources

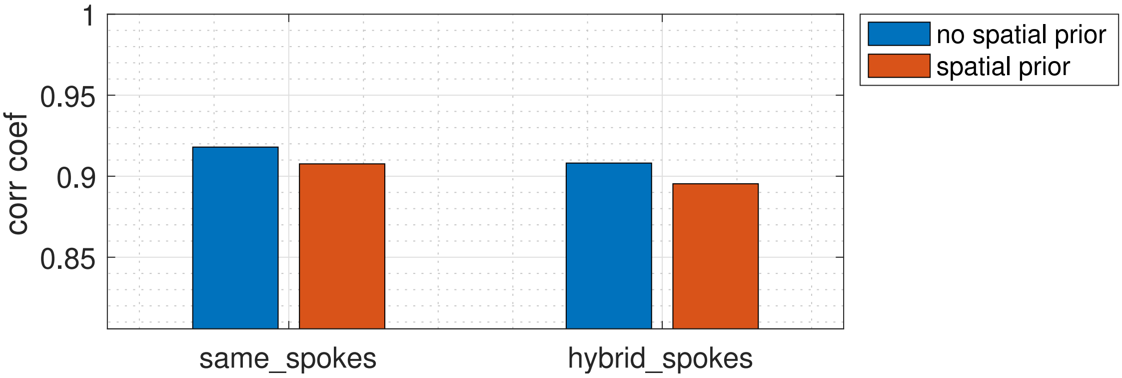

Figure 4.12: Image reconstruction quality with and without the use of a spatial prior using either the same spokes or hybrid method of sampling.

does not affect this as much as for example when aliased copies appear in a Cartesian framework. Having the coil sensitivities included in the reconstruction also helps with spatial localisation of the signal. There was simply not much to gain by such non-specific spatial regularisation. Why it made the results worse is a more puzzling question. Perhaps the bias produced by the spatial regularisation, even when broad and overlapping regions were used, was enough to lose some of the more subtle details around the edges of the penalised regions. However, this result suggests that for more complicated vasculature, with more overlap of vessels, this technique would be detrimental. For this reason, this type of cost function modification was abandoned.

Similarly, the temporal regularisation did not improve reconstruction quality beyond what the temporal smoothness constraint did, although it improved reconstruction quality compared to using no temporal regularisation (figure 4.13). The reason for these negative results is probably that the signal simply is not sufficiently compressible along the temporal dimension. Only 12 frames were used, and approximately four components were required to adequately represent the signal. A relative sparsity of $\frac{4}{12} = \frac{1}{3}$ is likely not sufficient for compressed sensing to work effectively.

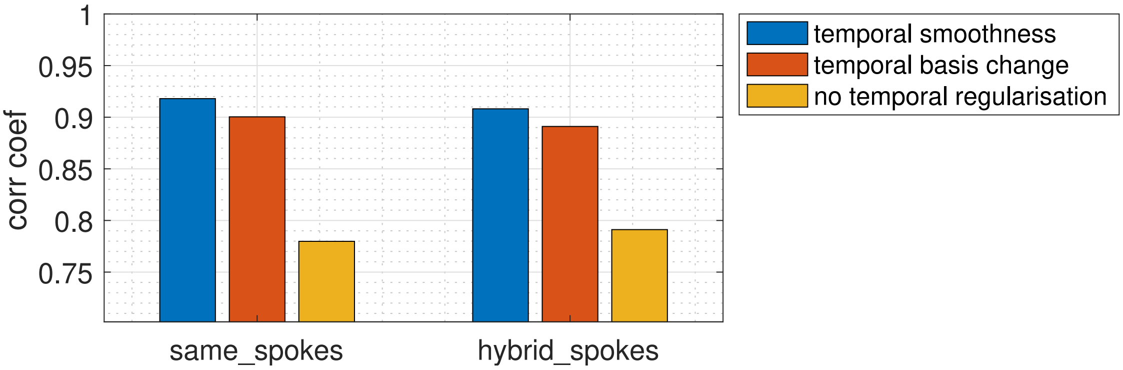

Figure 4.13: Image reconstruction quality with and without the use of temporal regularisation (either L2-smoothness or L1-sparse basis) using either the same spokes or hybrid method of sampling.

4.5 Combining the Extensions to Acquire 4D VE-ASL Angiography

Now, to combine the results of all of these extension experiments together and apply to a 4D acquisition.

The m-PSF analysis showed the benefits of sampling different spokes in different encodings as long as the static tissue can be dealt with. Of the two ways of "dealing with" the static tissue, sparsification and decoupling, decoupling worked better and more reliably when the static tissue was not artificially suppressed. The sparsification methods also increased reconstruction time, which is even more of a problem for 4D imaging. For these reasons, decoupling the static tissue through either the same spokes or hybrid method was chosen as the preferred approach for 4D imaging.

The move into a dynamic 2D phantom from the simple, single-frame, phantom, also showed that when temporal regularisation is used, it is important to use different spokes between frames in order to get any benefit from the hybrid spokes method over just using the same spokes for each encoding. So, for the move to 4D, the sampling was modified again to ensure good temporal incoherence, as well as complementary sampling across encodings.

The use of new spatial and temporal regularisation terms did not show any benefits in dynamic 2D simulations, and were thus not included in the 4D experiments. In particular, because the model based temporal sparsification method did not provide any benefits over the simpler temporal smoothness constraints in the simulations that had 12 frames, it was unlikely to provide benefits in-vivo where, to minimise the problem size, even fewer frames were reconstructed and the temporal compressibility was further constrained.

The main question for these experiments was how far 4D VE-ASL can be accelerated without meaningful loss of image quality. Can 4D VE-ASL angiography be performed in clinically feasible scan times (preferably < 10 min)?

4.5.1 Methods

Sampling Trajectory

The first change from dynamic 2D imaging to 4D was the change from the golden ratio method of sampling radial spokes, to the multidimensional golden means method [34]. This is a straightforward change as the multidimensional golden means method is a direct extension of 2D golden ratio sampling as discussed in chapter 2.

The other change with regards to sampling was to make the frames have as non-overlapping trajectories as possible when the hybrid sampling method was applied. In simulations this can simply be done by ordering the spokes by the golden means method across encodings and, along the time dimension, carry on from the last spoke used in the last encoding of the previous time frame (figure 4.14). With this ordering the hybrid approach can be achieved and no overlap between trajectories in subsequent frames occurs. However, in vivo, where the data is acquired in a segmented acquisition with multiple shots, the ability to retrospectively choose the undersampling factor by discarding later acquired shots and simultaneously have golden means distributed spokes in each frame, can only be maintained if subsequent repetitions carry on the ordering such that there is some overlap between frames. The first shot is acquired with the same ordering as in simulations, and the next subsequent shot carries on the ordering from the first shot, which results in the same spokes being acquired in the first shot, second frame, and the second shot, first frame, and so on. Figure 4.15 shows the ordering used for in vivo acquisitions.

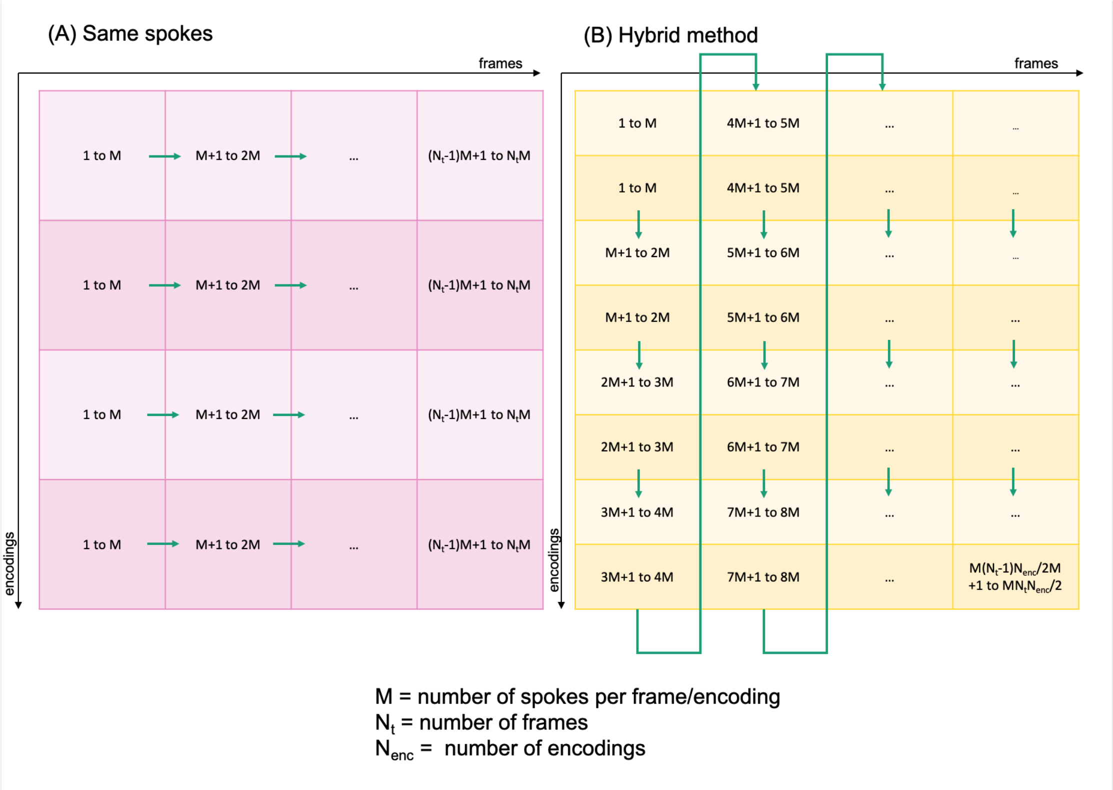

Figure 4.14: Spoke ordering in simulations. For the same spokes method, (A), the golden angle index increases just from frame to frame (green arrows). In the hybrid method, (B), the index is increased for every other encoding and the following frame carries on where the last encoding in the first frame finished.

Computational Limitation Considerations

As shown above, there are benefits with reconstructing and decoding jointly, as well as considering regularisation across both space and time. When moving to 4D, these benefits come at the cost of memory usage when running the reconstruction algorithm. The "image" that is reconstructed is a 5D dataset (three spatial dimensions, time, and vessel components), where high spatial and temporal resolution is required. This can be decoupled into smaller problems by decoding first and then reconstructing (only possible with same spokes sampling, and not as efficient as joint decoding as shown in section 4.2), or reconstructing each frame separately (no temporal regularisation possible, which showed large benefits in chapter 3). Because of the benefits of considering the full 5D dataset, rather than reducing the dimensionality of the problem, other methods were used to reduce memory burden.

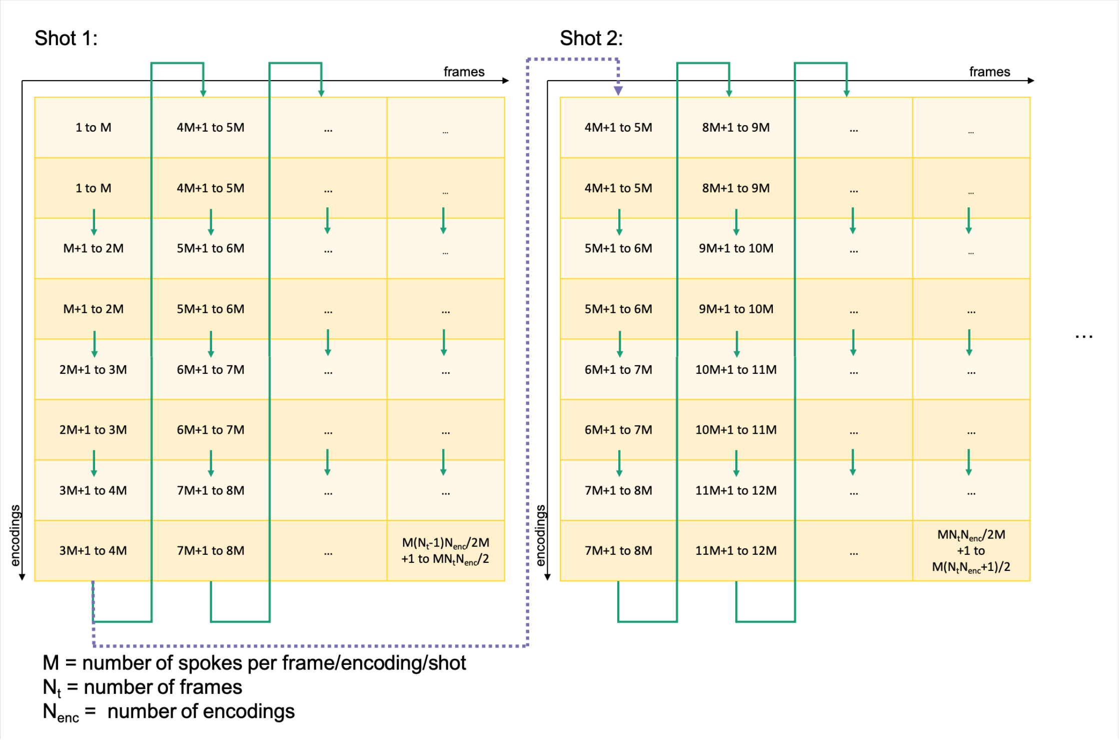

Figure 4.15: Ordering of spokes for the hybrid method in the multishot sequence used in vivo. Each shot follows the same ordering as was used in simulation (figure 4.14(B)), and following shots use spoke indicies increased by the number of spokes used in each frame. When spokes are combined across shots they have ideal golden ratio sampling within each frame, but subsequent frames have a number of overlapping spokes that depends on the number of shots combined.

The matrix needed to store the full 5D image, assuming the same resolution as was used for dynamic 2D in chapter 3, would be $192 \times 192 \times 192 \times 12 \times 4$. Assuming, each matrix entry needs to store a complex number to double precision the total memory size of the matrix would be $192 \times 192 \times 192 \times 12 \times 4 \times 16$ Bytes, which equals approximately 5 GB. To reduce the memory burden, the matrix size (and resolution of the image) was reduced both for simulation ($128 \times 128 \times 86 \times 6 \times 4$) and in vivo acquisitions ($192 \times 192 \times 192 \times 6 \times 4$).

However, it is not only the reconstruction that has to be stored and accessed from memory multiple times. The acquisition operator $\mathbf{E}$, which was implemented as a MATLAB object also had to carry a lot of data. For example, the coil sensitivity maps, that although coil compression was applied to reduce the number of coils from 32 to 8, required storage of complex data of the same dimensions as the image's spatial dimensions (up to 192 × 192 × 192 × 8 × 16 Bytes ≈ 0.8 GB). When the acquisition operator or its transpose is applied, it also has to generate large matrices internally to store the intermediate steps of the acquisition, for example the vessel encoded images, which for the hybrid method is double the size than for same spokes as the number of encodings are doubled. The raw k-space data also had to be stored in memory. However, the size of the data varied based on what undersampling factor was used.

One way to minimise memory usage is to use a lookup-table in the acquisition operator, instead of storing the gridding interpolation kernels for each k-space point. This reduces accuracy and speed of the operation slightly [67]. As in chapter 3, to speed up the iterative reconstruction, $\mathbf{E}^{\mathrm{H}}\mathbf{E}$ was implemented using Toeplitz embedding and the FFT instead of the NUFFT. However, the Toeplitz embedding requires the storage of a matrix with double the spatial dimensions of the image for each time point and encoding. It also generates large temporary matrices when applied (which happens every iteration). Both the lookup-table method and Toeplitz embedding were used in all 4D reconstructions.

The simulations were run on a personal laptop computer (MacBook Pro, 2018, 16 GB RAM, quad-core 2.7 GHz Intel Core i7), whereas the in vivo reconstructions were run on a computing cluster with jobs allocated two threads and 250 GB of RAM on a 2.6 GHz 2x20-core Intel Xeon E5-2660. The cluster was used so that multiple reconstructions could be run in parallel. Approximate memory usage was recorded for the simulations by periodically checking the inbuilt "Activity Monitor" application. In vivo reconstruction memory usage was also checked periodically as the reconstructions were running on the cluster.

Simulation Experiments

4D simulations were performed on a 128 × 128 × 86 VE-ASL angiogram with six temporal frames. This angiogram was acquired using an undersampled acquisition ($R \approx 20$, scan time ≈ 20 min) and reconstructed on a frame by frame basis using a lightly L1-regularised SENSE reconstruction [11]. The ground truth image was downsampled and its matrix size was reduced from a previous in vivo acquisition that had been reconstructed frame by frame (spatial resolution: 1.6 mm isotropic, temporal resolution: 210 ms).

The images were reconstructed with 1024 spokes per frame ($R \approx 25$) with either the same spokes or hybrid method of sampling as described above.

Apart from different ground truth and sampling, the simulation and reconstruction was identical to the dynamic 2D experiments in the above sections with only spatial sparsity and temporal smoothness constraints.

Since a ground truth was available for the simulated data, the different reconstructions were compared to it with a masked correlation coefficient as previously described.

In Vivo Acquisition

The in vivo acquisition was performed on one healthy volunteer with the spoke ordering as explained above. No ground truth was available since sampling to the Nyquist limit would require the subject to lie in the scanner for approximately 8.5 hours. Instead data was acquired for approximately 10 minutes ($R \approx 40$ for same spokes and $R \approx 80$ for the hybrid method) and retrospectively undersampled and reconstructed from the first 5 min of scanning ($R \approx 80$ for same spokes and $R \approx 160$ for the hybrid method) and 1 min of scanning ($R \approx 400$ for same spokes and $R \approx 800$ for the hybrid method).

Since no ground truth data was available, the analysis was limited to visual comparison of the reconstructions.

4.5.2 Results

Reconstruction Computational Requirements

The simulations required approximately 11 GB of memory and approximately 6.5 hours of reconstruction time (wall clock time) for the same spokes method. The hybrid method required approximately 15 GB of memory, and ran in approximately 14 hours.

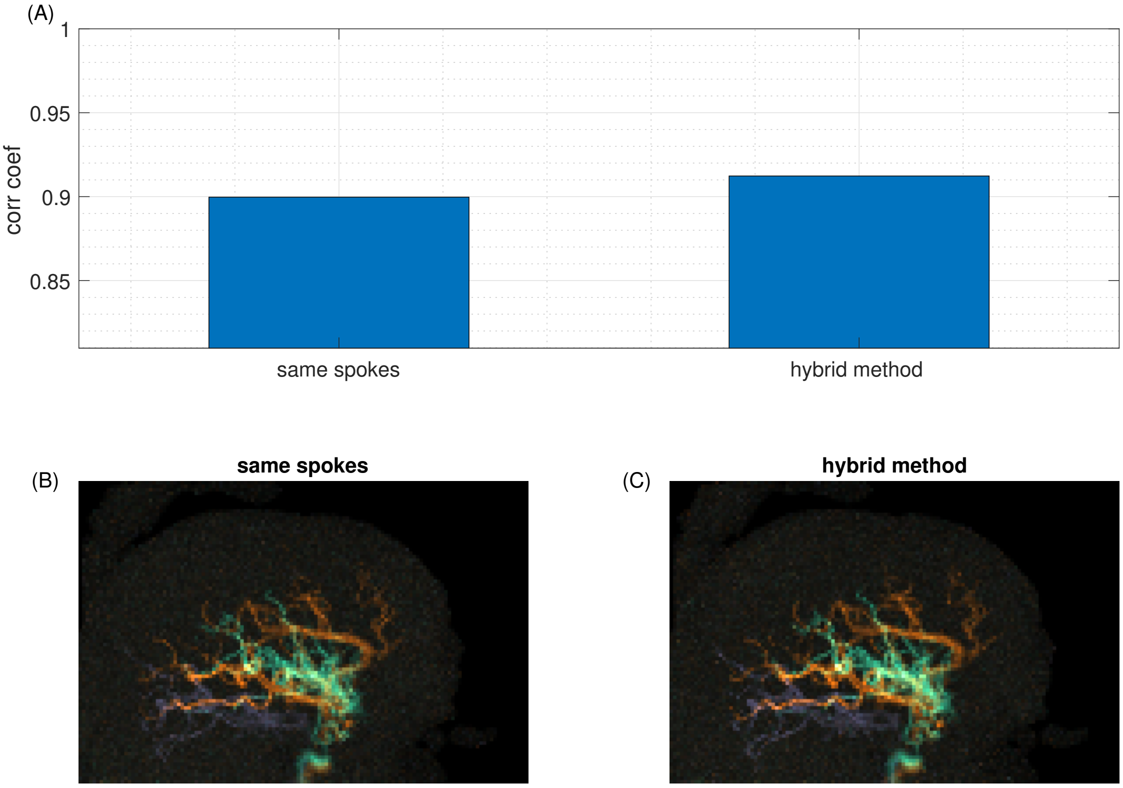

Figure 4.16: Image quality in 4D simulations.

The in vivo reconstructions required approximately 36 GB of memory for same spokes and took around 24 hours to reconstruct. The hybrid method required 54 GB and took approximately 50 hours to reconstruct. Again, the reported times are wall clock times.

Simulations

As hypothesised, the new sampling approach with no spoke overlap between frames resulted in a slight improvement by using the hybrid method over using the same spokes for each encoding (figure 4.16). The difference was, however, small (same spokes achieved a correlation coefficient of 0.90, the hybrid method 0.91), and the reconstructed images showed no large differences.

In Vivo

The in vivo data showed similar quality for both same spokes and the hybrid method for both 10 and 5 minutes of scanning (for the same spokes method $R \approx 40$ and

$R \approx 80$ and for the hybrid method $R = 80$ and $R \approx 160$ respectively), but with some degradation of signal in the 1 min ($R \approx 400$ for the same spokes method, and $R \approx 800$ for the hybrid spokes method) reconstruction (figure 4.17). The same spokes method showed more degradation (noise and loss of small and faint vessels) than the hybrid method did in the 1 minute case. Figure 4.18 shows additional views of the 5 min hybrid data as well as its time course.

4.6 Discussion

The examination of sampling-encoding interactions in section 4.2 has shown that spreading sampling incoherence across the encoding dimension is beneficial in sparse multi-dimensional reconstructions. While optimal sampling distributions are difficult to predict, m-PSFs are shown to be an intuitive tool for studying these interactions. Sampling-encoding interactions cannot be exploited when the decoding and reconstruction problems are treated separately, and this was shown in a simplified VE-ASL phantom system. This is a simple proof-of-principle application, and excludes considerations of coil sensitivity profiles that do interact with the sampling patterns. These ideas could be further applied to other multi-dimensional MR encoding techniques where traditionally multiple images are combined in post-processing, such as 4D flow imaging [88] or relaxometry [89].

In terms of sparsity, the vessels are already very sparse, and using either the same spokes every encoding, or the hybrid method, decouples the static tissue from the vessels such that the sparsity of the static tissue does not matter for reconstruction quality. To sample k-space without redundancies between encodings, the static tissue has to be considered more carefully. If the static tissue isn't (close to) perfectly reconstructed, its aliased signal will swamp the vessel signals. In section 4.3 it was demonstrated that sparsifying the static tissue using a patch based approach improved reconstruction quality only if the static tissue was suppressed to similar levels of intensity as thee vessels. The wavelet method did not provide the same benefit. In future work, the varying spokes method could be revisited combined with static tissue sparsification and better background suppression [90].

