Acquisition and Reconstruction Methods for Hybrid 2D/3D Diffusion MRI



Wenchuan Wu

St Cross College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2017

This thesis is dedicated to
my parents and my wife
for their constant support and unconditional love

Abstract

Diffusion magnetic resonance imaging (MRI) has been increasingly used in neuroscience studies, particularly for mapping white matter tracts in the brain. Typically, diffusion MRI is acquired using a two-dimensional (2D) single-shot echo planar imaging sequence, which allows rapid acquisition and reduced sensitivity to subject motion. However, conventional 2D diffusion MRI faces many limitations, such as long TR (repetition time) which lead to low SNR (signal-to-noise ratio) efficiency, and long scan times with advanced diffusion protocols which usually require a large number of diffusion directions and/or b values. These limitations become more acute at high spatial resolution. Three-dimensional (3D) methods have also been developed for diffusion MRI, but they are not widely used for acquiring data in vivo due to the challenges in subject motion. Hybrid 2D/3D methods, including 3D multi-slab acquisition and simultaneous multi-slice acquisition, have been recently proposed to address the limitations of conventional 2D and 3D methods. However, 3D multi-slab acquisition faces the problem of slab boundary artefacts, which could decrease the image quality and propagate into diffusion quantifications. Simultaneous multi-slice acquisition would suffer from significant noise amplification when in-plane under-sampling is also applied. The work in this thesis seeks to develop acquisition and reconstruction methods to improve hybrid 2D/3D diffusion MRI. A new method is proposed to correct slab boundary artefacts in 3D multi-slab imaging, which jointly estimates the slab profile and underlying image using a nonlinear reconstruction. Correction results demonstrate superior performance compared with previously proposed methods. A k-q acquisition and reconstruction approach is developed to accelerate diffusion MRI based on Gaussian process methods. Here, we target improvements to simultaneous multi-slice imaging, demonstrating high acceleration factors in combination with in-plane under-sampling. Combinations of hybrid 2D/3D acquisitions with the ultra-high field of 7T are also investigated, which could enable high resolution diffusion MRI without substantially compromising SNR. The methods developed in this work are expected to improve the data quality and scan efficiency of diffusion MRI.

Acknowledgements

I would like to thank all the people who have contributed in various ways to the work presented in this thesis.

Firstly, I am extremely grateful to my supervisors, Professor Karla Miller and Dr. Peter Koopmans for their unfailing support and guidance throughout my DPhil. It has been a pleasure to work with them. Karla's deep insights in the field of MRI helped me at various stages of my research, and her expertise, motivation and encouragement added considerably to my graduate experience. Peter has always been available to discuss research and explain MRI concepts, whose input at all levels of my research project deserves special mention.

I am very greatful for the generous funding received from FP7 Marie Curie Action of the European Commission that supported my DPhil in the first instance. I would like to thank Karla for the financial support I received to make it through the last year of my DPhil. Funding from the Wellcome Trust also supported the work in this thesis.

I would like to thank the whole FMRIB Physics group for giving advice on my research projects, being subjects in my noisy scan and offering suggestions on presentations and posters. The friendly atmosphere in the physics group has also made it a very pleasant time. I am particularly grateful to Dr. Robert Frost for his input in my first research project, which has been incredibly helpful.

I also would like to thank many people from FMRIB analysis group for explaining concepts in data analysis and giving valuable advice on FSL, machine learning, parallel computing, cluster, etc.

Lastly, but most importantly, I wish to thank my parents, Chengping Wu and Chengping Cao for their consistent support though my life. My gratitude to them is beyond words. I must acknowledge my wife and best friend, Lixia Yang, without whose love, encouragement and support, I would not finish this thesis.

Contents

Preface **xiii**

1 Introduction **1**
1.1 Motivation 1
1.2 Thesis Outline 4

2 Background **7**
2.1 MRI physics 7
2.1.1 MR Signal 8
2.1.2 Image Formation 18
2.2 Diffusion MRI 26
2.2.1 Diffusion MRI Signal 27
2.2.2 Diffusion MRI Sequence 28
2.2.3 Diffusion MRI Measurement 28
2.3 Partially Parallel Imaging 31
2.3.1 SENSE 32
2.3.2 GRAPPA 36

3 Image Formation in Diffusion MRI **39**
3.1 Introduction 39
3.2 Improving Image Fidelity 41
3.2.1 SSH-EPI Artefacts 41
3.2.2 Multi-shot Acquisition 45
3.2.3 Reduced FOV Method 53
3.3 Accelerating Diffusion MRI Acquisitions 55
3.3.1 Simultaneous Multi-slice Imaging 55
3.3.2 Fast Diffusion MRI with Compressed Sensing 59
3.4 Increasing Image SNR 61
3.4.1 Three-dimensional Acquisition 62
3.4.2 Ultra-high Field 64
3.5 Summary 65

4 Reducing Slab-Boundary Artefacts for 3D Multi-slab Diffusion MRI 67

4.1 Introduction . . . 68
4.1.1 Overview of Slab-Boundary Artefacts . . . 70
4.2 Bloch Simulations . . . 73
4.2.1 Slab Crosstalk in Multi-slab Imaging . . . 73
4.2.2 Propagation of Slab Crosstalk into Diffusion Quantification . 75
4.3 Methods . . . 78
4.3.1 Proposed Correction . . . 78
4.3.2 Numerical Simulation . . . 83
4.3.3 In vivo . . . 84
4.3.4 Reconstruction . . . 85
4.4 Results . . . 87
4.4.1 Reconstruction with Simulation Data . . . 87
4.4.2 Reconstruction with Data Acquired In Vivo . . . 92
4.5 Discussion . . . 95
4.6 Conclusion . . . 98

5 High-resolution Diffusion MRI at 7T using a 3D Multi-slab Acquisition 101

5.1 Introduction . . . 102
5.2 Theory . . . 105
5.3 Methods . . . 106
5.3.1 Sequence . . . 106
5.3.2 Data Acquisition . . . 108
5.3.3 Reconstruction . . . 110
5.3.4 Post-processing . . . 113
5.4 Results . . . 113
5.4.1 Comparison of Different ACS Acquisitions . . . 113
5.4.2 Slab-Boundary Artefacts Correction . . . 118
5.4.3 Diffusion Tensor Results . . . 118
5.4.4 Tractography Results . . . 121
5.5 Discussion . . . 123
5.6 Conclusion . . . 129

6 Accelerated k-q dMRI using Gaussian Processes 131
6.1 Introduction . . . 132
6.2 Theory . . . 133
6.2.1 Gaussian Process Regression . . . 133
6.2.2 Gaussian Process Model for dMRI Data . . . 136
6.2.3 dMRI reconstruction with a Gaussian Process Prior . . . 138
6.3 Methods . . . 144
6.3.1 k-q Sampling . . . 144
6.3.2 Simulation . . . 146
6.3.3 In Vivo . . . 148
6.3.4 Reconstruction . . . 150
6.3.5 Post-processing . . . 151
6.4 Results . . . 151
6.4.1 Simulation . . . 151
6.4.2 In Vivo . . . 159
6.5 Discussion . . . 167
6.6 Conclusion . . . 171

7 Summary and Future Work 173
7.1 Thesis Summary . . . 173
7.2 Future Work . . . 175

Appendices

A Quantitative Description of EPI Distortion 181

B Jacobian Matrix of the Nonlinear Operator in NPEN 183

C Posterior Probability Distribution with Gaussian Process Priors 187

References 189

Preface

List of Publications Arising from This Thesis

Journal Articles

1. **Wu W**, Koopmans PJ, Jesper A, Miller KL, Diffusion acceleration with Gaussian process estimated reconstruction (DAGER). In preparation.

2. **Wu W**, Miller KL, Image formation in diffusion MRI: a review of recent technique development. J Magnetic Resonance Imaging. 2017;46(3):646-662.

3. **Wu W**, Poser AB, Douaud G, Frost R, In MH, Speck O, Koopmans PJ, Miller KL. High-resolution diffusion MRI at 7T using a three-dimensional multi-slab acquisition. NeuroImage. 2016;143:1-14.

4. **Wu W**, Koopmans PJ, Frost R, Miller KL. Reducing slab boundary artifacts in three-dimensional multislab diffusion MRI using nonlinear inversion for slab profile encoding (NPEN). Magn Reson Med. 2016; 76:1183-1195.

Conference Proceedings

1. **Wu W**, Koopmans PJ, Andersson JR, Miller KL. Accelerated diffusion MRI using Gaussian Process. OHBM annual meeting, Vancouver, Canada, 2017.

2. **Wu W**, Koopmans PJ, Andersson JR, Miller KL. Accelerated k-q diffusion MRI reconstruction using Gaussian processes. ISMRM 25th Annual Meeting & Exhibition, Honolulu, USA, 2017.

3. **Wu W**, Koopmans PJ, Miller KL. Design of a locally incoherent sampling scheme for improved joint k-q reconstruction of diffusion MRI. ISMRM 25th Annual Meeting & Exhibition, Honolulu, USA, 2017.

4. **Wu W**, Koopmans PJ, Frost R, In Myung-Ho, Speck Oliver, Miller KL. High-resolution diffusion imaging at 7T using 3D multi-slab EPI. ISMRM 24th Annual Meeting & Exhibition, Singapore, 2016.

5. **Wu W**, Miller KL, Poser B, Koopmans PJ. Improved Robustness of Parallel Imaging in 3D Multi-Slab Diffusion Imaging Using an Adapted FLEET Approach. ISMRM 24th Annual Meeting & Exhibition, Singapore, 2016.

6. **Wu W**, Koopmans PJ, Frost R, Miller KL. Reducing Slab Boundary Artifacts in 3D Multi-Slab Diffusion Imaging by Jointly Estimating Slab Profile and Image. ISMRM 23rd Annual Meeting & Exhibition, Toronto, Canada, 2015.

7. **Wu W**, Koopmans PJ, Miller KL. Effects of Slab Boundary Artifacts on Diffusion Measures in 3D Multi-Slab Diffusion Imaging. ISMRM 23rd Annual Meeting & Exhibition, Toronto, Canada, 2015.

1 Introduction

This DPhil thesis concerns the development of acquisition and reconstruction methods for hybrid 2D/3D diffusion magnetic resonance imaging (MRI) of the brain. Diffusion MRI (dMRI) is a non-invasive technique that reveals tissue microstructural features by probing the diffusion of water molecules. Diffusion MRI has been widely used in clinics, particularly for the management of patients with acute ischemic stroke. In neuroscience research, dMRI is rapidly becoming a standard tool to study the structural connectivity of the brain due to its ability to characterize white matter fibre architectures. This Chapter describes the motivation of the research and an outline of the thesis.

1.1 Motivation

Diffusion MRI typically uses a two-dimensional (2D) acquisition, where the target volume is split into multiple thin slices and each slice is excited and encoded separately. A limitation of 2D acquisition is the coupling between the repetition time (TR) and the number of slices, which necessitates a long TR (typically 6-10s)

to achieve a full-brain coverage. For brain dMRI scans, the signal-to-noise ratio (SNR) efficiency of a spin-echo dMRI sequence at 3T is maximized when TR=1-2s. Hence, 2D acquisition is subject to a low SNR efficiency, particularly for acquiring small isotropic voxels. Three-dimensional (3D) acquisition removes this coupling by exciting the entire volume, defining voxels using a 3D k-space and acquiring the full image using multiple readouts. 3D acquisition can achieve thin slice thicknesses easily by increasing the number of phase encoding steps along the slice direction. However, to acquire an image within a reasonable scan time, 3D acquisition requires very short TR (<200ms), which is also far from the optimal TR. More importantly, the motion induced phase errors need to be accounted for in all three dimensions, but acquiring 3D phase-navigators is very challenging in the limited scan time.

The limitations of 2D and 3D acquisitions in diffusion MRI can be alleviated using a hybrid of these two methods. Two previously proposed hybrid 2D/3D approaches are 3D multi-slab acquisition and simultaneous multi-slice acquisition.

In 3D multi-slab acquisition, the entire volume is divided into multiple small slabs and each slab is separately excited and encoded. The TR in 3D multi-slab acquisition is proportional to the number of slabs (rather than the number of slices), which is typically 12-14 for full brain coverage, making it highly compatible with the optimal TR for maximizing SNR efficiency. In addition, recent studies show that the motion-induced 3D phase errors can be well approximated by a 2D navigator if a thin slab thickness is used, which can relieve the burden of acquiring 3D navigators. With 3D k-space encoding for each slab, high spatial resolution dMRI can be achieved using 3D multi-slab acquisition. However, 3D multi-slab acquisition suffers from slab-boundary artefacts, which appear as periodic signal modulation in dMRI images and can propagate into diffusion quantification. Therefore, a robust correction method for slab-boundary artefacts needs to be developed in order to apply 3D multi-slab dMRI in practice.

Simultaneous multi-slice acquisition can achieve a shorter TR by reducing the number of excitations needed to acquire the entire volume, which leads to improved SNR efficiency and reduced scan time. A major improvement on simultaneous multi-slice acquisition is the blipped-CAIPI technique [1, 2], which shifts slices in-plane to reduce the noise amplification caused by the ill-conditioning of the reconstruction. The simultaneous multi-slice acquisition represents a form of through-plane acceleration using parallel imaging. In-plane parallel imaging (under-sampling) has long been used in dMRI to reduce image distortion and blurring. These two forms of parallel imaging are compatible with each other and ideally one would like to combine them in data acquisition. However, simultaneous multi-slice acquisition with in-plane under-sampling faces significant challenges in acquiring dMRI data. The efficacy of blipped-CAIPI is compromised by the FOV folding effects caused by in-plane under-sampling, leading to strong noise amplification which is particularly problematic for diffusion MRI due to its intrisically low SNR. As a result, dMRI protocols usually use high through-plane acceleration without in-plane under-sampling [3, 4], or moderate through-plane acceleration with high in-plane under-sampling [5]. New acquisition and reconstruction methods need to be developed to enable high accelerations in both through-plane and in-plane directions.

Diffusion MRI can potentially benefit from ultra-high field (7T and above), which can provide stronger signals than more standard field strengths (e.g. 3T) that can be used to preserve the SNR at higher spatial resolution. However, combining ultra-high field with either 3D multi-slab acquisition or simultaneous multi-slice acquisition faces significant limitations. Due to the strong $\mathbf{B}_0$ inhomogeneity at ultra-high field, the slab profile of 3D multi-slab acquisition can be severely distorted, leading to stronger slab boundary artefacts. Therefore, correction methods for slab-boundary artefacts that work well at standard field strengths might need further improvements at 7T. Another downside of 7T is the increased specific absorption rate (SAR), which is a great challenge for spin-echo diffusion MRI, especially when simultaneous

multi-slice acquisition is used. Hence, sequence modifications are needed in order to achieve high slice acceleration at 7T. In-plane under-sampling is critical for dMRI at ultra-high field, particularly with a single-shot echo planar imaging readout, which suffers from image distortion and blurring due to the stronger $\mathbf{B}_0$ inhomogeneity and reduced tissue T_2^*. The performance of parallel imaging reconstruction is affected by the quality of the calibration data, which is usually sensitive to subject motion. Recent studies show that parallel imaging reconstruction can be improved by using a robust acquisition approach for the auto-calibration signal, which has been demonstrated in 2D and 3D. To improve the data quality and robustness of hybrid 2D/3D dMRI at ultra-high field similar strategies should be adopted.

1.2 Thesis Outline

This thesis is structured as follows:

Chapter 2 gives a brief background in MRI physics, generation and measurement of diffusion MRI signal, and parallel imaging methods that are commonly used in diffusion MRI.

Chapter 3 reviews the recent developments of image formation methods for diffusion MRI, covering three kinds of advances: improving image fidelity, accelerating acquisition and increasing SNR. How to use these developments in practice is also discussed.

Chapter 4 is the first research chapter on reducing slab-boundary artefacts using a nonlinear inversion for slab-profile encoding (NPEN). Bloch simulation is used to study the slab crosstalk effect, including its causes and how it affects diffusion measurements. Then a new method to minimize the slab-boundary artefacts for 3D multi-slab dMRI with SNR-optimal TRs ($\sim$2s) is proposed. Numerical simulation and in vivo data are used to validate NPEN, with comparisons with two other methods previously proposed.

Chapter 5 describes the combination of two highly compatible methods: ultra-high field and 3D multi-slab acquisitions to improve the SNR of high-resolution diffusion MRI. A recently published method is adapted in 3D multi-slab acquisition to improve the robustness of auto-calibration signal against subject motion. The NPEN method described in Chapter 4 is applied to reduce the slab boundary artefacts, with modifications to address the stronger $\mathbf{B}_0$ inhomogeneity at 7T.

Chapter 6 presents a method to accelerate k-q diffusion MRI based on Gaussian processes. The correlations between diffusion volumes are explored with a Gaussian process model, which is trained from the data and used as a prior to improve the reconstruction. Numerical simulations and data acquired in vivo are used to validate the proposed method.

Chapter 7 summarises the research work in this thesis and discusses possible future research directions.

Chapter 3, 4 and 5 are based on three peer-reviewed journal papers we have published in the Journal of Magnetic Resonance Imaging [6], Magnetic Resonance in Medicine [7] and NeuroImage [8], respectively.

2 Background

Contents

2.1 MRI physics 7
2.1.1 MR Signal 8
2.1.2 Image Formation 18
2.2 Diffusion MRI 26
2.2.1 Diffusion MRI Signal 27
2.2.2 Diffusion MRI Sequence 28
2.2.3 Diffusion MRI Measurement 28
2.3 Partially Parallel Imaging 31
2.3.1 SENSE 32
2.3.2 GRAPPA 36

2.1 MRI physics

This section presents a general introduction to MRI, which is based on standard text books [9–12].

2.1.1 MR Signal

The physical foundation of MRI is the nuclear magnetic resonance (NMR) phenomenon that nuclei within an external magnetic field absorb and re-emit radio frequency (RF) energy [13, 14].

A nucleus with an odd atomic number and/or odd neutron number possesses an intrinsic spin angular momentum, $\mathbf{L}$, or a nonzero spin. Those with even numbers for both have zero spin. As the proton carries a positive electric charge, the self-rotation of the nucleus produces a local magnetic field, which can be represented by a vector quantity called a magnetic moment, $\boldsymbol{\mu}$. These two properties are linearly related as:

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

where γ is the gyromagnetic ratio, a constant unique to each type of isotope. For example, $\gamma = 2\pi \times 42.577 \times 10^6\,\text{rad/s/T}$ for isotope Hydrogen-1 (^{1}H) and $\gamma = 2\pi \times 10.705 \times 10^6\,\text{rad/s/T}$ for isotope Carbon-13 (^{13}C).

Theoretically, any nucleus with nonzero spin can be used as the object of MRI studies, but in practice ^{1}H is the dominant choice for MRI examinations because it is the most abundant isotope for hydrogen, which is present in great abundance in the human body (i.e. hydrogen comprised in water and fat) and its high sensitivity to external magnetic field (stronger signal). ^{1}H is commonly referred to as a proton. Work in this thesis also focuses on proton MRI.

According to quantum theory, the magnitude of $\boldsymbol{\mu}$ is a certain value, but its direction is totally random in the absence of an external magnetic field (see Fig. 2.1a). As MRI studies the ensemble behaviours of a collection of protons (i.e. vector sum of magnetic moments from all protons), there is no net magnetization that can be observed.

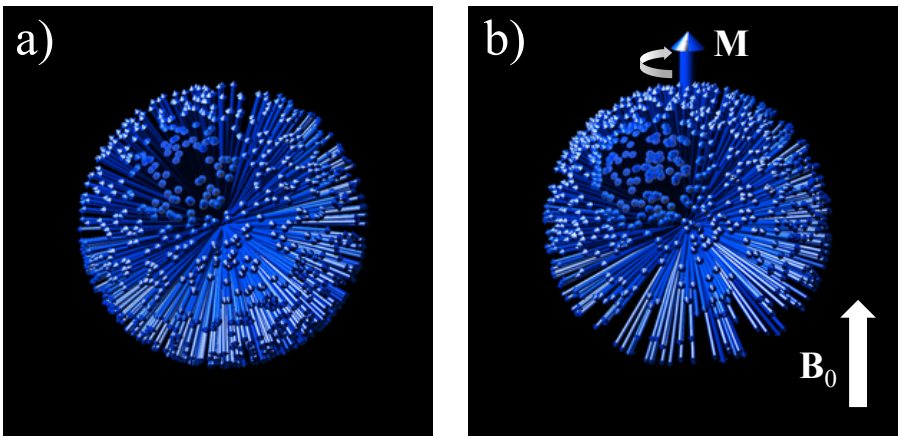


Figure 2.1: Net magnetization a) without an external magnetic field, where spins possess random orientations, resulting in zero net magnetization. b) in the presence of an external magnetic field $\mathbf{B}_0$, spin distribution is slightly skewed towards the direction of $\mathbf{B}_0$, leading to a net magnetization $\mathbf{M}$. Diagrams courtesy of Karla Miller adapted from Hanson [15].

Magnetization

In order to detect the magnetic moment of protons, the first step is to expose protons to an external magnetic field $\mathbf{B}_0$, in which the spin distribution will have a slight tendency towards the direction of $\mathbf{B}_0$, leading to a net magnetization $\mathbf{M}$ (Fig. 2.1b).

In the presence of $\mathbf{B}_0$, two energy states are created, leading to two spin populations: parallel-spins with lower energy and anti-parallel spins with higher energy. The numbers of spins in these two energy states are not identical and the difference is governed by Boltzmann distribution

$$\frac{n_\downarrow}{n_\uparrow} = \mathrm{e}^{-\Delta E/KT_s} \tag{2.2}$$

where $n_\downarrow$ and $n_\uparrow$ are the number of anti-parallel spins and parallel spins, respectively. K is the Boltzmann constant, T_s denotes the absolute temperature. ΔE is the energy difference between the two energy levels, which is determined according to quantum theory:

$$\Delta E = \gamma \hbar B_0 \tag{2.3}$$

$\hbar$ is the Planck's constant divided by 2π, B_0 is the magnitude of $\mathbf{B_0}$. Under typical experimental conditions, ΔE is much smaller than KT_s, thus the population

difference is only a small fraction of total spin numbers

$$n_{\uparrow} - n_{\downarrow} \approx n\frac{\gamma\hbar B_0}{2KT_{\mathrm{s}}} \tag{2.4}$$

where n is the total number of spins within the sample. For proton MRI at 3T, the amount of spin excess is ∼10 per million. Regardless, the net magnetization can still become detectable if there is an abundance of spins within the sample. Proton is one of the best choices for MRI scan in vivo based on this fact.

The net magnetization $\mathbf{M}$ is parallel to $\mathbf{B}_0$ and its magnitude M_0 is determined by

$$M_0 = \frac{\gamma^2\hbar^2\rho_0 B_0}{4KT_s} \tag{2.5}$$

where ρ_0 is the spin density (number of protons in the sample). M_0 is associated with the equilibrium state of $\mathbf{M}$ when the proton spins are at their lowest energy state. $\mathbf{M}$ will return to the equilibrium state after any external disturbances such as energy absorption from a RF excitation. Eq. 2.5 also indicates the strength of the net magnetization can be increased by manipulating several variables. However, for scanning in vivo, increasing the field strength B_0 is the only practical way to produce a stronger M_0. For this reason, the field strength of commercial scanners has been increasing since the invention of MRI, from low field (under 1 tesla) to high field (1.5 and 3 tesla) and ultra-high field (7+ tesla), to achieve stronger MRI signal.

Precession

From a classical perspective, the magnetic moment experiences a torque (rate of change of angular momentum) from $\mathbf{B_0}$, which can be expressed as

$$\frac{\mathrm{d}\mathbf{L}}{\mathrm{d}t} = \boldsymbol{\mu} \times \mathbf{B}_0 \tag{2.6}$$

By default, $\mathbf{B_0}$ is designed along the z-axis in the laboratory frame. In MRI literature, z-axis and x-y plane are often referred to as longitudinal and transverse,

respectively. Insert Eq. 2.1 into the Eq. 2.6:

$$\frac{\mathrm{d}\boldsymbol{\mu}}{\mathrm{d}t} = \gamma\boldsymbol{\mu} \times \mathbf{B}_0 \tag{2.7}$$

The behaviour of the net magnetization $\mathbf{M}$ can be considered as a sum of the magnetic moments

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \gamma\mathbf{M} \times \mathbf{B}_0 \tag{2.8}$$

which gives the solution

$$\begin{aligned} M_x(t) &= M_x(0)\cos(\gamma B_0 t) + M_y(0)\sin(\gamma B_0 t) \\ M_y(t) &= M_y(0)\cos(\gamma B_0 t) - M_x(0)\sin(\gamma B_0 t) \\ M_z(t) &= M_z(0) \end{aligned} \tag{2.9}$$

where vector $[M_x(0), M_x(0), M_x(0)]$ is the initial condition for $\mathbf{M}$. Eq. 2.9 describes a precession motion around the z-axis at a fixed frequency:

$$\omega_0 = \gamma B_0 \tag{2.10}$$

which is known as the Larmor frequency. A precessing spin is analogous to a gyroscope spinning around an axis parallel to the direction of gravity with its self-rotation axis tilted from the axis of gravity. The Larmor frequency is an intrinsic property of an isotope, which is proportional to the strength of the external magnetic field. For example, larmor frequency of ^{1}H is ∼128MHz at 3T and ∼298MHz at 7T.

Excitation

After spin polarization, the net magnetization still can not be detected, unless it is tipped away from the longitudinal direction. This step is known as excitation, achieved by applying a radio-frequency (RF) magnetic field $\mathbf{B}_1$, rotating in the transverse plane, to the polarized spin system. The $\mathbf{B}_1$ field is tuned at the Larmor frequency to satisfy the resonance condition. The $\mathbf{B}_1$ field is usually transmitted for a short duration ($<$ 10ms) and is commonly called an RF pulse.

During RF excitation, spins absorb energy from $\mathbf{B}_1$ and some transitions within the spin system take place. In a classical picture, the effect of the RF pulse can also be expressed as a torque on $\mathbf{M}$, which enforces $\mathbf{M}$ to tilt away from its equilibrium direction (i.e. longitudinal). Fig. 2.2a shows the behaviour of $\mathbf{M}$ during excitation in the laboratory frame, where $\mathbf{M}$ rotates about the z axis at the larmor frequency and at the same time gradually rotates away from the z axis. For ease of understanding the behaviour of $\mathbf{M}$ during excitation, it is helpful to use a frame rotating around the z axis at the Larmor frequency (Fig. 2.2b), in which precession about $\mathbf{B}_0$ disappears while flipping effects by $\mathbf{B}_1$ remain. This frame is conventionally called the rotating frame.

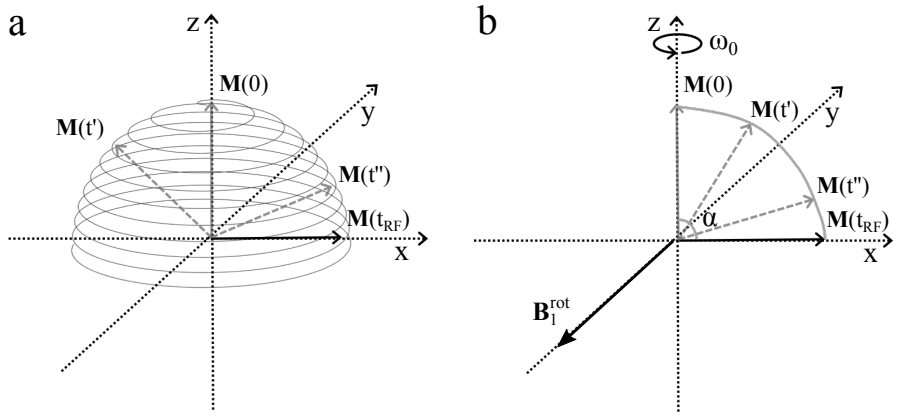


Figure 2.2: Schematics of $\mathbf{B}_1$ effects on the net magnetization in a) the laboratory frame and b) the rotating frame. $t = 0$ and $t = t_{\mathrm{RF}}$ represent the start and end timing of the RF pulse respectively, and t' and t'' are two time points in between. α is the flip angle of the RF pulse, which is 90° in this example.

An RF pulse can be characterized as $\mathbf{B}_1 = B_1^{\mathrm{env}}(t)\exp(-i(\omega_{\mathrm{RF}}t + \phi_0))$ in the laboratory frame, where $B_1^{\mathrm{env}}(t)$ is the envelope of the pulse, ω_{RF} is the RF frequency which is commonly tuned to the Larmor frequency. ϕ_0 is the initial phase of the RF pulse. i is the imaginary unit defined as $\sqrt{-1}$. In the rotating frame, the RF pulse is expressed as $\mathbf{B}_1^{\mathrm{rot}} = B_1^{\mathrm{env}}(t)(\cos(\phi_0)\hat{\mathbf{i}} - \sin(\phi_0)\hat{\mathbf{j}})$, which is a stationary vector in the x-y plane whose direction has a ϕ_0 angle away from the positive direction of x axis. $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$ are the unit vectors along x and y axes, respectively.

The effects of RF excitation in the rotating frame is illustrated in Fig. 2.2b, where the net magnetization $\mathbf{M}^{\text{rot}}$ precesses about the $\mathbf{B}_1^{\text{rot}}$. As a result, $\mathbf{M}^{\text{rot}}$ rotates away from z-axis, leading to a measurable transverse component. The angle of rotation enforced by the $\mathbf{B}_1^{\text{rot}}$ field is called flip angle α (as shown in Fig. 2.2b), which is determined by the magnitude and the duration of the RF pulse. For on-resonance excitation, the flip angle is given by

$$\alpha = \int_0^{t_{\text{RF}}} \gamma B_1^{\text{env}}(t)\,\mathrm{d}t \tag{2.11}$$

Signal reception

After excitation, the precession of $\mathbf{M}$ produces a time-varying magnetic flux in the transverse plane. According to Faraday's law, this induces voltage in a conductive receiver coil placed nearby, provided the coil plane is not parallel to the transverse plane. This voltage is the actual MR signal acquired from the scanner and used for image formation, which is known as the free induction decay (FID) signal. The strength of FID signal is proportional to the magnitude of transverse magnetization after excitation. If $\mathbf{M}$ is along the longitudinal direction, excitation with a 90° RF pulse will tip $\mathbf{M}$ to the transverse plane, which produces the maximal signal. A 180° RF pulse flips $\mathbf{M}$ along the longitudinal direction, leaving no magnetization in the transverse plane, thus no signal can be observed. $\mathbf{B}_1$ transmission and signal reception often use the same coil (i.e. T/R coil), which switches between two different modes at different sequence timing.

Relaxation

After excitation, spins will re-emit the absorbed energy and return to their equilibrium state. In the classical picture, $\mathbf{M}$ will gradually return to the longitudinal axis and recover its magnitude to M_0. This process is known as relaxation.

T_1 relaxation is the process of the excited spins transferring energy to its surroundings (lattice), which is also known as spin-lattice relaxation. The longitudinal

component of the net magnetization, M_z, recovers towards its equilibrium state (i.e. M_0) during T_1 relaxation. The recovery process is exponential, governed by a time constant T_1, which is equal to the duration needed for M_z to grow from 0 to 63%M_0 after a 90° excitation (Fig. 2.3a).

The strength of MRI signal is proportional to the magnitude of M_z before excitation, thus it is necessary to insert an interval between successive excitations to allow recovery of M_z. This interval is known as TR (repetition time). As shown in Fig. 2.3a, a very short TR means rapid repetition of RF excitations, which leads to a weak M_z and thus low SNR data. Using a very long TR allows greater recovery of M_z and guarantees a strong MRI signal, but a longer acquisition time is also required, which reduces the scan efficiency. In practice, a moderate TR is usually used such that M_z can reach a constant magnitude, which is lower than M_0 but sufficient for imaging.

The second type of relaxation is spin-spin relaxation, which is the process by which the excited spins gradually lose phase coherence ('de-phasing') due to the different precessional frequencies of spins, which leads to a decay of transversal magnetization, M_{xy}.

Many sources contribute to the loss of phase coherence. Firstly, due to molecular interactions like motion and rotation, the precessional frequency of a spin fluctuates over time. As a result, a phase difference is gradually built up between spins. The decay of M_{xy} caused by this effect is called T_2 decay, which is exponential and governed by a time constant T_2, whose value is equal to the duration needed for M_{xy} to decay from M_0 (immediately after a 90 degree excitation) to ∼37%M_0 (Fig. 2.3b).

Spin de-phasing can also be influenced by $\mathbf{B}_0$ inhomogeneity, local field inhomogeneity due to susceptibility variations and a gradient field which is deliberately applied for spatial encoding (see the next section). With all these effects, the decay rate of M_{xy} is much faster than the T_2 decay. The overall decay rate of M_{xy} is

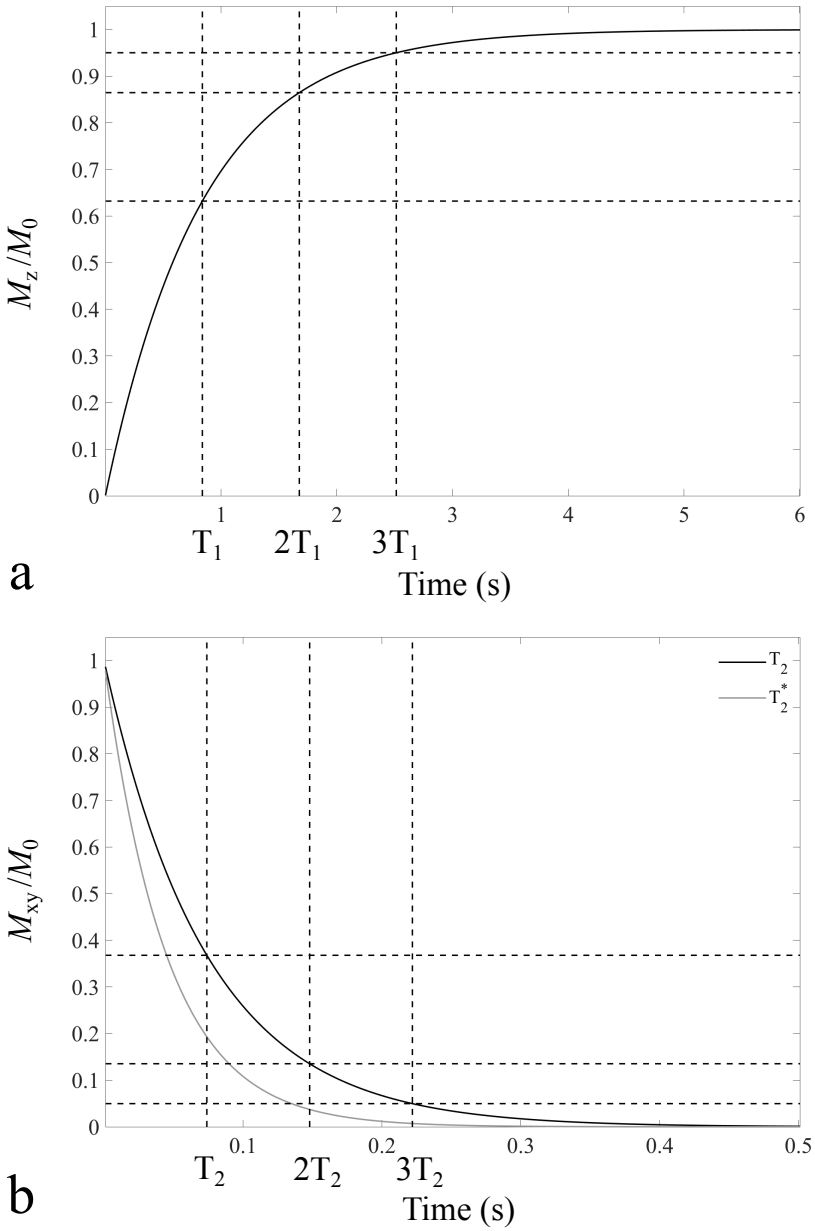


Figure 2.3: Spin relaxation after a 90° RF excitation. a) In the longitudinal axis, M_z is zero immediately after the excitation and then starts to recover towards the equilibrium (M_0). This is an exponential process with a time constant T_1. The recovered magnitude at time T_1 is $63\%M_0$. b) In the transverse plane, M_{xy} achieves its maximum magnitude right after the excitation and then starts to decay. T_2^* decay is more rapid than T_2 decay due to the spin de-phasing caused by field inhomogeneity, which can be reversed with a refocusing pulse. T_2 decay, however, is not reversible. The remaining transverse component at the time T_2 in a T_2 decay is $37\%M_0$. $T_1 = 838$ms, $T_2 = 74$ms and $T_2^* = 45$ms are used in the simulation, corresponding to white matter at 3T [16, 17].

characterized by another time constant $T_2^\star$:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{2.12}$$

where T_2' describes the signal decay induced by field inhomogeneity. As shown in Fig. 2.3b, FID signal undergoing T_2^* decay vanishes rapidly. An approach to mitigate this situation is to apply a refocusing pulse (e.g. 180°) after the excitation, which can reverse the de-phasing effects due to the field inhomogeneity (T_2') and form a spin echo which is weighted by a T_2 decay. Spin refocusing using a refocusing pulse is illustrated in Fig. 2.4. The formation time of the spin echo is called echo time (TE).

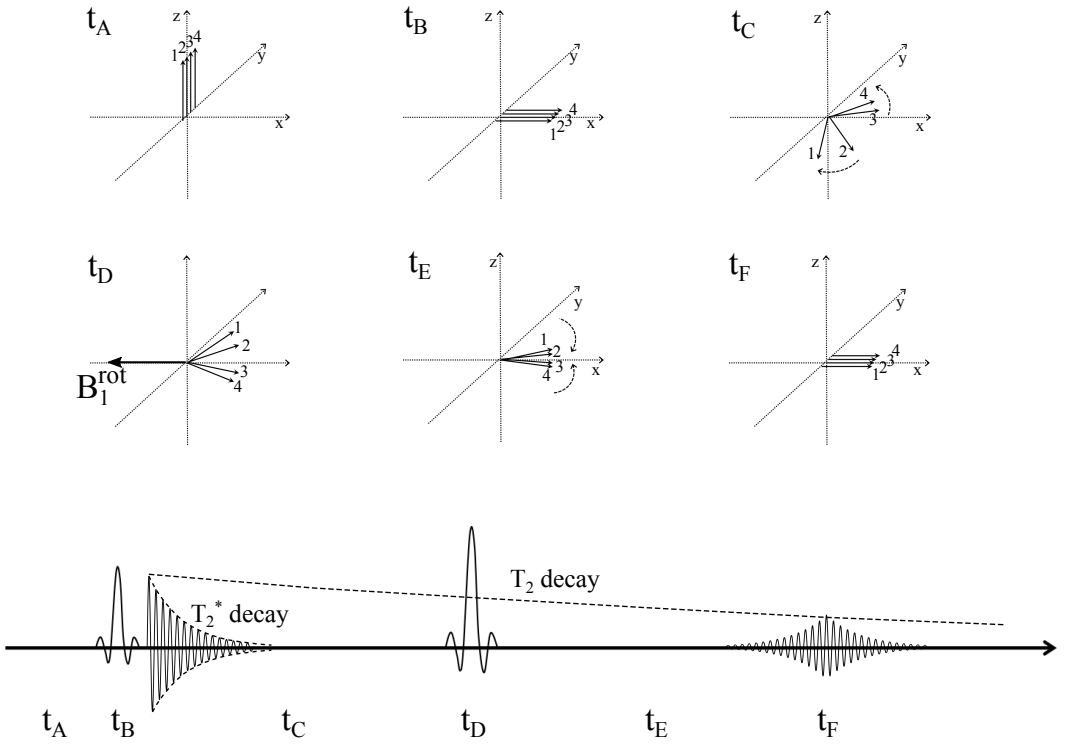


Figure 2.4: Schematic of spin echo formation. At the beginning (t_A), polarized spins are all along the longitudinal direction. A 90° excitation pulse flips the spins into the transverse plane at t_B. Due to the field inhomogeneity, each spin precesses at a different frequency, resulting in phase incoherence (t_C). A 180° refocusing pulse is applied at t_D, which flips the spin vectors about the x axis. The spin phase accumulated in the period between the excitation and refocusing is reversed. As the spin precession frequency is not changed, all spins start to refocus (t_E). At $t_F = t_D + (t_D - t_B)$, all spins are refocused, forming a spin echo. t_F is the echo time (TE). Signal acquired at the TE is weighted by a T_2 decay.

Steady state

If the same RF excitation and relaxation is repeated in a MRI scan, a steady state will be gradually established when the magnetization from one repetition is nearly identical to the next. The steady state signal is affected by many factors, such as field strength, tissue T_1, T_2 values, sequence parameters (e.g. RF flip angle, TR and TE). For different tissue types, signal amplitude can be manipulated by using different sequence parameters, leading to various tissue contrasts in MRI images. For instance, a short TR magnifies the contrast between tissues with different T_1, thus leading to a T_1-weighted image; a long TE enhances the contrast between tissues with different T_2, resulting in a T_2-weighted image; Using long TR and short TE at the same time can suppress the T_1 and T_2 contrasts, producing an image with proton density weight.

The steady state signal is also affected by the RF flip angle. If TR$\gg T_1$, a 90° pulse is a good choice as it enables strong MRI signal; However, if TR is close to or smaller than T_1, a 90° pulse might lead to a weak signal due to insufficient T_1 recovery. The flip angle that maximizes MRI signal is called the Ernst angle:

$$\alpha_E = \arccos(e^{-\mathrm{TR}/T_1}) \tag{2.13}$$

which is determined by the TR and tissue T_1. From this equation, $\alpha_E \approx 90°$ when TR$\gg T_1$; $\alpha_E \approx 68°$ when TR$= T_1$. Note, Eq. 2.13 is derived for sequences without a refocusing pulse. For spin echo sequence, the Ernst angle is given by $\alpha_E = \arccos(-e^{-\mathrm{TR}/T_1})$, which could be larger than 90°.

Bloch equation

The behaviour of the net magnetization $\mathbf{M}$ discussed above can be described using the Bloch equation [13]:

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \mathbf{M} \times \gamma\mathbf{B} - \frac{M_\mathrm{x}\hat{\mathbf{i}} + M_\mathrm{y}\hat{\mathbf{j}}}{T_2} - \frac{(M_\mathrm{z} - M_0)\hat{\mathbf{k}}}{T_1} \tag{2.14}$$

where $\mathbf{B}$ is the total magnetic field $\mathbf{M}$ observes, including the main field $\mathbf{B}_0$, the RF field $\mathbf{B}_1$ and the gradient field (see the next section). $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$ and $\hat{\mathbf{k}}$ are the unit vectors along x, y and z axes, respectively.

2.1.2 Image Formation

Gradient Field

In modern MRI systems, a spatially varying magnetic field known as gradient field is applied to achieve spatial encoding by registering spatial positions to different frequencies. The gradient field has the same direction as $\mathbf{B}_0$ but its magnitude varies linearly in space. When the gradient field is switched on, the total magnetic field becomes

$$\mathbf{B} = \mathbf{B}_0 + (\mathbf{G} \cdot \mathbf{r})\hat{\mathbf{k}} \tag{2.15}$$

where $\mathbf{G} = [G_\mathrm{x}, G_\mathrm{y}, G_\mathrm{z}]$ describes the amplitude of gradient ('slope') along three laboratory axes. According to Eq. 2.10, the Larmor frequency at location $\mathbf{r}$ becomes

$$\omega(\mathbf{r}) = \omega_0 + \gamma \mathbf{G} \cdot \mathbf{r} \tag{2.16}$$

which is a linear mapping between spatial position and larmor frequency (Fig. 2.5). This formulation enables the use of the Fourier transform to resolve voxel positions (see Eq. 2.22).

The gradient field is generated by a set of gradient coils. In modern MRI systems, there are typically three gradient coils, each generating a gradient field along one laboratory axis. On clinical scanners, the maximum gradient magnitude varies between 30 mT/m and 80 mT/m, corresponding to a bandwidth of 12.8 kHz and 34.1 kHz for a 1 cm spatial distance. The gradient field cannot reach its maximum amplitude immediately after being switched on. An interval is required for the gradient field to change from its current amplitude to an expected amplitude, which is determined by the slew rate of the gradient coil. Higher slew rate means

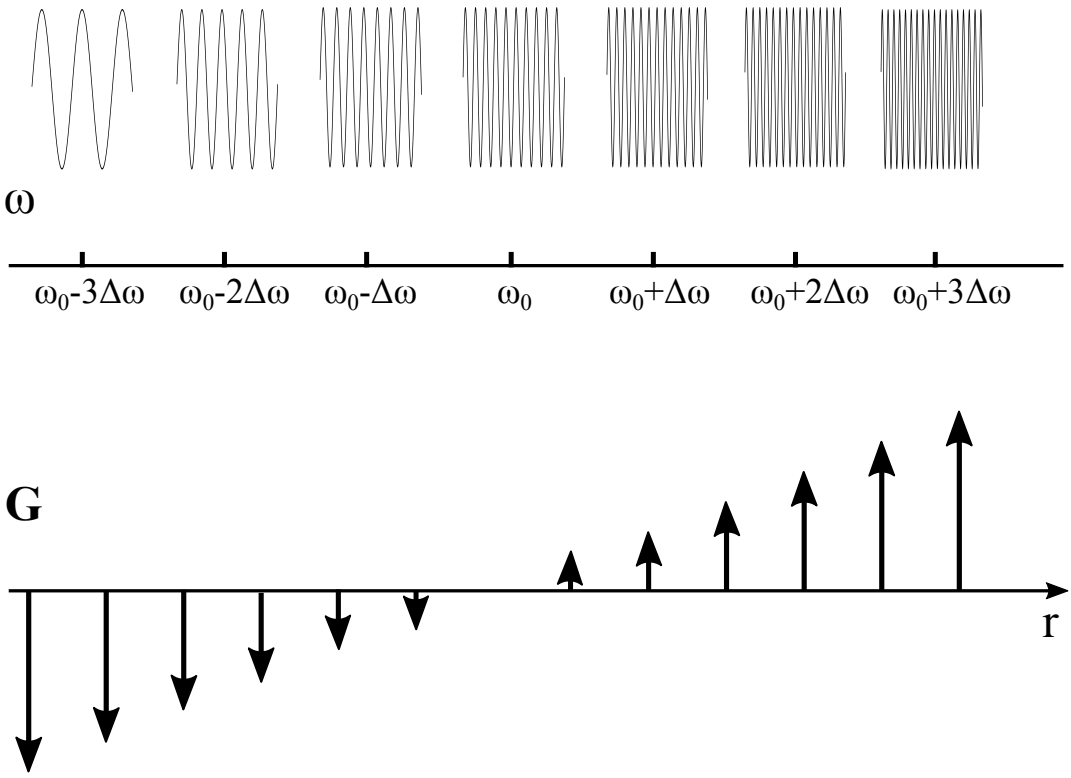


Figure 2.5: The gradient filed **G** is parallel to $\mathbf{B}_0$. Its magnitude varies linearly with spatial location **r**. With the gradient field, larmor frequency varies linearly in space with a slope governed by the gradient amplitude.

rapid gradient switching and thus fast acquisition speed. The highest slew rate available on a common clinical scanner is about 200mT/m/ms (require 0.2ms to reach 40mT/m gradient amplitude from zero). However, high slew rate can cause stronger eddy currents on the conducting surface in the magnet, which is associated with many MRI artefacts (e.g. image distortion). In addition, fast gradient switching might induce stimulation in subject's peripheral nerves. Therefore, the choice of slew rate in practice should be considered comprehensively instead of pushing to the hardware limit.

Selective Excitation

Volume selection is a basic requirement for diagnostic imaging. To selectively excite a slice or a slab, the RF pulse should be accompanied by a gradient field during excitation. This gradient field is known as a slice-selection gradient, which produces a

linear frequency ramp along the slice direction. The behaviour of $\mathbf{M}$ during selective excitation can be derived from the Bloch equation. For simplicity of notation, the relaxation effects are temporarily ignored, assuming a very short excitation. In the rotating frame, the Bloch equation for selective excitation is formulated as

$$\begin{bmatrix} \frac{\mathrm{d}M_{\mathrm{x}}^{\mathrm{rot}}(t)}{\mathrm{d}t} \\ \frac{\mathrm{d}M_{\mathrm{y}}^{\mathrm{rot}}(t)}{\mathrm{d}t} \\ \frac{\mathrm{d}M_{\mathrm{z}}^{\mathrm{rot}}(t)}{\mathrm{d}t} \end{bmatrix} = \begin{bmatrix} 0 & \gamma G_{\mathrm{z}} z & 0 \\ -\gamma G_{\mathrm{z}} z & 0 & \gamma B_1^{\mathrm{env}}(t) \\ 0 & -\gamma B_1^{\mathrm{env}}(t) & 0 \end{bmatrix} \begin{bmatrix} M_{\mathrm{x}}^{\mathrm{rot}}(t) \\ M_{\mathrm{y}}^{\mathrm{rot}}(t) \\ M_{\mathrm{z}}^{\mathrm{rot}}(t) \end{bmatrix} \tag{2.17}$$

where the RF pulse and the slice selection gradient are assumed along the x axis and z axis, respectively. $G_{\mathrm{z}}z$ is the magnetic field amplitude at position z. G_{z} is normally constant during excitation. Eq. 2.17 provides a direct approach to calculate the behaviour of the net magnetization during selective excitation. It is also the governing equation for RF pulse design.

In many cases, it is difficult to solve Eq. 2.17 analytically. An approximation is commonly used for small-flip-angle excitation, in which $M_{\mathrm{z}}^{\mathrm{rot}}$ is treated as a constant during excitation (i.e. $\mathrm{d}M_{\mathrm{z}}^{\mathrm{rot}}(t)/\mathrm{d}t = 0$) and an analytical solution for the transverse magnetization ($M_{\mathrm{xy}}^{\mathrm{rot}}(t) = M_{\mathrm{x}}^{\mathrm{rot}}(t) + iM_{\mathrm{y}}^{\mathrm{rot}}(t)$) can be derived:

$$M_{\mathrm{xy}}^{\mathrm{rot}}(t_{\mathrm{RF}}) = i\gamma M_0 \exp\left(-i\gamma z G_{\mathrm{z}} t_{\mathrm{RF}}\right) \int_0^{t_{\mathrm{RF}}} B_1^{\mathrm{env}}(t) \exp\left(-i\gamma z G_{\mathrm{z}} t\right) \mathrm{d}t \tag{2.18}$$

where an equilibrium initial condition ($M_{\mathrm{xy}}^{\mathrm{rot}}(0) = 0$, $M_{\mathrm{z}}^{\mathrm{rot}}(0) = M_0$) is used. t_{RF} denotes the time immediately after RF excitation. The solution in Eq. 2.18 actually describes a Fourier transform

$$M_{\mathrm{xy}}^{\mathrm{rot}}(t_{\mathrm{RF}}) = i\gamma M_0 \exp\left(i\gamma z G_{\mathrm{z}} t_{\mathrm{RF}}/2\right) \mathcal{F}^{-1}(B_1^{\mathrm{env}}(t)) \tag{2.19}$$

with the frequency defined as $f = (\gamma/2\pi) z G_{\mathrm{z}}$. The excitation profile is determined by the inverse Fourier transform of the pulse envelope $\mathcal{F}^{-1}(B_1^{\mathrm{env}}(t))$. As shown in Fig. 2.6, the slice thickness can be derived from the bandwidth of the RF pulse (typically defined as the full width at half maximum of the frequency profile) using the linear

relationship between position and frequency. Note in Eq. 2.19 there's a phase dispersion across the slice represented by the linear phase term $\exp(i\gamma z G_z t_{RF}/2)$, which must be removed to avoid signal cancellation. This is usually realized by adding a slice re-phasing gradient immediately after the slice selection gradient.

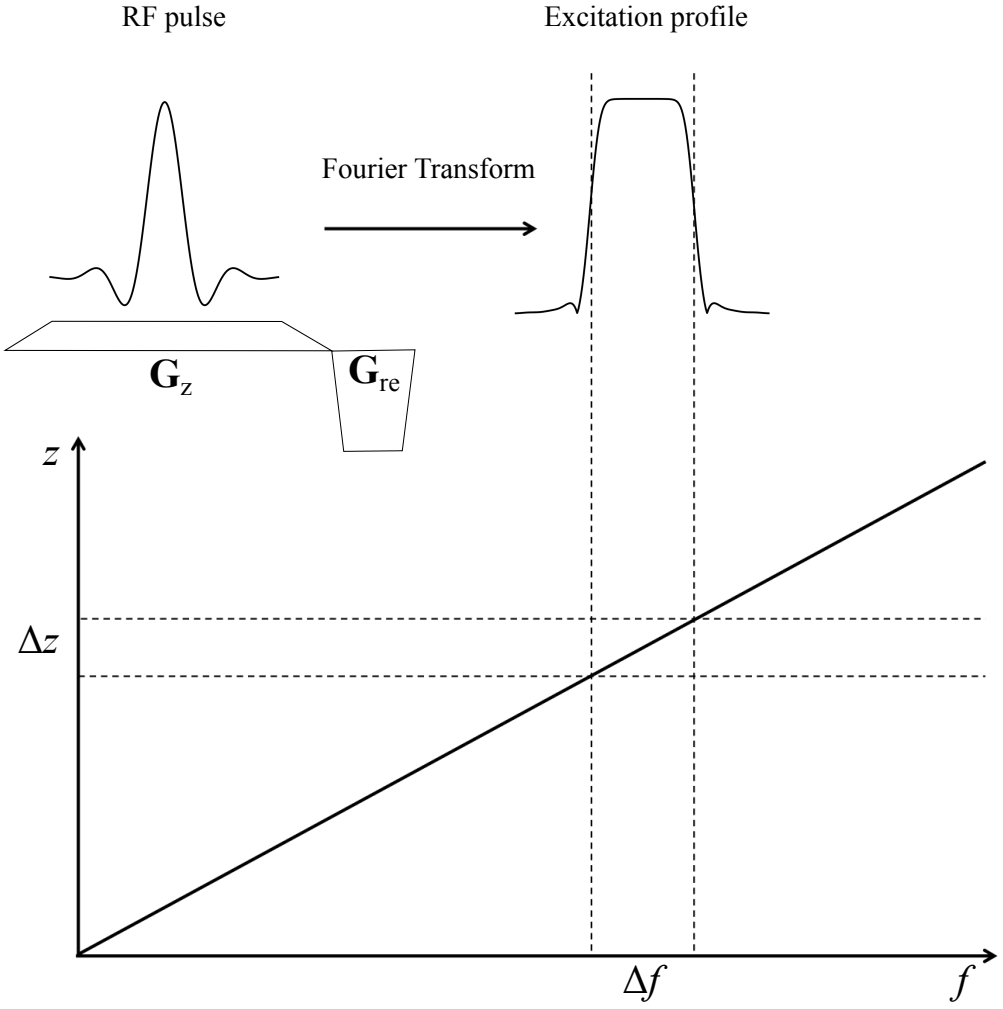


Figure 2.6: Selective excitation using an RF pulse accompanied by a slice-selection gradient $\mathbf{G}_z$. A re-phasing gradient $\mathbf{G}_{re}$ is added after $\mathbf{G}_z$ to reverse the phase dispersion. The excitation profile can be approximated by a Fourier transform of the envelop of the RF pulse, provided the flip angle is small. The slice thickness is linearly related to the bandwidth of the RF pulse as $\Delta f = (\gamma/2\pi)G_z\Delta z$.

The Fourier relationship defined in Eq. 2.19 is derived based on a small-flip-angle approximation. For large flip angle, the RF profile is usually calculated using a numerical solution of the Bloch equation. In practice, the Fourier method is sufficiently good for a flip angle as large as 90°.

A rectangular RF profile is always desirable for excitation, but according to Fourier analysis it is unlikely to happen, as it requires an infinitely long RF field. In practice, a short RF pulse is used, which inevitably introduces truncation artefacts, such as a non-flat profile within the slice, side-lobes outside of the slice, transition bands at slice edges and Gibbs ringing.

Time bandwidth product (TBWP) is a dimensionless parameter used to measure RF pulse selectivity. The TBWP of a windowed SINC pulse is equal to the total number of zero-crossings, and the pulse profile is fully specified by a given TBWP. If a higher flip angle (higher transmission power at constant pulse duration) is applied, the pulse can be played out longer to reduce power requirement, leading to a lower bandwidth. For two RF pulses with the same duration, the one with higher TBWP can produce a shaper excitation profile than the other. However, high-TBWP RF pulse requires high transmission power and can induce a high specific absorption rate (SAR) value. Therefore, the maximum TBWP allowed for subject scans is usually limited.

k-space

A modern MRI system utilizes two spatial encoding schemes, frequency encoding and phase encoding. The frequency encoding is accomplished by applying a gradient field during data readout, which registers the spatial locations along the gradient direction to different frequencies. The gradient used for frequency encoding is usually called 'readout gradient' and its direction is called 'readout direction'. Fourier transformation of the frequency encoded signal can resolve the spatial information along the readout direction, but it cannot separate voxels along the perpendicular directions. Phase encoding solves this problem with another gradient field ('phase encoding gradient') perpendicular to the readout gradient, which is typically applied for a short interval before the data readout. During phase

encoding, spins accumulate some initial phase, which is linearly related to the spatial location in the phase encoding direction.

With frequency and phase encoding, the MR signal acquired at a time point t can be expressed as an integral:

$$S(t) \propto \iint \rho(\mathbf{r}) e^{-i\gamma\left(B_0 t + \mathbf{r}\cdot\int_0^t \mathbf{G}(\tau)\mathrm{d}\tau\right)} \mathrm{d}x\,\mathrm{d}y \tag{2.20}$$

where a 2D acquisition is assumed here. ρ is the spin density, $\mathbf{r} = [x, y]$ is the spatial vector. Relaxation effects are ignored here for simplicity of notation. k-space is defined based on the 0^{th} order moment of the gradient

$$\mathbf{k}(t) = \frac{\gamma}{2\pi}\int_0^t \mathbf{G}(\tau)\mathrm{d}\tau \tag{2.21}$$

which is a spatial frequency space that $\mathbf{k}$ resides in. The carrier frequency (γB_0) is removed after signal demodulation. Insert Eq. 2.21 into Eq. 2.20 and express $S(t)$ using k-space coordinate $S(\mathbf{k})$, which gives

$$S(\mathbf{k}) \propto \iint \rho(\mathbf{r}) e^{-i2\pi\,\mathbf{r}\cdot\mathbf{k}(t)} \mathrm{d}x\,\mathrm{d}y \tag{2.22}$$

which describes a 2D Fourier relationship between the proton density map and the k-space formulation of the MRI signal ('k-space data'). Therefore, MRI images can be reconstructed with a 2D Fourier transform of the k-space data. Fourier transform has a highly efficient implementation (Fast Fourier transform, 'FFT'), which enables fast image reconstruction that is critical in clinical applications.

Many imaging properties can be derived from the Fourier relationship between the image space and k-space. According to Nyquist sampling criteria, if the sampling rate in k-space is dense enough, an object with limited size can be reconstructed without aliasing. For instance, let the object size along one dimension be L, the k-space sampling rate along that dimension should satisfy $\Delta k \leq 1/L$, otherwise aliasing (or 'foldover') artefacts will emerge. Therefore, Δk determines the FOV

of the acquisition. Any signal from outside of the FOV will be folded into the target FOV, leading to aliasing artefacts.

The spatial resolution of the reconstructed image is determined by the k-space coverage. For example, the spatial resolution in the x dimension is determined by $\Delta x = 1/|2k_{\mathrm{x}}^{\mathrm{max}}|$ where $k_{\mathrm{x}}^{\mathrm{max}}$ is the maximum k-space value along k_x. To achieve higher spatial resolution, a larger range of k-space should be sampled.

k-space can be sampled using different trajectories, which is accomplished by changing the gradient waveforms. Four commonly used trajectories are shown in Fig. 2.7. The Cartesian raster trajectory [18] (Fig. 2.7a) acquire a single k-space line after each excitation. Phase encoding is applied before data readout to move the trajectory to the start of the next readout line. Echo planar imaging (EPI) trajectory [19] (Fig. 2.7b) is a accelerated version of Cartesian raster trajectory, where multiple k-space lines are acquired consecutively following an excitation. Polarity of the readout gradient is flipped between successive readout lines. Phase-encoding gradients are applied as blips between readouts, which moves the trajectory along the phase-encoding direction.

Non-Cartesian trajectories can also be used to sample k-space. Unlike Cartesian trajectory, where frequency- and phase-encoding are independent functions, a non-Cartesian trajectory is usually designed with a 2D readout function. Two widely used non-Cartesian trajectories are Radial [20] and Spiral [21, 22]. The radial (Fig. 2.7c) trajectory is realised by rotating the readout lines about the k-space origin with different angles, such that k-space can be densely sampled to avoid aliasing. Each sampled line from the radial acquisition represents a projection of the image along the normal direction. A spiral trajectory (Fig. 2.7d) usually starts from the k-space centre and moves towards the outer k-space while rotating about the k-space origin. To achieve a spiral trajectory, two readout gradients with sinusoidal waveforms should be applied simultaneously along two orthogonal axes.

