

Acquisition and Reconstruction Methods for High-Resolution, High-Fidelity 3D Multi-Slab Diffusion MRI



Ziyu Li

Exeter College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2024

Acknowledgements

Reflecting on my DPhil journey, I am amazed at how swiftly time has passed over the three years. It has been a rewarding and enjoyable experience, made possible by the inspiration, support, and love of so many people.

My deepest gratitude goes to my supervisors, Profs. Wenchuan Wu and Karla Miller. Wenchuan is more than just a mentor to me — he has also been like an elder brother and a dear friend. The countless hours we have spent together in his office and in the scanner room are unforgettable, filled with learning, puzzle-solving, and a shared passion for MRI and science. Karla’s wisdom, insights, and her dedication to preparing me for a research career have been a constant source of inspiration. Her unwavering support has meant the world to me. They are, in every way, ideal supervisors, and if I ever have the opportunity to supervise others, I hope I can be even half as good as they have been.

If I could name a third supervisor, it would be Prof. Qiyuan Tian. I first met Qiyuan in 2020 during my internship at MGH, where he offered me the first systematic research training I ever received. Since then, he has been a guiding force in my life, both in science and in career development. I am deeply grateful for his generosity with his time and his continuous support throughout my journey. His influence is one of the key reasons I have chosen to pursue a career in academia toward professorship. He is a true role model for me.

A huge thanks goes to the entire FMRIB Physics Group for fostering such a friendly and inclusive atmosphere, as well as for the valuable advice and discussions during meetings and feedback sessions. Within FMRIB, a few individuals deserve recognition for their contributions to this thesis: Jesper Andersson and Rick Lange for their insights on registration and FSL’s “topup” and “eddy”; Prof. Mark Chiew for his guidance on MRI reconstruction methods; and my Transfer and Confirmation examiners, Profs. Aaron Hess and Tom Okell, for the enlightening discussions we had during the exams. Additionally, I extend my thanks to fellow students and friends — Xinyu Ye, James Kent, Zhiyu Zheng, Thijs de Buck, Cristiana Tisca, Xi Chen, Yang Ji, Daniel Kor, Minhao Hu, Hugh Simmons, Hossein Rafipoor, Ying-Qiu Zheng, Weikang Gong, and others — for making my DPhil journey more enjoyable and colourful. Special thanks go to those in the same cohort as me, Qijia Shen and Silei Zhu, for sharing many memorable DPhil experiences together. Here’s to a bright future for all of us.

I would also like to extend my gratitude to the wider MRI community beyond Oxford. Thank you to Prof. Susie Huang at MGH for the stimulating discussions on diffusion MRI and her support for my career development over the years; to Prof. Hua Guo at Tsinghua for his deep insights into diffusion MRI acquisition and reconstruction; to Yiming Dong at Leiden for both our scientific discussions and the short but pleasant time I spent with him and his adorable cat and dog; and to Profs. Klaas Pruessmann and Markus Weiger for their kind hospitality during my visit to their lab at ETH Zurich. I am also deeply grateful to the “Pulseq” and Harmonized MRI teams, led by Profs. Maxim Zaitsev, Jon-Fredrik Nielsen, Yogesh Rathi, and Berkin Bilgic, for their assistance with Pulseq programming and their dedication to the development, optimisation, and open-source sharing of Pulseq.

Finally, and most importantly, I wish to express my sincere thanks to my parents, Pu Li and Ruihua Jin, for their unconditional love, warmth, and support given to me throughout my life, which has been the foundation of my journey. The consistent companionship, encouragement, and love from my girlfriend and best friend, Zi Xin Ng, deserve special acknowledgement. After everything we have been through together, I feel we are ready for whatever comes next, and I look forward to all the future adventures we will share.

Abstract

Diffusion magnetic resonance imaging (dMRI) is a powerful tool for non-invasive mapping of brain microstructure and connectivity by measuring the random walk of water molecules. Conventional two-dimensional (2D) single-shot echo planar imaging (EPI) is widely used for dMRI due to its rapid acquisition and robustness against motion, but suffers from limitations such as long repetition times (TR) that reduce signal-to-noise ratio (SNR) efficiency and difficulty in achieving thin slice thickness, particularly for high-resolution imaging. To address these issues, three-dimensional (3D) multi-slab acquisition methods have been introduced, offering better SNR efficiency through shorter TR and the potential to achieve high, isotropic resolution dMRI for *in-vivo* human brain. However, key challenges remain for 3D multi-slab dMRI, including susceptibility to \mathbf{B}_0 inhomogeneity induced geometric distortions that compromise image anatomical fidelity, inefficiencies in current navigator-based methods for correcting motion-induced phase variations, and the need for careful considerations of blurring and noise to achieve true high spatial resolution.

This thesis develops novel acquisition and reconstruction methods to overcome these hurdles. First, sampling strategies and a joint reconstruction for blip-up/down data are designed to correct geometric distortions and slab boundary aliasing for 3D multi-slab imaging without increasing the scan time. Second, a self-navigated 3D multi-slab dMRI framework is developed, eliminating the need for navigator acquisitions through optimised self-navigation sampling and structured low-rank reconstruction for motion phase correction, which effectively shortens scan time and improves SNR efficiency. Lastly, we leverage in-plane segmented 3D multi-slab EPI and a denoiser-regularised reconstruction to achieve 0.5-0.6 mm isotropic resolutions for *in-vivo* dMRI with superior SNR, minimal blurring and distortions. These advances are expected to improve the image fidelity and scan efficiency of high-resolution 3D multi-slab dMRI, with the hope to ultimately benefit neuroscience research of the human brain.

Contents

List of Figures	ix
List of Tables	xi
List of Publications	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Thesis Outline	3
2 Background	7
2.1 MRI Physics	7
2.1.1 MR Signal	7
2.1.2 Image Formation	13
2.1.3 Parallel Imaging	19
2.2 Diffusion MRI	26
2.2.1 Introduction to Diffusion MRI	26
2.2.2 Diffusion MRI Acquisition Methods	29
2.2.3 Diffusion MRI Reconstruction Methods	33
3 Reducing Distortion and Slab-Boundary Aliasing for 3D Diffusion MRI	37
3.1 Introduction	38
3.2 Methods	41
3.2.1 Blip-Reversed Acquisition and Reconstruction	41
3.2.2 <i>In-Vivo</i> Experiments	45
3.2.3 Reconstruction Details	49
3.2.4 Post-Processing	51
3.2.5 Diffusion Analyses	52
3.3 Results	53
3.4 Discussion	64
3.5 Conclusions	70

4	Self-Navigated 3D Multi-Slab Diffusion MRI	71
4.1	Introduction	72
4.2	Methods	74
4.2.1	Theory	74
4.2.2	Experiments	83
4.3	Results	89
4.4	Discussion	101
4.5	Conclusion	105
5	Submillimetre 3D Multi-Slab Diffusion MRI	107
5.1	Introduction	108
5.2	Methods	111
5.2.1	Submillimetre 3D Diffusion Imaging Framework	111
5.2.2	In-Plane Segmented 3D Multi-Slab Acquisition	111
5.2.3	Denoiser-Regularised Reconstruction	114
5.2.4	<i>In-Vivo</i> Experiments	115
5.2.5	Reconstruction Details	117
5.2.6	Diffusion Analyses	119
5.3	Results	120
5.4	Discussion	132
5.5	Conclusion	137
6	Summary and Future Directions	139
6.1	Thesis Summary	139
6.2	Future Directions	141
6.3	Final Remarks	149
	References	151

List of Figures

2.1	Excitation of magnetisation with \mathbf{B}_1	9
2.2	Relaxation after a 90° RF excitation.	11
2.3	Illustration of a spin echo.	12
2.4	Gradient-encoding schemes and k-space trajectories.	17
2.5	Distortion artifacts in EPI.	18
2.6	Illustration of SENSE reconstruction.	21
2.7	Illustration of GRAPPA and SPIRiT reconstruction.	24
2.8	Diffusion preparation using the pulsed gradient spin echo sequence.	27
2.9	SNR efficiency comparison for spin echo diffusion acquisitions.	30
2.10	Phase inconsistency in multi-shot diffusion MRI acquisitions.	33
2.11	Sequence diagram for navigated diffusion MRI acquisitions.	35
2.12	Construction of block-Hankel matrix using multi-shot k-space data.	36
3.1	Proposed acquisition and reconstruction framework.	42
3.2	3D multi-slab acquisition and sampling patterns for blip-up data.	43
3.3	Sequence diagram of the proposed 3D multi-slab CAIPI-PF acquisition.	46
3.4	Illustration of 2D navigator acquisition trajectory	47
3.5	Impact of SPIRiT regularisation weight on reconstruction.	50
3.6	Boundary slice aliasing correction.	54
3.7	Stage 1 blip-up reconstruction.	55
3.8	Stage 1 blip-down sampling and reconstruction.	56
3.9	Field map estimation and joint reconstruction.	57
3.10	Comparisons of whole-brain diffusion images.	59
3.11	High-resolution whole-brain diffusion images.	60
3.12	Comparison of image anatomical fidelity.	61
3.13	Correction of eddy current induced distortion.	62
3.14	Comparison of tractography results.	63
3.15	Diffusion MRI results of multiple subjects using the proposed method.	64
3.16	Maximum intensity projections of fibre tracking results from three subjects.	65
3.17	Diffusion image results of the remaining subjects.	66
3.18	The $b=0$ images of a representative subject.	69

4.1	Comparison between the conventional navigated 3D multi-slab imaging and the proposed self-navigated imaging framework.	77
4.2	K-space sampling optimisation.	79
4.3	The block-Hankel matrix for the $k_z = 0$ plane reconstruction.	81
4.4	Image reconstruction results from the simulation experiment.	90
4.5	Phase map estimation from the simulation experiment.	91
4.6	Evaluation of sampling optimisation.	92
4.7	Impact of including the shot traversing the $k_z = 0$ plane.	93
4.8	Comparison of self-navigation methods.	94
4.9	Comparison of different reconstruction configurations.	95
4.10	Reconstruction of prospectively acquired data.	96
4.11	Impact of the $k_z = 0$ shot in the in-vivo experiment.	97
4.12	SNR maps of different diffusion encoding directions.	99
4.13	SNR efficiency and DTI comparisons.	100
4.14	Results at high b-values.	101
4.15	Tractography results.	102
5.1	Impact of the sampling order of in-plane segmented 3D EPI.	113
5.2	Masks for tracking the short association fibres.	121
5.3	Simulation of effective resolution and SNR for high-resolution diffusion-weighted EPI at 3T.	122
5.4	Simulation of effective resolution and SNR for high-resolution diffusion-weighted EPI at 7T.	123
5.5	Denoiser-regularised reconstruction results.	124
5.6	Retrospective under-sampled reconstruction.	125
5.7	Whole-brain data from 3T 0.65 mm Protocol.	127
5.8	Comparison of 3T 0.65 mm and 1.22 mm isotropic resolution DTI.	128
5.9	Whole-brain data from 3T 0.53 mm Protocol.	129
5.10	Example white matter tracts from the 0.53 mm and 1.22 mm data.	130
5.11	Gyral bias comparison between 0.53 mm and 1.22 mm data.	131
5.12	Fibre distributions at a large gyrus.	132
5.13	U-fibres comparison between 0.53 mm and 1.22 mm data.	133
5.14	Comparisons of 7T 0.61 mm and 1.05 mm diffusion data.	134
6.1	Simulation study for self-navigated, distortion-corrected 3D multi-slab dMRI.	143
6.2	FOV shifting and q-space modelling for improved slab boundary artifacts correction.	145

List of Tables

3.1	Tract mask volumes.	61
4.1	Acquisition parameters.	85
4.2	Group-level SNR values along different diffusion encoding directions.	98
5.1	Submillimetre dMRI acquisition parameters.	115

List of Publications

Part of This Thesis

Journal Articles

1. **Li Z**, Miller KL, Zhu S, Wu W, Submillimeter 3D diffusion MRI using in-plane segmented multi-slab EPI and a denoiser-regularized reconstruction. 2024. Preprint available on *bioRxiv*. doi: 10.1101/2024.10.10.617536.
2. **Li Z**, Miller KL, Chen X, Chiew M, Wu W, Self-navigated 3D diffusion MRI using an optimized CAIPI sampling and structured low-rank reconstruction estimated navigator. *IEEE Trans Med Imaging*. 2024 Sep 6;PP. doi: 10.1109/TMI.2024.3454994.
3. **Li Z**, Miller KL, Andersson JLR, Zhang J, Liu S, Guo H, Wu W. Sampling strategies and integrated reconstruction for reducing distortion and boundary slice aliasing in high-resolution 3D diffusion MRI. *Magn Reson Med*. 2023 Oct;90(4):1484-1501. doi: 10.1002/mrm.29741.