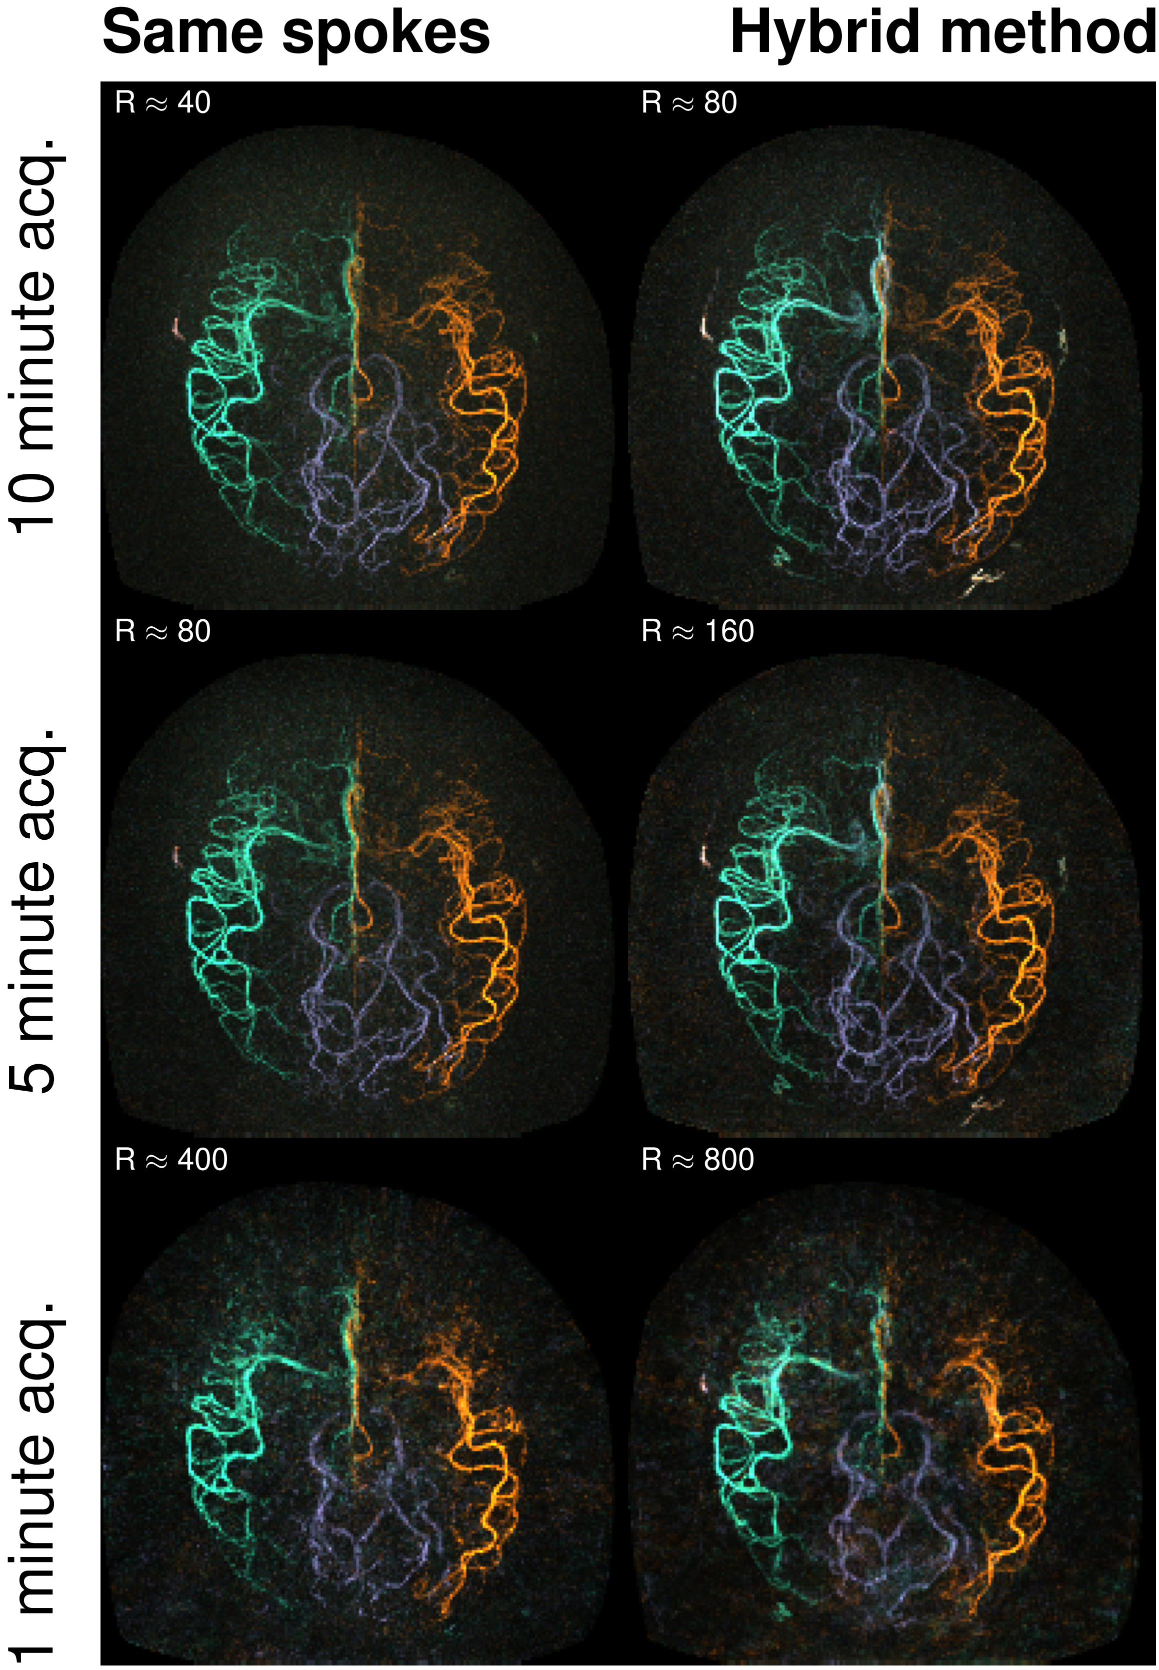

Figure 4.17: Image quality in 4D in vivo experiments.

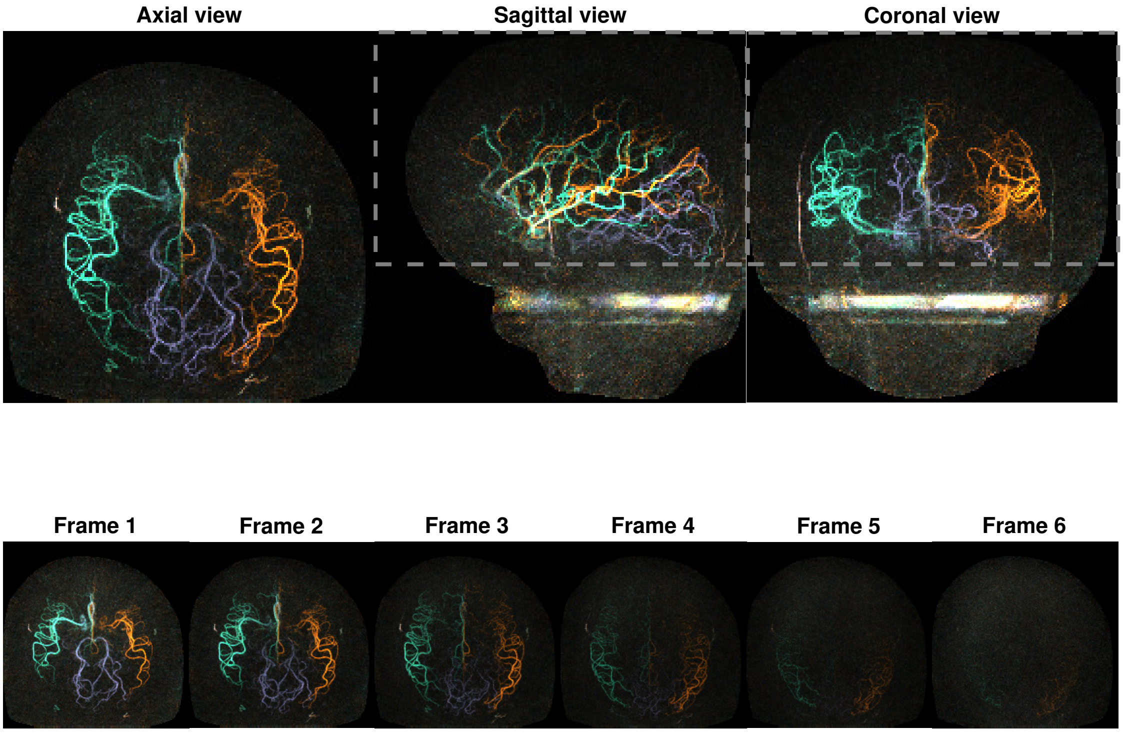

Figure 4.18: Different views of the 4D acquisition (5 minute acquisition with the hybrid method). The axial view is a MIP within the volume depicted by the dashed box in the sagittal and coronal view. This is to exclude the labelling plane artifact that can be seen in those views. The image series below shows the temporal evolution of the signal.

With current the acquisition method, decoupling the static tissue from the vessels was the better approach. However, in these experiments, it became evident that the same spokes method worked better than the hybrid spokes method when temporal regularisation was used. This incidental finding prompted a redesign of the sampling method for use in the 4D acquisitions to limit the amount of overlapping spokes in subsequent frames. The 4D simulation experiments confirmed that the hybrid method could bring small improvements when the frames used non-overlapping trajectories in combination with the L2-smoothness constraint.

Two new methods for including prior knowledge into the reconstruction method were trialled, one spatial and one temporal. Both could in theory have benefited the reconstruction, and similar methods have shown benefits in other applications [89, 91, 92]. However, the spatial method showed no benefit and carries the risk of producing a bias towards non-mixing vessel components that that will make the method even poorer if pathology is present. The temporal model could not constrain the signal sufficiently in a low dimensional subspace. Because four components were needed to explain 12 frames it did not sufficiently compress the signal and had inferior performance to L2-smoothness. For this reason it was not applied to the 4D data, where the number of frames was reduced to only six, further limiting the compressibility along that dimension. If the computational limitations can be overcome, the method could be trialled on acquisitions with more frames, where it would have a higher chance of producing benefits.

A large challenge with going to 4D imaging is the increase in data size, and thus processing power and time required to reconstruct the images. The current methods would not be suitable for clinical deployment due to the long reconstruction times although the scan time is well within a plausible range (5-10 min). In future work, the reconstruction code could be re-factored to use single precision throughout. Other improvements could also be achieved by optimising and parallelising portions of the code as currently it required looping through time points and encodings multiple times each iteration, or rewriting it in a more efficient language than MATLAB. The use of Toeplitz embedding should also be considered as the speed gained from using FFTs instead of the NUFFT in the iterations need to be weighed against the memory burden of the Toeplitz embedding.

The in vivo 4D acquisition was, nonetheless, promising. Especially the lack of meaningful differences between the 10 min and 5 min data suggests that both acquisitions fully managed to recover the signal and the difference in noise at these levels can be overcome by the denoising effect of the reconstruction. The slight improvement in quality of the hybrid method over the same spokes method at 1 min of scanning is also promising. It shows the benefit of considering all dimensions of the data, and how valuable it can be at very high undersampling factors.

The faster the scans can be, the less likely they are to be affected by patient motion due to, for example, discomfort during the scan. However, even in faster scans motion can still be an issue, and this method of imaging can be sensitive to subject motion. One way that motion can cause problems is if it occurs between the pre-scan that is used to set the labelling plane and vessel locations for encoding and the imaging sequence. Here, even relatively small changes in position can lead to reduced labelling efficiency. The other way motion can affect the scan is during imaging. Motion between encodings can cause imperfect decoding of the signal. However, the sequence interleaves encodings between shots, so the time between encodings is minimal. Motion correction methods [93, 94] could also be applied since the sequence is self-navigating by regularly sampling the centre of k-space, allowing a low resolution navigator to be reconstructed with a small number of spokes.

Both the 4D simulation and in vivo reconstructions were only assessed on a single subject/ground truth image, so more work would be required to confirm these results and explore extensions.

4.7 Conclusion

In this chapter, it has been shown that it is important to consider all dimensions of the data, especially when the undersampling factors are very high. In particular, regularisation across the temporal dimension was important for image quality, and with that the sampling needs to be adapted for the reconstruction approach.

Similarly, the spreading of artifacts across the component dimension showed to have small but visible benefits as long as the static tissue does not contaminate the vessel signal. This type of controlled aliasing across the components could also be applied in other types of MR imaging.

Finally, we have shown that it is possible to reach clinically feasible scan times for 4D VE-ASL angiography, however, improvements in reconstruction time still need to be achieved before it can take the step into clinical practice.

To infinity and beyond!

— Buzz Lightyear in Pixar's *Toy Story* 1995

5

Applications beyond VE-ASL angiography

Contents

5.1 Introduction 107
5.2 Time-encoded ASL angiography 109
  5.2.1 Methods 111
  5.2.2 Results 113
  5.2.3 Summary 114
5.3 Combined Angiography and Perfusion ASL 117
  5.3.1 Methods 117
  5.3.2 Results 119
  5.3.3 Summary 122
5.4 Discussion 122
5.5 Conclusion 125

5.1 Introduction

The methods presented in chapters 3 and 4 were developed in particular for VE-ASL angiography. However, many of the concepts used in those frameworks are generalisable to other types of cerebrovascular MR imaging. In this chapter, the basic reconstruction framework from chapter 3 is adapted and applied to two other flavours of ASL imaging in an attempt to achieve similar improvements in image quality at high acceleration factors that was found in VE-ASL angiography. The two domains are time encoded ASL angiography [95], and combined angiography and perfusion imaging [70].

Time encoded ASL [96, 97], or TEnc-ASL, is an alternative method to acquiring the temporal information of the bolus' passage through the arterial tree. The more common, and more straightforward method that has been used in previous chapters, is to acquire continuously after each labelling (tag or control) block for the desired readout duration and simply split up the data into a number of consecutively acquired frames. This is called the Look-Locker method [98], and requires small flip angles to prevent excessive signal attenuation from each RF pulse during the readout. In TEnc-ASL, on the other hand, the labelling period is modulated to carry the temporal information in a similar way to how the labelling plane is modulated in VE-ASL to carry spatial information. In TEnc, the signal can be acquired during a shorter readout period, and larger flip angles can therefore be used, which leads to a higher SNR, especially when combined with a variable flip angle scheme [99]. Using a simple re-gridding reconstruction, Woods et al. [99] demonstrated that TEnc-ASL and the variable flip angle scheme produced improved image quality compared to sequentially acquired angiograms even when the TEnc-ASL data was undersampled to match scan time of the sequentially acquired, fully sampled, data. In that study, only mild undersampling factors were used and the data was reconstructed linearly.

The introduction of non-linear reconstruction methods to TEnc-ASL angiography should allow for further acceleration, which will bring the method towards clinical feasibility. When limited by the Nyquist limit, TEnc-ASL takes longer than sequentially read out ASL to acquire, and for this reason this method is often not used despite its SNR benefit. In this chapter, an accelerated approach, again leveraging spatial sparsity and temporal smoothness of angiographic data, is developed based on the framework built for VE-ASL and assessed on TEnc-ASL angiograms with greatly reduced scan time.

Secondly, the power of golden ratio based sampling is extended to its full potential as the reconstruction methods are applied to combined angiography and perfusion imaging. Golden ratio sampling allows a single data set to be reconstructed with multiple spatial and temporal resolutions. The perfusion signal is much weaker than the angiographic signal because a longer PLD is needed to allow the signal to leave the arterial tree and perfuse into the brain tissue. Because a longer PLD is required, more T1 decay occurs. The perfusion signal in each voxel is also lower as the signal spreads out throughout the brain tissue. For these reasons, lower spatial and temporal resolution is generally needed to achieve sufficient SNR in perfusion imaging. The perfusion signal is, however, not as sparse in image space as the angiographic signal, so adaptations to the reconstruction method are needed.

The intrinsically low SNR of perfusion imaging makes accelerating it particularly challenging. Furthermore, the perfusion maps from simultaneous angiography and perfusion imaging suffer from further SNR-loss due to the early imaging pulses for angiography that suppresses the perfusion signal in later PLDs. Being able to remove undersampling artifacts, which add even more to the apparent noise level, is particularly important in this type of imaging.

5.2 Time-encoded ASL angiography

In "pure" TEnc-ASL, images with different PLDs are acquired by encoding the labelling period using a Hadamard encoding scheme, such that blood from different times can be decoded by combining all encoded images, similar to how blood from different vessels were encoded in VE-ASL. Figure 5.1 summarises three methods of acquiring dynamic ASL data; the conventional Look-Locker method, where all temporal information is contained in the readout period, pure time encoded ASL where all temporal information is contained in the labelling period, and a hybrid approach where some temporal information is contained in the labelling and some in the readout. In hybrid TEnc-ASL each readout frame is decoded using time encoding and labelling blocks have a longer duration to accommodate multiple readout frames in each encoding block.

The main limitation of time encoding (both pure and hybrid), similarly to vessel encoding, is that instead of just a tag and a control image, you need to acquire multiple encoded images. At least one more encoding than the number of temporal frames encoded in the labelling period is needed to also allow the separation of

Figure 5.1: Three different approaches to achieveing dynamic ASL imaging. The purple arrows between the ends of the labelling blocks to the centre of the readout blocks depict the PLDs.

signal from the static tissue. This means that in order to maintain the same scan time as a Look-Locker readout one needs to undersample each encoding to a greater extent by repeating the sequence fewer times (using fewer shots and thus filling k-space to a lesser amount). As seen in chapter 3, angiograms are spatially sparse even without multiple vessel components, and non-VE-ASL angiograms could be accelerated using sparsity based methods. TEnc can provide a small increase in sparsity compared to a conventional sequential readout as the sub-boli are shorter and, except for dispersion effects, fill up a smaller portion of the vascular tree in each decoded frame. Furthermore, TEnc-ASL requires us to think about how temporal smoothness is applied in a bit more detail, but otherwise it is well suited for acceleration using the same approach as used in chapter 3.

5.2.1 Methods

A single slab 2D TEnc-ASL dataset acquired for a previous project [99] on a 3T Siemens Verio system, with a 32-channel head-coil, in four healthy volunteers, was used in this study. The data had three readout frames (120 ms each) and three time-encoded blocks (360 ms each). When decoded, this resulted in nine frames with effective PLDs of 60 to 1020 ms in steps of 120 ms. Each readout frame was radially fully sampled (552 spokes) with uniformly distributed spokes, for a spatial resolution of $0.63 \times 0.63 \times 70$ mm^3, FOV = 220 mm^2, and slice thickness 70 mm. This required 46 repeats of each encoding (total scan time 4.5 min to reach the Nyquist Limit) as 12 spokes (TR = 10 ms) were acquired for each frame after each labelling period. The data was acquired with the variable flip angle scheme as proposed by Woods et al. [99] that compensates for signal attenuation from the previous RF pulses, meaning that the signal should vary smoothly from one TEnc block to the next, which does not happen if the flip angle is constant during readout. This is because with constant flip angle the signal is attenuated towards the end of the readout due to early excitations. So, unless this attenuation is counteracted by slowly increasing the flip angle during the readout, the images decoded from the latest block (frames 3, 6, and 9) will have lower baseline signal than the decoded images from the first block (frames 1, 4, and 7), which causes discontinuities between neighbouring frames (e.g. 3 and 4). VFA is therefore key for the ability to use temporal smoothness across blocks in the reconstruction.

The acquired data were retrospectively undersampled at $R = 12$, 23, and 46, which is the maximum acceleration factor, as the data is acquired in a single shot for each encoding. The undersampling was performed such that the resulting trajectory still had uniform coverage of k-space, and the three readout frames each had a different trajectory. For example, for $R = 12$, every 12th spoke was kept starting from the first spoke in the first readout frame, starting from the fifth in the second readout frame, and the ninth spoke in the third readout frame.

For simplicity, the same undersampling pattern was used for every encoding. In future work, varying the trajectory across encodings could be explored for benefits as were explored for VE-ASL.

Three methods were then used to reconstruct the data offline:

1. A naïve re-gridding of the non-Cartesian data using the NUFFT.
2. Non-Cartesian iterative SENSE parallel imaging reconstruction.
3. Compressed sensing reconstruction with sparsity in decoded image space and a temporal smoothness constraint.

The cost function for the CS reconstruction was defined as in previous chapters:

$$c(\mathbf{m}) = \frac{1}{2}\left|\mathbf{Em} - \mathbf{s}\right|_2^2 + \lambda_1 \left|\mathbf{m}\right|_1 + \lambda_2 \left|\nabla_t \mathbf{m}\right|_2^2 \tag{5.1}$$

with no static tissue sparsification because the same spokes were used each encoding. This cost function was minimised using FISTA, step size 0.01, 100 iterations. For the SENSE reconstruction, both λ_1 and λ_2 were set to 0, but otherwise the reconstruction was the same. For CS $\lambda_1 = 1 \times 10^{-6}$ and $\lambda_2 = 1$. These regularisation factors were chosen to be of similar magnitude as the high resolution data presented in chapter 3, that have similar appearance. For the naïve re-gridding, the data was simply reconstructed using the adjoint of the acquisition operator after pre-weighting for sampling density compensation using the weighting matrix $\mathbf{W}$: $\mathbf{m} = \mathbf{E}^{\mathrm{H}}\mathbf{W}^{-1}\mathbf{s}$. The coil sensitivity profiles included in $\mathbf{E}$ were estimated from the data averaged across readout frames and time-encodings using the adaptive combine approach [72].

A lightly regularised ($\lambda_1 = 10^{-7}$, $\lambda_2 = 0$) reconstruction using all the data was used as ground truth for analysis. Again, masks around the vessels were generated (by thresholding and dilating the masks) and correlation coefficients between the ground truth and the reconstructions were calculated.

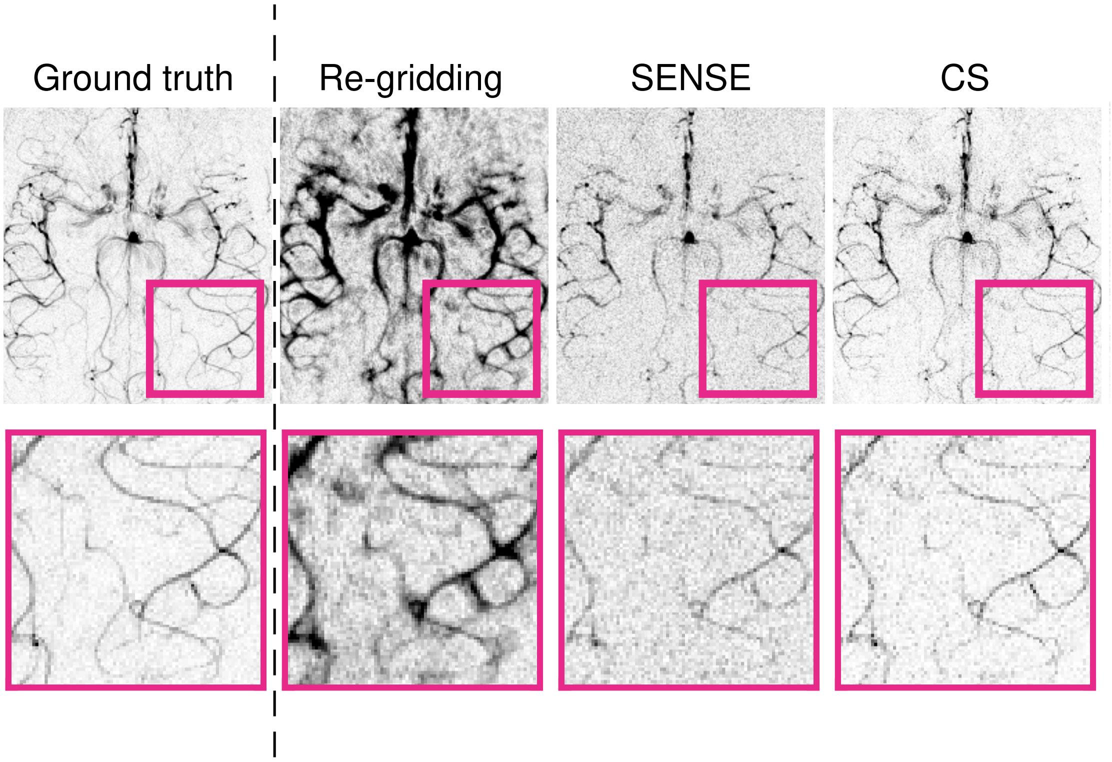

Figure 5.2: TEnc reconstructions (averaged across frames) at $R = 12$ compared with the fully sampled ground truth image.

5.2.2 Results

As figure 5.2 shows, the re-gridding reconstruction suffers from severe blurring due to the broadening of the PSF with undersampling. SENSE reconstructions are sharper but have taken a severe g-factor penalty compared with the re-gridded reconstruction. CS, on the other hand, retained both contrast and sharpness.

Figure 5.3 shows an example in another subject that again demonstrates the improved sharpness and higher apparent SNR in the CS reconstruction compared with SENSE and re-gridding.

In figure 5.4 the temporal reconstruction fidelity can be seen. The temporal dynamics are not biased due to the temporal smoothness constraint in the CS reconstruction framework. Again, the distal vessels in the late frames are both sharper than in the re-gridding reconstruction and have higher apparent SNR than the SENSE reconstruction.