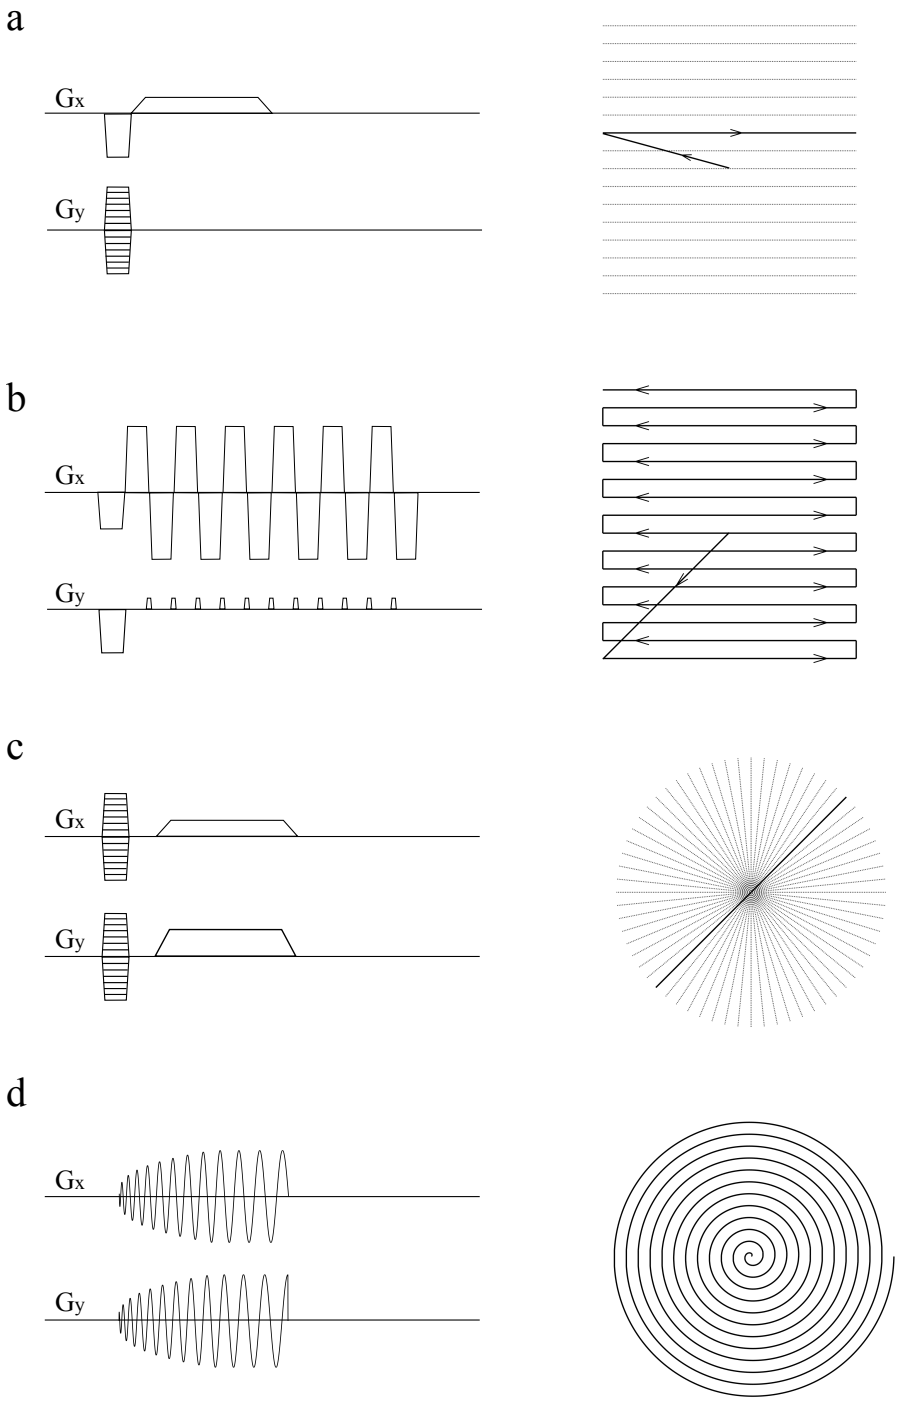


Figure 2.7: k-space trajectories (right) and gradient encoding schemes (left). a) basic Cartesian raster trajectory. b) EPI trajectory. c) Radial trajectory. d) Spiral.

The Cartesian trajectory is prevalent in MRI acquisition. The main reason is that artefacts with non-Cartesian trajectories are usually difficult to identify. Moreover, the gradient coil design on current commercial scanner is not optimal for sinusoidal gradients, which might cause inaccuracy in trajectory configuration, leading to artefacts in the image. Nonetheless, non-Cartesian trajectories have good features like fast sampling, robustness to motion and incoherent artefacts, which have been successfully explored in cardiac imaging and compressed sensing MRI.

2.2 Diffusion MRI

Diffusion is a spontaneous process in nature resulting from random collisions between liquid or gas molecules, which leads to molecular mixing without requiring bulk motion. Fick's laws [23] provide a description of diffusion in physics, which relate the diffusive flux linearly with the concentration gradient of the medium, and shows that rate of change of concentration is proportional to its second spatial derivative. The linear coefficient defined in Fick's laws is called the diffusion coefficient, which is an intrinsic property of the medium affected by many factors such as temperature, microenvironment and molecular size. According to Fick's laws, a prerequisite for diffusion to happen is the existence of a concentration difference in the medium, in other words, when the concentration difference vanishes, diffusion should stop. However, at the microscopic level, collisions between molecules always persist, even at thermodynamic equilibrium when flux 'seems' to disappear. This was first reported by Robert Brown when he observed random motions of pollen grains under a microscope, which is now known as 'Brownian motion', a result caused by the random collisions of water molecules. The two pictures of diffusion, depicted by Fick's laws and Brownian motion respectively, were unified in the early twentieth century by Albert Einstein [24], who developed a statistical framework to describe particles' diffusive motion by bridging the mean-squared displacement of a particle

with the diffusion coefficient defined in Fick's law:

$$\left\langle x^2 \right\rangle = 2D\Delta \tag{2.23}$$

where $\langle x^2 \rangle$ is the mean-squared displacement of a particle, D is the diffusion coefficient and Δ is the molecule diffusion time.

2.2.1 Diffusion MRI Signal

Diffusion effects in MR were first reported by Edwin Hahn [25] in his spin echo experiment, when he noticed that the measured echo signal was lower than expected and attributed it to spin de-phasing due to diffusive motion in an inhomogeneous magnetic field. Hahn mentioned the diffusion effects could be measurable, although it was Carr and Purcell [26] who later demonstrated spin diffusion could be directly measured from a spin echo signal. They added a constant magnetic gradient field across the spin echo sequence to amplify the de-phasing effects of spin diffusion. This gradient is commonly called a diffusion encoding gradient. They also derived a mathematical framework based on random walk theory to calculate the diffusion coefficient based on spin echo measurements. In 1956, Torrey revised the Bloch equation by adding a diffusion term, creating the Bloch-Torrey equation

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \mathbf{M} \times \gamma\mathbf{B} - \frac{M_{\mathrm{x}}\hat{\mathbf{i}} + M_{\mathrm{y}}\hat{\mathbf{j}}}{T_2} - \frac{(M_{\mathrm{z}} - M_0)\hat{\mathbf{k}}}{T_1} + D\nabla^2\mathbf{M} \tag{2.24}$$

which gives the same solution as derived by Carr and Purcell using random walk theory. Consider diffusion in an isotropic medium, the amplitude of transverse magnetization at the echo time can be expressed as

$$M_{xy}(\mathrm{TE}) = M_{xy}(0)e^{-\frac{\mathrm{TE}}{T_2}}e^{-bD} \tag{2.25}$$

where $M_{xy}(0)$ and $M_{xy}(\mathrm{TE})$ are the transverse magnetizations immediately after the 90° excitation pulse and at the echo time, respectively. The signal attenuation

caused by spin diffusion is determined by the decay term e^{-bD}, where the value of b is given by

$$b = \gamma^2 \int_0^{\mathrm{TE}} \left(\int_0^t |\mathbf{G}(\tau)| \mathrm{d}\tau \right)^2 \mathrm{d}t \tag{2.26}$$

where $\mathbf{G}$ includes both diffusion-encoding gradients and imaging gradients, although the latter is usually ignored in b value calculation due to its limited contribution to the diffusion-related signal attenuation.

2.2.2 Diffusion MRI Sequence

The most widely used diffusion preparation approach is the pulsed-gradient sequence (Fig. 2.8) proposed by Stejskal and Tanner in 1965 [27], which consists of a pair of diffusion gradient pulses with one applied between the 90° excitation and the 180° refocusing pulses and the other applied between the 180° RF pulse and the spin echo, respectively. For stationary spins, the de-phasing effects induced by the two gradient pulses cancel each other, leaving the echo signal untouched (apart from relaxation effects). For diffusive spins, however, the de-phasing caused by the first diffusion gradient pulse cannot be reversed by the second one due to the random nature of molecular diffusion. The b value of the Stejskal-Tanner preparation is given by

$$b = \gamma^2 |\mathbf{G}|^2 \delta^2 \left(\Delta - \frac{\delta}{3} \right) \tag{2.27}$$

here, the ramp-up and ramp-down time of diffusion encoding gradients are ignored.

2.2.3 Diffusion MRI Measurement

Diffusion MRI (dMRI) is a particular type of MRI method used to measure water diffusion in biological tissues. Water diffusion within biological tissue is restricted by cellular and sub-celluar structures (e.g. membranes, macromolecules), thus dMRI signal can be potentially used as a biomarker for tissue microstructure. The diffusion coefficient D can be estimated using dMRI, which requires at least

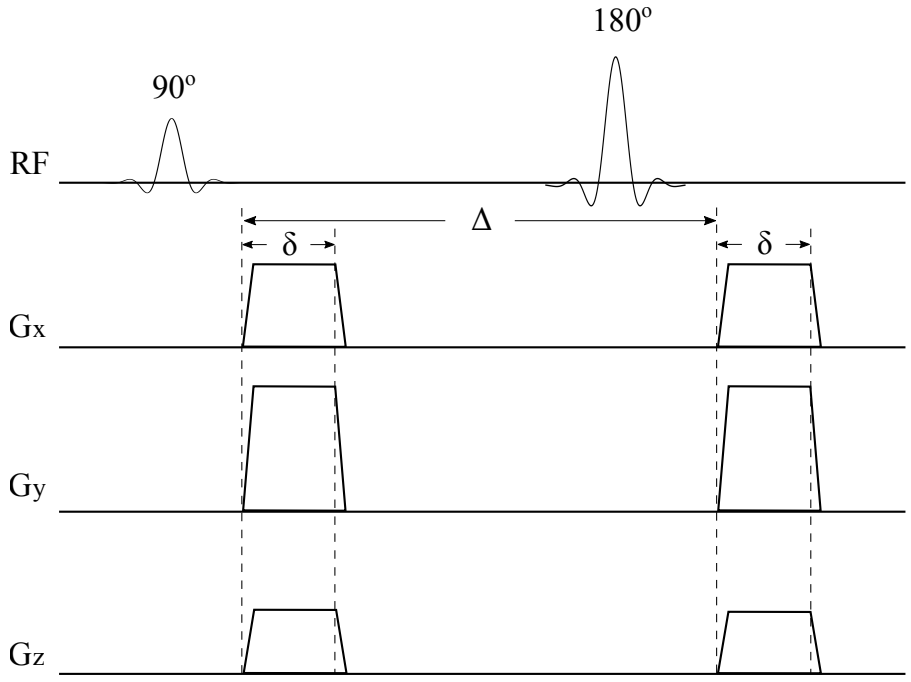


Figure 2.8: Diffusion preparation using the Stejskal-Tanner sequence. Note the scaling of gradient amplitude is arbitrary here, which will change on each axis depending on the desired diffusion encoding and weighting.

two measurements using different b values. Conventionally, a b=0 (T_2 weighted) volume $S(0)$ and a b>0 (T_2 and diffusion weighted) volume $S(b)$ are acquired. The diffusion coefficient can be calculated as

$$D = \frac{1}{b} \ln \frac{S(0)}{S(b)} \tag{2.28}$$

As the calculated diffusion coefficient is usually influenced by many factors such as restriction, perfusion and partial volumes effects, it is commonly called apparent diffusion coefficient (ADC).

Isotropic water diffusion assumes the restrictions imposed by tissue microstructures have a spherical symmetry, which is usually not the situation in biological tissues. Water molecules often diffuse more easily along one direction than the other, which is a phenomenon known as diffusion anisotropy. The scalar diffusion coefficient only captures the mean diffusivity, which cannot explain the directional difference of water diffusion. Therefore, a second-order tensor was proposed to

model the diffusion anisotropy [27–29]:

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \tag{2.29}$$

which describes the covariance of diffusion-induced displacements in three-dimensional space. The diagonal elements of a diffusion tensor represent the variances along x, y and z axes, respectively. The off-diagonal elements denote the covariances, which are symmetric ($D_{xy} = D_{yx},\ D_{xz} = D_{zx},\ D_{yz} = D_{zy}$) .

The diffusion tensor can be represented as a ellipsoid, which depicts the equal-probability contour of the diffusion displacement. The axes of the tensor ellipsoid are determined by the eigenvectors of the tensor ($\mathbf{v}_1$,$\mathbf{v}_2$,$\mathbf{v}_3$) and the semi-axes are determined by the eigenvalues (λ_1,λ_2,λ_3).

Diffusion tensor imaging (DTI) is a MRI technique measuring diffusion anisotropy based on the tensor model. Bloch-Torrey equation with anisotropic diffusion is expressed as

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \mathbf{M} \times \gamma\mathbf{B} - \frac{M_\mathrm{x}\hat{\mathbf{i}} + M_\mathrm{y}\hat{\mathbf{j}}}{T_2} - \frac{(M_\mathrm{z} - M_0)\hat{\mathbf{k}}}{T_1} + \nabla \cdot \mathbf{D}\nabla\mathbf{M} \tag{2.30}$$

the echo signal is given by

$$S(\mathrm{TE}) = S(0)e^{-\frac{\mathrm{TE}}{T_2}}e^{-b\,\mathbf{g}^\mathrm{T}\mathbf{D}\mathbf{g}} \tag{2.31}$$

where b value is the same as in Eq. 2.26 and $\mathbf{g}$ is the unit directional vector of the diffusion gradient $\mathbf{G}$. To calculate the diffusion tensor, at least six diffusion weighted images with non-coplanar diffusion encoding directions and one set of b=0 image are required. Once the diffusion tensor is obtained, some quantitative metrics can be derived:

$$\mathrm{MD} = \frac{1}{3}\sum\nolimits_i \lambda_i$$

$$\mathrm{FA} = \sqrt{\frac{3}{2}}\frac{\sqrt{\sum\nolimits_i (\lambda_i - \mathrm{MD})^2}}{\sqrt{\sum\nolimits_i {\lambda_i}^2}}$$

MD is the mean diffusivity. The fractional anisotropy (FA) is derived based on the standard deviation of the eigenvalues, which measure the anisotropy of water diffusion. If the water diffusion along three axes is similar, FA will be close to zero; if the diffusion is significantly higher along one direction than the other two, FA will be close to one. Both MD and FA are rotationally invariant. It should be noted that FA cannot specify the tensor shape, as different combinations of eigenvalues can give the same FA. But in practice, FA has shown high sensitivities under some pathological conditions and is widely used as a biomarker in clinical applications.

The eigenvectors of the diffusion tensor determine the orientation of the tensor ellipsoid. A principal diffusion direction (PDD) can be defined as the direction of the eigenvector corresponding to the largest eigenvalue, which is the most elongated direction of the tensor, or equivalently the direction with strongest diffusivity. One very useful application of PDD is fibre tractography, where the PDD is used as an estimation of fibre orientation in each voxel and structural connections between different brain regions are constructed by tracking the white matter fibre pathways.

2.3 Partially Parallel Imaging

Conventional MRI scanning is much slower than other common imaging modalities like CT and ultrasound, due to the formation of k-space which usually requires multiple excitations and an interval between excitations for the recovery of longitudinal magnetization. Single-shot approaches like single-shot EPI and single-shot spiral are fairly fast, allowing 'snapshot' acquisition ($<$ 100ms per image). However, the spatial resolution of single-shot imaging is typically low, restricted by the tissue T_2^* decay. Also, single-shot scanning is more vulnerable to hardware imperfections and local field inhomogeneity caused by tissue susceptibility variations, which results in more severe artefacts compared with multi-shot scanning.

Partially parallel imaging (PPI) is a class of approaches to accelerating MRI scanning by acquiring only a subset of k-space. Due to the Fourier relationship between image space and k-space, direct Fourier transformation of the under-sampled k-space data will result in aliasing in the image. PPI acquisitions typically use multiple receiver coils with each coil having localized RF sensitivity. The underlying image is reconstructed from the under-sampled multi-coil data using special reconstruction algorithms. Two widely used PPI reconstruction methods are sensitivity encoding (SENSE) [30] and generalized autocalibrating partially parallel acquisitions (GRAPPA) [31].

2.3.1 SENSE

SENSE reconstruction is an image domain method, which is based on a linear signal model:

$$\mathbf{MFSu} = \mathbf{d} \tag{2.32}$$

where $\mathbf{u}$ is a $N_i \times 1$ image vector, $N_i = N_x N_y$ is the total number of voxels in the image. $\mathbf{S}$ denotes the coil sensitivities for N_c channels, which is a sparse matrix with size $N_i N_c \times N_i$. $\mathbf{F}$ represents the Fourier transform operator, which is a sparse diagonal matrix consisting of N_c identical sub-matrices. Each sub-matrix is a $N_i \times N_i$ Fourier transform matrix for a single-channel image. $\mathbf{M}$ denotes k-space sampling, which is also a sparse diagonal matrix ($N_k N_c \times N_i N_c$) with only binary elements representing k-space sampling (one) and skipping operation (zero), respectively. $\mathbf{d}$ is the acquired k-space data, which has the size of $N_k N_c \times 1$. Structures of these matrices and vectors are illustrated in Fig. 2.9.

SENSE reconstruction estimates the unknown image by minimizing the mean squared errors between the model prediction and the acquired data:

$$\mathbf{u} = \arg\min_{\mathbf{u}} \left\| \mathbf{MFSu} - \mathbf{d} \right\|_2^2 \tag{2.33}$$

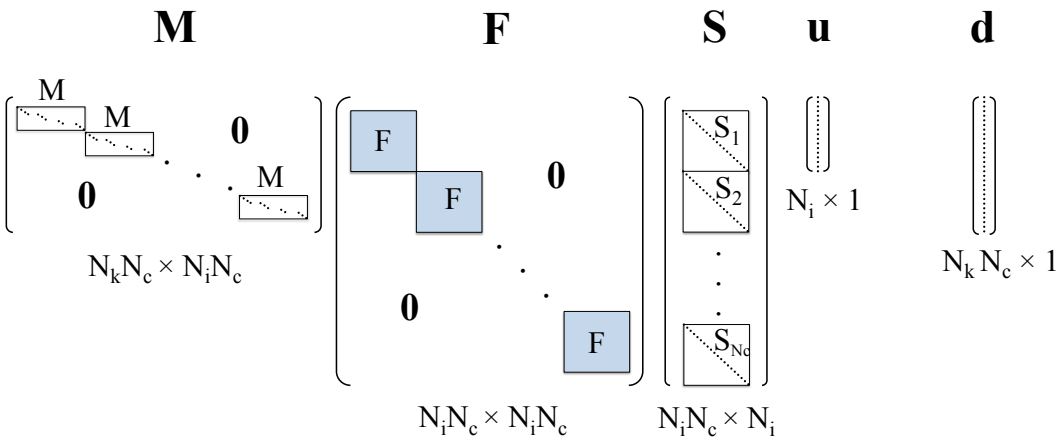


Figure 2.9: Illustration of matrix and vector structures in SENSE formulation.

which has a closed-form solution as $\mathbf{u} = \mathbf{A}^{\dagger}\mathbf{d}$ (define $\mathbf{A} = \mathbf{MFS}$ for simplicity of notation). $\mathbf{A}^{\dagger}$ is the pseudoinverse of $\mathbf{A}$ given by $(\mathbf{A}^{\mathrm{H}}\mathbf{A})^{-1}\mathbf{A}^{\mathrm{H}}$. H is the transpose conjugate.

For a Cartesian trajectory with uniform under-sampling, SENSE reconstruction has a fast implementation. Consider an equivalent form of Eq. 2.32, where $(\mathbf{MF})^{\mathrm{H}}$ is multiplied to the left-hand and right-hand sides simultaneously:

$$(\mathbf{MF})^{\mathrm{H}}\mathbf{MFSu} = (\mathbf{MF})^{\mathrm{H}}\mathbf{d} \tag{2.34}$$

The right-hand side is a direct Fourier transform of the under-sampled k-space data. On the left-hand side, $(\mathbf{MF})^{\mathrm{H}}\mathbf{MF}$ represents an aliasing operator, which folds different voxels into one position. For a Cartesian trajectory with an under-sampling factor of R, the effect of the aliasing operator can be incorporated into a revised reconstruction formulation (Fig. 2.10a):

$$\widetilde{\mathbf{S}}_p\mathbf{u}_p = \mathbf{I}_p \tag{2.35}$$

where $\mathbf{I}_p$ is a $N_c \times 1$ vector representing the p^{th} voxel in the aliased image $\mathbf{I}$, which is calculated by a direct Fourier transform of the under-sampled k-space data. $\mathbf{u}_p$ $(R \times 1)$ represents the R voxels which collapse to one voxel located at position p. $\widetilde{\mathbf{S}}_p$ $(N_c \times R)$ is the coil sensitivity matrix for the R voxels. With this formulation,

SENSE reconstruction can be applied to each set of R aliased voxels separately as $\mathbf{u}_p = (\widetilde{\mathbf{S}}_p^{\mathrm{H}}\widetilde{\mathbf{S}}_p)^{-1}\widetilde{\mathbf{S}}_p^{\mathrm{H}}\mathbf{I}_p$, which can be accelerated using parallel computing. To remove the noise correlation between channels, a pre-whitening matrix can be incorporated in the reconstruction as $\mathbf{u}_p = (\widetilde{\mathbf{S}}_p^{\mathrm{H}}\mathbf{\Phi}_p^{-1}\widetilde{\mathbf{S}}_p)^{-1}\widetilde{\mathbf{S}}_p^{\mathrm{H}}\mathbf{\Phi}_p^{-1}\mathbf{I}_p$, where $\mathbf{\Phi}$ is the noise covariance matrix which can be measured in a pre-scan.

A major challenge in parallel imaging is the reduced SNR, which is caused by the k-space under-sampling that results in a SNR reduction of $\sqrt{R}$, and the spatial correlation between coil sensitivities that leads to spatially varying noise amplification. The latter can be quantified using a g-factor (Fig. 2.10b). For Cartesian trajectory with uniform under-sampling, g-factor for SENSE reconstruction is calculated as

$$\mathrm{g}_n = \frac{1}{\sqrt{R}}\frac{\mathrm{SNR}_{\mathrm{full}}}{\mathrm{SNR}_{\mathrm{under}}} = \sqrt{(\widetilde{\mathbf{S}}_p^{\mathrm{H}}\mathbf{\Phi}_p^{-1}\widetilde{\mathbf{S}}_p)_{n,n}^{-1}(\widetilde{\mathbf{S}}_p^{\mathrm{H}}\mathbf{\Phi}_p^{-1}\widetilde{\mathbf{S}}_p)_{n,n}} \tag{2.36}$$

where g_n is g-factor for the n^{th} voxel in $\mathbf{u}_p$ $(R \times 1)$. A larger g-factor corresponds to stronger noise amplification.

The g-factor reflects the conditioning of the reconstruction problem. If the column vectors of $\widetilde{\mathbf{S}}$ are highly independent, the reconstruction problem will be well-conditioned, which is associated with a small g-factor; on the contrary, if the column vectors of $\widetilde{\mathbf{S}}$ are highly correlated, the reconstruction will be ill-conditioned, leading to a large g-factor and thus substantial noise amplification.

Generally, the g-factor related noise amplification can be alleviated using more receive coils and/or reducing the under-sampling factors. In addition, optimization of the under-sampling pattern can also be used to reduce g-factor, which has been widely used in simultaneous multi-slice acquisition.

For non-uniform under-sampling or non-Cartesian trajectory, the aliasing pattern might include a large number voxels and solving Eq. 2.35 is not efficient, thus iterative methods (e.g. conjugate gradient) are typically used. SENSE reconstruction requires prior information of coil sensitivity maps, which are usually measured from a calibration scan.

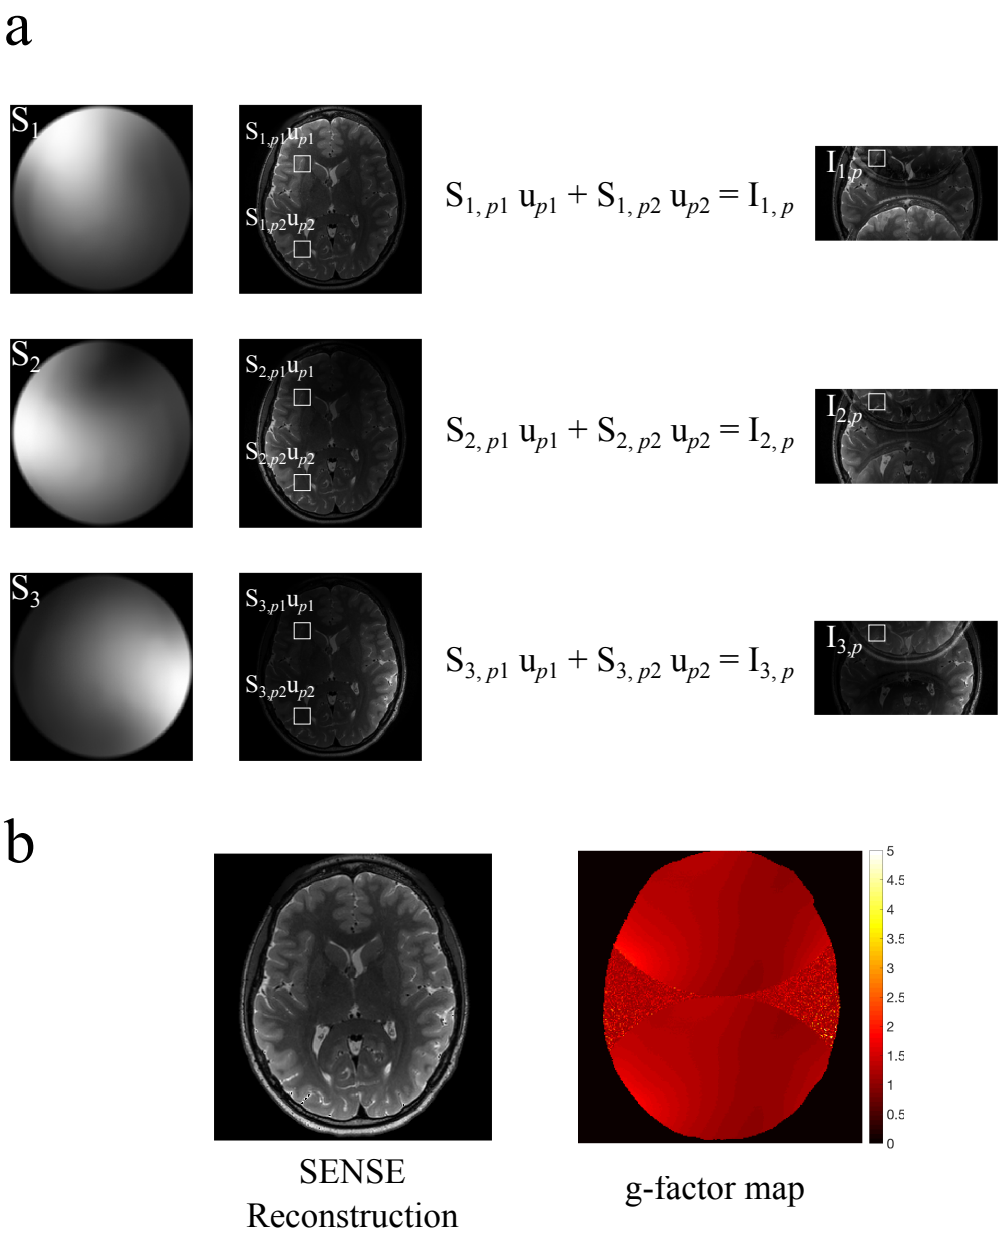


Figure 2.10: Demonstration of SENSE reconstruction for R=2 data using 3 receive coils. a) Formulation of SENSE reconstruction. Sensitivity maps measured from a phantom are shown on the left. For each coil, a linear equation can be constructed based on the SENSE model. u_{p1} and u_{p2} are two voxels from the original image, which collapsed together due to k-space under-sampling, forming the voxel I_p in the aliased image. $S_{i,v}$ is the coil sensitivity of the i^{th} channel at voxel v. b) Image reconstructed using SENSE (left), where aliasing is effecitively removed. The g-factor map is shown on the right, representing the noise amplification due to correlations between sensitivity maps.

2.3.2 GRAPPA

GRAPPA is a k-space based reconstruction method, which fills in missing k-space points for a target coil using data acquired at neighbouring points from all coils. A diagram of GRAPPA reconstruction is illustrated in Fig. 2.11. The reconstruction can be mathematically described as

$$d_i(k_x, k_y + m\Delta k_y) = \sum_{j=1}^{N_c} \sum_{l_x=-b_x}^{b_x} \sum_{l_y=-b_y}^{b_y} d_j(k_x + l_x, k_y + Rl_y\Delta k_y) w_m(l_x, l_y, j, i) \quad (2.37)$$

d_i and d_j are k-space data from the i^{th} and the j^{th}coil, R is the k-space under-sampling factor. $k_y + m\Delta k_y$ is the missing line which is $m\Delta k_y$ $(1 \leq m < R)$ away from an acquired line at k_y. Δk_y is the phase-encoding step. The GRAPPA kernel has the size of $(2b_x + 1) \times (2b_y + 1) \times N_c$, corresponding to a block in k-space with the size of $(2b_x + 1) \times (2b_y + 1) \times R \times N_c$. The acquired data within the kernel block from all channels are utilized to predict the missing data. w_m represents the GRAPPA kernel weight for the missing line that is $m\Delta k_y$ away from the nearby acquired line, resulting in m different sets of kernel weights in total. w_m should be calculated before GRAPPA reconstruction (kernel calibration). GRAPPA kernel weights are usually calculated from auto-calibration signal (ACS) data, which are typically fully sampled data acquired from the k-space centre.

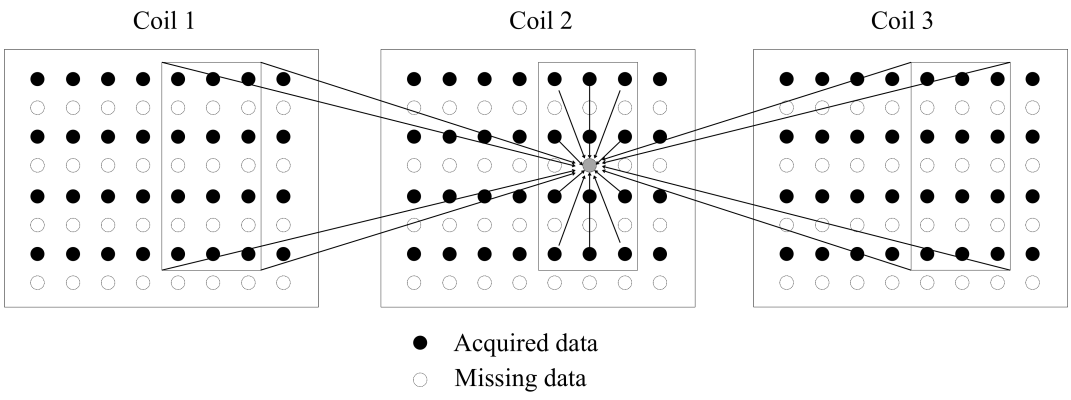


Figure 2.11: Illustration of GRAPPA reconstruction.

Unlike SENSE, GRAPPA is a coil-by-coil approach, which produces an individual image for each coil. The final image can be produced by combining all channel images using sum-of-squares or SENSE. A g-factor map, equivalent to that derived above, can also be calculated for GRAPPA [32].

3

Image Formation in Diffusion MRI

Contents

3.1 Introduction 39
3.2 Improving Image Fidelity 41
3.2.1 SSH-EPI Artefacts 41
3.2.2 Multi-shot Acquisition 45
3.2.3 Reduced FOV Method 53
3.3 Accelerating Diffusion MRI Acquisitions 55
3.3.1 Simultaneous Multi-slice Imaging 55
3.3.2 Fast Diffusion MRI with Compressed Sensing 59
3.4 Increasing Image SNR 61
3.4.1 Three-dimensional Acquisition 62
3.4.2 Ultra-high Field 64
3.5 Summary 65

This Chapter is based on a review paper (peer-reviewed) we published in the Journal of Magnetic Resonance Imaging [6].

3.1 Introduction

Single-shot echo planar imaging (SSH-EPI) has been used as the standard image formation method for diffusion MRI [33] on commercial scanners for more than

twenty years. This is mostly due to the fast acquisition speed of SSH-EPI (100 - 200ms per slice including diffusion preparation), which makes it fairly immune to subject motion and enables advanced diffusion protocols with a large number of diffusion directions and/or b-values within reasonable scan time. However, SSH-EPI is prone to several limitations, including image distortion due to B_0 inhomogeneity at tissue/air interfaces and T_2^* blurring, both of which place limitations on spatial resolution. High-resolution diffusion MRI provides the ability to resolve fine-scale structures, enabling detection of cortical anisotropy [34, 35], delineation of thin white matter tracts [36] and more accurate fibre tractography [37]. Although parallel imaging has improved the data quality of SSH-EPI, these problems still exist and become more pronounced at high field strength and high resolution.

Alternative acquisition schemes have been proposed to overcome the limitations of SSH-EPI, including segmented-EPI readout, non-EPI trajectories and reduced FOV. These methods have undergone rapid development in recent years, demonstrating significantly improved image quality compared to SSH-EPI. Conventional 2D acquisition schemes suffer from long scan time and low signal-to-noise ratio (SNR) efficiency when acquiring high isotropic-resolution diffusion MRI data with full brain coverage, which is increasingly needed in neuroscience studies. The recent development of simultaneous multi-slice techniques has dramatically changed this situation and diffusion MRI data can be acquired more rapidly and with higher SNR efficiency. Various three-dimensional (3D) diffusion MRI acquisitions have also been developed, which have high SNR efficiency and can provide more accurate slice definitions than 2D acquisition. Several studies have reported high-quality diffusion MRI data at an ultra-high field strength of 7T, which opens new possibilities for achieving higher spatial and angular resolution. Finally, developments for fast diffusion MRI using compressed sensing have been reported, which can further accelerate diffusion acquisition.

This Chapter provides a review of recent developments of image formation methods for diffusion MRI and discusses how these are likely to be used in practice. Three kinds of advances are covered: improving image fidelity, accelerating acquisition and increasing SNR. However, inevitably a method that impacts one of these metrics has consequences for the others, and these trade-offs are described throughout this Chapter. Some of the methods discussed here are fairly unique to diffusion imaging (e.g. navigated correction of motion-induced phase errors), while others have broader application (e.g. simultaneous multi-slice imaging). For conciseness, diffusion contrast mechanisms are not covered and examples are focused on diffusion MRI of the brain.

3.2 Improving Image Fidelity

3.2.1 SSH-EPI Artefacts

The two dominant image artefacts in SSH-EPI are image blurring and distortion. The spatial resolution of SSH-EPI is severely affected by the tissue T_2^* decay (Fig. 3.1), which results in intense signal loss at the outermost edges of k-space. Because outer k-space corresponds to fine spatial detail (high spatial frequencies), this weighting introduces image blurring. Thus, a short readout window is important for reducing EPI blurring.

Image distortion in SSH-EPI mainly happens in regions with strong local magnetic field inhomogeneity (e.g., tissue/air boundaries with fast susceptibility variation). This local field distortion can swamp the weak gradients used for phase encoding, resulting in misplaced signal that appears as image distortion (Fig. 3.2). Along the readout direction, the amount of distortion is usually negligible as the readout gradient is much stronger than the off-resonance field. In the phase-encoding direction, the scale of distortion is determined by the speed of k space transversal. A quantitative description of geometric distortion in SSH-EPI can

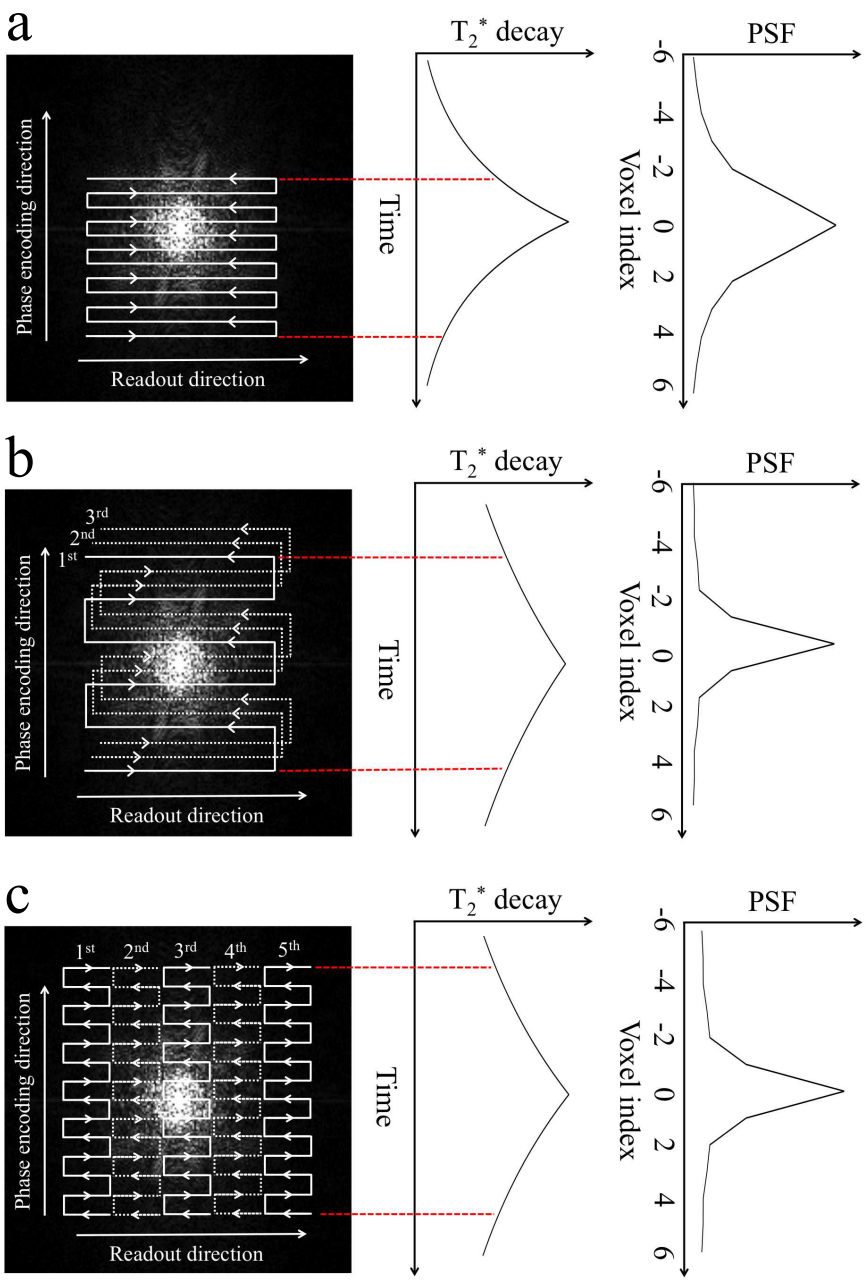


Figure 3.1: (a) Single-shot EPI trajectory samples k space very rapidly (30-50ms per image). However, tissue T_2^* decay causes signal loss at outer k space, which corresponds to high spatial frequencies, leading to image blurring as indicated by the point spread function (PSF). Partial Fourier acquisition is illustrated here, which can effectively reduce the echo time and hence increase SNR. (b) Conventional segmented EPI (three segments shown here) and (c) readout segmented EPI (five segments shown here) significantly reduces the effective echo spacing (e.g. from about 0.8ms in SSH-EPI to about 0.25ms in segmented EPI and 0.32ms in readout segmented EPI; parallel imaging is not considered here), leading to sharper PSF shape.

be found in Appendix A. Therefore, short echo spacing and under-sampling, both of which enable faster traversal along the phase-encoding direction, are desirable properties for improving SSH-EPI. The application of parallel imaging [30, 31] in SSH-EPI has been very successful, though this method faces challenges from noise amplification, particularly for diffusion MRI. "Effective echo spacing" accounts for the reduced distortion in parallel imaging by dividing the acquired echo spacing by the acceleration factor, which gives the echo spacing that would have been required to achieve this level of distortion without acceleration (i.e., for the same FOV, any EPI scan with the same effective echo spacing will have the same distortion). Reduced FOV methods can reduce distortion without noise penalty, but are limited to small coverage.

Another source of distortion in diffusion MRI is the eddy currents induced within the conducting surfaces of the magnet due to fast gradient switching. Modern systems include gradient "pre-emphasis" and shielded gradient coils that can substantially reduce eddy-current effects, but these corrections struggle to compensate for strong diffusion-encoding gradients with high slew-rate. The twice-refocused spin-echo diffusion preparation [38] can further reduce the effects of eddy currents, but at the cost of longer echo time. This is particularly problematic for high-resolution scans at ultra-high field, where this scheme requires 30ms longer echo time (TE) than conventional spin-echo diffusion preparation for 1mm resolution scan, exacerbating the already problematic T_2 signal loss. Another approach is to correct the distortion during post-processing [39]. More recently, correction of distortion based on non-parametric modelling of diffusion signal with respect to the diffusion-encoding direction has demonstrated excellent results [40].

Partial Fourier acquisition is often used to reduce the long echo time of SSH-EPI, which is another challenge for diffusion MRI due to greater T_2 signal decay (lower SNR). This is particularly problematic for high-resolution scans with longer echo time and smaller voxel size. Although partial Fourier reduces echo time, it increases

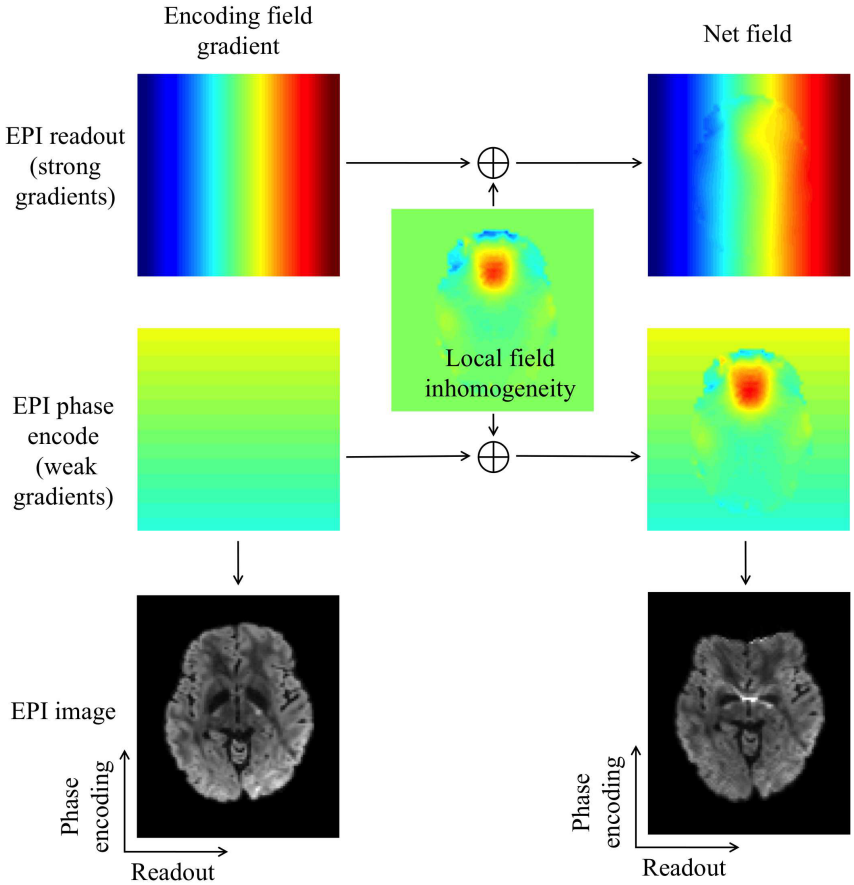


Figure 3.2: EPI distortions stem from inhomogeneity of the main magnetic field and are most pronounced at tissue-air and tissue/bone interfaces due to the large local field inhomogeneity caused by susceptibility variations. MRI image formation assumes the linear field gradient used for spatial encoding is achieved exactly as planned. However, due to the main field inhomogeneity, the net field deviates from the desired linear change, leading to incorrect mapping of voxels. The result is image distortion (e.g., compression in the frontal lobe as shown in the figure). As the encoding field for EPI phase encode is relatively weak compared to local field changes, distortions are severe along this direction. Encoding field for EPI readout is much stronger than local field changes and the image distortion along this direction is negligible.

the sensitivity to subject bulk motion during diffusion encoding [41], which induces echo shifting in k-space [42]. For strong motion, the echo is shifted toward the edge of k-space (higher spatial frequency), violating the fundamental assumptions underlying partial Fourier and causing the reconstruction to fail. This leads to image intensity oscillations and signal loss [43]. To alleviate this problem, adaptive partial Fourier reconstruction algorithms have been proposed [43, 44].

Other common EPI artefacts that may appear in diffusion MRI images include

Nyquist ghosting caused by hardware related odd-even echo misalignment and fat shifting along the phase encoding direction. Nyqusit ghosting is typically corrected by measuring the k-space shift between odd and even echoes using a reference scan and subsequently re-aligning the k-space data. Reference-less methods using image-entropy as a selection metric have also demonstrated robust correction of Nyquist ghosting [45]. Fat shifting occurs due to the difference of resonance frequencies between water and fat and the low bandwidth in the phase-encoding direction. For example, because fat differs from water by 440 Hz at 3T, fat signal will be shifted by $\sim$1/3 FOV for an EPI acquisition with echo spacing of 0.8ms = 1250Hz^{-1}. To eliminate fat-shifting artefacts, fat suppression is commonly implemented in EPI acquisition. Most fat suppression methods utilize special excitation schemes, including water-only spectral spatial excitation, fat saturation (and spoiling), and inversion-recovery preparation [46].

3.2.2 Multi-shot Acquisition

A number of multi-shot acquisition techniques have also been proposed, for which subject motion must be carefully handled, typically using a k-space "navigator" [47, 48]. Navigators provide an un-aliased, low-resolution image corresponding to a limited central k-space region that is used to predict effects from subject motion. Navigators are usually acquired immediately before/after the imaging data or extracted from the imaging data directly ("self-navigation"). Diffusion signals are intrinsically sensitive to subject motion because diffusion preparation gradients encode tiny (molecular) motions in signal phase. Even small subject motions (e.g., cardiac pulsation, respiration) during diffusion preparation can lead to substantial spatially-varying phase that is unrelated to diffusion. SSH-EPI is immune to these phase errors because it captures the entire image in one shot, such that the phase of the signal can be discarded, while diffusive motion is reflected in the signal magnitude. By comparison, multi-shot image acquisitions must retain

phase information in order to accurately combine data across the different k-space segments acquired in each shot. If motion-induced phase is not corrected before combining multi-shot segments, images become severely corrupted (Fig. 3.3). To correct motion-induced phase errors, it is common to acquire additional k-space measurements that can be used as a low-resolution navigator, using the phase of the navigator image to rectify the phase inconsistency between segments.

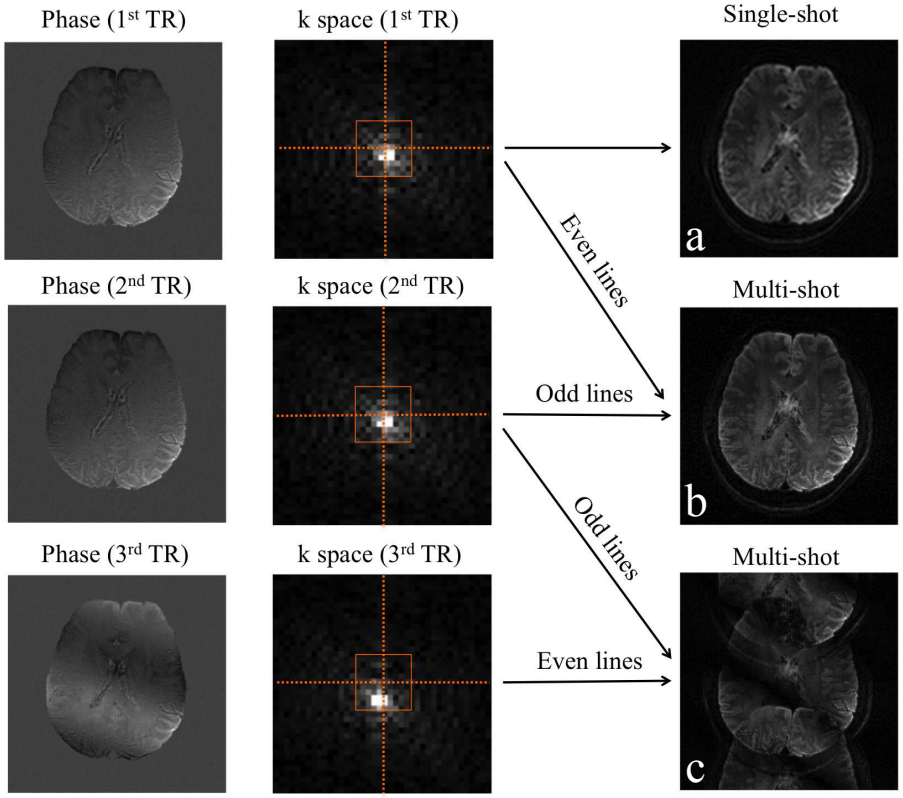


Figure 3.3: Subject motion during diffusion encoding can introduce substantial spatially-varying phase to the data. For single-shot EPI, this is not a problem, as only the magnitude is used and the phase is discarded (a). For multi-shot acquisition, however, this phase inconsistency must be properly handled. In this illustration, the subject is assumed to be steady during the first two TRs (no phase changes) and have a head rotation during the third TR, which introduces a linear phase offset (corresponding to a shift in k space). Multi-shot acquisition using even lines from the 1^{st}TR and odd lines from the 2^{nd}TR provides artefacts-free image (b). Whereas multi-shot acquisition using odd/even lines from the $2^{nd}$and the 3^{rd}TRs in a similar manner suffers from substantial image corruption due to the shot-to-shot phase inconsistency (c). One solution to this problem is navigation, which measures the motion-induced phase errors during each shot and corrects them before combining all segments.

Readout-segmented EPI

Readout-segmented EPI (rs-EPI) uses a series of EPI acquisitions to cover k-space in a mosaic pattern [49], as depicted in Fig. 3.1c. By acquiring concatenated k-space segments along the readout direction, rs-EPI can achieve much shorter echo spacing (e.g., echo spacing could be shortened from about 0.8ms in SSH-EPI to about 0.3ms in rs-EPI with 7 readout segments for diffusion scanning at 1.2mm isotropic-resolution [50]), and hence considerably reduce geometric distortion and T_2^* blurring [51–54]. The effective echo spacing can be further reduced via the combination with parallel imaging. The resulting TE of rs-EPI is also shorter than SSH-EPI (e.g., for 2mm^2 isotropic-resolution and b=1000s/mm^2 at 3T, TE could be reduced from about 87ms in SSH-EPI to about 73ms in rs-EPI (17), which is predicted to provide a 25% SNR improvement). The improvements offered by the reduced echo spacing (reduced distortion, blurring and TE) should be particularly beneficial at ultra-high field where tissue T_2^* and T_2 are very short and field inhomogeneity is worse [55].

As a multi-shot diffusion acquisition method, rs-EPI requires navigation to remove motion-induced phase errors. rs-EPI acquires a continuous k-space segment, enabling fast and robust correction using the nonlinear phase correction method [52, 53]. By comparison, in conventional segmented EPI (see 'segmented EPI' section below), the image reconstructed from each segment suffers from aliasing, which is difficult to correct robustly using the nonlinear phase correction. Iterative algorithms [56] can address this problem, but at the cost of long computational time.

The nonlinear phase correction only works if the motion induced phase errors can be accurately extracted from the navigator. In case of severe subject motion during diffusion preparation, the navigator fails to correct the phase errors when the center of k-space shifts out of the navigator acquisition window, making the motion information obtained from the navigator inaccurate. These corrupted data could severely degrade the image quality, and they cannot be simply removed as in

Propeller [57] (see "Propeller" section below) because there is no redundancy between readout segments. Instead, severely corrupted segments can be detected using the navigator and replaced by re-acquired segments at the same k-space location [53].

As with other multi-shot methods, the primary challenge of rs-EPI is the longer scan times required to form each image volume. Many of the approaches to reducing scan times discussed in the section on acceleration of diffusion MRI are compatible with rs-EPI, and both partial Fourier [58] and simultaneously multi-slice [50] have been proposed to reduce scan times. rs-EPI has also been demonstrated at 7T [55] and in conjunction with 3D multi-slab acquisitions [59].

Superior data quality using rs-EPI has been demonstrated compared to SSH-EPI, especially in regions with strong susceptibility variation, such as the temporal lobes and brainstem [52, 53]. Several clinical studies have investigated the performance of rs-EPI for clinical diagnosis, demonstrating high data fidelity and improved conspicuity of pathology, for applications including: breast cancer [60–62], liver tumors [63], pelvic [64, 65] and renal [66] diseases, and pediatric neuropathology [67].

Segmented EPI

Segmented EPI samples k-space over multiple EPI trajectories with broadly spaced lines that interleave to cover the full k-space, as depicted in Fig. 3.1b. In early implementations, a 2D navigator was used to correct the motion-induced phase errors [68, 69]. Recently, the multiplexed sensitivity encoding (MUSE) method [70] has been proposed to correct the motion-induced phase errors in segmented EPI without acquiring 2D navigators. MUSE reconstruction consists of three components: first, a phase navigator is calculated for each segment using parallel imaging to fill in the missing samples in central k-space; second, each segment is phase-error corrected using this navigator; third, all segments are combined to form the final image (in practice, the second and third parts are calculated simultaneously). Several refinements of MUSE were proposed to correct the rigid

motion between different segments (Fig. 3.4) [71, 72] and prospectively detect and reject severely motion-corrupted interleaves [73]. Extensions of MUSE with robust partial Fourier reconstruction, simultaneous multi-slice [44] and 3D multi-slab acquisition have also been proposed [74].

A disadvantage of the MUSE method is the limitation on the number of interleaves, since this equates to the acceleration factor in the first reconstruction step, which is constrained by the RF receive-coil design. A more recent method formulated the reconstruction as a low-rank matrix completion process without explicitly estimating the motion-induced phase errors [75]. These navigator-free methods improve the efficiency of segmented EPI acquisition (navigator acquisition takes 30-40ms for each excitation) and reduce the SAR (by 30% due to the removal of the refocusing pulse for the navigator echo), which may be critical at ultra-high field.

PROPELLER

One approach to reducing image distortion and T_2^* blurring uses a class of sequences known as rapid acquisition with refocused echoes (RARE) [76] (sometimes called FSE or TSE), which uses a series of refocusing pulses to create a train of spin echoes. This method can reduce artefacts by acquiring every k-space line at the center of a spin echo, thus avoiding the phase accumulation that leads to distortion and blurring in EPI. The most widely used implementation of FSE in diffusion MRI is PROPELLER (Periodically Rotated Overlapping ParallEL Lines with Enhanced Reconstruction), which acquires a strip ('blade') covering the center of k-space in each train of refocusing pulses [57, 77]. Multiple rotated blades are acquired to fully sample a circular k-space region over multiple shots, as shown in Fig. 3.5a. A key strength of PROPELLER is that it is self-navigating: all blades cover the k-space center, which can be used to estimate the motion induced phase errors for each shot.

Several key challenges for PROPELLER have been identified and addressed, including motion sensitivity (due to the conditions for formation of spin-echo trains),

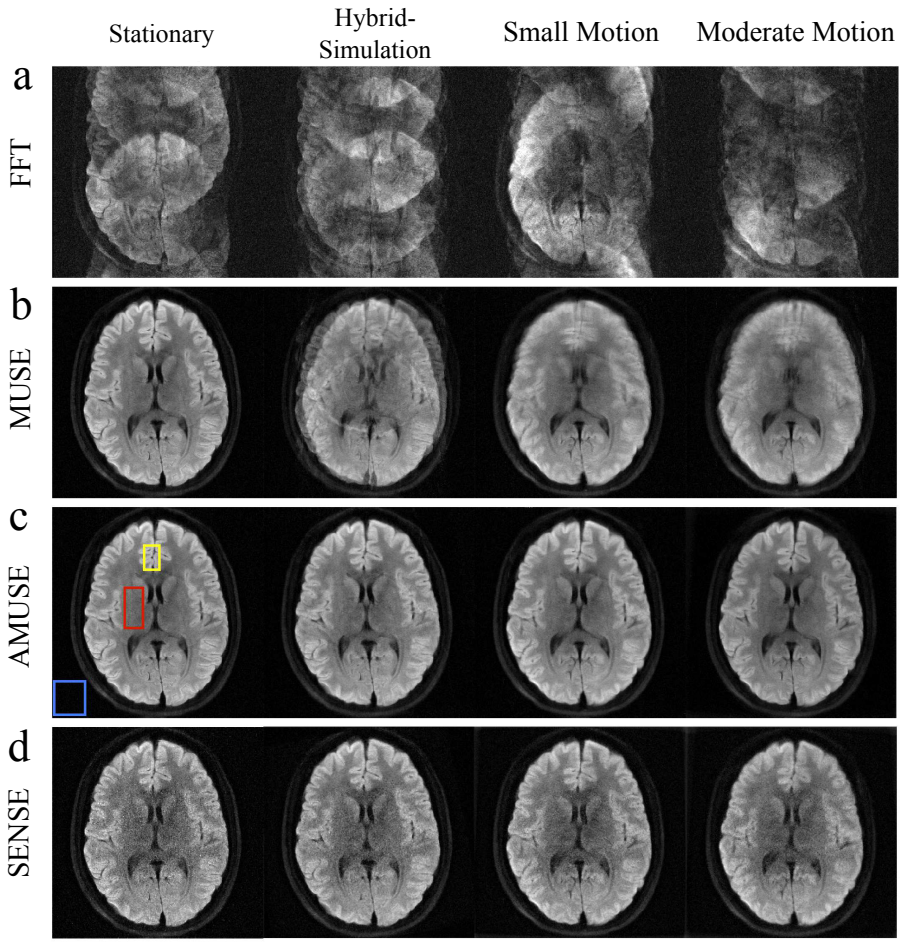


Figure 3.4: Reconstruction of multi-shot diffusion MRI data using direct fast Fourier transform (FFT), multiplexed sensitivity encoding (MUSE), augmented MUSE (AMUSE) and sensitivity encoding (SENSE) based motion correction. The AMUSE method simultaneously corrects motion induced phase errors and macroscopic motion. Four types of subject motion are evaluated, including stationary, hybrid-simulation (combining data from two scans, in which subject head is stationary during acquisition but rotates about 15° between scans), small motion (about ±5° rotation every 10-15s) and moderate motion (about ±10° rotation every 10-15s). FFT and MUSE are corrupted by subject motion, whereas both AMUSE and SENSE reduce the motion artefacts. AMUSE further provides higher SNR. Figure reproduced with permission from [71].

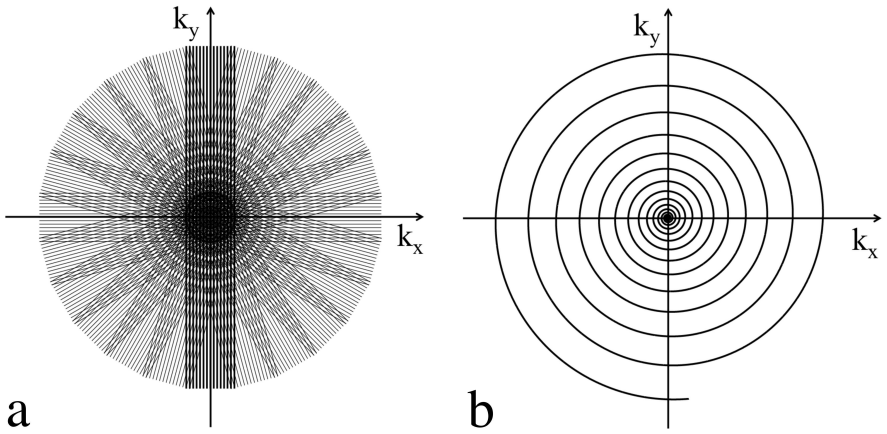


Figure 3.5: Two non-Cartesian diffusion acquisition trajectories: (a) PROPELLER. (b) Variable density spiral.