Conference Proceedings

1. **Li Z**, Miller KL, Wu W. Submillimeter diffusion MRI using in-plane segmented 3D multi-slab EPI and a denoiser-regularized reconstruction. The Annual Meeting of ISMRM, Singapore, 2024.
2. **Li Z**, Miller KL, Chen X, Chiew M, Wu W. Self-navigated 3D multi-slab EPI for SNR-efficient high-resolution diffusion MRI. The Annual Meeting of ISMRM, Singapore, 2024. (*Oral Presentation*)
3. **Li Z**, Chen X, Chiew M, Miller KL, Wu W. Self-navigated high-resolution 3D diffusion MRI using an extended blipped-CAIPI sampling and structured low-rank reconstruction. The Annual Meeting of ISMRM, Toronto, 2023.
4. **Li Z**, Miller KL, Wu W. Elimination of distortion and slice-aliasing in 3D diffusion MRI by integrating multiple sampling strategies into reconstruction. ISMRM Diffusion Workshop, Amsterdam, 2022. (*Oral Presentation*)
5. **Li Z**, Miller KL, Wu W. Integration of blip reversal with CAIPI sampling enables simultaneous correction of slice aliasing and distortion in 3D multi-slab diffusion MRI. The Annual Meeting of ISMRM, London, 2022. (*Oral Power Pitch*)

Not Part of This Thesis*

*Only journal articles are listed.

1. **Li Z**, Li Z, Li H, Fan Q, Miller KL, Wu W, Chaudhari AS, Tian Q, Enhance the image: super resolution using artificial intelligence in MRI. 2024. Preprint available on *arXiv*. doi: 10.48550/arXiv.2406.13625.
2. **Li Z**, Fan Q, Bilgic B, Wang G, Wu W, Polimeni JR, Miller KL, Huang SY, Tian Q. Diffusion MRI data analysis assisted by deep learning synthesized anatomical images (DeepAnat). *Med Image Anal.* 2023 May;86:102744. doi: 10.1016/j.media.2023.102744.
3. **Li Z**, Tian Q, Ngamsombat C, Cartmell S, Conklin J, Filho ALMG, Lo WC, Wang G, Ying K, Setsompop K, Fan Q, Bilgic B, Cauley S, Huang SY. High-fidelity fast volumetric brain MRI using synergistic wave-controlled aliasing in parallel imaging and a hybrid denoising generative adversarial network (HDnGAN). *Med Phys.* 2022 Feb;49(2):1000-1014. doi: 10.1002/mp.15427.
4. Suzuki Y, Koktzoglou I, **Li Z**, Jezzard P, Okell T. Improved visualization of intracranial distal arteries with multiple 2D slice dynamic ASL-MRA and super-resolution convolutional neural network. *Magn Reson Med.* 2024 Dec;92(6):2491-2505. doi: 10.1002/mrm.30245.
5. Yang H, Wang G, **Li Z**, Li H, Zheng J, Hu Y, Cao X, Liao C, Ye H, Tian Q. Artificial intelligence for neuro MRI acquisition: a review. *MAGMA.* 2024 Jul;37(3):383-396. doi: 10.1007/s10334-024-01182-7.
6. Li Z, **Li Z**, Bilgic B, Lee HH, Ying K, Huang SY, Liao H, Tian Q. DIMOND: Diffusion Model Optimization with Deep Learning. *Adv Sci.* 2024 Jun;11(24):e2307965. doi: 10.1002/advs.202307965.
7. Dong Y, Koolstra K, **Li Z**, Riedel M, van Osch MJP, Börnert P. Structured low-rank reconstruction for navigator-free water/fat separated multi-shot diffusion-weighted EPI. *Magn Reson Med.* 2024 Jan;91(1):205-220. doi: 10.1002/mrm.29848.
8. Tian Q, **Li Z**, Fan Q, Polimeni JR, Bilgic B, Salat DH, Huang SY. SDnDTI: Self-supervised deep learning-based denoising for diffusion tensor MRI. *Neuroimage.* 2022 Jun;253:119033. doi: 10.1016/j.neuroimage.2022.119033.
9. Zhang X, Yang J, Li J, Li W, Song D, Lu XA, Wu F, Li J, Chen D, Li X, Xu Z, Liu S, **Li Z**, Ying K, Lu P. Factors associated with treatment response to CD19 CAR-T therapy among a large cohort of B cell acute lymphoblastic leukemia. *Cancer Immunol Immunother.* 2022 Mar;71(3):689-703. doi: 10.1007/s00262-021-03009-z.

1

Introduction

Contents

1.1 Motivation	1
1.2 Thesis Outline	3

1.1 Motivation

Diffusion magnetic resonance imaging (dMRI) is a powerful, non-invasive tool for mapping brain microstructure and connectivity by characterising the diffusion of water molecules. High-resolution dMRI is desirable for resolving fine structures and lesions, such as ischemic changes and multiple sclerosis [1], and for linking dMRI findings with *post-mortem* examinations (e.g., histological staining) [2]. However, a significant barrier to the broader adoption of high-resolution dMRI in research and clinical practice is its intrinsically low signal-to-noise ratio (SNR), which requires longer scan times that may not be feasible in many settings.

For over three decades, two-dimensional (2D) single-shot echo planar imaging (EPI) has been the mainstay for dMRI acquisition due to its rapid acquisition speed (typically 70-150 ms per slice) and robustness against motion [3]. Despite these advantages, 2D EPI often suffers from limited SNR efficiency due to its

long repetition time (TR) necessary for covering a large number of slices for whole-brain imaging, especially for high-resolution dMRI which requires more slices. Typically, the TR of 2D EPI-based dMRI is longer than 4 seconds, whereas the SNR efficiency for spin-echo-based dMRI is maximised at TR=1-2 seconds. Additionally, achieving thin slices (≤ 1 mm) with 2D EPI is challenging due to the limitations in gradient strengths and radiofrequency (RF) power deposition to achieve precise thin slice selection, and cross-talk between adjacent slices, which further distorts the ideal slice profile.

To achieve SNR-efficient high-resolution dMRI, three-dimensional (3D) multi-slab acquisitions have been proposed [4–7]. In this method, the whole brain is divided into several thin slabs (typically 6-15 for full-brain coverage). The number of excitations within each TR is proportional to the number of slabs rather than much larger number of 2D slices, allowing for a shorter TR and increased SNR efficiency for dMRI [8]. Within each slab, 3D multi-shot EPI is typically used for efficient encoding, with each shot covering part of one k-space plane in the slice direction (k_z). By using RF pulses to excite thicker slabs rather than very thin slices, 3D multi-slab acquisitions reduce the gradient strength requirements while 3D k-space encoding within each slab enables high, isotropic resolution.

However, similar to 2D EPI, 3D multi-slab EPI suffers from \mathbf{B}_0 field inhomogeneity induced geometric distortions, compromising the anatomical fidelity of dMRI and potentially confounding downstream analyses such as fibre tracking (tractography). Previous studies have shown that acquiring a pair of blip-reversed phase-encoding (i.e., blip-up and blip-down) images and combining them for distortion correction can effectively improve anatomical fidelity for 2D acquisitions [9–14]. However, directly applying such acquisition strategies in 3D multi-slab EPI would double the scan time. Additionally, 3D multi-slab EPI is prone to slab boundary aliasing artifacts due to imperfect RF profile and limited field-of-view (FOV) along the slice direction for each slab. While oversampling along k_z can expand FOV and mitigate those artifacts, it would further prolong the scan. Therefore, a robust

acceleration framework is needed to correct distortions and boundary slice aliasing without increasing the scan time.

The multi-shot nature of 3D multi-slab EPI also makes it susceptible to shot-to-shot phase variations due to subject motion during the diffusion-encoding gradients. It has been demonstrated that 3D phase errors can be well approximated by 2D navigators if the slabs are sufficiently thin (i.e., ≤ 2 cm) [4]. However, acquiring 2D navigators requires an additional spin echo which prolongs the TR and increases the scan time by 25%-50% [15]. Furthermore, the additional refocusing RF pulse for the navigator echo increases the SAR, potentially necessitating further extensions to the TR, especially at high fields. A navigator-free 3D multi-slab EPI framework can further reduce the scan time and improve the SNR efficiency for dMRI by shortening the TR to the SNR-efficient regime.

With the potential to achieve optimal SNR efficiency and high, isotropic resolution, 3D multi-slab EPI holds great promise for enabling submillimetre dMRI to study the brain's fine-scale structures. However, the large matrix size associated with submillimetre resolution can lead to long echo spacing (strong distortions), readout times (T_2^* blurring) and long TE (low SNR). The intrinsically low SNR from the small voxel size might also hamper accurate analysis of the diffusion images. A novel imaging framework capable of providing submillimetre dMRI with minimal blurring, high SNR and anatomical fidelity has great potential to increase spatial resolution and resolve fine-detailed brain structures.

In summary, a few key challenges remain for 3D multi-slab dMRI to realise its potential benefits: (i) susceptibility to \mathbf{B}_0 -induced geometric distortions that compromise image anatomical fidelity; (ii) the inefficiencies of current navigator-based methods for correcting motion-induced phase artifacts; and (iii) the need for careful considerations of blurring and noise to achieve true high spatial resolution.

1.2 Thesis Outline

To address the challenges above, this thesis describes the development of acquisition and reconstruction methods for improving the imaging fidelity and SNR efficiency

for high-resolution 3D multi-slab dMRI. The structure of this thesis is as follows:

Chapter 2 gives a brief background in MRI physics, including MR signal generation, image formation, and parallel imaging methods. It also covers the basics of diffusion MRI, detailing commonly used dMRI acquisition and reconstruction methods.

Chapter 3 is the first research chapter on correcting distortions and boundary slice aliasing in 3D multi-slab dMRI without increasing scan time. Novel sampling strategies, such as k_z -blipped CAIPI and complementary partial Fourier, are designed to efficiently acquire both blip-up and blip-down data. Oversampling along k_z is used to reduce boundary slice aliasing. The reconstruction process involves two stages: first, reconstructing blip-up and blip-down images separately each using half of acquired shots, based on which a dynamic field map is estimated; second, incorporating all blip-up/down data and the dynamic field map into a joint reconstruction to produce a final image corrected for both distortion and slice aliasing. This method is validated in high-resolution *in-vivo* experiments at 7T, demonstrating improved anatomical fidelity compared to conventional methods with matched scan times.

Chapter 4 describes a self-navigated 3D multi-slab dMRI framework that eliminates the need for navigator acquisition, reducing scan times and improving SNR efficiency. A key challenge in this method is balancing the requirements for self-navigation and overall reconstruction quality. To achieve reliable self-navigation, each shot is designed to intersect with the central $k_z=0$ plane of each slab. At the same time, the multi-shot sampling trajectory is carefully optimised to provide sufficient k-space coverage across all shots to maintain high reconstruction quality. The $k_z=0$ intersections from all shots are jointly used to obtain a 2D phase map for each shot through a structured low-rank reconstruction. These phase maps are then used to eliminate shot-to-shot phase inconsistencies in the final 3D multi-shot reconstruction. *In-vivo* experiments at 7T show this method shortens scan time by 31.7% and improves SNR efficiency by 15.5% compared to conventional navigated methods.

Chapter 5 introduces a submillimetre dMRI framework using 3D multi-slab EPI with superior SNR, minimal blurring and distortions. Numerical simulations quantify the effective resolution and SNR for high-resolution diffusion-weighted EPI. Based on these simulations, a 3D Fourier encoding-based acquisition with in-plane segmented multi-slab EPI is designed to shorten the echo spacing (reducing distortions), readout time (reducing T_2^* blurring), and TE (improving SNR). A denoiser-regularised reconstruction is developed to further enhance SNR while preserving image fidelity. This method is implemented using the scanner-agnostic, open-source platform “Pulseq” to facilitate wider use. Comprehensive *in-vivo* experiments conducted at both 3T and 7T yield high-SNR, high-quality dMRI at isotropic resolutions of 0.53 mm, 0.61 mm, and 0.65 mm, revealing significantly richer microstructure information compared to diffusion data acquired at conventional resolutions of 1-1.2 mm.