Although the finer details are lost at very high undersampling factors, the CS

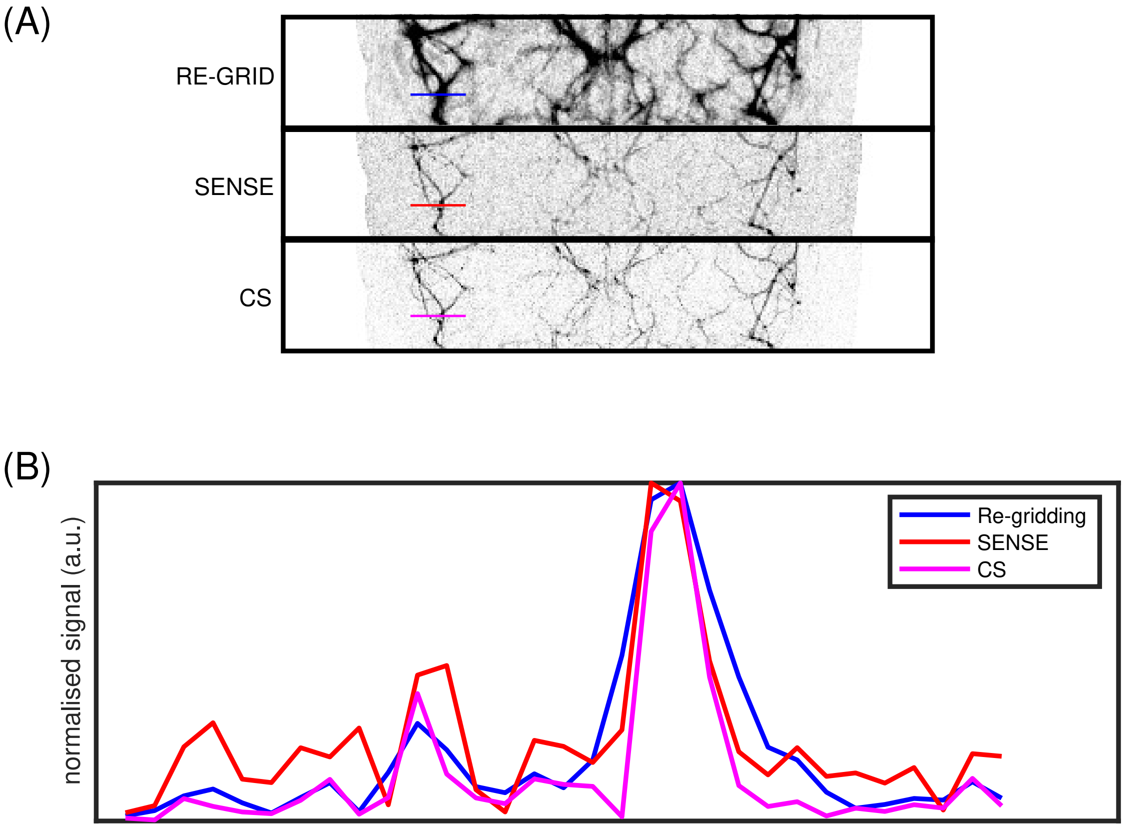

Figure 5.3: (A) shows a detail from the reconstructed images (R = 12, time averaged). (B) shows a line plot through two small vessels (marked with a line of matching colour in (A)). Re-gridding (blue) produces wider peaks and SENSE (red) has a higher noise floor than CS (magenta) that achieves both the narrowest peaks and low noise.

reconstruction method robustly reconstructs all the main arteries and showcases the range of image qualities available depending on how much scan time can be used. CS had consistently higher correlation coefficients than re-gridding and SENSE. Figure 5.5(A) shows how the image changes when the undersampling factor is pushed to a single shot acquisition approach. Figure 5.5(B), on the other hand, shows how the correlation coefficients vary as a function of R. The clinical significance of the reduced correlation coefficient at higher R, should be studied further by seeking expert opinions from radiologists.

5.2.3 Summary

This preliminary investigation into accelerating TEnc-ASL angiography has shown that CS outperforms re-gridding and SENSE type reconstructions, and that large

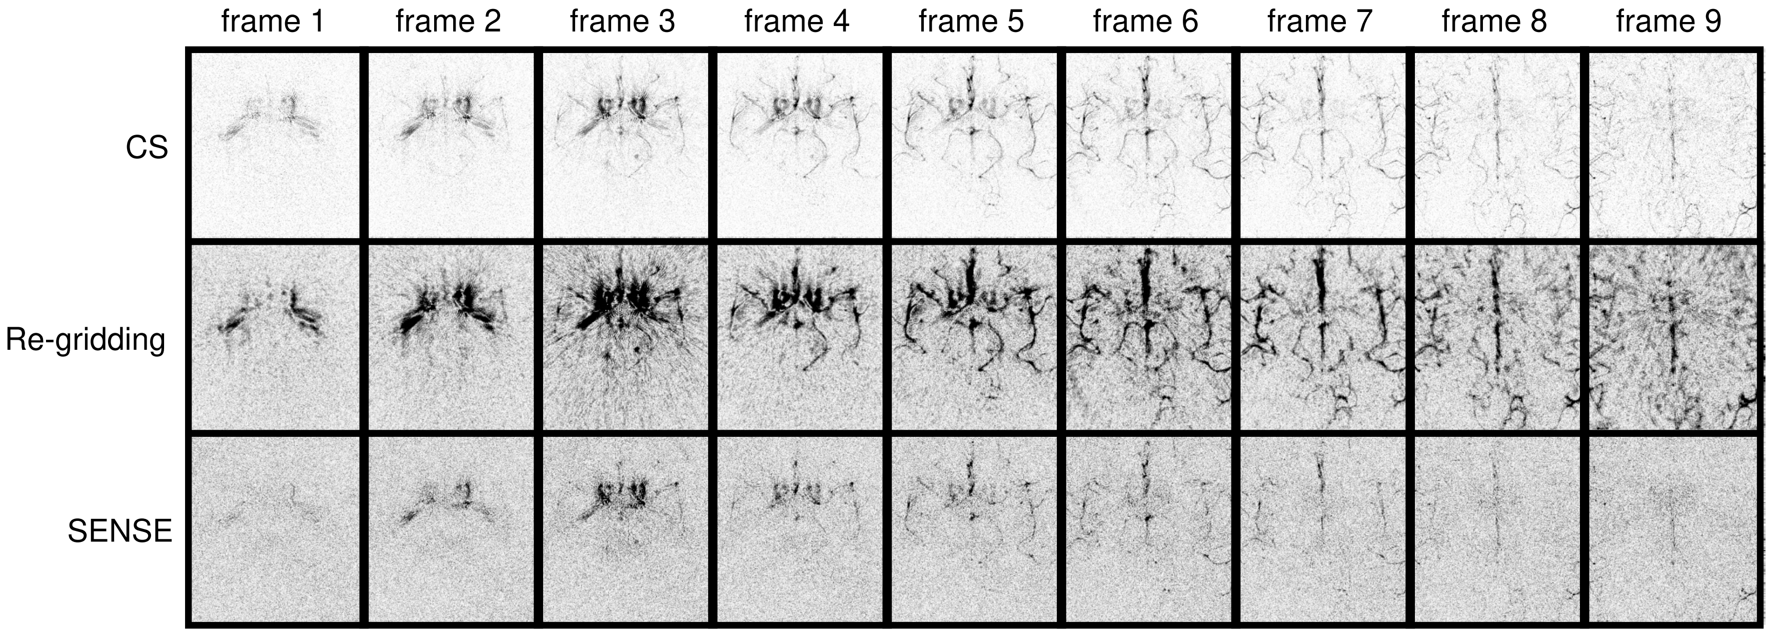

Figure 5.4: Frame-by-frame TEnc reconstructions (R = 12).

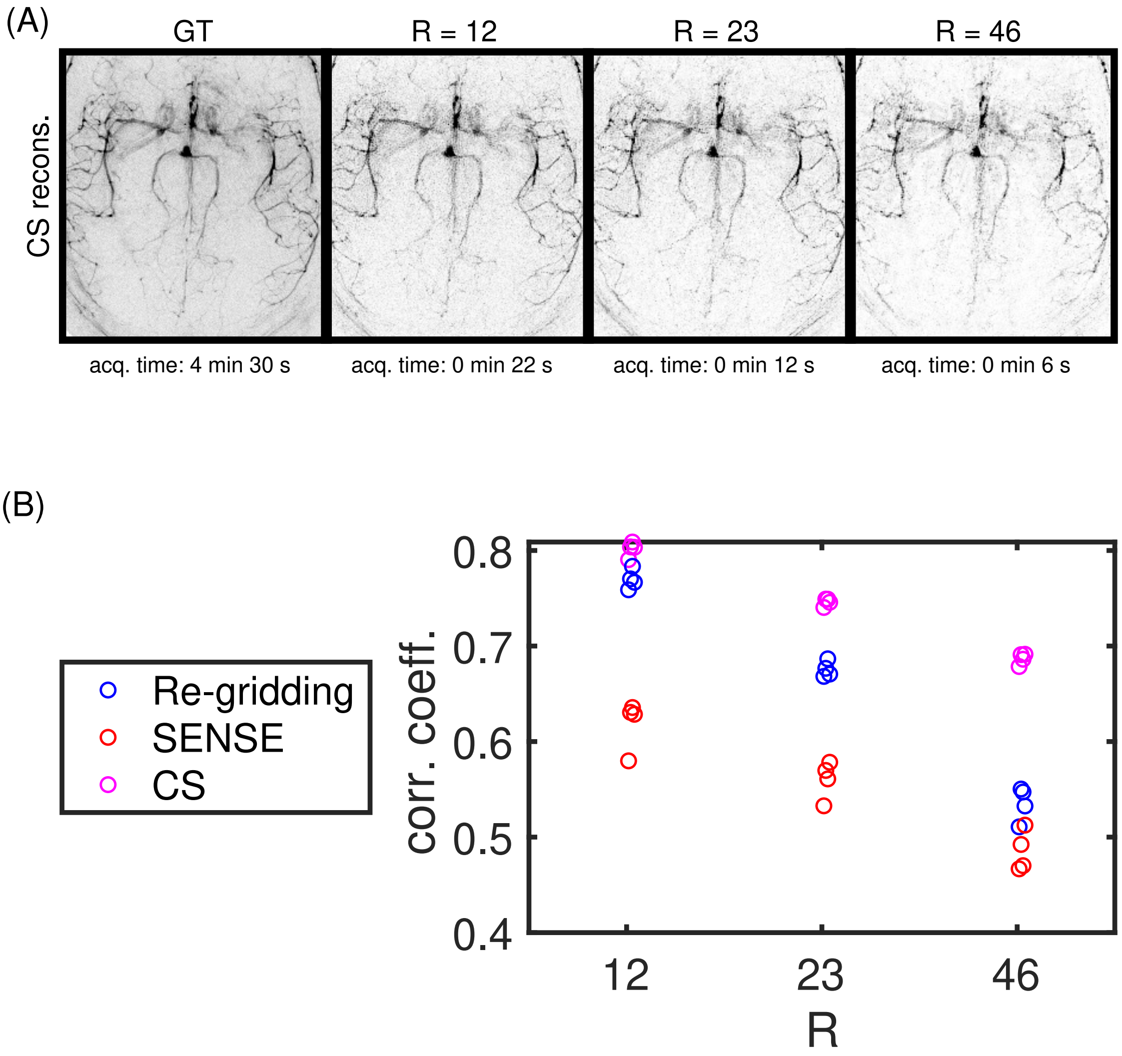

Figure 5.5: TEnc reconstruction quality at varying acceleration factors. (A) shows the CS time-averaged reconstructions (note that the ground truth image used a much lower regularisation factor in its reconstruction and thus has more residual background noise). (B) shows the measured correleation coefficients between the reconstructions and the ground truth. Each scatter point represents one subject.

acceleration factors can be achieved.

However, the CS method presented here was only assessed at a single combination of regularisation factors, which was chosen based on previous similar reconstructions. In future work the optimality of the regularisation factors would have to be assessed further in a larger cohort.

The method that was used in this section was also the simpler version of the reconstruction framework that was presented in chapter 3, but many of the extensions presented in chapter 4 could also be applied to TEnc data. This will be discussed further in the Discussion section (5.4) below.

5.3 Combined Angiography and Perfusion ASL

Combined angiography and perfusion using radial imaging and arterial spin labelling (CAPRIA) was first proposed by Okell [70], and initial work on reconstruction using both re-gridding [70, 100, 101] and non-linear methods [102, 103] have been presented. However, the non-linear reconstruction approaches have been limited to either reconstructing only the angiography part of a dynamic 2D CAPRIA dataset [102] or the perfusion of a 4D CAPRIA dataset [103]. In this chapter, the same dynamic 2D dataset will be reconstructed for both angiography and perfusion, and the limits of its acceleration factor to achieve good image quality for both types of images will be studied. The optimal regularisation factors to use in reconstruction of angiograms and perfusion imaging will also be studied; this aspect of non-linear reconstruction was not studied in previous work.

5.3.1 Methods

Four axial slices of non-VE pCASL data acquired with nominal resolution of $1.1 \times 1.1 \times 10$ mm^3 in a healthy volunteer were used as the test dataset for this study. The data was acquired on a 3T Siemens Verio system. Each slice was acquired with golden ratio sampling that was ordered such that each encoding had the same trajectory, but ideal golden ratio sampling was only achieved for a temporal resolution of 336 ms. The ordering is referred to as the t_{max}-method in the paper by Okell [70]. The labelling duration was 1.4 seconds, and readout duration 2 seconds. 588 spokes per frame were acquired at the optimal temporal resolution, 336 ms/frame (6 frames). Scan time of data acquisition was 2.5 minutes per slice (21 repeats of both the tag and control condition).

The CAPRIA data were reconstructed with the same three different approaches as used for TEnc-ASL; re-gridding, SENSE, and CS. Before reconstructing, subsets of the data were removed to simulate shorter scan times, and the remaining raw k-space data was split into either 18 frames (112 ms temporal resolution) for angiography, or kept at 6 frames (336 ms temporal resolution) for perfusion reconstruction. For angiographic reconstruction, the spatial resolution was kept at 1.1 mm in-plane (matrix size: 192×192), but the perfusion data were reconstructed with a spatial resolution of 3.3 mm in-plane (matrix size: 64×64) by discarding data in the periphery of k-space. The data were reconstructed with all 21 repeats for reference, and then reconstructed again with only 7, 2, or 1 repeat. This is equivalent to undersampling factors of $R_{192\times192} = \{1.5, 4.6, 17, 34\}$ and $R_{64\times64} = \{0.17, 0.51, 1.8, 3.6\}$.

For the CS based reconstruction, an L2 temporal smoothness constraint, as well as an L1 spatial sparsity constraint in image space, were applied to both angiography and perfusion data. The weighting between data consistency, image sparsity, and temporal smoothness was expected to be different for angiography and perfusion imaging, mainly due to lower sparsity and higher smoothness of the perfusion data. Therefore, a rough grid search of the regularisation factors was performed. The L1-sparsity regularisation factor, λ_1, was varied from 0 to 2×10^{-6} in steps of 2×10^{-7}, and the temporal smoothness regularisation factor, λ_2, was trialled at 0, to 6 in steps of 0.5. The grid search was performed to optimise the regularisation for the perfusion images that are very different to the angiograms that have been considered in previous chapters. The perfusion images are expected to have lower sparsity and higher temporal smoothness, and thus lower optimal λ_1 and higher λ_2 than for the angiograms.

Despite the lower inherent sparsity, the perfusion images were also reconstructed with sparsity in image space enforced by the CS algorithm. Previous preliminary work [103] has shown that perfusion images with only temporal regularisation can be reconstructed with CS methods. Other constraints, such as total variation [104] or total generalised variation [105], that encourages piece-wise smoothness in image space could also work in perfusion imaging. In this project, the simplest case, image sparsity, was trialled as an initial approach.

The data was again analysed by measuring the correlation coefficient between the 'ground truth' (fully sampled reference image reconstructed with a lightly regularised CS reconstruction, $\lambda_1 = 10^{-7}$, $\lambda_2 = 0$) and the reconstructed images

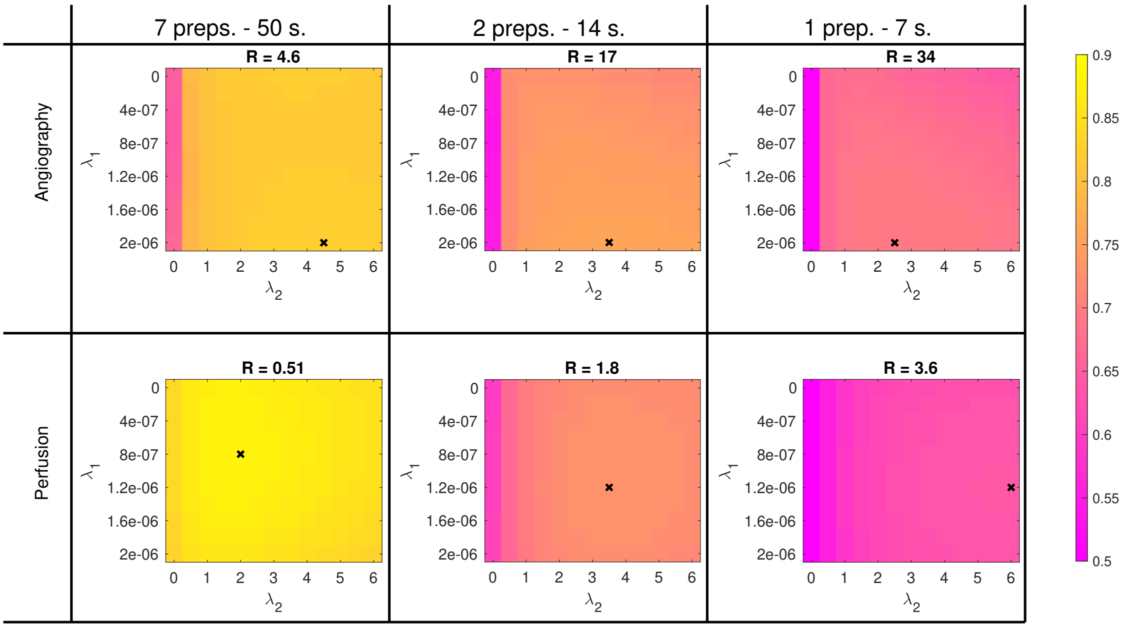

Figure 5.6: Regularisation factor grid search for CAPRIA. The colour represents the mean correlation coefficient for the four slices. The black cross in each sub-plot shows the optimum that was used in all further reconstructions.

in a brain mask for perfusion, and a dilated vessel mask for angiography. Only the frames with PLD > 1 s (frames 4 to 6) of the perfusion data were analysed, as earlier frames mainly contain blurred angiographic signal.

5.3.2 Results

The result of the regularisation factor grid search is shown in figure 5.6 averaged across the four slices. For the angiograms, the optimal temporal regularisation factor was $\lambda_2 = 4.5, 3.5$, and 2.5 respectively for increasing acceleration factors (R = 4.6, 17, and 34). λ_1 maxed out at 2×10^{-6} for all acceleration factors, which indicates that even higher λ_1 could be beneficial. The results are, however, similar to the results on non-VE angiograms in chapter 3 as expected. More interestingly, the regularisation optimisation for the perfusion data showed that even when $R < 1$, small amounts of regularisation for noise removal is beneficial. As expected, the optimal λ_1 for perfusion data was lower than for angiographic data.

With CS, for both perfusion and angiographic data, all acceleration factors and achieved an average correlation coefficient > 0.6 across the four slices.

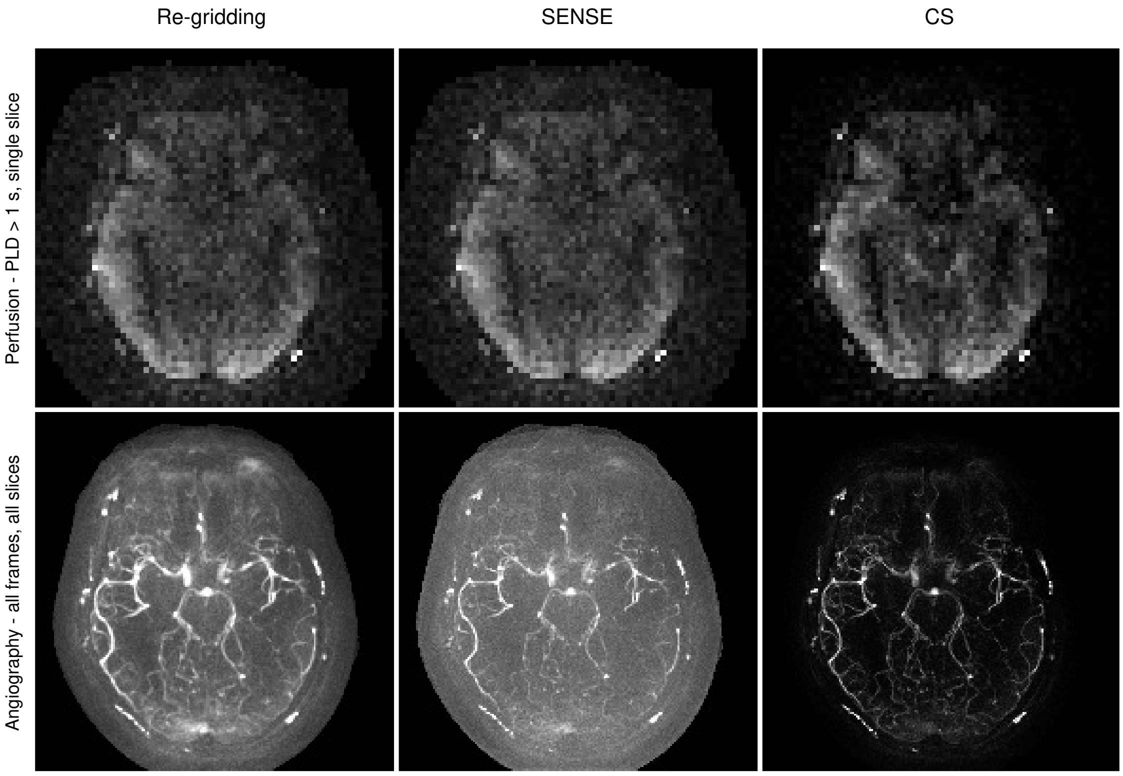

Figure 5.7: CAPRIA reconstructions for different reconstruction methods. All using 7 repeats (50 s. acquisition, R = 0.51 for angiography and 4.6 for perfusion)

Figure 5.7 shows examples of the achieved data quality for 7 preparations (scan time: approx 50 s per slice). The perfusion data is shown for a single slice averaged across frames 4 to 6, which have PLD's above 1 second. The angiographic data is combined across the four slices with a maximum intensity projection (MIP) and averaged across all 18 time frames. The denoising effect of CS is very clear in both the perfusion and angiographic images.

Comparison across the different acceleration factors are shown in figures 5.8 (qualitative assessment) and 5.9 (quantitative assessment). The ground truth angiography image has higher background noise level due to being reconstructed with minimal regularisation, and thus less denoising.

The temporal dynamics (three of the 18 frames of angiography MIPs and all three frames of single-slice perfusion weighted images) are shown in figure 5.10, showing that the CS reconstruction denoises the signal, but does not alter the temporal dynamics through excessive regularisation compared with the unbiased

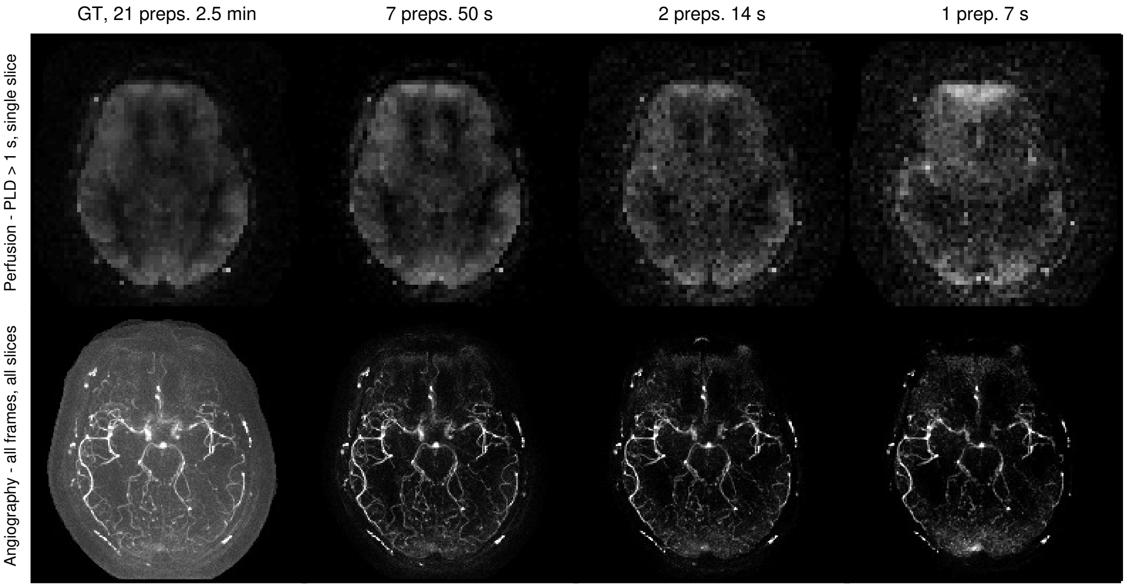

Figure 5.8: Qualitative result of reconstructing with different amounts of data. The ground truth (GT) angiogram has a darker background because it was reconstructed with a much lower regularisation factor, so has gone through less denoising.

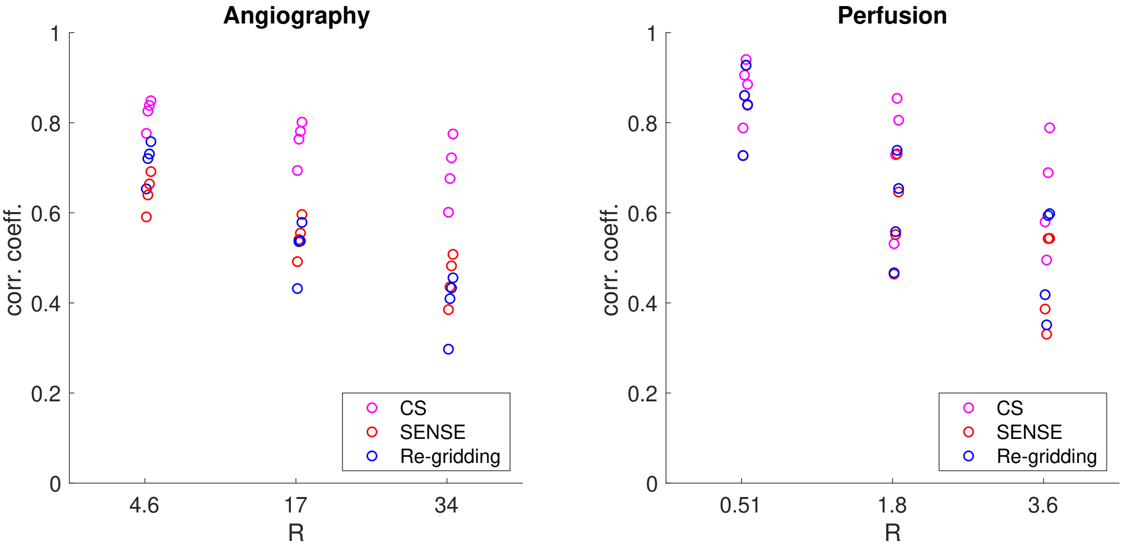

Figure 5.9: Quantitative result of reconstructing with different amounts of data. CS achieved consistently higher correlation coefficients than SENSE and re-gridding.

linear methods.