RF deposition (due to the large number of refocusing pulses in the FSE readout) and imaging speed (due to the multi-shot acquisition). Signal formation requires stable signal phase over the course of the spin-echo train [78], which is disrupted by motion, resulting in signal decay and oscillates. Alternating the phase of the refocusing pulses between x and y axes has been demonstrated to stabilize the signal [57, 79, 80]. To accelerate the acquisition and reduce RF deposition, Turboprop [81] collects multiple gradient echoes between refocusing pulse pairs. Turboprop is in effect a gradient- and spin-echo (GRASE [82]) sequence, providing a trade-off between image distortion and speed. Going further in this direction, PROPELLER-EPI [83] acquires each blade using a single EPI readout, combining the self-navigation of PROPELLER and the rapid acquisition and low-SAR properties of EPI. PROPELLER-EPI has greater blurring and distortion than the original FSE technique, which can be mitigated with parallel imaging [84]. Alternatively short-axis PROPELLER-EPI [85] places the EPI readout along the short axis of the blade (Fig. 3.6), leading to short echo spacing and therefore reduced blurring artefacts.

Spiral

The most commonly considered alternative to Cartesian sampling techniques like EPI are spiral trajectories, in which k-space is traced in a radiating pattern rather than

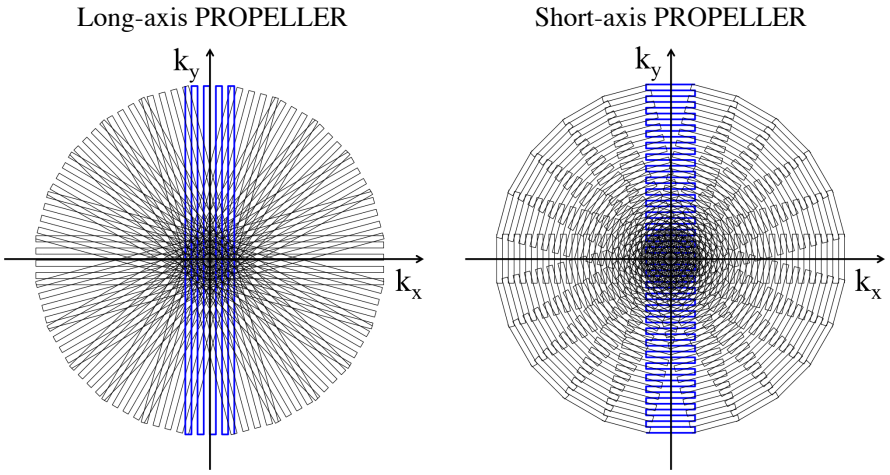


Figure 3.6: Two variants of PROPELLER-EPI: long-axis (LAP) PROPELLER and short-axis (SAP) PROPELLER. The readout directions for LAP and SAP are along the long- and the short-axis of the strip, respectively. The ky transverse speed in SAP acquisition is faster than that in LAP, resulting in less image blurring.

a line-by-line scan. Spiral imaging has the merit of intrinsic motion compensation through gradient-moment-nulling [86] and efficient use of gradient power [87]. Single-shot spiral imaging acquires diffusion MRI data with similar efficiency to SSH-EPI [88], but suffers from different artefacts due to off resonance and T_2^* decay, both of which result in image blurring.

Similar to segmented EPI, spiral acquisitions can be acquired in multiple shots using interleaved acquisitions to achieve high spatial resolution and reduced image blurring. As with all multi-shot diffusion imaging methods, motion-induced phase errors need to be corrected before combining interleaved spiral acquisitions. Variable density spirals (VDS) can be self-navigating by sampling central k-space densely with each interleave (Fig. 3.5b) [89]. Alternatively, constant density spirals can be used with a similar strategy as MUSE to extract motion navigators using a parallel imaging reconstruction [90]. Spirals are also pseudo-incoherent with respect to under-sampling artefacts, which makes it a preferable sampling method for compressed sensing reconstruction [91], as discussed below. Spirals have also been extended to 3D acquisitions using thin slabs [92]. The major challenge for spiral acquisition is image blurring, which compromises the spatial resolution that can be

achieved. Although de-blurring methods can alleviate this problem to some extent, their performance still needs to be improved in the presence of strong susceptibility variations or severe $\mathbf{B}_0$ inhomogeneity (e.g. at ultra-high field).

3.2.3 Reduced FOV Method

Similar to parallel imaging, reduced field of view (rFOV) techniques reduce distortion and T_2^* blurring by skipping phase encoding lines. In rFOV, aliasing is avoided by excluding signal from outside a limited volume within a region of extended tissue. Three main strategies for rFOV acquisition and their application in diffusion MRI are reviewed here: inner volume imaging, outer volume suppression and multi-dimensional RF excitation.

Inner volume imaging (IVI) uses orthogonal orientations for excitation and refocusing pulses such that only the overlapping regions of the excited volume and refocused volume create signal, enabling a reduced imaging FOV [93]. In its original form, IVI was limited to single-slice imaging because the refocusing pulse saturates parallel slices. A refinement of the IVI method places the refocusing pulse at a shallower angle to the excitation [94], which enables multiple slice acquisition, but requires gaps between slices. Alternatively, one can apply another refocusing pulse after the readout [95], returning the spins from the non-imaged slices to the positive longitudinal axis, which enables contiguous interleaved multi-slice acquisition.

Outer volume suppression (OVS) suppresses signal from outside the imaging volume using spatially selective RF pulses followed by de-phasing gradients [96]. The OVS pulses are applied prior to the imaging acquisition, resulting in signal only from the non-suppressed target region. OVS incurs increased SAR and longer scan time due to these suppression pulses, and is sensitive to RF transmit field inhomogeneity. Nevertheless, reduced FOV diffusion MRI with OVS has demonstrated superior structural details in spinal cord [96] and pons [97] compared to conventional SSH-EPI. In combination with parallel imaging, OVS has been used to address the severe

$\mathbf{B}_0$ inhomogeneity and short tissue T_2 value at ultra-high field (Fig. 3.7) [98–100]. OVS has also been combined with SMS for high-resolution diffusion MRI [101, 102].

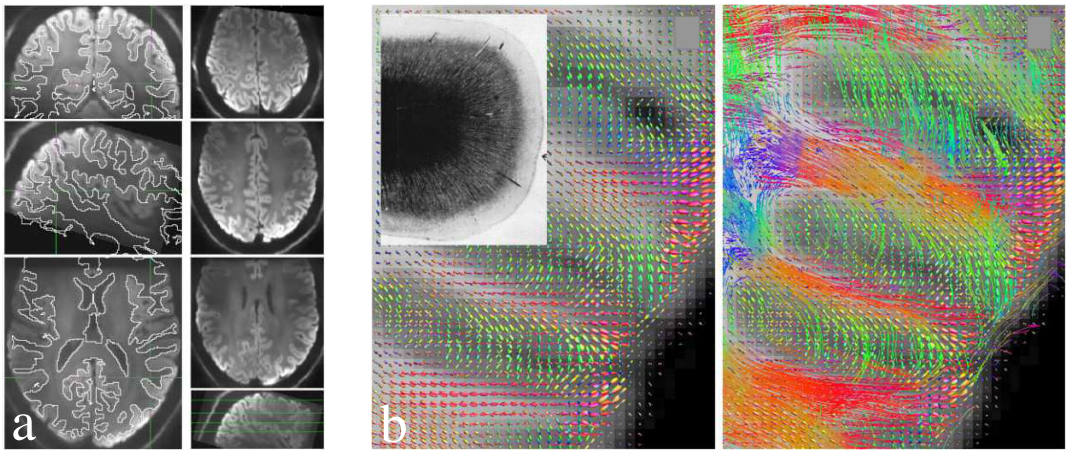


Figure 3.7: High-resolution (0.8 mm isotropic) diffusion MRI data acquired using a combination of reduced FOV and parallel imaging methods at 7T. (a). Left column: Trace-weighted images overlaid with white/grey matter boundaries obtained from an anatomical scan, demonstrating high geometric fidelity achieved by combining reduced FOV methods and parallel imaging. Right column: axial slices from different brain regions. (b). Fiber orientation density (left) and streamline tracking (right) based on the high resolution data, depicting white matter fibre tracts entering the cortex. Figure reproduced with permission from [99]

The third approach to reduced FOV uses 2D spatially selective RF pulse for excitation and a conventional slice-selective 180° pulse for refocusing [103]. The most common multi-dimensional pulses, using echo-planar gradients, result in a periodic excitation profile. This profile places limits on the orientation and number of slices [104]. A refinement of this method demonstrated the ability to simultaneously refocus two slices [105], doubling the number of slices that can be acquired in each scan. An alternate approach has been proposed that has virtually unlimited slice coverage, but which requires separate fat saturation [106]. The long pulse durations associated with multi-dimensional excitations can be reduced using parallel transmission [107]. Multi-dimensional excitation for rFOV diffusion imaging has been compared to the standard SSH-EPI method, demonstrating clinical feasibility and improved conspicuity for spinal cord [108, 109] and breast imaging [110–112].

3.3 Accelerating Diffusion MRI Acquisitions

As noted above, SSH-EPI is highly efficient during the signal readout period, providing all the spatial information for one slice in 20-40ms. However, regardless of readout, diffusion MRI sequences are generally inefficient, with $\leq 50\%$ of the sequence time dedicated to signal acquisition due to the need for a long diffusion preparation. In conventional 2D SSH-EPI sequences, this inefficiency is compounded by the fact that each slice is encoded independently in series, such that the acquisition time per volume scales with the number of imaging slices. Assuming fixed coverage, increased spatial resolution requires more slices, further inflating volume scan time. This inefficiency results in a difficult trade-off, particularly when a large number of diffusion directions are desired to improve the accuracy of angular information (e.g., for diffusion tractography). When total scan time is limited, there is a fundamental trade-off between spatial coverage, spatial resolution and angular resolution (density of sampling in the diffusion-encoding (directional) domain).

Several techniques have been introduced in the past few years that have the potential to dramatically reduce this trade-off. Simultaneous multi-slice (SMS) techniques have provided the ability to acquire multiple diffusion-encoded slices simultaneously, increasing the scan efficiency (as reflected in the number of slices acquired per unit time). Compressed sensing has shown the possibility to reconstruct MRI image from highly under-sampled data [91], which can benefit diffusion scans with a large number of diffusion directions and/or multi-shot k-space acquisition.

3.3.1 Simultaneous Multi-slice Imaging

The idea of exciting multiple slices simultaneously was proposed more than twenty years ago [113, 114]. However, the first SMS techniques required multiple excitations to separate slices and did not reduce scan time. A key advance was made when it was realized that multi-channel coil arrays enabled slice separation from a single

acquisition through a parallel imaging formulation, thereby accelerating volume acquisition [115].This approach was subsequently extended to SSH-EPI [116] and demonstrated to enable high spatial-angular resolution diffusion MRI [1, 117].

A major challenge faced by SMS is noise amplification when gaps between slices are small. Coil profiles in general vary slowly across space, meaning that closely separated slices tend to have similar profiles (i.e., the high-signal region from one slice overlaps with the high-signal region from another slice) (Fig. 3.8b). This problem is the SMS manifestation of the 'g-factor' (the noise amplification for a given image voxel, which reflects coil configuration, acquisition protocol and reconstruction algorithm) from conventional parallel imaging [30] and is particularly challenging for diffusion MRI due to its low intrinsic SNR. The "blipped-CAIPI" (controlled aliasing in parallel imaging [118]) method reduces the g-factor, representing a major improvement on SMS [1, 2]. Blipped-CAIPI introduces an apparent in-plane shift between the excited slices, such that a given coil profile is spatially separated in the overlapping slices. In this case, the aliased voxels corresponding to two slices can be more easily separated because the coil profiles appear more distinct (i.e., the high-signal region from an un-shifted slice overlaps with the low-signal region from a shifted slice) (Fig. 3.8c). SMS-EPI, in particular blipped-CAIPI and its variants, has greatly improved the quality of diffusion imaging studies by reducing the trade-off between spatial and angular resolution. Two high-profile examples in brain imaging include the Human Connectome Project, where it has enabled a protocol with high spatial and angular resolution at several diffusion weighting "shells" [37]; and the UK Biobank Project, where it has enabled multiple shells with reasonable angular resolution in very limited scan time [119]. Blipped-CAIPI SMS-EPI has also been incorporated with segmented-EPI [44, 50, 120] and reduced FOV (see above). An extension to 3D simultaneous multi-slab acquisition has also been reported [121].

Image reconstruction to separate slices in SMS builds strongly on the existing literature in parallel imaging. There are two main categories of reconstruction

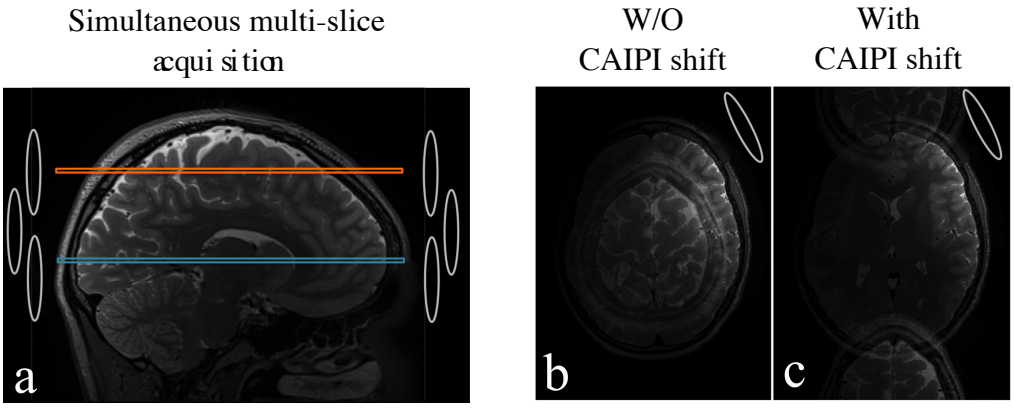


Figure 3.8: Simultaneous multi-slice imaging utilizes multiple-channel coil arrays to separate multiple slices excited simultaneously (shown in orange and blue in (a)). The collapsed slices can be separated using parallel imaging, given there are sufficient variations in the coil sensitivities. However, in the case of high slice acceleration, the distances between the aliased voxels become small, making it difficult to resolve the aliasing due to the lack of coil sensitivity variations. The result is residual artefacts and amplified noise (g-factor penalty) in the reconstructed images. With CAIPI scheme (e.g., blipped-CAIPI in SSH-EPI), the distances between slices are increased by an apparent shifting along the phase encoding direction (c), which can significantly improve the reconstruction.

methods for SMS data: SENSE-GRAPPA [122] and Slice-GRAPPA [1]. SENSE-GRAPPA treats the overlapping slices as if they are neighboring in the phase encode direction space over a larger FOV, which casts the slice separation problem in a form that can be solved by conventional parallel imaging reconstructions. With the blipped-CAIPI scheme, SENSE-GRAPPA contains artefacts at the concatenation points, which have been avoided by additional zero-padding [123–125] or concatenating along the readout direction [126]. By comparison, slice-GRAPPA trains slice-specific kernels that project k-space data to one corresponding slice. Slice-GRAPPA has been shown to be dependent on coil sensitivity rather than the image contrast [1], which is a desirable property for diffusion MRI, where the reconstruction is trained on data with no diffusion weighting. Several further refinements to SMS reconstruction have been proposed. The reconstruction kernel in Slice-GRAPPA has been improved to reduce "leakage" between slices [127] by training the kernel to block signal from all but one slice. This modification has been crucial for simultaneous slice and in-plane acceleration, both of which play an important role

in data quality. However, the interaction between these two accelerations is still challenging, given that both methods rely on multi-channel coils in a similar manner.

RF pulses that excite multiple slices simultaneously are generally referred to as "multiband" (MB) pulses, since they deposit energy at several separate frequency bands. These pulses in general require a higher energy deposition, which is necessarily limited by patient safety considerations. A basic multiband pulse is a superposition of multiple conventional RF pulses, leading to an N^2 increase in the peak RF power for N simultaneously excited slices. This problem is particularly challenging for diffusion MRI due to the use of high-energy 180° refocusing pulses. RF power can be reduced by optimizing the phases of the superimposed RF pulses [128] or using a time-shift scheme [129]; however, these methods only reduce the peak RF power, not the SAR level. The variable-rate selective excitation (VERSE) algorithm [130] can be used to improve multi-band pulses by modifying the gradients to slow down k-space transversal speed during peak power deposition; however, VERSE suffers from RF profile distortion in the presence of strong field inhomogeneity. Another approach is PINS (power independent of number of slices) pulses [131], which under-sample conventional single-band pulses in a manner that excites equally separated slices while retaining power deposition comparable to a single-slice excitation. Drawbacks of PINS pulse include poor slice profiles and limitations on the achievable slice orientation. Improvements include combination of PINS with more conventional multiband pulses and methods that improve robustness to $\mathbf{B}_1$ inhomogeneity [132–134]. Parallel transmission (pTx) has also been explored for improving SMS acquisition. Using a full pTx-MB model, a significant reduction on total RF power can be achieved [135]. Alternatively, a dual-ring RF array design [136, 137] has been proposed, but suffers from the RF discontinuity between SMS slice stacks.

In summary, SMS enables full-brain diffusion MRI with acquisition time of 3-7s per volume (dependent on resolution and acceleration factor), enabling high spatial

resolution and high angular resolution diffusion acquisition, provided reconstruction challenges such as compatibility with in-plane acceleration can be addressed. The challenges from RF inhomogeneity and high SAR level are likely to be addressed by novel RF pulse design and pTx techniques.

3.3.2 Fast Diffusion MRI with Compressed Sensing

The total acquisition time of a diffusion MRI exam is determined by the number of diffusion volumes (referred to as "q-space samples") and scan time for each volume (k-space). Recently, compressed sensing (CS) methods have been applied to accelerated diffusion acquisition by exploring data sparsity in these two independent domains: q-space and k-space (Fig. 3.9).

Density of sampling in q-space is a key determinant of data quality for diffusion MRI methods that aim to characterize the orientational microstructure of tissue, for example to enable diffusion tractography (delineation of white matter pathways achieved by following the preferential direction of water diffusion from voxel-to-voxel) in the brain. This presents an important challenge for these methods because scan times are in general proportional to the number of q-space samples acquired (since each direction represents one imaging volume). Recently, CS has been used to improve the reconstruction for acquisitions with a small number of q-space samples (i.e., q-space under-sampling, generally with a pseudo-random sampling pattern). For example, several groups have reported CS methods for reconstructing data acquired with a single b-value (q-space "shell") to capture complex fibre architecture based on a multi-tensor model [138] or spherical ridgelets basis [139]. Other CS methods aim to capture the additional information available from multiple q-space shells using adaptive dictionaries or a sparsity prior on the diffusion probability density function [140, 141]. These CS methods have reported high fidelity recovery of fibre orientation properties from data with under-sampling factors ranging from three to eight, where higher accelerations are more readily

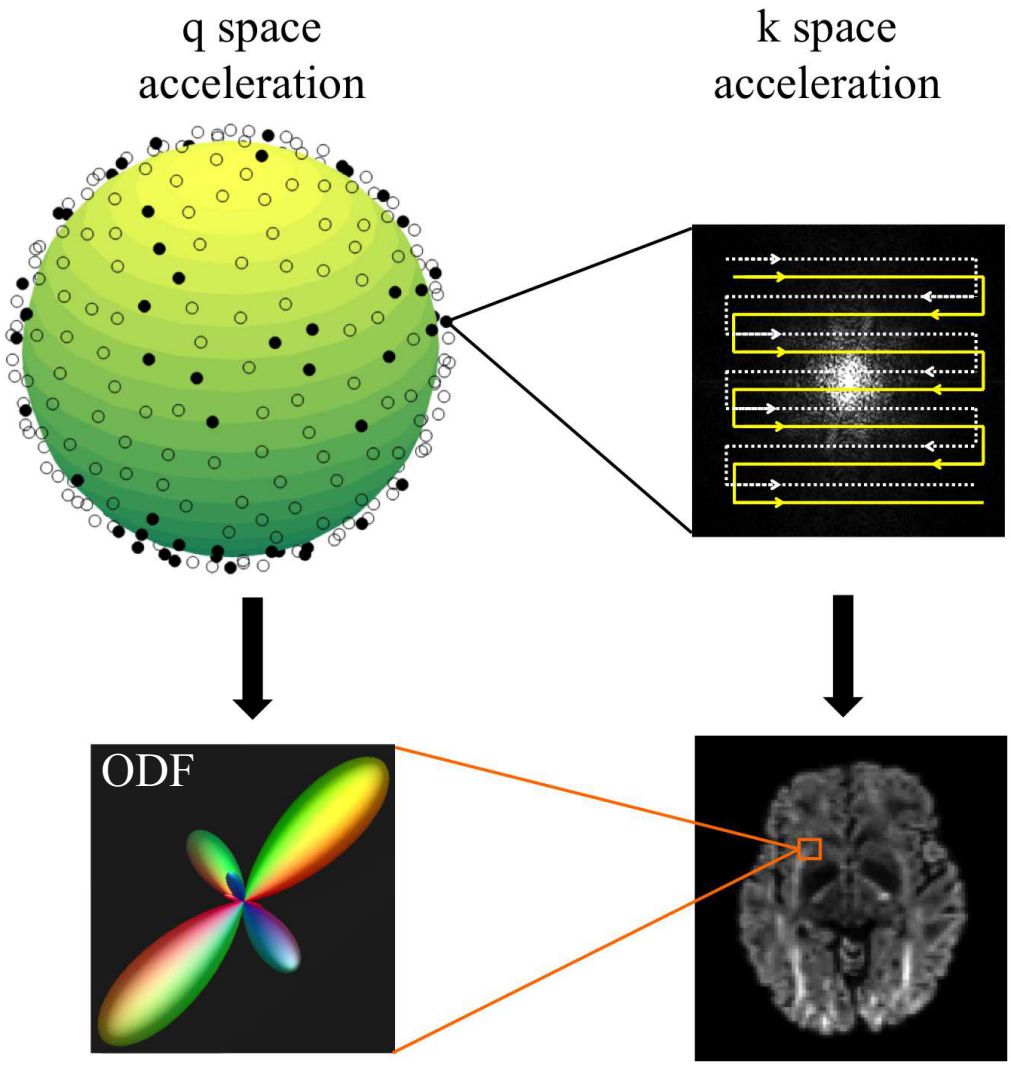


Figure 3.9: In diffusion MRI examinations, a full k-space is acquired at each q-space sample. The total scan time is determined by the number of q-space samples and the k-space acquisition time. Therefore, diffusion MRI can be accelerated in both q-space and k-space, although their objectives are different: in k-space acceleration, it is important to robustly recover the image information from reduced k-space data, whereas in q-space acceleration, the main target is to achieve accurate fibre orientation estimation (e.g. orientation distribution function (ODF)) with fewer q-space samples.

achievable for simple models (e.g., multi-tensor fits) compared to less constrained models (e.g., orientation distribution functions). Note that there are some other developments for CS diffusion MRI [142–144].

In general, k-space under-sampling is only effective for reducing scan time in multi-shot methods, since 2D single shot techniques like SSH-EPI acquire the entire k-space required for one slice in a very brief time window. To improve the reconstruction fidelity with highly under-sampled data, CS reconstructions have been proposed using sparsity priors on diffusion anisotropy [145], wavelet representations of the image data [146] and Gaussian mixture models of the diffusion process [147].

Joint k-q space accelerations, which under-sample in both domains, have also been proposed [147–149]. These methods use circulated or randomized k-space under-sampling patterns for different q-space locations to increase the incoherence of signal aliasing, and used CS or low-rank models to recover the underlying images. A critical step in the k-q joint reconstruction is to correct the motion induced phase errors, which are inconsistent among q-space samples. Navigator correction [147, 149] and the MUSE scheme [148] discussed in previous sections were applied to address this issue.

3.4 Increasing Image SNR

The third major challenge of diffusion MRI is the intrinsic low SNR, which makes it particularly difficult to achieve high spatial resolution diffusion MRI data. Two dominant methods to enhance SNR of diffusion MRI without significantly increasing the scan time include using imaging sequences with higher SNR per unit time (efficiency) (e.g. 3D acquisitions), and the use of ultra-high field scanner (7 tesla and above). In this section, the developments of these two approaches for diffusion acquisition are discussed.

3.4.1 Three-dimensional Acquisition

As detailed later on in this section, 3D acquisition methods are attractive because they enable imaging with high SNR efficiency. A further benefit of 3D imaging is the ability to define thin slices using gradient encoding (3D k-space), which produces slices with similar quality of definition to in-plane voxels, where 2D slices tend to have poor slice definition (e.g., warping).

3D diffusion MRI is challenging because k-space encoding in three dimensions can only be accomplished with a multi-shot approach, meaning that motion induced phase errors need to be corrected before combining across shots. In general, phase errors will vary spatially in all directions, implying that a full correction of these errors would require navigation in three dimensions; however, the acquisition of such a 3D navigator would require a prohibitively long scan time. Several approaches for avoiding the need for navigation in 3D imaging have been proposed. Driven equilibrium returns the magnetization to the longitudinal axis after diffusion preparation in order to enable spoiling of the phase error prior to the 3D readout [150]. Although driven equilibrium avoids the profound image artefacts associated with phase cancellation, this method still has signal dependence on motion due to corruption of the longitudinal magnetization, which is highly problematic for quantification. Similarly, cardiac gating can reduce the effects of pulsatile motion in 3D imaging [151], but residual motion artefacts persist.

Many of the 3D acquisitions for diffusion imaging were first proposed in the context of diffusion weighted SSFP (DW-SSFP) [152], which is an intrinsically 3D sequence. Early work explored the use of 1D [153] and 2D [154, 155] navigation for correcting 3D data. Subsequent work proposed a rotated-EPI acquisition scheme (TURBINE), which enabled a composite 3D navigator based on cardiac-synchronized "batching" of 2D readouts [156]. This approach was extended by including rigid-body motion correction [157]. However, a major outstanding challenge faced by

DW-SSFP methods is the disruption of signal formation by motion-induced phase errors, which results in altered signal magnitude that confounds quantification.

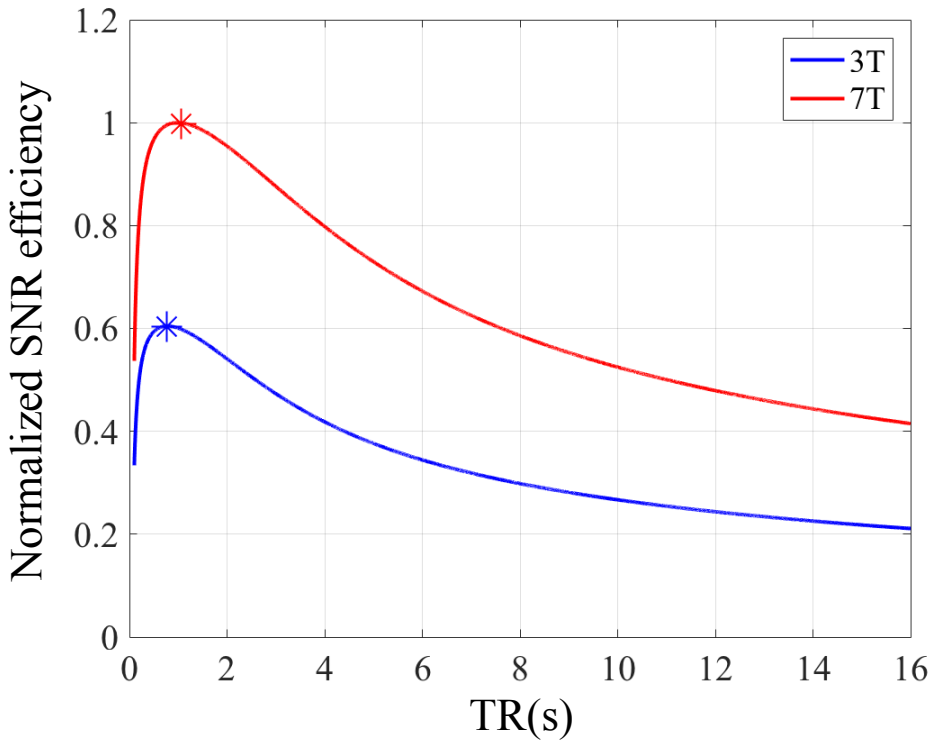


Figure 3.10: SNR efficiency for spin-echo diffusion acquisition, which is calculated as signal/$\sqrt{\text{TR}}$. T_1 values for white matter at 3T (838ms) and 7T (1220ms) are used. The optimal TR for white matter acquisition is 1-2s (indicated by the '$*$' on the plot).

Recently, development in this area has focused on 3D multi-slab acquisitions, in which a series of slabs are excited in sequence (similar to 2D multi-slice imaging but many thick slices), within which slices are defined using gradient encoding (i.e., kz encoding as in 3D imaging). As a hybrid 2D/3D method, 3D multi-slab acquisition can achieve thin slices like 3D imaging while being compatible with the moderate TR (1-2s) associated with high SNR efficiency (Fig. 3.10) [59, 158]. When a thin slab is applied, the motion induced phase error along the slice direction can be assumed to be slowly varying, thus a 2D navigator can provide sufficiently good correction [92, 158]. The use of a 2D navigator for phase-error correction is a key enabler of 3D multi-slab diffusion acquisition as it relieves the demanding

requirement of 3D navigator acquisition. Several groups have demonstrated the use of 3D multi-slab acquisitions to achieve high-resolution diffusion MRI data [8, 59, 74, 158, 159]. A further challenge with 3D multi-slab acquisition is slab-boundary artefacts, which arises because the magnetization experiences variable saturation at the overlap of slab profiles. Several reconstruction methods have been proposed to address this issue, with increasing degree of sophistication [7, 160, 161]. Correction of slab-boundary artefacts will be discussed in greater detail in Chapter 4.

3.4.2 Ultra-high Field

Because it provides greater signal, ultra-high field strength enables higher intrinsic SNR, which has great potential for high resolution and/or high b-value diffusion imaging. However, in order to fully benefit from the increased signal, several challenges must be addressed, including: shorter T_2, which reduces signal level dramatically at the long echo times typical in diffusion MRI; shorter T_2^*, which increases blurring artefacts in EPI images; the strong RF transmit field inhomogeneity, which is particularly problematic for refocusing pulses in spin-echo sequences; increased $\mathbf{B}_0$ inhomogeneity, which is associated with EPI distortion; and significantly increased SAR level, also associated with refocusing and/or SMS pulses. Despite these challenges, several studies have demonstrated high quality diffusion MRI data at 7T [5, 8, 35, 101, 162]. Due to shorter T_2, achieving short echo time is critical for realizing the SNR benefit at ultra-high field [3]; this has been demonstrated through the combination of parallel imaging and partial Fourier acquisition [5, 8]. The use of rs-EPI [50, 55] and reduced FOV [99, 101] can further reduce the echo time, as well as reducing EPI distortion. RF transmit field inhomogeneity can be mitigated using dielectric pads [5], although this approach inevitably improves some regions while worsening homogeneity in others [163]. A more comprehensive solution to transmit inhomogeneity would be to use parallel transmission, which can also reduce the SAR level. The combination of 3D multi-slab acquisition and

7T can further enhance the SNR performance [8], which is introduced in Chapter 4. In all, diffusion MRI at ultra-high field is a rapidly advancing research area. Given the intrinsic low-SNR of diffusion MRI and the increasing desire for higher spatial resolution and b-value, ultra-high field is expected to play an important role in future diffusion MRI research.

3.5 Summary

This Chapter presented an overview of the major technical advances for diffusion MRI acquisition in recent years. Compared to standard 2D SSH-EPI acquisition, these methods have demonstrated capabilities to enhance data fidelity, reduce acquisition time and improve SNR. Some of these techniques (rs-EPI, rFOV, PROPELLER) have already been applied in clinical settings with the aim of minimizing image distortion and blurring for improved diagnosis or disease monitoring in clinically feasible scan times. In contrast, basic neuroscience research applications typically require many diffusion directions with full brain coverage, for which accelerated acquisitions like SMS have already demonstrated major impact. For high isotropic resolution (e.g., sub-millimetre) diffusion scans, 3D acquisition has advantages including the accurate slice profile and SNR efficiency. Finally, diffusion MRI is starting to take advantage of advances in ultra-high field technology, including improved hardware and RF transmit techniques. While most of these methods have been developed in the brain, diffusion MRI throughout the body has become a major area of research focus, building on these advances for entirely new applications. Further improvement and combination of these techniques should enable more robust diffusion MRI acquisition with higher data quality, providing a powerful tool for new diagnostic, monitoring and mechanistic studies.

4

Reducing Slab-Boundary Artefacts for 3D Multi-slab Diffusion MRI

Contents

4.1 Introduction 68
4.1.1 Overview of Slab-Boundary Artefacts 70
4.2 Bloch Simulations 73
4.2.1 Slab Crosstalk in Multi-slab Imaging 73
4.2.2 Propagation of Slab Crosstalk into Diffusion Quantification 75
4.3 Methods 78
4.3.1 Proposed Correction 78
4.3.2 Numerical Simulation 83
4.3.3 In vivo 84
4.3.4 Reconstruction 85
4.4 Results 87
4.4.1 Reconstruction with Simulation Data 87
4.4.2 Reconstruction with Data Acquired In Vivo 92
4.5 Discussion 95
4.6 Conclusion 98

This Chapter describes a method we proposed to reduce the slab-boundary artefacts in 3D multi-slab diffusion MRI. Bloch simulation is used to investigate the effects of multiple factors on slab-boundary artefacts, including characterisation of

residual errors on diffusion quantification. A nonlinear inversion method is proposed to simultaneously estimate the slab profile and the underlying (corrected) image. Correction results of numerical phantom and in vivo data demonstrate that the method can effectively remove slab-boundary artefacts for diffusion data. Notably, the nonlinear inversion is also successful at short TR, a regime where previously proposed methods (slab-profile encoding and weighted average) retain residual artefacts in both diffusion-weighted images and diffusion metrics (mean diffusion coefficient and fractional anisotropy). In conclusion, the nonlinear inversion for removing slab-boundary artefacts provides improvements over existing methods, particularly at the short TRs required to maximise SNR efficiency. This Chapter is based on a peer-reviewed journal paper we published in Magnetic Resonance in Medicine [7].

4.1 Introduction

Improvements in spatial resolution for diffusion MRI of the brain can provide the ability to resolve small structures, enabling investigation of finer neuroanatomical features and detailed white matter fibre architecture. Although 2D acquisition can increase the in-plane resolution significantly [53, 54, 57, 69, 89, 99, 155, 164], this approach faces limitations when small isotropic voxels are required. Most importantly, 2D methods couple the repetition time (TR) to the number of slices, such that whole-brain coverage necessitates long TR (typically 6-10s). In the brain, signal-to-noise ratio (SNR) efficiency of spin-echo sequences is maximised when TR=1-2s (see Fig. 3.10 in Chapter 3), meaning that 2D acquisitions are far from optimal for even modest resolution [59]. This effect becomes acute at high resolution due to the competing needs for a large number of slices to cover the brain (requiring longer TR) and high SNR efficiency to support reduced voxel volume (favouring shorter TR). While simultaneous multi-slice methods mitigate

these issues to some degree, the coupling between number of slices and TR will remain an issue, particularly at high resolution.

Three-dimensional (3D) acquisitions remove this coupling by exciting the entire volume and defining voxels based on 3D k-space, acquiring the entire image over multiple readouts (multiple shots). Nevertheless, the utilization of 3D acquisitions for dMRI faces several limitations. Most importantly, in order to produce image volumes reasonably quickly, single-slab 3D dMRI requires very short TR ($\sim$100ms), which is also very SNR inefficient [59]. Another challenge comes from the motion-induced phase errors, which can degrade image quality in multi-shot acquisitions if not properly addressed. Applying a full 3D nonlinear phase correction method [155] requires time-consuming 3D phase-navigators, which would significantly diminish the scan efficiency. Acquisition of 3D phase-navigators could be avoided by using approaches like driven equilibrium diffusion preparation [150], acquisition with a small field of view (FOV) [151, 165, 166], cardiac gating [151, 159, 167], estimating 3D phase errors using data from multiple shots [156, 157]. However, these methods still retain residual artefacts, especially when diffusion encoding is along the superior-inferior direction, the direction in which brain pulsation is most significant [168–170].

To overcome the limitations faced by 2D and 3D acquisition, 3D multi-slab acquisition has been proposed for high-resolution dMRI [59, 158, 159]. In 3D multi-slab acquisition, the whole volume is divided into multiple smaller slabs, where each slab is excited and encoded separately (i.e. each slab is in effect a reduced-FOV volume along the slice direction). As a 3D imaging method, it can easily achieve thin slices by increasing phase encoding lines in the slice direction. With proper choice of slab thickness, motion-induced 3D phase errors in each slab can be well approximated by a 2D navigator [92, 158]. This hybrid method partly removes the coupling between number of slices and TR: although TR is proportional to the number of slabs, the slice definition within each slab is accomplished using 3D k-space encoding. The number of slabs required to cover a given volume is much less

than the number of slices and is primarily dictated by the thickest slab that can be phase corrected with a 2D navigator. In practice, the whole brain can be covered in 12-14 slabs, for which the TR is highly compatible with maximising SNR efficiency.

4.1.1 Overview of Slab-Boundary Artefacts

A remaining challenge for 3D multi-slab acquisition is the slab-boundary artefact [171], which stems from the shape of the RF profile and appears as periodic signal modulation in the slice direction. Ideally, the RF profile of a slab would have a rectangular shape and the width of the slab thickness, but this requires RF pulses to be infinitely long. In reality, RF pulses are truncated, resulting in a non-rectangular RF profile, including magnitude variation inside the slab, transition bands and sidelobes. Therefore, the slab edges have lower signal magnitude than the slab centre. The transition bands and sidelobes can extend to the neighbouring slabs, introducing slab crosstalk and aliasing (Fig. 4.1).

Several methods have been proposed to reduce the slab-boundary artefacts. The most widely used methods include oversampling in the slice direction, deliberately overlapping the adjacent slabs[82, 171, 172]), separating acquisitions for odd/even slabs [159] and applying a (weighted) slab combination at the slab-boundary region [59, 82, 158, 159, 171–173]. These methods work well in dMRI if long TRs ($\sim$4s) are used together with large amounts of slab overlap, but exhibit residual artefacts in shorter TR acquisitions [161]. A variation of the overlapping method shifts the slab by one slice position once all kz phase encoding steps for a designated subset of ky phase encoding steps are acquired, to transform the signal modulation in the slice direction to the ky phase encoding direction, using signal demodulation to reduce ghosting artefacts [174, 175]. This method has also been demonstrated in non-Cartesian sequences [176–178]. However, if the signal modulation is not properly corrected, this method may retain ghosting artefacts. Another penalty with all these methods is the reduction of scan efficiency due to oversampling,

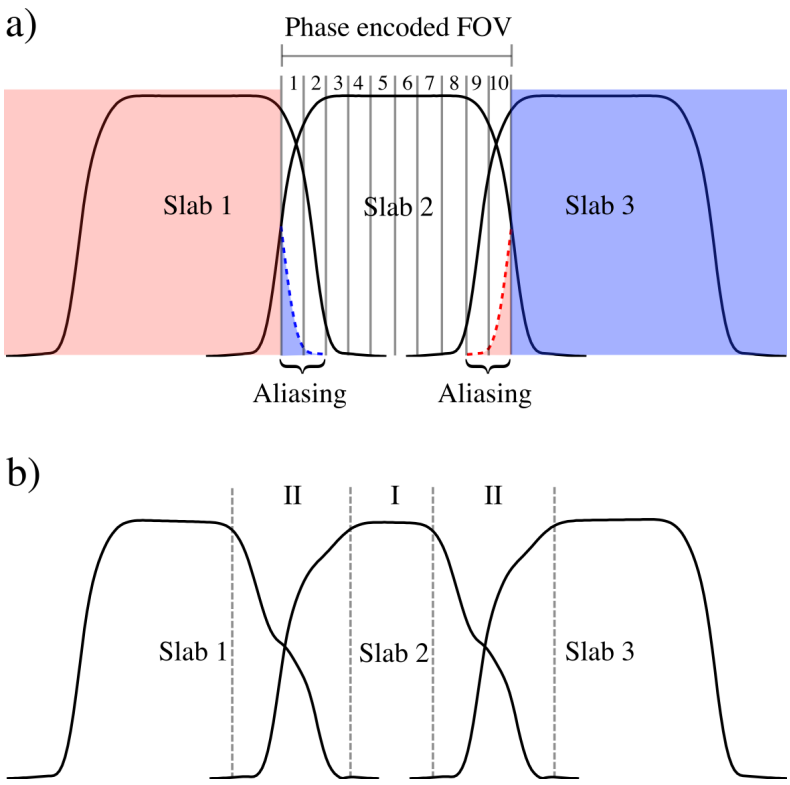


Figure 4.1: (a). Illustration of slab aliasing. The multi-slab acquisition shown here excites 3 slabs with 10 k_z phase-encodes for each slab. Extension of the slab profile is wider than the phase-encoded FOV (also the predefined slab thickness), leading to aliasing artefacts. (b). Illustration of slab-crosstalk effect. The slab profile extends to its adjacent slabs, exacerbating T_1 saturation effects at slab boundaries where adjacent slabs overlap. In this example with 3 slab acquisition, magnetization at slab centre (region I) is excited every TR, while magnetization at slab boundary (region II) is excited with intervals of 2/3TR - 1/3TR - 2/3TR - 1/3TR $\cdots$, leading to decreased signal at slab boundaries and asymmetry in the shapes. Slab profiles in (a) and (b) were generated using Bloch simulation with TR=20s and 2s respectively, and interleaved acquisition order. $T_1/T_2 = 840/70$ms (WM at 3T is assumed). Details of other simulation parameters are described in the Methods.

overlapping or discarding some slices. A histogram-matching method was proposed to correct the slab-boundary artefacts without increasing scan time, which assumes the slab-boundary artefact is a global intensity variation in the slice direction [179]. However, when this assumption is not valid, residual artefacts remain after correction. Recently, a slab-profile encoding (PEN) method was proposed to minimize slab-boundary artefacts with a minor increase in scan time [160]. PEN treats multi-slab acquisition as a linear encoding problem, which can be solved by linear inversion

approaches (e.g. least squares/pseudo inverse) as used in SENSE [30], with the slab profiles effectively substituting coil sensitivity profiles. With slab profiles estimated from calibration data and long TRs (∼4s), PEN can effectively correct aliasing, but some residual artefacts remain due to slab crosstalk, especially at shorter TRs required to maximize SNR efficiency(Fig. 4.2).

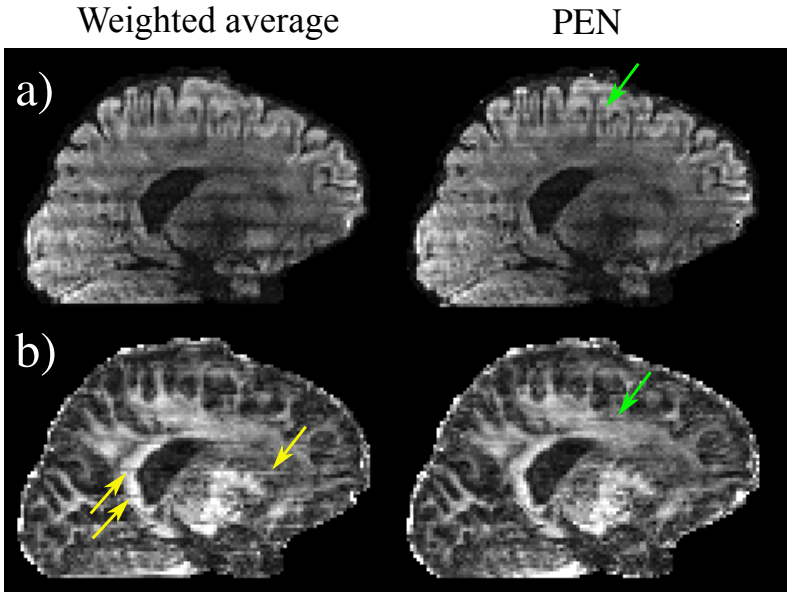


Figure 4.2: Slab-boundary artefacts in DWI images and FA maps. A short TR of 2s is used in data acquisition here, which is associated with strong slab crosstalk. Weighted average of all slabs results in strong artefacts in both the DWI image and FA map, which can be reduced by PEN correction (e.g. the artefacts indicated by the yellow arrows in the FA maps are reduced by PEN correction). However, due to the strong slab crosstalk, some artefacts remain in PEN results (indicated by the green arrows).

Here, we present a method to minimize the slab-boundary artefacts for 3D multi-slab dMRI with SNR-optimal TRs (∼2s). We first present detailed simulations demonstrating the slab-boundary artefact is a combination of slab crosstalk and aliasing. While aliasing has been explained and addressed in PEN [160], slab crosstalk effects have not been studied in detail. We characterize the effect of multiple factors (acquisition scheme, TR, T_1, $\mathbf{B}_0$ inhomogeneity) on slab crosstalk and how it may alter the measurement of diffusion metrics, despite signal normalization by b=0 images as part of quantification. We then introduce a nonlinear inversion extension of PEN, which we call NPEN, where the slab profile encoding is formulated

as a nonlinear optimization. Two constraints are added, which enforce in-plane smoothness on the slab profile and suppress frequency components corresponding to inter-slab distance. Results from simulation and data acquired in vivo demonstrate the proposed method can effectively reduce the slab-boundary artefacts for 3D multi-slab dMRI with a short TR of 2s, a regime in which weighted average (WA) and PEN reconstructions contain significant residual slab-boundary artefacts.

4.2 Bloch Simulations

This section is focused on simulation of slab crosstalk, including its causes and how it affects diffusion measurements. These results are crucial motivation for the correction method presented in the following section, where aliasing effects are also discussed.

4.2.1 Slab Crosstalk in Multi-slab Imaging

Fig. 4.1b depicts a Bloch simulation of a multi-slab acquisition, including both RF profile and T_1 saturation effects. The central part of each slab is not overlapped with adjacent slabs, so the magnetization in this region experiences an RF pulse every TR. At slab boundaries, the magnetization is excited every $\sim$TR/2, but with both the timing and flip angle varying for the target and adjacent slab excitations, which exacerbates T_1 saturation and results in lower signal. This effect is called slab crosstalk.

The precise nature of slab crosstalk is affected by a range of factors, including RF profiles, excitation order, TR, T_1 and $\mathbf{B}_0/\mathbf{B}_1$ inhomogeneity. To obtain a quantitative understanding about how these factors affect the slab crosstalk, we performed Bloch simulations of the spin-echo sequence used in this work [53]. The simulations utilized 600 isochromats in each voxel, which are summed to estimate the MR signal. We simulated a range of TRs and acquisition schemes (sequential and interleaved slab ordering), the three major types of tissue in human brain

(white matter=WM, grey matter=GM and cerebrospinal fluid=CSF), and a range of off-resonance frequencies. Simulations were matched to the pulse-sequence design used for the experiments performed in vivo described in later sections. The sequence used 90° and 180° RF pulses designed with the SLR algorithm [180] with time-bandwidth product of 20 and 8 respectively, which were chosen to achieve a sharp composite RF profile within RF power limits (imposed by hardware and/or SAR). The RF pulse waveforms are plotted in Fig. 4.3. The duration for excitation pulse and refocusing pulse were 7.18ms and 10.24ms, respectively. The same RF pulses are also used for the acquisition of in vivo data. The echo times for the first and second echo (phase navigator) were 78ms and 122ms. T_1 and T_2 values for 3T were assumed: WM 840/70ms, GM 1320/110ms, CSF 3000/2000ms [16, 181, 182]. In all simulations, a 15-mm slab thickness was chosen as that allows whole-brain coverage with TR=2s while the slabs remain thin enough to be corrected with 2D navigators [158]. Each slab contained 10 slices, 20% of which were overlapped with slices from adjacent slabs to favour (N)PEN reconstruction, i.e. two slices on each side.

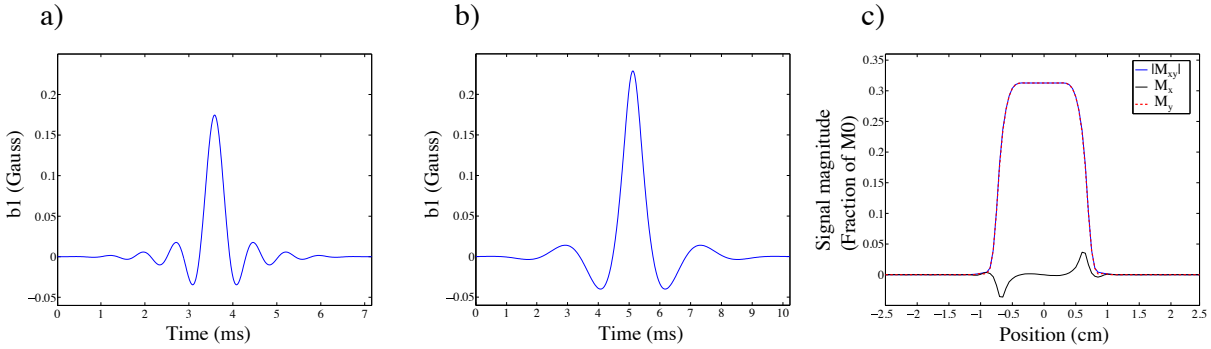


Figure 4.3: The waveforms of the excitation pulse (a) and the refocusing pulse (b) used in the simulations and scans performed in vivo. (c). The simulated signal profile with 1 slab, 2s TR.

The effects of TR, acquisition schemes and $\mathbf{B}_0$ inhomogeneity on slab crosstalk are demonstrated in Fig. 4.4. Fig. 4.4a compares the simulated signal profile with different TR, showing reduced saturation effect at slab boundaries with longer TR, although in practice this comes at the significant cost of SNR efficiency. Fig. 4.4b

compares acquisitions with different slab-ordering schemes. Sequential acquisition suffers from substantial signal loss at one side of the slab, while interleaved acquisition shows high signal intensity at both sides of the slab. Interleaved acquisition with an odd number of slabs produces almost identical signal profile for the odd and even slabs (Fig. 4.4c), while in the acquisition with an even number of slabs, the signal profile for the odd slab varies from the even slab (Fig. 4.4d). In the following simulations and in vivo scans, interleaved acquisition with an odd number of slabs was used to achieve homogeneous slab profiles with a well-defined periodicity, which facilitates the NPEN reconstruction (described below). Fig. 4.4e compares the slab profiles for WM, GM and CSF, each normalized by the signal magnitude of the centre slice to emphasize differences in shape due to T_1 saturation effects.

To simulate the effects of $\mathbf{B}_0$ inhomogeneity, Bloch simulation was performed with off-resonance frequencies starting from -250Hz to +250 Hz with 5 Hz steps. The signal profile of one slab in WM is shown in Fig. 4.5. The profile not only shifts but also becomes distorted, resulting in more signal loss at the slab boundaries (Fig. 4.4f). The distortion is caused by the difference in bandwidths of the excitation and refocusing pulses. It should be noted that a constant off-resonance frequency for each slab is assumed in the $\mathbf{B}_0$ inhomogeneity simulation, but in reality, the off-resonance frequencies are spatially varying, which can distort the signal profile even with matched bandwidths of the excitation and refocusing pulses.

4.2.2 Propagation of Slab Crosstalk into Diffusion Quantification

The practical importance of variations in signal intensity at slab boundaries will depend on whether these artefacts alter diffusion quantification, based on signal attenuation relative to a non-diffusion weighted (b=0) reference. Multiplicative slab artefacts that attenuate diffusion-weighted and un-weighted signals equally would be removed when normalizing diffusion-weighted images by the b=0 reference.

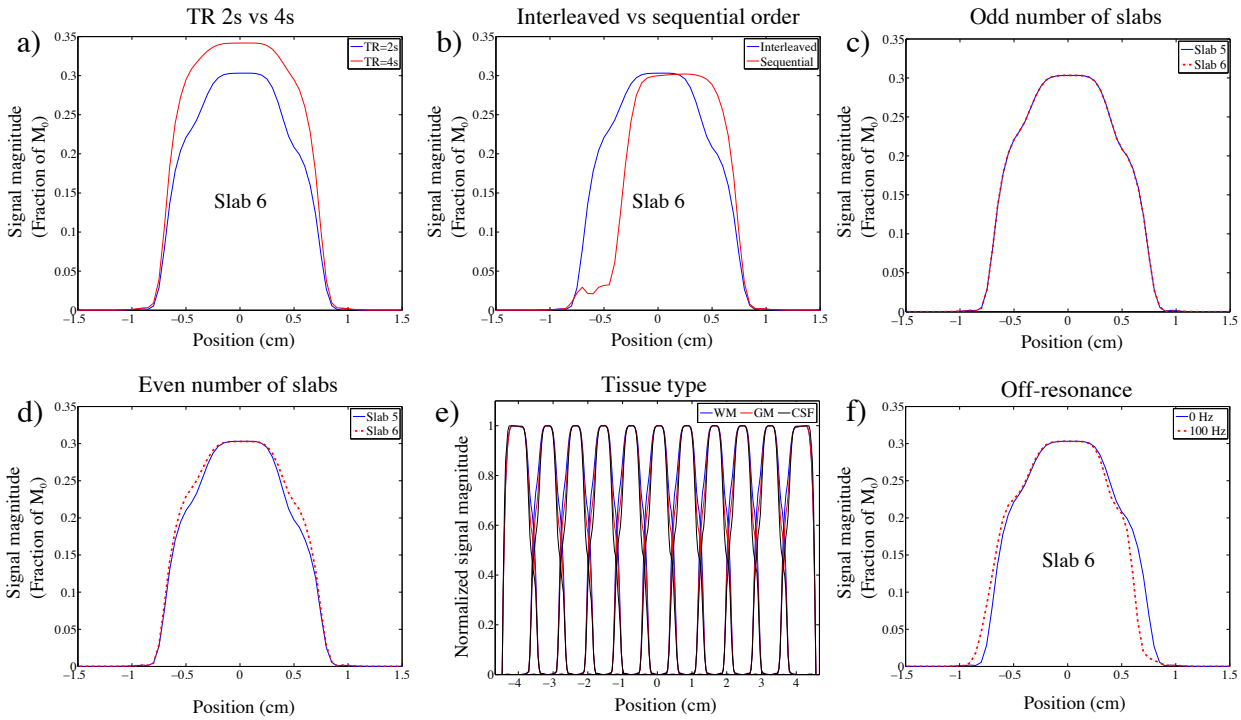


Figure 4.4: Results of Bloch simulation. (a) - (d) compare the signal profiles of WM generated from Bloch simulation using different TR, acquisition scheme and number of slabs: (a) 11 slabs, sequential acquisition with TR = 2s (blue) and TR = 4s (red). (b) TR = 2s, 11 slabs with interleaved acquisition (blue) and sequential acquisition (red). (c) TR = 2s, interleaved acquisition with 11 slabs, the signal profile of slab 5 (blue) and slab 6 (red) are overlaid for better comparison. (d) TR = 2s, interleaved acquisition with 12 slabs, the signal profile of slab 5 (blue) and slab 6 (red) are overlaid. The effects of TR, acquisition scheme and number of slabs on GM and CSF are similar to WM and not shown here. (e). Slab profiles for WM, GM and CSF, which were generated by normalizing the MR signal profile with the signal magnitude of the centre slice. (f) shows centred signal profiles simulated with 0 Hz and 100 Hz off-resonance frequencies.

In this case, the only effect on diffusion quantification would be SNR reduction at slab boundaries. Conversely, non-multiplicative effects could create a more significant artefact by introducing bias in diffusion metrics at different regions of slab profiles. For example, mean diffusivity (MD) or fractional anisotropy (FA) could exhibit a consistent bias at slab boundaries. One potential cause of non-multiplicative attenuation is "partial volume" effects, in which both relaxation and diffusion alter the signal to create a more complicated signal dependence than either effect on its own.

To demonstrate non-multiplicative effects in diffusion metrics due to partial

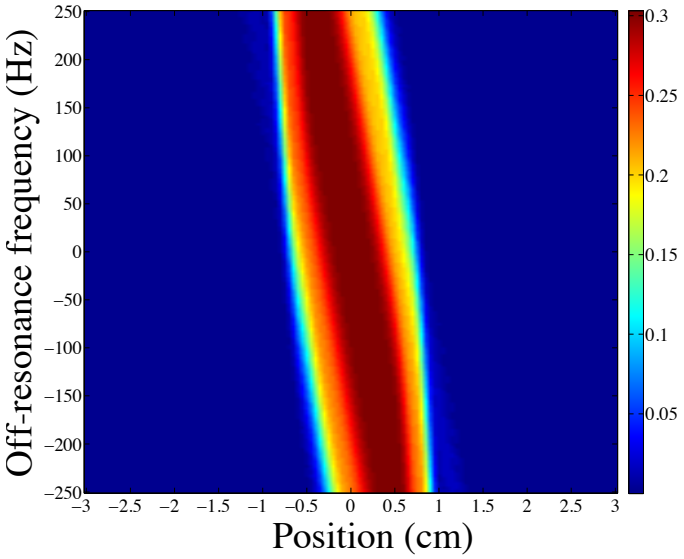


Figure 4.5: WM signal profile of one slab simulated at different off-resonance frequencies.

volume, we simulated a two-compartment diffusion-weighted signal corresponding to WM/GM and WM/CSF respectively [183]:

$$S(b, \mathbf{v}) = \sum_{i=1}^{2} S_i(b=0) f_i \exp(-b\mathbf{v}^T \mathbf{D}_i \mathbf{v}) \tag{4.1}$$

Where f_i and $\mathbf{D}_i (i = 1, 2)$ are the volume fractions and diffusion tensors for the two compartments and $\mathbf{v}$ is the diffusion gradient vector. The b=0 signal $S(b = 0)$ is calculated for each tissue type seperately using Bloch simulations and the diffusion-weighted signal $S(b, \mathbf{v})$ is subsequently calculated from the partial-volume model. The diffusion tensor eigenvalues $[\lambda_1, \lambda_2, \lambda_3]$ for WM $[1.4, 0.35, 0.35]$ (FA=0.707), GM $[0.7, 0.7, 0.7]$ and CSF $[3.0, 3.0, 3.0]$ were used [183]. A set of 6 non-collinear directions was assumed for $\mathbf{v}$. A wide range of b values (1 – 10000 s/mm^2) was investigated. We calculated MD and FA for a simulated voxel with a given partial volume, considering these as broadly representative of bias in diffusion metrics (although specific model parameters will vary in their sensitivity to slab-boundary artefacts). The goal of these simulations is to identify any consistent bias in diffusion metrics for the slab boundary relative to the slab centre. A relative error was calculated to analyze the slab crosstalk effect on MD ($\epsilon_r = |\theta_c - \theta_b|/\theta_c$, where θ_c and θ_b are the MD values at

slab centre and boundary). As the FA values of CSF and GM approach zero when their fractions go to 1, leading to inflated measure of relative errors, we calculated the absolute FA errors between slab centre and boundary instead.

Fig. 4.6 presents the results of FA simulation from the partial-volume mode. The absolute errors depend on the partial volume ratio of the tissues and the b-values used in the simulation. For WM/CSF partial-volume (Fig. 4.6a,b), the absolute FA error approaches its maximum value of ∼0.045 with WM percentage of ∼7% and b-value smaller than 1000 s/mm^2. For WM/GM partial-volume (Fig. 4.6c,d), the maximum absolute FA error is ∼0.02. In both cases, the FA errors decrease when b-value goes up, as more isotropic-diffusion compartments (GM and CSF) are suppressed due to their higher diffusion coefficients.

Fig. 4.7 shows the relative errors for MD from WM/CSF partial-volume, which are on the order of a few percent with maximum value of ∼8%. For WM/GM partial-volume, the relative errors for MD are negligible (image not shown). Importantly, these simulations of partial volume effects demonstrate that slab-profile effects in voxels with partial volume cannot be trivially removed by normalization during the calculation of diffusion parameters.

4.3 Methods

This section describes a method proposed to correct slab-boundary artefacts. Correction results are demonstrated using numerical simulation and in vivo data.