Chapter 6 summarises the research work in this thesis and discusses potential future research directions, including preliminary attempts to integrate the methods developed in Chapters 3 and 4.

2

Background

Contents

2.1 MRI Physics	7
2.1.1 MR Signal	7
2.1.2 Image Formation	13
2.1.3 Parallel Imaging	19
2.2 Diffusion MRI	26
2.2.1 Introduction to Diffusion MRI	26
2.2.2 Diffusion MRI Acquisition Methods	29
2.2.3 Diffusion MRI Reconstruction Methods	33

2.1 MRI Physics

This section covers the basic principles of MRI physics, much content of which is based on standard text books [16, 17].

2.1.1 MR Signal

MRI relies on the principles of nuclear magnetic resonance (NMR), where nuclei within an external magnetic field absorb and re-emit radiofrequency (RF) energy.

In short, the generation of MR signal involves four key steps:

- **Polarisation:** In the presence of an external magnetic field \mathbf{B}_0 , nuclei with nonzero spins tend to align with the field and create net magnetisation.
- **Excitation:** A brief RF pulse (magnetic field \mathbf{B}_1) that satisfies the resonance condition is applied, perturbing the magnetisation and tipping it out of equilibrium.
- **Precession:** Once out of alignment with \mathbf{B}_0 , the magnetisation precesses about the axis of \mathbf{B}_0 , during which MR signals can be detected.
- **Relaxation:** The excited magnetisation gradually returns to its equilibrium state.

The MRI scanner's receive coils detect the emitted signal. The signal strength depends on the environment of the detected nuclei, such as the tissue properties. This dependency allows for the generation of detailed images of the body's internal structures based on differences in tissue composition and characteristics.

Polarisation

Only nuclei with an odd number of protons or neutrons possess an intrinsic spin angular momentum, resulting in a nonzero spin. This spin gives rise to an associated magnetic moment, akin to a tiny magnet, due to the proton's positive charge. When placed in an external magnetic field \mathbf{B}_0 , the spins tend to slightly align with the field, with a marginally higher number of spins aligning parallel to \mathbf{B}_0 than anti-parallel (e.g., ~ 10 per million at 3T), which leads to a net magnetic moment \mathbf{M} (known as "magnetisation"). This process is referred to as polarisation.

The created magnetisation \mathbf{M} has a resonance frequency, at which it tends to oscillate. This frequency is proportional to the strength of the magnetic field B_0 , as given by the following equation:

$$\omega_0 = \gamma B_0, \quad (2.1)$$

where γ represents the gyromagnetic ratio of the nucleus (e.g., $\gamma = 2.675 \times 10^8$ rad/s/T for ^1H and $\gamma = 7.075 \times 10^7$ rad/s/T for ^{31}P). The resonance frequency

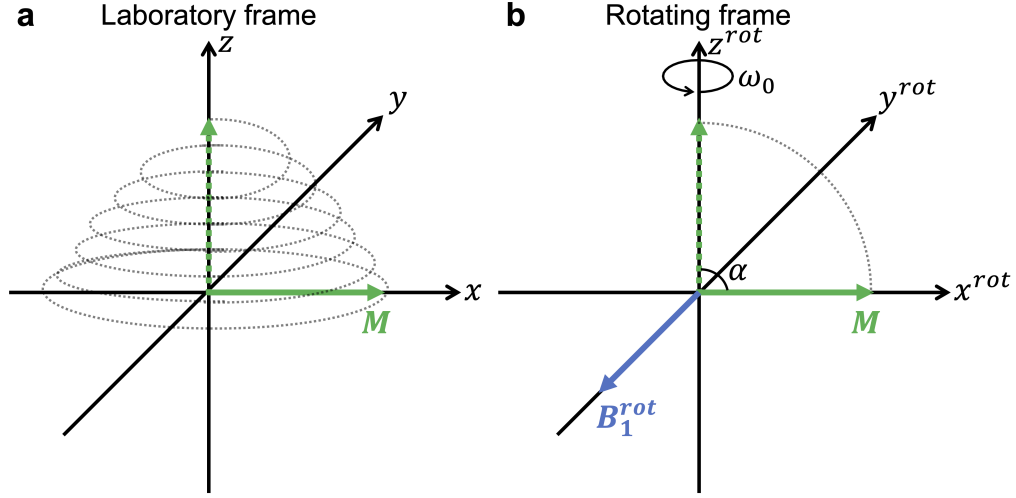


Figure 2.1: Excitation of magnetisation with \mathbf{B}_1 . (a) In the laboratory frame of reference, the magnetisation \mathbf{M} tilts away from alignment with \mathbf{B}_0 field. (b) In the frame that rotates with \mathbf{B}_1 (indicated by the direction of ω_0), \mathbf{M} flips about the rotation axis defined by the RF field (i.e., $\mathbf{B}_1^{\text{rot}}$). α is known as the flip angle ($\alpha = 90^\circ$ in this case).

ω_0 is known as the Larmor frequency. In practice, ^1H is the most commonly investigated nuclei due to its great abundance in the human body (e.g., hydrogen comprised in water and fat) and its high sensitivity to external magnetic field (i.e., higher γ and stronger signal).

Excitation

To detect the MR signal, the magnetisation along the longitudinal axis must be perturbed out of alignment with \mathbf{B}_0 through a process known as excitation. This involves applying a secondary magnetic field, \mathbf{B}_1 , which oscillates at the Larmor frequency in the transverse plane to satisfy the resonance condition. Since the Larmor frequency for protons is within radiofrequency and \mathbf{B}_1 is only applied for a short duration (less than 15 ms), \mathbf{B}_1 is commonly referred to as an RF pulse. The RF pulse exerts a torque on \mathbf{M} , causing the magnetisation to tilt away from its equilibrium direction along the longitudinal axis (z-axis) and begin to rotate about that axis. The x-y plane is termed the transverse plane.

In the laboratory frame of reference, the RF pulse causes \mathbf{M} to rotate about the z-axis at the Larmor frequency while gradually tilting away from it (Fig. 2.1a). To better understand this process, it is helpful to consider the rotating frame,

which rotates with \mathbf{B}_1 at the Larmor frequency, where \mathbf{M} simply flips about the axis defined by the RF field (Fig. 2.1b). The angle of rotation, known as the flip angle, is determined by the magnitude and the duration of the RF pulse — greater magnitude and longer duration result in a larger flip angle.

Precession

Once the magnetisation is perturbed out of alignment with \mathbf{B}_0 , it experiences a torque from \mathbf{B}_0 , which tends to restore it back to alignment. However, due to the magnetisation's inherent angular momentum, this torque causes \mathbf{M} to precess around the longitudinal axis at the Larmor frequency instead of directly rotating into alignment.

The precession of \mathbf{M} generates a time-varying magnetic field that induces a voltage in a receiver, allowing the signal to be detected. This receiver may be the same coil used for excitation, as is common in pre-clinical systems and NMR spectrometers, or an array of receive coils in most human MRI scanners, which enables parallel imaging as detailed in Section 2.1.3. The strength of the detected MRI signal is proportional to the magnitude of transverse magnetisation.

Relaxation

The process by which the excited magnetisation \mathbf{M} returns to its equilibrium state after excitation is called relaxation. Relaxation encompasses two components: the recovery of longitudinal magnetisation and the decay of transverse magnetisation, described by the time constants T_1 and T_2 , respectively.

Relaxation is essentially a process where the excited spins release the energy they absorbed. During T_1 relaxation, the magnetisation transfers energy to its surroundings, and the longitudinal component (M_z) gradually returns to its equilibrium value, M_0 . As shown in Fig. 2.2a, following a 90° RF pulse, the recovery of M_z can be described by:

$$M_z(t) = M_0(1 - e^{-t/T_1}). \quad (2.2)$$

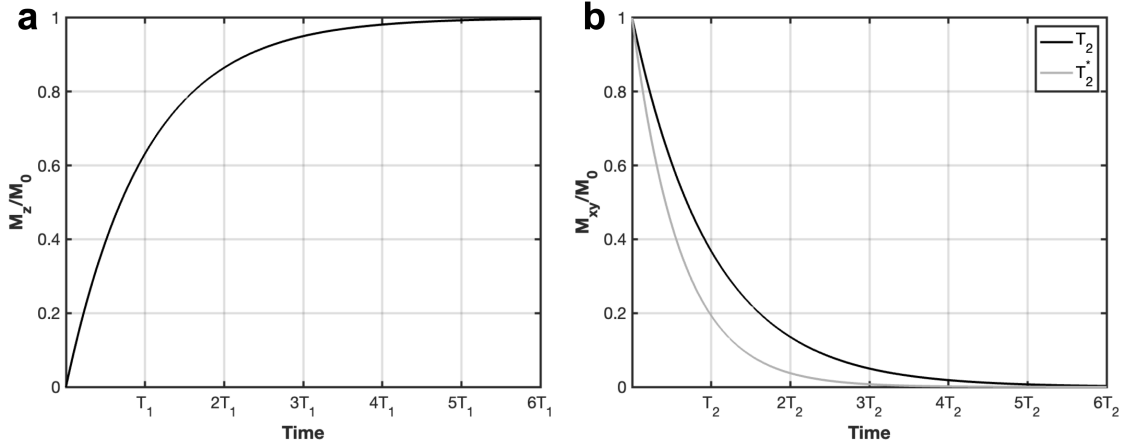


Figure 2.2: Relaxation after a 90° RF excitation. (a) During T_1 relaxation, longitudinal magnetisation (M_z) gradually recovers. (b) Magnetisation in the transverse plane (M_{xy}) gradually decays during T_2 and T_2^* relaxation. T_2^* decay is more rapid than T_2 due to field inhomogeneity-induced spin de-phasing. Here, $T_2 = 69\text{ms}$ and $T_2^* = 45\text{ms}$ are assumed (parameters for white matter at 3T).

T_1 is determined by the local environment of the nuclei, such as tissue properties (e.g., surrounding lipids, proteins, and macromolecules [18]). It represents the time required for M_z to recover to approximately 63% of M_0 after a 90° excitation and typically increases with the strength of the magnetic field. The interval between successive excitations, known as the repetition time (TR), is chosen to allow for T_1 recovery. A longer TR enables a more complete recovery of M_z and stronger MR signal at a cost of longer scan time.

During T_2 relaxation, spins lose coherence in the transverse plane due to small variations in their precession frequencies, caused by interactions with neighbouring molecules, such as random thermal motions. This leads to the exponential decay of transverse magnetisation (M_{xy}), as illustrated in Fig. 2.2b. After a 90° RF pulse, this decay is described by:

$$M_{xy}(t) = M_0 e^{-t/T_2}. \quad (2.3)$$

T_2 is influenced by interactions between the protons themselves [18] and represents the time it takes for M_{xy} to decay to approximately 37% of M_0 . Typically, T_2 decreases as the field strength increases and is shorter than T_1 . For white matter at 3T, T_1 is about 1084 ms, while T_2 is around 69 ms [19].




Figure 2.3: Illustration of a spin echo. The 90° excitation RF pulse is followed by a 180° refocusing RF pulse in a spin echo. The signal experiences T_2^* decay after 90° excitation. At the echo time (2τ), only a T_2 decay is observed.

Other factors, such as \mathbf{B}_0 inhomogeneities, local susceptibility variations, and gradient fields used for diffusion encoding (see Section 2.2), can result in a more rapid decay of M_{xy} , which is referred to as T_2^* relaxation. The relationship between T_2^* and T_2 is described by:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}, \quad (2.4)$$

where T_2' is the signal decay induced by field inhomogeneity. To mitigate the signal loss due to T_2' , a 180° refocusing RF pulse can be applied after the 90° excitation pulse (Fig. 2.3). This pulse reverses the phase of the spins in the transverse plane, allowing them to rephase and maximise the signal. This process is known as a “spin echo”. The time from excitation to the refocused signal maximum is called the echo time (TE), represented as 2τ in Fig. 2.3.

Bloch Equation

The effect discussed above can be expressed in a single differential equation, known as Bloch equation, which describes the behaviour of the magnetisation with the presence of an external magnetic field:

$$\frac{d\mathbf{M}}{dt} = (\mathbf{M} \times \gamma\mathbf{B}) - \frac{M_x\hat{\mathbf{i}} - M_y\hat{\mathbf{j}}}{T_2} + \frac{(M_0 - M_z)\hat{\mathbf{k}}}{T_1}, \quad (2.5)$$

where \mathbf{B} is the total magnetic field including \mathbf{B}_0 , \mathbf{B}_1 , and the gradient field (see Section 2.1.2), $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$, and $\hat{\mathbf{k}}$ are the unit vectors along x, y, and z axes, respectively.