5.3.3 Summary

The proposed reconstruction framework works well on CAPRIA data. It appears that although perfusion data is not as sparse in image space as angiography, it benefits from the denoising properties as well as temporal regularisation in the framework.

The results showed that the single preparation (7 second acquisition) perfusion data is unlikely to have sufficient SNR for any meaningful interpretation, but already with two (14 second acquisition) or seven (50 second acquisition) preparations anatomical and physiological details are discernible.

5.4 Discussion

In this chapter the reconstruction methods developed in chapter 3 have been applied to new two applications: TEnc ASL angiography and CAPRIA imaging. Both methods achieved large acceleration factors using the proposed reconstruction framework without a large loss in image quality.

This is, however, very preliminary work, and for optimal performance both TEnc and CAPRIA imaging could be combined with, for example, ideas from previous chapters and with each other.

Since the mathematical framework for time-encoding is identical to that of vessel-encoding, the same hybrid encoding and sampling scheme that was presented in chapter 4 could be applied to time-encoded ASL angiography. Time-encoding boosts the SNR of the images compared to the conventional Look-Locker readout [99], and could thus provide benefits in other ASL methods. A varying spokes or hybrid approach could similarly also be applied to CAPRIA.

Combining vessel-encoding with time-encoding would be an interesting extension. However, to do full time- and vessel-encoding, the number of ASL preparations required would grow again, which further increases the need for these kinds of acceleration methods. Luckily, combination of TEnc with VE would increase relative sparsity of the images in the same way as shown in chapter 3, and thus likely, again,

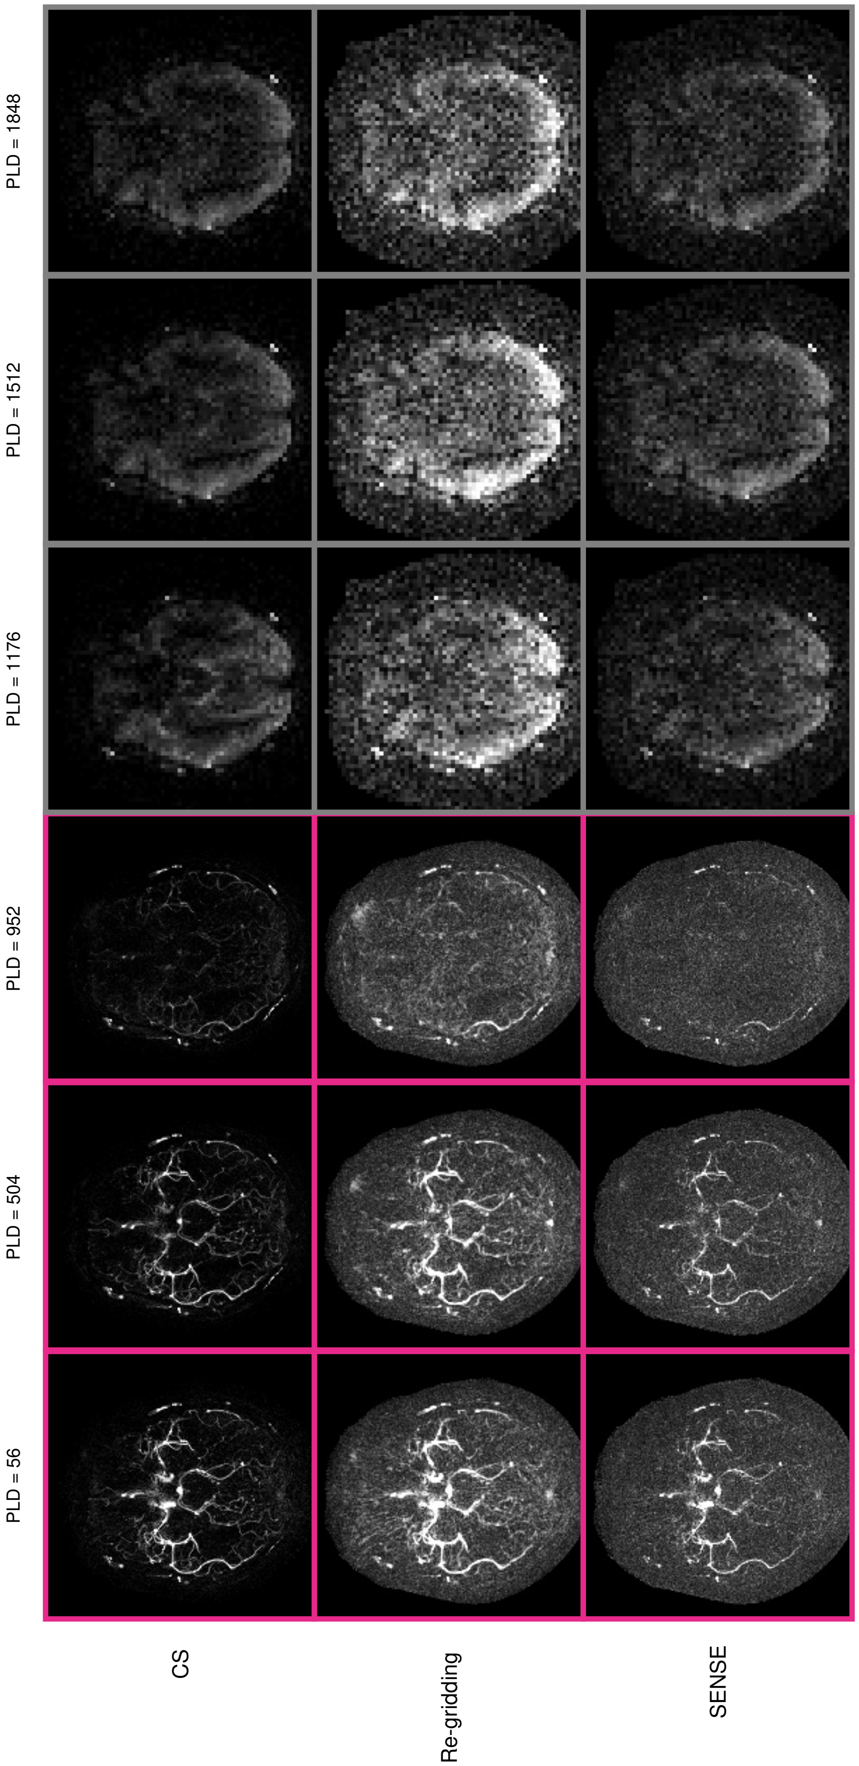

Figure 5.10: Example frames from a 50 s (7 preps.) CAPRIA acquisition. Three frames of angiography with PLD < 1 s (pink border), and three frames of the perfusion image with PLD > 1 s (gray border).

allow for increased acceleration. Opportunities to control aliasing across space, time, and vessel components by choosing trajectories and encoding approach jointly could further boost the achievable acceleration for such an approach. An alternative suggestion on how to combine vessel- and time-encoding was presented by Okell et al. [106], where only a single sub-bolus was additionally vessel-encoded. The increased SNR due to acquiring the data using time encoding could potentially push CAPRIA type reconstructions further as well. ASL perfusion data are notoriously low SNR and by optimising the acquisition protocol with time encoding, the sampling pattern for optimal sampling between encodings and k-space coverage (as will be discussed in the next chapter), and the reconstruction approach jointly, many small increases in image quality could boost its performance, especially when high acceleration is required. Improved SNR and/or reduced scan time for images that contain (potentially vessel selective) angiographic and perfusion information gives a very large amount of information from a single scan. This information would allow for detailed examinations of, for example, blood supply to lesions [107], distinction between arterial and venous phases in complicated blood supply [108], distinction of major feeding arteries [7], and the presence of "steal" phenomena affecting nearby tissue perfusion [109].

So far, the perfusion images have only been perfusion weighted, and no attempt at perfusion quantification was performed. Non-linear reconstruction alters the noise characteristics in the images, such that the noise cannot be assumed to be Gaussian, which common quantification methods [110] assume. One way of getting around this problem would be to include a perfusion model in reconstruction and use it to fit perfusion model parameters directly.

TEnc and CAPRIA are just two additional ASL methods the reconstruction framework was extended into. Further work could explore even more ASL methods.

An interesting direction could be to see how much acceleration could be achieved in a standard perfusion ASL acquisition without the additional angiography of CAPRIA that reduces the SNR of the perfusion images. For example, data could be acquired following the white paper ASL perfusion protocol recommended by the

ISMRM Perfusion Study Group and the *European ASL in Dementia Consortium* [111], and the achievable acceleration factors could be explored. However, the benefit of this method on data acquired with the white paper protocol might be limited due to the recommendation of using a single PLD, and much of the benefit of the proposed reconstruction comes from temporal regularisation.

Another method the reconstruction method could be trialled on is super-selective ASL angiography [7]. VE-ASL has SNR benefits over sequential super-selective acquisitions of multiple arteries, however, if only one feeding artery is studied, this labelling approach could have benefits. Again, this could be combined with time encoding and/or CAPRIA.

5.5 Conclusion

This chapter shows that the methods that were developed specifically for VE-ASL angiography, can be generalised and applied in different, although related, settings. The results presented here are preliminary but promising, and further work is required to optimise and explore these extensions fully.

As good as gold . . . and better.

— Charles Dickens in *A Christmas Carol* 1843

6

Efficient Radial Sampling Beyond the Golden Angle Method

Contents

6.1 Introduction 126

6.2 Theory 128

6.2.1 Properties of set increment sampling 128

6.2.2 Sampling and SNR 129

6.2.3 Measuring sampling efficiency 133

6.3 Methods 135

6.3.1 Optimisation 135

6.3.2 Simulation experiments 136

6.3.3 In vivo experiments 139

6.4 Results 141

6.5 Discussion 145

6.6 Conclusion 150

6.1 Introduction

Radial k-space sampling is often used in highly accelerated and/or dynamic imaging modalities, including in the work presented in the previous chapters. Among radial sampling methods, uniformly distributed spokes are the most efficient in terms of signal-to-noise ratio (SNR) because SNR scales with uniformity of sampling density [112]. However, uniform sampling does not allow for flexibility in reconstructing the same dataset at multiple temporal resolutions as done, for example, in CAPRIA (chapter 5). As shown in chapter 5, reconstructing the same data with multiple temporal resolutions can be useful for retrospectively analysing the data at different scales, for example through the angiograms and perfusion images. If all temporal frames in a readout period are sampled with differently oriented spokes the data can often be combined across those frames to create a fully sampled temporal average image that can be used, for example, for coil sensitivity estimation [78]. In other types of imaging applications, it can be impossible to know a priori how many spokes will be combined to reconstruct an image. An example of this is when the number of spokes per frame is unknown at the time of acquisition because data is retrospectively binned based on respiratory or cardiac phases. This makes uniform radial sampling unsuited for use in these applications.

In all previous chapters, versions of a commonly used alternative to uniform sampling, the radial golden ratio (GR) method [35], have been used for flexibility in both temporal resolution and retrospective undersampling. In GR based methods, the direction of each k-space sampling spoke is determined as a set angle increment from the previous spoke, such that each new spoke intersects the largest gap in k-space by the golden ratio. This sampling approach results in relatively high uniformity for any number of subsequently acquired spokes, with peaks in uniformity at a number of spokes equal to a member of the Fibonacci sequence (1, 1, 2, 3, 5, 8, 13, 21, 35, 56...). This allows for good flexibility in reconstructing a dataset with any number of spokes in each window. Additionally, subsequent frames are simply sampled with a rotated version of the previous frame's k-space sampling pattern, so they have the same sampling SNR efficiency when considering the trajectory alone, and ignoring coil sensitivity considerations. Subsequent non-overlapping frames contain no repeats of the same spoke, which is beneficial when regularisation or view sharing is applied in the temporal domain for dynamic imaging reconstructions.

In this chapter, a method of sampling is presented that maintains the favourable aspects of golden ratio sampling (flexible temporal resolution, constant sampling SNR efficiency for sliding windows, and non-overlapping subsequent frames), but optimises it for a user defined set of window sizes (numbers of spokes per frame) instead of being limited to the Fibonacci numbers. GR sampling methods get approximate uniformity for all possible window-sizes, however, this is unnecessarily general for most applications. Imaging experiments are generally designed for reconstruction with a set number of specific temporal resolutions, or, in the case of retrospectively binned data, a small range of potential window sizes. The method presented here is thus an alternative to GR sampling, where a restricted set of window sizes is optimised for. The main aim of this work was to develop a method of choosing the angle increment between subsequent spokes such that it maximises the minimum uniformity of the sampling within a pre-defined, application specific, set of window sizes. By relaxing the requirement for 'near-uniformity' to only apply to a specific set of window sizes, it was hypothesised that one could find a tradeoff between the flexibility of GR sampling with the SNR efficiency of uniform sampling. The proposed method is here compared with both GR and with uniform sampling in both simulations and in vivo non-VE-ASL angiography data. The proposed sampling method will be referred to as SILVER, the Set Increment with Limited Views Encoding Ratio method.

6.2 Theory

6.2.1 Properties of set increment sampling

A set increment for a 2D radial trajectory is simply a constant angle increment ($\theta = \alpha \times 180°$) from the previous spoke. How the data sampled with a set increment fills k-space depends on whether this step is a rational or irrational fraction of the whole circle.

When the step ratio, α, is a rational fraction, the exact same spoke will eventually be repeated, whereas if the step is irrational or rational with a very large denominator in its simplest form (and thus practically irrational), no two spokes acquired within the duration of the experiment will be the same. In dynamic imaging experiments where the length of a frame is unknown a-priori, irrational increments are preferred to avoid acquiring data from the same spoke location multiple times within the frame or leaving large gaps in k-space, both of which result in suboptimal image SNR. Similarly, if multiple temporal resolutions are required it is also beneficial to use irrational sampling to avoid duplicate spokes in the different window sizes, especially if data are combined across frames to be reconstructed such that the Nyquist criterion is met at some longer temporal resolution. For this reason, the golden ratio, which is often referred to as the most irrational number [113], is often used to sample when no periodicity is wanted, and every new spoke should fill k-space with near-optimal uniformity.

As mentioned above, by acquiring radial k-space data with a set angular increment, regardless of whether the step ratio is rational or irrational, images with the same sampling efficiency can be reconstructed with any set of N subsequently acquired spokes. Each k-space trajectory will simply be a rotated version of the N previous spokes, rotated by $N\alpha \times 180°$ (figure 6.1). This allows for sliding window [114] and view sharing [8] reconstructions with complete flexibility in where to start and end each frame.

Uniform radial sampling with full width spokes can be achieved with set increment sampling by choosing the angular increment between subsequently acquired spokes to be $\alpha_{uniform} = 1/N$, where N is the number of spokes used to reconstruct one frame. In GR sampling, on the other hand, the step is instead fixed at approximately 111.25°, which corresponds to $\alpha_{GR} = 1/\phi = (\sqrt{5} - 1)/2 \approx 0.6180$. Alternatives to the GR method, such as the tiny golden angle method [115], use a smaller increment but still achieves equivalent sampling to GR if the desired window size is above a certain minimum number of spokes. The aim of SILVER is to generalise further, and find a more optimal increment, $0 < \alpha < 1$, for an arbitrary set of window sizes.

6.2.2 Sampling and SNR

It is well known that sampling with non-uniform density leads to noise amplification in MRI [112, 116]. This section will revise why and how that noise amplification occurs.

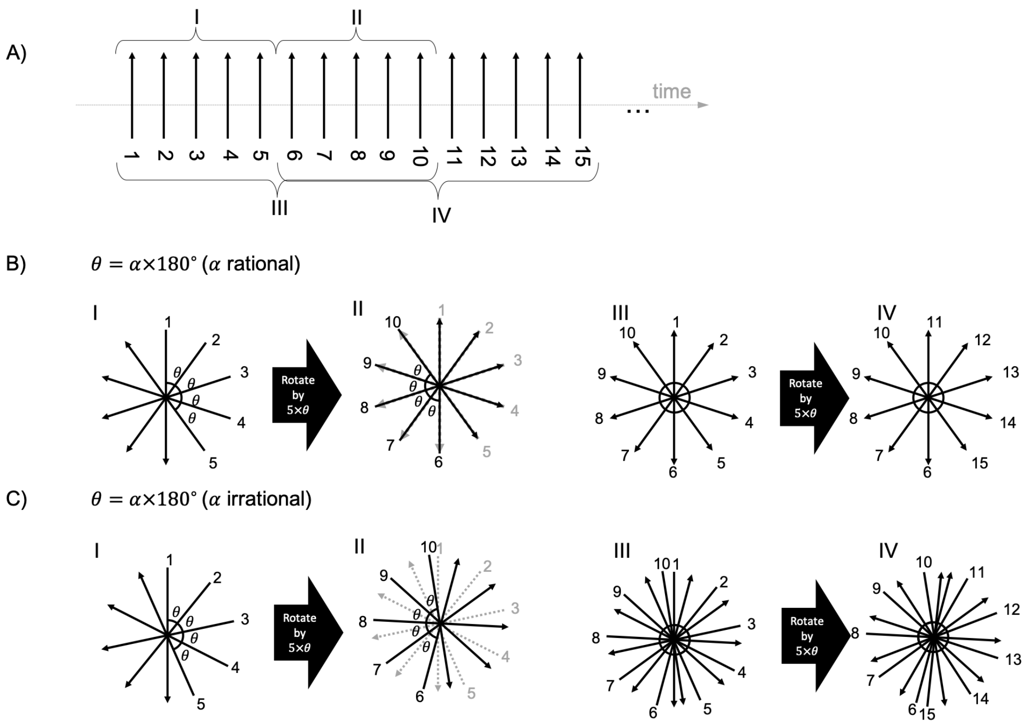

Figure 6.1: The effect of set increment radial sampling with rational and irrational increments. (A) demonstrates four different groupings of subsequent spokes (I, II, III, and IV) and gives each spoke a number. (B) shows how the subsequent spokes in the different frames relate if α is rational (in this case 1/5), showing how both non-overlapping and overlapping frames have repeats of the same spokes. (C) shows the effect of an irrational α; Non-overlapping subsequent frames have no repeat spokes, and overlapping frames have some of the same spokes. In both (B) and (C) a subsequent frame (whether overlapping or not) is simply a rotation of the previous frame and will thus have the same level of uniformity in each case.

When sampling k-space in a non-uniform manner there are two main approaches to reconstructing the signal. The first is to re-grid the signal onto a Cartesian grid and then proceed to use the conventional inverse fast Fourier transform for efficient reconstruction. When re-gridding, the data need to be weighted based on their local sampling density to avoid effectively filtering the output image by the sampling density function. Assuming that each data point is sampled with uncorrelated Gaussian noise with a variance σ^2, the image variance, σ_{im}^2 depends on the local sampling densities as they are used as weighting factors on each sampled point:

$$\sigma_{im}^2 \propto \sum_i^N \frac{\sigma^2}{D(s_i)^2} = \sum_i^N w_i^2 \sigma^2 \tag{6.1}$$

where $D(s_i)$ is the local sampling density around the i^{th} k-space sample, N is the number of sampled points, and $w_i = \frac{1}{D(s_i)}$ is thus the weighting factor applied to each sample.

Assuming the signal can be recovered fully, image SNR is fully determined by the standard deviation of the image noise. Thus:

$$\begin{aligned} SNR_{im} &\propto 1/\sigma_{im} \\ &\propto 1/\sqrt{\sum_i^N w_i^2 \sigma^2} \\ &\propto 1/\sqrt{\sum_i^N w_i^2} \end{aligned} \tag{6.2}$$

It can be derived from the Cauchy-Schwartz inequality in combination with the assumption that $\sum_i^N w_i = 1$, that $\sum_i^N w_i^2$ always is greater or equal to 1/N (which leads to the familiar relationship between SNR and the square-root of the number of samples). However, the equality can only be achieved if $w_i = 1/N$ for all $i \in [1, 2, 3, ..., N]$. Since non-uniform density leads to non-constant weights, non-uniform sampling results in worse SNR assuming that the total number of samples remains constant. The derivation is available in appendix B.

When the acquisition is undersampled (some gaps in k-space are so large that aliased signal ends up in the field of view) the situation is more complex as we may no longer be guaranteed to recover the signal perfectly.

Re-gridding modifies the sampled data such that it can be assumed to have been sampled on a regular grid in k-space. This makes the transform between image and k-space orthogonal (regularly sampled DFT), and SNR penalty is completely determined by the weighting that is applied to the data. Methods such as the Pipe-Menon method for determining weighting factors [117], optimise reconstruction fidelity for non-uniformly sampled data but take the SNR penalty described by equation 6.1.

The other option is to invert (or pseudo-invert) the actually used acquisition operator, either directly or through iterative methods. If k-space is sampled off a

Cartesian grid, the acquisition operator, $\mathbf{E}$, will be a non-orthogonal matrix, and this method of reconstructing leads to noise amplification in a similar manner to how non-orthogonal coil sensitivity maps lead to noise amplification in parallel imaging.

Assuming a pseudo-inverse reconstruction:

$$\tilde{\mathbf{m}} = (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{s} \tag{6.3}$$

with $\tilde{\mathbf{m}}$ being the reconstructed image, $\mathbf{E}$ the acquisition operator and $\mathbf{s}$ the sampled signal, and the signal being generated using the Gaussian noise model:

$$\mathbf{s} = \mathbf{E}\mathbf{m} + \mathbf{n} \tag{6.4}$$

with $\mathbf{m}$ being the actual object, and $\mathbf{n}$ white Gaussian noise. Combining equations 6.3 and 6.4, we get:

$$\begin{aligned}\tilde{\mathbf{m}} &= (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}(\mathbf{E}\mathbf{m} + \mathbf{n}) \\ &= (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{E}\mathbf{m} + (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{n}\end{aligned} \tag{6.5}$$

The variance of $\tilde{\mathbf{m}}$, that determines image SNR is then:

$$\sigma^2_{\tilde{\mathbf{m}}} = E[(\tilde{\mathbf{m}} - E[\tilde{\mathbf{m}}])(\tilde{\mathbf{m}} - E[\tilde{\mathbf{m}}])^{\mathrm{H}}] \tag{6.6}$$

where $E[\tilde{\mathbf{m}}]$ is is the expectation value of $\tilde{\mathbf{m}}$ and is:

$$E[\tilde{\mathbf{m}}] = (\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{E}\mathbf{m} = \mathbf{m} \tag{6.7}$$

Now, combining this result with equation 6.5 and 6.6. We can see that the variance becomes:

$$\sigma^2_{\tilde{\mathbf{m}}} = E[((\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{n})((\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1}\mathbf{E}^{\mathrm{H}}\mathbf{n})^{\mathrm{H}}] \tag{6.8}$$

which, for uncorrelated noise ($E[\mathbf{n}\mathbf{n}^{\mathrm{H}}]$ is identity scaled by the noise variance $\sigma^2_{\mathbf{n}}$) becomes:

$$\sigma^2_{\tilde{\mathbf{m}}} = \sigma^2_{\mathbf{n}}(\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1} \tag{6.9}$$

which demonstrates how $\sigma^2_{\tilde{\mathbf{m}}}$ is dependent on the conditioning of $\mathbf{E}$.

Radial sampling in a linear acquisition and reconstruction framework thus has intrinsically lower SNR than Cartesian sampling due to the variable density with which k-space is sampled, whether or not it is reconstructed via re-gridding or (pseudo-)inversion. The SNR predictions, regardless of which linear reconstruction approach is taken, are comparable, with differences boiling down to practical implementation relating to things like the interpolation kernel used in re-gridding and the de-apodisation function, etc. This was confirmed using simple Monte-Carlo simulations of a Shepp-Logan phantom using three different reconstruction methods: density compensated re-gridding, pseudo-inversion, and an interative reconstruction (figure 6.2). The simulations showed similar SNR for all three methods, but the re-gridded image showed considerably more systematic errors. The fact that image SNR depends primarily on the sampling density [116] can also be understood by considering that all implementations approximate the same linear transformation from k-space to image space, and any method that produces a nearly unbiased image of the original object can be expected to transform noise in a similar way.

Among radial sampling methods, radially uniform sampling achieves the highest SNR, and performance for a fixed number of spokes degrades the more non-uniform the sampling is. This can be understood based on the arguments regarding uniformity shown above. Along the spokes all radial trajectories have the same uniformity, whereas for a set radius, the local sampling density will vary based on the chosen spoke directions.