4.3.1 Proposed Correction

The simulations presented in the previous sections demonstrate that slab-boundary artefacts depend on the details of the tissue environment, including relaxation times and field inhomogeneity. Furthermore, these effects cannot be assumed to be removed by normalization with a b=0 image in the likely case of tissue partial volume. Finally, these effects are considerably worse in the SNR-optimal regime of

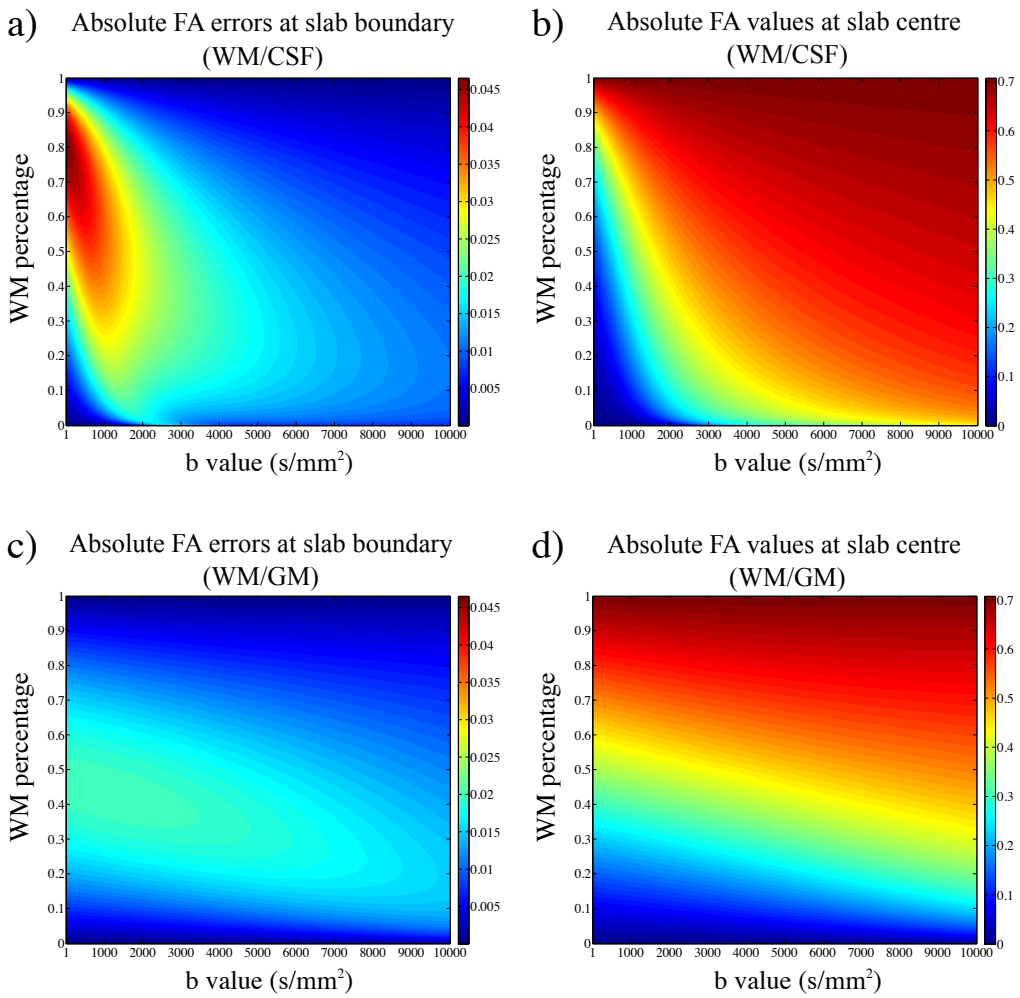


Figure 4.6: Simulation of FA errors caused by slab crosstalk. (a) and (b) are simulated based on WM/CSF partial volume model: (a) The absolute FA errors (differences of FA values between slab boundary and slab centre) as a function of WM percentage and b value. (b) The absolute FA values at slab centre. (c) and (d) are simulated from WM/GM partial volume model:(c) The absolute FA errors as a function of WM percentage and b value. (d) The absolute FA values at slab centre.

1-2s. The accuracy of 3D multi-slab dMRI thus depends on finding a method for reducing the impact of these artefacts on quantified diffusion values. It is useful to begin by reviewing slab aliasing and the previously proposed PEN method [160]. We then present a new method based on PEN that aims to deal with the additional slab crosstalk issues raised by the simulations above.

The PEN method treats the 3D multi-slab acquisition as a linear encoding

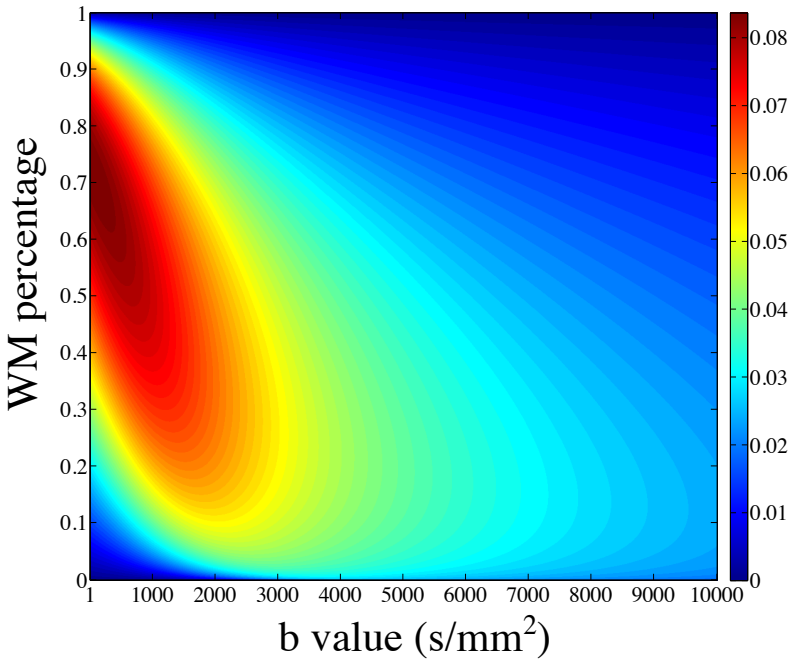


Figure 4.7: Relative MD errors as a function of WM percentage and b value simulated with WM/CSF partial volume. For WM/GM partial-volume, the relative errors for MD are negligible (image not shown).

process, which can be expressed as:

$$\mathbf{MFCSu} = \mathbf{d} \tag{4.2}$$

where **u** is the unknown 3D image, **S** describes the slab profiles, **C** describes the coil sensitivities, **F** is the Fourier transform operator, **M** captures the k-space sampling trajectory, and **d** is the acquired multi-slab k-space data. This is very similar to the general parallel imaging model, with the difference that, here, the image is weighted by both coil sensitivity and slab profile. Aliasing artefacts arise as a result of "undersampling" in the slice direction, where we here consider the effect of magnetization excited by a given pulse throughout the entire brain volume. The quality of PEN reconstruction is critically dependent on the ability to estimate the slab profiles accurately. In several studies, simulated slab profile was used to correct slab-boundary artefacts [160, 184]. However, as shown in Fig. 4.4, the slab crosstalk is affected by multiple factors, making it difficult to estimate slab profiles accurately using Bloch simulation. An alternative approach is to measure

slab profiles from a calibration scan, which is a multi-slab scan at the same slab locations but oversampled in the slice direction to minimize aliasing. The slab images were first zero padded to the full FOV, and then a combined volume was generated by sum-of-squares of all the zero-padded slabs. Finally, the slab profile is estimated by normalizing each of the zero-padded slab images by this combined single-slab image [160]. This method will not increase the scan time significantly, but the reconstructed images will be modulated by the sum-of-squares of the slab profiles [30, 160], which could be relatively flat in the slice direction with a long TR (∼4s), but will result in significant signal variation with a short TR (∼2s).

As originally proposed, the PEN method relies on the assumption that slab-profile estimates are accurate, which is increasingly difficult to ensure at short TR. To overcome this problem, we propose a new nonlinear reconstruction that jointly estimates the slab profile and the image. In this method, both slab profile and image are treated as unknowns

$$\mathbf{x} = \begin{bmatrix} \mathbf{u} \\ \mathbf{s}_1 \\ \mathbf{s}_2 \\ \vdots \\ \mathbf{s}_{N_{slab}} \end{bmatrix} \tag{4.3}$$

which satisfies:

$$E(\mathbf{x}) = \begin{bmatrix} \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_1 \cdot \mathbf{u}\} \\ \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_2 \cdot \mathbf{u}\} \\ \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_3 \cdot \mathbf{u}\} \\ \vdots \\ \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_{N_{slab}} \cdot \mathbf{u}\} \end{bmatrix} = \mathbf{d} \tag{4.4}$$

Where E is a nonlinear operator mapping image ($\mathbf{u}$), slab profile ($\mathbf{s}$) and coil sensitivity ($\mathbf{c}$) to the acquired data ($\mathbf{d}$). The dot operator " $\cdot$ " represents element by element multiplication. This poses the 3D multi-slab reconstruction as a nonlinear inversion problem, which we refer to as nonlinear inversion for slab profile encoding (NPEN). As shown in [185–187], this problem can be solved using

an iteratively regularised Gauss-Newton algorithm (IRGN) [188–190]. In step n, Eq. 4.4 is first linearized around x_n:

$$E'(\mathbf{x}_n)\Delta\mathbf{x}_n = \mathbf{d} - E(\mathbf{x}_n) \tag{4.5}$$

where $E'(\mathbf{x}_n)$ is the Jacobian of E at the current guess $\mathbf{x}_n$, $\Delta\mathbf{x}_n$ is the update, and the new guess can be obtained by $\mathbf{x}_{n+1} = \mathbf{x}_n + \Delta\mathbf{x}_n$. Then the linear problem described by Eq. 4.5 is solved by the following Tikhonov regularized minimization:

$$\Delta\mathbf{x}_n = \arg\min_{\Delta\mathbf{x}_n} \left\| E'(\mathbf{x}_n)\Delta\mathbf{x}_n - (\mathbf{d} - E(\mathbf{x}_n)) \right\|_2^2 + \alpha_n \left\| \mathbf{x}_n + \Delta\mathbf{x}_n - \mathbf{x}_0 \right\|_2^2 \tag{4.6}$$

Where the first term is the data fidelity term and the second term is the regularisation term. x_0 is an initial guess and α_n is a regularisation parameter, which is reduced in each step to benefit from the robustness at the initial steps of iterations (Gradient-Descent like) and the fast convergence when $\mathbf{x}$ is close to the solution (Gauss-Newton like). Eq. 4.6 can be solved with the Conjugate Gradient algorithm [191]. Calculation of the Jacobian matrix $E'(\mathbf{x}_n)$ and its adjoint $E'(\mathbf{x}_n)^{\mathrm{H}}$ can be found in Appendix B.

As shown in previous work [185, 187], more accurate estimation of the solution to Eq. 4.4 can be achieved by adding further constraints in the reconstruction. In this work, we consider two constraints. First, the slab profile is constrained to be smooth (in-plane), which is similar to the polynomial fitting process in the slab profile estimation in PEN. The justification for this constraint is the observation that the slab profiles for WM, GM and CSF are very similar (Fig. 4.4e). To realise this constraint, the preconditioning method used in [185] is applied, where $\mathbf{x}$ is replaced by $\hat{\mathbf{x}} = \mathbf{U}\mathbf{x}$ with $\mathbf{U}$ defined as

$$\mathbf{U} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_s\mathbf{F} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{W}_s\mathbf{F} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{W}_s\mathbf{F} \end{bmatrix} \tag{4.7}$$

$\mathbf{W}_s$ is a matrix used to penalize high spatial frequency k_x-k_y-components of $\mathbf{s}$ in each k_z plane. The preconditioned system equation becomes $E(\mathbf{U}^{-1}\hat{\mathbf{x}}) = \mathbf{d}$, which

can be written as $G(\hat{\mathbf{x}}) = \mathbf{d}$ for simplicity. The preconditioned reconstruction for image update is then expressed as

$$\Delta\hat{\mathbf{x}}_n = \underset{\Delta\hat{\mathbf{x}}_n}{\arg\min} \left\| G'(\hat{\mathbf{x}}_n)\Delta\hat{\mathbf{x}}_n - (\mathbf{d} - G(\hat{\mathbf{x}}_n)) \right\|_2^2 + \alpha_n \left\| \hat{\mathbf{x}}_n + \Delta\hat{\mathbf{x}}_n - \hat{\mathbf{x}}_0 \right\|_2^2 \tag{4.8}$$

The smoothness constraint is achieved using the regularization term [185], which can be expressed explicitly as:

$$\|\hat{\mathbf{x}} - \hat{\mathbf{x}}_0\|_2^2 = \|\mathbf{u} - \mathbf{u}_0\|_2^2 + \sum_{i=1}^{N_{slab}} \|\mathbf{W}_s\mathbf{F}(\mathbf{s}_i - \mathbf{s}_{0,i})\|_2^2 \tag{4.9}$$

We use a smooth initialization for the slab profile and choose a $\mathbf{W}_s$ with high weighting at outer k-space (high spatial frequency) and low weighting at k-space center (low spatial frequency) to impose a smoothness constraint on the slab profiles.

The second constraint exploits the fact that slab-boundary artefacts are approximately periodic in image space. Therefore, residual artefacts would correspond to spikes in at known frequencies. By suppressing these spurious k-space signals, images with reduced slab-boundary artefacts can be reconstructed. This constraint was transformed as an additional penalty term in the reconstruction:

$$\Delta\hat{\mathbf{x}}_n = \underset{\Delta\hat{\mathbf{x}}_n}{\arg\min} \left\| G'(\hat{\mathbf{x}}_n)\Delta\hat{\mathbf{x}}_n - (\mathbf{d} - G(\hat{\mathbf{x}}_n)) \right\|_2^2 + \alpha_n \left\| \hat{\mathbf{x}}_n + \Delta\hat{\mathbf{x}}_n - \hat{\mathbf{x}}_0 \right\|_2^2 + \beta_n \left\| \mathbf{W}_u\mathbf{F}\mathbf{u}_n \right\|_2^2 \tag{4.10}$$

where $\mathbf{W}_u$ is a weighting matrix with large values at the spike frequencies and small values at other frequencies such that $\|\mathbf{W}_u\mathbf{F}\mathbf{u}_n\|_2^2$ penalizes residual slab-boundary artefacts. β_n is a regularisation parameter controlling the strength of this constraint.

4.3.2 Numerical Simulation

To test the reconstruction algorithm under controlled conditions, we generated realistic simulations of spin-echo images based on the 'MNI-Colin27' high-resolution multi-modal brain atlas (very high-SNR T_1- and T_2-weighted scans of a single subject) [192]. Artificial 3D multi-slab data were generated based on the following

pipeline. Tissue partial volume maps were estimated using FAST [193] based on the T_1-weighted images. The tissue-specific slab profiles calculated from Bloch simulations were weighted by the partial volume maps and summed to generate the integrated slab profile. The T_2-weighted images were weighted by the integrated slab profile and Fourier transformed to produce k-space data. Finally, the k-space data were sampled in a way that is consistent with the acquisition of 11 slabs with 10 slices per slab to simulate aliasing effects. This simulated 3D multi-slab data were then reconstructed using the candidate methods (WA, PEN and NPEN, details below) and compared to the ground truth (the original T_2-weighted atlas).

To compare the performance of PEN and NPEN at different TR, we performed a Bloch simulation at TRs of 2s, 4s, and 7s respectively using SLR pulses designed with the same TBWP as used in the PEN work [160]. The calculated slab profiles were used to generate simulation data in the way described above, and then the data were reconstructed using PEN and NPEN.

4.3.3 In vivo

A diffusion weighted readout-segmented EPI sequence [49, 53, 54] was modified to acquire 3D multi-slab data. This sequence was further modified to oversample the central segment of b=0 data by a factor of two. This oversampled b=0 data could be used to estimate slab profiles for PEN or calculate initial guesses for NPEN, alleviating the need for a separate set of calibration data.

Data were acquired from four healthy subjects with informed consent in accordance with local ethics. For each subject, 9 slabs with 10 slices per slab were acquired and adjacent slabs were overlapped by 2 slices (20% overlapping), which resulted in 74 slices in the final reconstruction. k-space coverage of each slab was achieved using a 3D adaptation of the readout-segmented EPI approach [59]. The matrix size in each k_z plane was 146×146, using 5 readout segments and 3/5 partial Fourier (i.e. only 3 segments were acquired) [58]. Isotropic spatial resolution 1.5×1.5×1.5 mm^3

was acquired with a FOV of 220×220×111 mm^3. The same echo times (78ms, 122ms) and RF pulses were used as in the Bloch simulation and TR was 2s. The scan time for one diffusion direction was 1.3 min with central-segment oversampling (the first b=0 scan only) and 1 min without central-segment oversampling (all other volumes).

Diffusion-weighted images were acquired with b=1000 s/mm^2 using a Stejskal-Tanner diffusion preparation and 20 isotropically-distributed diffusion encoding directions, and two b=0 images. Only the first b=0 data were oversampled for estimating the slab profile, which was then applied to the reconstruction of all other data (such that calibration data represent a constant overhead equivalent to a single volume scan). Total acquisition time was 22.3 min. Note that readout-segmented EPI is used here, which trades off reduced distortion and blurring for longer scan times. If in-plane single-shot EPI is applied, the scan time will be reduced to ∼20s per volume .

4.3.4 Reconstruction

The reconstruction of readout-segmented EPI data followed procedures described in previous work [54, 59], except that the oversampled central segments were reconstructed separately. Two sets of images were obtained from this reconstruction: high-resolution 3D multi-slab images and low-resolution 3D multi-slab images with doubled FOV in the slice direction, which were used as calibration images. Three slab-boundary-artefact correction methods, WA, PEN and NPEN, were compared on simulation data and data acquired in vivo.

In NPEN reconstruction, the initial guess for the image was set to zero. The calibration images were used to generate the initial guess for slab profiles as follows. These images contain no significant aliasing artefacts due to oversampling in the slice direction, such that the variation of signal intensity along the slice direction is dominated by slab profile, including sidelobes that lead to saturation effects.

First, each slab image was zero-padded to the full FOV, resulting in image $\mathbf{I}$. The averaged signal for each slice was calculated as

$$\mathbf{I}_{ave}(z) = \sum_{x,y} |\mathbf{I}(x,y,z)\mathbf{Q}(x,y,z)| / \sum_{x,y} |\mathbf{Q}(x,y,z)| \tag{4.11}$$

where $\mathbf{Q}$ is a brain mask generated by FSL's Brain Extraction Tool [194], (x, y) indicates in-plane locations and z indicates the locations in the slice direction. Then the averaged signal was normalized by the signal magnitude at the slab centre

$$\mathbf{I}'_{ave}(z) = \mathbf{I}_{ave}(z) / \mathbf{I}_{ave}(z_c) \tag{4.12}$$

where z_c denotes the location of the centre slice in the slab. Finally, the 3D slab profile estimate S_{est} was generated by repeating $\mathbf{I}'_{ave}(z)$ for all the in-plane locations of S_{est}.

The initial regularisation parameters α_0 and β_0 were set to ensure that after the first iteration, the residual $\|E(\mathbf{x}) - \mathbf{d}\|$ was approximately 3/4 of the initial residual [187]. This resulted in $\alpha_0 = 0.2$ and $\beta_0 = 0.4$ for both simulated data and data acquired in vivo. The decay rates of α and β were set to 1.5, which performed well in the preliminary testing. To achieve better suppression of slab-boundary artefacts, the decrease of β was stopped at 0.03.

PEN reconstruction was implemented as described in [160]. The slab profile was estimated from the calibration images, in which the single-slab image was divided by the sum-of-squares of all slab images. A 2D low-pass Hamming filter was applied in k-space to reduce noise in the estimated slab profiles. In WA reconstruction, a 1D Fermi filter function was applied to each slab by multiplication in the spatial domain to reduce the signal from the edge slices with aliasing, and then the weighted slab images were averaged to generate a 3D single slab image [158, 160].

Performance of different methods with the simulation data was evaluated using root-mean-square-error (RMSE) and difference maps from the ground truth image. To assess the reconstruction performance on different tissue types, RMSE and

difference maps were also calculated for WM, GM and CSF separately using binary segmentation maps generated from partial-volume maps with a threshold of 0.8.

For the in vivo data, eddy current correction was performed in FSL using the first b=0 volume as reference, after which DTI fitting was applied to generate diffusion parameter maps using FSL's diffusion toolbox [195].

4.4 Results

4.4.1 Reconstruction with Simulation Data

The simulations enable us to explore and optimize the reconstruction under realistic but well-defined conditions. For example, to develop suitable stopping criteria. Ideally, we would like to use the minimum RMSE as the stopping criteria of the iteration. However, for in vivo scanning, there is no ground truth image for RMSE calculation, so we need to find other stopping criteria. The RMSEs and image update ($\|\Delta\mathbf{u}\|$) from the atlas-based simulation reconstructions are shown in Fig. 4.8. After the first 3 iterations, $\|\Delta\mathbf{u}\|$ decreases towards its minimum and increases afterwards (Fig. 4.8a). The iteration number corresponding to the minimum $\|\Delta\mathbf{u}\|$ also approximately corresponds to the lowest RMSE for the full image (Fig. 4.8b). This suggests that a stopping criterion based on $\|\Delta\mathbf{u}\|$ would be appropriate in practice (when the ground truth image is not available to calculate RMSE).

However, as shown in Fig. 4.8b, the iteration numbers corresponding to the minimal RMSE of the full image, WM and GM are different, which probably stems from the T_1 dependence of the slab profile. To achieve the best reconstruction of WM, the iteration should be stopped before the minimal RMSE of the full image, or the minimal $\|\Delta\mathbf{u}\|$ in practice. Therefore, in this work, the iteration was stopped with the smallest image update $\|\Delta\mathbf{u}\|$, and the best reconstruction for WM was retrospectively chosen based on visual inspection. Although the final reconstruction

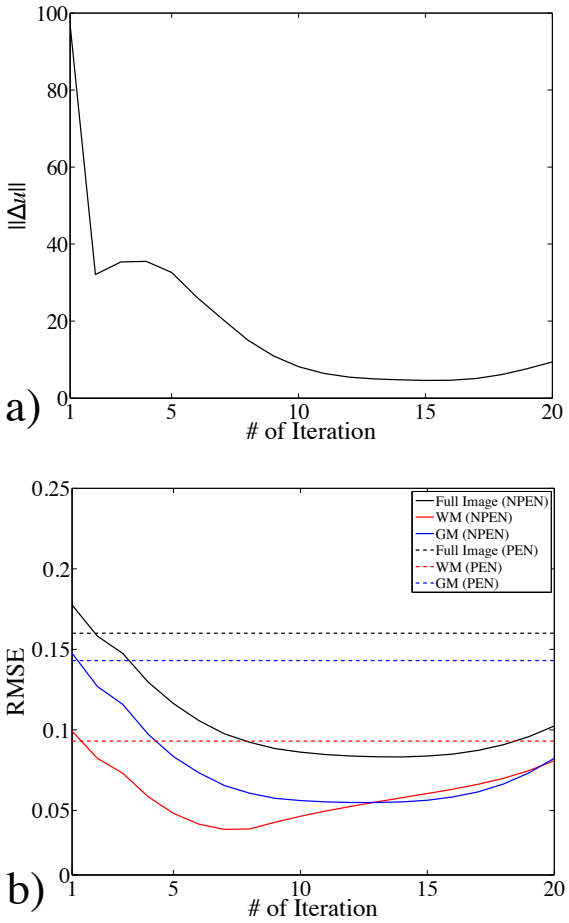


Figure 4.8: (a). The image update at each iteration. (b) RMSE for the full image (black), WM (red) and GM (blue) calculated between the reconstructed image and the ground truth at each iteration: solid lines show the RMSE of NPEN reconstruction and dashed lines indicate the RMSE of PEN reconstruction.

is determined by visual inspection, the stopping criteria based on $\|\Delta\mathbf{u}\|$ can limit the number of iterations and reduce the computation time.

Fig. 4.9 shows the correction results on simulated data. WA and PEN results still retain strong artefacts at slab boundaries, while the NPEN result shows significantly reduced artefact level, revealed by smaller RMSE and minor residual artefacts in the difference map. It is worth noting that the literature on PEN has much lower artefact level [160], which is due to the use of longer TR corresponding to less T_1 saturation but which results in lower SNR efficiency and longer scan time.

Fig. 4.10 shows the tissue-specific comparison between three reconstruction methods. As shown in Fig. 4.10a, the artefacts in WM have been mostly removed

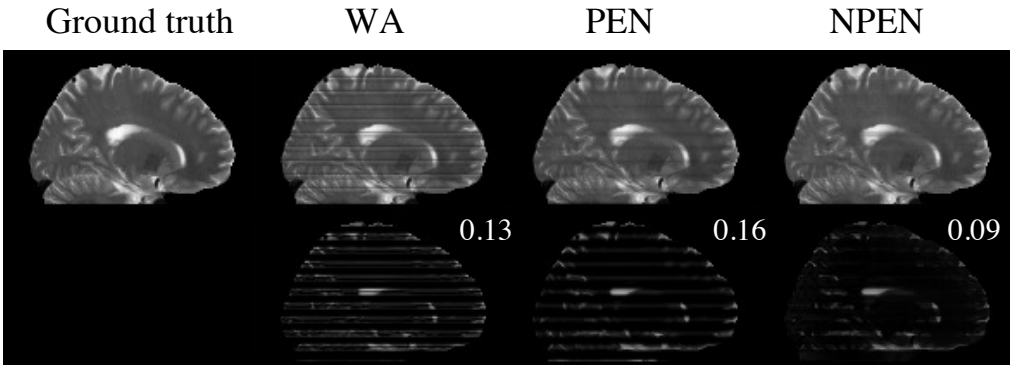


Figure 4.9: Reconstruction results from simulated data. Upper row: ground-truth image and images reconstructed by WA, PEN and NPEN. Lower row: difference maps between the reconstructed images and the ground truth, RMSE is shown on the upper right corner. Difference map is rescaled for better depiction of the error distributions.

by NPEN, which results in a RMSE of 0.03, while WA and PEN results retain high artefact levels, with RMSE of 0.1 and 0.09 respectively. The GM results (Fig. 4.10b) also indicate NPEN performs better than WA and PEN. Fig. 4.10c shows the CSF results, where none of the three methods provides a satisfactory correction, but NPEN still gives the lowest RMSE. The tissue-specific comparison shows that most residual artefacts in NPEN are located in CSF regions, which is not problematic for most dMRI applications.

For a clearer comparison, the WM signal is averaged in each slice and shown in Fig. 4.11 a. NPEN method reduces the signal drops at slab boundaries and generates a signal profile similar to the ground truth, while signal variation along the slice direction is still strong in WA and PEN results. A 1D signal profile in the slice direction passing through the left putamen is shown in Fig. 4.11b, NPEN reconstruction provides the best approximation to the ground truth. Fig. 4.11c shows another 1D signal profile in the slice direction passing through the left ventricle, which is chosen as a worst case for NPEN reconstruction.

The NPEN result exhibits large errors around slice 10, slice 40 and slice 58 (Fig. 4.11c, black arrows), which correspond to CSF regions, while in non-CSF regions, the NPEN result nicely matches the ground truth. In comparison, WA and PEN

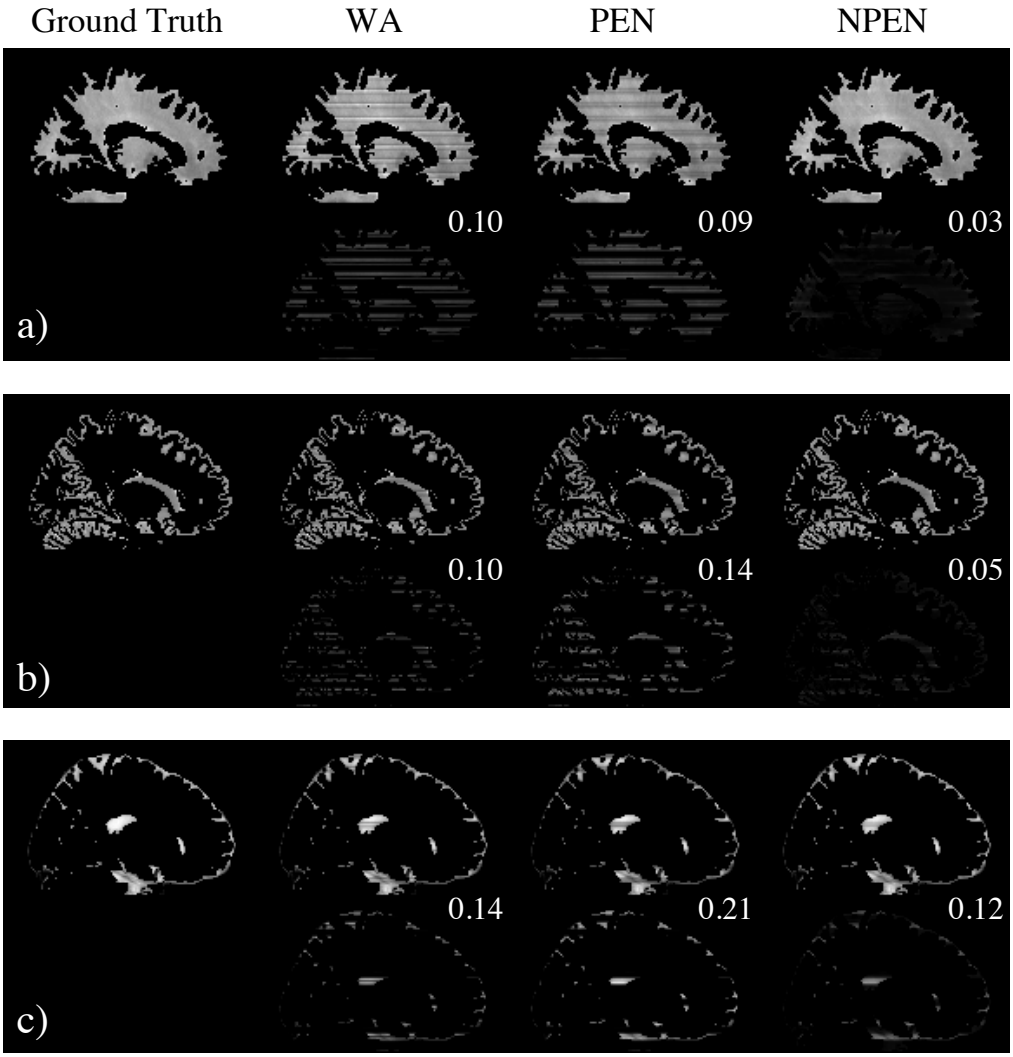


Figure 4.10: Tissue-specific comparison of the reconstruction results from simulation data: (a) WM (b) GM and (c) CSF. For each tissue type, upper row shows the reconstruction results, lower row shows the difference maps between the reconstructed images and the ground truth, RMSE is shown on the upper right corner of the difference map. Difference map is rescaled for better depiction of the residual errors.

results differ more from the ground truth in either CSF regions or non-CSF regions.

Fig. 4.12 shows results of PEN and NPEN at different TR. It can be seen that the residual artefacts in PEN reconstruction at short TR (2s) are very strong but can be reduced at longer TRs (4s and 7s). NPEN can effectively reduce slab-boundary artefacts regardless of TR values.

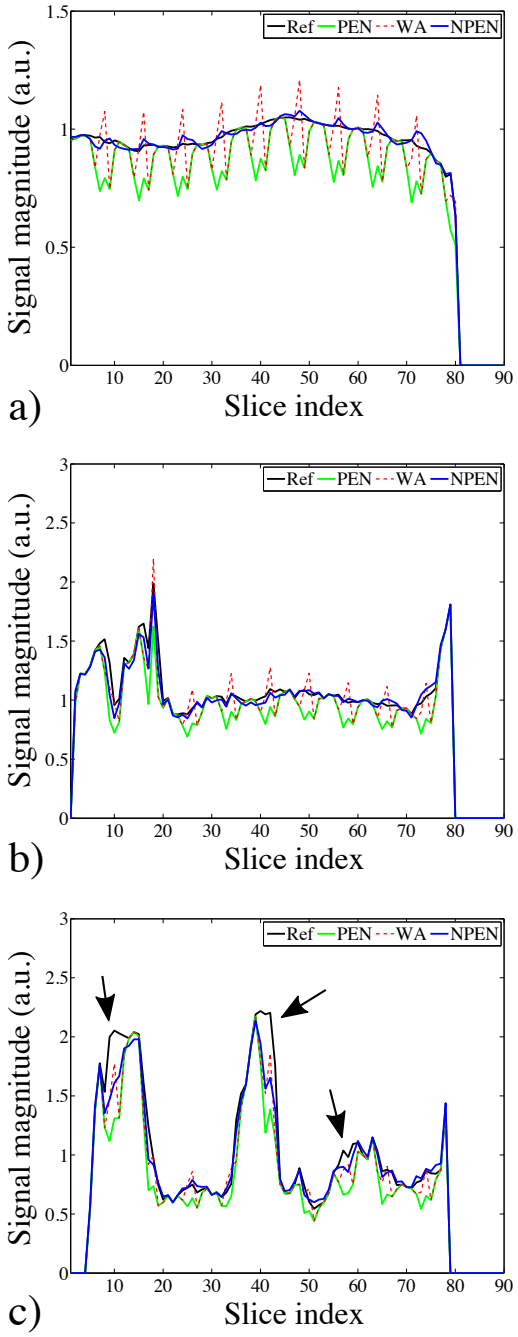


Figure 4.11: (a). Averaged WM signal at each slice. (b). Signal profile along a line passing through left putamen in the slice direction. (c). Signal profile along a line passing through left ventricle in the slice direction.

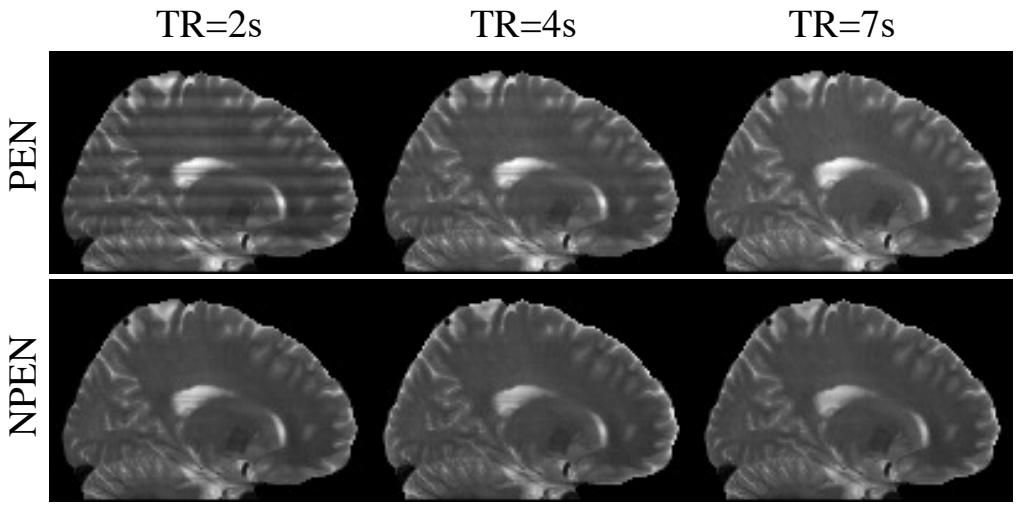


Figure 4.12: PEN and NPEN reconstruction at different TR.

4.4.2 Reconstruction with Data Acquired In Vivo

Fig. 4.13 compares the reconstruction methods on data acquired in vivo. As shown in b=0 images (Fig. 4.13a,e) and diffusion-weighted images (Fig. 4.13b), WA and PEN results still contain strong slab-boundary artefacts, while NPEN effectively reduces the artefact level with only minor residual artefacts visible in CSF regions. As shown in the FA maps (Fig. 4.13c) and color-coded maps of the principal eigenvector (Fig. 4.13d), the WA result still has strong artefacts at slab boundaries even after normalization by the b=0 image. The PEN result shows reduced artefacts in the FA map compared to WA, but retains some residual errors, as indicated by yellow arrows in the zoomed-in FA map (Fig. 4.13f). Given that the primary goal of PEN is to remove aliasing, this suggests that the effect of aliasing artefacts in WA reconstruction is significant. NPEN achieves the best results with only a few visible artefacts in the b=0 images (green arrows), suggesting that NPEN is effective at removing slab-boundary artefacts associated with the RF profile, aliasing and saturation. Reconstructions based on data acquired from a different subject are shown in Fig. 4.14, demonstrating similar performance differences between the methods.

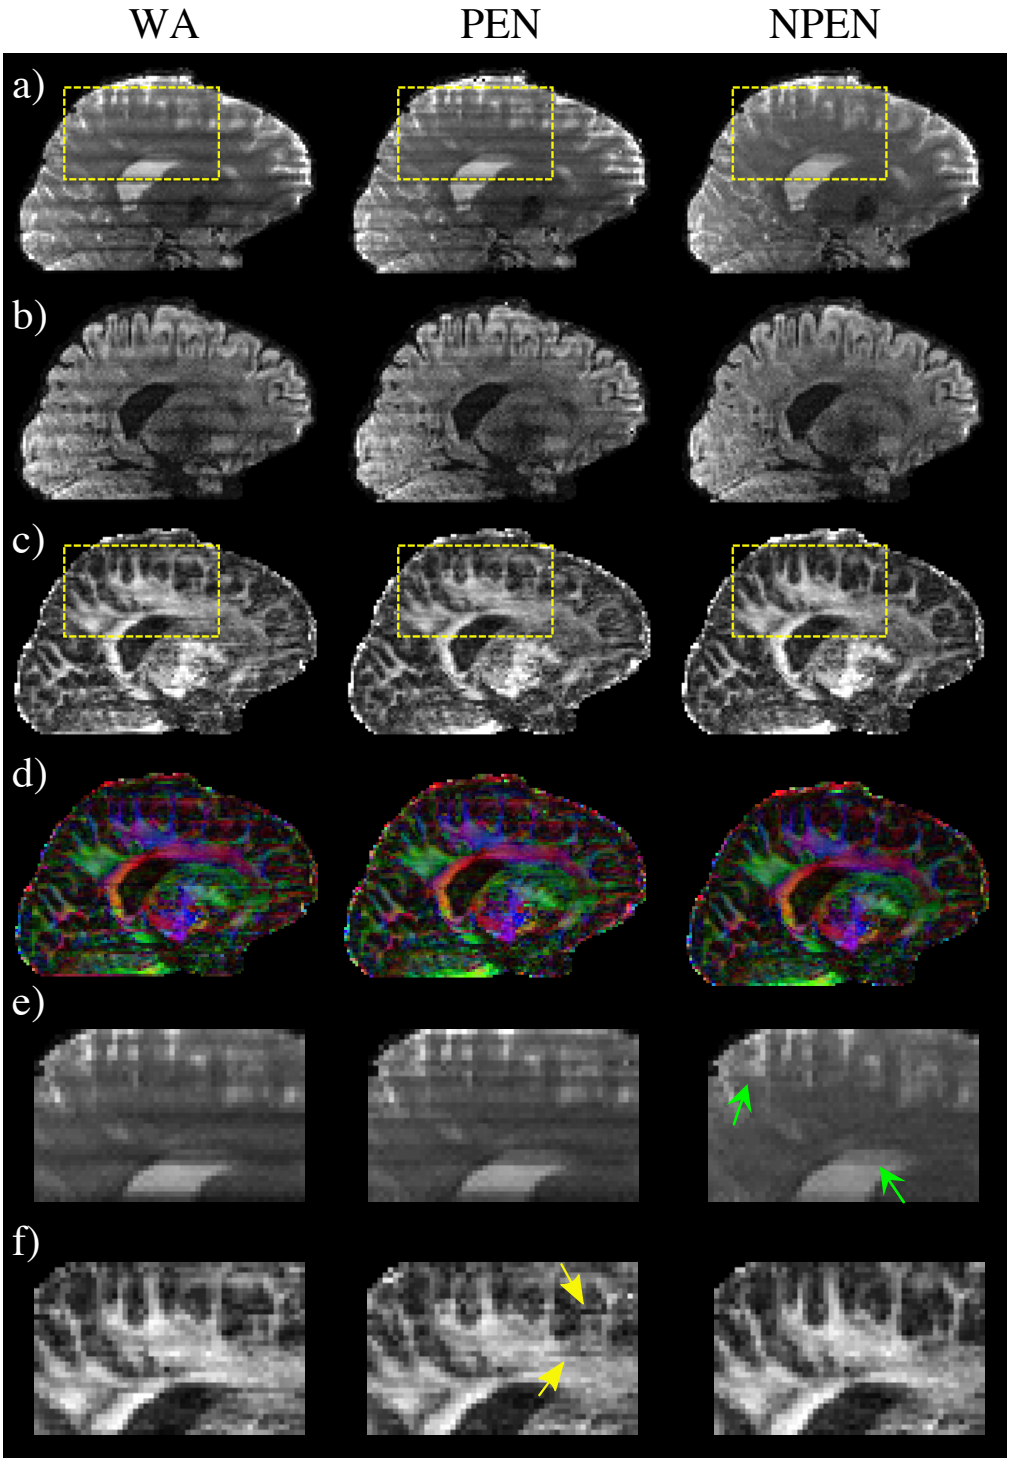


Figure 4.13: Comparison of three reconstruction methods based on data acquired in vivo. (a) b = 0 image. (b) DWI image (b=1000 s/mm^2). (c) FA maps. (d) Color-coded maps of the principal eigenvector (red: right-left, green: anterior-posterior, blue: superior-inferior). (e) and (f) are the zoomed-in regions specified by the yellow rectangles in (a) and (c) respectively. The yellow arrows indicate the residual artefacts in the FA map of the PEN correction, which are effectively reduced in the FA map of the NPEN correction. A few residual artefacts are visible in the b=0 image of the NPEN result, which are indicated by the green arrows.

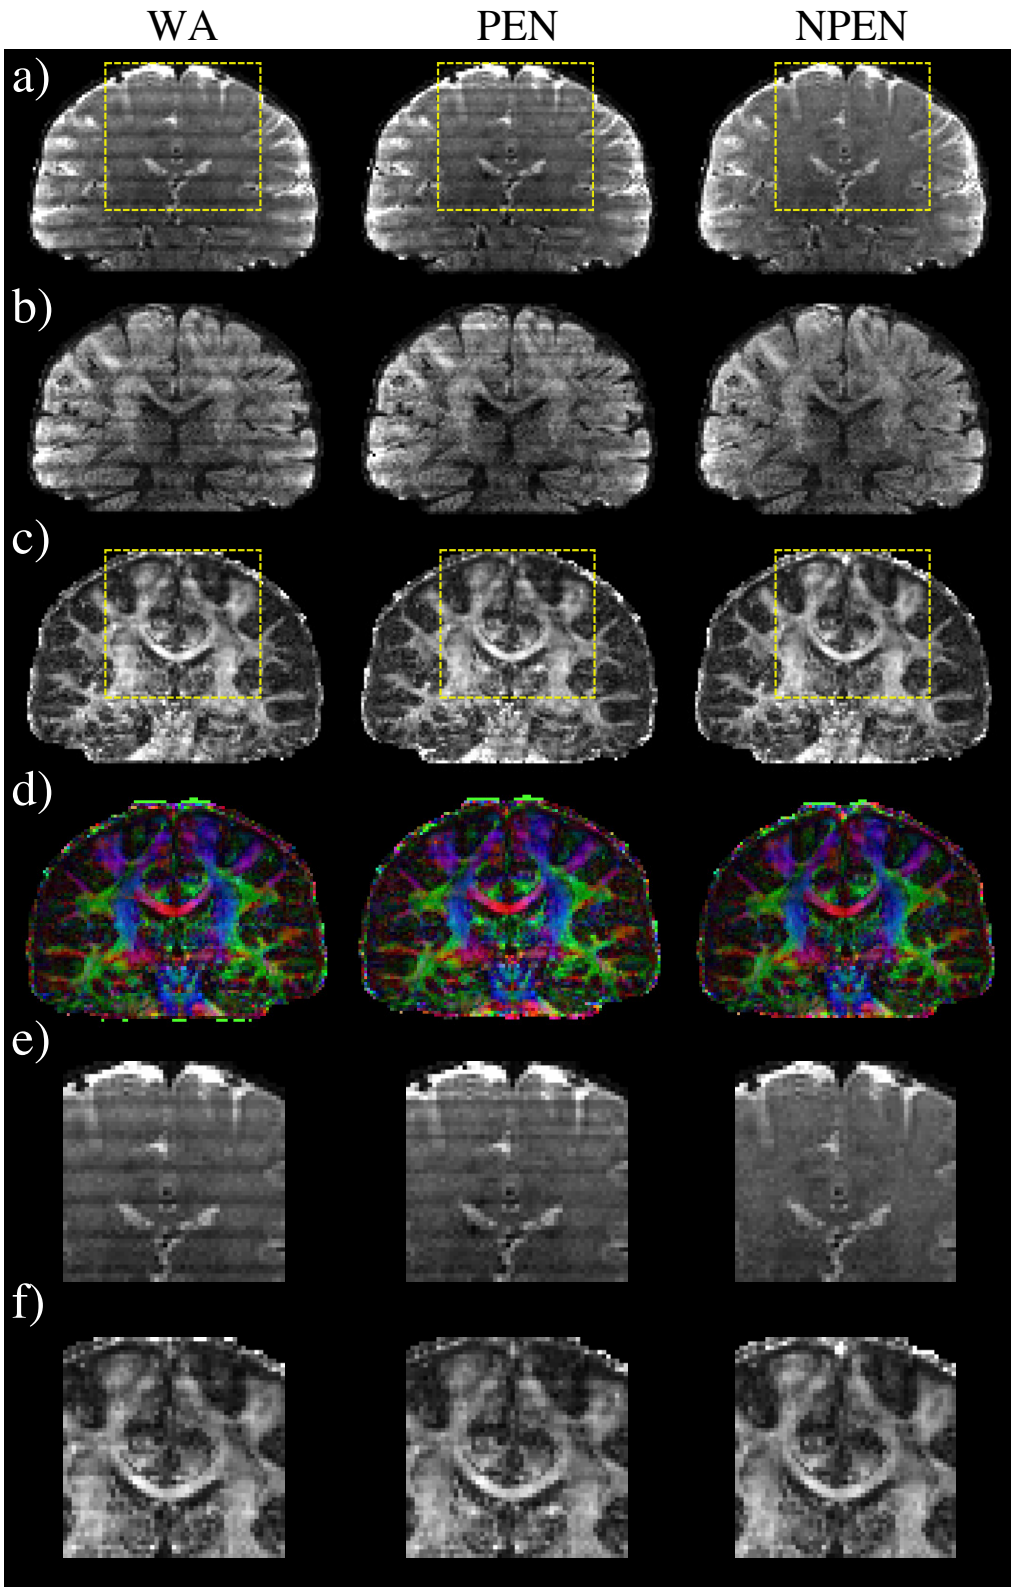


Figure 4.14: A coronal slice of the reconstruction results based on the data acquired from a different subject. (a) b = 0 image. (b) DWI image (b=1000 s/mm^2). (c) FA maps. (d) Color-coded maps of the principal eigenvector (red: right-left, green: anterior-posterior, blue: superior-inferior). (e) and (f) are the zoomed-in regions specified by the yellow rectangles in (a) and (c) respectively.

4.5 Discussion

Several methods have been proposed to correct slab-boundary artefacts in dMRI, including the weighted average and the PEN methods considered here. Weighted average avoids aliasing at slab boundaries by oversampling along the slab direction. The goal of PEN is to remove aliasing artefacts, with no attempt to address the saturation effects. Thus PEN works well under the condition that the sum-of-squares image used in slab profile estimation has a relatively flat signal profile. If not, residual signal variation remains in the reconstructed images and can propagate into diffusion parameter maps such as MD and FA. Our NPEN method extends PEN to deal with both aliasing and slab crosstalk, which enables correction of slab-boundary artefacts under conditions of strong T_1 saturation due to short TR (1-2 s) associated with optimal SNR efficiency.

The conventional coupling between slice coverage and TR is lessened in multi-slab acquisition. Encoding of slices through a combination of slab excitation and within-slab phase encoding confers flexibility in achieving a desired coverage and TR without major impact on scan time. With short TR, we can only accommodate a small number of excitation periods. At longer TR, we can achieve the same coverage by increasing the number of slabs while reducing the within-slab phase encoding, avoiding significant increase in total scan time. Moreover, one could also deploy simultaneous multi-slab acceleration [121] to increase coverage without requiring more excitations per TR. As such, the need for longer TR to improve reconstruction in PEN and WA does not in itself require longer scan time; and conversely, the key benefit conferred by NPEN in enabling short TR is increased SNR efficiency – not shorter scan times.

Nevertheless, one drawback to all slab-boundary corrections is a reduction in scan-time efficiency. At its most efficient, the number of k_z phase encodings would simply be equal to the number of reconstructed z slices (although this would require

a perfect rectangular slab profile, which cannot be achieved in practice). We can define scan-time efficiency in terms of the fractional increase in scan time beyond this idealized case. All slab-boundary corrections have reduced scan-time efficiency due to additional measurements associated with either: (a) deliberate overlapping of adjacent slab FOVs in the transition bands; or (b) oversampling of excited-slab FOVs into their transition regions. WA uses both overlapping and oversampling, with the original implementation requiring 0.4× more kz encodings than reconstructed slices, or 71.4% scan time efficiency (21). Both PEN and NPEN overlap slab FOVs by ∼20%, and also incur a fixed overhead to acquire calibration data (here, equivalent to acquiring one additional volume). Hence, our protocol with two b=0 and 20 diffusion-weighted volumes had 79.7% scan-time efficiency.

Two constraints were applied in NPEN. The in-plane smoothness constraint imposed on the slab profile is used to improve the conditioning of the problem. This constraint seems reasonable given that Bloch simulations of the slab profiles for WM, GM and CSF are quite similar and $\mathbf{B}_0/\mathbf{B}_1$ profiles vary smoothly over space, although WM and CSF profiles do exhibit consistent differences (Fig. 4.4e). Nevertheless, the tissue-specific comparison (Fig. 4.10) suggests the artefacts in WM and GM components are nicely corrected, with greatest residual artefacts being in CSF, which is only of interest in areas of large partial volume. The second constraint was used to suppress the spurious frequency components corresponding to slab-boundary artefacts. A similar approach was adopted in [161], where a band-pass filter is used in the estimation of a 1D banding function, which is applied to correct slab-boundary artefacts. The main difference between the two approaches is that NPEN allows the spatial variation of the slab profile within each slice plane, capturing the effects caused by different tissue types and $\mathbf{B}_0$ inhomogeneity, while the 1D banding function used in the previous method can only provide a global correction. The effects of these constraints on NPEN reconstruction are demonstrated in Fig. 4.15. The reconstruction RMSE increases after a certain number of iterations. This

most likely stems from overfitting, which may occur due to the high complexity of the nonlinear reconstruction model. In the reconstruction, the algorithm tries to improve the fitting of the unknowns based on the reconstruction model. However, after a certain point, the improvement on the fitting of the unknowns comes at the expense of increased errors, leading to inflated RMSE. As the RMSE curves for WM and GM are different, in some cases it is possible the error in GM decreases while error in WM increases, but as shown Fig. 4.8 this is not a significant effect, which is not clearly visible in our results.

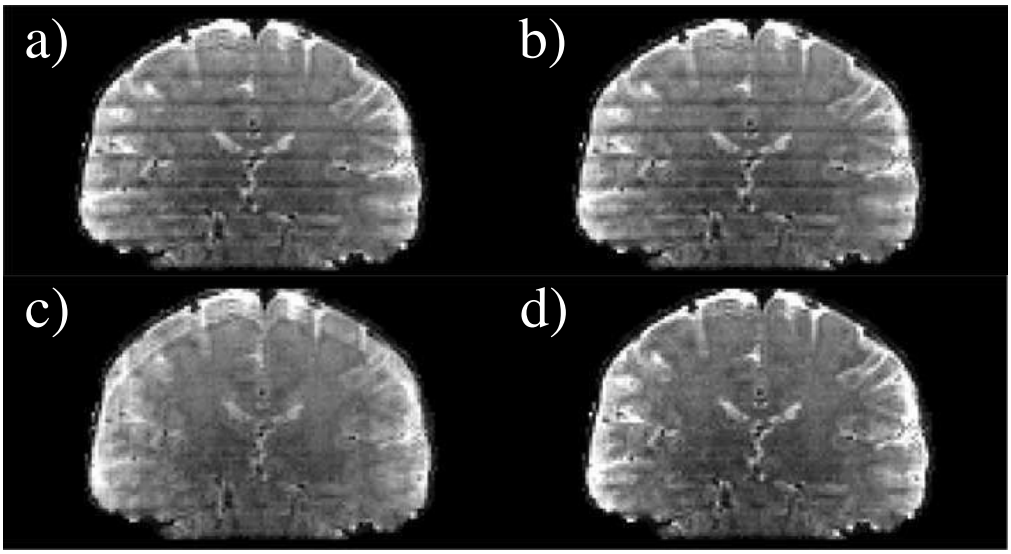


Figure 4.15: Reconstruction results using iteratively regularised Gauss-Newton method: (a) without additional constraints, (b) with in-plane smoothness constraints on slab profile, (c) with frequency constraints on image, (d) with both in-plane smoothness constraints on slab profile and frequency constraints on image. All other reconstruction parameters (e.g. number of iterations, regularisation parameters) are the same in the reconstructions.

In the PEN method, subject motion could be an outstanding problem as the slab profile estimated from the calibration scan might be misaligned with the images acquired later, especially for DTI where a large number of volumes are required and scan time is long. In comparison, NPEN is a data-driven method, in which the calibration image is only used for initialization and its effect is reduced during the iterative reconstruction, so NPEN is expected to be less sensitive to subject motion.

The long scan times (1 min per volume) used in the data acquisition are dictated by the use of readout segmentation and are the tradeoff associated with reduced distortion and T_2^* blurring. It should be noted that acquisition with single-shot per kz plane (with equivalent distortion and T_2^* blurring to 2D single-shot) can achieve ∼20s per volume, which is about 70% longer than a conventional 2D single-shot EPI due to the need for navigator echoes and slab overlap, but would still have about 50% higher SNR efficiency.

A drawback of NPEN is the computational cost. This is in part due to the operations on large matrices and the iterative nature of the method. The computation time for all diffusion directions was ∼1h using MATLAB (2014a, The MathWorks, Natick, MA) on a distributed computing cluster. The cluster consists of 19 computers, 8 of them have AMD Opteron Processors 6328 (32 cores, 3.2GHz) and 512GB RAM, and the others have Intel Xeon Processors E5640 (16 cores, 2.67GHz) and 64GB RAM. It should be noted that the computation time could be definitely reduced with optimized C/C++ implementation.

The NPEN method is investigated here in the context of diffusion measurements. However, this method may be applicable to other MRI applications using 3D multi-slab acquisition, such as 3D fast spin echo imaging and time-of-flight angiography.

4.6 Conclusion

We present a new approach to reduce the slab-boundary artefacts in 3D multi-slab dMRI. Slab-boundary artefacts were evaluated using Bloch simulations for a range of tissue properties and acquisition parameters, with particular interest on the potential for these artefacts to bias estimates of diffusion-based quantification (e.g. FA and MD). The presented reconstruction method, NPEN, can correct both the aliasing artefacts and signal variation caused by slab crosstalk even at short TR (∼2s). Reconstruction results on simulated data and data acquired in

vivo demonstrate that the proposed method improves upon previous methods, including the original PEN technique.

5

High-resolution Diffusion MRI at 7T using a 3D Multi-slab Acquisition

Contents

5.1 Introduction 102
5.2 Theory 105
5.3 Methods 106
5.3.1 Sequence 106
5.3.2 Data Acquisition 108
5.3.3 Reconstruction 110
5.3.4 Post-processing 113
5.4 Results 113
5.4.1 Comparison of Different ACS Acquisitions 113
5.4.2 Slab-Boundary Artefacts Correction 118
5.4.3 Diffusion Tensor Results 118
5.4.4 Tractography Results 121
5.5 Discussion 123
5.6 Conclusion 129

High-resolution diffusion MRI can provide the ability to resolve small brain structures, enabling investigations of detailed white matter architecture. A major challenge for high-resolution diffusion MRI in vivo is the low signal-to-noise ratio. In the work described in this Chapter, we combine two highly compatible methods,

ultra-high field and three-dimensional multi-slab acquisition to improve the SNR of high-resolution diffusion MRI. As each kz plane is encoded using a single-shot echo planar readout, scan speeds of the proposed technique are similar to the commonly used two-dimensional diffusion MRI. In-plane parallel acceleration is applied to reduce image distortions. To reduce the sensitivity of auto-calibration signal data to subject motion and respiration, several new adaptions of the fast low angle excitation echo-planar technique (FLEET) that are suitable for 3D multi-slab echo planar imaging are proposed and evaluated. A modified reconstruction scheme is proposed for auto-calibration with the most robust method, Slice-FLEET acquisition, to make it compatible with navigator correction of motion-induced phase errors. Slab-boundary artefacts are corrected using the nonlinear slab profile encoding method described in Chapter 3. Results in vivo demonstrate that high-quality diffusion MRI data with ∼1mm isotropic spatial resolution can be achieved with this combination of ultra-high field and 3D multi-slab acquisition. Content of this Chapter is based on a peer-reviewed journal paper we published in NeuroImage [8].

5.1 Introduction

Diffusion MRI has been widely used in the investigation of neuroanatomical connections. Although the possibility of high-resolution diffusion MRI in vivo is compelling, the feasibility of this technique in practice is still limited. The primary challenges include motion sensitivity and the reduced SNR at high spatial resolution, in which the latter is particularly problematic for diffusion MRI due to its low intrinsic SNR, which often means imaging near the noise floor. To improve the SNR of high-resolution diffusion MRI without substantially increasing the scan time, the two dominant approaches are the use of ultra-high field (7T and above) scanners [5, 99, 101, 196] and acquisition schemes with higher SNR per unit time

[2, 3, 59, 158, 159, 173] compared to the current standard, two dimensional (2D) single-shot echo planar imaging (EPI) [19, 33].

Ultra-high field MRI can provide stronger signals than more standard field strengths (e.g. 3T), which could be used to preserve the SNR at higher spatial resolution. The downside of ultra-high field however is that T_2 and T_2^* become shorter, which to some extent counteracts the SNR benefit of high field because of the long echo time (TE) associated with spin-echo diffusion preparation [3, 197]. Short T_2^* and strong $\mathbf{B}_0$ inhomogeneity are the causes of two major EPI artefacts, image blurring and distortion, which can dramatically degrade the image quality, particularly at high field. Hence, parallel acceleration [30, 31, 198] becomes crucial for single-shot EPI at 7T to reduce the echo train length (less T_2^* blurring) and decrease the effective echo spacing (less distortion) [5]. Segmented EPI could serve as an alternative solution for reducing EPI artefacts [53, 68], which has been primarily used in clinical dMRI. However, the long scan time and practical challenges associated with segmented EPI acquisition due to motion sensitivity and phase errors might hinder its application in many diffusion MRI studies. Another downside of 7T is the increased specific absorption rate (SAR), which is a great challenge for spin-echo diffusion MRI, especially when multi-band acceleration is applied. Even without multi-band acceleration, the increased SAR could still limit the RF pulses (especially the refocusing pulse) that can be used at 7T.

As discussed in Chapter 4, SNR can also be improved by addressing fundamental inefficiencies in the acquisition. The TR of 3D multi-slab acquisition is proportional to the number of slabs, which can be adjusted to favour a TR associated with optimal SNR efficiency, making this a very attractive method for maximising SNR [59, 158].

The higher signal at 7T and optimal SNR efficiency from 3D multi-slab acquisition are complementary, such that the combination of these two methods should be able to provide high-quality diffusion MRI images. However, there

are still two problems to be solved, which could otherwise lead to substantial deterioration of the image quality.

First, parallel imaging methods that require auto-calibration signal (ACS) measurements [31] are greatly affected by the quality of ACS data, particularly for EPI-based acquisitions [199]. Single-shot ACS data have different geometric distortions than the imaging data they are applied to, leading to residual artefacts; conversely, segmented ACS acquisitions match distortions but are sensitive to subject motion and respiration, again leading to artefacts. Recently, it has been shown that a re-ordering of the segmented ACS acquisition scheme [200][199] can provide more robust ACS data in the presence of motion and respiration as well as proving more robust to strong $\mathbf{B}_0$ offset [199]. In this work, we extend these alternative ACS acquisitions to 3D multi-slab acquisition to improve the robustness and performance of the GRAPPA reconstruction.

Second, 3D multi-slab exhibits slab boundary artefacts because the edges of the slab profiles cannot be perfectly sharp and abutting, resulting in periodic signal modulations due to saturation at the boundaries and slab aliasing [7, 160]. In Chapter 4, we proposed the NPEN method which corrects both slab aliasing and signal modulations, while being compatible with optimal TR [7]. The performance of NPEN correction has already been demonstrated at 3T; here, NPEN is applied at 7T to reduce the slab boundary artefacts for 3D multi-slab diffusion MRI, with modification to address the stronger $\mathbf{B}_0$ inhomogeneity at 7T.

In the work presented in this Chapter, we leverage the compatibility of 7T and 3D multi-slab acquisition to achieve high-resolution, high-SNR diffusion MRI. We compare different ACS acquisition strategies for 3D multi-slab imaging. Slab boundary artefacts are corrected using the NPEN method, which has required further improvements to handle larger $\mathbf{B}_0$ offsets at 7T. With all these advances, we demonstrate in vivo diffusion MRI data at 7T with high resolution and high

quality. Diffusion analysis and tractography on these data are presented alongside published post-mortem results, demonstrating consistent findings.

5.2 Theory

The fast low angle excitation echo-planar technique (FLEET) is a hybrid of FLASH and segmented-EPI [199, 201]. Unlike conventional segmented EPI, which uses an inner slice loop and cycles through segments in the outer loop to favour recovery of longitudinal magnetisation, FLEET swaps the order of slice and segment loops, such that all segments for a given slice are fully acquired before moving to the next slice. This reduces the temporal separation between segments of a single slice, minimizing the effects of subject motion and respiration. These favourable properties of FLEET-type segmented acquisition make it highly suitable for use as an ACS module in 2D EPI sequences [199] and this also extends to 3D single-slab EPI [202], where a more robust 3D reference can be obtained with all segments for a given kz partition acquired in immediate succession.