2.1.2 Image Formation

Gradient Fields

In addition to \mathbf{B}_0 and \mathbf{B}_1 essential for the generation of MR signals, another set of magnetic fields, known as the gradient fields, are applied to create spatial variation in the resonance frequency that enables the localisation of MR signals. The gradient fields are aligned with \mathbf{B}_0 but varies in magnitude in a linear manner along the x, y, and z axes. These variations are produced by three separate gradient coils whose strengths can be independently manipulated in time. With the presence of the gradient field, the total magnetic field becomes:

$$\mathbf{B}(\mathbf{r}) = \mathbf{B} + \mathbf{G} \cdot \mathbf{r}, \quad (2.6)$$

where \mathbf{r} is the 3D spatial coordinate and $\mathbf{G} = [G_x, G_y, G_z]$ represents the gradient field strength along the x, y, and z axes, respectively. This gradient introduces spatial variations to the resonance frequency, $\omega(\mathbf{r}) = \omega_0 + \gamma(\mathbf{G} \cdot \mathbf{r})$, which facilitates the creation of an image as described below. Gradient fields are very small compared to \mathbf{B}_0 . For example, the maximum gradient strength for a Siemens 3T Prisma scanner is 80 mT/m.

Slice Selection

Slice-selective excitation, which restricts the excited signal to a specific slice or slab, can be achieved using the gradient field introduced above. To achieve this, an RF pulse needs to be applied with the gradient field simultaneously.

One simple and effective method to design a slice-selective RF pulse is the Fourier transform approach. With the gradient field, the spatial selection problem is converted to a frequency selection problem, i.e., selecting the range of frequencies ω that correspond to the location and thickness of the desired slice/slab. Denoting this frequency selection as $p(\omega) = \Pi(\frac{\omega - \omega_c}{\Delta\omega})$, where Π is a ‘‘boxcar’’ function of width $\Delta\omega$ centred at $\omega = \omega_c$, the key assumption of the Fourier approach is the RF pulse $B_1(t)$ is related to $p(\omega)$ by the Fourier transform: $B_1(t) \propto \int_{-\infty}^{+\infty} p(\omega)e^{-i2\pi\omega t} d\omega$.

Given the Fourier transform relationship that a “boxcar” function corresponds to a sinc function in the Fourier domain, we obtain:

$$B_1(t) \propto \int_{-\infty}^{+\infty} p(\omega) e^{-i2\pi\omega t} d\omega = \Delta\omega \operatorname{sinc}(\pi\Delta\omega t) e^{-i2\pi\omega_c t}. \quad (2.7)$$

The excitation frequency of the RF pulse ω_c is the resonance frequency at the centre of the excited slice or slab, and the pulse envelope function is $A \operatorname{sinc}(\pi\Delta\omega t)$, where A is a constant determined by the desired flip angle. The width $\Delta\omega$ is determined by $\Delta\omega = \gamma G_z \Delta z$ where Δz is the desired slice/slab thickness. In practice, $\Delta\omega$ is usually held constant (on the order of 1000-2000 Hz), and slice thickness Δz is varied by adjusting G_z : thinner slices require stronger gradients [20].

However, the Fourier transform approach relies on the assumption of small flip angles, where the longitudinal magnetisation is only slightly perturbed, allowing a linear response between the RF pulse and the resulting magnetisation. In practice, this method works well for flip angles up to about 90° . Beyond this, the response becomes nonlinear due to the Bloch equation’s intrinsic nonlinearity, and more advanced methods, such as direct solutions to the Bloch equations, are needed for accurate pulse design.

Additionally, as a perfect rectangular RF profile requires an infinitely long RF pulse which is not achievable, the use of short, truncated RF pulse in practice inevitably introduces artifacts such as Gibbs ringing and non-uniformity within the excited slice/slab. These effects are particularly pronounced near the slice/slab boundaries. To mitigate these effects, RF pulses are often modulated with window functions (e.g., Hamming windows) to balance the trade-off between reducing ringing artifacts and maintaining a sharp excitation profile.

Frequency and Phase Encoding

Within the excited slice or slab, the MR signal $S(t)$ with the presence of the gradient field $\mathbf{G}(t)$ can be expressed as:

$$\begin{aligned}
S(t) &\propto \int M_{xy}(\mathbf{r}, t) e^{-i \int_0^t \omega(\tau) d\tau} d\mathbf{r} \\
&= \int M_{xy}(\mathbf{r}, t) e^{-i \int_0^t [\omega_0 + \gamma(\mathbf{G}(\tau) \cdot \mathbf{r})] d\tau} d\mathbf{r} = \int M_{xy}(\mathbf{r}, t) e^{-i\gamma(B_0 t + \mathbf{r} \cdot \int_0^t \mathbf{G}(\tau) d\tau)} d\mathbf{r}.
\end{aligned} \tag{2.8}$$

By defining

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(\tau) d\tau, \tag{2.9}$$

and demodulating the signal at the Larmor frequency ($\omega_0 = \gamma B_0$), Eq. 2.8 becomes

$$S(\mathbf{k}) \propto \int M_{xy}(\mathbf{r}, t) e^{-i2\pi \mathbf{r} \cdot \mathbf{k}(t)} d\mathbf{r}, \tag{2.10}$$

which indicates the Fourier relationship between the transverse magnetisation and the signal formulation defined in Eq. 2.9 (often referred to as “k-space”).

This Fourier relationship leads to several imaging properties of MRI. For example, according to the Nyquist sampling theorem, the sampling rate in k-space needs to be sufficiently dense to reconstruct the image without aliasing. Specifically, for Cartesian sampling, to image an object with length of L along one dimension, the gaps between k-space samples (Δk) must satisfy $\Delta k \leq 1/L$ to prevent aliasing. Therefore, the field-of-view (FOV) of MRI is determined by Δk (i.e., $\text{FOV} \propto 1/\Delta k$) and signal outside FOV will lead to aliasing artifacts. The image spatial resolution along a dimension (Δx) depends on the maximal extent of the k-space, assuming sampling is performed around $k_x = 0$ and extends to $\pm k_x^{\text{max}}$: $\Delta x = 1/2k_x^{\text{max}}$. A broader coverage of k-space corresponds to a higher spatial resolution (i.e., smaller Δx) in MRI.

For a typical 2D Cartesian acquisition where $\mathbf{r} = [x, y]$, the directions along the x and y axes are usually referred to as the “frequency encoding” and “phase encoding” directions, respectively. These are represented as k_x and k_y in k-space, respectively. The gradient used for frequency encoding is usually called the “readout gradient”, which assigns a unique frequency to each spatial location the readout gradient direction (x). Notably, the readout gradient is turned on during the signal acquisition. Phase encoding resolves the spatial location along the perpendicular

direction (y). This is typically achieved over multiple repeated signal readouts by applying short intervals of gradient before the readout gradients and signal acquisitions, which causes spins at different locations along the phase encoding direction to accumulate different initial phases.

K-Space Trajectories

Equation 2.9 shows that k-space trajectories are fundamentally the integral of the gradient waveforms. Figure 2.4 presents three common k-space trajectories for 2D acquisitions and the gradient-encoding schemes to achieve them. For the basic Cartesian raster trajectory, each k-space line can be acquired within one excitation using different phase encoding gradient strengths (G_y) to acquire different phase encoding lines in k-space (Fig. 2.4a). However, the large number of excitations needed for covering numerous phase encoding lines renders this acquisition method rather slow.

A faster version of the Cartesian raster trajectory is Echo Planar Imaging (EPI) trajectory (Fig. 2.4b). EPI collects multiple phase encoding lines in a single excitation by alternating the polarity of the readout gradient between successive lines and using “blips” (short positive gradients in Fig. 2.4b) for the phase encoding gradients. EPI can be further categorised into single-shot and multi-shot approaches. Single-shot EPI traverses the entire k-space in one acquisition, while multi-shot EPI splits the k-space acquisition over several excitations. Single-shot EPI is highly efficient and robust to motion artifacts, making it ideal for applications such as diffusion MRI and functional MRI.

However, single-shot EPI is prone to distortion artifacts (Fig. 2.5). These artifacts arise from local inhomogeneity (e.g., at tissue/air boundaries) of the main magnetic field \mathbf{B}_0 . Field inhomogeneity causes variations in the resonance frequencies of nearby spins, resulting in misplaced signals that manifest as image distortions. These distortions are most apparent along the phase encoding direction, where the encoding gradient is weak (Fig. 2.4b) and the frequency differences due to field inhomogeneity are relatively large. In contrast, along the frequency

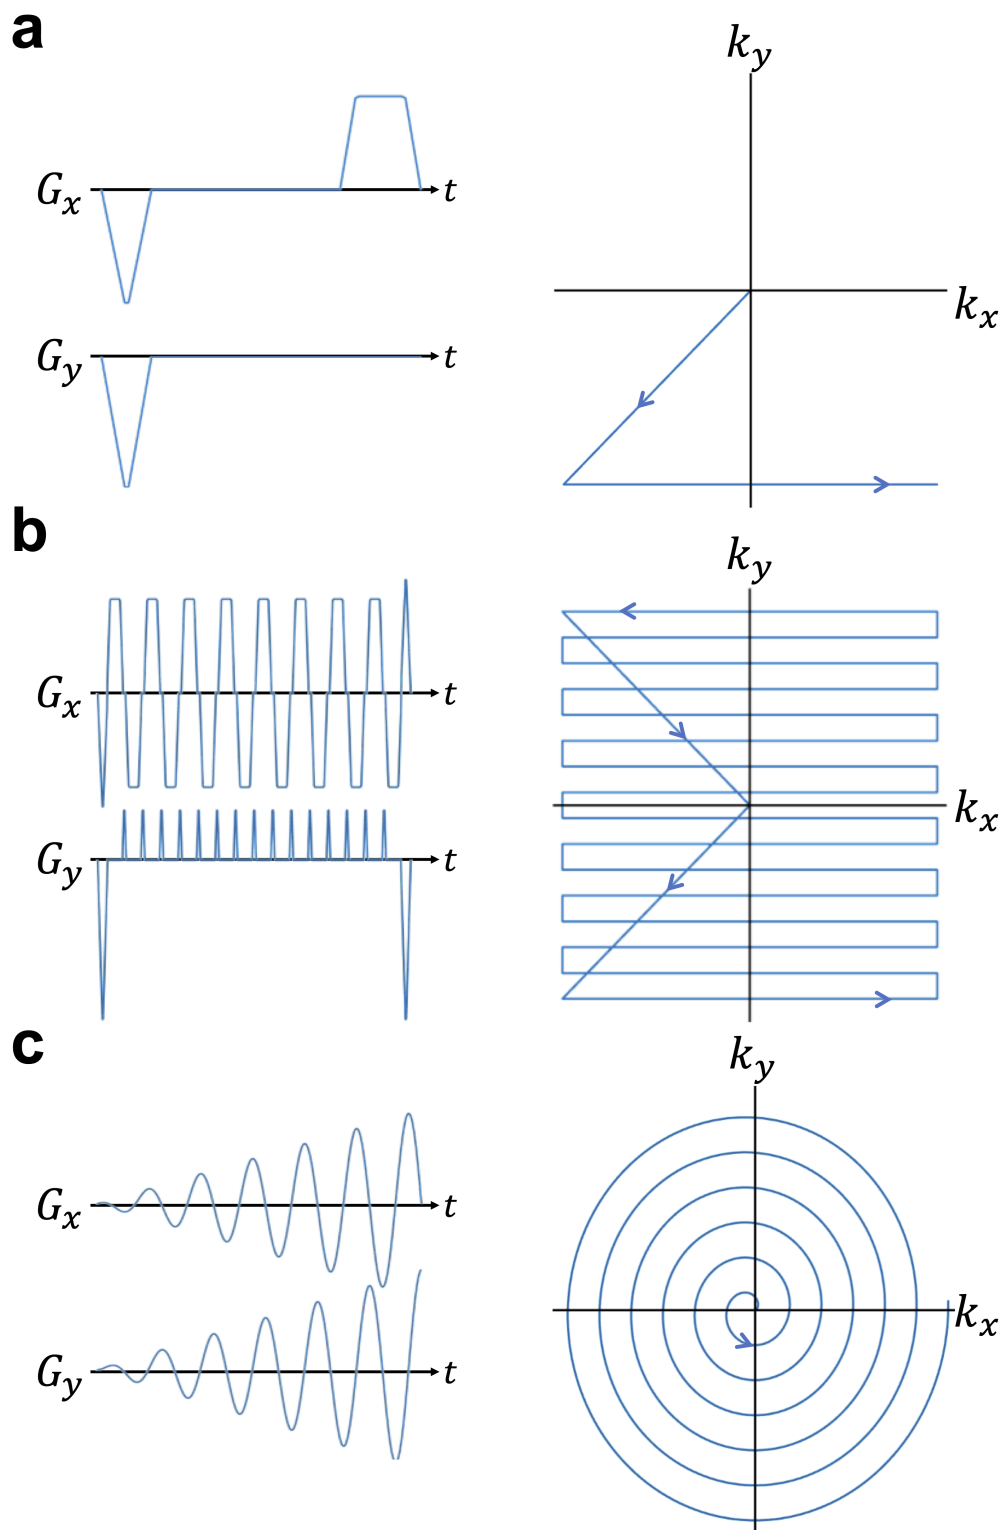


Figure 2.4: Gradient encoding schemes (left) and k-space trajectories (right). (a) One line of basic Cartesian raster trajectory. (b) EPI trajectory. (c) Spiral trajectory.