6.2.3 Measuring sampling efficiency

Because SNR intrinsically depends on the uniformity of the local sampling density, different methods for estimating the local sampling density have been proposed. For 2D radial sampling in particular, Winkelmann et al. defined sampling density by the inverse of the average azimuthal distance between adjacent spokes [35]. For spiral imaging the sampling density has been estimated using the speed of spiralling around k-space [118], and, for example, using voronoi cells [119]. For 3D radial sampling numerically defined approaches using voronoi cells on spheres have also been used

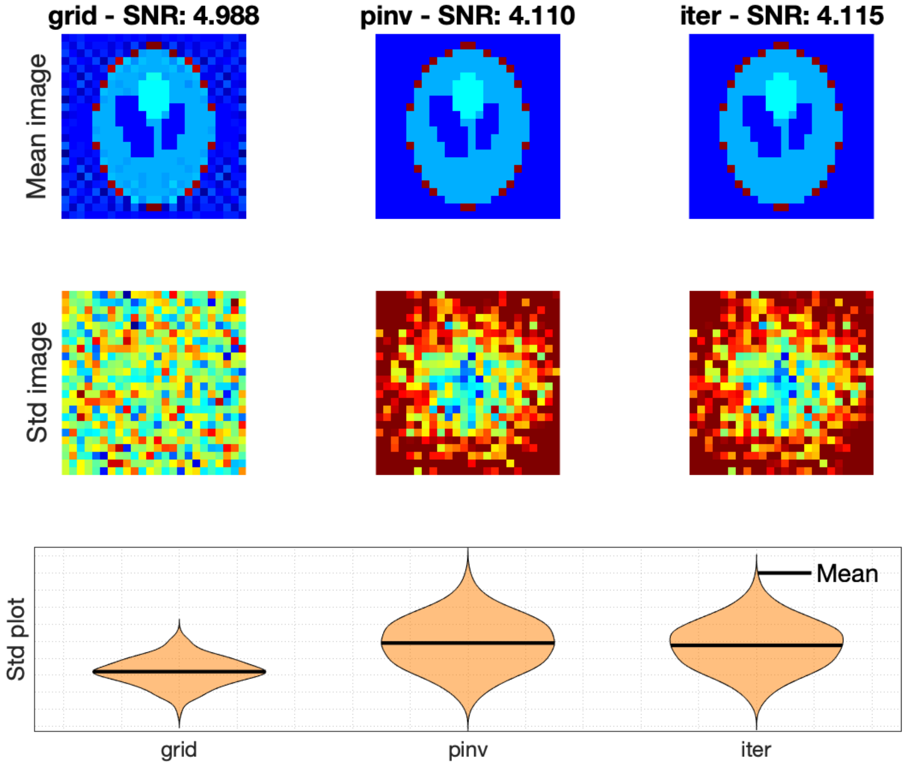

Figure 6.2: Monte Carlo simulations to measure SNR using different reconstruction methods. The data was simulated with 36 uniformly sampled spokes, reconstructed onto a 24 × 24 grid using only one uniform sensitivity coil. 1000 simulations were run to measure the noise standard deviation. Re-gridding has the highest SNR, but considerable bias. Pseudo-inversion (pinv) and iterative reconstruction (iter) solves the exact same mathematical problem and thus have near-identical results.

[34]. An alternative to these methods of considering density distributions, specific to radial sampling, is to consider a physical model based on electrostatic force or potential between charges placed on the spoke tips. This constitutes a version of the Thomson problem (how charges would naturally be distributed on a sphere with minimal potential energy) where each charge has a pair on the opposite side of the sphere [120, 121]. Methods like this are extensively used to choose diffusion sampling directions in diffusion tensor imaging [122, 123]. Electrostatic methods, although more commonly used in 3D sampling, can be generalised to 2D sampling. An electrostatic potential minimising model was used in this version of SILVER because of its easy extension to 3D in the future, and its high penalty for overlapping spokes.

In the electrostatic potential model the radial sampling pattern is treated as an ensemble of unit charges placed on both ends of each spoke, constrained to the unit circle (or unit sphere if extended to 3D). The total energy of the system is thus defined as:

$$U = \sum_{i,j=1}^{2N(i \neq j)} 1/r_{ij} \tag{6.10}$$

where r_{ij} is the distance between the $\mathrm{i^{th}}$ and $\mathrm{j^{th}}$ points, and N is the number of spokes ($2N$ is therefore the number of spoke tips).

Sampling efficiency compared to radially uniform sampling, η, was then defined as the ratio of total electrostatic potential stored in the system of point charges, U, to a system with the same number of spokes in the lowest possible energy state (uniformly distributed spokes), U_{ref}:

$$\eta = U_{ref}/U \tag{6.11}$$

U_{ref} is included in the efficiency metric η as a normalising factor, so that comparison of efficiencies at different N is possible. For sampling with a set increment, U generalises to a function of only the set increment, α, and the number of spokes, N. For a uniform radial distribution $\alpha = 1/N$ so U_{ref} is simply a function of N. The points are symmetrically distributed because each sampled spoke has one perceptual charge on each end.

Combining this with equation 6.11 we get:

$$\eta(\alpha, N) = (U_{ref}(N))/U(\alpha, N) \tag{6.12}$$

6.3 Methods

6.3.1 Optimisation

The SILVER method was formulated as an optimisation problem, where the task is to maximise the minimum efficiency, η, for a pre-defined set of window sizes, $S = \{N_1, N_2, N_3, \dots\}$. The objective function was therefore defined as:

$$\max_{\alpha}(\min_{N \in S}(\eta(\alpha, N)) \tag{6.13}$$

with η defined as in equation 6.11, N is the number of spokes in the window, and α is the set increment (as defined in section 6.2.1). The optimization was performed in MATLAB R2018b (The MathWorks, Inc., Natick, Massachusetts, United States) using optimcon() in the Optimization Toolbox using the interior-point minimization algorithm. To avoid local minima the optimisation algorithm was restarted 100 times with 99 initial values of α drawn from a uniform probability distribution between 0 and 1 and one run with the golden ratio as the starting value for α.

SILVER was compared to GR sampling for a large range of plausible sets of window sizes, S. The efficiency of both methods was measured using the electrostatic potential method as described in section 6.2.3. First, continuous ranges of window sizes were explored. These sets contained a minimum window size of M spokes, and all intermediate window sizes up to a maximum window size, $M + P$, such that $S = \{M, M + 1, \ldots, M + P\}$. M was set to 4, 16, and 32, and P was set to all integers from 1 to 100.

Similarly, the effect of optimising for multiple specific temporal resolutions was examined. The efficiency for sets with $S = \{M, 2M\}$, and $S = \{M, 2M, 3M\}$ were studied. Again, M was set to 4, 16, and 32. As a final case, where the GR approach is expected to perform optimally, S was set to consist of Fibonacci numbers ($S = \{5, 8, 13, 21, 34\}$).

6.3.2 Simulation experiments

To study how the theoretical benefit of a lower cost function (electrostatic potential) translates into image quality, a range of simulations were performed.

First, a simple experiment with a simple 64 × 64 digital dynamic phantom (figure 6.3) as ground truth was simulated with an oversampled acquisition on a single channel system. This allows for analysis of the sampling efficiency in the absence of undersampling artifacts and noise-amplification due to the coil sensitivity profiles. Sampling was done either with radially uniform, GR, or SILVER sampling.

Figure 6.3: Dynamic phantom used for SILVER experiments. It is temporally symmetric with the first 30 frames being reversed in the last 30 frames. The "tongue-like" and "jaw-like" structures move between frames, everything else is stationary.

The phantom had 60 frames, and each frame was reconstructed with 125, or 150 spokes ($R \approx 0.81, 0.67$ respectively). For uniform sampling α was thus set to either 1/125, or 1/150. For GR α was by definition $1/\phi \approx 0.6180\ldots$, and the SILVER optimisation for the set S = {125,150} resulted in $\alpha = 0.2080....$ In k-space, complex white noise with a standard deviation of 1000, $SNR_k = 37$ (calculated the same way as in chapter 3, equation 3.4) was added. Reconstruction was performed iteratively with no regularisation until convergence (the L2-distance between the forward transformed reconstructed image and the simulated data changed less than 1% between consecutive iterations).

Then, an undersampled, multi-coil acquisition and reconstruction of the dynamic phantom was performed. A set of eight coil sensitivity maps (compressed from a 32-channel phantom measurement) was used. The images were reconstructed with iterative SENSE, again running the algorithm to convergence. This time, a slightly more stringent convergence criterion than for the oversampled acquisitions was used; to ensure convergence of the more poorly conditioned problem, the change in data consistency was required to go below 0.1% (based on preliminary experiments). For these reconstructions, 48, or 64 spokes per frame were used, corresponding to undersampling factors $R = 2.10$ or 1.58 respectively. The SILVER optimised α was 0.3539 (rounded). Everything else was kept the same as for the single coil experiment.

The reconstruction quality of the simulations was compared by measuring their SNR using Monte-Carlo simulations. 100 different noise instances were simulated for the oversampled single coil case, and 20 for the undersampled multi-coil data.

Maps of noise amplification due to acquisition operator non-orthogonality (similar to g-factor [1]) were produced for SILVER, GR and uniform sampling trajectories. To create noise amplification maps, the linear acquisition operator, $\mathbf{E} = \mathbf{FS}$, with $\mathbf{S}$ being a set of eight coil sensitivities (compressed from a 32-channel sensitivity map measured in a phantom) and $\mathbf{F}$ being the discrete Fourier transform operator for samples along the trajectory, was calculated. The noise amplification maps were then generated by analytic evaluation of the acquisition operator in the same way as g-factors are calculated for SENSE as described by Pruessmann et al., which is feasible for small matrix sizes with a small number of coils.

$$g_\rho = ((\mathbf{E}^{\mathrm{H}}\mathbf{E})^{-1})_{\rho,\rho} \tag{6.14}$$

with ρ signifying the pixel under consideration.

Noise amplification maps were calculated for five different window sizes, $N = 32, 48, 64, 128$. For each window size the map mean intensity was evaluated for a GR trajectory, a uniformly sampled radial trajectory, and for SILVER optimized for seven different SILVER sets: $S = \{N, 2N\}, \{N, 2N, 3N\}, \{N-1, \ldots, N+1\}, \{N-2, \ldots, N+2\}, \{N-3, \ldots, N+3\}, \{N-4, \ldots, N+4\}, \{N-5, \ldots, N+5\}$. For example for $N = 32$, SILVER was optimised for $S_{32_1} = \{32, 64\}$, $S_{32_2} = \{32, 64, 96\}$, $S_{32_3} = \{31, \ldots, 33\}$, $S_{32_4} = \{30, \ldots, 34\}$, $S_{32_5} = \{29, \ldots, 35\}$, $S_{32_6} = \{28, \ldots, 36\}$, $S_{32_7} = \{27, \ldots, 37\}$.

6.3.3 In vivo experiments

Finally, in vivo dynamic non-VE-ASL angiography datasets were acquired from three healthy volunteers using a 3 Tesla Verio scanner (Siemens Healthineers, Erlangen, Germany). The matrix size was 192 × 192. Data was acquired with 5 different protocols:

1. Uniform sampling, 68 spokes per frame, 27 frames, $\alpha = 1/68$ (acceleration factor, $R \approx 4$)
2. Uniform sampling, 153 spokes per frame, 12 frames, $\alpha = 1/153$ ($R \approx 2$)
3. Uniform sampling, 306 spokes per frame, 6 frames, $\alpha = 1/306$ ($R \approx 1$)
4. GR sampling, $\alpha = 1/\phi \approx 0.6180\ldots$
5. SILVER optimized for $S = \{68, 153, 306\}$, $\alpha \approx 0.2770\ldots$

For all protocols, a 600ms labelling phase was followed by a 1288 ms continuous GRE Look-Locker readout (TE = 5.95 ms, TR = 11.7 ms, FA = 7°) where 108 spokes were acquired. A total of 1836 spokes (17 shots) were acquired for the tag and control conditions respectively. The spokes were ordered such that combining the 17 shots resulted in the expected set increment trajectory [124], similar to the ordering described in chapter 4, figure 4.15(B), but with the same spokes for both the tag and control condition (figure 6.4). This method of combining spokes does not allow for retrospective undersampling by discarding repeats, as has been done in previous chapters, but instead allows for flexible combination of subsequent frames such as the total trajectory in each frame remains the same as if the spokes had been acquired sequentially. Total scan time for each protocol was 1 min 8 s.

All in vivo data was first phase corrected and then the tag and control acquisitions were pre-subtracted for increased reconstruction speed (less data to be held in memory). The dynamic 2D images were reconstructed with both linear SENSE and a non-linear CS reconstruction using a sparsity prior, but no temporal regularisation because no optimisation had been performed temporally (i.e. the radially uniform

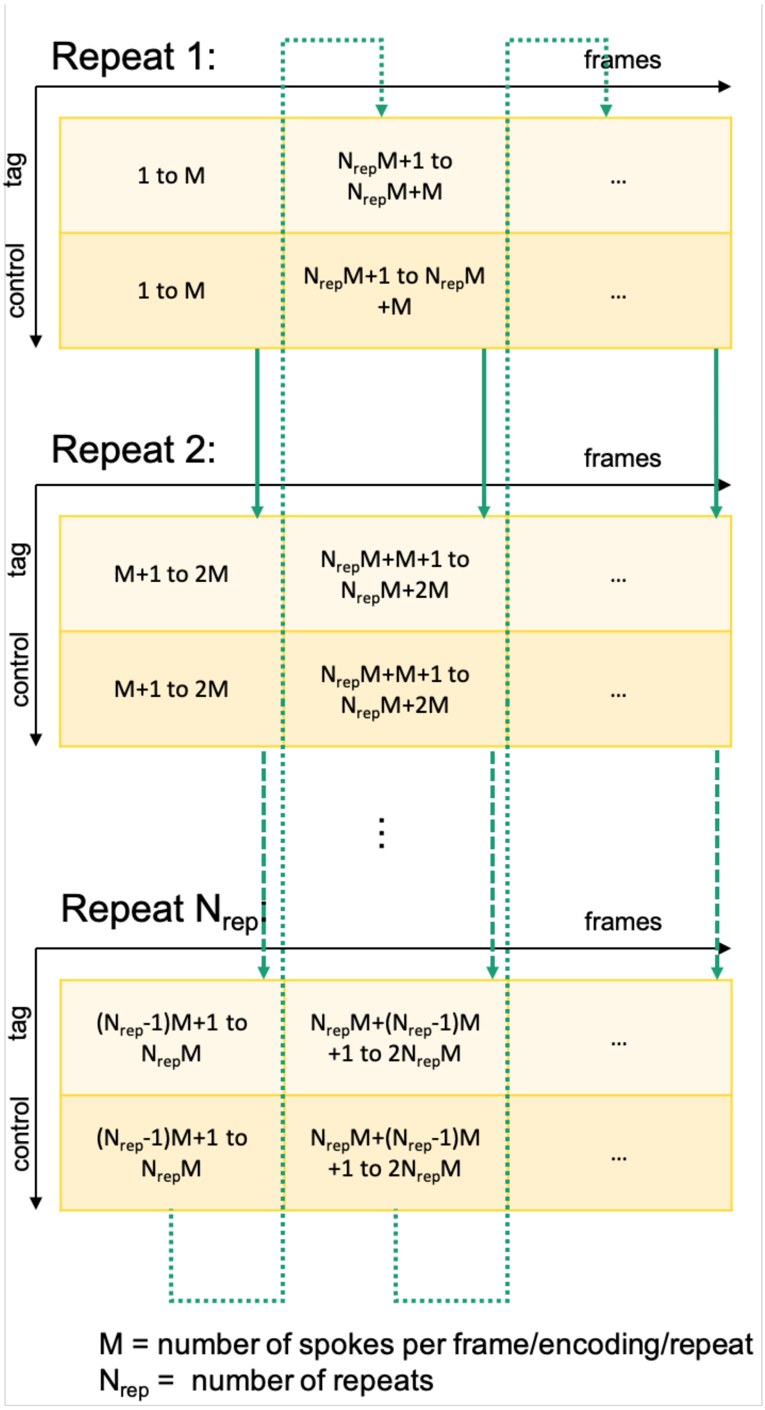

Figure 6.4: Spoke ordering used in SILVER in vivo experiments.

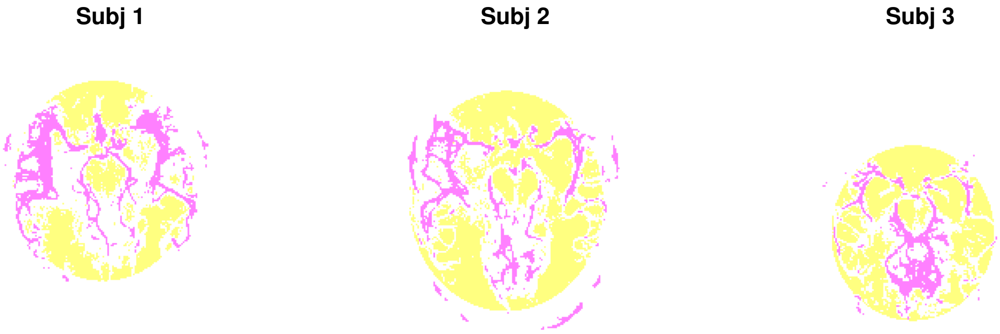

Figure 6.5: The masks used in vivo for SNR estimation. Signal was estimated as mean intensity in the region shown in pink. Noise was defined as standard deviation of signal in the region shown in yellow.

sampling pattern was the same for each frame). The sparsity term in the CS reconstruction had a weighting factor of 10^{-5}. The iterative reconstructions were run until convergence ($< 1\%$ change in cost from previous iteration). The linearly reconstructed in vivo reconstructions were also compared using SNR, whereas the CS reconstructions were only studied qualitatively due to the lack of ground truth and bias introduced by non-linear reconstruction, which makes SNR an unsuitable metric for image quality assessment. SNR was measured by (i) applying a mask to the vessels and averaging for signal, and (ii) estimating noise based on the standard deviation in a mask covering non-vessel regions within the brain. The masks are shown in figure 6.5. They were generated by thresholding the fully sampled (306 spokes/frame) reconstructions with radially uniform sampling at one level for the vessels and the inverse of thresholding at a lower level for the noise (to remove partial vessel signal). An ellipse was inscribed within the brain to remove areas of the noise mask such as the eyes where there can be significant physiological noise from movement.

6.4 Results

Two examples of SILVER and GR theoretical efficiencies for specific window sizes are shown in figure 6.6. The minimum efficiency of the SILVER trajectory within each targeted set was higher or equivalent to the trajectory produced by the GR method.

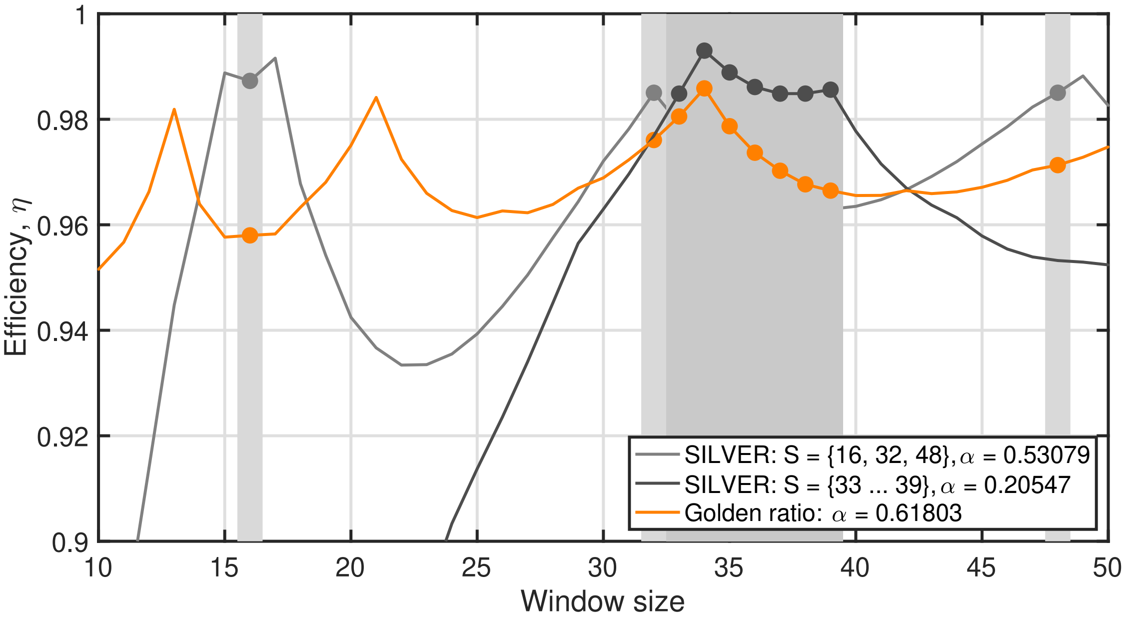

Figure 6.6: Two examples of SILVER efficiencies (light and dark grey) compared with the golden ratio (orange) in a range of window sizes. The golden ratio has typical peaks at the Fibonacci numbers, SILVER has much more irregular peaks. The minimum efficiency for SILVER within the optimised range (shaded, round markers) is higher than for the golden ratio.

SILVER minimum efficiency was higher than the GR method for both continuous sets (figure 6.7a)), and multiple temporal resolutions (figure 6.7b)), although the improvement was negligible for large window ranges and for the set containing five Fibonacci numbers of spokes, as expected. The maximum observed improvement over GR was 4.7% (for $S = \{4, 5\}$). For a minimum window size, M, of 16 spokes the maximum efficiency increase was 3.8% (for $S = \{16, 17\}$), and remained above 1% for continuous sets with up to 10 members ($S = \{16, \ldots, 25\}$). Similarly, for a minimum window size, M, of 32, the maximum improvement was 2.2% (for $S = \{32, 33\}$), and remained above 1% for continuous sets with up to 14 members ($S = \{32, \ldots, 45\}$). All multiple temporal resolution experiments except the set of Fibonacci numbers had improvements of more than 1.8%, with the maximum improvement being 4.2% (for $S = \{4, 8\}$).

The simulations showed SNR differences that were accurately predicted by the electrostatic efficiency metric for the oversampled reconstructions, and although the scaling of the SNR for GR, SILVER and uniform sampling did not closely follow

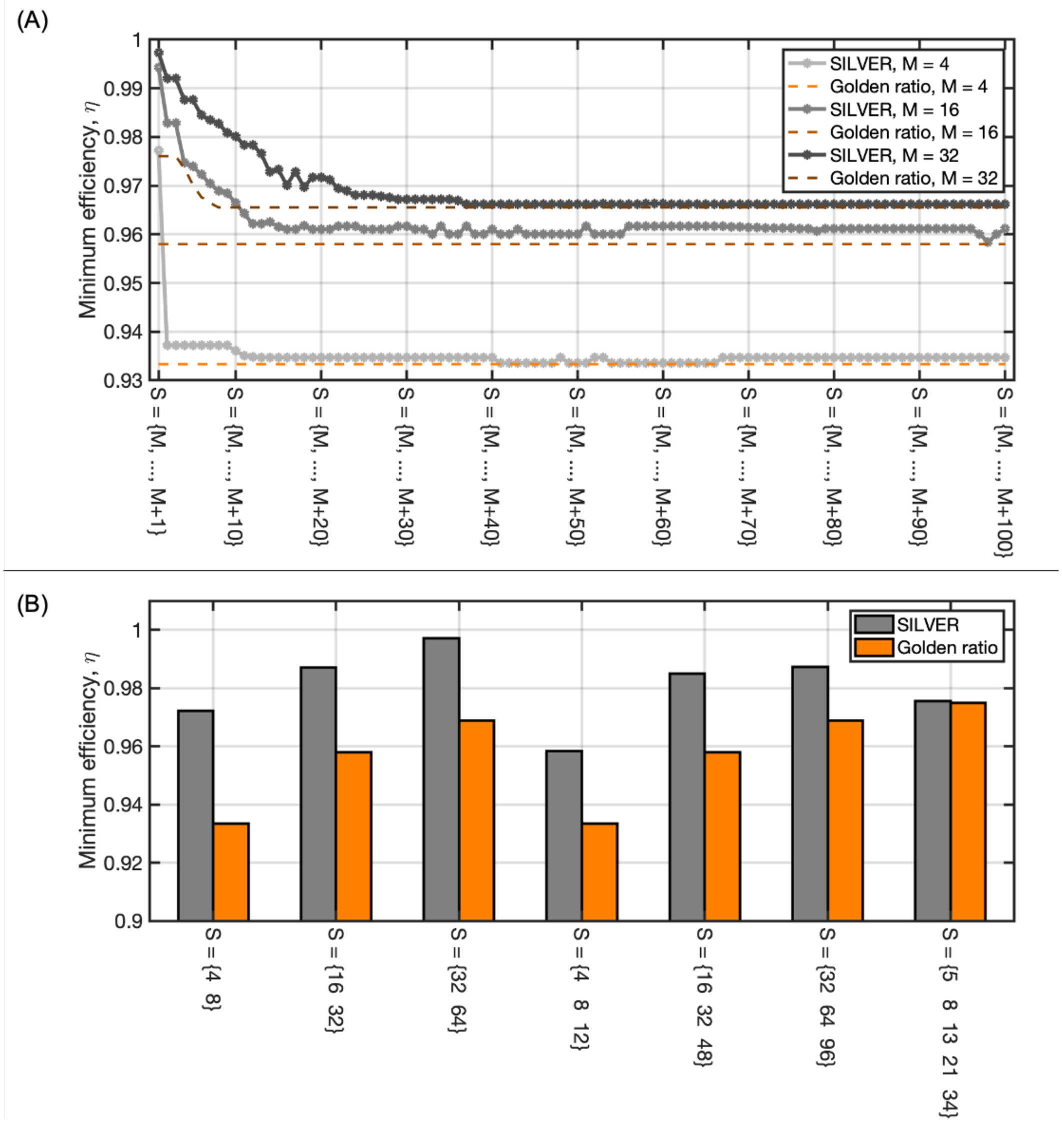

Figure 6.7: For sets of continuous ranges of window sizes SILVER can outperform the golden ratio if the range is short enough. How short the range has to be depends on the minimum number of spokes as shown in (A). For pairs and triples of window sizes SILVER performed much better than the golden ratio, but in the golden ratio optimal case (Fibonacci numbers) they were equally good as shown in (B). The y-axis in both graphs, η, represents the minimum efficiency within the SILVER targeted range. Note that efficiency measured with the electrostatic potential method is not asymptotic and rises slowly for more spokes per window (see figure 6.6) and thus the three different values of M in (A) give rise to different minima for the GR efficiency.