In 3D multi-slab acquisition, only in-plane GRAPPA is considered, as the slab thickness is too thin, which is not compatible with acceleration along kz. To combine FLEET-ACS and 3D multi-slab acquisition, two possible schemes are immediately apparent (Fig. 5.1a): 'Slab-FLEET', which excites slabs aligned to the centre of the subsequent imaging slab data, and 'Slice-FLEET', which excites 2D slices with locations matching to the final imaging volumes. Slice-FLEET obtains unique reference data for each slices, whereas Slab-FLEET acquires one set of reference for all slices in a given slab, resulting in a shorter calibration scan. The two schemes are expected to differ in three key aspects: (i) intra-voxel dephasing, (ii) coil profile mismatch, and (iii) acquisition time, which are discussed below:

(i) FLEET is a gradient-echo based sequence, which suffers from signal dephasing inside a voxel, exacerbated by the strong $\mathbf{B}_0$ inhomogeneity at 7T and thick

excitations. Because Slab-FLEET excites a thick slab (without kz phase encoding), it is worth considering slab excitations that are thinner than the imaging slab (which is kz phase encoded) (Fig. 5.1a). Slice-FLEET will have considerably less intra-voxel dephasing due to the thin slices.

(ii) Slab-FLEET essentially averages the coil sensitivity of the whole FLEET slab. This strategy assumes the coil sensitivity is relatively constant along the slab direction, but in reality the averaged coil profile is not optimal for GRAPPA reconstruction at each individual slice. If thinner Slab-FLEET excitation is utilised to reduce intra-voxel dephasing, slices at slab boundaries would have more inaccurate GRAPPA kernel calibration. In contrast, Slice-FLEET can provide matched coil profile for each slice, potentially enabling more accurate GRAPPA reconstruction.

(iii) Despite the potential disadvantages in coil profile mismatch and intra-voxel dephasing, Slab-FLEET requires only one set of ACS data per slab. By contrast, Slice-FLEET acquires one set of ACS for each slice, requiring longer acquisition time.

5.3 Methods

5.3.1 Sequence

A diffusion weighted readout-segmented EPI sequence [53] was modified to acquire 3D multi-slab data with the Stejskal-Tanner preparation scheme [27] and a single-shot EPI readout. Further modification was implemented to allow oversampling along kz by a factor of two, which could be used to produce an initial guess for NPEN correction [7]. The sequence used the 90° excitation and 180° refocusing pulses provided by the vendor with a time-bandwidth product of 24 and 5 and durations of 7.68ms and 10.24ms, respectively. After acquisition of each k-space ("imaging") segment, a second refocusing pulse was applied to acquire a 2D navigator at lower resolution to correct the phase errors induced by subject bulk motion during diffusion encoding [155, 158]. The imaging data were phase-encoded along ky and

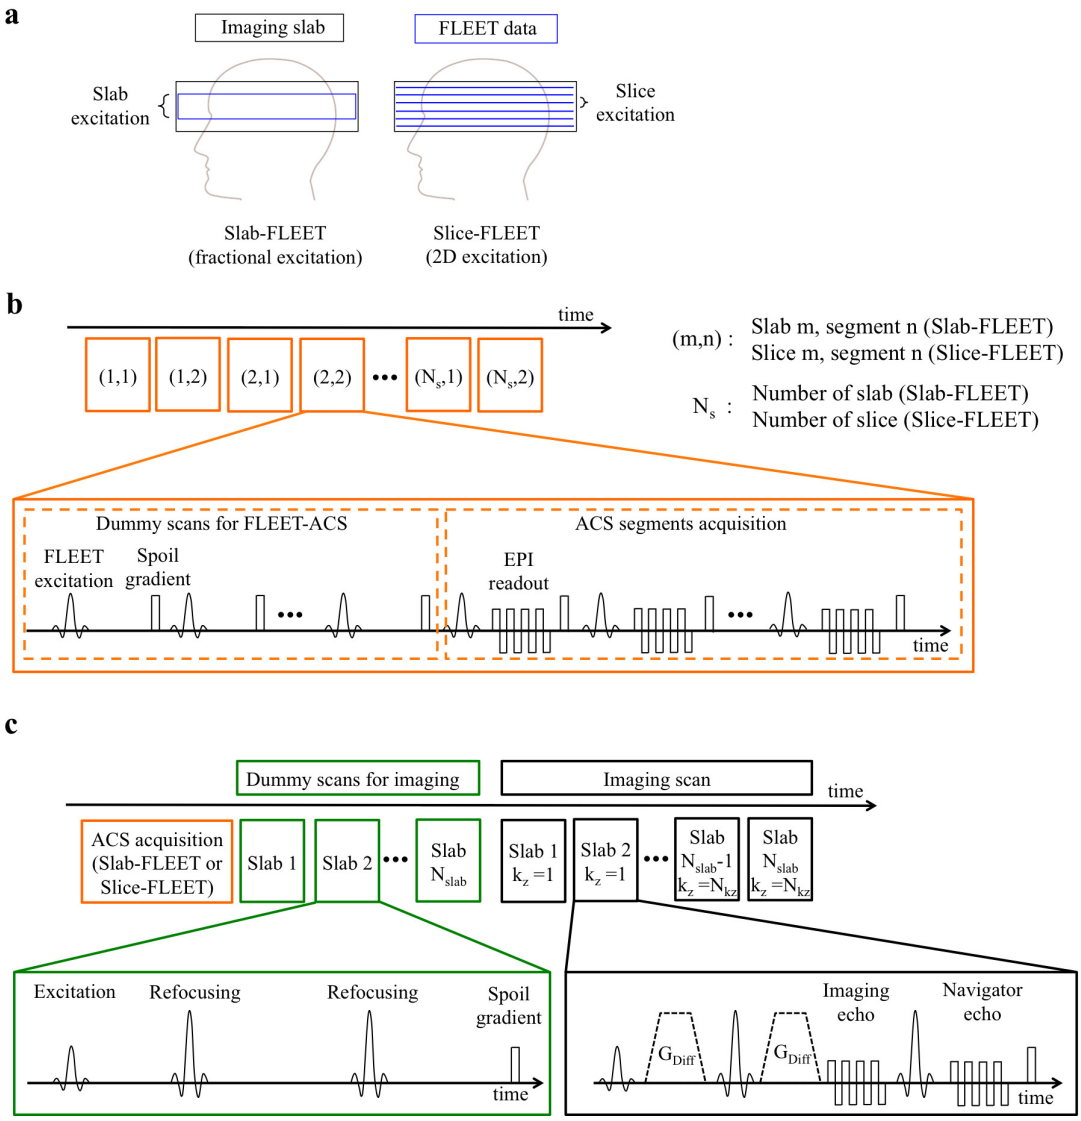


Figure 5.1: Illustration of 3D multi-slab acquisition with FLEET-ACS. (a) Two possible schemes for implementing FLEET-ACS in 3D multi-slab acquisition. (b) Schematics for FLEET-ACS acquisition, GRAPPA factor of 2 is assumed in this example. Both Slab-FLEET and Slice-FLEET utilize consecutive-segment acquisition order, such that the interval between the acquisition of the first and the last segments for a given slice/slab can be minimised. (c) Schematics for imaging data acquisition, the orange block represents the full FLEET-ACS acquisition shown in (b). N_{slab} is the number of slabs. N_{kz} is the number of kz phase encodes for each slab.

kz to generate a 3D k-space, whereas the navigator data was only phase encoded along ky, forming a 2D k-space, in which the signal of the excited slab is collapsed across the slab [59]. Use of a 2D navigator has been demonstrated to provide satisfactory correction of phase errors in thin (up to 30 mm) 3D slabs [92, 158]. As the navigator is not kz phase-encoded, Slab-FLEET ACS is a more suitable option for GRAPPA reconstruction than Slice-FLEET ACS.

Each kz plane was sampled with single-shot EPI (SSH-EPI), enabling a similar sampling efficiency to conventional 2D SSH-EPI sequences. Parallel imaging acceleration using GRAPPA [31] was applied along the ky phase-encoding direction. Partial Fourier was used to achieve a short TE. The navigator data were also acquired with GRAPPA accelerated SSH-EPI, but the ky dimension of the acquisition matrix was kept to 64 lines to reduce the scan time. No partial Fourier was used in the navigator acquisition.

In ACS acquisition, all segments for a given slab (Slab-FLEET) or slice (Slice-FLEET) were acquired consecutively (Fig. 5.1b). A constant flip angle was used for excitation. Before ACS acquisition, a series of dummy scans were applied to achieve a steady-state magnetization, such that all segments would have equal signal level. RF spoiling and gradient spoiling were applied as described previously [199]. The TE of the ACS acquisition was minimised to reduce the acquisition time. Fat saturation was applied before RF excitations to suppress the fat signal, which was not shown in the figure. After ACS acquisition, a series of imaging dummy scans were applied, followed by the 3D multi-slab image acquisition (Fig. 5.1c).

5.3.2 Data Acquisition

Six healthy subjects were scanned on a Siemens 7T scanner using a 32-channel RX Nova coil and a SC72D gradient set with maximum gradient amplitude of 70 mT/m on each axis and maximum slew rate of 200 mT/m/ms. Written informed consent in accordance with local ethics was obtained from each subject before experiments.

Three subjects were scanned using an ACS evaluation protocol to evaluate the performance of different ACS acquisition methods. The other three subjects were scanned using a high-resolution diffusion protocol.

ACS Evaluation Protocol

For each subject, 14 slabs with 10 slices per slab were acquired and neighbouring slabs were overlapped by 20% (2 slices), resulting in 114 slices in the final reconstruction. The slab acquisition was interleaved while kz encoding inside each slab was sequential. Each kz plane was acquired using SSH-EPI, such that the number of segments is equal to the number of kz encodings in each slab. 51 echoes were collected with an echo spacing of 0.82 ms and a bandwidth of 1442 Hz/pixel. GRAPPA acceleration factor 3 was applied along ky phase encoding, resulting in an effective echo spacing of 0.27 ms. Partial Fourier acquisition of 6/8 was used, resulting in a final in-plane matrix of 204x204. The TEs for the EPI readouts were 71 ms (imaging) and 141 ms (navigator); TR was 2600 ms. Isotropic resolution of $1.03\times1.03\times1.03$ mm^3 was acquired with a FOV of $210\times210\times117$ mm^3. Diffusion-weighted data were acquired using three b-values 0, 1000, 2000 s/mm^2 with diffusion encoding along the left-right direction. The two high b values were used to evaluate the performance of the FLEET-based methods at difference noise levels. As the GRAPPA reconstruction is relatively independent of the diffusion encoding direction, the evaluation results of ACS acquisition should be similar for all diffusion directions. ACS data were acquired using spin-echo SSH-EPI, conventional spin-echo segmented EPI, Slab-FLEET and Slice-FLEET respectively, where spin-echo SSH-EPI and conventional spin-echo segmented EPI only acquire one set of ACS data for all slices in a given slab, as with Slab-FLEET. To evaluate the effects of slab/slice thickness on FLEET-ACS acquisition, four different slab thicknesses (100%, 50%, 20% and 10%) were measured for the Slab-FLEET, and two slice thicknesses (1× and 2× with respect to the final

imaging slice) were measured for the Slice-FLEET, resulting in 8 sets of ACS data. The total acquisition time for the ACS data and the image data was about 6 minutes.

High-resolution Diffusion Protocol

Diffusion-weighted images were acquired with b = 1500 s/mm^2 and 64 diffusion directions uniformly sampled on a sphere, and eight b = 0 images (six of these acquired with anterior-posterior phase encoding to match the high-b scan, and two acquired with opposite phase encoding for distortion correction). Other sequence parameters were the same as the ACS evaluation protocol except a longer TR of 2700 ms was used due to SAR limitation. One set of b = 0 data was oversampled by a factor of 2 along the kz direction for the initial estimation of slab profile required by NPEN. Two sets of ACS data were acquired using Slab-FLEET with 50% slab thickness and Slice-FLEET with 1× slice thickness, which were used to reconstruct navigator signal and imaging signal respectively. The total scan time for all diffusion-weighted data was ∼35 minutes.

5.3.3 Reconstruction

GRAPPA Reconstruction

GRAPPA reconstructions using ACS data acquired with spin-echo SSH-EPI, conventional spin-echo segmented-EPI and Slab-FLEET followed the procedures described previously [158]. In short, for each slab, a 2D GRAPPA kernel was calibrated and applied to synthesize the missing k-space lines for all kz encodes and the navigator. After the GRAPPA reconstruction, motion-induced phase errors for each kz segment were corrected using the phase of the navigator images. Finally, a 3D Fourier transform produced the image.

The reconstruction pipeline needs to be altered when using ACS data acquired with Slice-FLEET, which uses a GRAPPA kernel calibrated for each individual slice. In order to reconstruct the imaging data, a Fourier transform along kz

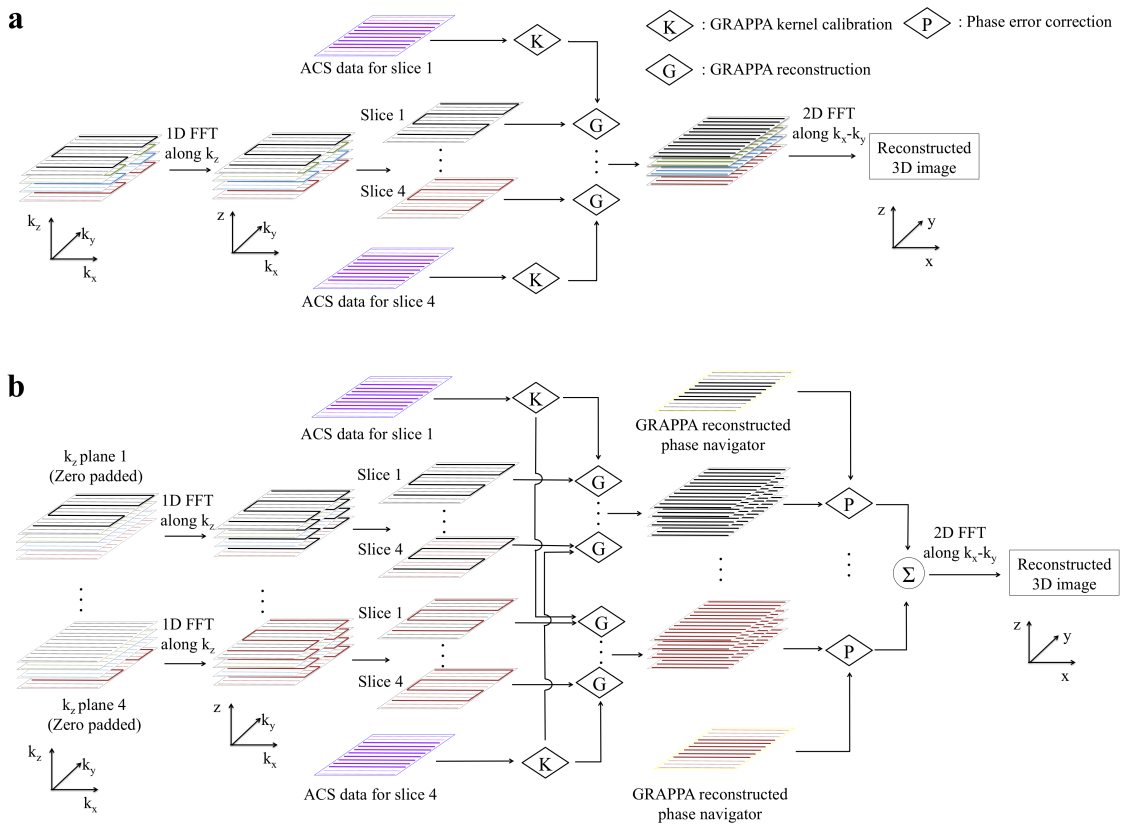


Figure 5.2: GRAPPA reconstruction and phase error correction for 3D multi-slab MRI with Slice-FLEET ACS: (a) b = 0 data, (b) Diffusion-weighted data.

should be applied to transform from the kx-ky-kz space to the kx-ky-z space before the slice-by-slice in-plane GRAPPA reconstructions (Fig. 5.2a). For b = 0 data this works fine, but for diffusion-weighted images this initial kz Fourier transform becomes problematic due to phase-inconsistencies between segments. These phase errors need to be corrected using the navigator echoes, but the navigator correction itself needs full-FOV images (i.e. images that already have been GRAPPA reconstructed). GRAPPA reconstruction therefore is a circular problem when applied to diffusion-weighted images in combination with Slice-FLEET ACS data. To solve this problem, a new scheme (Fig. 5.2b) was proposed by exploiting the linearity of each reconstruction stage, which meant that the contribution of each 3D k-space segment to each z plane (slice) could be calculated independently. Specifically, every individual kz segment was first placed in an otherwise zero-filled 3D k-space, and then

Fourier transformed along kz. This avoids the different levels of phase corruption between segments to interfere with one another, in other words: it is safe to perform a kz Fourier transform despite the phase errors not having been corrected yet. We then applied slice-by-slice GRAPPA reconstruction, followed by phase navigator correction. Finally, the N_{kz} k-spaces were summed and Fourier transformed to generate a 3D slab image (i.e. invoking the linearity of individual stages).

For Slice-FLEET with 2x slice thickness, GRAPPA reconstruction for an image slice used kernels calibrated from the nearest FLEET slice.

NPEN Correction

The method for correcting slab-boundary artefacts using NPEN has been described in Chapter 4. Initial estimations are required for nonlinear inversion, which was previously based on the b=0 image [7]. A potential drawback of this method is that if the $\mathbf{B}_0$ field is very inhomogeneous, the slab profile can be locally distorted and shifted, and as a result, the initial 3D slab profile estimated from the averaged signal within the slice may not be suitable at all slice locations, which can lead to biased reconstructions. At 3T, this effect was not seen due to relatively homogeneous $\mathbf{B}_0$ and high time-bandwidth product of RF pulses. However, this does become problematic at 7T due to increased $\mathbf{B}_0$ inhomogeneity and the use of lower time-bandwidth product pulses due to SAR limitations.

To mitigate this problem, slab-profile initialisation was modified to include a rough estimation of in-plane profile variation caused by $\mathbf{B}_0$ inhomogeneity. First, a Bloch simulation of the spin-echo sequence used in this work was performed with a range of off-resonance frequencies to generate a lookup table for MRI signals with voxel location and off-resonance frequency as entries, which only needs to be done once if the sequence parameters are not changed. A field map was estimated from the eight b = 0 images with both normal and reversed phase using FSL's TOPUP tool [203]. With the field map and the lookup table, the distorted slab

profile could be produced within seconds. Then a 2D weighting map was generated for each slice by normalizing the simulated profile with the averaged magnitude within slice. After that, the initially estimated 3D slab profile was multiplied with the weighting maps to formulate the final slab profile estimation. It should be noted that, as with any field map based correction method, with substantial subject motion, the slab profile estimation would become inaccurate if the field map was not updated. To alleviate this problem, the multiple b = 0 image pairs (with normal and reversed phase) could be arranged at different stages of the scan to generate field maps at different time points.

5.3.4 Post-processing

Image distortions induced by susceptibility and eddy currents were estimated and corrected using TOPUP and EDDY [40]. We then fit the tensor model using the DTIFIT tool [195] and multiple fibre populations within a single voxel using BedpostX [204]. White matter tractography was performed in the diffusion space using FSL's probabilistic tracking tool [204, 205]. Two major association fibre tracts (cingulum bundle and superior longitudinal fasciculus, SLF) were investigated with seed masks drawn on a single coronal slice in the diffusion space (Fig. 5.3), no target or stopping masks were used. An exclusion mask at mid-brain was applied to avoid cross-hemispheric tracts. Step length of 0.2 mm and 5000 streamlines were used in the tractography with no curvature threshold. To evaluate the quality of tractography, we compared these results to high-resolution post-mortem data presented in previous work [206].

5.4 Results

5.4.1 Comparison of Different ACS Acquisitions

Fig. 5.4 compares GRAPPA reconstructions for b = 0 data (Fig. 5.4a) and diffusion-weighted data (Fig. 5.4b, b = 1000 s/mm^2) with ACS data acquired using different

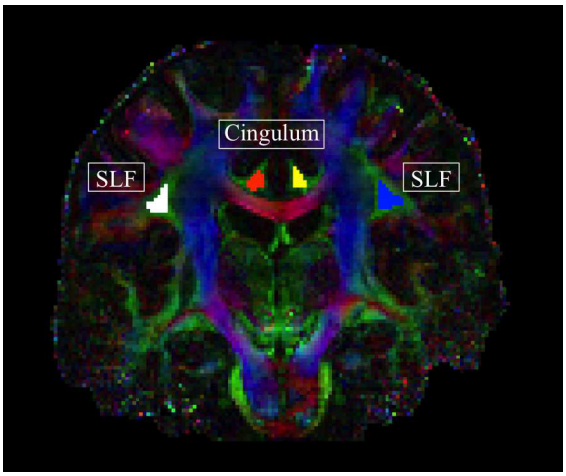


Figure 5.3: Seed masks used in the tractography for cingulum bundle and SLF.

methods. ACS data acquired with SSH-EPI suffers from strong residual aliasing and noise amplification, most likely due to the mismatched distortions between ACS data and imaging data. ACS data acquired using segmented EPI match distortion, but nevertheless display strong residual artefacts and amplified noise, which is likely due to subject motion and/or respiration during the long ACS acquisition time. All FLEET-ACS methods are observed to have reduced artefact. The choice of slab thickness in Slab-FLEET is not trivial: a thin slab thickness (e.g. 20%, 10%) allows good GRAPPA reconstruction at the slab centre, but fails to provide good results at slab boundaries, presumably because these regions are not considered during GRAPPA kernel calibration. This is avoided when a thick slab (e.g. 50%, 100%) is used, which however can lead to strong noise amplification, especially in the diffusion-weighted images (e.g. the noise level of the GRAPPA reconstruction using Slab-FLEET (50%) ACS is ∼20% higher than that using Slice-FLEET (1x), see Fig. 5.5b). This might be caused by the strong intra-voxel de-phasing effects within thick slabs and inaccurate estimation of local coil sensitivity. In contrast, Slice-FLEET performs well in both scenarios with 1x and 2x slice thicknesses showing good reconstructions.

A quantitative comparison of different ACS acquisition methods for three subjects

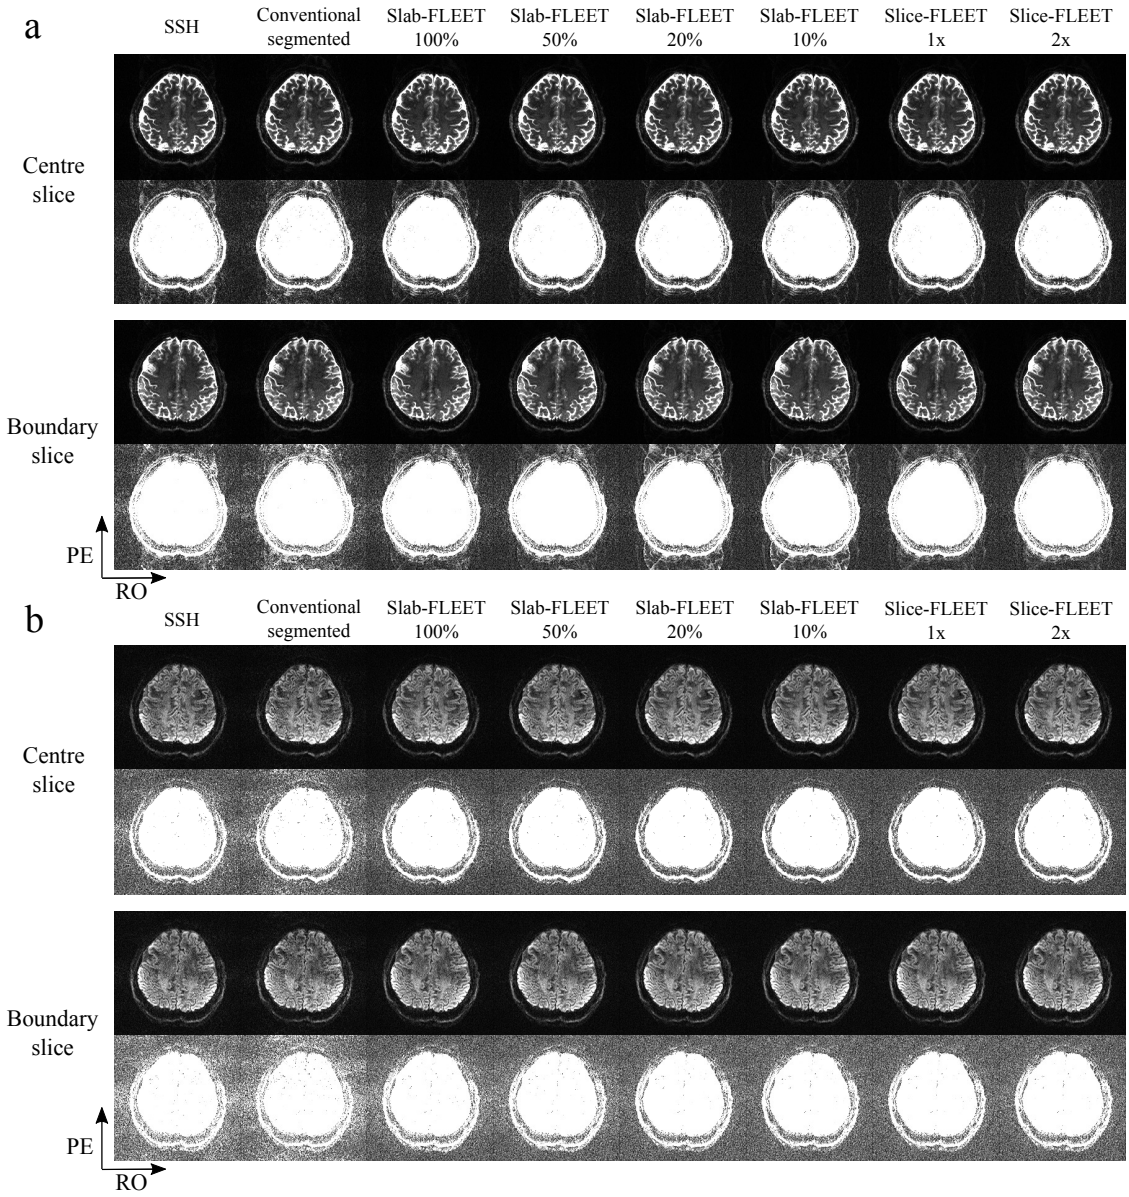


Figure 5.4: Comparison of GRAPPA reconstructions using different ACS acquisition methods. (a) b = 0 images. (b) Diffusion-weighted images acquired with b = 1000 s/mm^2. In each figure, the top two rows show a slice at slab centre and the bottom two rows show a slice at slab boundary. The phase encoding (PE) direction is along anterior-posterior.

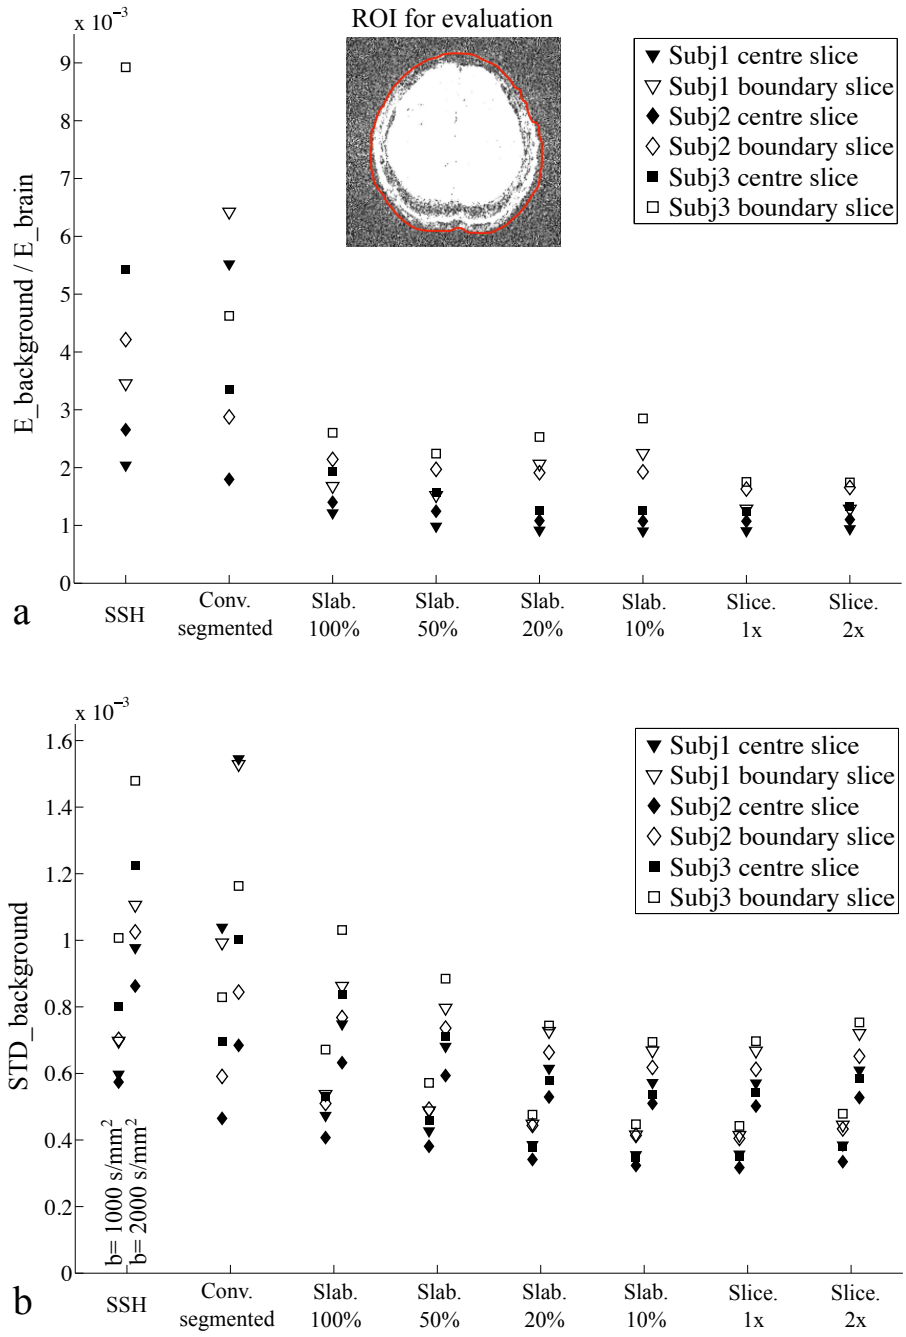


Figure 5.5: Quantitative comparison of different ACS acquisition methods for three subjects. (a) Ratio between signal energy in the background region (E_background) and signal energy in the brain region (E_brain) for b = 0 data, used as a measurement of residual aliasing. (b) Standard deviation of background signal for b = 1000 s/mm ^{2} and b = 2000 s/mm ^{2} data, used as a measurement of noise level. A region of interest (ROI) for the brain of one subject is shown in (a) and the background region is defined accordingly, which are used to evaluate ghosting artefacts and noise level.

is shown in Fig. 5.5. For the b = 0 data, the ratio between signal energy in the background region and signal energy in the brain region was calculated and used as a measurement of residual ghosting artefacts (Fig. 5.5a). For the diffusion-weighted data, the standard deviation of background signal was used as a measurement of noise level (Fig. 5.5b). It can be seen that FLEET-ACS methods enable more robust GRAPPA reconstruction for all subjects compared to SSH-EPI and conventional segmented EPI. Among all FLEET-ACS schemes, Slice-FLEET exhibits superior performance with low aliasing artefacts and low noise amplification for both centre slice and boundary slice.

Fig. 5.6 shows a more specific comparison between Slab-FLEET and Slice-FLEET. A slice at slab centre with a b-value of 2000 s/mm^2 is reconstructed using ACS data acquired with Slab-FLEET (50% slab thickness) and Slice-FLEET (1x slice thickness) separately. As shown in the results, Slice-FLEET ACS enables superior SNR in GRAPPA reconstruction and better preservation of detailed brain structures than Slab-FLEET method, for example, the diffusion-weighted contrast where white matter fibres enter a gyrus on the medial surface of the right hemisphere is clearer in the Slice-FLEET result (yellow rectangle).

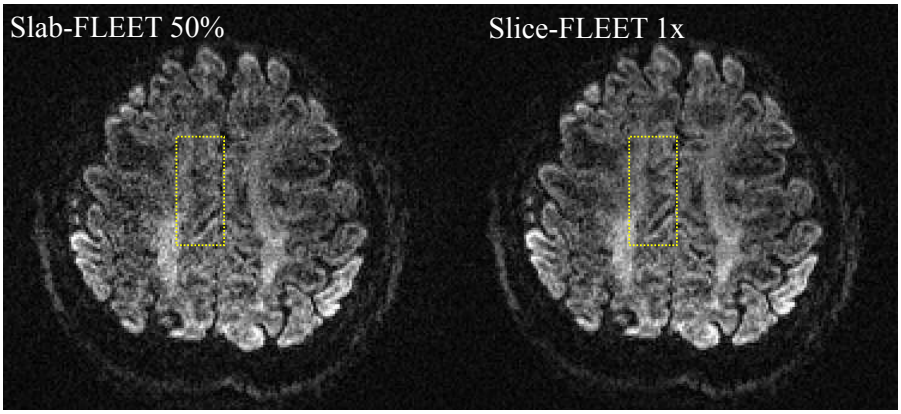


Figure 5.6: Comparison of GRAPPA reconstruction using ACS data acquired with Slab-FLEET (50% slab thickness) and Slice-FLEET (1× slice thickness). b = 2000s/mm^2. The Slice-FLEET method demonstrates superior SNR and better preservation of detailed brain structures than Slab-FLEET (yellow rectangle).

Based on the results of the comparison, the diffusion-weighted data acquired with the high-resolution diffusion protocol were reconstructed using Slice-FLEET ACS data.

5.4.2 Slab-Boundary Artefacts Correction

Fig. 5.7 shows the results of slab-boundary artefacts correction. Three methods are compared here: direct slab combination (the first and the last slices of each slab are discarded and the remaining slices from all slabs are concatenated), NPEN and improved NPEN. The direct slab combination retains strong artefacts at slab boundaries. NPEN effectively reduces the artefact level. However, NPEN correction still contains some residual artefacts in the region of strong $\mathbf{B}_0$ inhomogeneity (yellow arrows in the zoomed-in image). With improved NPEN, these residual artefacts are further reduced (zoomed-in image).

5.4.3 Diffusion Tensor Results

Fig. 5.8 shows high-resolution ($\sim$1 mm isotropic) b = 0 and diffusion-weighted images (b = 1500 s/mm^2) acquired at 7T with a scan time of 27s per volume. The images show high SNR, high contrast and good anatomical detail thanks to the high resolution. It should be noted that the low signal intensity in the temporal lobes (yellow circles) is due to the inhomogeneity of transmission field ($\mathbf{B}_1^+$) at 7T.

Fig. 5.9 shows colour-coded maps of the principle eigenvectors from three subjects. High data quality is evident, for example, in the cingulum bundle as it descends into the temporal lobe toward hippocampal projections (Fig. 5.9a, yellow arrows), a thin tract but here it is several voxels thick. Fig. 5.9b shows the white matter tracts near the posterior horn of the lateral ventricles: the superior longitudinal fasciculus, posterior thalamic radiations and tapetum are nicely distinguished in the high-resolution DTI data acquired in vivo.

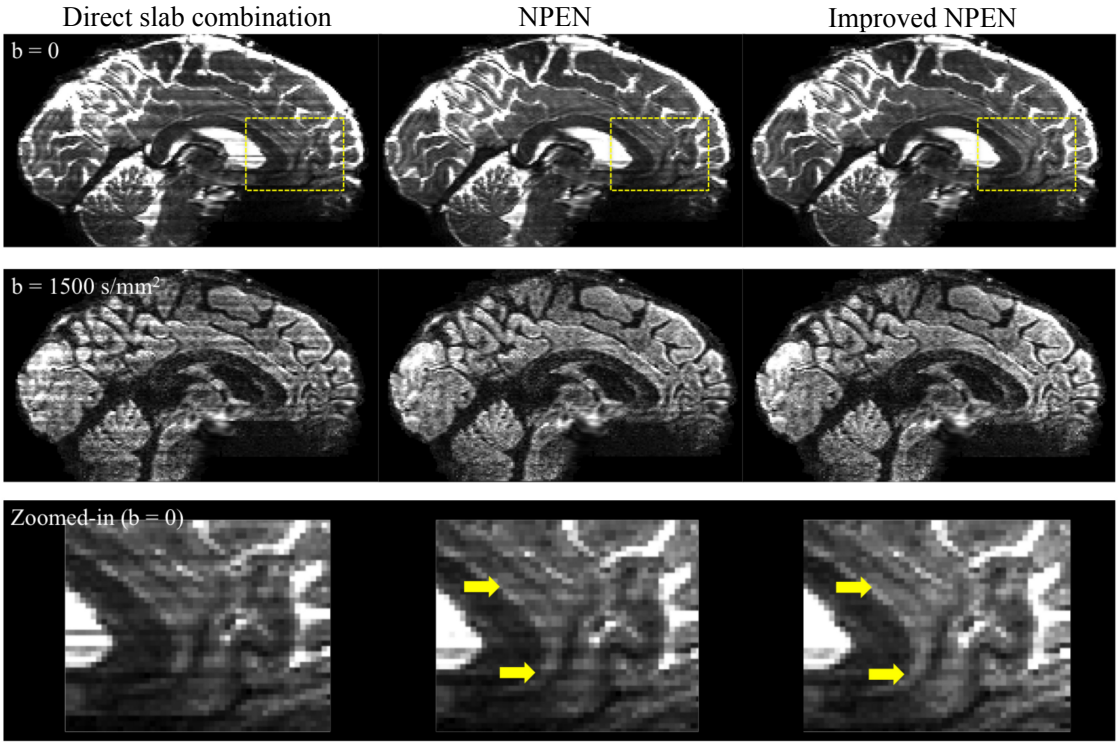


Figure 5.7: Correction of slab-boundary artefacts using direct slab combination, NPEN and improved NPEN. Slab-boundary artefacts are effectively reduced using NPEN correction, although there are still some visible residual artefacts in the areas with strong $\mathbf{B}_0$ inhomogeneity (yellow arrows in the zoomed-in version). Improved NPEN further reduces these artefacts. Note that due to the low intensity of the diffusion-weighted images in the inferior frontal region, the differences between the improved NPEN and previously proposed NPEN are less clearly visible.

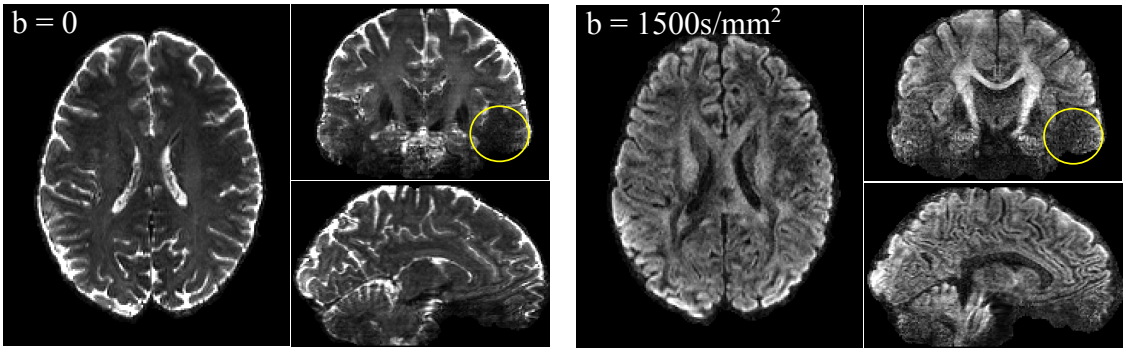


Figure 5.8: High-resolution data acquired at 7T using the high-resolution diffusion protocol. The yellow circles at temporal lobes indicate signal drop caused by the severe $\mathbf{B}_1^+$ inhomogeneity at 7T.

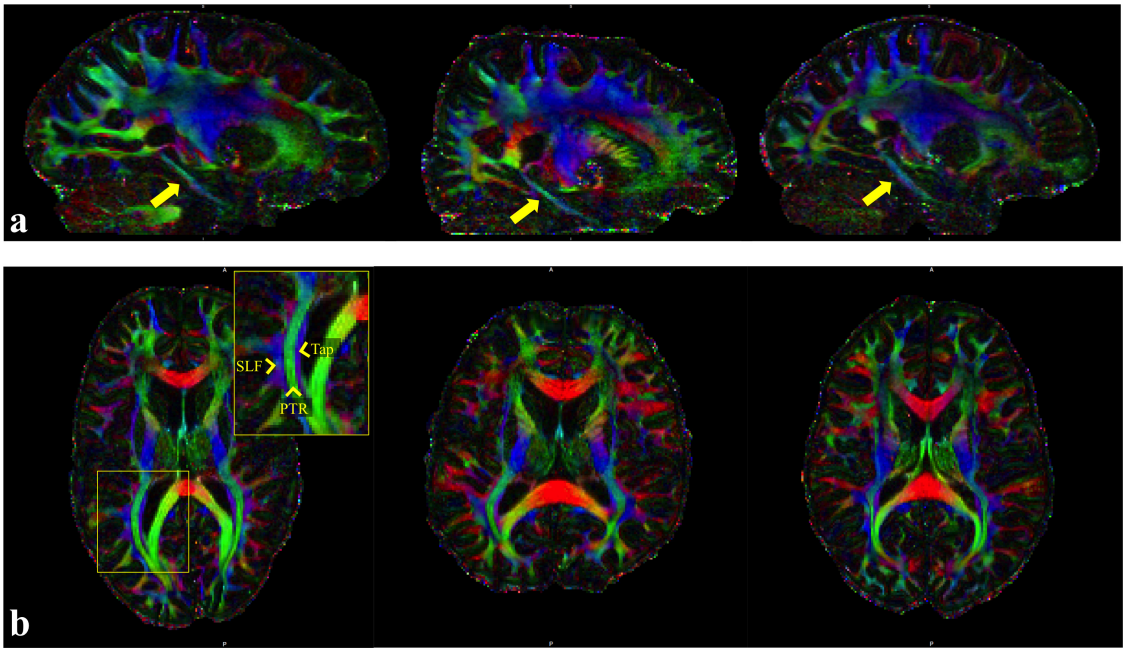


Figure 5.9: Colour-coded maps of the principle eigenvectors for the high-resolution DTI data (red, right-left; green, anterior-posterior; blue, superior-inferior). The high spatial resolution enables (a) clear detection of thin fibre tracts such as the cingulum bundle descending into the temporal lobe (yellow arrows) and (b) crisp disambiguation of some fibres which can be difficult to separate using low resolution diffusion MRI (superior longitudinal fasciculus, SLF, posterior thalamic radiation, PTR, the tapetum of the corpus collosum, Tap).

Fig. 5.10 shows the colour maps and fibre reconstruction of the pons of one subject, where the corticospinal tract (running superior-inferior, blue) interdigitate with the decussation of the pons (crossing right-left, red). The top row shows color maps from a single fibre reconstruction, in which the corticospinal tracts dominate (blue, indicated by the arrows). Nevertheless, the multi-fibre reconstructions on the bottom row are able to robustly estimate both fibre populations within a voxel, demonstrating high contrast and high SNR of the data.

Fig. 5.11 demonstrates high-resolution depiction of cortical anisotropy, focusing on the central sulcus. For all three subjects, the high-resolution data consistently supports visualisation of fibre architecture within gyri. Here, the precentral gyrus is depicted in an axial plane just above the superior-most arch of the cingulum bundle. In all subjects, fibres can be seen entering the gyrus, turning slightly

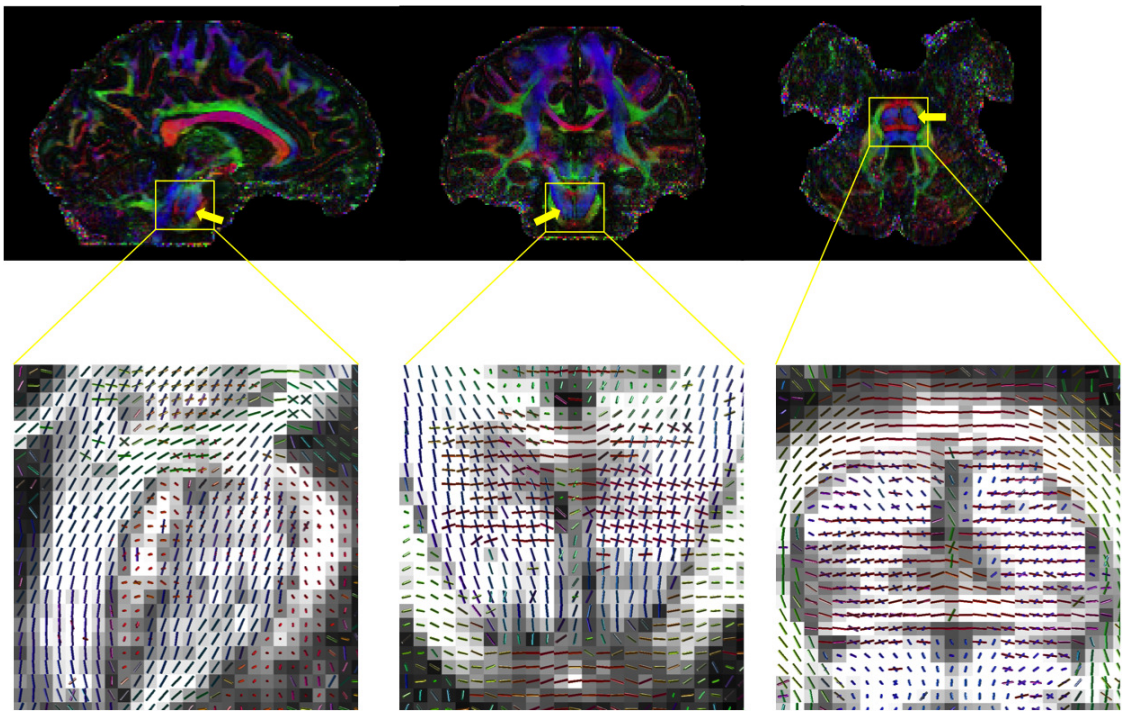


Figure 5.10: Complex fibre architecture in the pons is robustly detected in our high-resolution diffusion data. Top row: colour-coded maps of principle eigenvectors (red, right-left; green, anterior-posterior; blue, superior-inferior). Bottom row: Zoomed region of the pons showing multiple fibre populations estimated using the ball and sticks model, demonstrating the ability of the high-resolution data to resolve crossing fibres. In particular, tensor fits in the regions indicated by the arrows in the top row are dominated by tracts running superior-inferior, but additional right-left fibre populations are accurately captured with the crossing fibre models.

and entering cortex, with particularly clear delineation of radial orientation on the anterior bank of the central sulcus (i.e. primary motor cortex) and tangential orientation on the posterior bank of the central sulcus (i.e. primary somatosensory), consistent with previous reports [207].

5.4.4 Tractography Results

Tractography of the cingulum bundle and SLF are shown in Fig. 5.12 and Fig. 5.13, respectively. For comparison, tractography from post-mortem data in a previous study are also shown [206] (6 hours scan at 3T using a diffusion-weighted steady state free precession sequence with 0.94mm×0.94mm×0.94mm spatial resolution [208]) and used as gold standard. Even with these relatively simple tractography analyses

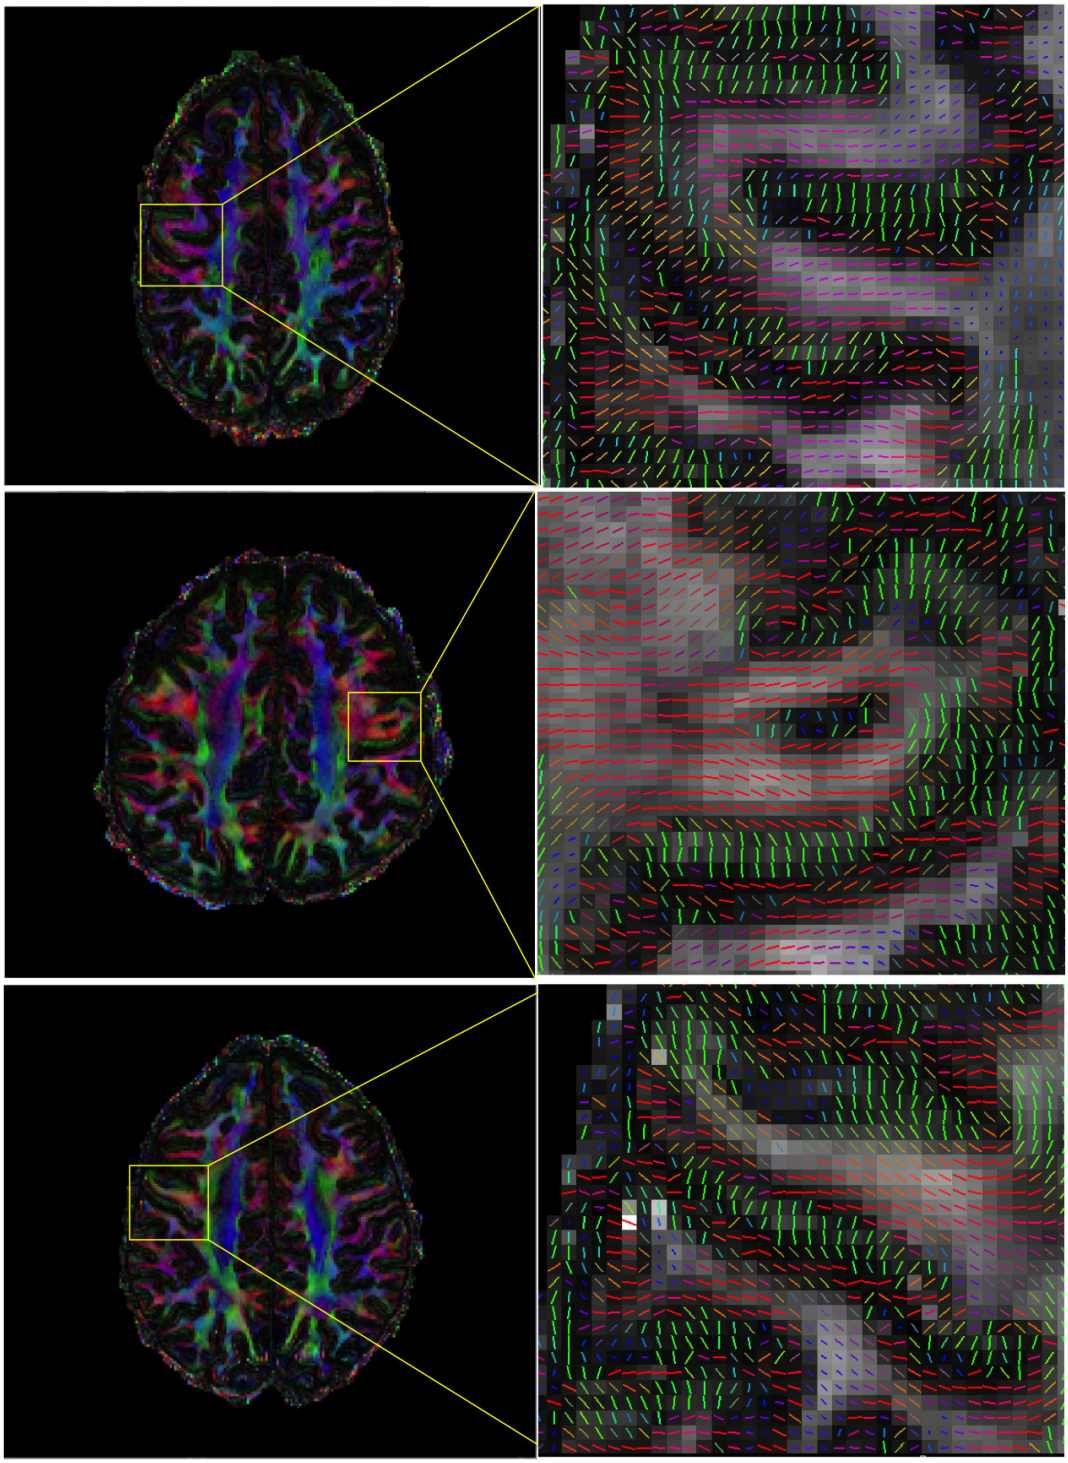


Figure 5.11: Depiction of cortical anisotropy using high-resolution DTI data (red, right-left; green, anterior-posterior; blue, superior-inferior).

(using a single seed mask, a mid-line exclusion mask and no waypoint masks), the in vivo tractography captures essentially the entire extent of the cingulum bundle, including temporal and frontal lobe tracts and cortical projections into cingulate cortex along its entire extent. The SLF tractography on the in vivo data also demonstrates consistency with the post-mortem data, including tracks into the pars triangularis and pars opercularis of the inferior frontal gyrus, and tracks into the supramarginal gyrus and angular gyrus in the parietal lobe.

5.5 Discussion

We have demonstrated the acquisition of excellent quality, high-resolution (1 mm isotropic) diffusion imaging through the combination of 7T and 3D multi-slab imaging, the latter enabling TRs in the optimal range with respect to SNR efficiency. While the multi-slab approach considerably reduces the coupling between coverage and TR, the shortest possible TR for 3D multi-slab acquisition is to an extent limited by the FOV along the slab direction and the slab thickness. Thin slabs are preferred as they allow better phase-error correction with 2D navigators, while large FOV is necessary when full brain coverage is required. As a result, the shortest TR we could use in practice is slightly longer than the TR associated with optimal SNR efficiency. Nevertheless, the TR used in this work (2.6/2.7s) provides a considerably higher SNR efficiency than the typical TR (>7s) that is commonly used for 2D diffusion MRI at 7T. The use of thin slabs might be less robust to truncation artefacts due to the small number of kz encodes, but this is not clearly visible in our results. The lengthening of T_1 at 7T is a disadvantage for SNR, but its effect is not substantial in this application compared to the signal gains from the increase in field strength.

Another advantage of 3D multi-slab acquisition is its intrinsically higher SNR for each diffusion volume, which is particularly beneficial in high-resolution diffusion MRI scan. Compared to simultaneous multi-slice imaging, the primary disadvantage

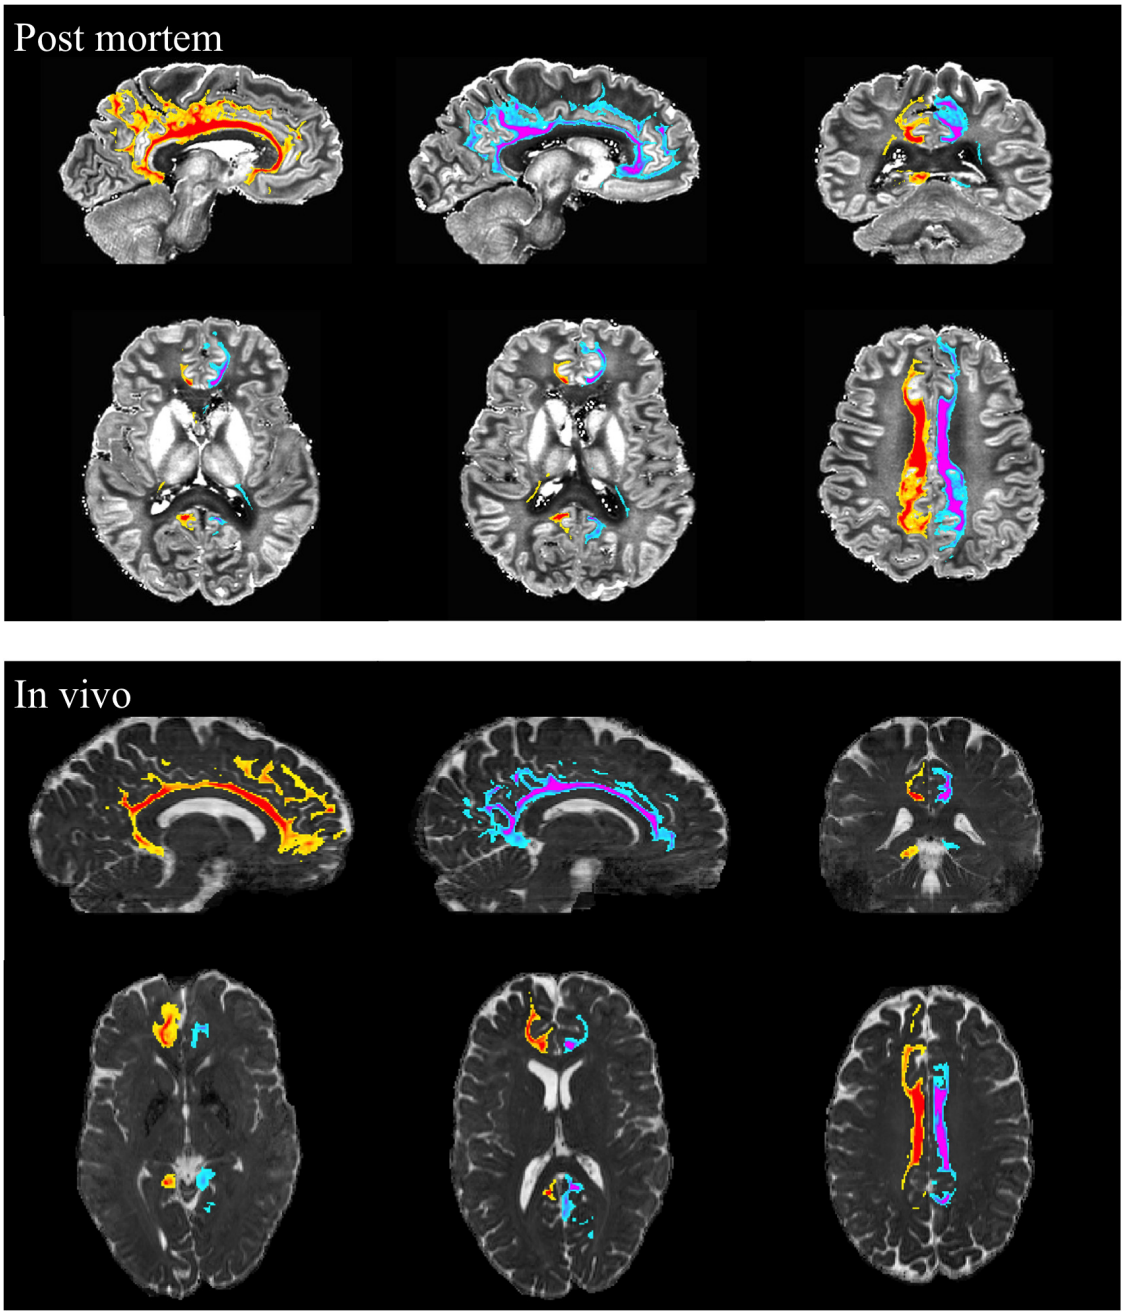


Figure 5.12: Tractography results obtained from the left (blue-magenta) and right (yellow-red) seed regions in the cingulum bundle. Top row: for reference, post-mortem data acquired with 0.94 mm^3 isotropic resolution using 3D diffusion weighted steady state free precession at 3T. Bottom row: in-vivo data acquired with 1.03 mm^3 isotropic resolution using spin-echo 3D multi-slab EPI at 7T. All results are rendered with a threshold of 100 streamlines and overlaid on mean diffusivity maps.