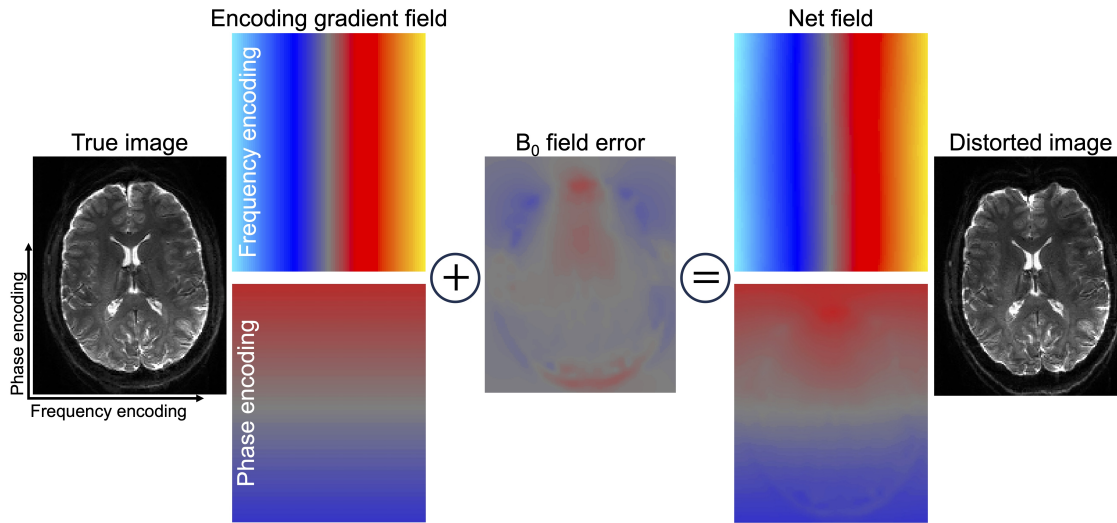


Figure 2.5: Distortion artifacts in EPI. The distortions arise from \mathbf{B}_0 field inhomogeneities. As the phase encoding gradient field for EPI is relatively weak compared to local field changes, distortions are severe along this direction. EPI frequency encoding field is much stronger than local field changes and the image distortion along the readout direction is negligible.

encoding direction, where the encoding gradient is strong, the distortion is usually negligible. The scale of distortion depends on the speed of k-space traversal, with slower traversal (weaker encoding gradients) leading to more pronounced distortions.

In addition to distortion, single-shot EPI suffers from T_2 and T_2^* blurring, which arises from signal decay at the outer regions of k-space, where high-frequency information is encoded. This blurring is more pronounced in high-resolution acquisitions with large matrix sizes and long readout times.

Distortion and blurring artifacts generally scale with the durations of the readout. Longer readouts result in slower traversal of k-space in the phase encoding direction, leading to greater susceptibility to field inhomogeneity-induced distortions and increased T_2/T_2^* signal decay. A straightforward approach to mitigate these artifacts is to shorten the readout time. This can be achieved through parallel imaging, which speeds up k-space traversal by under-sampling (i.e., acquiring fewer phase encoding lines) (see Section 2.1.3 for details), or by using multi-shot acquisitions that divide k-space into segments, thereby shortening the readout per shot. A widely employed multi-shot acquisition is segmented EPI, where k-space is sampled in multiple interleaved segments to ensure complete coverage. Each segment is

under-sampled along the phase encoding direction, resulting in a shortened readout duration and reduced distortion and blurring.

Non-Cartesian trajectories, such as the spiral trajectory (Fig. 2.4c), sample k-space differently. Spiral trajectories start from the centre of k-space and move outward while rotating about the k-space origin. Unlike Cartesian sampling, non-Cartesian techniques do not rely on conventional phase encoding. Instead, a 2D readout function is designed to assign a time-varying frequency at each spatial location throughout the trajectory. Reconstructing images from non-Cartesian data commonly involves resampling onto a Cartesian grid. This process is known as “gridding”.

It is also worth noting that these k-space trajectories can be adapted for 3D acquisitions by incorporating an additional phase encoding gradient along the slice selection direction (z). Short gradient pulses (commonly called k_z -blips) can be similarly applied along the z -axis before the readout to achieve spatial encoding in this third dimension.

2.1.3 Parallel Imaging

The slow acquisition speed is one of the major hurdles for the wider adoption of MRI. To address this challenge, modern MRI systems are often equipped with multiple receive coils, each with spatially distinct but complementary sensitivity profiles. This configuration allows for accelerated MRI scanning by acquiring only a subset of k-space data. The missing k-space data can then be estimated from the acquired data, leveraging the sensitivity information provided by multiple coils. This approach is known as parallel imaging. Four widely adopted parallel imaging reconstruction methods are introduced below: sensitivity encoding (SENSE) [21], generalised autocalibrating partially parallel acquisitions (GRAPPA) [22], iterative self-consistent parallel imaging reconstruction (SPIRiT) [23], and Eigenvector-based SPIRiT (ESPIRiT) [24].

SENSE and ESPIRiT

SENSE provides the fundamental signal model for image-based parallel imaging reconstruction, while ESPIRiT demonstrates robust performance in accurately estimating sensitivity maps essential for SENSE reconstruction. In the SENSE model, the acquired signal y is related to the image signal x through a coil sensitivity matrix S , Fourier transform operator F , and k-space sampling matrix Ψ :

$$\Psi FSx = y, \quad (2.11)$$

Given Ψ , S , y , the unknown image can be estimated by solving:

$$x = \arg \min_x \|\Psi FSx - y\|_2^2, \quad (2.12)$$

which has a closed form solution $x = (A^H A)^{-1} A^H y$, where A is the forward model ($A = \Psi FS$), $(A^H A)^{-1} A^H$ is the pseudoinverse of A , and H is the transpose conjugate. However, directly computing the pseudoinverse can be computationally expensive for non-Cartesian trajectories. To address this, iterative methods like the conjugate gradient algorithm are often employed to solve Eq. 2.12 more efficiently [25], which avoid the explicit inversion and significantly reduce computational complexity. Regularisations (e.g., sparsity [26, 27] and low-rank [28, 29] constraints) can be incorporated into the optimisation to further improve reconstruction performance. While regularisation can also be applied to the closed-form pseudoinverse solution [30], this typically comes at the cost of greater computational demand compared to the iterative methods.

Figure 2.6 provides a simplified illustration of SENSE reconstruction with two receive coils and uniform under-sampling with a factor of two (often denoted as $R=2$). The measured signals from two coils, y_1 and y_2 , are the superposition of voxel intensities (x_1 and x_2) weighted by the respective coil sensitivity profiles $S_{1,1}$, $S_{1,2}$, and $S_{2,1}$, $S_{2,2}$. Specifically, $y_1 = S_{1,1}x_1 + S_{1,2}x_2$ and $y_2 = S_{2,1}x_1 + S_{2,2}x_2$. SENSE reconstruction separates x_1 and x_2 from these combined signals by solving a system of linear equations using the known sensitivity profiles. Modern

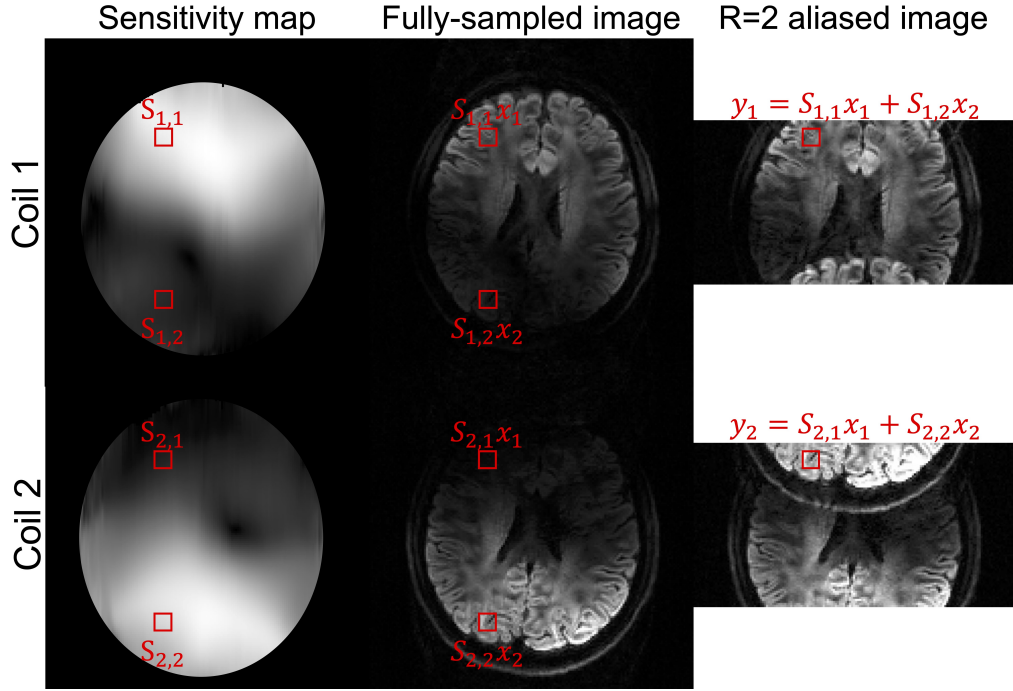


Figure 2.6: Illustration of SENSE reconstruction with 2 receive coils and R=2 data. $S_{i,j}$ is the coil sensitivity of j^{th} voxel in the i^{th} coil, x_j and y_j are the desired full image signal and the acquired image signal of j^{th} voxel, respectively.

MRI systems are often equipped with coil arrays containing 4 to 64 receive coils, depending on the target anatomy. For brain imaging, 32- or 64-channel receive coils are commonly used, theoretically enabling higher under-sampling factors. However, the maximum achievable acceleration factor in practice is usually limited to R=4-6 due to SNR reduction.

The SNR reduction from parallel imaging acceleration has two major sources: the intrinsic \sqrt{R} SNR loss due to fewer acquired measurements (i.e., k-space samples) and the noise amplification from solving the inverse problem in Eq. 2.12. The noise amplification is usually characterised by g-factor ($g = \frac{1}{\sqrt{R}} \frac{\text{SNR}_{\text{full}}}{\text{SNR}_{\text{under}}}$), where SNR_{full} and $\text{SNR}_{\text{under}}$ denote the SNR of fully- and under-sampled images, respectively. A higher g-factor represents stronger noise amplification. The g-factor depends on the conditioning of the forward model $A = \Psi FS$, where an ill-posed inversion of A leads to a strong noise amplification. Thus, both the sampling matrix Ψ and the coil sensitivity profiles S contribute to the g-factor. For SENSE reconstruction with Cartesian uniform under-sampling, the g-factor for n^{th} voxel g_n can be calculated

as $\sqrt{(S\Phi^{-1}S)_{n,n}^{-1}(S\Phi^{-1}S)_{n,n}}$, where Φ is the noise covariance matrix, which can be measured in a pre-scan. Intuitively, the reconstruction suffers from significant noise amplification in regions where coil sensitivity profiles are highly correlated (i.e., when S lacks linear independence due to overlapping sensitivities from multiple coils).