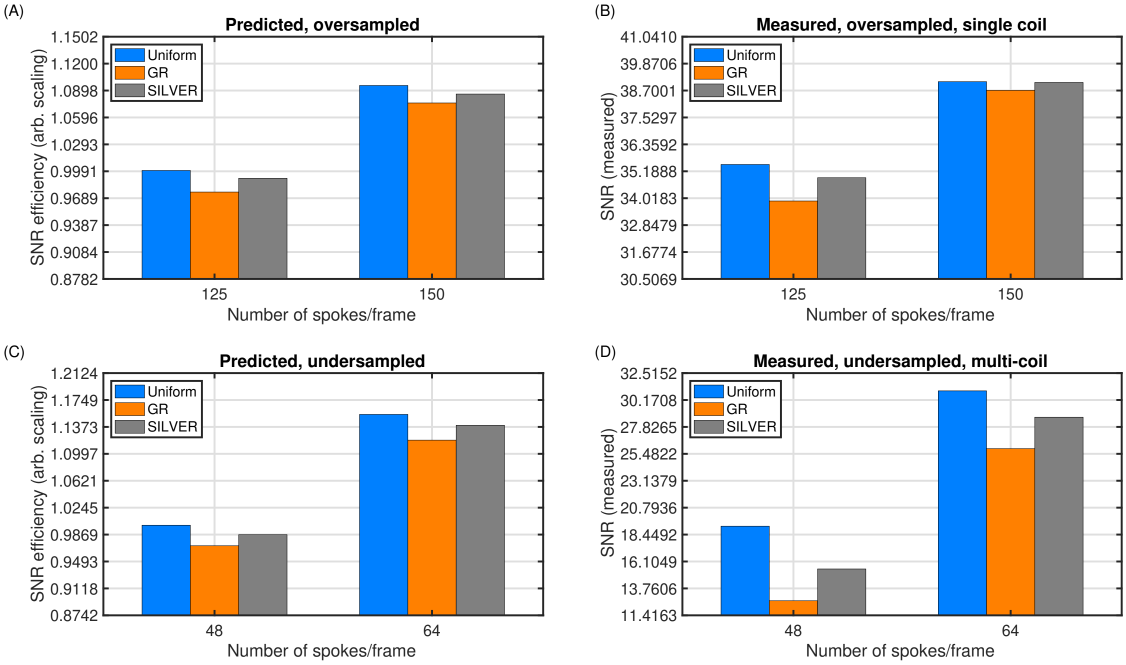

Figure 6.8: SNR predictions (A) and (C), and measurements in the dynamic phantom (B) and (D).

the prediction for undersampled multi-coil reconstruction, the trend $\text{SNR}_{\text{uniform}} > \text{SNR}_{\text{SILVER}} > \text{SNR}_{\text{GR}}$ was followed (figure 6.8).

The SNR was temporally stable for all simulations except the multi-coil GR reconstruction, where the direction of the sampling pattern clearly interacted with the spatial distribution of coil sensitivity maps. This can be seen both in figure 6.9(A) that shows SNR across frames, and figure 6.9(H) that shows the reconstructed image.

The improvement of SILVER over GR was reflected in the noise amplification maps as well. The noise amplification maps produced by the trialled SILVER trajectories had on average only 7% higher noise amplification than uniform sampling, but were on average only 55% of the GR noise amplification. Figure 6.10(A) shows the result for the seven SILVER trials compared to uniform and GR. Figure 6.10(B) shows an example of noise amplification maps (log-scaled) for a SILVER, GR, and uniformly radial trajectories for the four different windows.

In vivo SNR measured across the three subjects (figure 6.11) follow the expected pattern only for 68 spokes per frame, where the differences reached statistical significance. For the other window sizes SNR was more variable among the

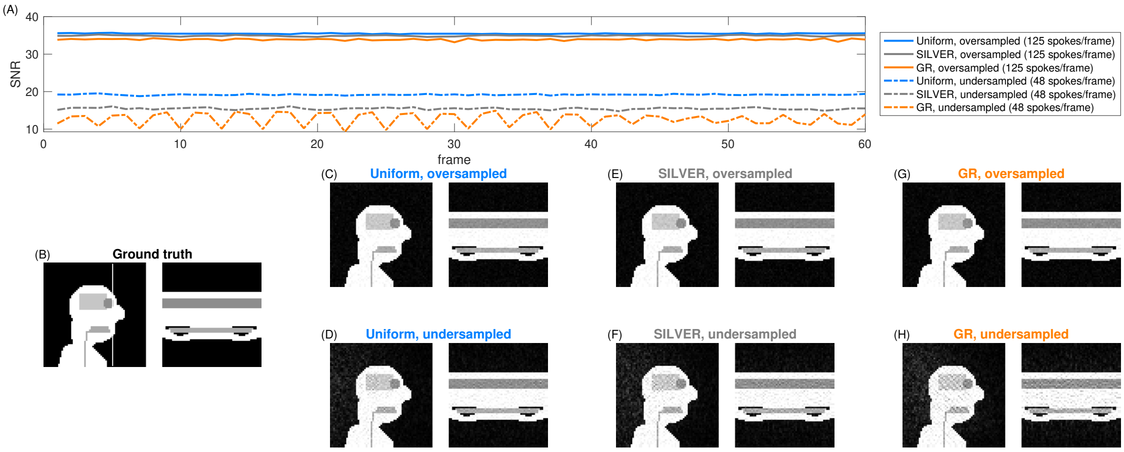

Figure 6.9: SNR measurements across time (A), and the last frame from the reconstruction along with a time series of the phantom experiments (C-H). The white line on the ground truth image (B) shows where the image was sliced through time (right part of every subfigure).

three subjects.

Figure 6.12 shows the time averaged angiograms of all subjects reconstructed with both SENSE and CS. The reconstruction with uniform sampling exhibits streaking when averaged across frames because each frame had the exact same trajectory and thus even small amounts of residual streaking artifacts add up. This is artifact is notably missing from both the SILVER and GR image whose trajectory rotates between frames and have no overlap in trajectory. At 153 and 306 spokes/frame some of the subjects exhibit significant artifacts, potentially due to motion. These artifacts might explain the lack of significant results at these window sizes.

6.5 Discussion

We have presented a method of choosing an optimal angular increment for dynamic radial MRI when the set of window sizes to consider is constrained. We have showed that this method results in small but measurable and potentially valuable increases in SNR, as well as large reductions in noise amplification, compared to the more general golden ratio method, with only a minor change required to the acquisition protocol.

The in vivo experiment showed inconclusive results, especially at 153 and 306 spokes per frame. At 68 spokes/frame, however, the reconstructions followed the

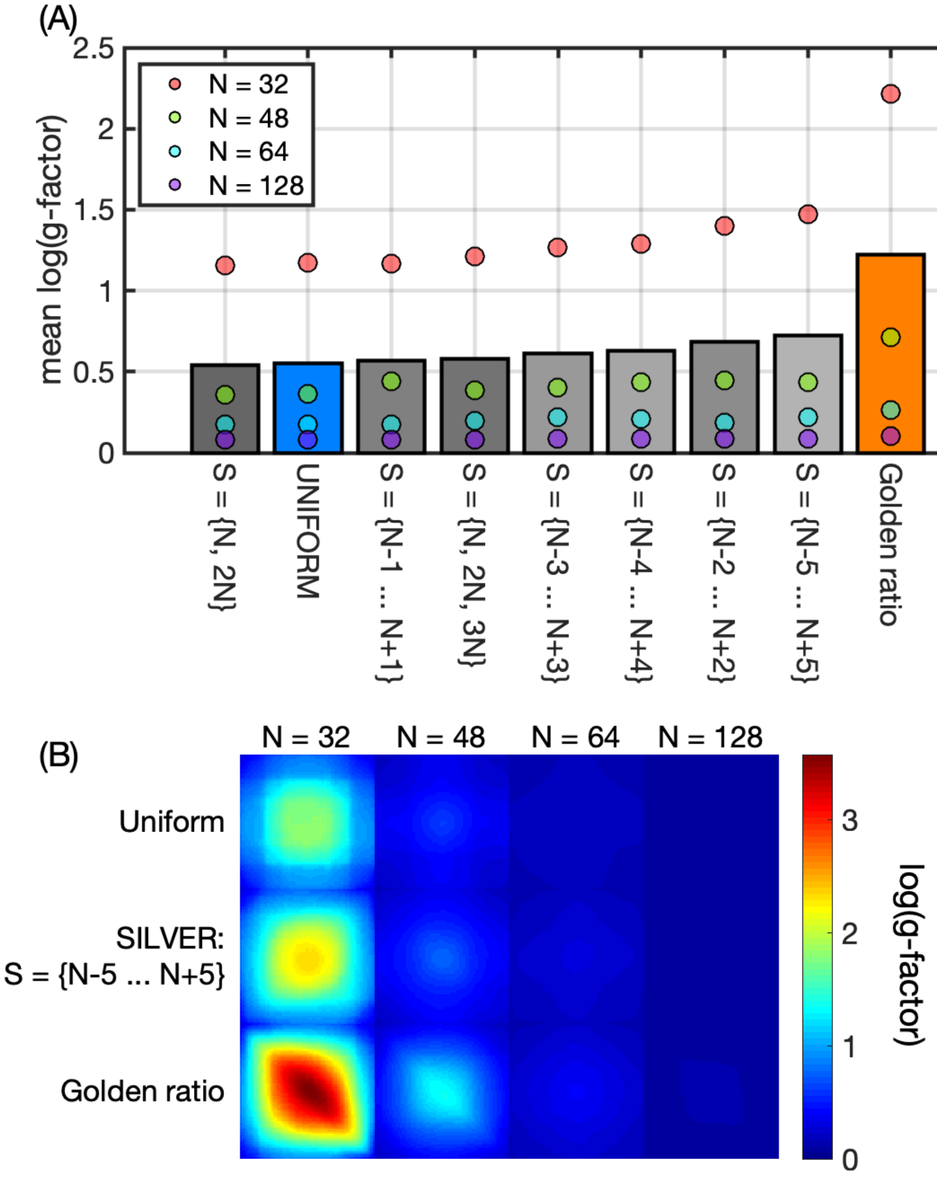

Figure 6.10: Noise amplification in multi-coil SILVER experiments, log-transformed for visibility, as otherwise the difference between numbers of spokes overshadows any differenced between the sampling methods. (A) shows the average result for different optimisation sets and numbers of spokes per window, N. (B) shows an example of the noise amplification maps for one specific SILVER optimisation range compared with radially uniform sampling and GR.

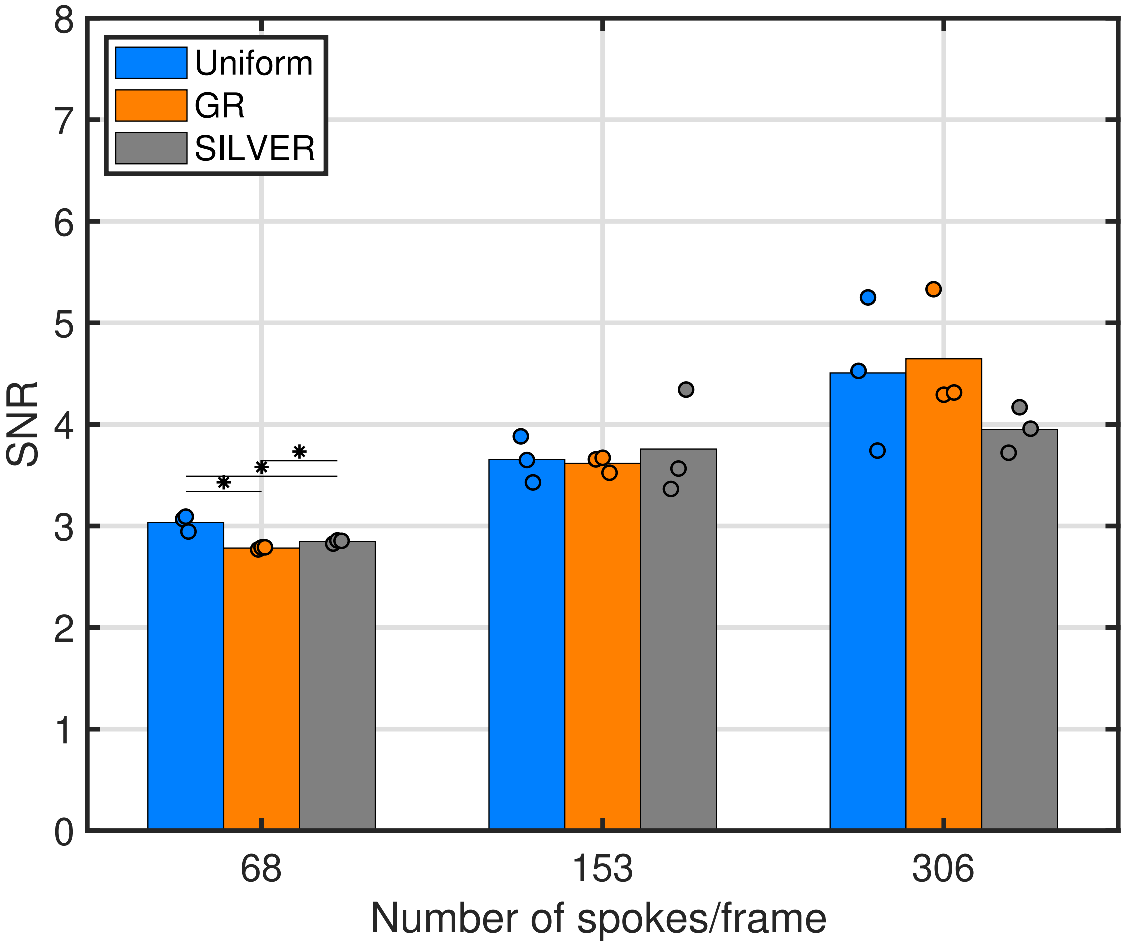

Figure 6.11: In vivo SNR measurements for the three subjects (each shown with a round marker) and their averages (bars).

expected pattern ($\text{SNR}_{\text{uniform}} > \text{SNR}_{\text{SILVER}} > \text{SNR}_{\text{GR}}$) and the differences reached statistical significance. Multiple factors could give rise to the inconclusive results for 153 and 306 spokes/frame. First, for the reconstructions with more spokes per frame, the images have fewer frames, and thus fewer pixels to measure SNR over, which can explain the larger variance than for the 68 spokes per frame case. Additionally, the less controlled real-life conditions might also contribute to the uncertainty in the results. For example the placement of the subject with regards to the coil sensitivity profiles, physiological noise, and subject motion that produced prominent artifacts in some of the data sets. Future experiments in a larger cohort might show the image improvements expected based on the theory and simulation results.

Because orthogonality of the acquisition operator depends on the coil sensitivities,

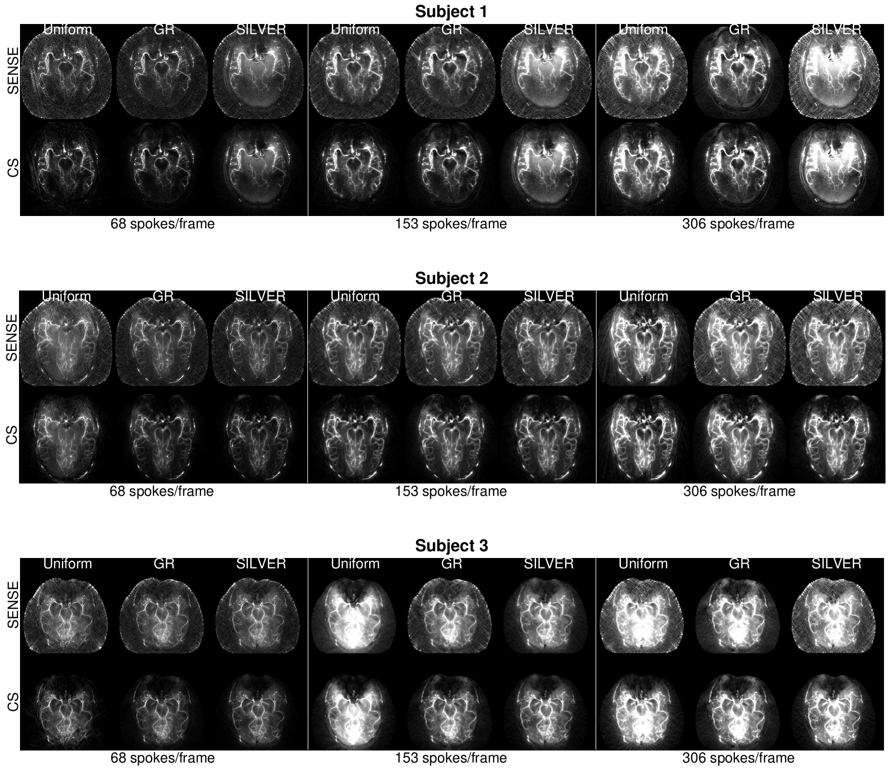

Figure 6.12: In vivo uniform, GR, and SILVER reconstructions.

as well as the sampling trajectory, optimising for sampling uniformity alone does not always produce the optimal trajectory. This combination of coils and trajectory could explain why, for $S = \{N, 2N\}$ the noise amplification of SILVER were unexpectedly marginally lower than for uniform spokes (figure 6.10). It can also explain the results seen in figure 6.9 - when coils are included in the reconstruction method the SNR for consecutive frames are no longer constant as the direction of the spokes combined with the spatial distribution of coil sensitivities interact closely. The golden angle method showed prominent drops in SNR approximately every three frames. This could be explained by that for 48 spokes/frame every three frames the sampling pattern is rotated a net $179.4° \approx 180°$.

However, without a-priori knowledge of coil sensitivity profiles, optimising for uniformity alone is a tractable step towards optimised imaging. In future work, the SILVER framework could be extended to optimise for properties other than trajectory uniformity through design of alternative cost functions, including cost functions that consider the incoherence of the operator for compressed sensing reconstructions [4] or include coil sensitivity maps acquired in a pre-scan for SENSE like reconstructions.

One thing to note about the current implementation of SILVER, is that not all windows in the optimised set of window sizes are guaranteed to have higher efficiency than GR, as the optimisation was set up as a max-min problem. Within the set, GR might have better performance than SILVER at certain windows, for example if a Fibonacci number is included in the set. One could imagine instead optimising for average performance within the set or uniformity of efficiency. Similarly, cost functions that optimise the trajectory across time could be considered, for example to have as little overlap as possible. The effect of different cost functions would be an interesting follow-on experiment on the work so far.

Both SILVER and the golden ratio method only optimise for subsequently sampled spokes, however, there is no guarantee that binned data from multiple repeats of radial acquisitions are distributed optimally. This problem was not directly addressed in this work, but suggestions for overcoming this problem have been proposed for golden ratio sampling with e.g. sector wise sampling [125]. With sector wise sampling the number of repeats to sample and retrospectively bin based on e.g. cardiac phase is pre-determined and rather than sampling over the whole k-space, a single sector of k-space is sampled each repeat such that golden ratio sampling is maintained within that section. When the sections are combined to a single k-space higher uniformity than standard binned GR is achieved. Because SILVER can achieve higher uniformity than GR when the number of spokes per phase is only approximately known, it ought to work even better or equally well as GR in sector wise sampled binned data.

To extend SILVER to 3D sampling, the definition of a set increment needs further consideration. Chan et al. [34] proposed the multi-dimensional golden means method, a 3D analogy of the golden ratio method with a constant azimuthal angle increment and a constant k_z-axis increment of the tip of the spoke constrained to the surface of a sphere. However, subsequent frames generated with this method are not simply a rotation of the spokes from the previous frame, and therefore not all image frames are guaranteed to have the same sampling efficiency. Because of the increased complexity of 3D radial sampling and the ambiguity of how to define a set increment, we constrained the scope of this chapter to focus on the more commonly used 2D case, although extension of SILVER to 3D trajectories will be considered in future work.

6.6 Conclusion

SILVER is a method that generalises set increment sampling beyond golden ratio derived methods when either a set number of temporal resolutions required, or an approximate range of numbers of spokes to reconstruct jointly is known a priori, as is commonly the case. It offers a trade-off between flexibility of golden ratio based methods and efficiency of uniformly spaced sampling when the requirement to be close to uniform for any window size is relaxed.

The experiments conducted on SILVER sampling with multiple coils reminds us again that just like in previous chapters, one must not forget to consider all dimensions of a problem, as optimising the sampling alone might not always produce the ideal sampling pattern for a specific set of coils sensitivity profiles. It is, however, a tractable step in the right direction. Especially considering that scan sequences that were previously set up with golden ratio, or other set increment sampling methods, can easily incorporate SILVER. The incorporation is simple because SILVER only changes a single parameter of those sequences.

SILVER applied to a large set of window sizes, whose uniformity cannot be improved upon, simply returns the golden ratio (or a tiny golden ratio). Therefore, one can conclude that sometimes SILVER is better than gold, and when it is not, it is just as good.

For me, becoming isn't about arriving somewhere or achieving a certain aim. I see it instead as forward motion, a means of evolving, a way to reach continuously toward a better self. The journey doesn't end.

— Michelle Obama in her Autobiography *Becoming*
2018

7

Summary and Future Work

Contents

7.1 Summary 152
  7.1.1 Technical Development 153
  7.1.2 Benefits and Limitations of Novel Methods 154
  7.1.3 Extensions and Generalisations 155
7.2 Directions of Future Work 156
  7.2.1 Multidimensional Analysis 156
  7.2.2 High Dimensional Reconstruction with Tensors 158
  7.2.3 Static Tissue Suppression Methods 158
  7.2.4 Combining Vessel and Time Encoding 158
  7.2.5 Clinical Protocol Development 159
  7.2.6 Faster Reconstruction 159
  7.2.7 Quantitative Analysis of Physiological Parameters . . . 160
  7.2.8 Generalising the SILVER Framework Further 160
7.3 Final Remarks 161

In this chapter, I will summarise and reflect on the work that has been presented in this thesis, and share my ideas for future work that could push this line of research further.

7.1 Summary

The main aims of this thesis, as presented in chapter 1, were to i) investigate the technical hurdles and opportunities of applying non-linear reconstruction methods and efficient, information sharing sampling methods to VE-ASL angiography, ii) explore the potential benefits and limitations of these types of reconstruction and sampling methods, and iii) extend and generalise the methods developed for VE-ASL angiography to other MRI methods. Now, have these aims been achieved?

7.1.1 Technical Development

We started off in chapter 3 by building a reconstruction framework based on current but already established methods of MRI reconstruction, a parallel imaging and compressed sensing approach. An additional temporal smoothness constraint was also added, and the method was applied to a data type that had not been accelerated using these types of methods before; vessel encoded ASL angiography. The regularisation parameters in the method were calibrated and the impact of the different terms in the cost function explored and the method was assessed both in simulations and in vivo. It was found that including the vessel-decoding into the reconstruction process made the reconstruction sparser, which allowed for higher acceleration of VE than non-VE data.

Since the results in chapter 3 were encouraging, extensions of the framework were considered, in order to optimise it and allow for even higher acceleration factors and translation into 4D imaging (chapter 4). In this chapter, modifications to both the acquisition and the reconstruction protocol were proposed, implemented, and assessed, with the aim of taking the relatively general approach to image reconstruction that was used in chapter 3 to a method more optimised for VE-ASL angiography in particular. However, of the proposed techniques (varying the trajectory across encodings, sparsifying the static tissue using wavelets or a data driven approach, adding additional terms to the cost function to regularise the image spatially or temporally), only varying the spokes across encodings using the hybrid method brought benefits to VE-ASL angiographic imaging beyond the simple method presented in chapter 3. This development, is, however, an interesting one, since it could be extended to other methods of MRI, where the property of interest is encoded into multiple acquisitions. In order to study this type of interaction between encoding and sampling, the m-PSF, was proposed as a tool for visualisation and analysis of the interaction.