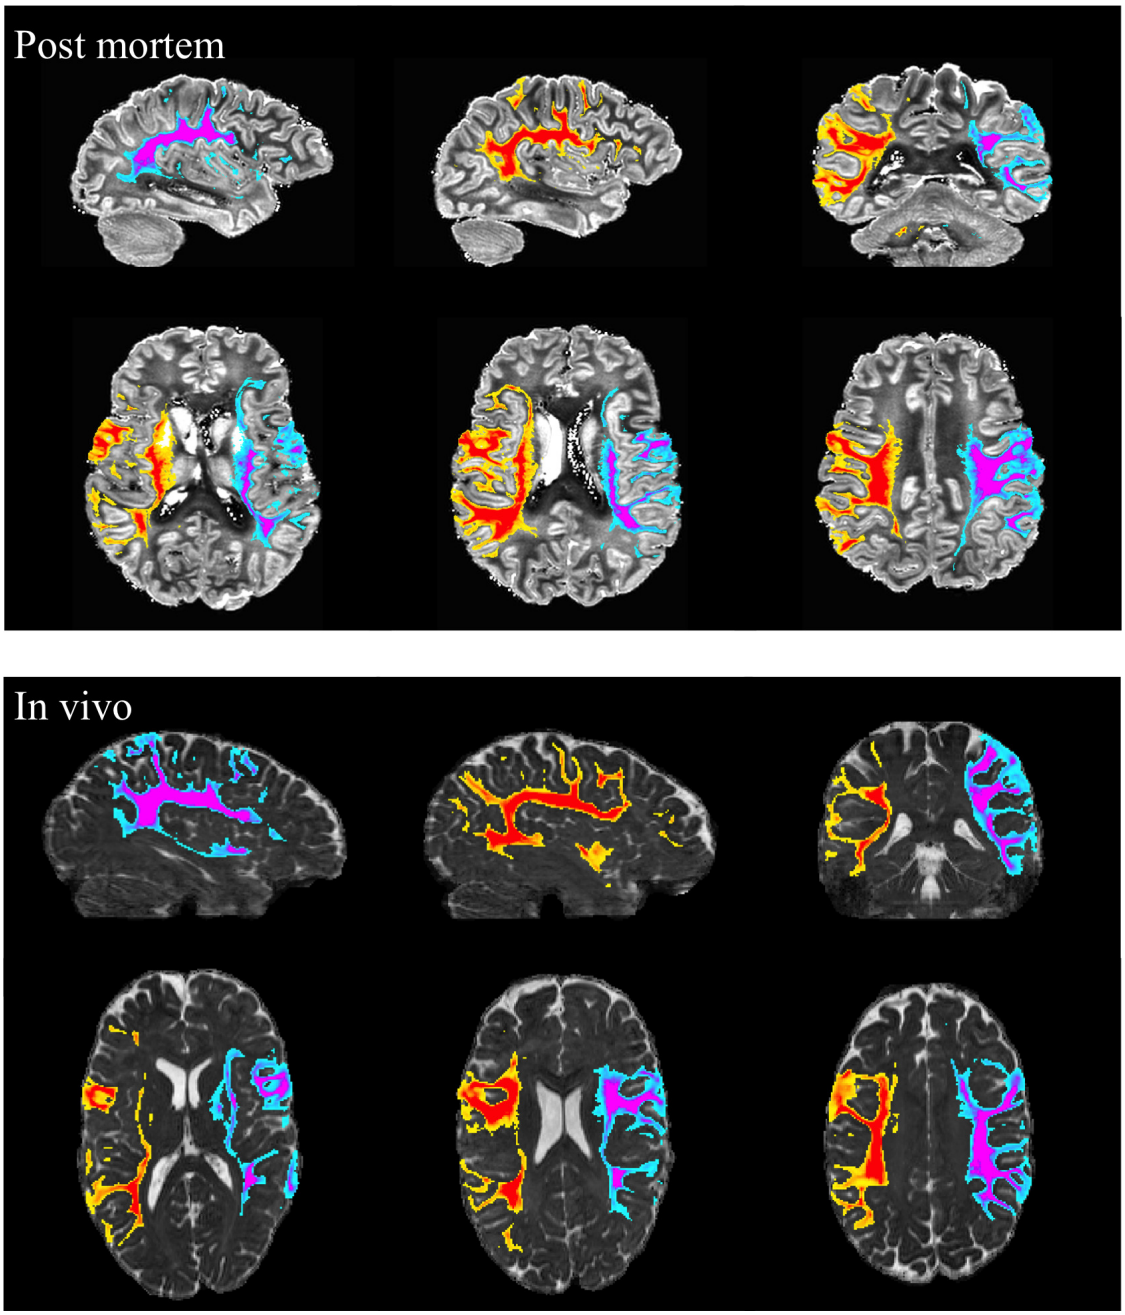


Figure 5.13: Tractography results obtained from the left (blue-magenta) and right (yellow-red) seed regions in SLF bundle for post-mortem data (top row) and in-vivo data (bottom row). All results are rendered with a threshold of 100 streamlines and overlaid on mean diffusivity maps.

of 3D multi-slab imaging is the increased scan time per volume, which limits the number of diffusion directions that can be achieved using the same scan time. The scan protocol used in this work allows more than 50 diffusion directions in a 30-minute scan, which should be sufficient for many diffusion MRI studies in vivo. One promising further improvement to this sequence would be the acquisition of multiple non-contiguous slabs simultaneously (i.e. the 3D extension of simultaneous multi-slice imaging), which would both reduce the excitation TR and the total scan time [121, 209].

It should be noted that the SAR predicted by the system is affected by some subject-specific factors, such as subject's weight and coil load. In the ACS evaluation protocol, we used RF pulses with a high time-bandwidth product to achieve a good slab profile, which results in a high SAR value when scanning the three subjects involved in this experiment. Three different subjects were subsequently scanned with the high-resolution diffusion MRI protocol, at which point the SAR limits were exceeded for one subject. We decided to increase the TR slightly (100 ms) to satisfy the SAR constraint without changing other sequence parameters.

Due to the severe $\mathbf{B}_0$ inhomogeneity and short tissue T_2^* at 7T, in-plane parallel acceleration is critical for SSH-EPI. The choice of acceleration factor is a tradeoff between the reduction in image distortion and blurring versus the amplification of noise. In this work, we used an in-plane GRAPPA factor of 3, which provides substantial reduction in image distortion and sufficient image SNR [5]. It should be noted that, due to the T_2^* decay during the EPI echo train, the ky phase encoding direction exhibits the largest blurring effect. With the acquisition protocol used in this work, the additional blurring in WM induced by the T_2^* ($\sim$ 25 ms [210]) decay is $\sim$20% (only in the in-plane phase-encoding direction), indicated by the full-width at half maximum (FWHM) of the point spread function [9].

In accordance with previous work [199], both Slab-FLEET ACS and Slice-FLEET ACS demonstrate greatly improved image quality compared with the

conventional methods (SSH-EPI, conventional segmented EPI), including reduced residual artefacts and better-preserved SNR, which are crucial for achieving high-quality diffusion MRI data. The difference in contrast between the gradient-echo FLEET ACS data and spin-echo imaging data does not appear to introduce substantial effects in the reconstruction, consistent with previous reports that GRAPPA is relatively insensitive to contrast differences [211]. Slice-FLEET acquires ACS for each slice, requiring considerably longer scan time (∼58s for 114 slices using 5 dummy scans) than Slab-FLEET (∼7s for 14 slabs using 5 dummy scans) but better captures local coil sensitivity and reduces intra-voxel de-phasing. This time penalty can be halved by using doubled slice thickness, without any visible differences in the GRAPPA reconstruction results (Fig. 5.4). Considering also that only one ACS acquisition is required for all directions, Slice-FLEET has only a minimal impact on the total scan time while providing the best reconstructions. The navigator signal is acquired without kz phase encoding, so the Slab-FLEET ACS is an apparent option for GRAPPA reconstruction and was used in this work, which added ∼7s on the total scan time. Nevertheless, the need for Slab-FLEET ACS could be also avoided by using slab-wise averaged Slice-FLEET ACS data, which should be able to provide similar results. Another possible FLEET-ACS scheme for multi-slab acquisition is to acquire 3D ACS slabs with full kz encodes. However, this scheme is expected to be very sensitive to subject motion and respiration due to the long interval between the acquisitions of the first and the last segments (i.e. precisely the reason why FLEET-ACS was proposed for 2D EPI), we therefore didn't consider it in this work.

It is worth noting that in this work the GRAPPA reconstruction was applied to each kz plane individually instead of using a true 3D GRAPPA reconstruction, which might further improve the reconstruction. Further studies will be needed to investigate the performance of 3D GRAPPA reconstruction on the 3D multi-slab data.

Slab-boundary artefacts are one major challenge with 3D multi-slab sequences. The recently proposed NPEN method has demonstrated its ability to reduce slab boundary artefacts effectively and robustly at 3T [7]. However, at 7T, this method has to contend with increased $\mathbf{B}_0$ inhomogeneity and reduced RF bandwidths due to SAR limitations, which distort and shift the slab profile. In this work, we modified the initialisation of the slab profile by incorporating $\mathbf{B}_0$ effects, which could provide a more accurate starting point for the iterative reconstruction. However, even with this modification, there are still some residual boundary artefacts visible in the region of strong $\mathbf{B}_0$ inhomogeneity, presumably because the slab shifting and distortion interfere with the NPEN assumption that artefacts are periodic along the slice dimension. A possible solution to this problem is to estimate a warping matrix based on the field map and include it in the reconstruction to restore the periodicity of the artefacts, such that it satisfies the assumption of NPEN. However, due to the slab aliasing, a single warping matrix cannot correctly restore both the underlying signal and the aliasing signal from one voxel, which might lead to incorrect reconstruction. Further investigation is needed to improve the NPEN correction under strong $\mathbf{B}_0$ inhomogeneity. Nevertheless, despite these residual artefacts, we are able to reconstruct high quality visualisations of fibre architecture from the data.

Subject motion is a major challenge for multi-shot diffusion MRI. In this work, a navigator based nonlinear correction method was applied to correct the motion induced phase errors. In general, the acquisition of 3D navigator is infeasible due to the long time required to traverse 3D k-space, although some groups have considered methods for batching 3D navigators across multiple repetitions [212]. In the current work, a 2D navigator was used in the correction, assuming that the phase error along the slice direction is constant if a thin slab is acquired [158]. To the extent that this assumption doesn't hold, there could be residual artefacts in the data, which might alter the quantification of diffusion anisotropy as they are dependent on diffusion directions. Further studies are needed to evaluate the effect

of these residual artefacts. It should be noted that if subject motion is very strong, neither 2D nor 3D nonlinear corrections could address the phase errors. In this case, methods like real-time image rejection and re-acquisition could be helpful [53].

5.6 Conclusion

Improving the resolution of diffusion MRI has been a topic of interest recently due to its ability to provide a more accurate depiction of complex fibre architectures. However, the reduced SNR at high resolution presents a substantial challenge to diffusion MRI, severely degrading the quality of high-resolution data. In this work, we utilize two powerful tools, 7T and 3D multi-slab acquisitions, to improve the SNR of diffusion MRI. Robust imaging required modification of two existing techniques: FLEET-ACS for improved parallel imaging and NPEN correction of slab-boundary artefacts. Two FLEET-ACS schemes for use with 3D multi-slab acquisition, Slab-FLEET and Slice-FLEET, were compared, with Slice-FLEET demonstrating better reconstruction quality. A new reconstruction scheme was also proposed for Slice-FLEET ACS data to make it compatible with navigator correction of phase errors. Results with $\sim$1 mm isotropic resolution demonstrate very high quality data, revealing fine anatomical features and supporting detailed delineation of white-matter tracts using tractography.

6

Accelerated k-q dMRI using Gaussian Processes

Contents

6.1 Introduction 132
6.2 Theory 133
6.2.1 Gaussian Process Regression 133
6.2.2 Gaussian Process Model for dMRI Data 136
6.2.3 dMRI reconstruction with a Gaussian Process Prior . . 138
6.3 Methods 144
6.3.1 k-q Sampling 144
6.3.2 Simulation 146
6.3.3 In Vivo 148
6.3.4 Reconstruction 150
6.3.5 Post-processing 151
6.4 Results 151
6.4.1 Simulation 151
6.4.2 In Vivo 159
6.5 Discussion 167
6.6 Conclusion 171

6.1 Introduction

Parallel imaging has become increasingly important for diffusion MRI. As discussed in Chapter 3, in-plane under-sampling is mainly used to reduce image artefacts (e.g. distortion and blurring) and achieve a shorter echo time, which are critical for achieving high-quality data, particularly at ultra-high field. In general, in-plane under-sampling is not used to speed up single-shot EPI as the diffusion preparation occupies more than half of the sequence time, although it can be used to accelerate segmented acquisitions. Through-plane acceleration using the SMS technique can effectively reduce scan time, particularly with the blipped-CAIPI scheme [1], which shifts image slices in-plane to suppress g-factor related noise amplification. These two forms of parallel imaging achieve different goals, both of which are important for diffuson MRI, and therefore one would ideally like to use both together. However, the combination of SMS and in-plane under-sampling faces significant challenges in dMRI. The efficacy of blipped-CAIPI will be penalized by the in-plane under-sampling due to the FOV folding effects, which results in significant noise amplification that is particularly problematic for dMRI due to its intrinsically low SNR. Consequently, typical dMRI protocols either use a high multiband factor with no in-plane under-sampling (e.g. MB=3 with R=1 [3, 4]) or use a low multiband factor to incorporate in-plane under-sampling (e.g. MB=2 with R=3 [5]).

Diffusion MRI typically acquires a number of diffusion volumes (q-space samples) and conventionally each volume is reconstructed independently. As reviewed in Chapter 3, the shared information between q-space samples has been used for q-space acceleration, where complex fibre architectures can be resolved with a smaller number of q-space samples using a compressed sensing based reconstruction [138–141]. In this work, we leverage the shared information between diffusion volumes to improve the reconstruction of k-space data. Specifically, correlations

between q-space samples are learned from the data using Gaussian process methods and used as a prior to improve the estimation of the underlying image.

Gaussian processes have been successfully exploied in dMRI data analysis with main applications in removing eddy-current induced image distortions and subject movements [213, 214]. The covariance between q-space samples can be modelled as a spherical function, which has been shown to have good coherence with the covariance observed from real dMRI data [213]. Hence, the spherical model is also used in the proposed method. The 'hyper-parameters' which specify the covariance function are estimated from the data using an iterative approach.

Numerical simulations and data acquired in vivo are used to validate the proposed method. Two acceleration schemes are evaluated, including using a very high in-plane under-sampling factor (i.e. reaching the theoretical limit of conventional parallel imaging in simulation) and using a high SMS acceleration factor with in-plane under-sampling (i.e. target the actual application in practice). Both simulation results and results in vivo demonstrate superior performance of the proposed method, compared with the conventional parallel imaging reconstruction methods.

6.2 Theory

6.2.1 Gaussian Process Regression

A Gaussian process (GP) is a statistical model defined as a probability distribution over functions $f(x)$, such that the function values evaluated at any arbitrary set of points $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_N]$ have a joint Gaussian distribution [215]. Hence, Gaussian processes can also be seen as an generalization of multivariate Gaussian distribution to an infinite dimensional space.

The joint distribution of the N random variables $\mathbf{f} = [f(\mathbf{x}_1), f(\mathbf{x}_2), \cdots, f(\mathbf{x}_N)]$ is completely specified by its mean and the covariance between any two variables:

$$f(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{6.1}$$

where $\boldsymbol{\mu}$ is the mean vector and $\boldsymbol{\Sigma}$ is the covariance matrix with each element representing the covariance between two variables.

The covariance matrix is symmetric with $N(N+1)/2$ unknown elements. A key feature of Gaussian processes is the covariance is modelled as a kernel function C, which is characterised by a few parameters $\boldsymbol{\Theta}$. The element of the covariance matrix can be determined as

$$\boldsymbol{\Sigma}_{i,j} = C(\mathbf{x}_i, \mathbf{x}_j) \tag{6.2}$$

The choice of the kernel function depends on applications. A basic property of a kernel function is that the covariance between two variables, $f(\mathbf{x}_i)$ and $f(\mathbf{x}_j)$, should be higher when the input points $\mathbf{x}_i$ and $\mathbf{x}_j$ are closer to each other. The parameters in the kernel function are called 'hyper-parameters', which have to be estimated as part of the model but in the end it is the $f(\mathbf{x})$ values that we actually care about.

Let's assume the hyper-parameters are known for the time being and we'll discuss their estimation later. We acquire a set of data $\mathbf{y} = [y(\mathbf{x}_1), y(\mathbf{x}_1), \cdots, y(\mathbf{x}_N)]$ which contain additive noise:

$$y(\mathbf{x}_i) = f(\mathbf{x}_i) + n_i \tag{6.3}$$

The noise n_i is further assumed to follow an independent, identically distributed (i.i.d) Gaussian distribution with a zero mean and a variance of σ_n. This gives the likelihood function

$$p(\mathbf{y}|\mathbf{f}, \mathbf{X}) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi {\sigma_n}^2}} \exp\Big(- \frac{(y(\mathbf{x}_i) - f(\mathbf{x}_i))^2}{2{\sigma_n}^2} \Big) \tag{6.4}$$

According to the Bayes' theorem, the posterior of the image $p(\mathbf{f}|\mathbf{y}, \mathbf{X})$ can be related to the likelihood $p(\mathbf{y}|\mathbf{f}, \mathbf{X})$ and the prior $p(\mathbf{y})$ as

$$p(\mathbf{f}|\mathbf{y}, \mathbf{X}) = \frac{p(\mathbf{y}|\mathbf{f}, \mathbf{X}) p(\mathbf{f})}{p(\mathbf{y}|\mathbf{X})} \tag{6.5}$$

with $p(\mathbf{y}|\mathbf{X})$ be the model evidence, which is independent of function values. With a GP prior (Eq. 6.1), the posterior is given by (see Appendix C for details):

$$\mathbf{f}\,|\,\mathbf{y},\mathbf{X} \sim \mathcal{N}(\sigma_n^{-2}(\sigma_n^{-2}+\boldsymbol{\Sigma}^{-1})^{-1}\mathbf{y}, (\sigma_n^{-2}+\boldsymbol{\Sigma}^{-1})^{-1}) \tag{6.6}$$

The posterior mean expressed as $\sigma_n^{-2}(\sigma_n^{-2}+\boldsymbol{\Sigma}^{-1})^{-1}\mathbf{y}$ is a point estimator known as the maximum a posteriori (MAP) estimation, which is a linear function of the observed data. The GP regression aims to predict the underlying image by filtering out the noise and the filter parameters are specified by the signal covariance and the noise variance. Note, the prior mean is set to zero here for simplicity of notation. Commonly, if prior knowledge of the mean is not available, the prior mean is set to zero by symmetry. Using a zero-mean prior is not a strict restriction as the posterior mean is not confined to be zero, which will be determined by the observation (Eq. 6.6).

The Gaussian process model can also be used to make predictions. Let $\mathbf{f}^* = [f^*(\mathbf{x}_{u_1}), f^*(\mathbf{x}_{u_2}), \cdots, f^*(\mathbf{x}_{u_N})]$ be the un-observed function values at points $\mathbf{X}^* = [\mathbf{x}_{u_1}, \mathbf{x}_{u_2}, \cdots, \mathbf{x}_{u_N}]$. The joint distribution of the observed data $\mathbf{y}$ and the unknown values $\mathbf{f}^*$ is a Gaussian:

$$\begin{bmatrix}\mathbf{y}\\ \mathbf{f}^*\end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix}\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}}+\sigma_n^2\mathbf{I} & \boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}^*}\\ \boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}} & \boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}^*}\end{bmatrix}\right) \tag{6.7}$$

where the covariance matrix is split into four sub-matrices: $\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}}$ is the covariance matrix evaluated using the observed points; $\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}^*}$ and $\boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}}$ are the covariance matrices evaluated using all pairs between the observed and un-observed points. $\boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}^*}$ is the covariance matrix derived using the un-observed points. The prediction of the unknown function values $\mathbf{f}^*$ can be obtained by conditioning the posterior on the observation:

$$\mathbf{f}^*\,|\,\mathbf{y},\mathbf{X},\mathbf{X}^* \sim \mathcal{N}(\boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}}(\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}}+\sigma_n^2\mathbf{I})^{-1}\mathbf{y}, \boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}^*}-\boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}}(\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}}+\sigma_n^2\mathbf{I})^{-1}\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}^*}) \tag{6.8}$$

where the MAP estimator, $\boldsymbol{\Sigma}_{\mathbf{X}^*,\mathbf{X}}(\boldsymbol{\Sigma}_{\mathbf{X},\mathbf{X}}+\sigma_n^2\mathbf{I})^{-1}\mathbf{y}$, is also a linear combination of the observed data as in Eq. 6.6.

Training a Gaussian process includes the selection of the covariance function and the estimation of the hyper-parameters, $\mathbf{\Theta}$. Here, we just consider the case that the covariance function has been determined, then the optimal hyper-parameters can be estimated by maximising the log marginal likelihood of observing values $\mathbf{y}$ at points $\mathbf{X}$ given the hyper-parameters $\mathbf{\Theta}$:

$$\log p(\mathbf{y}|\mathbf{X}, \mathbf{\Theta}) = -\frac{1}{2}\mathbf{y}^{\mathrm{T}}(\sigma_n^2\mathbf{I} + \mathbf{\Sigma})^{-1}\mathbf{y} - \frac{1}{2}\log|\sigma_n^2\mathbf{I} + \mathbf{\Sigma}| + const \tag{6.9}$$

which can be solved iteratively using methods like conjugate gradient.

6.2.2 Gaussian Process Model for dMRI Data

The dMRI signal can be seen as the output of a real-valued function $S(\mathbf{q})$ with the input $\mathbf{q} = [\mathbf{q}_1, \mathbf{q}_2, \cdots, \mathbf{q}_{N_d}]$ defined in q-space. Each input point $\mathbf{q}$ is specified by a diffusion-encoding direction $\mathbf{g}$ and a b value. Modelling dMRI signal as GPs has been previously proposed and successfully implemented in data analysis [213]. As dMRI data are typically acquired on one or several q-space spheres/shells, a proper choice for the covariance function should be defined based on the angular differences between diffusion directions and the distances between shells (for multi-shell data acquired with different b values). It has been previously demonstrated that the 'Spherical model' offers superior performance over other models for depicting dMRI data [213]. For single-shell acquisition (single b value for all dMRI data), the covariance function based on the Spherical mode is given by

$$C_\theta(\theta_{i,j}; a) = \begin{cases} 1 - \dfrac{3\theta_{i,j}}{2a} + \dfrac{{\theta_{i,j}}^3}{2a^3} & \text{if } \theta \leq a \\ 0 & \text{if } \theta > a \end{cases} \tag{6.10}$$

where $\theta_{i,j}$ denotes the angle between the two vectors pointing from q-space origin to q-space points $\mathbf{q}_i$ and $\mathbf{q}_j$, defined as $\theta_{i,j} = \arccos\langle \mathbf{g_i}, \mathbf{g_j} \rangle$. If the spherical symmetry of dMRI data is taken into account, the definition is changed to $\theta_{i,j} = \arccos|\langle \mathbf{g_i}, \mathbf{g_j} \rangle|$. The covariance function defined in Eq. 6.10 gives a large value when the two points are closely located and a smaller value when they are distant from each other.

The decay rate of the covariance is controlled by the hyper-parameter a. The covariance is set to zero if θ is smaller than a.

For multi-shell dMRI data, the covariance function described in Eq. 6.10 should be modified to model the covariance between signal from different shells. This can be achieved by constructing an integrated covariance function based on two basic covariance functions:

$$C(\mathbf{q}_i, \mathbf{q}_j) = C_{\theta_{i,j}}(\theta; a) C_b(b_i, b_j) \tag{6.11}$$

where b_i and b_j are the b values corresponding to the q-space points $\mathbf{q}_i$ and $\mathbf{q}_j$, respectively. $C_b(b_i, b_j)$ is a function representing the signal covariance between shells (b values). Previous work has advocated modelling C_b as [213]:

$$C_b(b_1, b_2; l) = e^{-\frac{(\log b_1 - \log b_2)^2}{2l^2}} \tag{6.12}$$

where l is a hyper-parameter characterising the length scale, which determines how far away two points need to be in order to become irrelevant.

The current work is focused on single-shell data and thus only the covariance function described by Eq. 6.10 is considered in the rest of this Chapter. In order to model the signal variability and measurement noise, two other hyper-parameters can be added to the model, which gives the covariance matrix as

$$\mathbf{\Sigma}_{i,j} = \lambda C(\theta_{i,j}; a) + \sigma^2 \delta_{i,j} \tag{6.13}$$

where λ is a signal scaling factor capturing the variability of the signal and σ is the variance of the noise. δ is the Dirac function. There are three hyper-parameters, a, λ, σ in the Spherical model for single-shell dMRI signal, which can be estimated from the observed data by maximizing the marginal likelihood as described by Eq. 6.9. In this work, the hyper-parameter fitting is applied globally, i.e. estimating one set of hyper-parameters for all imaging voxels as in [213].

6.2.3 dMRI reconstruction with a Gaussian Process Prior

This section describes the method we proposed to reconstruct under-sampled dMRI data based on Gaussian processes.

Signal Model

As introduced in 2.3.1, a linear MRI signal model can be expressed as

$$\mathbf{d}_c = \mathbf{MFS}_c\mathbf{u}_0 + \mathbf{n}_c \tag{6.14}$$

with $c = 1, \cdots, N_c$, where $\mathbf{u}_0$ $(N_i \times 1)$ is the unknown image vector, $\mathbf{d}_c$ $(N_k \times 1)$ is the k-space data for one image volume acquired from the c^{th} channel. $\mathbf{n}_c$ $(N_k \times 1)$ is the measurement noise, which is assumed to be additive and follow an i.i.d zero-mean Gaussian distribution with a variance of σ_n^2. $\mathbf{S}_c$ $(N_i \times N_i)$ denotes the coil sensitivity map, which can be measured from the scan. $\mathbf{F}$ $(N_i \times N_i)$ and $\mathbf{M}$ $(N_k \times N_i)$ are the Fourier transform matrix and k-space sampling matrix, respectively.

In diffusion MRI, multiple image volumes are acquired and thus the linear model can be reformulated as

$$\mathbf{d}_{\mathbf{q},c} = \mathbf{M}_\mathbf{q}\mathbf{FS}_c\mathbf{u}_\mathbf{q} + \mathbf{n}_{\mathbf{q},c} \tag{6.15}$$

where $\mathbf{u}_\mathbf{q}$ and $\mathbf{d}_\mathbf{q}$ are the unknown images and the acquired k-space data for diffusion volume $\mathbf{q} = [\mathbf{q}_1, \mathbf{q}_2, \cdots, \mathbf{q}_{N_d}]$ with N_d denoting the total number of diffusion volumes. $\mathbf{M}_\mathbf{q}$ is the k-space sampling matrix, which is usually identical for all diffusion volumes but can also be different between diffusion volumes to favour particular reconstruction methods (see the Methods section).

For simplicity of notation, we define $\mathbf{A}_{\mathbf{q},c} = \mathbf{M}_\mathbf{q}\mathbf{FS}_c$ and let $\mathbf{A}$ be the full system matrix including all diffusion volumes and channels, which has the size of $N_kN_cN_d \times N_iN_d$. In addition, we define $\mathbf{d}$ $(N_kN_cN_d \times 1)$ as the full k-space data vector concatenated across all channels and diffusion volumes and $\mathbf{u}$ $(N_iN_d \times 1)$ as the full image vector concatenated across all diffusion volumes.

The linear signal model and noise characteristics give the likelihood function

$$p(\mathbf{d}|\mathbf{u}) = \prod_{c=1}^{N_c}\prod_{i=1}^{N_d}\frac{1}{\sqrt{2\pi{\sigma_n}^2}}\exp\Big(-\frac{(\mathbf{d}_{\mathbf{q}_i,c}-\mathbf{A}_{\mathbf{q}_i,c}\mathbf{u}_{\mathbf{q}_i})^2}{2{\sigma_n}^2}\Big) \tag{6.16}$$

The maximal likelihood (ML) estimation of the unknown image is given by

$$\mathbf{u}_\mathbf{q} = \arg\min_{\mathbf{u}_\mathbf{q}}\sum_{c=1}^{N_c}\Big\|\mathbf{A}_{\mathbf{q},c}\mathbf{u}_\mathbf{q}-\mathbf{d}_{\mathbf{q},c}\Big\|_2^2 \tag{6.17}$$

which is equivalent to the volume-by-volume SENSE reconstruction. The reconstruction of all diffusion volumes can be formulated as a single equation using the full system matrix $\mathbf{A}$:

$$\mathbf{u} = \arg\min_{\mathbf{u}}\Big\|\mathbf{A}\mathbf{u}-\mathbf{d}\Big\|_2^2 \tag{6.18}$$

As discussed in Chapter 2, if the k-space data are highly under-sampled, the system matrix will become ill-conditioned. In such case, solving Eq. 6.17 directly will result in strong noise amplification and residual aliasing artefacts. A common approach to alleviate this problem is to apply regularization in the reconstruction, which imposes particular constraints in the reconstruction based on some prior knowledge of the solution. A widely used regularization method in MRI reconstruction is the Tikhonov regularization, which has a typical form of

$$\mathbf{u} = \arg\min_{\mathbf{u}}\left(\Big\|\mathbf{A}\mathbf{u}-\mathbf{d}\Big\|_2^2+\beta\Big\|\mathbf{u}\Big\|_2^2\right) \tag{6.19}$$

In the bracket is the cost function, which consists of two terms: the data fidelity term (left) and the regularization term (right). Minimizing the cost function will lead to a solution that achieves a trade-off between fitting the data and having a small L_2 norm for the solution. The regularization parameter β controls this trade-off.

In the Bayesian formulation, image reconstruction with a Tikhonov regularization corresponds to the MAP estimation with a Gaussian prior on the image:

$$\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \omega^2\,\mathbf{I}) \tag{6.20}$$

which is a Gaussian distribution with a zero mean and a variance of ω^2. Invoking Bayes' theorem, the MAP estimation is given by

$$\mathbf{u} = \arg\min_{\mathbf{u}} \left(\frac{1}{2\sigma_n^2} (\mathbf{Au} - \mathbf{d})^{\mathrm{T}} (\mathbf{Au} - \mathbf{d}) + \frac{1}{2\omega^2} \mathbf{u}^{\mathrm{T}} \mathbf{u} \right) \tag{6.21}$$

which is equivalent to the SENSE reconstruction with Tikhonov regularization (Eq. 6.19).

The prior considered in Eq. 6.20 assumes no correlations between the elements of $\mathbf{u}$, which can also be modelled to improve the reconstruction. The image vector $\mathbf{u}$ contains image data for all diffusion volumes and thus priors can be added in both image space and q-space. The current work focuses on the investigation of the data correlation in q-space. Nevertheless, priors in the image space can be also included, which will be studied in future work.

Gaussian Process Prior

As introduced in section 6.2.2, diffusion MRI signal can be modelled as Gaussian processes, such that the signal observed at points $\mathbf{q} = [\mathbf{q}_1, \mathbf{q}_2, \cdots, \mathbf{q}_{N_d}]$ follows a multi-variate Gaussian distribution:

$$\mathbf{u} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma} \otimes \mathbf{I}_{N_i}) \tag{6.22}$$

where $\mathbf{u}$ is real-valued diffusion signal. $\boldsymbol{\mu}$ $(N_i N_d \times 1)$ is the mean vector. $\boldsymbol{\Sigma}$ $(N_d \times N_d)$ is the covariance matrix constructed based on the spherical model (Eq. 6.10). $\mathbf{I}_{N_i}$ is a $N_i \times N_i$ identity matrix. $\otimes$ is the Kronecker product, which is used to produce a $N_d N_i \times N_d N_i$ covariance matrix as

$$\boldsymbol{\Sigma} \otimes \mathbf{I}_{N_i} = \begin{bmatrix} \Sigma_{1,1}\mathbf{I}_{N_i} & \Sigma_{1,2}\mathbf{I}_{N_i} & \cdots & \Sigma_{1,N_d}\mathbf{I}_{N_i} \\ \Sigma_{2,1}\mathbf{I}_{N_i} & \Sigma_{2,2}\mathbf{I}_{N_i} & \cdots & \Sigma_{2,N_d}\mathbf{I}_{N_i} \\ \vdots & \vdots & \cdots & \vdots \\ \Sigma_{N_d,1}\mathbf{I}_{N_i} & \Sigma_{N_d,2}\mathbf{I}_{N_i} & \cdots & \Sigma_{N_d,N_d}\mathbf{I}_{N_i} \end{bmatrix} \tag{6.23}$$

This formation assumes the same covariance matrix for all image voxels, as in the Eddy implementation [213].

With the GP prior, the MAP solution has the form of

$$\mathbf{u} = \arg\min_{\mathbf{u}} \left(\frac{1}{2\sigma_n^2} (\mathbf{Au} - \mathbf{d})^{\mathrm{T}} (\mathbf{Au} - \mathbf{d}) + \frac{1}{2} (\mathbf{u} - \boldsymbol{\mu})^{\mathrm{T}} (\boldsymbol{\Sigma} \otimes \mathbf{I}_{N_i})^{-1} (\mathbf{u} - \boldsymbol{\mu}) \right) \quad (6.24)$$

To solve this equation, one needs to first determine the covariance matrix $\boldsymbol{\Sigma}$, which is specified by three hyper-parameters, λ, a and σ_n. The optimal hyper-parameters can be estimated by maximizing the log marginal likelihood:

$$\log p(\mathbf{d}|\lambda, a, \sigma_n) = -\frac{1}{2}\mathbf{d}^{\mathrm{T}}(\sigma_n^2 + \mathbf{A}(\boldsymbol{\Sigma} \otimes \mathbf{I}_{N_i})\mathbf{A}^{\mathrm{T}})^{-1}\mathbf{d} - \frac{1}{2}\log|\sigma_n^2 + \mathbf{A}(\boldsymbol{\Sigma} \otimes \mathbf{I}_{N_i})\mathbf{A}^{\mathrm{T}}| + const \quad (6.25)$$

However, solving this minimization problem requires the inversion of a large matrix $[\sigma_n^2 + \mathbf{A}(\boldsymbol{\Sigma} \otimes \mathbf{I}_{N_i})\mathbf{A}^{\mathrm{T}}]$ $(N_k N_c N_d \times N_k N_c N_d)$, which could be computational expensive and numerically unstable.

Hyper-parameter Estimation and Image Reconstruction

In this work, we use an iterative scheme to jointly estimate the optimal hyper-parameters and the underlying images: in each iteration, a set of hyper-parameters is calculated from the images reconstructed in the previous iteration, and then a new image reconstruction is conducted with the updated hyper-parameters. As the iteration goes on, it is expected both the hyper-parameter estimation and the image reconstruction can be improved.

The noise in the reconstructed images (σ_{im}) generally have different characteristics from the noise in the k-space data (σ_n). Therefore, we cannot directly use the noise variance estimated from the reconstructed images in the MAP reconstruction (Eq. 6.24). Nevertheless, as shown in 2.3.1, the propagation of noise in parallel imaging can be quantitatively measured (g-factor) based on sensitivity maps. Hence, an estimate of σ_n could be derived from σ_{im}.

The pipeline of the proposed method is summarised below:

Iteration $n = 1$: Apply conventional SENSE to reconstruct all diffusion volumes, which produces a image set $\mathbf{u}^0$. A large number of voxels are extracted from $\mathbf{u}^0$, which are used as training data to estimate the GP hyper-parameters (λ^0, a^0 and σ_{im}^0). Here, 1000 randomly sampled brain voxels are used as in [213]. The final estimates for λ^0 and a^0 are an average of the estimates from all voxels. A covariance matrix Σ^0 is then constructed with a^0 and λ^0. The posterior mean $\boldsymbol{\mu}^0$ is also calculated. The noise variance in the k-space data is estimated from σ_{im}^0 as

$$\sigma_n = \frac{1}{1000} \sum_{\rho=1}^{1000} \frac{\sigma_{im}^0(\rho)}{g^\star(\rho)} \tag{6.26}$$

where $g^\star$ is noise variance propagation function, which can be derived in a similar way as the g-factor calculation (see 2.3.1). Particularly, for Cartesian sampling with an under-sampling factor of R, the values of $g^\star$ for R collapsed voxels are given by

$$g^\star = \mathrm{diag}[(\widetilde{\mathbf{S}}^{\mathrm{H}} \boldsymbol{\Phi}^{-1} \widetilde{\mathbf{S}})^{-1}] \tag{6.27}$$

where $\widetilde{\mathbf{S}}$ ($N_c \times R$) denotes the coil sensitivities for the R voxels and $\boldsymbol{\Phi}$ ($N_c \times N_c$) is the coil noise matrix.

Iteration $n = 2$ to $\mathbf{N}_{iter}$: MAP reconstruction of all diffusion volumes using the covariance matrix and posterior mean derived from previous iteration as priors in the current iteration:

$$\mathbf{u}^n = \arg\min_{\mathbf{u}} \left(\left\| \mathbf{A}\mathbf{u} - \mathbf{d} \right\|_2^2 + {\sigma_n}^2 (\mathbf{u} - \boldsymbol{\mu}^{n-1})^{\mathrm{T}} (\Sigma^{n-1} \otimes \mathbf{I}_{N_i})^{-1} (\mathbf{u} - \boldsymbol{\mu}^{n-1}) \right) \tag{6.28}$$

and then another GP is trained using 1000 randomly sampled brain voxels from $\mathbf{u}^n$, which gives the new hyper-parameters a^n, λ^n and σ_{im}^n. Calculate the new covariance matrix Σ^n and the posterior mean $\boldsymbol{\mu}^n$, which will be used as priors in the next iteration.

Output: Iteration stops when a fixed number of iterations is reached or the normalized image update $\|\mathbf{u}^n - \mathbf{u}^{n-1}\| / \|\mathbf{u}^{n-1}\|$ is below a preset threshold. The output of the last MAP reconstruction (Eq. 6.28) will be the final result.

Dealing with Phase Errors

In the previous section, the dMRI image is considered as real-valued. However, as mentioned in previous Chapters, phase errors usually arise in dMRI data due to eddy currents, physiology effects and subject motion. These phase errors have no relation to the q-space locations and thus violate the GP assumption that information is relatively consistent in local q-space. Therefore, direct reconstruction of the phase-error corrupted dMRI data using the proposed method will lead to strong artefacts. Phase from other sources such as $\mathbf{B}_0$ field inhomogeneity and local field inhomogeneity due to tissue susceptibility variations, is highly consistent across diffusion volumes (both locally and globally in q-space), so it is still compatible with the proposed reconstruction.

In this work, we address the phase-error problem by modelling it explicitly:

$$\tilde{\mathbf{u}} = \mathbf{P}\mathbf{u} \tag{6.29}$$

where $\mathbf{u}$ and $\tilde{\mathbf{u}}$ denote the underlying image and the phase-error corrupted image, respectively. $\mathbf{P}$ is a diagonal matrix representing the phase errors. The iterative reconstruction described in Eq. 6.28 can be modified into two equivalent forms:

$$\mathbf{u}^n = \arg\min_{\mathbf{u}} \left(\left\| \mathbf{A}\mathbf{P}\mathbf{u} - \mathbf{d} \right\|_2^2 + {\sigma_n}^2 (\mathbf{u} - \boldsymbol{\mu}^{n-1})^{\mathrm{H}} (\Sigma^{n-1} \otimes \mathbf{I}_{N_i})^{-1} (\mathbf{u} - \boldsymbol{\mu}^{n-1}) \right) \tag{6.30}$$

and

$$\tilde{\mathbf{u}}^n = \arg\min_{\tilde{\mathbf{u}}} \left(\left\| \mathbf{A}\tilde{\mathbf{u}} - \mathbf{d} \right\|_2^2 + {\sigma_n}^2 (\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{u}} - \boldsymbol{\mu}^{n-1})^{\mathrm{H}} (\Sigma^{n-1} \otimes \mathbf{I}_{N_i})^{-1} (\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{u}} - \boldsymbol{\mu}^{n-1}) \right) \tag{6.31}$$

where the first model estimates the uncorrupted images and the second estimates the corrupted ones. The corrupted and uncorrupted images can be trivially derived from each other with the estimated phase errors $\mathbf{P}$. The second model is implemented in the current work. The phase errors are assumed to be slowly varying in space, which can be extracted from a low-resolution navigator acquired following each imaging

readout. Alternatively, as demonstrated in MUSE method [70], the phase errors can also be estimated from the imaging data itself without a navigator, provided the conventional parallel imaging reconstruction is able to eliminate most of the aliasing artefacts. In the rest of this Chapter, these two approaches are referred to as navigator-based and navigator-free approaches, respectively. The navigator-free approach could be very useful at ultra-high field, particularly with SMS application, as it enables a lower SAR by avoiding the acquisition of navigators, which usually requires another refocusing pulse. The scan efficiency can also be improved by the disposal of the navigator acquisition. It should be noted that the phase extracted from either the navigator or the imaging data contains not only the phase errors but also the low-frequency components of the phase induced by the field inhomogeneity. Removing both phase terms realigns the image phase across all diffusion volumes, making the images data close to real-valued.

The proposed method is called Diffusion Acceleration with Gaussian-process Estimated Reconstruction (DAGER) in the rest of this Chapter.

6.3 Methods

6.3.1 k-q Sampling

Like in conventional parallel imaging methods, the precise pattern of sampling can impact on the quality of reconstruction. In this section, we consider how sampling in k-q space can be designed to aid DAGER reconstruction.

For joint reconstruction of all diffusion volumes, it has been reported that improved reconstructions can be achieved by using different k-space sampling patterns for different diffusion volumes [147–149]. This sampling strategy allows the reconstruction to utilize complementary k-space information across diffusion volumes, which is expected to improve the recovery of missing k-space data.

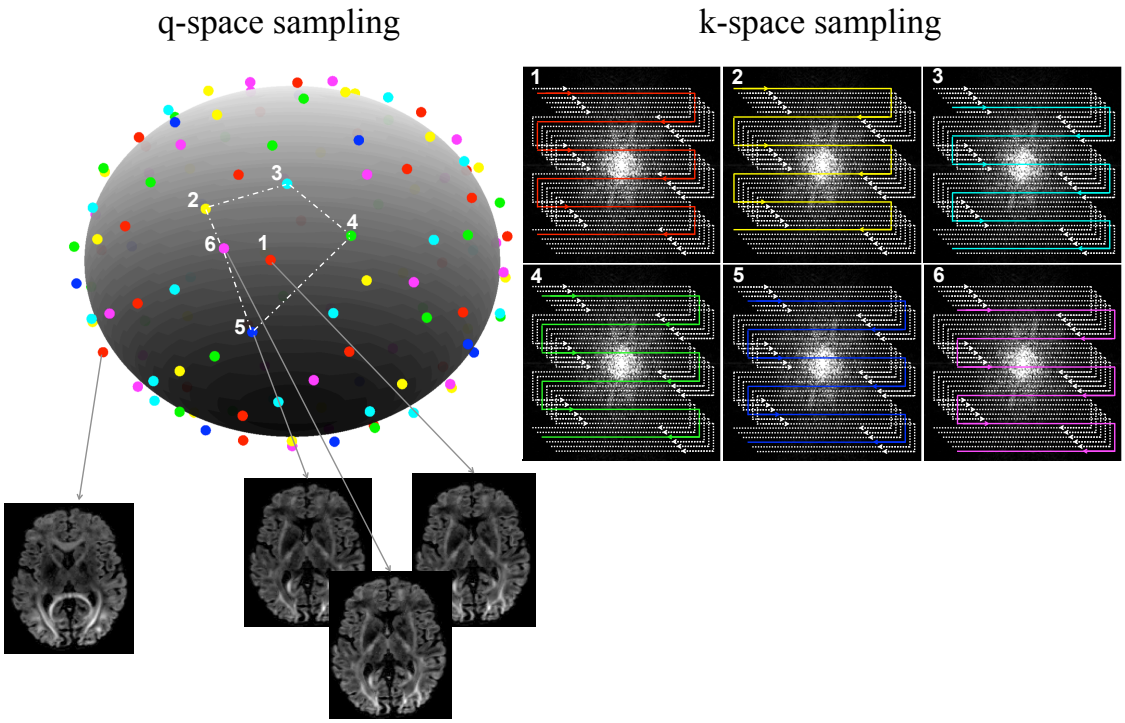


Figure 6.1: DAGER reconstructions take advantage of local smoothness in q-space: dMRI images that are near to each other in q-space look very similar (as shown in the inset dMRI images). The q-space sampling is shown on the left, where each colour corresponds to a different under-sampling pattern in k-space as shown on the right. The local neighborhood of q-space point 1 (red) is defined as the area delineated by dash lines, which include points 1 - 6.

Nearby q-space points will in general share more common features (i.e. the images are more similar) than distant points (Fig. 6.1). We therefore hypothesize we can improve the joint reconstruction by incoherently under-sampling k-space within local neighborhood in q-space. In this work, a graph model approach is proposed to design locally incoherent k-space under-sampling.

For SSH-EPI with an under-sampling factor of R, there exist R unique sampling patterns (produced by shifting the under-sampled trajectory along the phase-encoding direction) that can be applied at different q-space points. For each point, a neighborhood is defined including itself and its R-1 nearest neighbors (measured by angular distance). To maximize the local incoherence of k-space under-sampling patterns, it is desirable to assign different k-space sampling patterns for the q-space points within each q-space neighborhood (Fig. 6.1). This can be equivalently

described as a graph-coloring problem, where each q-space point corresponds to a vertex (colored in the q-space sampling map in Fig. 6.1) and an edge between two vertexes exists if the corresponding two q-space points are within the same neighborhood. If the constructed graph can be colored using R different colors while obeying the constraint that no connected vertexes share the same color, the resultant coloring scheme provides locally-incoherent k-space sampling.

Graph coloring is an "NP complete" problem, which means no efficient algorithms are available that provide the exact optimal solution. We applied a greedy algorithm [216] to approximate the solution. However, for high reduction factors (highly connected graph), the greedy algorithm fails to find a solution. To address the problem, the size of neighborhoods was slightly reduced when constructing the graph (e.g. use a neighborhood size of R-1).

For comparison, we also implemented three other sampling methods used in previous work: fixed k-space sampling for all q-space points (i.e. conventional sampling); random sampling [147]; linear circulant sampling [148, 149]. Random sampling of k-space is widely used for compressed sensing, but is not very compatible with EPI due to hardware limitation and reconstruction challenges; instead, we randomize the distribution of EPI under-sampling pattern across different q-space points to achieve incoherent sampling (i.e. the sampling patterns depicted in Fig. 6.1). For linear circulant sampling, the angular distances between all q-space points and a pre-selected reference q-space point is calculated and sorted in an ascending order, then a circulant sampling pattern is constructed following the sorted order.

6.3.2 Simulation

The simulations were constructed based on the HCP data, which had already been extensively pre-processed to account for eddy currents and distortions [37, 217].

In-plane Under-sampling

Diffusion MRI data sets with 128 directions and b=1000s/mm^2 were simulated using a ball-and-stick model ([204, 205]). Multi-channel data were generated with coil sensitivity maps measured from an 8-channel coil. Noisy measurements were produced by adding complex random noise to the data. Three in-plane under-sampling factors R= 4,6,10 were evaluated with R=10 exceeding the theoretical parallel imaging limits. To investigate the effects of image phase on DAGER reconstruction, motion induced phase errors were simulated, which consisted of a linear phase term corresponding to a shift along ky by a random distance between [0, 2Δky] and a constant phase term randomly selected between [$-\pi$, π]. The effects of non-linear phase errors were not studied in the simulation, which were demonstrated in the experiments performed in vivo.

SMS with In-plane Under-sampling

An SMS data set was simulated from four slices of single-band image data (MB=4) with 192 directions and b=1000s/mm^2. The slice gap is 31.5mm. The blipped-CAIPI scheme was used to introduce a FOV/4 shift between adjacent SMS slices. In-plane under-sampling factor R=3 was applied, leading to a total acceleration factor of 12. Noise and phase errors were added to the data as described in the previous section. Coil sensitivity maps were measured using a 32-channel coil.

The in-plane under-sampling simulations, which used a small number of channels, aimed to test the limits of the proposed method by using extreme acceleration along one dimension (which was not necessarily what one would do in practice), including the case where the number of channels was exceeded by the under-sampling factor (and therefore the information supporting the reconstruction had to come from q-space). The SMS simulations were more targeting the actual application we would like to achieve in practice, which was to improve the ability to accelerate in-plane

and through-plane simultaneously. Therefore, the same number of channels were used for the SMS simulation as for the in vivo data acquisition.

6.3.3 In Vivo

Four subjects were scanned on a Siemens 7T scanner. Informed consent in accordance with local ethics was obtained before each scan.

In-plane Under-sampling

Diffusion MRI data were acquired from two subjects using a 32-channel coil. Bipolar diffusion preparation (twice-refocused spin echo [38]) scheme was used to reduce eddy current effects. Data with acceleration factors R=3 and R=6 were acquired using SSH-EPI. After each imaging echo, a 64×64 navigator was acquired. For the R=6 imaging data, the navigator was under-sampled by a factor of 4 to satisfy the conventional parallel imaging reconstruction and produce similar distortions as the imaging data. For subject one, 20 slices were acquired with $1.3 \times 1.3\text{cm}^2$ resolution, matrix size 192×192, slice thickness 2mm, b=1000s/mm^2, 128 directions, TR/TE=3500ms/90ms, total scan time is ∼7.5min; For subject two, 40 slices were acquired with 1.5 mm^3 isotropic resolution, matrix size 168×168, b=1000s/mm^2, 128 directions, TR/TE=6500ms/82ms, total scan time ∼13.8min. The R=3 data were acquired with partial Fourier 6/8, which were used as references. 30 sets of b=0 data were acquired, including 24 'blipped-down' volumes and 6 'blipped-up' volumes.

SMS with In-plane Under-sampling

One of the challenges for SMS at 7T is that many 'standard' protocol settings, like the use of the bipolar diffusion preparation, carry a heavy SAR burden that is completely exacerbated by SMS. Thus, a modified monopolar diffusion encoding scheme [218] was used in this work for SMS dMRI acquisition, which was based on the conventional monopolar preparation ('Stejskal-Tanner') , where part of the diffusion encoding from the second diffusion gradient (following the 180° pulse) was placed

before the 180° pulse with amplitude negated. This preparation approach produced a lower SAR and a shorter TE compared with the bipolar preparation scheme and enabled reduced eddy-current distortions compared with the conventional monopolar scheme. A MB pulse with TBWP/duration=2.52/10.9ms was used for excitation. A MultiPINS [132] refocusing pulse with the same TBWP and duration was implemented for refocusing. The MultiPINS pulse could achieve substantially reduced SAR compared with the conventional MB refocusing pulse, allowing SMS acquisition with a high MB factor and a relatively short TR. From simulation and in vivo results, it was found that the initial SENSE reconstruction was able to provide a good estimation of the phase errors. Therefore, we decided to eliminate the navigator in SMS acquisition, which could further reduce the SAR burden (navigator acquisition required another RF refocusing pulse) and improve the scan efficiency. With all these improvements, the final sequence for the MB=4 acquisition could work at a relatively low SAR level (<50% as monitored on the scanner), which made it possible to use higher multi-band acceleration and/or better RF pulses.

SMS data were acquired from four subjects, including the two subjects involved in the in-plane under-sampling experiments described above. Four slices (MB=4) were simultaneously acquired using the blipped-CAIPI scheme, where a FOV/4 shift was introduced between adjacent SMS slices. 21 sets of MB slices with 31.5mm slice gap were acquired, resulting in a total of 84 slices (126mm) for full brain coverage. In-plane under-sampling R=3 was applied. Other scan parameters are 1.5 mm^3 isotropic resolution, matrix size 156×156, TR/TE=3500ms/71ms, b=1000s/mm^2, 192 diffusion directions, total scan time is ∼11mins. 36 b=0 volumes were acquired (24 blipped-down volumes and 12 blipped-up volumes). For one subject, three repetitions of single-band (SB) data were acquired using the same sequence parameters (TE, TR, diffusion encoding, etc.) as in the SMS protocol but containing only 21 slices to match the TR, leading to a total acquisition time of ∼33min. This set of data was used as a high SNR reference. For the other

three subjects, a set of single-band (SB) data was acquired with the same FOV (i.e. 84 slices) as the SMS protocol, which was used as a full FOV reference. A longer TR (12000ms) and a smaller number of diffusion directions (50) were used to achieve similar acquisition time.

6.3.4 Reconstruction

In addition to DAGER reconstruction, images were also reconstructed using conventional SENSE. The implementation details of DAGER reconstruction are as follows:

Stopping criteria

Iteration stopped when either the maximum iteration number (set to 20) was reached or the normalized solution update ($\|\mathbf{u}^n - \mathbf{u}^{n-1}\|/\|\mathbf{u}^{n-1}\|$) was below 0.01.

Phase error estimation

In simulation, the phase navigator was constructed using the central k-space (32×32) of the imaging data.

If a navigator was available in the data acquired in vivo, a conventional SENSE calculation was applied to reconstruct a navigator image, which was then smoothed by a 2D hamming filter to suppress ringing artefacts and noise. The phase of the navigator image was used as an estimate of the phase errors.

For in vivo data without explicitly acquired navigator, the phase errors were estimated from the imaging data, where a conventional SENSE calculation was applied to reconstruct the imaging data, followed by a total variation [219] based de-noising process. The phase of the de-noised image was used as the phase-error estimate.

Covariance matrix calculation

The symmetric property of the diffusion signal was considered in simulation, where the covariance function was calculated with the angle θ defined as $\theta_{i,j} = \arccos |\langle \mathbf{g_i}, \mathbf{g_j} \rangle|$. For in vivo data, to reduce the effects of eddy currents and Maxwell

terms in DAGER reconstruction the signal symmetry was not incorporated, which led to a definition of θ as $\theta_{i,j} = \arccos\langle \mathbf{g_i}, \mathbf{g_j} \rangle$.

6.3.5 Post-processing

For the simulation data, root-mean-squared-error (RMSE) values were calculated between the reconstructed images and the noise-free ground-truth with a brain mask to emphasize the reconstruction errors within brain regions. The in vivo data were first processed using Topup and Eddy [214] to correct image distortions and then fitted with a tensor model using DTIFIT [195] and multiple fibre populations within a single voxel using BEDPOSTX([204, 205]). Probability based tractography of 27 major white matter tracts was conducted using the AutoPtx plugin available in FSL (https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/AutoPtx) [220], which are a set of processing scripts and pre-defined start/stop masks that allow one to run probabilistic tractography automatically.

6.4 Results

6.4.1 Simulation

Fig. 6.2a shows the reconstruction results of simulated data with in-plane under-sampling factors of 4 and 6 using fixed k-q sampling. The ground truth is shown on the left. Conventional SENSE cannot reconstruct the images correctly for either under-sampling factor. The DAGER reconstruction enables an improved reconstruction for R=4 data, but fails to work for the R=6 data.

The failure of the reconstruction on the R=6 data relates to the use of a fixed k-q sampling pattern. The reconstruction is expected to be improved with the optimized sampling strategy. Fig. 6.2b shows the results, where the reconstruction is substantially improved using the DAGER method, which reduces the RMSE by $\sim$9 times compared to the SENSE reconstruction, and the residual artefacts are mostly noise.

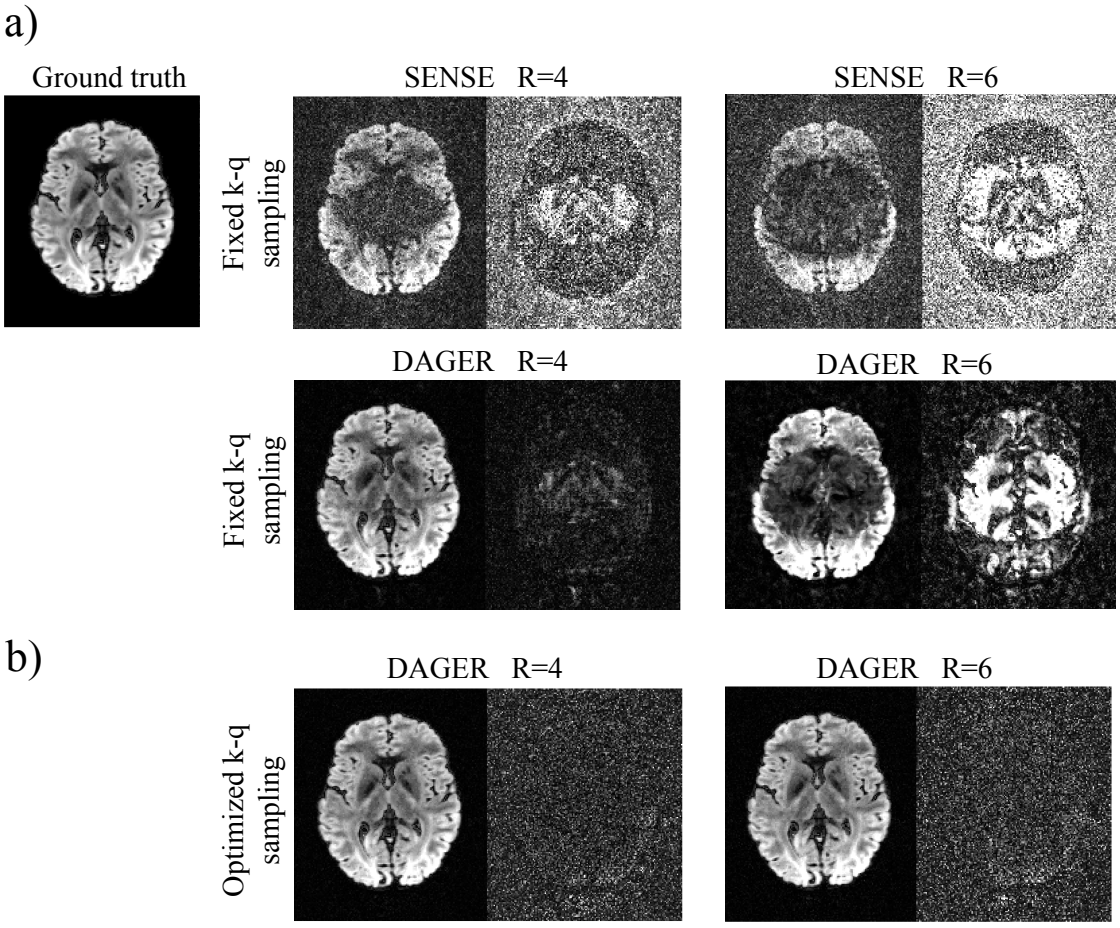


Figure 6.2: Simulation of in-plane under-sampling. Two under-sampling factors R=4 and R=6 are evaluated. The difference images are shown next to the reconstructed images. a) Using a fixed k-q sampling. Difference images are scaled by a factor of 3 for better visualization. RMSE values for SENSE reconstruction with R=4 and R=6 are 0.31 and 0.48, respectively. RMSE values for DAGER reconstruction with R=4 and R=6 are 0.09 and 0.37, respectively. b) Using the optimized k-q sampling, The different images are scaled by a factor of 10. RMSE for DAGER with R=4 and R=6 are 0.04 and 0.05, respectively.

Fig. 6.3 shows the convergence of the DAGER reconstruction with R=6 data, where each point in Fig. 6.3a represents one diffusion volume from one iteration. The RMSEs for all diffusion volumes decrease monotonically as the iteration goes on and reach a minimum value after about 10 iterations. Although RMSE is a useful metric to determine the stopping criteria, ground truth is commonly not available in practice. Alternatively, the update of the solution can be used to monitor the reconstruction progress. As shown in Fig. 6.3b, the normalized solution update

has similar decay behaviour as the RMSE shown in Fig. 6.3a. After 7 iterations, the image update is below 1% and the iteration is stopped.

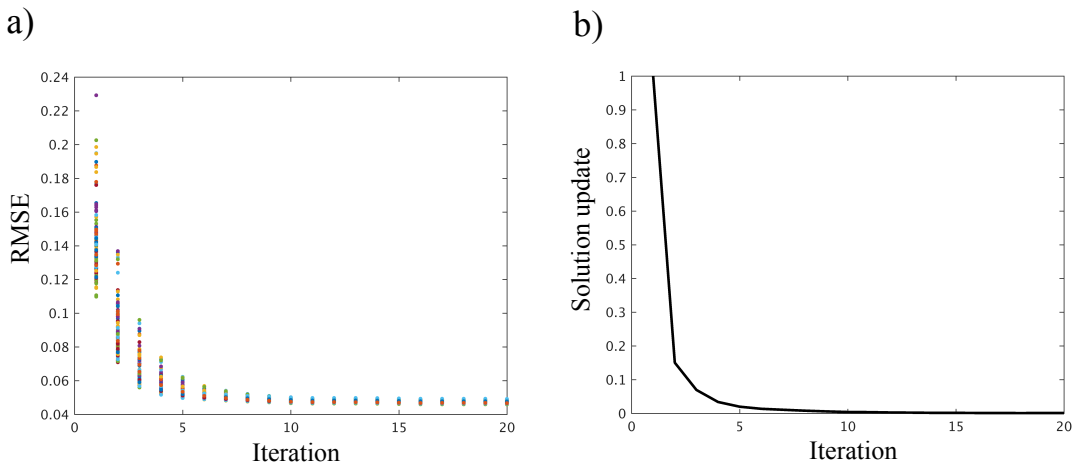


Figure 6.3: Convengece of the DAGER reconstruction. In-plane under-sampling with R=6 is used here. Left: RMSE of all diffusion directions in 20 iterations (each point represents on diffusion direction). Right: solution update $\|\mathbf{u}^n - \mathbf{u}^{n-1}\|/\|\mathbf{u}^{n-1}\|$ in 20 iterations.

To validate the robustness of DAGER reconstruction at different noise level and evaluate the accuracy of noise variance estimation (which is used as a reconstruction parameter in DAGER), experiments are conducted on five data sets added with different levels of noise. The noise-free image is first normalized to have a mean signal of 40, and then five noise variances are tested, corresponding to a SNR of 6,10,20,30,60, respectively. For each data set, reconstructions are applied with 14 testing noise variance values σ^2_{test}, including $[2^{-10}\,2^{-7}\,2^{-4}\,2^{-3}\,2^{-2}\,2^{-1}\,2^{0}\;2^{1}\;2^{2}\;2^{3}\;2^{4}\;2^{7}\;2^{10}] \times \sigma^2_{\text{true}}$ and the noise variance estimated from the first iteration of DAGER, σ^2_{est}. Data were under-sampled by R=6. As shown in the result (Fig. 6.4), optimal reconstructions (minimal RMSE values) at all noise level are achieved using the ground-truth noise variance σ^2_{true} (testing points with $\log_2(\sigma^2_{\text{test}}/\sigma^2_{\text{true}}) = 0$ in each plot). The RMSE increases when the testing noise variance σ^2_{test} deviates from σ^2_{true}. All plots have an area of low RMSE for test noise variances close to the true variance. This area is wider if the true noice variance is low (e.g. Fig. 6.4d,e), suggesting the choice of

noise variance is less critical for high SNR data. However, for low SNR data, small deviation from σ^2_{true} can lead to a sharp increase of reconstruction errors (Fig. 6.4a). Nevertheless, the DAGER method can provide a fairly good estimation (red circles) of the true noise variance ($1/2\sigma^2_{\text{true}} - 8\sigma^2_{\text{true}}$), enabling a robust reconstruction.