Reducing the g-factor can be achieved by using more receive coils or optimising sampling trajectories to better leverage coil sensitivity variations. A notable example is controlled aliasing in parallel imaging (CAIPI) [31–33]. CAIPI introduces additional gradient blips to modify the k-space trajectory, directing sampling positions to control aliasing patterns and prevent them from overlapping. By shifting aliasing to more efficiently exploit the spatial sensitivity variations of the receive coils, CAIPI improves the conditioning of the reconstruction problem, thereby reducing noise amplification and enhancing SNR.

The performance of SENSE reconstruction depends significantly on the accuracy of the sensitivity maps, which are usually measured from a calibration scan (a fully-sampled central portion of k-space data). Conventionally, the sensitivity maps can be obtained from the low-resolution calibration data by normalising each coil image using the sum-of-squared (SOS) coil-combined image to remove the underlying anatomy while leaving the coil’s unique magnitude and phase profile. A recently developed method ESPIRiT is now a popular choice for sensitivity map estimation thanks to its improved accuracy and robustness. ESPIRiT operates on the observation that k-space blocks of the calibration data for different coils are correlated. If a block-Hankel matrix H_0 is constructed using the calibration data, this correlation implies H_0 has a null space. H_0 can be decomposed into $H_0 = U\Sigma V^H$ using singular value decomposition (SVD), and V can be further separated into V_{\parallel} (the row space of H_0) and V_{\perp} (the null space of H_0). ESPIRiT assumes the k-space data $\hat{x} = FSx$ satisfies consistency with the calibration data, such that:

$$V_{\parallel}V_{\parallel}^H R_r \hat{x} = R_r \hat{x} \text{ or } V_{\perp}^H R_r \hat{x} = 0, \quad (2.13)$$

where R_r is a patch extraction operator from location r . Equation 2.13 can be written as:

$$\mathcal{W}\hat{x} = \hat{x}, \quad (2.14)$$

where $\mathcal{W} = (\sum_r R_r^H R_r)^{-1} \sum_r R_r^H V_{\parallel} V_{\parallel}^H R_r$. With $\hat{x} = FSx$, we obtain:

$$F^{-1}\mathcal{W}FSx = Sx. \quad (2.15)$$

This implies that the vector of coil images is an eigenvector of $F^{-1}\mathcal{W}F$ with an eigenvalue of 1. To obtain the coil sensitivities, SVD is performed on $F^{-1}\mathcal{W}F$. Because \mathcal{W} functions as a convolution operator in k-space, it can be decoupled into point-wise matrix operation \mathcal{G}_q at each position q in the image space (i.e., $F^{-1}\mathcal{W}F|_q = \mathcal{G}_q$). This allows us to simplify Eq. 2.15 into matrix multiplications at each position q : $\mathcal{G}_q S(q)x(q) = S(q)x(q)$. The explicit sensitivity maps can then be derived at locations where $x(q) \neq 0$ by performing an eigenvalue decomposition on each \mathcal{G}_q , selecting only the eigenvectors corresponding to eigenvalues close to 1. Intuitively, an eigenvalue of 1 signifies that the multi-coil data is fully explained by the corresponding eigenvector, making it a reliable sensitivity profile.

GRAPPA and SPIRiT

GRAPPA and SPIRiT are k-space-based reconstruction methods which leverage the correlations between different coils in k-space data to fill in the unacquired k-space points for each coil.

GRAPPA aims to recover the unacquired k-space data in each coil using a linear combination of acquired neighbouring k-space data from all coils (Fig. 2.7a). The synthesis of the k-space data $\hat{x}_i(r)$ for i^{th} coil at location r can be described as:

$$\hat{x}_i(r) = \sum_j g_{r\ ji}^* (\tilde{R}_r \hat{x}_j), \quad (2.16)$$

where $g_{r\ ji}$ is a vector set of weights (known as a kernel) obtained by calibration for the particular sampling pattern around position r between coils i and j , $g_{r\ ji}^*$ is transpose conjugate, and \tilde{R}_r selects only the acquired k-space locations in the neighbourhood of position r . $g_{r\ ji}$ can be trained by minimising $\|\sum_j g_{r\ ji}^* (\tilde{R}_\rho \hat{x}_j) - \hat{x}_i(\rho)\|_2^2$ where $\rho \subseteq$ calibration data.




Figure 2.7: Illustration of GRAPPA and SPIRiT reconstruction. GRAPPA (a) synthesises missing k-space points from neighbouring acquired data while SPIRiT (b) uses all data points in an iterative manner to approach the desired full k-space.

The original GRAPPA algorithm [22] is constrained by its requirement for a fixed under-sampling pattern in the k-space neighbourhood of the missing points. This restriction limits the flexibility of GRAPPA for accommodating various sampling patterns. Subsequent implementations (e.g., <https://people.eecs.berkeley.edu/~mlustig/Software.html>) have supported arbitrary Cartesian sampling by training different kernels $g_{r,ji}$ for k-space neighbourhoods r with distinct sampling patterns, which is, however, quite time-consuming. Furthermore, unlike SENSE (Eq. 2.12), incorporating image priors or adjusting the forward model in GRAPPA reconstruction remains challenging, further restricting its applicability across a broader range of applications.

To address these challenges, SPIRiT has been proposed to allow for a more general and flexible k-space-based reconstruction. Specifically, SPIRiT formulates the reconstruction as an inverse problem and approaches the desired full k-space in an iterative manner by enforcing consistencies with the calibration data and with the acquired data.

For the consistency with the calibration data, unlike GRAPPA which only enforces consistency between missing points and the acquired points, SPIRiT enforces consistency between *every* point in the associated neighbourhood (Fig.

2.7b). This formulation also allows SPIRiT to accommodate arbitrary sampling patterns. The synthesis of $\hat{x}_i(r)$ can be expressed as:

$$\hat{x}_i(r) = \sum_j g_{ji}^*(R_r \hat{x}_j), \quad (2.17)$$

where g_{ji} is a full kernel independent of the k-space sampling pattern and is the same for all k-space positions, and R_r selects the whole k-space neighbourhood of position r . Equation 2.17 can be simply written as:

$$\hat{x} = G\hat{x}, \quad (2.18)$$

where G is a series of convolution operators that convolve the entire k-space with the appropriate calibration kernels, which can be trained similarly to a GRAPPA kernel using the calibration data. Note that this calibration consistency is similar to that in ESPIRiT (Eq. 2.14). SPIRiT uses a different operator G and enforces consistency iteratively by minimising $\|(G - I)\hat{x}\|_2^2$.

The consistency with the acquired data is enforced by minimising $\|\Psi\hat{x} - y\|_2^2$. Taken together, SPIRiT formulates the reconstruction problem as:

$$\hat{x} = \arg \min_{\hat{x}} \|\Psi\hat{x} - y\|_2^2 + \lambda \|(G - I)\hat{x}\|_2^2, \quad (2.19)$$

where λ controls the weight of the calibration consistency term. Similar to Eq. 2.12, Eq. 2.19 can be solved using a conjugate gradient algorithm, and regularisations can be incorporated to further improve the reconstruction. It is also feasible to integrate off-resonance corrections into the reconstruction by altering the data-consistency term $\|\Psi\hat{x} - y\|_2^2$.

As k-space-based reconstruction methods, GRAPPA and SPIRiT implicitly leverage the coil information and are advantageous in regions where accurate explicit sensitivity maps are challenging to obtain due to low SNR or subject motion [34, 35]. The reconstruction produces images for individual coils, after which a final image can be obtained by combining all channels of $F^{-1}\hat{x}$ using SOS or sensitivity maps. Notably, g-factors can also be calculated for such k-space-based

reconstruction methods [36], which has been found generally lower than those of image-based reconstruction methods, likely due to the implicit conditioning of noise for k-space-based reconstruction [37].

2.2 Diffusion MRI

This section begins by delving into the fundamental principles of diffusion MRI (dMRI), followed by an overview of the commonly used acquisition and reconstruction methods for dMRI.

2.2.1 Introduction to Diffusion MRI

This introduction is largely based on *Diffusion MRI: From Quantitative Measurement to In vivo Neuroanatomy* [38].

According to Fick’s first law, in regions where the concentration of molecules varies, particles naturally move from areas of high concentration to areas of low concentration, creating a net flux. The rate of concentration change is proportional to its second spatial derivative, scaled by a constant known as the diffusion coefficient, which depends on factors such as temperature, molecular size, and the surrounding microenvironment. This process, called *diffusion*, is a fundamental mass transport mechanism that facilitates the mixing of molecules or particles without requiring bulk motion. Later, Brown observed that diffusion also occurs at the molecular level even in the absence of a concentration gradient, in a state of thermodynamic equilibrium, which is now commonly referred to as self-diffusion. This molecular motion, driven by random collisions between atoms or molecules in a liquid or gas, is known as “Brownian motion”. In the early twentieth century, Einstein unified the diffusion phenomena described by Fick’s law and Brownian motion, deriving an explicit relationship between the mean squared displacement of a particle and the diffusion coefficient defined in Fick’s law:

$$\langle d^2 \rangle = 2D\Delta, \quad (2.20)$$




Figure 2.8: Diffusion preparation using the pulsed gradient spin echo sequence. δ denotes the diffusion encoding pulse duration, Δ is the time between the application of the two gradient pulses (diffusion time), and G_d is the diffusion encoding gradient strength.

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

Measuring diffusion becomes particularly interesting in biological tissues, where water diffusion is shaped by complex cellular structures. Unlike relaxation times, which reflect the chemical environment, diffusion measurements provide insights into microstructural properties that affect water movement, offering a promising method for studying cellular architecture and pathology.

The capability of measuring the effects of random molecular spreading caused by diffusion using a spin echo MR sequence was first demonstrated by Carr and Purcell. They proposed adding a constant magnetic gradient field, known as a “diffusion encoding gradient”, during the spin echo sequence to amplify the de-phasing effects of spin diffusion. Modern diffusion measurements by MRI were later made possible by Stejskal and Tanner, who introduced the pulsed gradient spin echo (PGSE) sequence (Fig. 2.8). The PGSE sequence replaced Carr and Purcell’s constant magnetic field with a pair of short-duration diffusion gradient pulses: one applied between the 90° excitation and the 180° refocusing pulses, and the other applied between the 180° refocusing pulse and the spin echo.

These diffusion gradients induce phase shifts in the spins of water molecules based on their positions. If a molecule remains stationary between the two gradient pulses, the induced phase shifts cancel out, leading to minimal signal loss. However,

if the molecule diffuses during this time, the phase shifts do not fully cancel, resulting in dephasing and signal attenuation. This attenuation depends on factors such as the strength and duration of the diffusion gradients, the time between the two gradient pulses (diffusion time), and the diffusion coefficient of the medium. The signal attenuation factor E after applying PGSE is characterised by:

$$E = e^{-bD} = e^{-(\gamma\delta G_d)^2 D(\Delta - \delta/3)}, \quad (2.21)$$

where $b = (\gamma\delta G_d)^2(\Delta - \delta/3)$ is called the b -value, which characterises the induced sensitivity to diffusion, δ denotes the diffusion encoding pulse duration, Δ is the diffusion time, and G_d is the diffusion encoding gradient strength (Fig. 2.8). Notably, the orientation of G_d can be flexibly controlled in three dimensions (x, y, z) on modern MRI scanners, which operate with three gradient axes. As the measured diffusion in biological tissues is affected by the complex cellular environment, D is commonly referred to as *apparent* diffusion coefficient (ADC), rather than the true, unrestricted diffusion seen in homogeneous media.

To characterise anisotropic diffusion — where, in certain biological tissues, diffusion is more restricted in some directions than others due to the underlying cellular structure, in contrast to isotropic self-diffusion — Stejskal introduced *tensor*, a 3×3 matrix that represents the natural orientation of anisotropic diffusion with respect to three directions (x, y, and z):

$$D = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix}, \quad (2.22)$$

where the diagonal elements represent the variances along the x, y, and z axes, and the off-diagonal elements denote the covariances, which are symmetric (i.e., $D_{xy} = D_{yx}$, $D_{xz} = D_{zx}$, $D_{yz} = D_{zy}$). This tensor can be visualised as an ellipsoid, with the direction and length of the ellipsoid's axes determined by the eigenvectors ($\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$) and eigenvalues ($\lambda_1, \lambda_2, \lambda_3$) of the matrix D , respectively. Building on this concept, Basser et al. developed a general scheme known as diffusion tensor imaging (DTI) to achieve voxel-wise measurement of the diffusion tensor, which

requires the acquisition of at least six diffusion weighted images with non-coplanar diffusion encoding directions and one $b=0$ image.