Chapter 5 was about application of the techniques discussed previously. The reconstruction framework was optimised for use in two different scenarios: time-encoded angiography and CAPRIA, both of which had specific technical requirements. TEnc-ASL required the smoothness constraint to be altered to bridge across labelling blocks, and the perfusion part of CAPRIA required recalibration of the regularisation factors due to the very different data characteristics of perfusion images.

Throughout the thesis, the golden ratio radial sampling method had been extensively used and modified to optimally spread aliasing across encodings and the temporal dimension. In chapter 6, on the other hand, the optimality of this sampling method was challenged, and an alternative, SILVER, was developed and presented. SILVER is a simple method of choosing a fixed angular step in a radial sampling trajectory such that the uniformity of the trajectory is optimal for the reconstruction window sizes actually used. Trajectory development is often intricate and requires optimisation of multiple parameters, the simplicity of SILVER (only optimising for the step size) makes the optimisation simple to perform and the result simple to apply to existing protocols.

7.1.2 Benefits and Limitations of Novel Methods

The methods presented in this thesis were assessed in both simple simulations, in realistic simulations on real images, and on real scan data acquired in vivo.

The main benefit of these acceleration methods, is that imaging can be brought to clinically feasible scan times. In chapter 4 high resolution 4D angiograms acquired in only 5 minutes were presented. If the Nyquist limit would have had to be reached, the scan would have had to last over 8 hours, which is clearly impossible even for healthy subjects, let alone patients with neurovascular disorders. The main limitation is now at the other end of the imaging pipeline, reconstruction time.

In terms of applying non-linear reconstruction methods to VE-ASL angiography in particular, the data itself provides benefits to the reconstruction by increasing the relative sparsity compared to non-VE-ASL. This produced the net benefit of being able to achieve higher acceleration of VE-ASL to match scan time with non-VE-ASL, and thus making vessel-selectivity available for "free".

The static tissue component of a vessel-encoded angiogram was shown to be a limitation for utilising efficient sampling (different trajectories each encoding), however, the methods presented in this thesis showed that the static tissue reconstruction problem can be decoupled from the vessel components by using a hybrid approach of sampling and encoding. This method provided small benefits over using the same spokes for each encoding. Some promise of using static tissue sparsification methods in combination with varying the spokes each encoding was shown in simulations where the static tissue intensity was reduced. This shows the potential of further improvements if better static tissue suppression can be achieved.

The methods presented have shown that considering data, acquisition, and reconstruction jointly generally improves results. So, whilst being very focused on VE-ASL angiography in chapters 3 and 4, the theoretical frameworks in this thesis, such as considerations of the m-PSF or the sampling uniformity discussion related to SILVER in chapter 6, benefit from being very generalisable and can be readily applied to other types of MRI. In chapter 5 the methods developed for VE-ASL angiography were already successfully applied to other ASL methods, but there are also many possibilities for extension outside the field of ASL.

7.1.3 Extensions and Generalisations

The final aim of developing generalisable and extendible methods were, as mentioned above, shown in chapters 5 and 6.

In chapter 5 the simpler approach from chapter 3 was applied to time-encoded ASL angiography as well as combined perfusion and angiographic imaging.

The maths governing time encoding are the same as vessel encoding, so a direct translation of the method was straightforward, with the only addition being that the smoothness constraint had to be applied between decoded images as well as readout frames directly. Shorter sub-boli means that less of the arterial tree will be filled in each frame compared with sequential acquisition, but because of dispersion effects, this effect is likely smaller than the difference between VE and non-VE. Time encoding does, however, improve intrinsic SNR of the data compared to data acquired with sequential readout.

The big question when applying the framework that was developed for angiography to perfusion imaging was whether the sparsity constraint would hinder the reconstruction. A grid search of regularisation factors showed that a small L1-constraint can act as a denoiser, and be beneficial in perfusion imaging too. The benefits from the temporal smoothness constraint, as shown in the angiographic data earlier showed similar benefits in perfusion data. Overall, the framework was very applicable to this, slightly different, type of imaging, and produced high quality images even when the scan time was reduced to well below a minute per slice.

SILVER was the presented as a final generalisation of the acquisition method used in the previous chapters. It is clearly not an ASL specific method, as the trajectory improvement could easily benefit any dynamic MRI method that currently uses the golden ratio by increasing performance, or any method with uniform sampling by increasing flexibility. It was assessed in vivo for ASL angiography, but was overall presented and developed as a general method.

7.2 Directions of Future Work

Based on the work presented in this thesis, a number of future projects can be envisioned. I will list some of my ideas below, but I am sure that someone else could think of many more directions this research could move in.

7.2.1 Multidimensional Analysis

The m-PSF analysis presented in chapter 4 only scratched the surface of optimising the sampling and encoding jointly. Only slightly different versions of radial sampling were considered, along with only Hadamard-like encoding matrices. Thus, only methods that fully suppressed or mixed signals from the different components were studied. The fact that permuting which spokes were acquired in which encoding did not affect the diagonal blocks in the m-PSF, but did change the off-diagonal elements, raises questions regarding whether the reconstruction could be optimised by choosing the right spokes in the right encoding, and whether the m-PSF can be used to predict what spoke ordering will work best. Could analysis of the off-diagonal elements in the m-PSF predict reconstruction quality based on, for example, the energy distribution between components? Also, are there implications for noise amplification if each encoded image is acquired with non-ideal trajectories but the combination of trajectories (diagonal blocks) is good?

In future work on applying the m-PSF to ASL data, a more general joint optimisation of sampling and encoding could be studied, for example to see whether partial suppression of the static tissue signal is possible and desirable. Could a sampling and encoding scheme be designed such that the mixing between components can be controlled based on the energy of each component? What features of the m-PSF could be used to predict reconstruction quality?

Should the m-PSF framework be extended to consider higher dimensional data? As we saw in chapter 4, if the temporal dimension is neglected, the results may not follow the predictions based on the m-PSF that at the moment only considers spatial sampling.

There are many MRI methods beyond ASL that image a property that is inferred from multiple encoded images. Examples include phase contrast imaging [59], and different types of multiecho imaging for quantification of tissue parameters (e.g. T_1- or quantitative susceptibility mapping [126, 127]). Efficiency in acquiring these types of images could likely be improved by jointly decoding the property of interest and reconstructing the signal. Could these types of imaging methods also benefit from m-PSF analysis?

7.2.2 High Dimensional Reconstruction with Tensors

As we have seen throughout the thesis, it is important to consider the full dimensionality of the data in order to optimise its acquisition and remove redundant data sampling (e.g. acquiring data from the same k-space location for multiple encodings or frames). The development of the m-PSF is a step in that direction, but the reconstruction framework presented in this thesis could potentially benefit from a low-rank tensor formalism. Low-rank methods can be used to consider space-time correlations, or space-component correlations, or time-component correlations, which have not been considered here. These can be encapsulated within a [space × time × component] tensor model of the data. Some recent reconstruction methods have been proposed using explicit tensor structure of the data within the reconstruction framework rather than considering each dimension separately [128, 129].

7.2.3 Static Tissue Suppression Methods

In chapter 4 we saw potential benefits in using varying spokes every encoding if the static tissue intensity was significantly lower than in the real scan. Tissue sparsification methods also provided benefits in this regime. The question now is, could the acquisition be modified such that the background intensity is low enough to use these methods in vivo? Examples of improved background suppression methods include using multiple inversion pulses [90] or interleaving background suppression pulses with the labelling [130] or readout to maintain good BGS across a reasonably long readout period, rather than optimising for just the beginning or middle of the readout, as is done now.

7.2.4 Combining Vessel and Time Encoding

VE-ASL angiography and/or perfusion imaging could benefit from the SNR boost of a time-encoded acquisition. In this thesis the two modalities were assessed separately, but combining the two could provide mutual benefits. If both vessel and time encoding were fully encoded, the number of encoding preparations would be at least $(N_{vessels}) \times (N_{frames}) + 1$, which would quickly make the number required encodings, and thus scan time, rise.

For 2D imaging with some temporal resolution attained from the readout (as in chapter 5), this could still be feasible, especially for angiography that can tolerate higher undersampling. Another alternative is to just vessel encode one of the time encoded blocks as was proposed by Okell et al. previously [106]. This has the drawback of not providing temporal information for the individual vessels/perfusion territories.

Again, the combination of encoding the two properties of interest (originating vessel and PLD) and trajectory could be explored using the m-PSF framework to minimise acquisition of redundant information and control aliasing across the encodings. Now, aliasing across time frames would also be possible to consider, and optimal regularisation to remove them could be developed further.

7.2.5 Clinical Protocol Development

A fast, vessel-encoded, and dynamic combined angiography and perfusion scan in 3D would provide a very large amount of clinically relevant information. An interesting project to pursue, would be to combine all the pieces presented in this thesis and put them together into a clinical protocol in collaboration with clinicians to assess the added value of such a scan in patients with, for example, arteriovenous malformations or carotid stenosis. The acceleration methods presented here could form the basis for accelerating the acquisition, but much work on reconstruction speed would be required.

7.2.6 Faster Reconstruction

One large limitation with the current method, especially for 4D imaging, is the large memory footprint and long reconstruction times.

The memory requirements could likely be reduced by optimisation of the code, store data with single precision, avoid creating temporary copies within the methods etc.

The long reconstruction times are mainly driven by the large number of iterations required by the iterative methods, as well as the size of the data and images. Part of the method would likely be possible to parallelise and run on modern graphics processing units (GPUs). For example, the non-uniform fast Fourier transform, (NUFFT) has been implemented for GPU use [131]. Deep learning based approaches have also been shown to accelerate reconstructions by 'unrolling' iterative methods, and reconstruct images with similar or superior quality in fewer iterations [132].

7.2.7 Quantitative Analysis of Physiological Parameters

ASL perfusion imaging is often referred to as a quantitative perfusion measurement technique, from which absolute physiological measures, such as cerebral blood flow and arterial transit time can be derived. Similarly, physiological measures related to angiography like blood volume within arteries, transit times, and dispersion parameters can be studied. Okell et al. [66] showed that the signal seems to be both delayed and more dispersed distal to a stenosis in a patient with moyamoya disease.

The methods used for quantitative ASL analysis [64] normally assume unbiased Gaussian noise statistics, so images reconstructed with non-linear methods might not provide accurate results in these frameworks. As an alternative, the fitting to biophysical models could be built into the reconstruction itself. Rather than first reconstructing images that are then analysed, quantitative maps could be reconstructed directly by modifying the acquisition operator to model the transforms from physiological parameters of interest to the scanner's sensor domain. However, these models often include non-linear transforms, which could provide additional challenges to the reconstruction.

7.2.8 Generalising the SILVER Framework Further

SILVER, in its current form assumes radial 2D sampling and it assumes that optimality can be calculated using an electrostatic potential model. There is no reason to limit the SILVER method to these constraints. 3D imaging could be optimised, other cost functions could be trialled and adapted for other trajectories used in dynamic imaging.

Some cost functions that could be trialled include other methods of measuring uniformity, e.g. Pipe-Menon weights [117], or not optimising for the minimum uniformity in the set of window sizes, and instead optimising for the average within the set.

Another direction for extension of the SILVER framework for optimising for set windows could be extended to imaging with other trajectories previously derived from the golden ratio, for example: multi-shot spirals [133, 134] or cones [135], where SILVER could optimise the rotation of the interleaves. Or even Cartesian sampling [136], if uniformity is beneficial (PI reconstructions) but the temporal resolution is partially unknown. For these types of trajectories the cost function for measuring uniformity (or any other property) would have to be redesigned accordingly.

7.3 Final Remarks

Research is often a painfully slow process, and often the steps into the unknown are small and seemingly insignificant at the time they are taken. Parallel imaging and compressed sensing were big leaps towards more efficient MRI acquisition. This thesis, however, presents many little steps that together can chip off on the long road towards fast and efficient MRI.

I started off by applying PI and CS, the "big leap" methods, to a specific problem; advanced ASL methods, and vessel-encoded ASL in particular. Then I started pushing forward myself and made forays into multiple different possible directions. Steps towards sampling efficiency were taken in developing a new way to vary the spokes across encodings to benefit from mutual information in the images (the hybrid method), and a new way to choose the spoke directions more efficiently by trading flexibility for uniformity (SILVER). Along the way, many more smaller steps have been taken, from optimising regularisation factors, to trying things that didn’t work (data driven temporal and spatial priors etc.).

The main lessons I have learnt from this work, and I hope that you too can appreciate, having made it to the end, is that in order to optimise, one needs to consider the full dimensionality of the data, and that optimising one little bit at a time often yields poorer results than considering the full problem. Start with what you want to know, e.g. how blood from different feeding arteries fills the cerebrovasculature, and work your way back to optimise the reconstruction and acquisition jointly for your specific problem.

Appendices

A

Toeplitz Embedding

In ISTA and FISTA, each iteration the multiplication $\mathbf{E}^{\mathrm{H}}(\mathbf{Em}_n - \mathbf{s})$ has to be performed. This can be expanded to $\mathbf{E}^{\mathrm{H}}\mathbf{Em}_n - \mathbf{E}^{\mathrm{H}}\mathbf{s}$. That is three NUFFT's every iteration. However, $\mathbf{E}^{\mathrm{H}}\mathbf{s}$ is a constant and can be pre-calculated. The other two $\mathbf{E}$'s can be partially replaced by fast Fourier transforms, FFT's, via the process of Toeplitz embedding. It works because $\mathbf{E}$ is built up of three parts, the vessel-encoding ($\mathbf{H}$), the coil sensitivity encoding ($\mathbf{S}$) and the NUFFT ($\mathbf{F}$):

$$\begin{aligned}
&\mathbf{E} = \mathbf{FSH} \\
&\mathbf{E}^{\mathrm{H}} = (\mathbf{FSH})^{\mathrm{H}} = \mathbf{H}^{\mathrm{H}}\mathbf{S}^{\mathrm{H}}\mathbf{F}^{\mathrm{H}} \\
&\mathbf{E}^{\mathrm{H}}\mathbf{E} = \mathbf{H}^{\mathrm{H}}\mathbf{S}^{\mathrm{H}}\mathbf{F}^{\mathrm{H}}\mathbf{FSH}
\end{aligned} \tag{A.1}$$

$\mathbf{F}^{\mathrm{H}}\mathbf{F}$ is the non-uniform transform from image space to k-space and back. It has a block Toeplitz structure (constant diagonals) which can be padded into a block circulant matrix, $\mathbf{B}$. Circulant matrices can be diagonalised by a discrete Fourier transform matrix, $\mathbf{U}$, which can be easily implemented using the FFT and its inverse, the iFFT. The eigenvalues are easily calculated as a circulant matrix is fully defined by its first column $\mathbf{b}$. The entries of the diagonal matrix can be pre computed ($\mathbf{Ub} = \mathbf{z}$) and applied as a point-wise multiplication between the FFT and iFFT.

$$\begin{aligned}\mathbf{B} &= \mathbf{U}^{-1}\mathbf{D}\mathbf{U} \\ &= \mathbf{U}^{-1}diag(\mathbf{U}\mathbf{b})\mathbf{U} \qquad (A.2) \\ &= \mathbf{U}^{-1}diag(\mathbf{z})\mathbf{U}\end{aligned}$$

In conclusion, $\mathbf{F}^{\mathrm{H}}\mathbf{F}$ can be replaced by $\mathbf{U}^{-1}diag(\mathbf{z})\mathbf{U}$ by adding an embedding step before the first iteration that calculates the first column of $\mathbf{B}$, and each iteration can be sped up by using FFT's instead of NUFFT's.

B

SNR Optimal Weighting Factors

As laid out in chapter 6, in a re-gridding type reconstruction the image SNR depends on the weighting factors, w_i, that are added to the N sampling points that all have an intrinsic uncorrelated variance, σ^2:

$$\begin{aligned} SNR_{im} &\propto 1/\sigma_{im} \\ &\propto 1/\sqrt{\sum_i^N w_i^2 \sigma^2} \\ &\propto 1/\sqrt{\sum_i^N w_i^2} \end{aligned} \tag{B.1}$$

Furthermore, we assume that the weighting factors add to one, such that there is no net amplification of the signal:

$$\sum_i^N w_i = 1 \tag{B.2}$$

We shall now show that optimal SNR is achieved when $w_i = 1/N$ for all $i \in (1, 2, 3, ..., N)$.

The Cauchy-Schwartz inequality states that:

$$\left|\sum_{k=1}^n u_k v_k\right|^2 \leq \sum_{j=1}^n |u_j|^2 \sum_{k=1}^n |v_k|^2 \tag{B.3}$$

By setting u_j to w_i, and v_i to $1/\sqrt{N}$ we can rewrite the previous equation to:

$$\left(\sum_i^N \frac{w_i}{\sqrt{N}}\right)^2 \leq \sum_i^N w_i^2 \sum_i^N \frac{1}{N} \tag{B.4}$$

Expansion and combination with equation B.2 gives us:

$$\begin{aligned} \left(\sum_i^N \frac{w_i}{\sqrt{N}}\right)^2 &\leq \sum_i^N w_i^2 \\ \left(\frac{1}{\sqrt{N}} \sum_i^N w_i\right)^2 &\leq \sum_i^N w_i^2 \\ \frac{1}{N} &\leq \sum_i^N w_i^2 \end{aligned} \tag{B.5}$$

which, when inserted into the equation B.1 gives:

$$SNR_{im} \propto \frac{1}{\sqrt{\sum_i^N w_i^2}} \leq \sqrt{N} \tag{B.6}$$

Now, let's explore when the equality holds. We have two simultaneous equations:

$$\sum_i^N w_i^2 = \frac{1}{N} \tag{B.7}$$

and

$$\sum_i^N w_i = 1 \tag{B.8}$$

These can be combined to:

$$\sum_i^N w_i^2 = \frac{\sum_i^N w_i}{N} \tag{B.9}$$

which can be rearranged to:

$$\sum_i^N N w_i^2 - w_i = 0 \tag{B.10}$$

Since we know that $0 < w_i \leq 1$, we know that $w_i^2 \leq w_i$. Because of equation B.7, we also know that $0 < w_i^2 \leq \frac{1}{N}$. So, $N w_i^2 - w_i \leq 0$. Which means that in order for equation B.10 to hold, every term must equal zero:

$$\begin{aligned} N w_i^2 - w_i &= 0 \\ w_i(N w_i - 1) &= 0 \end{aligned} \tag{B.11}$$

which has the trivial solution $w_i = 0$ and the expected solution $w_i = \frac{1}{N}$.

Do you not know that a man is not dead while his name is still spoken?

— Terry Pratchett in *Going Postal* 2004

References

[1] K. P. Pruessmann et al. “SENSE: Sensitivity Encoding for Fast MRI”. In: *Magn. Reson. Med.* 42.5 (1999), pp. 952–962. DOI: `10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S`.

[2] M. A. Griswold et al. “Generalized autocalibrating partially parallel acquisitions (GRAPPA)”. In: *Magn. Reson. Med.* 47.6 (2002), pp. 1202–1210. DOI: `10.1002/mrm.10171`.

[3] D. L. Donoho. “Compressed sensing”. In: *IEEE T. Inform. Theory.* 52.4 (2006), pp. 1289–1306. DOI: `10.1109/TIT.2006.871582`.

[4] M. Lustig, D. Donoho, and J. M. Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magn. Reson. Med.* 58.6 (2007), pp. 1182–1195. DOI: `10.1002/mrm.21391`.

[5] E. C. Wong. “Vessel-encoded arterial spin-labeling using pseudocontinuous tagging”. In: *Magn. Reson. Med.* 58.6 (2007), pp. 1086–1091. DOI: `10.1002/mrm.21293`.

[6] P. M. Robson et al. “Time-resolved Vessel-selective Digital Subtraction MR Angiography of the Cerebral Vasculature with Arterial Spin Labeling”. In: *Radiology* 257.2 (2010), pp. 507–515. DOI: `10.1148/radiol.10092333`.

[7] U. Jensen-Kondering et al. “Superselective pseudo-continuous arterial spin labeling angiography”. In: *Eur. J. Radiol.* 84.9 (2015), pp. 1758–1767. DOI: `10.1016/j.ejrad.2015.05.034`.

[8] Z. Zhou et al. “Accelerated noncontrast-enhanced 4-dimensional intracranial MR angiography using golden-angle stack-of-stars trajectory and compressed sensing with magnitude subtraction”. In: *Magn. Reson. Med.* 79.2 (2018), pp. 867–878. DOI: `10.1002/mrm.26747`.

[9] Z. Lin et al. “Clinical feasibility study of 3D intracranial magnetic resonance angiography using compressed sensing”. In: *J. Magn. Reson. Im.* 50.6 (2019), pp. 1843–1851. DOI: `10.1002/jmri.26752`.

[10] S. S. Schauman, M. Chiew, and T. W. Okell. “Highly accelerated vessel-selective arterial spin labeling angiography using sparsity and smoothness constraints”. In: *Magn. Reson. Med.* 83.3 (2020), pp. 892–905. DOI: `10.1002/mrm.27979`.

[11] S. S. Schauman, M. Chiew, and T. W. Okell. “Highly Accelerated Dynamic 2D and 3D Vessel-Encoded Arterial Spin Labelling Angiography”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 27. 2019, p. 744.

[12] S. S. Schauman, T. W. Okell, and M. Chiew. “Radial sampling interactions in multi-dimensional sparse encoding problems using a joint decoding-reconstruction framework”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2020, p. 3683.

[13] S. S. Schauman, T. W. Okell, and M. Chiew. “High resolution 4D vessel selective angiography in under 5 minutes using a constrained reconstruction”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2020, p. 1073.

[14] S. S. Schauman et al. “Highly Accelerated Time-Encoded Dynamic ASL Angiography”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2020, p. 1620.

[15] P. C. Lauterbur. “Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance”. In: *Nature* 242.5394 (1973), pp. 190–191. DOI: `10.1038/242190a0`.

[16] M. A. Bernstein, K. F. King, and Z. J. Zhou. *Handbook of MRI pulse sequences.* Amsterdam ; Boston: Academic Press, 2004.

[17] D. G. Nishimura. *Principles of Magnetic Resonance Imaging.* 1.2. Stanford University: Self Published, 2010.

[18] D. I. Hoult and B Bhakar. “NMR signal reception: Virtual photons and coherent spontaneous emission”. In: *Concept. Magn. Reson. A.* 9.5 (1997), pp. 277–297.

[19] U. Katscher et al. “Transmit SENSE: Transmit SENSE”. In: *Magn. Reson. Med.* 49.1 (2003), pp. 144–150. DOI: `10.1002/mrm.10353`.

[20] J. Pauly, D. Nishimura, and A. Macovski. “A k-space analysis of small-tip-angle excitation”. In: *J. Magn. Reson.* 81.1 (1989), pp. 43–56. DOI: `10.1016/0022-2364(89)90265-5`.

[21] S. Ljunggren. “A simple graphical representation of fourier-based imaging methods”. In: *J. Magn. Reson.* 54.2 (1983), pp. 338–343. DOI: `10.1016/0022-2364(83)90060-4`.

[22] D. B. Twieg. “The *k* -trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods: *k* -trajectory formulation of NMR imaging”. In: *Med. Phys.* 10.5 (1983), pp. 610–621. DOI: `10.1118/1.595331`.

[23] Y. E. Esin and F. N. Alpaslan. “MRI image enhancement using Biot–Savart law at 3 tesla”. In: *Turk. J. Electr. Eng. Co.* 25.4 (2017), pp. 3381–3396.

[24] J. W. Cooley and J. W. Tukey. “An algorithm for the machine calculation of complex Fourier series”. In: *Math. Comp.* 19.90 (1965), pp. 297–297. DOI: `10.1090/S0025-5718-1965-0178586-1`.

[25] D. Gallichan et al. “Simultaneously driven linear and nonlinear spatial encoding fields in MRI”. In: *Magn. Reson. Med.* 65.3 (2011), pp. 702–714. DOI: `10.1002/mrm.22672`.

[26] D. A. Feinberg et al. “Halving MR imaging time by conjugation: demonstration at 3.5 kG.” In: *Radiology* 161.2 (1986), pp. 527–531. DOI: `10.1148/radiology.161.2.3763926`.

[27] M. Uecker et al. “ESPIRiT-an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA”. In: *Magn. Reson. Med.* 71.3 (2014), pp. 990–1001. DOI: `10.1002/mrm.24751`.

[28] K. P. Pruessmann et al. “Advances in sensitivity encoding with arbitrary *k* -space trajectories: SENSE With Arbitrary *k* -Space Trajectories”. In: *Magn. Reson. Med.* 46.4 (2001), pp. 638–651. DOI: `10.1002/mrm.1241`.

[29] A. Beck and M. Teboulle. “A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems”. In: *SIAM J. Imaging Sci.* 2.1 (2009), pp. 183–202. DOI: `10.1137/080716542`.

[30] D. G. Nishimura, J. I. Jackson, and J. M. Pauly. “On the nature and reduction of the displacement artifact in flow images”. In: *Magn. Reson. Med.* 22.2 (1991), pp. 481–492. DOI: `10.1002/mrm.1910220255`.

[31] M. Lustig et al. “Compressed Sensing MRI”. In: *IEEE Signal Proc. Mag.* 25.2 (2008), pp. 72–82. DOI: `10.1109/MSP.2007.914728`.