The effects of phase error on reconstruction is shown in Fig. 6.5, where a corrupted image is produced by the reconstruction if the phase errors are not corrected (Fig. 6.5a) while a high-quality result can be achieved by incorporating the phase errors correction in the reconstruction(Fig. 6.5b).

Fig. 6.6 shows a comparison between different k-q sampling strategies. When phase errors (Fig. 6.6a) are not considered, all methods except the fixed sampling produce a good reconstruction, suggesting using incoherent k-q sampling can effectively improve reconstruction performance.

When phase errors are added to the data (Fig. 6.6c), the fixed sampling can also achieve good reconstructions, as the random phase errors shift the k-space such that the 'fixed' sampling pattern to some extent samples the k-space incoherently. Nevertheless, the reconstruction with fixed sampling is still inferior to the other non-fixed sampling approaches. Fig. 6.6c and Fig. 6.6d are the zoomed-in plot of Fig. 6.6a and Fig. 6.6b, which suggest the proposed sampling approach still performs slightly better than the other two approaches (less outliers). Therefore, the proposed sampling approach is used in all other simulations and in vivo acquisitions.

Fig. 6.7 shows the reconstruction results with a very high acceleration factor R=10, which has exceeded the theoretical limit of conventional parallel imaging methods. The DAGER method still manages to recover the underlying image, indicating the efficacy of GP prior in dMRI reconstruction.

The results shown above provide a validation of the DAGER reconstruction for in-plane under-sampling. SMS acceleration is essentially the same as in-plane under-sampling in terms of signal model and reconstruction formulation. Next, the

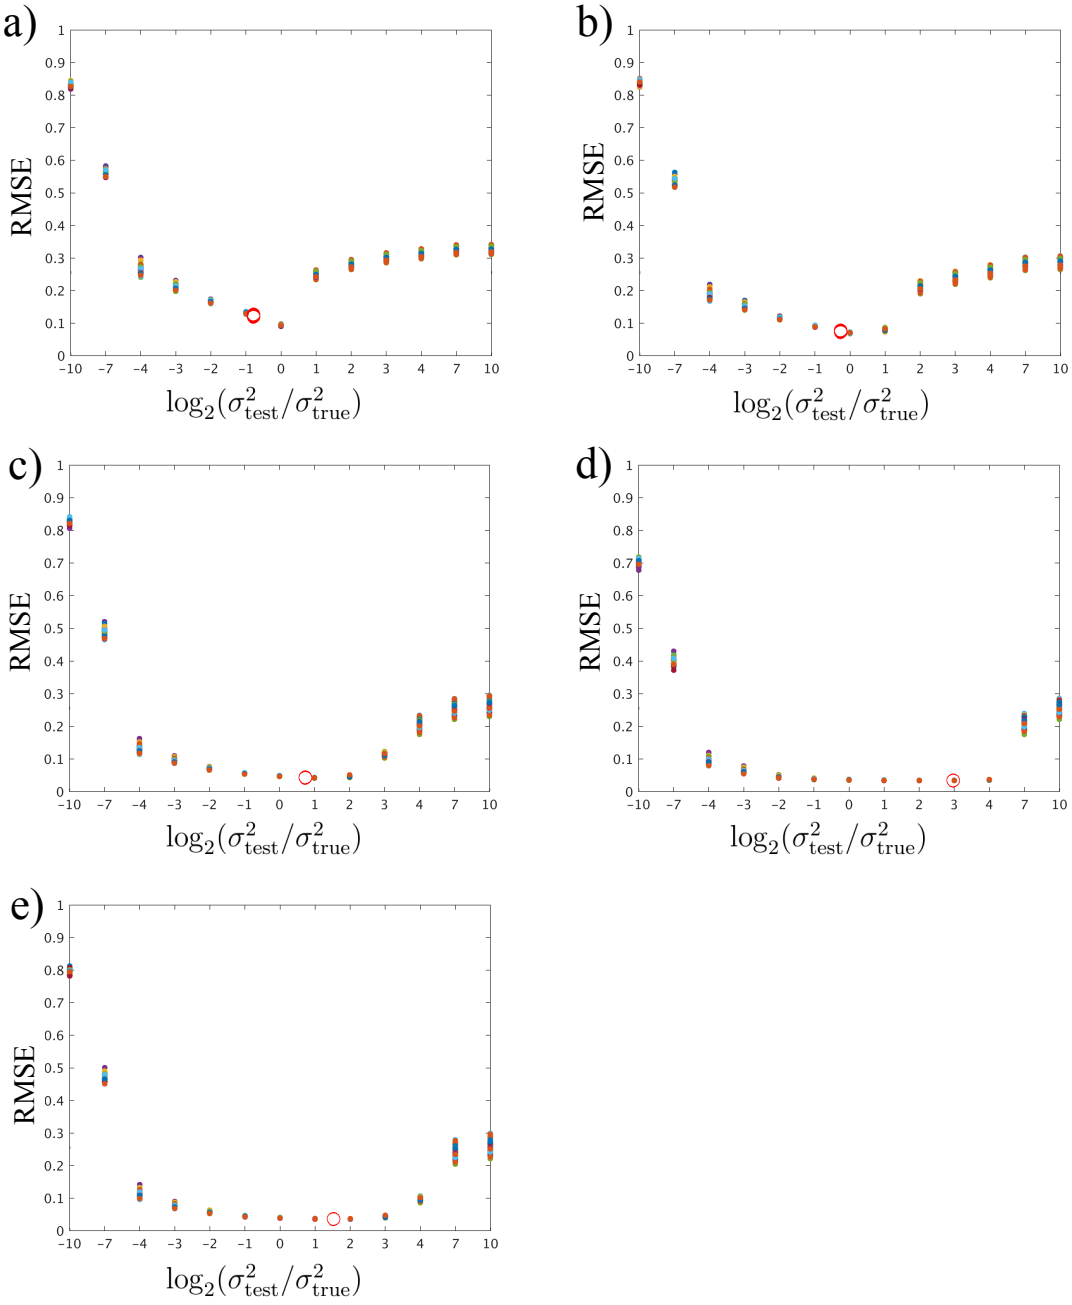


Figure 6.4: DAGER reconstruction at different noise level. a) - e) correspond to a SNR of 6,10,20,30,60, respectively. Each solid point represents the RMSE of DAGER reconstruction for one diffusion direction using a specific testing noise covariance. The red circles (highly overlapped due to consistent reconstruction performance for all directions) represent the RMSEs of DAGER reconstructions using the estimated noise variances, which are close to the true noise variances (testing points corresponding to $\log_2(\sigma^2_{\text{test}}/\sigma^2_{\text{true}}) = 0$ in each plot) and enable a low RMSE in all plots (all noise levels). In-plane acceleration with R=6 is used here.

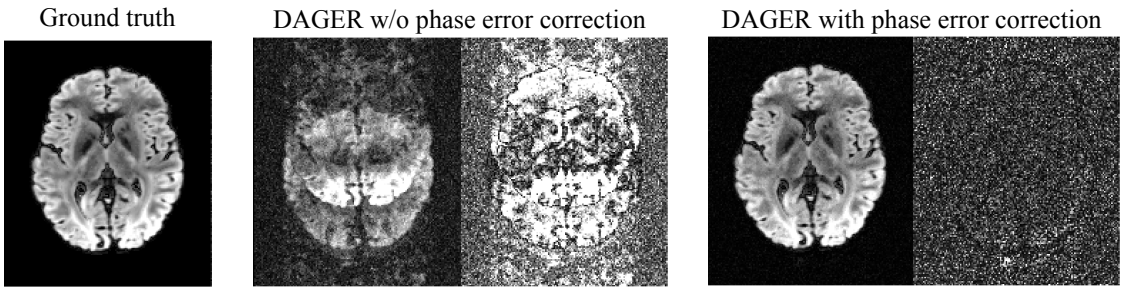


Figure 6.5: Demonstration of phase error correction in DAGER reconstruction. Without phase error correction (middle), DAGER reconstrution results in strong artefacts with a RMSE of 0.58. The different image is scaled by a factor of 3. With phase error correction (right), DAGER reconstruction achieves high-quality results with just noise-like residual artefacts. The RMSE is 0.05 and the different image is scaled by a factor of 10. In-plane acceleration with R=6 is used here.

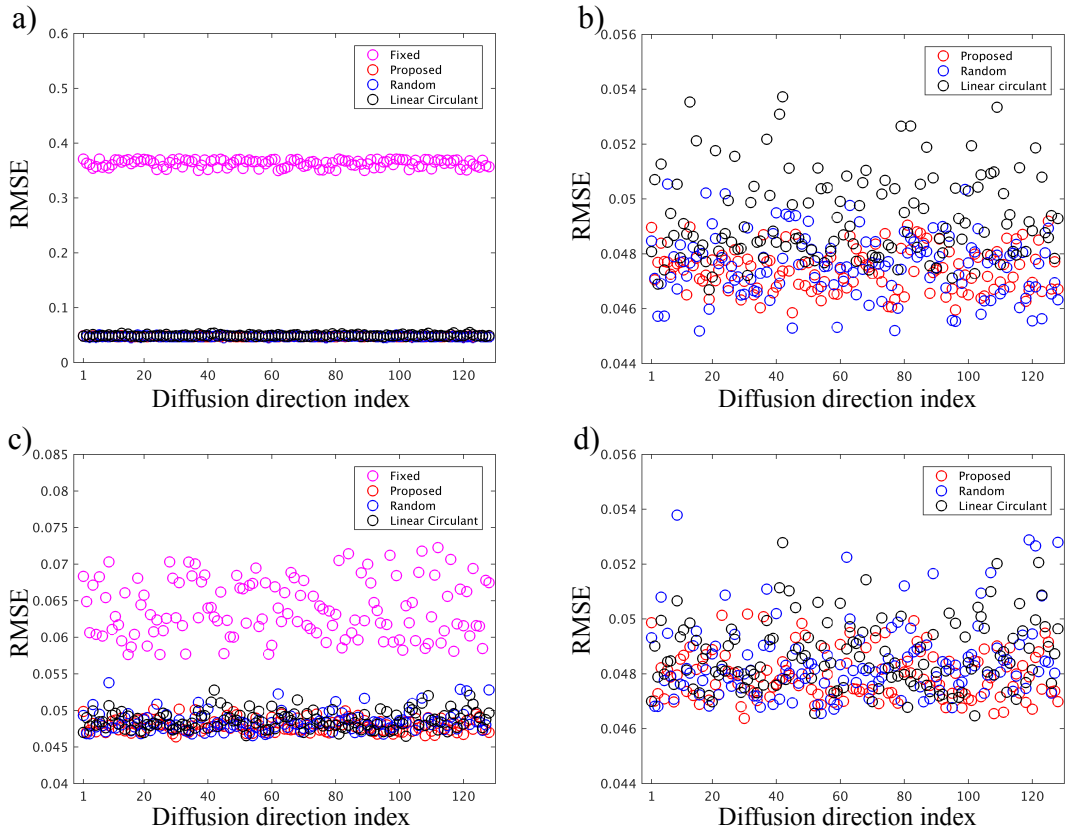


Figure 6.6: Comparison of different k-q sampling for DAGER reconstruction. In-plane acceleration with R=6 is used here. a) DAGER reconstruction of data without phase errors using four different sampling approaches. c) DAGER reconstruction of data with phase errors using four different sampling approaches. Phase error correction is applied. b) and d) are zoomed-in plot of a) and c) to show the slightly different performance between random sampling, linear circulant sampling and the proposed sampling method.

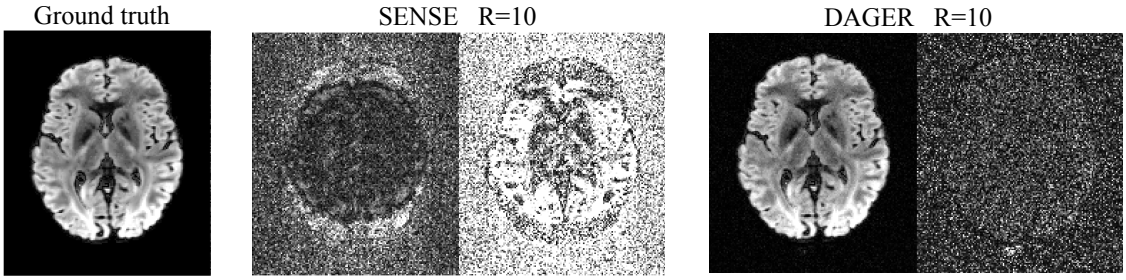


Figure 6.7: Simulation of in-plane under-sampling with a factor of 10, which has exceeded the theoretical limit of conventional parallel imaging, leading to corrupted reconstruction using SENSE (middle). The corresponding RMSE is 0.61 and the difference image is scaled by a factor of 3. Nevertheless, DAGER reconstruction (right) still manages to recover the underlying image with a RMSE of 0.06. The difference image is scaled by a factor of 10.

improvement on the combination of in-plane under-sampling and SMS acceleration using SMS-DAGER (tailored DAGER reconstruction for SMS data) is demonstrated.

A set of SMS data is simulated with MB=4 and in-plane under-sampling R=3. Phase errors are randomly generated for each image slice and each diffusion volume separately. Fig. 6.8 shows the reconstruction results, where SMS-DAGER achieves a considerably improved reconstruction compared with SMS-SENSE (tailored SENSE reconstruction for SMS data [221]). Two phase-error-estimation approaches, navigator-based and navigator-free, are applied, which produce comparable reconstructions. Note the strong artefacts in the navigator-free SMS-DAGER reconstruction (middle) are mostly located in the background area, which stem from the incorrect configuration of phase information in the background. Nevertheless, the errors within the brain region are very small and the RMSE is similar to the navigator-based SMS-DAGER reconstruction. The marginal difference between the navigator-based method and navigator-free method suggests the initial SENSE reconstruction provides a good estimation of the phase errors (for brain region). Therefore, the navigator-free method can be adopted to reduce SAR (by $\sim$50% for monopolar preparation).

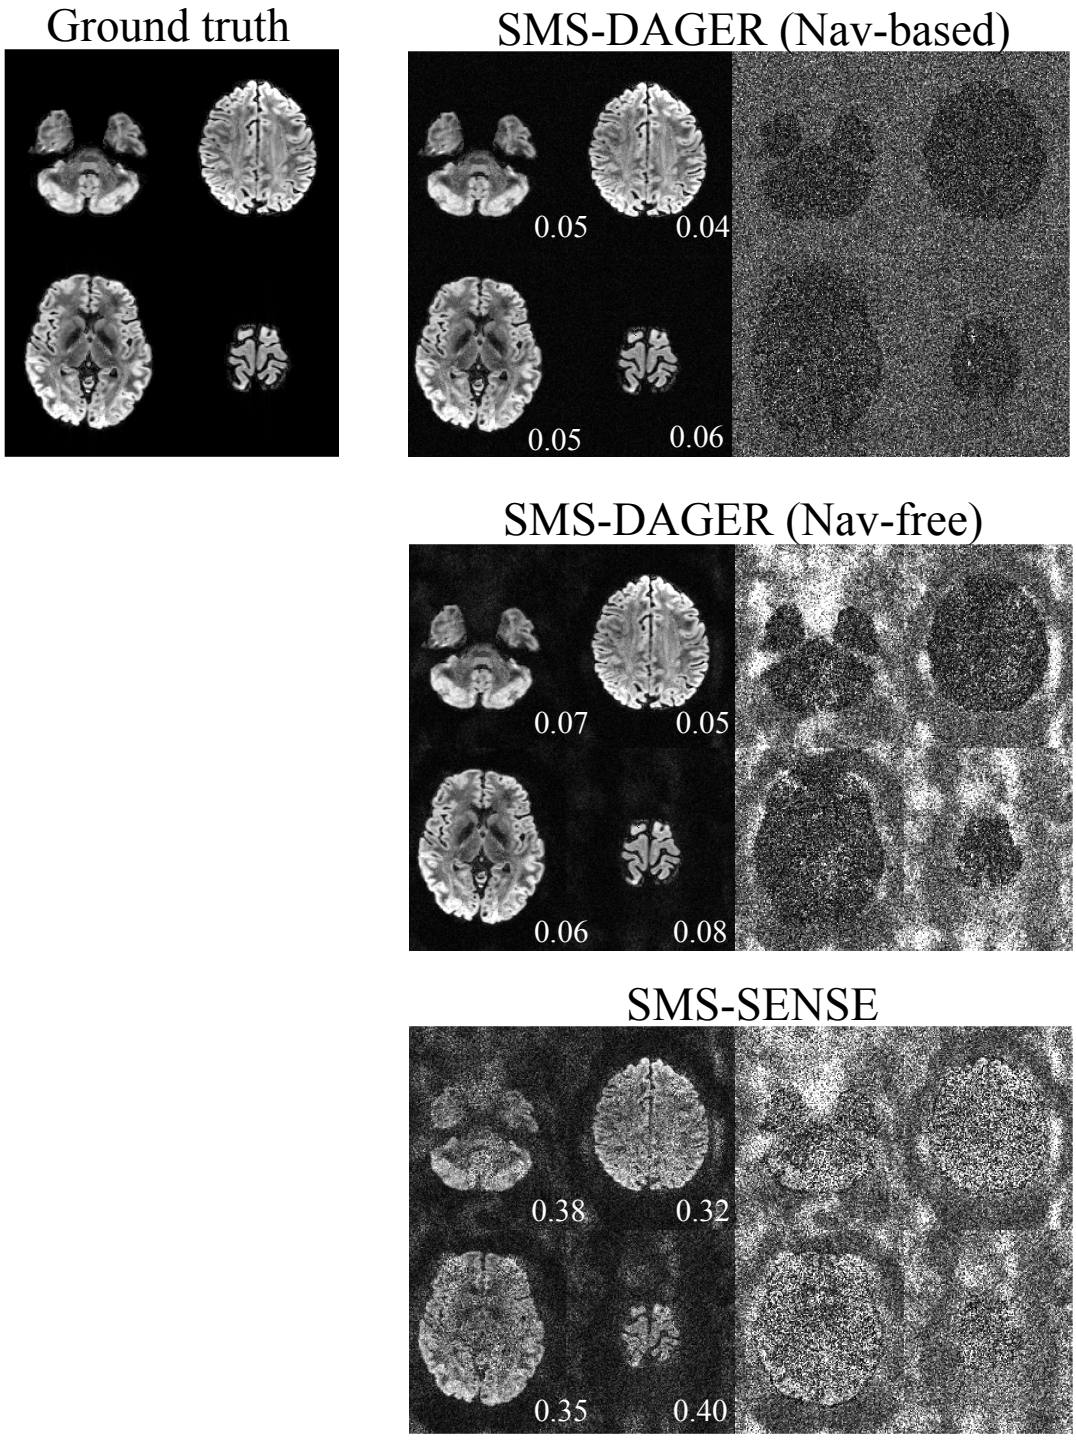


Figure 6.8: Simulation of SMS data with MB=4 and in-plane under-sampling factor of 3. Three reconstructions approaches are evaluated here: SMS-SENSE reconstruction (tailored SENSE reconstruction for SMS data), SMS-DAGER reconstruction with navigator acquisition (Nav-based) and witout navigator acquisition (Nav-free). The RMSE values are shown on the bottom right of the reconstructed images. Note the RMSE is calculated for the brain region only. The difference images are scaled by a factor of 3 for SMS-SENSE and 10 for SMS-DAGER.

6.4.2 In Vivo

In-plane Under-sampling

Fig. 6.9 shows results from subject one. Conventional parallel imaging reconstruction at R=6 suffers from amplified noise and residual aliasing even with a large number of coil elements (32ch). With the Gaussian processes prior, reconstructions are considerably improved. Phase-error correction with and without an explicitly acquired phase navigator demonstrate very similar performance, indicating it is possible to dispose of the navigator acquisition, which can reduce the scan time and SAR burden. Compared with the R=3 image, the R=6 reconstruction with DAGER demonstrates reduced image distortion and improved delineation of structural boundaries.

Propagation of noise and residual artefacts in diffusion quantification are demonstrated in Fig. 6.10. At R=6, conventional parallel imaging fails to capture many secondary fibre populations, most likely due to the low SNR. This information is well preserved in the DAGER reconstruction. The SENSE reconstruction of R=3 data provides a good SNR with negligible residual artefacts, which is used as a reference here.

SMS with in-plane under-sampling

Fig. 6.11 shows the reconstruction of SMS data acquired in vivo with the combination of MB=4 and R=3 (total acceleration factor 12). Conventional SENSE reconstruction suffers from substantially amplified noise due to the high acceleration applied in both through-plane and in-plane. The DAGER method manages to reconstruct a high-quality image with effectively suppressed noise and aliasing artefacts.

Fig. 6.12 shows four diffusion directions of a transverse slice. Images reconstructed by SMS-DAGER are placed alongside the high-SNR single-band reference, which is an average of three repetitions (SB-3ave). The SMS images reconstructed with the DAGER method shows consistent contrast with the single-band reference.

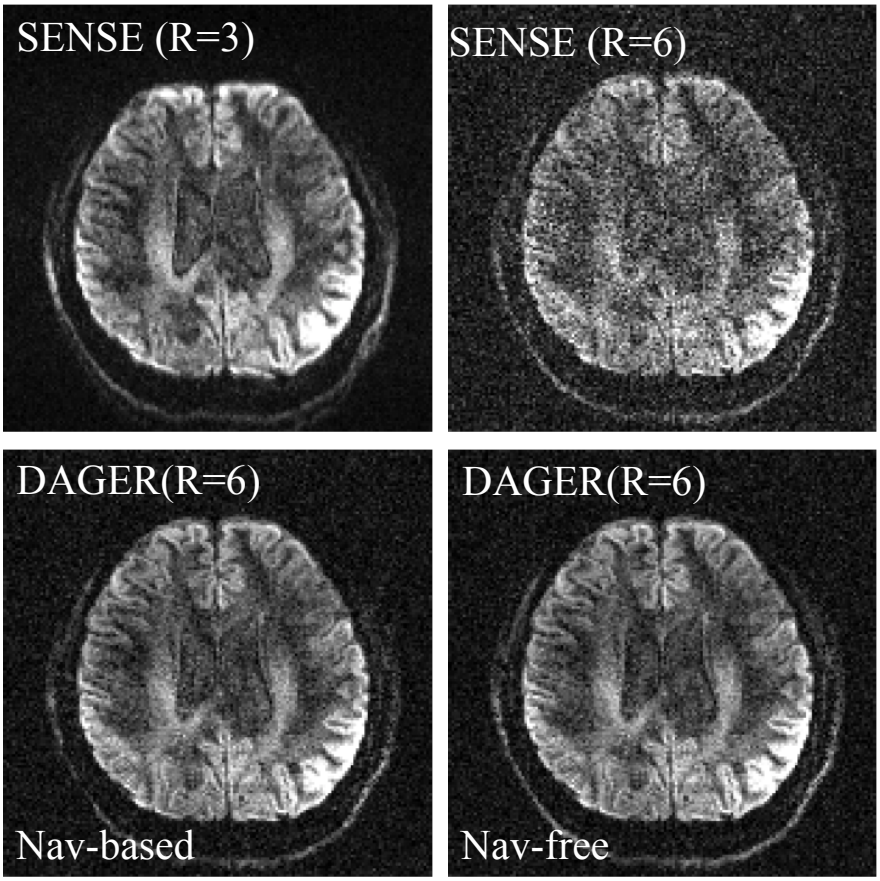


Figure 6.9: Reconstruction of in-plane under-sampled data acquired from subject one. The R=6 data were reconstructed with SENSE and DAGER (with and without navigator acquisition). The R=3 SENSE reconstruction is used as a reference here, although the distortion is stronger than the R=6 data.

In addition, the noise level of the SMS image is similar to that of the single-band reference, suggesting GP prior can effectively eliminate random noise by exploring the common features between diffusion volumes. It should be noted that the single-band images shown here were acquired using ∼33min covering only 21 slices; In comparison, the SMS images were acquired using ∼11min covering the full brain with 84 slices.

The tensor fitting results are shown in Fig. 6.13. Tensor metrics derived from SMS-DAGER reconstruction shows good consistency with the high-SNR single-band reference, while the conventional SENSE results in a poor estimation.

Individual images could still look good even if significant angular blurring is

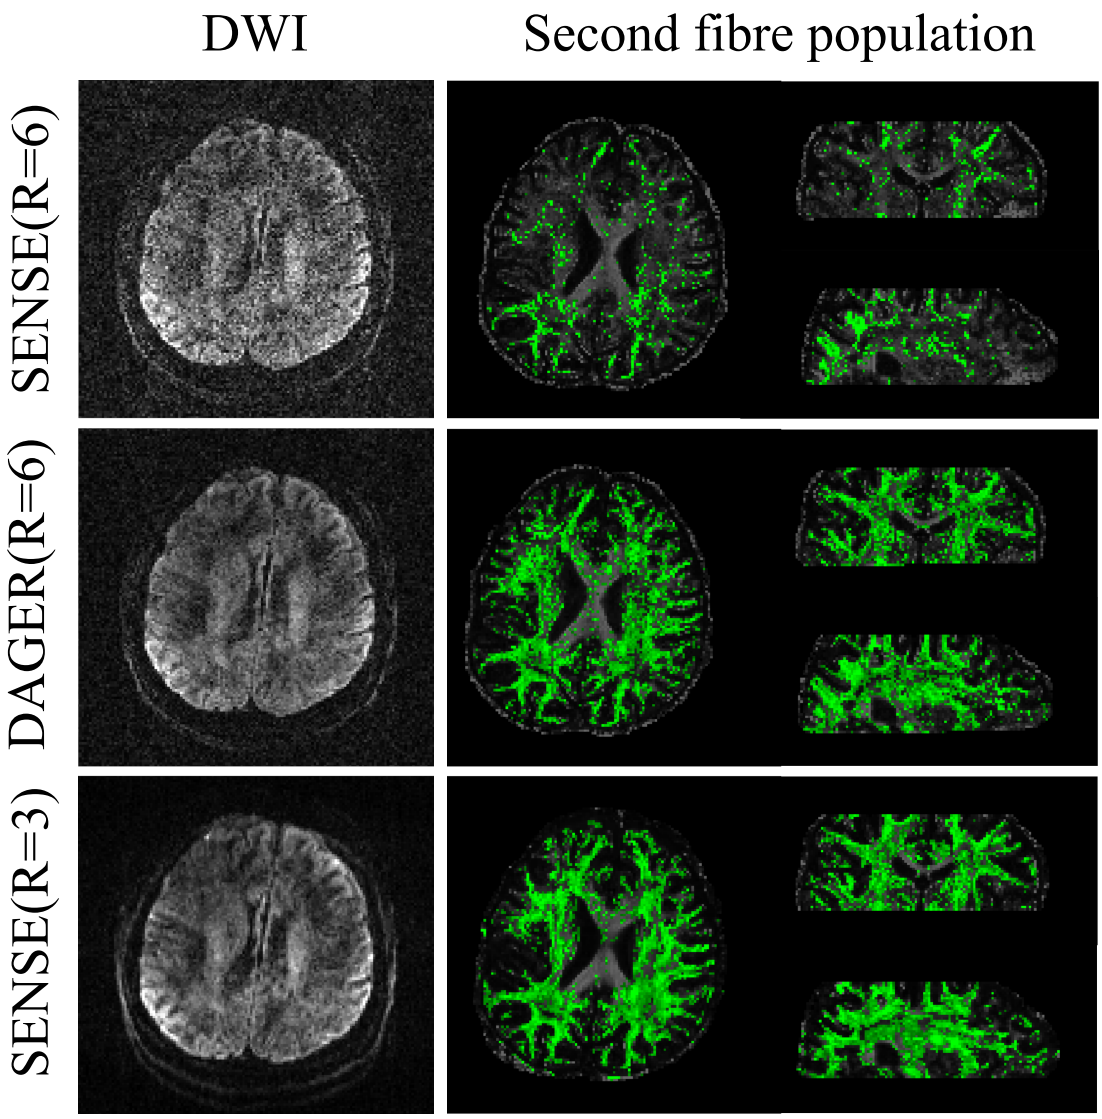


Figure 6.10: In-plane under-sampled data acquired from subject two. Multiple fibre populations within a single voxel are estimated from R=6 data reconstructed with SENSE (top) and DAGER (middle), and R=3 data reconstructed with SENSE (bottom). The R=3 data are used as a reference here.

happening. The DAGER approach aims to identify the right hyper-parameters to represent the characteristics of the signal and noise without introducing blurring in reconstruction. However, it is critical to evaluate whether the reconstruction gets this right in practice. Here, we investigate the effects of SMS-DAGER reconstruction on angular resolution by looking at the covariance of the in vivo dMRI signal between diffusion directions. White matter voxels are used in the calculation within a white

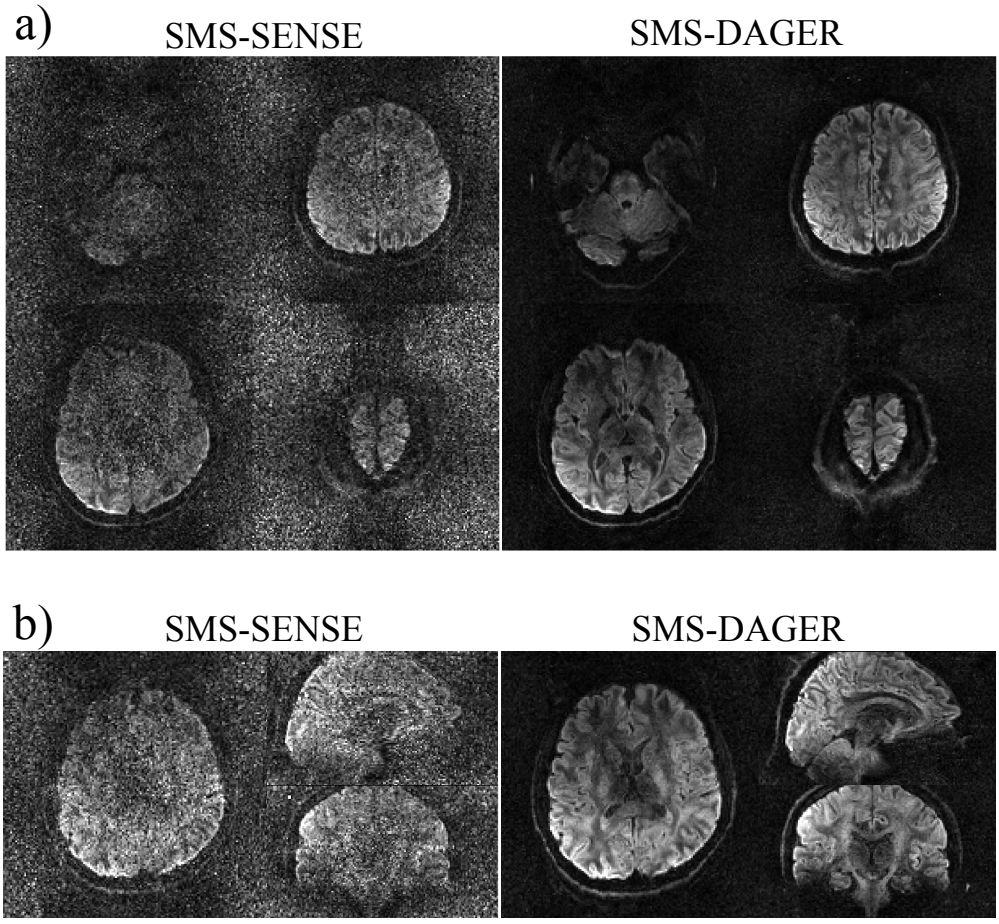


Figure 6.11: Reconstruction of in vivo SMS data. a) shows SMS-SENSE reconstruction and SMS-DAGER reconstruction of four simultaneously acquired slices. b) shows three orthogonal planes of the reconstructed images using the two methods.

matter mask generated from the FA map with a threshold of 0.3. Fig. 6.14 shows the result, where each point represents a diffusion volume. Points with angular distance of 0° are the signal variances. Other points with larger angles depict how rapidly diffusion contrast changes as a function of the distance in q-space. A smaller contrast difference corresponds to a lower covariance, whereas a larger contrast difference corresponds to a higher covariance. Compared to the high-SNR single-band data (Fig. 6.14a), SMS images reconstructed with SMS-DAGER yield a very similar covariance (Fig. 6.14b and Fig. 6.14d), indicating a good preservation of angular resolution in the DAGER reconstruction. With conventional SENSE reconstruction, the image signal is dominated by noise, resulting in a low covariance

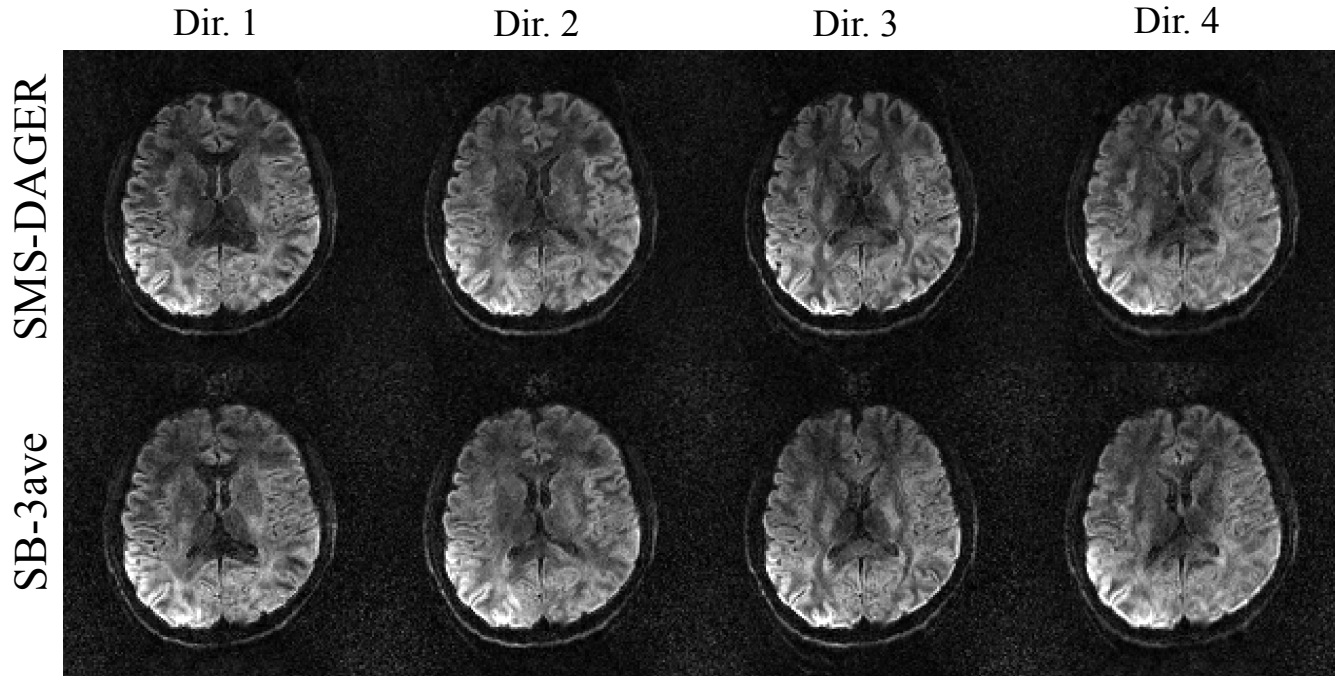


Figure 6.12: Comparison of SMS-DAGER reconstruction (top) and high-SNR single-band reference (bottom), which is generated by averaging three single-band data sets. Four diffusion directions are shown here.

between diffusion directions (Fig. 6.14c).

To further evaluate the angular resolution of the images reconstructed by SMS-DAGER, a direct comparison of the fibre orientations is conducted between the SMS-DAGER images and a reference. Two fibre populations within each voxel are estimated using BEDPOSTX and the mean of the fibre orientation distribution is used for comparison. Two sets of reference data are tested, one is the average of the three repetitions single-band data (SB-3ave), the other is the average of two repetitions single-band data (SB-2ave). The SB-3ave data has a higher SNR than the SB-2ave data. The angular differences are calculated between the SMS-DAGER images and the two references for the two fibre populations separately, leading to four comparison results: fibre 1, SMS-DAGER vs SB-2ave; fibre 1, SMS-DAGER vs SB-3ave; fibre 2, SMS-DAGER vs SB-2ave; fibre 2, SMS-DAGER vs SB-3ave. The same comparisons are also conducted between a set of sing-band data (SB-1ave, using the repetition left out of the SB-2ave images) and the two references.

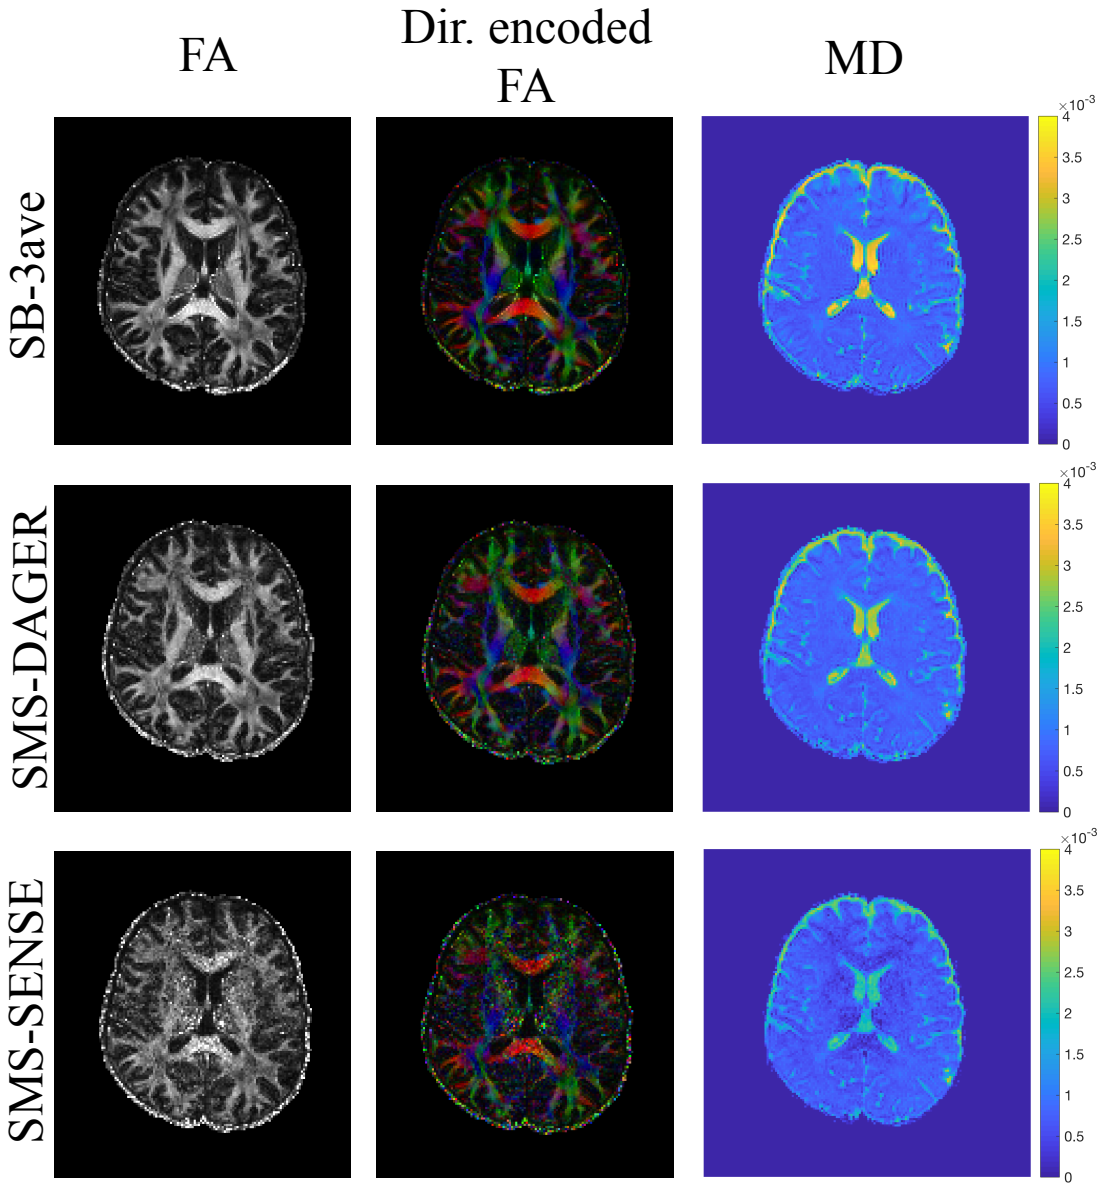


Figure 6.13: Diffusion tensor analysis.

The results of the comparisons are shown in Fig. 6.15. The angular differences derived from the SMS-DAGER images are comparable to those derived from the SB-1ave data with slightly higher angular differences ($< 4°$, measured by the median (same below)). Compared to the SB-2ave reference (Fig. 6.15a), the SB-3ave (Fig. 6.15b) reference results in similar angular differences ($< 0.8°$) for fibre 1 (from either the SB-1ave data or the SMS-DAGER images), but lower angular differences for fibre 2 (by 3.2° for the SB-1ave data and 1.5° for the SMS-DAGER data).

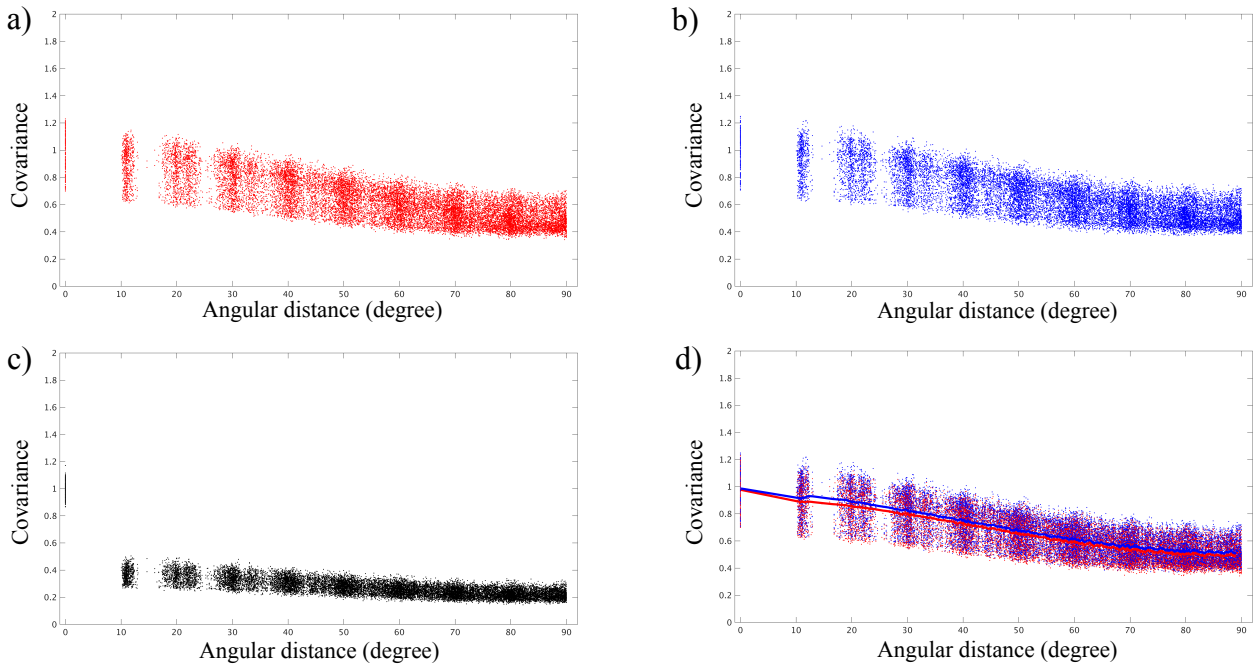


Figure 6.14: Signal covariance of the reconstructed dMRI images. a) High-SNR single-band reference. b) SMS-DAGER reconstruction. c) SMS-SENSE reconstruction. d) Overlay plot of a) and b) with the mean covariance of each plotted as solid line. Each point represent one diffusion direction. The covariance plot is normalized by the median of the variances for all diffusion directions (angular distance=0°).

The reduced angular differences with the SB-3ave reference suggests the noise in the reference data might also contribute to the angular differences. Thus, a further reduction of angular differences might be achievable using a reference with a higher SNR. Note the comparison between the SB-1ave data and the SB-3ave data might be biased due to the inherent correlation between the two data sets (SB-1ave is used in the calculation of SB-3ave).

We conducted automatic mapping of 14 major white matter tracts using AutoPtx [220] on the full-FOV data sets. To match the scan time (11min), SMS-data were acquired with 192 directions and SB data were acquired with 50 directions. The SB data have high SNR due to the long TR required to match the number of slices (12s). As shown in Fig. 6.16, all of the fiber bundles are robustly identified from the SMS-DAGER data. The single-band results capture many fiber bundles, but its performance is compromised by the limited number of diffusion directions

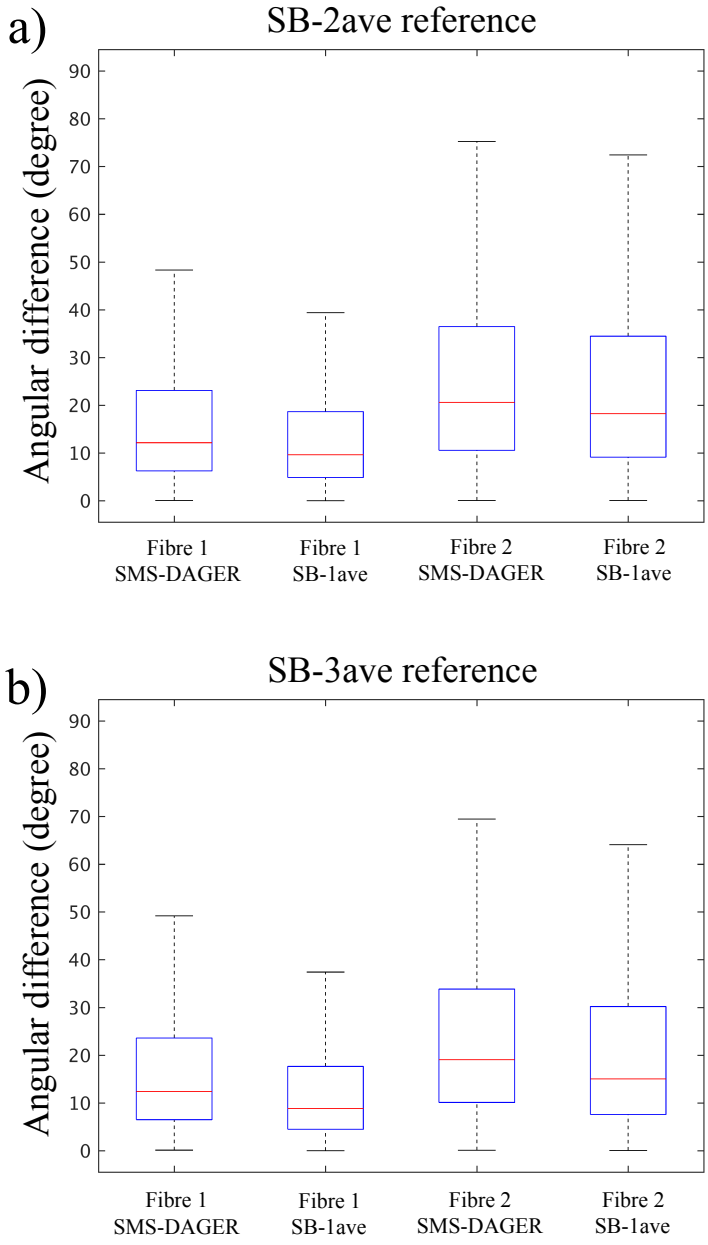


Figure 6.15: Evaluation of fibre orientations estimated from SMS-DAGER and SB-1ave data. Angular differences are calculated based on a) SB-2ave reference and b) SB-3ave reference.

achievable in this scan time. Tracts are in general less abundant in the single-band data than the SMS-DAGER data. In some cases, tracts are totally missing at the threshold used in rendering. For example, the acoustic radiation tracts (Fig. 6.16, middle column, in purple) are clearly captured by the SMS-DAGER data, but essentially absent in the single-band results. The SMS data reconstructed with SENSE performed the worst, failing to capture many fiber tracts, and producing thinner and shorter tracts overall. This reflects the higher level of noise in the SMS-SENSE reconstructions, which lead to greater uncertainty in the fiber estimates and therefore loss of coherence in streamlines. Note that these renderings used the same threshold across the different datasets but required these to be determined manually. However, where tracts are largely absent, there was no threshold at which decent tract reconstructions could be produced, and the overall conclusions about the relative performance of the datasets does not change with threshold level.

6.5 Discussion

Recently, an increasing tendency has arisen in diffusion MRI to acquire more directions, multiple b values and higher spatial resolutions, which are expected to improve the estimation of white matter architecture and thus enable more accurate depiction of structural connectivity in the brain. However, these measurements typically require very long scan time, making it infeasible in practice. Therefore, scan acceleration is becoming more and more important in dMRI. Conventional MRI acceleration using parallel imaging suffers from amplified noise due to the ill-conditioning of the reconstruction problem, which is a great challenge for diffusion MRI due to its intrinsically low SNR. As a result, conventional parallel imaging is typically applied with a small acceleration factor in dMRI acquisitions.

The ill-conditioning of the reconstruction problem can be mitigated by adding some prior information, which can potentially enable a high acceleration factor for

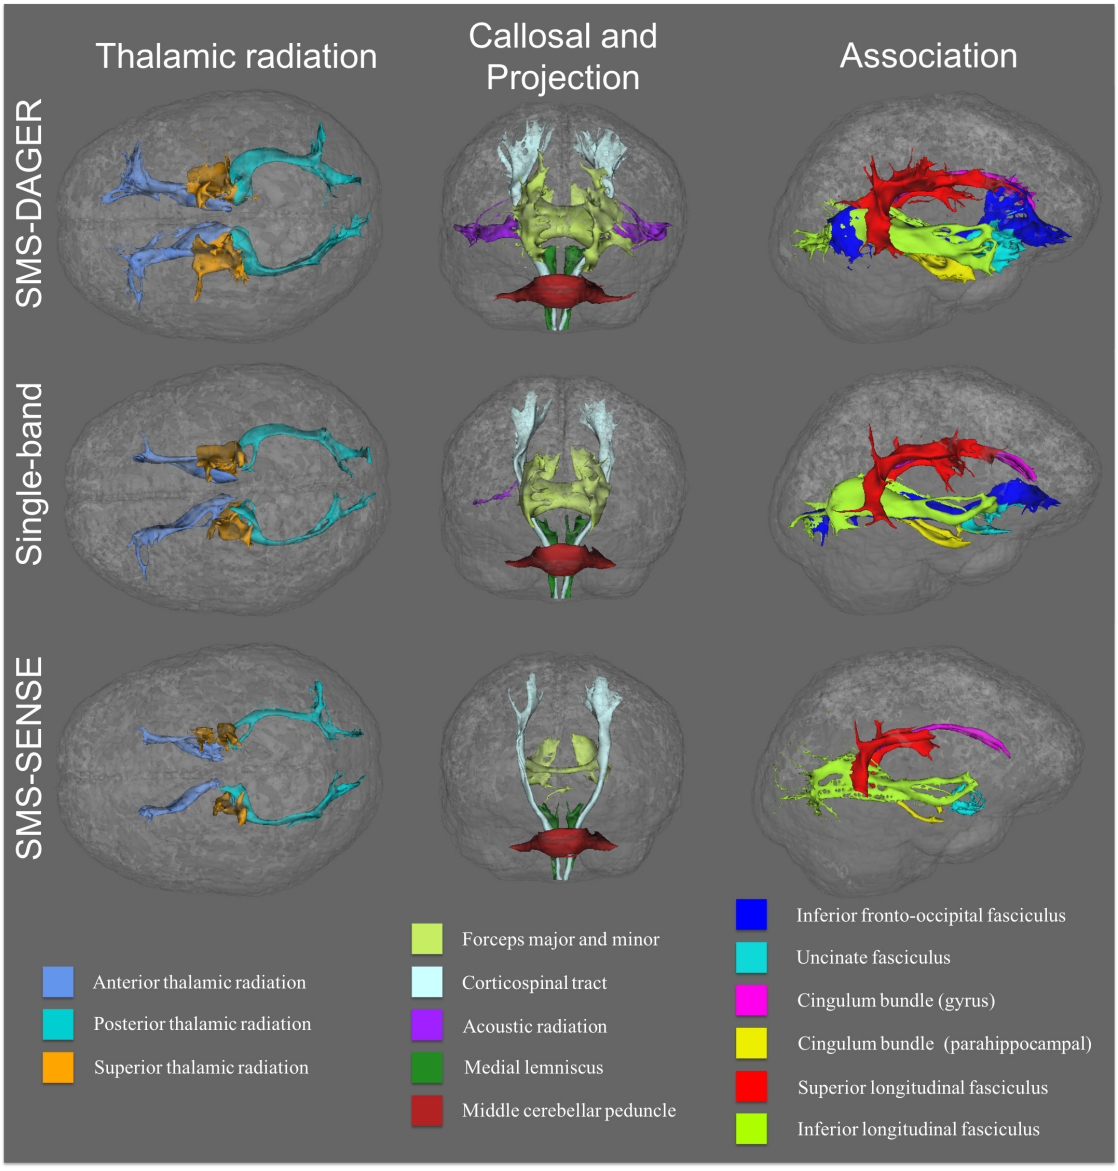


Figure 6.16: White matter fiber tracts generated using AutoPtx. 14 major white matter pathways (constituting 27 separate tracts in right and left hemispheres) are rendered in superior (left, thalamic radiation), anterior (middle, projection and callosal fibers) and lateral (right, association fibers) views. To match the scan time, SMS data and single-band data contain 192 and 50 directions, respectively. SMS data reconstructed with DAGER support the delineation of all tracts, whereas the single-band data and the SMS data reconstructed with SENSE fail to capture some tracts. Most notable are the acoustic radiations (missing from both SB and SENSE reconstructions) and inferior fronto-occipital fasciculus (missing from SENSE reconstruction).

dMRI. In this work, Gaussian processes are applied to model diffusion MRI signals, which have already been used in dMRI data pre-processing. Compared with other modelling methods, the Gaussian process method is a non-parametric approach, which leaves the hyper-parameters related to the signal smoothness in q-space determined by the data itself. In addition, the GP approach provides an elegant way to model multi-shell dMRI data, allowing the estimation of data from one shell using data from different shells. Using a GP prior in dMRI reconstruction can benefit from all these advantages, enabling a reconstruction with high quality and high fidelity. Moreover, the use of GP in both data reconstruction and analysis provides a consistency that can potentially enable an integration of image reconstruction and data pre-processing, which is expected to benefit both procedures and further improve the dMRI data quality.

The hyper-parameters in GP model are estimated from the data based on a Bayesian formulation. To address the detrimental effects of image aliasing on hyper-parameter estimation, an iterative updating scheme is adopted in DAGER, where the hyper-parameter estimation and the image reconstruction are applied back-and-forth, such that errors in both procedures can be gradually suppressed as the iteration goes on. As shown in Fig. 6.4, the calibration of noise variance is critical for reconstruction, especially for low SNR data, which are very common in diffusion MRI. In DAGER, the signal noise variance is automatically estimated from the data itself, which is fairly accurate as demonstrated in the simulation and allows a high-quality reconstruction for both simulated data and in vivo data. The automatic determination of reconstruction parameters is very helpful for practical applications as no manual adjustment is needed.

As DAGER pulls information from adjacent q-space locations, it could easily introduce angular smoothing if the hyper-parameters overestimate how smooth the actual signal is. Therefore, we conduct several experiments to investigate the effects of DAGER reconstruction on angular resolution. First, the signal covariance is

calculated, which shows the DAGER reconstruction enables a consistent covariance to that of the high-SNR reference data. Second, a direct comparison of fibre orientations is conducted, which shows the angular differences introduced by DAGER will not be considerably larger than those derived from the conventional single-band data. Finally, tractography of major white matter tracts are compared, where all data sets are acquired using the same scan time. Neither single-band data nor SENSE-reconstructed SMS data support a good estimation of fibre tracts, whereas the SMS data reconstructed by DAGER support an abundance of fibre tracts. These evaluations and investigations suggest DAGER reconstruction can provide a good preservation of angular resolution.

Apart from the reconstruction, the SMS acquisition of dMRI data itself faces significant challenges due to the SAR limit, particularly at ultra-high field with a high MB factor. To address this issue, several compatible optimizations have been implemented in the dMRI sequence, including the modified diffusion preparation, the improved RF pulse and the elimination of navigator acquisition, which allow a balance between SAR, distortion and reconstruction performance. It should be noted that the navigator-free method used in DAGER is valid only if a good estimate of image phase could be obtained from the initial reconstruction using conventional parallel imaging methods. Otherwise, a phase navigator is still needed.

Joint k-q reconstructions have been proposed before. In comparison, the DAGER approach proposed in this work uses a non-parametric model rather than a parametric model (e.g. tensor-based models used in [147] and [145]), allows automatic determination of reconstruction parameters rather than manually adjustment, and for the first time demonstrates dMRI acceleration with SMS acquisition, which is the most effective way to accelerate dMRI in practice.

Although the current work only demonstrates accelerated 2D and 2D-SMS dMRI using DAGER, other acquisition methods like 3D and 3D multi-slab are also compatible provided an incoherent k-q sampling scheme is incorporated. Recently,

several studies have demonstrated the use of 'simultaneous multi-slab' acquisition for dMRI, which enables optimal SNR efficiency with shorter scan time [121, 184]. It is expected that the combination of DAGER and simultaneous multi-slab acquisition can achieve even higher scan efficiency with improved image quality, which will be investigated in the future.

The DAGER framework can be readily extended to include image-space constraints (e.g. total variation) to further improve the reconstruction, which will also be studied in future work.

6.6 Conclusion

In this Chapter, a new acquisition and reconstruction approach for faster diffusion MRI is proposed. Both simulation and data in vivo demonstrate superior performance of the proposed method compared with conventional parallel imaging. More validations will be conducted in the next step. Possible extensions of the proposed methods (e.g. combination of DAGER and 3D multi-slab acquisition, incorporation of image-space priors) will be investigated in the future.

7 Summary and Future Work

7.1 Thesis Summary

Hybrid 2D/3D acquisition methods have been previously proposed to address the limitations of conventional 2D and 3D acquisitions in dMRI. 3D multi-slab acquisition can achieve thin slice thickness and enable optimal SNR efficiency, which is ideal for high-resolution dMRI. However, a prerequisite for using 3D multi-slab for dMRI acquisition in practice is the robust correction of slab-boundary artefacts, which would otherwise not only deteriorate the image quality but also propagate into diffusion quantification. Several approaches have been proposed to correct the slab-boundary artefacts, among which the PEN method shows effective reduction of slab aliasing without significantly increasing scan time. However, residual artefacts might remain in PEN results due to slab crosstalk, particularly at short TRs which we aim to achieve for maximising SNR efficiency. The NPEN method proposed in this thesis can correct both slab aliasing and the signal variation induced by slab crosstalk even at short TR ($\sim$2s). Simulation and results from data acquired in vivo demonstrate the superior performance of NPEN correction compared with

PEN and the direct slab combination (weighted average).

Simultaneous multi-slice acquisition is the other hybrid 2D/3D method that has been used to improve diffusion MRI. The key benefit conferred by SMS is reduced scan time without significant SNR loss (using a blipped-CAIPI scheme). However, when in-plane under-sampling is applied the g-factor noise amplification in SMS reconstruction becomes severe, even with blipped-CAIPI acquisition. This challenge stems from the ill-conditioning of the parallel imaging reconstruction, which can potentially be improved with some prior information. In the DAGER method proposed in Chapter 6 correlations between q-space samples are explored with a Gaussian process model, which is used as a prior to improve the reconstruction of the underlying image. Numerical simulation and data acquired in vivo demonstrate the good performance of the DAGER reconstruction, particularly for the data highly accelerated both in-plane and through-plane. Apart from providing a new reconstruction algorithm, DAGER also offers an approach for designing locally incoherent k-q sampling, which can potentially benefit any other k-q acceleration methods.