Once the diffusion tensor is known, some useful quantitative metrics can be derived. Mean diffusivity (MD) and fractional anisotropy (FA) are two widely used parameters derived from DTI. $MD = \frac{1}{3} \sum_i \lambda_i$ reflects the average magnitude of molecular displacement by diffusion. $FA = \sqrt{\frac{3}{2}} \sqrt{\frac{\sum_i (\lambda_i - MD)^2}{\sum_i \lambda_i^2}}$ quantifies the degree of anisotropy in diffusion, ranging from 0 (isotropic diffusion, such as free water) and 1 (infinite anisotropic diffusion). MD is one of the most widely used metrics in clinical practice, particularly for diagnosing conditions such as brain ischaemia and tumours [39]. FA is a valuable biomarker in clinical research as it shows high sensitivity under certain pathological conditions. Another important metric is the direction of the eigenvector corresponding to the largest eigenvalue (often denoted as V1), which estimates the fibre orientation within each voxel. By tracking these white matter fibre pathways, known as tractography, structural connections between different brain regions can be constructed.

2.2.2 Diffusion MRI Acquisition Methods

Below is an overview of EPI-based 2D and 3D dMRI acquisition methods. It is worth noting that, in addition to EPI, alternative acquisition techniques such as PROPELLER [40], and spiral imaging [41, 42] have also been employed for dMRI. However, these methods are beyond the scope of this thesis.

Two-Dimensional Acquisition

One of the most widely used acquisition methods for dMRI is 2D single-shot EPI [3]. This method captures an entire 2D plane in a single EPI readout, typically taking 70-150 ms per slice, including diffusion preparation, for a common resolution of ~ 2 mm. The full set of slices for a whole-brain scan is acquired within a TR of 10-15 seconds. The rapid acquisition speed and single-shot nature of 2D EPI also make it highly robust against motion artifacts.




Figure 2.9: SNR efficiency comparison for spin echo diffusion acquisitions. For spin-echo diffusion MRI, the optimal SNR efficiency ($\text{SNR}/\sqrt{\text{TR}}$) for white matter is achieved when $\text{TR}=1\text{-}2$ seconds, which can be delivered by 3D multi-slab acquisition.

However, 2D single-shot EPI suffers from several limitations for dMRI acquisition [3]. In addition to its susceptibility to distortion and T_2^* blurring, as discussed in Section 2.1.2, 2D single-shot EPI suffers from limited SNR efficiency for dMRI due to its long TR required for covering a large number of slices in a volume. The SNR efficiency is defined as $\text{SNR}/\sqrt{\text{TR}}$, which captures the trade-off between image quality and acquisition time, making it a more effective indicator of overall imaging performance compared to simple SNR. For spin echo diffusion acquisition of the brain white matter, the optimal SNR efficiency ($\propto \frac{1-e^{-\text{TR}/T_1}}{\sqrt{\text{TR}}}$) is achieved when $\text{TR}=1\text{-}2$ seconds (Fig. 2.9).

Furthermore, the ability of 2D single-shot EPI to achieve high-resolution dMRI is intrinsically limited. In addition to the T_2^* blurring that compromises in-plane resolutions, the through-plane resolution (i.e., slice thickness) is limited by the challenges in achieving high slice selection gradient strengths required to resolve thin slices (≤ 1 mm). Moreover, reducing the slice thickness requires a larger number of slices and thus a longer TR, further compromising the SNR efficiency.

More advanced 2D acquisition schemes have been developed for dMRI to address these challenges. Simultaneous multi-slice (SMS) techniques [32, 43, 44] have been introduced to shorten the TR and improve the SNR efficiency. SMS achieves this by exciting multiple slices simultaneously in a single acquisition using specially designed RF pulses that deposit energy at multiple frequency bands. During reconstruction, the excited slices are separated based on coil sensitivity variations along the slice direction. The RF pulses used for SMS are commonly referred to as “multiband” (MB) pulses, and the number of simultaneously excited slices is denoted as the MB factor. SMS enables TR and total scan time reduction by the MB factor (typically 2-3 for dMRI), thereby boosting SNR efficiency (Fig. 2.9). However, the TR for 2D SMS acquisitions typically remains longer than 4 seconds [3, 45], suggesting further opportunities for increasing SNR efficiency.

Multi-shot acquisitions have been utilised for 2D dMRI to reduce the distortion and blurring artifacts and enable higher in-plane resolutions. Segmented EPI introduced in Section 2.1.2 can shorten the effective echo spacing (reduced distortion due to effectively stronger phase encoding gradients), readout time (reduced blurring due to less T_2^* decay), and echo time (improved SNR due to less T_2 decay). Another widely used multi-shot dMRI acquisition method is readout-segmented EPI [46], which acquires concatenated k-space segments along the readout direction in a mosaic pattern, resulting in similarly reduced effective echo spacing, readout time, and echo time. However, because diffusion preparation gradients encode tiny molecular motions in signal phase, even minor physiological motions (e.g., cardiac pulsations) introduce phase variations across different shots. If not corrected properly, these inconsistencies can significantly degrade image quality. Methods to address phase inconsistencies during image reconstruction will be discussed in Section 2.2.3.

Recent advances have enabled 2D dMRI acquisitions with thin reconstructed slices and high isotropic resolutions with the use of super-resolution techniques. For instance, generalised slice-dithered enhanced resolution (gSlider) [47] excites a thick slice with thickness s multiple times (denoted as N) using N distinct RF pulses with minimally overlapped and complementary spatial profiles. Using the

acquired signal along with the known RF pulse profiles (represented by an RF encoding matrix), thin slices with thickness s/N can then be resolved by solving a system of linear equations. Methods that achieve super-resolution dMRI using multiple thick-slice acquisitions with rotated FOV have also been proposed [48–50]. However, these super-resolution approaches may be susceptible to blurring due to the use of regularisations and error propagation in solving the inverse problem. They also face the intrinsic limitations of 2D acquisitions, such as suboptimal SNR efficiency resulting from long TR.

Three-Dimensional Acquisition

The primary advantages of 3D acquisition methods include the ability to shorten the TR, thereby enhancing SNR efficiency, and the capability to achieve thin slice thickness through gradient encoding in 3D k-space. Gradient encoding offers orthogonal Fourier bases for resolving thin slices, which improves SNR and produces high-fidelity voxel shapes. Typically, 3D acquisitions add a second phase-encoding along the slice selection direction, forming a 3D k-space. Multi-shot EPI is often used to cover this 3D k-space efficiently.

While 3D EPI has demonstrated success in achieving high-resolution, high-SNR dMRI in *post-mortem* brain imaging [2], its application in *in-vivo* settings presents challenges. As mentioned earlier, multi-shot dMRI necessary for 3D acquisitions is sensitive to motion-induced phase error, necessitating precise phase correction when combining data from different shots. Accurately capturing phase errors that vary spatially in all three dimensions is exceedingly difficult to achieve in practice (see Section 2.2.3). Additionally, efficient 3D single-slab acquisitions typically have a TR of less than 100 ms, which is suboptimal for spin echo dMRI’s SNR efficiency in *in-vivo* brain imaging (Fig. 2.9).

To address these challenges, 3D multi-slab acquisitions have been developed [4–7]. This technique divides the brain into multiple thin slabs (usually 6 to 15) to facilitate whole-brain coverage. Within each slab, 3D EPI is commonly employed to encode 3D k-space. A key assumption for 3D multi-slab acquisitions is that

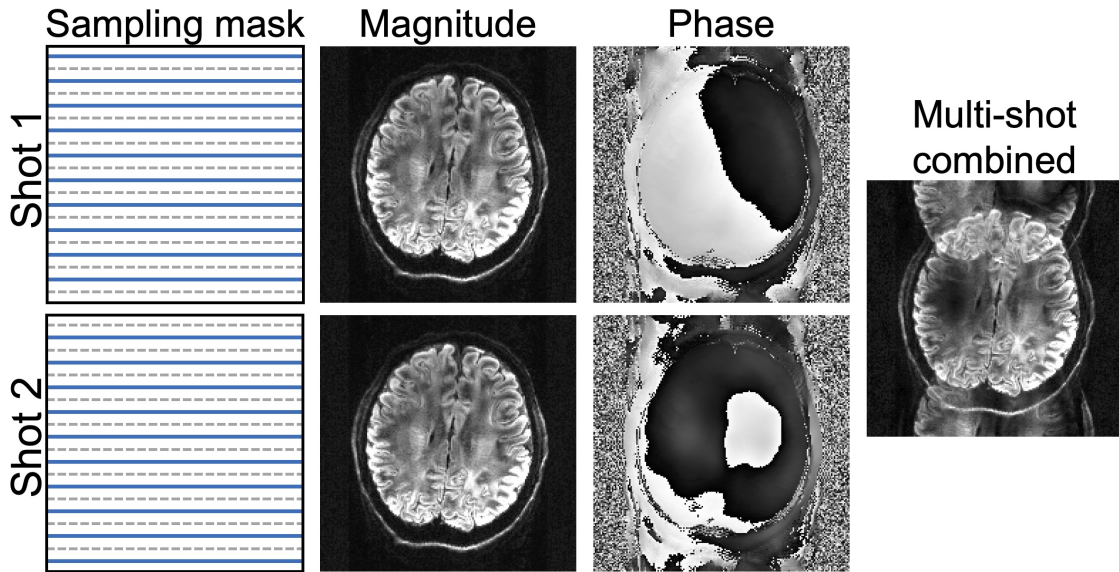


Figure 2.10: Phase inconsistency in multi-shot diffusion MRI acquisitions. A two-shot acquisition is simulated by retrospectively adding realistic motion-induced phase to a fully sampled 2D k-space. In the sampling masks, the acquired and unacquired data are illustrated as blue lines and grey dashed lines, respectively. Two shots share the same underlying magnitude images while having different motion-induced phases.

within thin slabs (≤ 2 cm), motion-induced phase errors along the slice direction vary slowly and can be considered negligible, making 2D phase error correction for 3D multi-slab dMRI sufficient (see Section 2.2.3). Since 2D phase correction is practical, 3D multi-slab dMRI becomes feasible for *in-vivo* imaging. Moreover, for 3D multi-slab acquisitions, the number of excitations within each TR is proportional to the number of slabs. By adjusting the number of slabs, TR can be tuned to the SNR-efficient regime (i.e., 1-2 seconds), as depicted in Fig. 2.9.

2.2.3 Diffusion MRI Reconstruction Methods

The major focus of this section is the correction methods for motion-induced phase variations for multi-shot dMRI reconstruction. For single-shot dMRI reconstruction, motion-induced phase can be simply ignored, and methods introduced in Section 2.1.3 can be readily applied.

As discussed in Sections 2.2.1 and 2.2.2, diffusion gradients make the phase of diffusion images highly sensitive to even small movements. Minor physiological motions, such as cardiac pulsations, can introduce phase variations, leading to

inconsistencies between shots in multi-shot dMRI, which degrade the reconstructed image quality if not corrected (Fig. 2.10). One established approach to correct these phase inconsistencies involves acquiring a navigator after the imaging echo (Fig. 2.11). This navigator echo traverses 2D k -space in a single-shot manner, allowing for the reconstruction of a phase map. Since phase corruption is primarily introduced by the diffusion gradients during the diffusion preparation module, the phase remains mostly stable afterward until another diffusion gradient is applied. As a result, the navigator phase can accurately represent the motion-induced phase, which can then be subtracted from the imaging data or incorporated into a forward model for more effective correction [51]. For pure 3D acquisitions, phase correction poses significant challenges, as it is nearly impossible to use a single-shot 3D navigator to cover the entire 3D volume and provide reliable phase estimation. Nevertheless, in 3D multi-slab acquisitions, 2D navigators sampling the central $k_z = 0$ plane (i.e., a slab-averaged 2D phase map) have demonstrated sufficient accuracy for phase error correction, provided the slabs are sufficiently thin (≤ 2 cm) [4]. However, the additional spin echo required for navigator acquisition extends the TR and scan time, thereby compromising SNR efficiency. Moreover, the use of an additional refocusing pulse increases the specific absorption rate (SAR).

To overcome these disadvantages, phase error correction methods that do not require navigator acquisition have been developed for 2D multi-shot dMRI. Two notable methods are multiplexed sensitivity-encoding (MUSE) [52] and multi-shot sensitivity encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS) [53].