[32] D. Piccini et al. “Spiral phyllotaxis: The natural way to construct a 3D radial trajectory in MRI: Spiral Phyllotaxis Radial 3D Trajectory”. In: *Magn. Reson. Med.* 66.4 (2011), pp. 1049–1056. DOI: `10.1002/mrm.22898`.

[33] C. A. Corum, S. J. Kruger, and V. A. Magnotta. “HEALPix View-order for 3D Radial Self-Navigated Motion-Corrected ZTE MRI”. In: *arXiv:1910.10276 [astro-ph, physics:physics]* (2019). arXiv: 1910.10276. DOI: `arXiv:1910.10276`.

[34] R. W. Chan et al. “Temporal stability of adaptive 3D radial MRI using multidimensional golden means”. In: *Magn. Reson. Med.* 61.2 (2009), pp. 354–363. DOI: `10.1002/mrm.21837`.

[35] S. Winkelmann et al. “An Optimal Radial Profile Order Based on the Golden Ratio for Time-Resolved MRI”. In: *IEEE Trans. Med. Imaging* 26.1 (2007), pp. 68–76. DOI: `10.1109/TMI.2006.885337`.

[36] P. G. Anderson. “Multidimensional Golden Means”. In: *Applications of Fibonacci Numbers*. Ed. by G. E. Bergum, A. N. Philippou, and A. F. Horadam. Dordrecht: Springer Netherlands, 1993, pp. 1–9. DOI: `10.1007/978-94-011-2058-6_1`.

[37] H. K. Song, L. Dougherty, and M. D. Schnall. “Simultaneous acquisition of multiple resolution images for dynamic contrast enhanced imaging of the breast”. In: *Magn. Reson. Med.* 46.3 (2001), pp. 503–509. DOI: `10.1002/mrm.1220`.

[38] T. Akasaka et al. “Optimization of Regularization Parameters in Compressed Sensing of Magnetic Resonance Angiography: Can Statistical Image Metrics Mimic Radiologists’ Perception?” In: *PLOS ONE* 11.1 (2016), e0146548. DOI: `10.1371/journal.pone.0146548`.

[39] K. Hammernik et al. “Learning a variational network for reconstruction of accelerated MRI data”. In: *Magn. Reson. Med.* 79.6 (2018), pp. 3055–3071. DOI: `10.1002/mrm.26977`.

[40] M. Oehler and T. M. Buzug. “Statistical Image Reconstruction for Inconsistent CT Projection Data”. In: *Methods Inf. Med.* 46.03 (2007), pp. 261–269. DOI: `10.1160/ME9041`.

[41] V. Yanamadala et al. “Non-contrast 3D time-of-flight magnetic resonance angiography for visualization of intracranial aneurysms in patients with absolute contraindications to CT or MRI contrast”. In: *J. Clin. Neurosci.* 20.8 (2013), pp. 1122–1126. DOI: `10.1016/j.jocn.2012.12.005`.

[42] D. R. Hadizadeh et al. “Cerebral Arteriovenous Malformation: Spetzler-Martin Classification at Subsecond-Temporal-Resolution Four-dimensional MR Angiography Compared with That at DSA”. In: *Radiology* 246.1 (2008), pp. 205–213. DOI: `10.1148/radiol.2453061684`.

[43] T. Adla and R. Adlova. "Multimodality Imaging of Carotid Stenosis". In: *Int J Angiol* 24.03 (2015), pp. 179–184. DOI: `10.1055/s-0035-1556056`.

[44] M. van Vliet et al. "MR Angiography of Tumor-related Vasculature: From the Clinic to the Micro-environment". In: *RadioGraphics* 25.suppl_1 (2005), S85–S97. DOI: `10.1148/rg.25si055512`.

[45] R. Takeda et al. "Is MRA at 3.0 Tesla sufficient for preoperative planning for aneurysmal clipping in patients with contraindicated condition of contrast media?" In: *Interdisciplinary Neurosurgery* 15 (2019), pp. 6–10. DOI: `10.1016/j.inat.2018.09.003`.

[46] J. H. Kim et al. "Dynamics of the cerebral blood flow response to brief neural activity in human visual cortex". In: *J Cereb Blood Flow Metab* 40.9 (2020), pp. 1823–1837. DOI: `10.1177/0271678X19869034`.

[47] T. van Seeters et al. "CT angiography and CT perfusion improve prediction of infarct volume in patients with anterior circulation stroke". In: *Neuroradiology* 58.4 (2016), pp. 327–337. DOI: `10.1007/s00234-015-1636-z`.

[48] M. Wintermark et al. "Comparative Overview of Brain Perfusion Imaging Techniques". In: *Stroke* 36.9 (2005), e83–e99. DOI: `10.1161/01.STR.0000177884.72657.8b`.

[49] B. Laviña. "Brain Vascular Imaging Techniques". In: *Int. J. Mol. Sci.* 18.1 (2016), p. 70. DOI: `10.3390/ijms18010070`.

[50] W. Hanafee and P. Stout. "Subtraction Technic". In: *Radiology* 79.4 (1962), pp. 658–661. DOI: `10.1148/79.4.658`.

[51] T. J. Kaufmann et al. "Complications of Diagnostic Cerebral Angiography: Evaluation of 19 826 Consecutive Patients ^{1}". In: *Radiology* 243.3 (2007), pp. 812–819. DOI: `10.1148/radiol.2433060536`.

[52] R. Klingebiel. "Comparative evaluation of 64-slice CT angiography and digital subtraction angiography in assessing the cervicocranial vasculature". In: *VHRM* Volume 4 (2008), pp. 901–907. DOI: `10.2147/VHRM.S2807`.

[53] M. Tekşam, B. Çakır, and M. Coşkun. "CT perfusion imaging in the early diagnosis of acute stroke". In: *Diagn. Interv. Radiol.* 11.4 (2005), pp. 202–205.

[54] H. Kim. "Variability in Quantitative DCE-MRI: Sources and Solutions". In: *J Nat Sci.* 4.1 (2018), e484.

[55] S. J. Riederer, E. G. Stinson, and P. T. Weavers. "Technical Aspects of Contrast-enhanced MR Angiography: Current Status and New Applications". In: *Magn. Reson. Med. Sci.* 17.1 (2018), pp. 3–12. DOI: `10.2463/mrms.rev.2017-0053`.

[56] T. Grobner. "Gadolinium – a specific trigger for the development of nephrogenic fibrosing dermopathy and nephrogenic systemic fibrosis?" In: *Nephrology Dialysis Transplantation* 21.4 (2006), pp. 1104–1108. DOI: `10.1093/ndt/gfk062`.

[57] T. Kanda et al. "High Signal Intensity in the Dentate Nucleus and Globus Pallidus on Unenhanced T1-weighted MR Images: Relationship with Increasing Cumulative Dose of a Gadolinium-based Contrast Material". In: *Radiology* 270.3 (2014), pp. 834–841. DOI: `10.1148/radiol.13131669`.

[58] W. G. Bradley and V Waluch. "Blood flow: magnetic resonance imaging." In: *Radiology* 154.2 (1985), pp. 443–450. DOI: `10.1148/radiology.154.2.3966131`.

[59] P. R. Moran. "A flow velocity zeugmatographic interlace for NMR imaging in humans". In: *Magn. Reson. Med.* 1.4 (1982), pp. 197–203. DOI: `10.1016/0730-725X(82)90170-9`.

[60] W. T. Dixon et al. "Projection angiograms of blood labeled by adiabatic fast passage". In: *Magn. Reson. Med.* 3.3 (1986), pp. 454–462. DOI: `10.1002/mrm.1910030311`.

[61] D. S. Williams et al. "Magnetic resonance imaging of perfusion using spin inversion of arterial water." In: *Proc. Natl. Acad. Sci. USA* 89.1 (1992), pp. 212–216. DOI: `10.1073/pnas.89.1.212`.

[62] J. A. Detre et al. "Tissue specific perfusion imaging using arterial spin labeling". In: *NMR in Biomedicine* 7.1-2 (1994), pp. 75–82. DOI: `10.1002/nbm.1940070112`.

[63] W. Dai et al. "Continuous Flow Driven Inversion for Arterial Spin Labeling Using Pulsed Radiofrequency and Gradient Fields". In: *Magn Reson Med* 60.6 (2008), pp. 1488–1497. DOI: `10.1002/mrm.21790`.

[64] M. A. Chappell et al. "A general framework for the analysis of vessel encoded arterial spin labeling for vascular territory mapping". In: *Magn. Reson. Med.* 64.5 (2010), pp. 1529–1539. DOI: `10.1002/mrm.22524`.

[65] E. Candes, J. Romberg, and T. Tao. "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information". In: *IEEE T. Inform. Theory.* 52.2 (2006), pp. 489–509. DOI: `10.1109/TIT.2005.862083`.

[66] T. W. Okell et al. "A kinetic model for vessel-encoded dynamic angiography with arterial spin labeling". In: *Magn Reson Med* 68.3 (2012), pp. 969–979. DOI: `10.1002/mrm.23311`.

[67] J. A. Fessler and B. P. Sutton. "Nonuniform fast Fourier transforms using min-max interpolation". In: *IEEE Trans. Signal Process.* 51.2 (2003), pp. 560–574. DOI: `10.1109/TSP.2002.807005`.

[68] J. A. Fessler. *Michigan Image Reconstruction Toolbox (http://web.eecs.umich.edu/~fessler/code/index.html).*

[69] T. W. Okell et al. "Optimization of 4D vessel-selective arterial spin labeling angiography using balanced steady-state free precession and vessel-encoding". In: *NMR Biomed.* 29.6 (2016), pp. 776–786. DOI: `10.1002/nbm.3515`.

[70] T. W. Okell. "Combined angiography and perfusion using radial imaging and arterial spin labeling". In: *Magn. Reson. Med* 81.1 (2019), pp. 182–194. DOI: `10.1002/mrm.27366`.

[71] J. Brenner. "The Hadamard Maximum Determinant Problem". In: *The American Mathematical Monthly* 79.6 (1972), p. 626. DOI: `10.2307/2317092`.

[72] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin. "Adaptive reconstruction of phased array MR imagery". In: *Magn. Reson. Med.* 43.5 (2000), pp. 682–690. DOI: `10.1002/(SICI)1522-2594(200005)43:5<682::AID-MRM10>3.0.CO;2-G`.

[73] M. Buehrer et al. "Array compression for MRI with large coil arrays". In: *Magn. Reson. Med.* 57.6 (2007), pp. 1131–1139. DOI: `10.1002/mrm.21237`.

[74] J. Fessler et al. “Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity”. In: *IEEE Trans. Signal Process.* 53.9 (2005), pp. 3393–3402. DOI: `10.1109/TSP.2005.853152`.

[75] Z. Wang et al. “Image Quality Assessment: From Error Visibility to Structural Similarity”. In: *IEEE T. Image. Process.* 13.4 (2004), pp. 600–612. DOI: `10.1109/TIP.2003.819861`.

[76] E. S. K. Berry, P. Jezzard, and T. W. Okell. “An Optimized Encoding Scheme for Planning Vessel-Encoded Pseudocontinuous Arterial Spin Labeling: Optimized Encoding Scheme for VEASL”. In: *Magn. Reson. Med.* 74.5 (2015), pp. 1248–1256. DOI: `10.1002/mrm.25508`.

[77] A. V. Barger et al. “Time-resolved contrast-enhanced imaging with isotropic resolution and broad coverage using an undersampled 3D projection trajectory”. In: *Magn. Reson. Med.* 48.2 (2002), pp. 297–305. DOI: `10.1002/mrm.10212`.

[78] L. Feng et al. “Golden-angle radial sparse parallel MRI: Combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI”. In: *Magn. Reson. Med.* 72.3 (2014), pp. 707–717. DOI: `10.1002/mrm.24980`.

[79] M Lustig et al. “k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 14. 2006, p. 2420.

[80] Jian Zhang et al. “Improved total variation based image compressive sensing recovery by nonlocal regularization”. In: *2013 IEEE International Symposium on Circuits and Systems (ISCAS2013)*. Beijing: IEEE, 2013, pp. 2836–2839. DOI: `10.1109/ISCAS.2013.6572469`.

[81] L. Zhao et al. “Rapid 3D dynamic arterial spin labeling with a sparse model-based image reconstruction”. In: *NeuroImage* 121 (2015), pp. 205–216. DOI: `10.1016/j.neuroimage.2015.07.018`.

[82] S. Nakajima et al. “Computer-assisted Surgical Planning for Cerebrovascular Neurosurgery”. In: *Neurosurgery* 41.2 (1997), pp. 403–410. DOI: `10.1097/00006123-199708000-00013`.

[83] J. Tsao, P. Boesiger, and K. P. Pruessmann. “k-t BLAST andk-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations”. In: *Magn. Reson. Med.* 50.5 (2003), pp. 1031–1042. DOI: `10.1002/mrm.10611`.

[84] H. Jung et al. “k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI”. In: *Magn. Reson. Med.* 61.1 (2009), pp. 103–116. DOI: `10.1002/mrm.21757`.

[85] M. Chiew et al. “k-t FASTER: Acceleration of functional MRI data acquisition using low rank constraints”. In: *Magn. Reson. Med.* 74.2 (2015), pp. 353–364. DOI: `10.1002/mrm.25395`.

[86] F. A. Breuer et al. “Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging”. In: *Magn. Reson. Med.* 53.3 (2005), pp. 684–691. DOI: `10.1002/mrm.20401`.

[87] S. Ravishankar and Y. Bresler. "MR Image Reconstruction From Highly Undersampled k-Space Data by Dictionary Learning". In: *IEEE Transactions on Medical Imaging* 30.5 (2011), pp. 1028–1041. DOI: `10.1109/TMI.2010.2090538`.

[88] D. S. Goolaub et al. "Multidimensional fetal flow imaging with cardiovascular magnetic resonance: a feasibility study". In: *J Cardiovasc Magn Reson* 20.1 (2018), p. 77. DOI: `10.1186/s12968-018-0498-z`.

[89] Z. Dong et al. "Echo planar time-resolved imaging with subspace reconstruction and optimized spatiotemporal encoding". In: *Magn. Reson. Med* 84.5 (2020), pp. 2442–2455. DOI: `10.1002/mrm.28295`.

[90] N. Maleki, W. Dai, and D. C. Alsop. "Optimization of background suppression for arterial spin labeling perfusion imaging". In: *Magn Reson Mater Phy* 25.2 (2012), pp. 127–133. DOI: `10.1007/s10334-011-0286-3`.

[91] M. Chiew, N. N. Graedel, and K. L. Miller. "Recovering task fMRI signals from highly under-sampled data with low-rank and temporal subspace constraints". In: *NeuroImage* 174 (2018), pp. 97–110. DOI: `10.1016/j.neuroimage.2018.02.062`.

[92] H. Mason et al. "Improving k-t PERRI: a low-rank data-driven fMRI k-t acceleration method". In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2019, p. 365.

[93] A. G. Anderson et al. "Adaptive retrospective correction of motion artifacts in cranial MRI with multicoil three-dimensional radial acquisitions: Adaptive 3D Motion Correction". In: *Magn. Reson. Med.* 69.4 (2013), pp. 1094–1103. DOI: `10.1002/mrm.24348`.

[94] G. Vaillant et al. "Retrospective Rigid Motion Correction in k-Space for Segmented Radial MRI". In: *IEEE Trans. Med. Imaging* 33.1 (2014), pp. 1–10. DOI: `10.1109/TMI.2013.2268898`.

[95] Y Suzuki et al. "Improving 4D pCASL angiography by combining Hadamard time-encoding with Look-Locker readout." In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 22. 2014, p. 6460.

[96] M. Günther. "Encoded Continuous Arterial Spin Labeling". In: *ISMRM Workshop on Cerebral Perfusion and Brain Function: Novel Techniques and Applications.* Salvador da Bahia, Brazil, 2007, p. 1.

[97] J. A. Wells et al. "In vivo hadamard encoded continuous arterial spin labeling (H-CASL)". In: *Magn. Reson. Med* 63.4 (2010), pp. 1111–1118. DOI: `10.1002/mrm.22266`.

[98] D. C. Look and D. R. Locker. "Time Saving in Measurement of NMR and EPR Relaxation Times". In: *Rev. Sci. Instrum.* 41.2 (1970), pp. 250–251. DOI: `10.1063/1.1684482`.

[99] J. G. Woods et al. "Optimization of time-encoded pseudo-continuous ASL angiography with a variable flip-angle scheme". In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 27. 2019, p. 844.

[100] T. W. Okell. "Combined Angiography and Perfusion using Radial Imaging and Arterial Spin Labeling". In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 24. 2016, p. 1001.

[101] T. W. Okell. “4D Combined Angiography and Perfusion using Radial Imaging and Arterial Spin Labeling”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 25. 2017, p. 675.

[102] M. Chiew and T. W. Okell. “Improved Golden Ratio Radial Arterial Spin Labeling Angiography Reconstruction using k-t Sparsity Constraints”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 26. 2018, p. 3351.

[103] T. W. Okell and M. Chiew. “High Resolution Perfusion Imaging using Golden Angle Radial Arterial Spin Labelling”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 27. 2019, p. 685.

[104] K. T. Block, M. Uecker, and J. Frahm. “Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint”. In: *Magn. Reson. Med.* 57.6 (2007), pp. 1086–1098. DOI: `10.1002/mrm.21236`.

[105] F. Knoll et al. “Second order total generalized variation (TGV) for MRI: Second Order TGV for MRI”. In: *Magn. Reson. Med.* 65.2 (2011), pp. 480–491. DOI: `10.1002/mrm.22595`.

[106] T. W. Okell et al. “Time- and vessel encoded pCASL: a free lunch with all the trimmings”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 23. Toronto, Canada, 2015, p. 264.

[107] T. W. Okell et al. “Visualizing artery-specific blood flow patterns above the circle of Willis with vessel-encoded arterial spin labeling”. In: *Magn. Reson. Med* 81.3 (2019), pp. 1595–1604. DOI: `10.1002/mrm.27507`.

[108] Y. Iryo et al. “Intracranial Dural Arteriovenous Fistulas: Evaluation with 3-T Four-dimensional MR Angiography Using Arterial Spin Labeling”. In: *Radiology* 271.1 (2014), pp. 193–199. DOI: `10.1148/radiol.13122670`.

[109] M. Garcia et al. “Feasibility of Flat Panel Detector CT in Perfusion Assessment of Brain Arteriovenous Malformations: Initial Clinical Experience”. In: *AJNR Am J Neuroradiol* 38.4 (2017), pp. 735–739. DOI: `10.3174/ajnr.A5091`.

[110] M. Chappell et al. “Variational Bayesian Inference for a Nonlinear Forward Model”. In: *IEEE Trans. Signal Process.* 57.1 (2009), pp. 223–236. DOI: `10.1109/TSP.2008.2005752`.

[111] D. C. Alsop 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”. In: *Magn. Reson. Med.* 73.1 (2015), pp. 102–116. DOI: `10.1002/mrm.25197`.

[112] J.-R. Liao et al. “Reduction of motion artifacts in cine MRI using variable-density spiral trajectories”. In: *Magn. Reson. Med.* 37.4 (1997), pp. 569–575. DOI: `10.1002/mrm.1910370416`.

[113] T. Phillips. *The Most Irrational Number - American Mathematical Society Feature Column (http://www.ams.org/publicoutreach/feature-column/fcarc-irrational1)*. 1999.

[114] H. J. Jeong et al. “Radial sliding-window magnetic resonance angiography (MRA) with highly-constrained projection reconstruction (HYPR)”. In: *Magn. Reson. Med.* 61.5 (2009), pp. 1103–1113. DOI: `10.1002/mrm.21888`.

[115] S. Wundrak et al. “A Small Surrogate for the Golden Angle in Time-Resolved Radial MRI Based on Generalized Fibonacci Sequences”. In: *IEEE Trans. Med. Imaging* 34.6 (2015), pp. 1262–1269. DOI: `10.1109/TMI.2014.2382572`.

[116] C.-M. Tsai and D. G. Nishimura. “Reduced aliasing artifacts using variable-density k-space sampling trajectories”. In: *Magn. Reson. Med.* 43.3 (2000). Publisher: John Wiley & Sons, Ltd, pp. 452–458. DOI: `10.1002/(SICI)1522-2594(200003)43:3<452::AID-MRM18>3.0.CO;2-B`.

[117] J. G. Pipe and P. Menon. “Sampling density compensation in MRI: Rationale and an iterative numerical solution”. In: *Magn. Reson. Med.* 41.1 (1999), pp. 179–86. DOI: `10.1002/(sici)1522-2594(199901)41:1<179::aid-mrm25>3.0.co;2-v`.

[118] D. M. Spielman, J. M. Pauly, and C. H. Meyer. “Magnetic resonance fluoroscopy using spirals with variable sampling densities”. In: *Magn. Reson. Med.* 34.3 (1995), pp. 388–394. DOI: `10.1002/mrm.1910340316`.

[119] V. Rasche et al. “Resampling of data between arbitrary grids using convolution interpolation”. In: *IEEE Trans. Med. Imaging* 18.5 (1999), pp. 385–392. DOI: `10.1109/42.774166`.

[120] E. B. Saff and A. B. J. Kuijlaars. “Distributing many points on a sphere”. In: *The Mathematical Intelligencer* 19.1 (1997), pp. 5–11. DOI: `10.1007/BF03024331`.

[121] C. G. Koay. “A simple scheme for generating nearly uniform distribution of antipodally symmetric points on the unit sphere”. In: *J. Comput. Sci.-Neth.* 2.4 (2011), pp. 377–381. DOI: `10.1016/j.jocs.2011.06.007`.

[122] K. M. Jansons and D. C. Alexander. “Persistent angular structure: new insights from diffusion magnetic resonance imaging data”. In: *Inverse Probl.* 19.5 (2003), pp. 1031–1046. DOI: `10.1088/0266-5611/19/5/303`.

[123] D. K. Jones. “The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: A Monte Carlo study”. In: *Magn. Reson. Med.* 51.4 (2004), pp. 807–815. DOI: `10.1002/mrm.20033`.

[124] H. K. Song et al. “Noncontrast enhanced four-dimensional dynamic MRA with golden angle radial acquisition and k-space weighted image contrast (KWIC) reconstruction”. In: *Magn. Reson. Med.* 72.6 (2014), pp. 1541–1551. DOI: `10.1002/mrm.25057`.

[125] A. Fyrdahl et al. “Sector-wise golden-angle phase contrast with high temporal resolution for evaluation of left ventricular diastolic dysfunction”. In: *Magn Reson Med* (2019), mrm.28018. DOI: `10.1002/mrm.28018`.

[126] K. Shmueli et al. “Magnetic susceptibility mapping of brain tissue in vivo using MRI phase data”. In: *Magn. Reson. Med.* 62.6 (2009), pp. 1510–1522. DOI: `10.1002/mrm.22135`.

[127] M. Sabati and A. A. Maudsley. “Fast and high-resolution quantitative mapping of tissue water content with full brain coverage for clinically-driven studies”. In: *Magn. Reson. Med.* 31.10 (2013), pp. 1752–1759. DOI: `10.1016/j.mri.2013.08.001`.

[128] J. D. Trzasko and A. Manduca. “A Unified Tensor Regression Framework for Calibrationless Dynamic, Multi-Channel MRI Reconstruction”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 21. 2013, p. 603.

[129] R. Ramb et al. “Low-rank plus sparse tensor reconstruction for high-dimensional cardiac MRI”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 25. 2017, p. 1199.

[130] J. G. Woods. “Time-Efficient Methods for Non-Invasive Brain Blood Flow Imaging Using Arterial Spin Labelling MRI”. PhD thesis. University of Oxford, 2019.

[131] D. S. Smith et al. “Trajectory optimized NUFFT: Faster non-Cartesian MRI reconstruction through prior knowledge and parallel architectures”. In: *Magn. Reson. Med* 81.3 (2019), pp. 2064–2071. DOI: `10.1002/mrm.27497`.

[132] D. Liang et al. “Deep Magnetic Resonance Image Reconstruction: Inverse Problems Meet Neural Networks”. In: *IEEE Signal Process. Mag.* 37.1 (2020), pp. 141–151. DOI: `10.1109/MSP.2019.2950557`.

[133] Y.-C. Kim, S. S. Narayanan, and K. S. Nayak. “Flexible retrospective selection of temporal resolution in real-time speech MRI using a golden-ratio spiral view order: Speech MRI using Golden-Ratio Spiral”. In: *Magn. Reson. Med.* 65.5 (2011), pp. 1365–1371. DOI: `10.1002/mrm.22714`.

[134] G. T. Kowalik et al. “Golden ratio stack of spirals for flexible angiographic imaging: Proof of concept in congenital heart disease”. In: *Magn. Reson. Med* 81.1 (2019), pp. 90–101. DOI: `10.1002/mrm.27353`.

[135] E. J. Zucker et al. “Free-breathing pediatric chest MRI: Performance of self-navigated golden-angle ordered conical ultrashort echo time acquisition: Free-Breathing Pediatric UTE Chest MRI”. In: *J. Magn. Reson. Imaging* 47.1 (2018), pp. 200–209. DOI: `10.1002/jmri.25776`.

[136] S. Li et al. “Dynamic magnetic resonance imaging method based on golden-ratio cartesian sampling and compressed sensing”. In: *PLoS ONE* 13.1 (2018), e0191569. DOI: `10.1371/journal.pone.0191569`.