The ultra-high field of 7T enables a higher signal than standard field strength (3T), making it possible to improve the resolution of dMRI without compromising the SNR. In Chapter 5 and 6 we leverage the compatibility of 7T and hybrid 2D/3D acquisition to achieve high resolution, high SNR diffusion MRI. The NPEN method is modified to address the stronger $\mathbf{B}_0$ inhomogeneity at 7T. Several improvements and modifications have been implemented in the SMS dMRI sequence to enable a high multi-band factor (MB=4) and a relatively short TR ($\sim$3.5s) without exceeding the SAR limit. To improve the robustness of auto-calibration data, the recently proposed FLEET approach is adopted in hybrid 2D/3D dMRI, which enables superior reconstruction performance compared with conventional auto-calibration-signal acquisition methods.

In summary, this thesis presents the acquisition and reconstruction methods we developed for hybrid 2D/3D dMRI, which are expected to improve the data quality and scan efficiency of diffusion MRI.

7.2 Future Work

Sub-millimetre diffusion MRI using simultaneous multi-slab acquisition and DAGER reconstruction

Isotropic voxels at sub-millimetre resolution are only achievable with 3D acquisitions, as 2D approaches are limited by the selectivity of RF pulses, which is worse for thin slice thickness. To achieve optimal SNR efficiency, 3D multi-slab acquisition would be the ideal method, particularly at such high resolution. Full brain coverage with sub-millimetre resolution necessitates a large number of slices, which can be either acquired using a small number of relatively 'thick' slabs (with a short TR and assuming the slab thickness is still compatible with a 2D navigator correction) with many kz segments in each slab, or acquiring many thin slabs (with a long TR) with a small number of kz segments in each slab. The scan time will be very long for both methods, particularly when a dense q-space sampling is also required.

The first approach, despite being compatible with the optimal SNR efficiency, will be more sensitive to subject motion as more segments are included in each slab. More importantly, this approach faces limitations in scan acceleration. As the optimal TR has been achieved and so should not be reduced, the only way to accelerate the scan is to apply intra-slab under-sampling by skipping some kz segments. However, this might result in an ill-conditioned reconstruction as the variations of coil sensitivities are quite small along the slab direction.

The second approach faces the limitation of long TR, which could be alleviated using simultaneous multi-slab acquisition, an adaption of the multi-band method

in 3D multi-slab acquisition which has been previously implemented to reduce the TR and scan time of dMRI [121, 184]. The signal model and reconstruction formation are the same for simultaneous multi-slab and simultaneous multi-slice acquisitions, thus the scan efficiency demonstrated by the latter approach should also be achievable by the former approach. However, simultaneous multi-slab acquisition will also face g-factor challenges if in-plane under-sampling is jointly applied. This could be alleviated using the DAGER reconstruction presented in Chapter 6. Hence, the combination of DAGER and simultaneous multi-slab acquisition is expected to achieve sub-millimetre dMRI data within a reasonable scan time and without significant g-factor noise amplification.

Integrated image reconstruction and distortion/motion correction based on Gaussian processes

The Gaussian processes used in DAGER reconstruction are also applied in the 'Eddy' framework to correct eddy-current-induced distortion and subject motion. DAGER reconstruction and Eddy correction are mutually beneficial: reduced distortion and motion (from Eddy) enables more accurate characterization of dMRI signal with Gaussian processes, which can be used to further improve DAGER reconstruction; Eddy correction can also benefit from DAGER, particularly for highly accelerated data, as the results of conventional parallel imaging usually suffer from amplified noise and strong residual aliasing that might decrease the efficiency of Eddy correction. Therefore, it is expected that the integration of these two methods can further improve dMRI data quality. This will be studied in the future.

DAGER with image space priors

Currently, DAGER only considers q-space priors and makes no assumption on the data correlation in image space, which could be utilized to further improve the reconstruction. Total variation (TV)-based approaches have been widely used for de-noising and improving image reconstruction (e.g. compressed sensing methods discussed in Chapter 3), which uses a prior assumption that images contain some areas that are piece-wise constant. Incorporation of TV constraints into the DAGER reconstruction should enable a higher acceleration factor with better image SNR.

Appendices

A

Quantitative Description of EPI Distortion

This Appendix is a recap of the theory section in [222].

With field inhomogeneity, the k space signal becomes

$$S(m\Delta k_x, n\Delta k_y) = \iint \rho(x,y) e^{-i2\pi(m\Delta k_x x + n\Delta k_y y) + i\gamma\Delta B_0(nt_{es} + (-1)^n m\Delta t) + i\Phi} \mathrm{d}x\mathrm{d}y \quad \text{(A.1)}$$

where S is the MRI signal, Δk_x and Δk_y are k space increments along readout and phase encoding directions respectively, m and n are the sampling points in k space ($-N_y/2 \leq n \leq N_y/2 - 1$ and $-N_x/2 \leq m \leq N_x/2 - 1$). ΔB_0 is the spatial varying field inhomogeneity. t_{es} is the echo spacing (duration between two consecutive echoes). Δt is the sampling interval in data readout. Φ represents a constant phase term which does not vary during data readout. The $(-1)^n$ term stems from the zig-zag sampling order of EPI. Inverse Fourier transform of the k-space data gives the image

$$\rho(x', y') = \rho(x - \frac{\gamma\Delta B_0(x,y)(-1)^n\Delta t}{2\pi\Delta k_x}, y - \frac{\gamma\Delta B_0(x,y)t_{es}}{2\pi\Delta k_y}) e^{i\Phi} \quad \text{(A.2)}$$

Let $\Delta k_x = \gamma/2\pi G_x \Delta t$ and $\Delta k_y = \gamma/2\pi \overline{G_y} t_{PE}$, where $\overline{G_y}$ is the mean gradient amplitude during phase encoding and t_{PE} is the phase encoding interval (blip duration). Eq. (A.2) can be written as

$$\rho(x', y') = \rho(x - \frac{\Delta B_0(x,y)(-1)^n}{G_x}, y - \frac{\Delta B_0(x,y) t_{es}}{\overline{G_y} t_{PE}}) e^{i\Phi} \tag{A.3}$$

where the magnitude image is distorted. The EPI distortion along readout direction is negligible as EPI readout gradient (G_x) is typically very strong. Along the phase encoding direction, the distortion is determined by the field inhomogeneity (ΔB_0) and the k-space transverse speed along ky ($\overline{G_y} t_{PE}/t_{es}$). Therefore, a shorter echo spacing is desired, which enables reduced EPI distortion.

B

Jacobian Matrix of the Nonlinear Operator in NPEN

This appendix presents the Jacobian matrix of the nonlinear operator E and its adjoint.

$$E(\mathbf{x}) = \begin{bmatrix} \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_1 \cdot \mathbf{u}\} \\ \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_2 \cdot \mathbf{u}\} \\ \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_3 \cdot \mathbf{u}\} \\ \vdots \\ \mathbf{MF}\{\mathbf{c} \cdot \mathbf{s}_{N_{slab}} \cdot \mathbf{u}\} \end{bmatrix} \tag{B.1}$$

For simplicity of notation, a single receive coil is assumed here. Let $\mathbf{u}$ be the $N \times 1$ image vector, the Fourier transform operator $\mathbf{F}$ can be represented as a $N \times N$ matrix:

$$\mathbf{F} = \begin{bmatrix} f_{1,1} & f_{1,2} & f_{1,3} & \cdots & f_{1,N} \\ f_{2,1} & f_{2,2} & f_{2,3} & \cdots & f_{2,N} \\ f_{3,1} & f_{3,2} & f_{3,3} & \cdots & f_{3,N} \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ f_{N,1} & f_{N,2} & f_{N,3} & \cdots & f_{N,N} \end{bmatrix} \tag{B.2}$$

First, just look at the encoding matrix with the m^{th} slab profile s_m, which is

$$\mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_m\cdot\mathbf{u}\}=\mathbf{M}\begin{bmatrix}\sum_{k=1}^{N} f_{1,k}\,c_k\,s_{m,k}\,u_k\\ \sum_{k=1}^{N} f_{2,k}\,c_k\,s_{m,k}\,u_k\\ \sum_{k=1}^{N} f_{3,k}\,c_k\,s_{m,k}\,u_k\\ \vdots\\ \sum_{k=1}^{N} f_{N,k}\,c_k\,s_{m,k}\,u_k\end{bmatrix} \tag{B.3}$$

where $s_{m,k}$ is the k^{th} element of $\mathbf{s}_m$. The derivative of this matrix with respect to s_m is

$$\frac{\partial\mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_m\cdot\mathbf{u}\}}{\partial\mathbf{s}_m}=\mathbf{M}\begin{bmatrix}f_{1,1}\,c_1\,u_1 & f_{1,2}\,c_2\,u_2 & f_{1,3}\,c_3\,u_3 & \cdots & f_{1,N}\,c_N\,u_N\\ f_{2,1}\,c_1\,u_1 & f_{2,2}\,c_2\,u_2 & f_{2,3}\,c_3\,u_3 & \cdots & f_{2,N}\,c_N\,u_N\\ f_{3,1}\,c_1\,u_1 & f_{3,2}\,c_2\,u_2 & f_{3,3}\,c_3\,u_3 & \cdots & f_{3,N}\,c_N\,u_N\\ \vdots & \vdots & \vdots & \cdots & \vdots\\ f_{N,1}\,c_1\,u_1 & f_{N,2}\,c_2\,u_2 & f_{N,3}\,c_3\,u_3 & \cdots & f_{N,N}\,c_N\,u_N\end{bmatrix} \tag{B.4}$$

which is the same for any $\mathbf{s}_m$. And the derivative with respect to $\mathbf{s}_n(n\neq m)$ is zero:

$$\frac{\partial\mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_m\cdot\mathbf{u}\}}{\partial\mathbf{s}_n}=\mathbf{0} \tag{B.5}$$

The derivative with the image vector $\mathbf{u}$ can also be calculated as

$$\frac{\partial\mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_m\cdot\mathbf{u}\}}{\partial\mathbf{u}}=\mathbf{M}\begin{bmatrix}f_{1,1}\,c_1\,s_{m,1} & f_{1,2}\,c_2\,s_{m,2} & f_{1,3}\,c_3\,s_{m,3} & \cdots & f_{1,N}\,c_N\,s_{m,N}\\ f_{2,1}\,c_1\,s_{m,1} & f_{2,2}\,c_2\,s_{m,2} & f_{2,3}\,c_3\,s_{m,3} & \cdots & f_{2,N}\,c_N\,s_{m,N}\\ f_{3,1}\,c_1\,s_{m,1} & f_{3,2}\,c_2\,s_{m,2} & f_{3,3}\,c_3\,s_{m,3} & \cdots & f_{3,N}\,c_N\,s_{m,N}\\ \vdots & \vdots & \vdots & \cdots & \vdots\\ f_{N,1}\,c_1\,s_{m,1} & f_{N,2}\,c_2\,s_{m,2} & f_{N,3}\,c_3\,s_{m,3} & \cdots & f_{N,N}\,c_N\,s_{m,N}\end{bmatrix} \tag{B.6}$$

For simplicity of notation, define $\mathbf{D}_s=\dfrac{\partial\mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_m\cdot\mathbf{u}\}}{\partial\mathbf{s}_m}$ and $\mathbf{D}_u(m)=\dfrac{\partial\mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_m\cdot\mathbf{u}\}}{\partial\mathbf{u}}$, then the Jacobian of E can be expressed as

$$E'(\mathbf{x})=\begin{bmatrix}\mathbf{D}_u(1) & \mathbf{D}_s & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0}\\ \mathbf{D}_u(2) & \mathbf{0} & \mathbf{D}_s & \mathbf{0} & \cdots & \mathbf{0}\\ \mathbf{D}_u(3) & \mathbf{0} & \mathbf{0} & \mathbf{D}_s & \cdots & \mathbf{0}\\ \vdots & \vdots & \vdots & \cdots & \vdots & \\ \mathbf{D}_u(N_{slab}) & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{D}_s\end{bmatrix} \tag{B.7}$$

and the adjoint $E'(\mathbf{x}_n)^{\mathrm{H}}$ is given by

$$E'(\mathbf{x})^{\mathrm{H}} = \begin{bmatrix} \mathbf{D}_u(1)^{\mathrm{H}} & \mathbf{D}_u(2)^{\mathrm{H}} & \mathbf{D}_u(3)^{\mathrm{H}} & \cdots & \mathbf{D}_u(N_{slab})^{\mathrm{H}} \\ \mathbf{D}_s^{\mathrm{H}} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_s^{\mathrm{H}} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{D}_s^{\mathrm{H}} \end{bmatrix} \tag{B.8}$$

NPEN reconstruction can be solved using conjugate gradient algorithm, which involves the calculation of $E'(\mathbf{x})\Delta\mathbf{x}$ and $E'(\mathbf{x})^{\mathrm{H}}\mathbf{d}$:

$$E'(\mathbf{x})\Delta\mathbf{x} = E'(\mathbf{x}) \begin{bmatrix} \mathrm{d}\mathbf{u} \\ \mathrm{d}\mathbf{s}_1 \\ \mathrm{d}\mathbf{s}_2 \\ \vdots \\ \mathrm{d}\mathbf{s}_{N_{slab}} \end{bmatrix} = \begin{bmatrix} \mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_1\cdot\mathrm{d}\mathbf{u} + \mathbf{c}\cdot\mathrm{d}\mathbf{s}_1\cdot\mathbf{u}\} \\ \mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_2\cdot\mathrm{d}\mathbf{u} + \mathbf{c}\cdot\mathrm{d}\mathbf{s}_2\cdot\mathbf{u}\} \\ \mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_3\cdot\mathrm{d}\mathbf{u} + \mathbf{c}\cdot\mathrm{d}\mathbf{s}_3\cdot\mathbf{u}\} \\ \vdots \\ \mathbf{MF}\{\mathbf{c}\cdot\mathbf{s}_{N_{slab}}\cdot\mathrm{d}\mathbf{u} + \mathbf{c}\cdot\mathrm{d}\mathbf{s}_{N_{slab}}\cdot\mathbf{u}\} \end{bmatrix} \tag{B.9}$$

and the adjoint is given by

$$E'(\mathbf{x})^{\mathrm{H}}\mathbf{d} = E'(\mathbf{x})^{\mathrm{H}} \begin{bmatrix} \mathbf{d}_1 \\ \mathbf{d}_2 \\ \mathbf{d}_3 \\ \vdots \\ \mathbf{d}_{N_{slab}} \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^{N_{slab}} \mathbf{s}_i^* \cdot \mathbf{c}^* \cdot \mathbf{F}^{-1}\mathbf{M}^{\mathrm{H}}\mathbf{d}_i \\ \mathbf{u}^* \cdot \mathbf{c}^* \cdot \mathbf{F}^{-1}\mathbf{M}^{\mathrm{H}}\mathbf{d}_1 \\ \mathbf{u}^* \cdot \mathbf{c}^* \cdot \mathbf{F}^{-1}\mathbf{M}^{\mathrm{H}}\mathbf{d}_2 \\ \vdots \\ \mathbf{u}^* \cdot \mathbf{c}^* \cdot \mathbf{F}^{-1}\mathbf{M}^{\mathrm{H}}\mathbf{d}_{N_{slab}} \end{bmatrix} \tag{B.10}$$

where $*$ is the complex conjugate.

C

Posterior Probability Distribution with Gaussian Process Priors

The likelihood function (Eq. 6.4) can be expressed as

$$\begin{aligned} p(\mathbf{y}|\mathbf{f},\mathbf{X}) &= \prod_{i=1}^{N} p(y(\mathbf{x}_i)|f(\mathbf{x}_i)) = \prod_{i=1}^{N} \frac{1}{(2\pi\sigma_n^2)^{1/2}} \exp\left(-\frac{(y(\mathbf{x}_i)-f(\mathbf{x}_i))^2}{2\sigma_n^2}\right) \\ &= \frac{1}{(2\pi\sigma_n^2)^{N/2}} \exp\left(-\frac{(\mathbf{y}-\mathbf{f})^{\mathrm{T}}(\mathbf{y}-\mathbf{f})}{2\sigma_n^2}\right) \end{aligned} \tag{C.1}$$

A GP prior can be expressed as

$$p(\mathbf{f}|\mathbf{X}) = \frac{1}{(2\pi)^{N/2}} \frac{1}{|\mathbf{\Sigma}|^{1/2}} \exp\left(-\frac{1}{2}\mathbf{f}^{\mathrm{T}}\mathbf{\Sigma}^{-1}\mathbf{f}\right) \tag{C.2}$$

where a zero mean is used for simplicity of notation. $\mathbf{\Sigma}$ is the covariance matrix.

The posterior is given by

$$p(\mathbf{f}\,|\,\mathbf{y},\mathbf{X}) = \frac{p(\mathbf{y}|\mathbf{f},\mathbf{X})p(\mathbf{f}|\mathbf{X})}{p(\mathbf{y}|\mathbf{X})} \tag{C.3}$$

where the denominator is the marginal likelihood ('model evidence'), which is independent of the function values. The posterior can be expressed with only the

terms dependent on the function values:

$$\begin{aligned} p(\mathbf{f}\,|\,\mathbf{y},\mathbf{X}) &\propto p(\mathbf{y}|\mathbf{f},\mathbf{X})p(\mathbf{f}|\mathbf{X}) \\ &\propto \exp\left(-\frac{(\mathbf{y}-\mathbf{f})^{\mathrm{T}}(\mathbf{y}-\mathbf{f})}{2\sigma_n^2}\right)\exp\left(-\frac{1}{2}\mathbf{f}^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}\mathbf{f}\right) \\ &\propto \exp\left(-\frac{1}{2}(\mathbf{f}-\bar{\mathbf{f}})^{\mathrm{T}}(\sigma_n^{-2}+\boldsymbol{\Sigma}^{-1})(\mathbf{f}-\bar{\mathbf{f}})\right) \end{aligned} \tag{C.4}$$

The posterior is also a Gaussian distribution, where $\bar{\mathbf{f}} = \sigma_n^{-2}(\sigma_n^{-2}+\boldsymbol{\Sigma}^{-1})^{-1}\mathbf{y}$ is the posterior mean, which is a linear combination of observations. The variance of the posterior is $(\sigma_n^{-2}+\boldsymbol{\Sigma}^{-1})^{-1}$.

References

[1] Kawin Setsompop et al. “Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty”. In: *Magn Reson Med* 67 (2012), pp. 1210–1224.

[2] K. Setsompop et al. “Improving diffusion MRI using simultaneous multi-slice echo planar imaging”. In: *NeuroImage* 63 (2012), pp. 569–580.

[3] Kamil Uğurbil et al. “Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project”. In: *Neuroimage* 80 (2013), pp. 80–104.

[4] Stamatios N Sotiropoulos et al. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project”. In: *Neuroimage* 80 (2013), pp. 125–143.

[5] A. T. Vu et al. “High resolution whole brain diffusion imaging at 7T for the Human Connectome Project”. In: *NeuroImage* 122 (2015), pp. 318–331.

[6] Wenchuan Wu and Karla L Miller. “Image formation in diffusion MRI: A review of recent technical developments”. In: *Journal of Magnetic Resonance Imaging* (2017).

[7] W. Wu et al. “Reducing slab boundary artifacts in three-dimensional multislab diffusion MRI using nonlinear inversion for slab profile encoding (NPEN)”. In: *Magn Reson Med* 76 (2016).

[8] W. Wu et al. “High-resolution diffusion MRI at 7T using a three-dimensional multi-slab acquisition”. In: *NeuroImage* 143 (2016).

[9] E Mark Haacke et al. *Magnetic resonance imaging: physical principles and sequence design*. Vol. 82. Wiley-Liss New York: 1999.

[10] Dwight G Nishimura. *Principles of magnetic resonance imaging*. Standford Univ., 2010.

[11] Matt A Bernstein, Kevin F King, and Xiaohong Joe Zhou. *Handbook of MRI pulse sequences*. Elsevier, 2004.

[12] Zhi-Pei Liang and Paul C Lauterbur. *Principles of magnetic resonance imaging: a signal processing perspective*. SPIE Optical Engineering Press, 2000.

[13] Felix Bloch. “Nuclear induction”. In: *Physical review* 70.7-8 (1946), p. 460.

[14] Edward M Purcell, H Co Torrey, and Robert V Pound. “Resonance absorption by nuclear magnetic moments in a solid”. In: *Physical review* 69.1-2 (1946), p. 37.

[15] Lars G Hanson. “Is quantum mechanics necessary for understanding magnetic resonance?” In: *Concepts in Magnetic Resonance Part A* 32.5 (2008), pp. 329–340.

[16] Janaka P Wansapura et al. “NMR relaxation times in the human brain at 3.0 tesla”. In: *Journal of magnetic resonance imaging* 9.4 (1999), pp. 531–538.

[17] Tie-Qiang Li et al. “Characterization of T2* heterogeneity in human brain white matter”. In: *Magnetic resonance in medicine* 62.6 (2009), pp. 1652–1657.

[18] William A Edelstein et al. “Spin warp NMR imaging and applications to human whole-body imaging”. In: *Physics in medicine and biology* 25.4 (1980), p. 751.

[19] Peter Mansfield. “Multi-planar image formation using NMR spin echoes”. In: *Journal of Physics C: Solid State Physics* 10.3 (1977), p. L55.

[20] PC Lauterbur. “Image formation by induced local interactions: examples employing nuclear magnetic resonance”. In: *Nature* 242 (1973), pp. 190–191.

[21] CB Ahn, JH Kim, and ZH Cho. “High-speed spiral-scan echo planar NMR imaging-I”. In: *IEEE transactions on medical imaging* 5.1 (1986), pp. 2–7.

[22] Craig H Meyer et al. “Fast spiral coronary artery imaging”. In: *Magnetic resonance in medicine* 28.2 (1992), pp. 202–213.

[23] Adolf Fick. “Ueber diffusion”. In: *Annalen der Physik* 170.1 (1855), pp. 59–86.

[24] Albert Einstein. “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen”. In: *Annals of physics* 322.8 (1905), pp. 549–560.

[25] Erwin L Hahn. “Spin echoes”. In: *Physical review* 80.4 (1950), p. 580.

[26] Herman Y Carr and Edward M Purcell. “Effects of diffusion on free precession in nuclear magnetic resonance experiments”. In: *Physical review* 94.3 (1954), p. 630.

[27] Edward O Stejskal and John E Tanner. “Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient”. In: *The journal of chemical physics* 42.1 (1965), pp. 288–292.

[28] Peter J Basser, James Mattiello, and Denis LeBihan. “Estimation of the effective self-diffusion tensor from the NMR spin echo”. In: *Journal of Magnetic Resonance, Series B* 103.3 (1994), pp. 247–254.

[29] Peter J Basser, James Mattiello, and Denis LeBihan. “MR diffusion tensor spectroscopy and imaging”. In: *Biophysical journal* 66.1 (1994), pp. 259–267.

[30] K. P. Pruessmann et al. “SENSE: sensitivity encoding for fast MRI”. In: *Magn Reson Med* 42.5 (1999), pp. 952–62.

[31] M. A. Griswold et al. “Generalized autocalibrating partially parallel acquisitions (GRAPPA)”. In: *Magn Reson Med* 47.6 (2002), pp. 1202–10.

[32] Felix A Breuer et al. “General formulation for quantitative G-factor calculation in GRAPPA reconstructions”. In: *Magnetic Resonance in Medicine* 62.3 (2009), pp. 739–746.

[33] Robert Turner, Denis Le Bihan, and A Scott Chesnicks. “Echo-planar imaging of diffusion and perfusion”. In: *Magnetic resonance in medicine* 19.2 (1991), pp. 247–253.

[34] J. A. McNab et al. “Surface based analysis of diffusion orientation for identifying architectonic domains in the in vivo human cortex”. In: *Neuroimage* 69 (2013), pp. 87–100.

[35] M. Kleinnijenhuis et al. “Diffusion tensor characteristics of gyrencephaly using high resolution diffusion MRI in vivo at 7T”. In: *Neuroimage* 109 (2015), pp. 378–87.

[36] K. L. Miller et al. “Diffusion imaging of whole, post-mortem human brains on a clinical MRI scanner”. In: *Neuroimage* 57.1 (2011), pp. 167–81.

[37] Stamatios N Sotiropoulos et al. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project.” In: *NeuroImage* 80 (2013), pp. 125–43.

[38] TG Reese et al. “Reduction of eddy current induced distortion in diffusion MRI using a twice refocused spin echo”. In: *Magn Reson Med* 49.1 (2003), pp. 177–182.

[39] P. Jezzard, A. S. Barnett, and C. Pierpaoli. “Characterization of and correction for eddy current artifacts in echo planar diffusion imaging”. In: *Magn Reson Med* 39.5 (1998), pp. 801–812.

[40] J. L. R. Andersson and S. N. Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *Neuroimage* 125 (2016), pp. 1063–1078.

[41] Matthew D. Robson and David A. Porter. “Reconstruction as a source of artifact in nongated single-shot diffusion-weighted EPI”. In: *Magnetic Resonance Imaging* 23.9 (2005), pp. 899–905.

[42] A. W. Anderson and J. C. Gore. “Analysis and correction of motion artifacts in diffusion weighted imaging”. In: *Magn Reson Med* 32.3 (1994), pp. 379–87.

[43] P. Storey et al. “Partial k-space reconstruction in single-shot diffusion-weighted echo-planar imaging”. In: *Magn Reson Med* 57.3 (2007), pp. 614–9.

[44] Hing Chiu Chang, Shayan Guhaniyogi, and Nan Kuei Chen. “Interleaved diffusion-weighted improved by adaptive partial-Fourier and multiband multiplexed sensitivity-encoding reconstruction”. In: *Magn Reson Med* 73 (2015), pp. 1872–1884.

[45] Stuart Clare. “Iterative Nyquist ghost correction for single and multi-shot EPI using an entropy measure”. In: *Proc 11th Annual Meeting ISMRM, Toronto*. 2003, p. 1041.

[46] E. M. Delfaut et al. “Fat suppression in MR imaging: techniques and pitfalls”. In: *Radiographics* 19.2 (1999), pp. 373–82.

[47] Roger J Ordidge et al. “Correction of motional artifacts in diffusion-weighted MR images using navigator echoes”. In: *Magnetic resonance imaging* 12.3 (1994), pp. 455–460.

[48] Adam W Anderson and John C Gore. “Analysis and correction of motion artifacts in diffusion weighted imaging”. In: *Magnetic resonance in medicine* 32.3 (1994), pp. 379–387.

[49] Matthew D. Robson, Adam W. Anderson, and John C. Gore. “Diffusion-weighted multiple shot echo planar imaging of humans without navigation”. In: *Magn Reson Med* 38 (1997), pp. 82–88.

[50] Robert Frost et al. “Scan time reduction for readout-segmented EPI using simultaneous multislice acceleration: Diffusion-weighted imaging at 3 and 7 Tesla”. In: *Magn Reson Med* 74 (2015), pp. 136–149.

[51] D Porter and E Mueller. “Multi-shot diffusion-weighted EPI with readout mosaic segmentation and 2D navigator correction”. In: *Proc 12th Annual Meeting of ISMRM, Kyoto*. 2004, p. 442.

[52] Samantha J. Holdsworth et al. “Readout-segmented EPI for rapid high resolution diffusion imaging at 3T”. In: *European Journal of Radiology* 65 (2008), pp. 36–46.

[53] David A. Porter and Robin M. Heidemann. “High resolution diffusion-weighted imaging using readout-segmented echo-planar imaging, parallel imaging and a two-dimensional navigator-based reacquisition”. In: *Magn Reson Med* 62 (2009), pp. 468–475.

[54] Samantha J Holdsworth et al. “Robust GRAPPA-accelerated diffusion-weighted readout-segmented (RS)-EPI.” In: *Magn Reson Med* 62 (2009), pp. 1629–40.

[55] Robin M. Heidemann et al. “Diffusion imaging in humans at 7T using readout-segmented EPI and GRAPPA”. In: *Magn Reson Med* 64 (2010), pp. 9–14.

[56] Chunlei Liu, Michael E Moseley, and Roland Bammer. “Simultaneous phase correction and SENSE reconstruction for navigated multi-shot DWI with non-cartesian k-space sampling.” In: *Magn Reson Med* 54 (2005), pp. 1412–22.

[57] James G. Pipe, Victoria G. Farthing, and Kirsten P. Forbes. “Multishot diffusion-weighted FSE using PROPELLER MRI”. In: *Magn Reson Med* 47 (2002), pp. 42–52.

[58] Robert Frost et al. “Implementation and assessment of diffusion-weighted partial Fourier readout-segmented echo-planar imaging”. In: *Magn Reson Med* 68 (2012), pp. 441–451.

[59] Robert Frost et al. “3D multi-slab diffusion-weighted readout-segmented EPI with real-time cardiac-reordered k-space acquisition”. In: *Magn Reson Med* 72 (2014), pp. 1565–1579.

[60] Wolfgang Bogner et al. “Readout-segmented echo-planar imaging improves the diagnostic performance of diffusion-weighted MR breast examinations at 3.0 T”. In: *Radiology* 263.1 (2012), pp. 64–76.

[61] Wolfgang Bogner et al. “Bilateral diffusion-weighted MR imaging of breast tumors with submillimeter resolution using readout-segmented echo-planar imaging at 7 T.” In: *Radiology* 274 (2015), pp. 74–84.

[62] Dorota J. Wisner et al. “High-resolution diffusion-weighted imaging for the separation of benign from malignant BI-RADS 4/5 lesions found on breast MRI at 3T”. In: *J Magn Reson Imaging* 40 (2014), pp. 674–681.

[63] Hirokazu Tokoro et al. “Usefulness of free-breathing readout-segmented echo-planar imaging (RESOLVE) for detection of malignant liver tumors: Comparison with single-shot echo-planar imaging (SS-EPI)”. In: *European Journal of Radiology* 83 (2014), pp. 1728–1733.

[64] Warren D. Foltz et al. “Readout-segmented echo-planar diffusion-weighted imaging improves geometric performance for image-guided radiation therapy of pelvic tumors”. In: *Radiotherapy and Oncology* 117 (2014), pp. 525–531.

[65] Yee Liang Thian et al. “Readout-segmented Echo-planar imaging for diffusion-weighted imaging in the pelvis at 3T-A feasibility study”. In: *Academic Radiology* 21 (2014), pp. 531–537.

[66] Iris Friedli et al. “Improvement of renal diffusion-weighted magnetic resonance imaging with readout-segmented echo-planar imaging at 3T.” In: *Magnetic resonance imaging* 33 (2015), pp. 701–8.

[67] Kristen W. Yeom et al. “Comparison of readout-segmented echo-planar imaging (EPI) and single-shot epi in clinical application of diffusion-weighted imaging of the pediatric brain”. In: *American Journal of Roentgenology* 200 (2013), pp. 437–443.

[68] K. Butts et al. “Diffusion-weighted interleaved echo-planar imaging with a pair of orthogonal navigator echoes”. In: *Magn Reson Med* 35.5 (1996), pp. 763–770.

[69] David Atkinson et al. “Nonlinear phase correction of navigated multi coil diffusion images”. In: *Magn Reson Med* 56.5 (2006), pp. 1135–1139.

[70] Nan kuei Chen et al. “A robust multi-shot scan strategy for high-resolution diffusion weighted MRI enabled by multiplexed sensitivity-encoding (MUSE)”. In: *NeuroImage* 72 (2013), pp. 41–47.

[71] Shayan Guhaniyogi et al. “Motion immune diffusion imaging using augmented MUSE for high-resolution multi-shot EPI”. In: *Magn Reson Med* 652 (2015), pp. 639–652.

[72] Michael Herbst et al. “Prospective motion correction of segmented diffusion weighted EPI”. In: *Magn Reson Med* 74 (2015), pp. 1675–1681.

[73] Zhe Zhang et al. “Self-feeding MUSE: A robust method for high resolution diffusion imaging using interleaved EPI”. In: *NeuroImage* 105 (2015), pp. 552–560.

[74] Hing-Chiu Chang et al. “Human brain diffusion tensor imaging at submillimeter isotropic resolution on a 3 Tesla clinical MRI scanner”. In: *NeuroImage* 118 (2015), pp. 667–675.

[75] Merry Mani et al. “Multishot sensitivity encoded diffusion data recovery using structured low rank matrix completion (MUSSELS)”. In: *Magn Reson Med* (2016. doi:10.1002/mrm.26382).

[76] J. Hennig, A. Nauerth, and H. Friedburg. “RARE imaging: a fast imaging method for clinical MR”. In: *Magn Reson Med* 3.6 (1986), pp. 823–33.

[77] J. G. Pipe. “Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging”. In: *Magn Reson Med* 42.5 (1999), pp. 963–9.

[78] S. Meiboom and D. Gill. “Modified Spin-Echo Method for Measuring Nuclear Relaxation Times”. In: *Review of Scientific Instruments* 29.8 (1958), pp. 688–691.

[79] M. H. Levitt and R. Freeman. “Compensation for Pulse Imperfections in Nmr Spin-Echo Experiments”. In: *Journal of Magnetic Resonance* 43.1 (1981), pp. 65–80.

[80] P. Le Roux. “Non-CPMG fast spin echo with full signal”. In: *Journal of Magnetic Resonance* 155.2 (2002), pp. 278–292.

[81] J. G. Pipe and N. Zwart. “Turboprop: Improved PROPELLER imaging”. In: *Magn Reson Med* 55.2 (2006), pp. 380–385.

[82] K. Oshio and D. A. Feinberg. “GRASE (Gradient- and spin-echo) imaging: a novel fast MRI technique”. In: *Magn Reson Med* 20.2 (1991), pp. 344–9.

[83] Fu-Nien Wang et al. “PROPELLER EPI: An MRI technique suitable for diffusion tensor imaging at high field strength with reduced geometric distortions”. In: *Magn Reson Med* 54.5 (2005), pp. 1232–1240.

[84] Tzu-Chao Chuang et al. “PROPELLER-EPI with parallel imaging using a circularly symmetric phased-array RF coil at 3.0 T: Application to high-resolution diffusion tensor imaging”. In: *Magn Reson Med* 56.6 (2006), pp. 1352–1358.

[85] Stefan Skare et al. “Propeller EPI in the other direction”. In: *Magn Reson Med* 55.6 (2006), pp. 1298–1307.

[86] C. H. Meyer et al. “Fast Spiral Coronary-Artery Imaging”. In: *Magn Reson Med* 28.2 (1992), pp. 202–213.

[87] GH Glover. “Basic and advanced concepts of spiral imaging”. In: *International Society for Magnetic Resonance in Medicine Fast MRI Workshop, Asilomar, CA*. 1997, pp. 115–119.

[88] R Bammer, GH Glover, and ME Moseley. “Diffusion tensor spiral imaging”. In: *Proc 10th Annual Meeting of ISMRM, Honolulu*. 2002, p. 1111.

[89] Chunlei Liu et al. “Self-navigated interleaved spiral (SNAILS): application to high-resolution diffusion tensor imaging.” In: *Magn Reson Med* 52 (2004), pp. 1388–96.

[90] Hua Guo et al. “POCS-enhanced inherent correction of motion-induced phase errors (POCS-ICE) for high-resolution multishot diffusion MRI”. In: *Magn Reson Med* 75 (2016), pp. 169–180.

[91] Michael Lustig et al. “Compressed sensing MRI”. In: *IEEE Signal Processing Magazine* 25.2 (2008), pp. 72–82.

[92] Lawrence R Frank et al. “High efficiency, low distortion 3D diffusion tensor imaging with variable density spiral fast spin echoes (3D DW VDS RARE).” In: *NeuroImage* 49 (2010), pp. 1510–23.

[93] David A Feinberg et al. “Inner volume MR imaging: technical concepts and their application”. In: *Radiology* 156.3 (1985), pp. 743–747.

[94] Claudia A M Wheeler-Kingshott et al. “ADC mapping of the human optic nerve: Increased resolution, coverage, and reliability with CSF-suppressed ZOOM-EPI”. In: *Magn Reson Med* 47 (2002), pp. 24–31.

[95] Eun Kee Jeong et al. “High-resolution DTI with 2D interleaved multislice reduced FOV single-shot diffusion-weighted EPI (2D ss-rFOV-DWEPI)”. In: *Magn Reson Med* 54 (2005), pp. 1575–1579.

[96] Bertram J. Wilm et al. “Reduced field-of-view MRI using outer volume suppression for spinal cord diffusion imaging”. In: *Magn Reson Med* 57 (2007), pp. 625–630.

[97] Dimitrios C Karampinos et al. “High-resolution diffusion tensor imaging of the human pons with a reduced field-of-view, multishot, variable-density, spiral acquisition at 3 T.” In: *Magn Reson Med* 62 (2009), pp. 1007–16.

[98] Cornelius von Morze et al. “Reduced field-of-view diffusion-weighted imaging of the brain at 7 T.” In: *Magnetic resonance imaging* 28 (2010), pp. 1541–5.

[99] Robin M Heidemann et al. “k-space and q-space: combining ultra-high spatial and angular resolution in diffusion imaging using ZOOPPA at 7T”. In: *Neuroimage* 60.2 (2012), pp. 967–978.

[100] Christopher J. Wargo and John C. Gore. “Localized high-resolution DTI of the human midbrain using single-shot EPI, parallel imaging, and outer-volume suppression at 7T”. In: *Magnetic Resonance Imaging* 31 (2013), pp. 810–819.

[101] Cornelius Eichner et al. “Slice accelerated diffusion-weighted imaging at ultra-high field strength”. In: *Magn Reson Med* 71.4 (2014), pp. 1518–1525.

[102] Kawin Setsompop et al. “SLIce Dithered Enhanced Resolution Simultaneous MultiSlice (SLIDER-SMS) for high resolution (700 um) diffusion imaging of the human brain”. In: *Proc 23rd Annual Meeting ISMRM, Toronto*. 2015, p. 339.

[103] S. Rieseberg, J. Frahm, and J. Finsterbusch. “Two-dimensional spatially-selective RF excitation pulses in echo-planar imaging”. In: *Magn Reson Med* 47.6 (2002), pp. 1186–93.

[104] Emine Ulku Saritas et al. “DWI of the spinal cord with reduced FOV single-shot EPI”. In: *Magn Reson Med* 60 (2008), pp. 468–473.

[105] Emine Ulku Saritas et al. “Hadamard slice encoding for reduced-FOV diffusion-weighted imaging”. In: *Magn Reson Med* 72 (2014), pp. 1277–1290.

[106] Jürgen Finsterbusch. “High-resolution diffusion tensor imaging with inner field-of-view EPI”. In: *Journal of Magnetic Resonance Imaging* 29 (2009), pp. 987–993.

[107] Johannes T. Schneider et al. “Inner-volume imaging in vivo using three-dimensional parallel spatially selective excitation”. In: *Magn Reson Med* 69 (2013), pp. 1367–1378.

[108] G Zaharchuk et al. “Reduced field-of-view diffusion imaging of the human spinal cord: comparison with conventional single-shot echo-planar imaging.” In: *AJNR. American journal of neuroradiology* 32 (2011), pp. 813–820.

[109] J B Andre et al. “Clinical evaluation of reduced field-of-view diffusion-weighted imaging of the cervical and thoracic spine and spinal cord”. In: *Ajnr: American Journal of Neuroradiology* 33 (2012), pp. 1860–1866.

[110] Jin Young Park et al. “Comparison of readout segmented echo planar imaging (EPI) and EPI with reduced field-of-vIew diffusion-weighted imaging at 3t in patients with breast cancer”. In: *J Magn Reson Imaging* 42 (2015), pp. 1679–1688.

[111] Lisa Singer et al. “High-resolution Diffusion-weighted Magnetic Resonance Imaging in Patients with Locally Advanced Breast Cancer”. In: *Academic Radiology* 19.5 (2012), pp. 526–534.

[112] Lisa J. Wilmes et al. “High-Resolution Diffusion-Weighted Imaging for Monitoring Breast Cancer Treatment Response”. In: *Academic Radiology* 20.5 (2013), pp. 581–589.

[113] S. P. Souza et al. “Sima - Simultaneous Multislice Acquisition of Mr Images by Hadamard-Encoded Excitation”. In: *Journal of Computer Assisted Tomography* 12.6 (1988), pp. 1026–1030.

[114] S. Muller. “Multifrequency Selective Rf Pulses for Multislice Mr Imaging”. In: *Magn Reson Med* 6.3 (1988), pp. 364–371.

[115] David J. Larkman et al. “Use of multicoil arrays for separation of signal from multiple slices simultaneously excited”. In: *Journal of Magnetic Resonance Imaging* 13 (2001), pp. 313–317.

[116] RG Nunes et al. “Simultaneous slice excitation and reconstruction for single shot EPI”. In: *Proc 14th Annual Meeting ISMRM, Seattle, Washington*. 2006, p. 293.

[117] David A Feinberg et al. “Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging”. In: *PloS one* 5.12 (2010), e15710.

[118] Felix A Breuer et al. “Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging”. In: *Magnetic resonance in medicine* 53.3 (2005), pp. 684–691.

[119] K. L. Miller et al. “Multimodal population brain imaging in the UK Biobank prospective epidemiological study”. In: *Nat Neurosci* (2016).

[120] Erpeng Dai et al. “Simultaneous multislice accelerated interleaved EPI DWI using generalized blipped-CAIPI acquisition and 3D K-space reconstruction”. In: *Magn Reson Med* 00 (2016), pp. 1–13.

[121] R. Frost et al. “Simultaneous multi-slab acquisition in 3D multislab diffusion-weighted readout-segmented echo-planar imaging”. In: *Proc 21st Annual Meeting ISMRM, Salt Lake City*. 2013, p. 3176.

[122] Martin Blaimer et al. “Accelerated volumetric MRI with a SENSE/GRAPPA combination”. In: *J Magn Reson Imaging* 24.2 (2006), pp. 444–450.

[123] Daniel Stäb et al. “CAIPIRINHA accelerated SSFP imaging”. In: *Magn Reson Med* 65.1 (2011), pp. 157–164.

[124] Martin Blaimer et al. “Multiband phase-constrained parallel MRI”. In: *Magn Reson Med* 69.4 (2013), pp. 974–980.

[125] Kangrong Zhu, Adam Kerr, and J Pauly. “Autocalibrating CAIPIRINHA: reformulating CAIPIRINHA as a 3D problem”. In: *Proc 20th Annual Meeting ISMRM, Melbourne*. 2012, p. 518.

[126] S Moeller et al. “RO extended FOV SENSE/GRAPPA for multiband imaging with FOV shift”. In: *Proc 22nd Annual Meeting ISMRM, Milan*. 2014, p. 4396.

[127] Stephen F. Cauley et al. “Interslice leakage artifact reduction technique for simultaneous multislice acquisitions”. In: *Magn Reson Med* 72 (2014), pp. 93–102.

[128] E Wong. “Optimized phase schedules for minimizing peak RF power in simultaneous multi-slice RF excitation pulses”. In: *Proc 20th Annual Meeting ISMRM, Melbourne*. 2012, p. 2209.

[129] Edward J. Auerbach et al. “Multiband accelerated spin-echo echo planar imaging with reduced peak RF power using time-shifted RF pulses”. In: *Magn Reson Med* 69.5 (2013), pp. 1261–1267.

[130] Steven Conolly et al. “Variable-rate selective excitation”. In: *Journal of Magnetic Resonance (1969)* 78.3 (1988), pp. 440–458.

[131] David G. Norris et al. “Power independent of number of slices (PINS) radiofrequency pulses for low-power simultaneous multislice excitation”. In: *Magn Reson Med* 66 (2011), pp. 1234–1240.

[132] Cornelius Eichner, Lawrence L. Wald, and Kawin Setsompop. “A low power radiofrequency pulse for simultaneous multislice excitation and refocusing”. In: *Magn Reson Med* 72 (2014), pp. 949–958.

[133] Anuj Sharma et al. “kT-PINS RF pulses for low-power field inhomogeneity-compensated multislice excitation”. In: *Proc 21st Annual Meeting ISMRM, Salt Lake City*. 2013, p. 73.

[134] Rebecca E. Feldman et al. “A SEmi-Adiabatic matched-phase spin echo (SEAMS) PINS pulse-pair for B1-insensitive simultaneous multislice imaging”. In: *Magn Reson Med* 75.2 (2016), pp. 709–717.

[135] Xiaoping Wu et al. “Simultaneous multislice multiband parallel radiofrequency excitation with independent slice-specific transmit B1 homogenization”. In: *Magn Reson Med* 70 (2013), pp. 630–638.

[136] Benedikt A. Poser et al. “Simultaneous multislice excitation by parallel transmission”. In: *Magn Reson Med* 71 (2014), pp. 1416–1427.

[137] Xiaoping Wu et al. “Distributing coil elements in three dimensions enhances parallel transmission multiband RF performance: A simulation study in the human brain at 7 Tesla”. In: *Magn Reson Med* 2472 (2016), pp. 2464–2472.

[138] Bennett A. Landman et al. “Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI”. In: *NeuroImage* 59 (2012), pp. 2175–2186.

[139] Oleg Michailovich, Yogesh Rathi, and Sudipto Dolui. “Spatially regularized compressed sensing of diffusion MRI data”. In: *Psychiatry Interpersonal and Biological Processes* 1 (2010), pp. 1100–15.

[140] Marion I. Menzel et al. “Accelerated diffusion spectrum imaging in the human brain using compressed sensing”. In: *Magn Reson Med* 66 (2011), pp. 1226–1233.

[141] Berkin Bilgic et al. “Accelerated diffusion spectrum imaging with compressed sensing using adaptive dictionaries”. In: *Magn Reson Med* 68 (2012), pp. 1747–1754.

[142] Sylvain Merlet, Emmanuel Caruyer, and Rachid Deriche. “Parametric dictionary learning for modeling EAP and ODF in diffusion MRI.” In: *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention* 15 (2012), pp. 10–7.

[143] Julio M. Duarte-Carvajalino et al. “Estimation of the CSA-ODF using bayesian compressed sensing of multi-shell HARDI”. In: *Magn Reson Med* 72 (2014), pp. 1471–1485.

[144] Lipeng Ning et al. “A joint compressed-sensing and super-resolution approach for very high-resolution diffusion imaging”. In: *NeuroImage* 125 (2016), pp. 386–400.

[145] Xinwei Shi et al. “Parallel imaging and compressed sensing combined framework for accelerating high-resolution diffusion tensor imaging using inter-image correlation”. In: *Magn Reson Med* 73 (2015), pp. 1775–1785.

[146] Xiaoping Wu et al. “Z-stacked RF array design enhances parallel transmit multiband RF performance in whole brain simultaneous multislice”. In: *Proc Intl Soc Mag Reson Med* 8 (2014), p. 543.

[147] Merry Mani et al. “Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data”. In: *Magn Reson Med* 73 (2015), pp. 126–138.

[148] Tzu-Cheng Chao et al. “Fast diffusion imaging with high angular resolution”. In: *Magn Reson Med* (2016. doi:10.1002/mrm.26163).

[149] Hao Gao et al. “PCLR: Phase-constrained low-rank model for compressive diffusion-weighted MRI”. In: *Magn Reson Med* 72 (2014), pp. 1330–1341.

[150] Eun Kee Jeong, Seong Eun Kim, and Dennis L. Parker. “High-resolution diffusion-weighted 3D MRI, using diffusion-weighted driven-equilibrium (DW-DE) and multishot segmented 3D-SSFP without navigator echoes”. In: *Magn Reson Med* 50 (2003), pp. 821–829.

[151] X Golay et al. “High-resolution isotropic 3D diffusion tensor imaging of the human brain”. In: *Magn Reson Med* 47 (2002), pp. 837–843.

[152] Richard B Buxton. “The diffusion sensitivity of fast steady state free precession imaging”. In: *Magn Reson Med* 29.2 (1993), pp. 235–243.

[153] Elyakim Bosak and Paul R. Harvey. “Navigator motion correction of diffusion weighted 3D SSFP imaging”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 12 (2001), pp. 167–176.

[154] Karla L Miller et al. “Steady state diffusion weighted imaging of in vivo knee cartilage”. In: *Magn Reson Med* 51.2 (2004), pp. 394–398.

[155] Karla L Miller and John M Pauly. “Nonlinear phase correction for navigated diffusion imaging.” In: *Magn Reson Med* 50 (2003), pp. 343–353.

[156] Jennifer A. McNab, Daniel Gallichan, and Karla L. Miller. “3D steady state diffusion weighted imaging with trajectory using radially batched internal navigator echoes (TURBINE)”. In: *Magn Reson Med* 63 (2010), pp. 235–242.

[157] R. L. O’Halloran et al. “3D isotropic high resolution diffusion weighted MRI of the whole brain with a motion corrected steady state free precession sequence”. In: *Magn Reson Med* 70 (2013), pp. 466–478.

[158] Mathias Engström and Stefan Skare. “Diffusion-weighted 3D multislab echo planar imaging for high signal-to-noise ratio efficiency and isotropic image resolution”. In: *Magn Reson Med* 70 (2013), pp. 1507–1514.

[159] Anh T Van, Diego Hernando, and Bradley P Sutton. “Motion-induced phase error estimation and correction in 3D diffusion tensor imaging.” In: *IEEE transactions on medical imaging* 30 (2011), pp. 1933–40.

[160] Anh T Van et al. “Slab profile encoding (PEN) for minimizing slab boundary artifact in three-dimensional diffusion-weighted multislab acquisition.” In: *Magn Reson Med* 00 (2014), pp. 1–9.

[161] M. Engström et al. “On the Signal-to-Noise Ratio Efficiency and Slab-Banding Artifacts in Three-Dimensional Multislab Diffusion-Weighted Echo-Planar Imaging”. In: *Magn Reson Med* 73.2 (2015), pp. 718–725.

[162] Ha Kyu Jeong, John C. Gore, and Adam W. Anderson. “High-resolution human diffusion tensor imaging using 2-D navigated multishot SENSE EPI at 7 T”. In: *Magn Reson Med* 69 (2013), pp. 793–802.

[163] W. M. Teeuwisse, W. M. Brink, and A. G. Webb. “Quantitative assessment of the effects of high-permittivity pads in 7 Tesla MRI of the brain”. In: *Magn Reson Med* 67.5 (2012), pp. 1285–93.

[164] Kim Butts et al. “Isotropic diffusion-weighted and spiral-navigated interleaved EPI for routine imaging of acute stroke”. In: *Magnetic resonance in medicine* 38.5 (1997), pp. 741–749.

[165] Jiun Jie Wang et al. “3D DT-MRI Using a Reduced-FOV Approach and Saturation Pulses”. In: *Magn Reson Med* 51 (2004), pp. 853–857.

[166] Eun Kee Jeong et al. “High-resolution DTI of a localized volume using 3D single-shot diffusion-weighted stimulated echo-planar imaging (3D ss-DWSTEPI)”. In: *Magnetic Resonance in Medicine* 56 (2006), pp. 1173–1181.

[167] Youngkyoo Jung et al. “3D Diffusion Tensor MRI With Isotropic Resolution Using a Steady State Radial Acquisition”. In: *Journal of Magnetic Resonance Imaging* 29 (2009), pp. 1175–1184.

[168] DR Enzmann and NJ Pelc. “Brain motion: measurement with phase-contrast MR imaging.” In: *Radiology* 185.3 (1992), pp. 653–660.

[169] Dan Greitz et al. “Pulsatile brain movement and associated hydrodynamics studied by magnetic resonance phase imaging”. In: *Neuroradiology* 34.5 (1992), pp. 370–380.

[170] BP Poncelet et al. “Brain parenchyma motion: measurement with cine echo-planar MR imaging.” In: *Radiology* 185.3 (1992), pp. 645–651.

[171] Dennis L Parker, Chun Yuan, and Duane D Blatter. “MR angiography by multiple thin slab 3D acquisition”. In: *Magnetic resonance in medicine* 17.2 (1991), pp. 434–451.

[172] James W Murakami et al. “Multislab three-dimensional T2-weighted fast spin-echo imaging of the hippocampus: Sequence optimization”. In: *Journal of Magnetic Resonance Imaging* 5.3 (1995), pp. 309–315.

[173] Hing Chiu Chang et al. “Human brain diffusion tensor imaging at submillimeter isotropic resolution on a 3Tesla clinical MRI scanner”. In: *NeuroImage* 118 (2015), pp. 667–675.

[174] Kecheng Liu and Brian K Rutt. “Sliding interleaved kY (SLINKY) acquisition: a novel 3D MRA technique with suppressed slab boundary artifact”. In: *Journal of Magnetic Resonance Imaging* 8.4 (1998), pp. 903–911.

[175] Kecheng Liu, Yansun Xu, and Mark Loncar. “Reduced slab boundary artifact in multi-slab 3D fast spin-echo imaging”. In: *Magnetic resonance in medicine* 44.2 (2000), pp. 269–276.

[176] Dennis L Parker et al. “Magnetic resonance angiography with sliding interleaved projection reconstruction (SLIPR) acquisition”. In: *Journal of Magnetic Resonance Imaging* 10.4 (1999), pp. 569–575.

[177] Kie Tae Kwon et al. “Non-contrast-enhanced peripheral angiography using a sliding interleaved cylinder acquisition”. In: *Magnetic resonance in medicine* 74.3 (2015), pp. 727–738.

[178] Zhiqiang Li et al. “Sliding-slab three-dimensional TSE imaging with a spiral-In/Out readout”. In: *Magnetic resonance in medicine* 75.2 (2016), pp. 729–738.

[179] Eugene G Kholmovski, Andrew L Alexander, and Dennis L Parker. "Correction of slab boundary artifact using histogram matching". In: *Journal of Magnetic Resonance Imaging* 15.5 (2002), pp. 610–617.

[180] John Pauly et al. "Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm (NMR imaging)". In: *IEEE transactions on medical imaging* 10.1 (1991), pp. 53–65.

[181] Neil Gelman et al. "Interregional variation of longitudinal relaxation rates in human brain at 3.0 T: relation to estimated iron and water contents". In: *Magnetic resonance in medicine* 45.1 (2001), pp. 71–79.

[182] Eleanor F Cox and Penny A Gowland. "Simultaneous quantification of T2 and T2' using a combined gradient echo-spin echo sequence at ultrahigh field". In: *Magnetic resonance in medicine* 64.5 (2010), pp. 1440–1445.

[183] Andrew L Alexander et al. "Analysis of partial volume effects in diffusion-tensor MRI". In: *Magnetic Resonance in Medicine* 45.5 (2001), pp. 770–780.

[184] Iain P Bruce et al. "3D-MB-MUSE: A robust 3D multi-slab, multi-band and multi-shot reconstruction approach for ultrahigh resolution diffusion MRI". In: *NeuroImage* 159 (2017), pp. 46–56.

[185] Martin Uecker et al. "Image reconstruction by regularized nonlinear inversion—joint estimation of coil sensitivities and image content". In: *Magnetic Resonance in Medicine* 60.3 (2008), pp. 674–682.

[186] Martin Uecker et al. "Real-time MRI at a resolution of 20 ms". In: *NMR in Biomedicine* 23.8 (2010), pp. 986–994.

[187] Florian Knoll et al. "Parallel imaging with nonlinear reconstruction using variational penalties". In: *Magnetic Resonance in Medicine* 67.1 (2012), pp. 34–41.

[188] Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of inverse problems*. Vol. 375. Springer Science & Business Media, 1996.

[189] Thorsten Hohage. "Logarithmic convergence rates of the iteratively regularized Gauss-Newton method for an inverse potential and an inverse scattering problem". In: *Inverse problems* 13.5 (1997), p. 1279.

[190] Anatolii Borisovich Bakushinsky and M Yu Kokurin. *Iterative methods for approximate solution of inverse problems*. Vol. 577. Springer Science & Business Media, 2005.

[191] Magnus Rudolph Hestenes and Eduard Stiefel. *Methods of conjugate gradients for solving linear systems*. Vol. 49. 1. NBS, 1952.

[192] Berengere Aubert-Broche, Alan C Evans, and Louis Collins. "A new improved version of the realistic digital brain phantom". In: *NeuroImage* 32.1 (2006), pp. 138–145.

[193] Yongyue Zhang, Michael Brady, and Stephen Smith. "Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm". In: *IEEE transactions on medical imaging* 20.1 (2001), pp. 45–57.

[194] Stephen M Smith. "Fast robust automated brain extraction". In: *Human brain mapping* 17.3 (2002), pp. 143–155.

[195] Stephen M Smith et al. “Advances in functional and structural MR image analysis and implementation as FSL”. In: *Neuroimage* 23 (2004), S208–S219.

[196] Barbara Strotmann et al. “High-resolution MRI and diffusion-weighted imaging of the human habenula at 7 tesla”. In: *Journal of Magnetic Resonance Imaging* 39.4 (2014), pp. 1018–1026.

[197] David C Van Essen and Kamil Ugurbil. “The future of the human connectome”. In: *Neuroimage* 62.2 (2012), pp. 1299–1310.

[198] Daniel K Sodickson and Warren J Manning. “Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays”. In: *Magnetic Resonance in Medicine* 38.4 (1997), pp. 591–603.

[199] Jonathan R Polimeni et al. “Reducing sensitivity losses due to respiration and motion in accelerated echo planar imaging by reordering the autocalibration data acquisition”. In: *Magnetic resonance in medicine* 75.2 (2016), pp. 665–679.

[200] S Lalith Talagala et al. “Improvement of temporal signal-to-noise ratio of GRAPPA accelerated echo planar imaging using a FLASH based calibration scan”. In: *Magnetic resonance in medicine* 75.6 (2016), pp. 2362–2371.

[201] B Chapman et al. “Real-time movie imaging from a single cardiac cycle by NMR”. In: *Magnetic resonance in medicine* 5.3 (1987), pp. 246–254.

[202] D Ivanov et al. “Robust ACS acquisition for 3D echo planar imaging”. In: *Proceedings of the 23rd Annual Meeting of ISMRM, Toronto, Canada (Abstract 2059)*. 2015.

[203] Jesper LR Andersson, Stefan Skare, and John Ashburner. “How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging”. In: *Neuroimage* 20.2 (2003), pp. 870–888.

[204] Timothy EJ Behrens et al. “Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?” In: *Neuroimage* 34.1 (2007), pp. 144–155.

[205] Timothy EJ Behrens et al. “Characterization and propagation of uncertainty in diffusion-weighted MR imaging”. In: *Magnetic resonance in medicine* 50.5 (2003), pp. 1077–1088.

[206] Karla L Miller et al. “Diffusion tractography of post-mortem human brains: optimization and comparison of spin echo and steady-state free precession techniques”. In: *Neuroimage* 59.3 (2012), pp. 2284–2297.

[207] Jennifer A McNab et al. “Surface based analysis of diffusion orientation for identifying architectonic domains in the in vivo human cortex”. In: *Neuroimage* 69 (2013), pp. 87–100.

[208] Jennifer A McNab et al. “High resolution diffusion-weighted imaging in fixed human brain using diffusion-weighted steady state free precession”. In: *Neuroimage* 46.3 (2009), pp. 775–785.

[209] Jenni Schulz, Rasim Boyacioğlu, and David G Norris. “Multiband multislab 3D time-of-flight magnetic resonance angiography for reduced acquisition time and improved sensitivity”. In: *Magnetic resonance in medicine* 75.4 (2016), pp. 1662–1668.

[210] Essa Yacoub et al. “Imaging brain function in humans at 7 Tesla”. In: *Magnetic Resonance in Medicine* 45.4 (2001), pp. 588–594.

[211] Mark A Griswold et al. “Autocalibrated coil sensitivity estimation for parallel imaging”. In: *NMR in Biomedicine* 19.3 (2006), pp. 316–324.

[212] Jennifer A McNab, Daniel Gallichan, and Karla L Miller. “3D steady-state diffusion-weighted imaging with trajectory using radially batched internal navigator echoes (TURBINE)”. In: *Magnetic resonance in medicine* 63.1 (2010), pp. 235–242.

[213] Jesper LR Andersson and Stamatios N Sotiropoulos. “Non-parametric representation and prediction of single-and multi-shell diffusion-weighted MRI data using Gaussian processes”. In: *Neuroimage* 122 (2015), pp. 166–176.

[214] Jesper LR Andersson and Stamatios N Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *Neuroimage* 125 (2016), pp. 1063–1078.

[215] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*. Vol. 1. MIT press Cambridge, 2006.

[216] Dominic JA Welsh and Martin B Powell. “An upper bound for the chromatic number of a graph and its application to timetabling problems”. In: *The Computer Journal* 10.1 (1967), pp. 85–86.

[217] David C Van Essen et al. “The WU-Minn human connectome project: an overview”. In: *Neuroimage* 80 (2013), pp. 62–79.

[218] John N Morelli et al. “Evaluation of a modified Stejskal-Tanner diffusion encoding scheme, permitting a marked reduction in TE, in diffusion-weighted imaging of stroke patients at 3 T”. In: *Investigative radiology* 45.1 (2010), pp. 29–35.

[219] Leonid I Rudin, Stanley Osher, and Emad Fatemi. “Nonlinear total variation based noise removal algorithms”. In: *Physica D: Nonlinear Phenomena* 60.1-4 (1992), pp. 259–268.

[220] Marius de Groot et al. “Improving alignment in tract-based spatial statistics: evaluation and optimization of image registration”. In: *Neuroimage* 76 (2013), pp. 400–411.

[221] Kangrong Zhu et al. “Hybrid-Space SENSE Reconstruction for Simultaneous Multi-slice MRI”. In: *IEEE Trans Med Imaging* 35(8) (2016), pp. 1824–1836.

[222] Huairen Zeng and R Todd Constable. “Image distortion correction in EPI: comparison of field mapping with point spread function mapping”. In: *Magnetic Resonance in Medicine* 48.1 (2002), pp. 137–146.