MUSE reconstructs a 2D phase map from each under-sampled shot. Although the reconstruction for each shot might suffer from noise amplification due to the high under-sampling factor, motion-induced phase variations, which are expected to be spatially smooth, can still be reliably estimated using appropriate smoothing and denoising operators (e.g., total variation). In the original work [52], MUSE has demonstrated robustness in 4-shot dMRI reconstruction. However, MUSE may

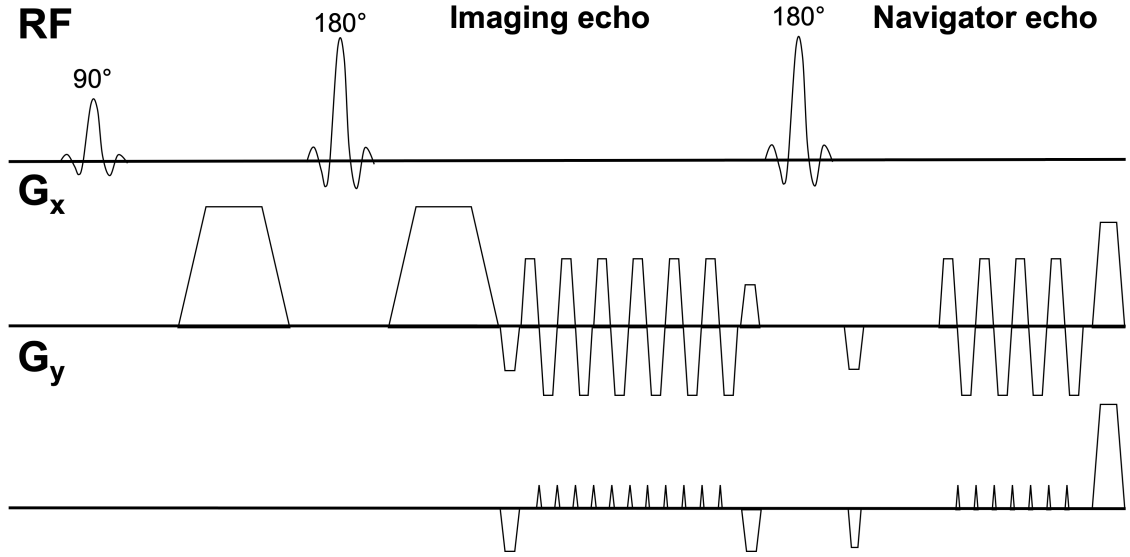


Figure 2.11: Sequence diagram for navigated diffusion MRI acquisitions. A second refocusing pulse is added to form a navigator echo to record the phase variations after the imaging echo. Here, the diffusion encoding is applied along the readout direction (G_x).

encounter difficulties when the under-sampling factor R of each shot is too large (e.g., $R > 6$) due to the inherent limitations of parallel imaging reconstructions.

MUSSELS leverages the magnitude consistency between shots to create a structured Hankel matrix and estimate the missing data in each shot using structured low-rank (SLR) matrix completion. It assumes different shots share the same underlying magnitude image m despite their different phases: $m = x_i \Phi_i^H = x_j \Phi_j^H$, where x_i, x_j and Φ_i, Φ_j denote the complex images and phase maps of the i^{th} and j^{th} shot, respectively, and Φ^H denotes the conjugate of Φ . An annihilation relation can be derived after applying the Fourier transform:

$$\hat{x}_i * \hat{\Phi}_j - \hat{x}_j * \hat{\Phi}_i = 0, \quad (2.23)$$

where \hat{x} and $\hat{\Phi}$ are the Fourier transforms of x and Φ , respectively, and $*$ is the convolution operation. Because the phase maps Φ are spatially smooth with limited k-space support, a block-Hankel matrix $H(\hat{x})$ using the multi-shot k-space data (Fig. 2.12) can be constructed which replaces the convolution operation with matrix multiplication. Since Eq. 2.23 holds true for all shots, the annihilation relation can be written in a matrix form:

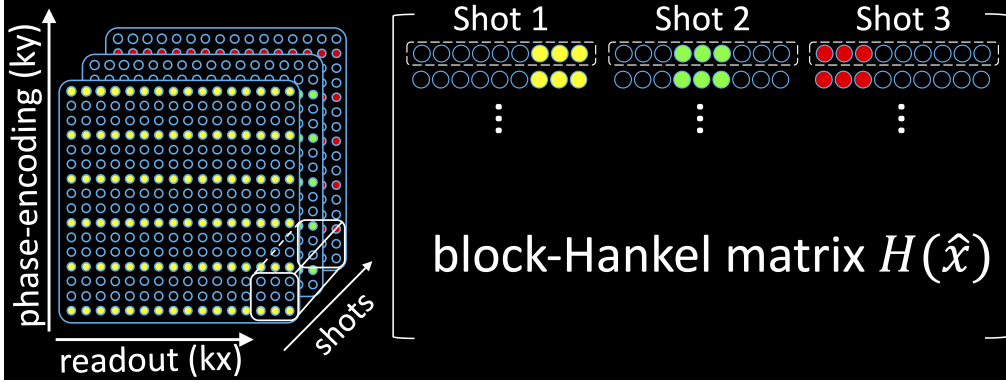


Figure 2.12: Construction of block-Hankel matrix using multi-shot k-space data. The kernel denoted by the white box traverses the whole k-space in a sliding window fashion.

$$\begin{bmatrix} H(\hat{x}_1) & H(\hat{x}_2) & \dots & H(\hat{x}_{N_{\text{shot}}}) \end{bmatrix} \begin{bmatrix} \hat{\Phi}_2 & 0 & 0 & \hat{\Phi}_3 \\ -\hat{\Phi}_1 & \hat{\Phi}_3 & 0 & 0 \\ 0 & -\hat{\Phi}_2 & 0 & -\hat{\Phi}_1 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \hat{\Phi}_{N_{\text{shot}}} & 0 \\ 0 & 0 & -\hat{\Phi}_{N_{\text{shot}}-1} & 0 \end{bmatrix} = \mathbf{0}, \quad (2.24)$$

where N_{shot} is the total number of shots. This indicates the structured matrix $H(\hat{x}) = [H(\hat{x}_1) \ H(\hat{x}_2) \ \dots \ H(\hat{x}_{N_{\text{shot}}})]$ exhibits a low-rank property due to the existence of a non-trivial null space. This low-rank property can be enforced during reconstruction, usually by performing singular value shrinkage to $H(\hat{x})$, to enable an accurate recovery of the multi-shot data without requiring explicit phase maps. The phase maps can also be extracted from the recovered multi-shot data. Because MUSSELS leverages the redundancy in the shot dimension, it enables reconstruction with a higher shot number and a higher under-sampling factor per shot compared to MUSE.

For 3D multi-slab acquisitions, navigator-free imaging is intrinsically more challenging. The challenges have been discussed in detail in Chapter 4, where we develop an acquisition and reconstruction framework that eliminates the requirement for the navigator acquisition and enables self-navigated 3D multi-slab dMRI.

3

Sampling Strategies and Integrated Reconstruction for Reducing Distortion and Boundary Slice Aliasing in High-Resolution 3D Diffusion MRI

Contents

3.1	Introduction	38
3.2	Methods	41
3.2.1	Blip-Reversed Acquisition and Reconstruction	41
3.2.2	<i>In-Vivo</i> Experiments	45
3.2.3	Reconstruction Details	49
3.2.4	Post-Processing	51
3.2.5	Diffusion Analyses	52
3.3	Results	53
3.4	Discussion	64
3.5	Conclusions	70

In this chapter, we develop a new method for high-fidelity, high-resolution 3D multi-slab diffusion MRI (dMRI) with minimal distortion and boundary slice aliasing. Our method modifies 3D multi-slab imaging to integrate blip-reversed acquisitions for distortion correction and oversampling in the slice direction (k_z) for reducing boundary slice aliasing. Our aim is to achieve robust acceleration

to keep the scan time the same as conventional 3D multi-slab acquisitions, in which data are acquired with a single direction of blip traversal and without k_z -oversampling. We employ a two-stage reconstruction. In the first stage, the blip-up/down images are respectively reconstructed and analysed to produce a field map for each diffusion direction. In the second stage, the blip-reversed data and the field map are incorporated into a joint reconstruction to produce images that are corrected for distortion and boundary slice aliasing. We conducted experiments at 7T in six healthy subjects. Stage 1 reconstruction produces images from highly under-sampled data ($R=7.2$) with sufficient quality to provide accurate field map estimation. Stage 2 joint reconstruction substantially reduces distortion artifacts with comparable quality to fully-sampled blip-reversed results ($2.4\times$ scan time). Whole-brain *in-vivo* results acquired at 1.22 mm and 1.05 mm isotropic resolutions demonstrate improved anatomical fidelity compared to conventional 3D multi-slab imaging. Data demonstrate good reliability and reproducibility of the proposed method over multiple subjects. In conclusion, the proposed acquisition and reconstruction framework provide major reductions in distortion and boundary slice aliasing for 3D multi-slab dMRI without increasing the scan time, which can potentially produce high-quality, high-resolution dMRI. This chapter is based on a peer-reviewed journal paper we published in *Magnetic Resonance in Medicine* [54].

3.1 Introduction

Diffusion MRI probes tissue at the microscopic scale non-invasively [55, 56], providing information about healthy and pathological changes to neural architecture. High-resolution dMRI can depict microstructure details in the brain, facilitating tracking of thin fibres and accurate depiction of complex fibre configurations [2, 57, 58]. Because of these advantages, high-resolution dMRI is compelling for neuroscientific research and clinical diagnosis.

Three-dimensional (3D) multi-slab acquisition has great potential to achieve high-resolution *in-vivo* dMRI, which can produce optimal signal-to-noise ratio (SNR) efficiency for spin-echo-based dMRI due to its compatibility with a short $TR=1-2$

s and achieve thin slices using a 3D k-space encoding [4–7, 59]. 3D multi-slab acquisitions divide the whole imaging volume into multiple thin slabs (typically with 10-20 slices in each slab) and encode each slab with a 3D k-space readout, typically using a 3D echo-planar imaging (EPI) trajectory for an efficient acquisition. However, the image quality of 3D multi-slab imaging with EPI-based trajectories is compromised by two major image artifacts: distortion and boundary slice aliasing.

First, similar to conventional 2D EPI, 3D EPI suffers from distortions from the static magnetic field (B_0) inhomogeneity [60]. These distortions occur along the phase-encoding direction due to its low bandwidth, especially near tissue/bone and tissue/air interfaces with large susceptibility differences. EPI distortions are conventionally corrected after image reconstruction using a field map acquired from a separate scan, using either a GRE field mapping sequence [60] or a pair of EPI images acquired with opposite phase-encoding directions [9]. However, field mapping-based corrections are inadequate, particularly if distortions become sufficiently severe that the signal from multiple voxels overlap in the reconstructed image. Moreover, static field maps cannot capture dynamic B_0 field changes due to subject motion, eddy currents, and field drift across different diffusion directions [61].

EPI distortion corrections using a pair of blip-reversed phase-encoding images have also been developed, which have demonstrated superior performance compared to the field mapping-based correction in dMRI [9–12, 14]. While most of these approaches perform distortion correction using separately reconstructed blip-up and blip-down images [9–11], methods that jointly reconstruct blip-up and blip-down data have also been developed [12–14]. The recently proposed BUDA (blip-up/down acquisition) EPI [14] method performed interleaved blip-reversed acquisitions for each diffusion direction and incorporated field maps estimated from separate blip-up/down images into a distortion corrected joint reconstruction. The method demonstrated improved distortion correction, dynamic B_0 field mapping capability, and reduced g-factor penalty due to the combination of blip-up/down data in the joint reconstruction. However, existing blip-reversed acquisition methods require

doubled scan time compared with conventional single-shot EPI using a single phase encoding acquisition.

A second source of artifacts in 3D multi-slab imaging is slice aliasing near the slab boundary. Boundary slice aliasing arises due to the inability to achieve a sharp excitation profile, which means some tissue beyond the targeted slab is excited. The transition bands and sidelobes of the RF profile thus extend to adjacent slabs, leading to aliasing of signal from one end of the slab into the other [62]. To avoid aliasing, conventional methods expand the field-of-view (FOV) along the slice direction for each slab via oversampling or increase the overlapping between adjacent slabs [4]. However, these methods inevitably require additional scan time.

In this chapter, we aim to incorporate the blip-reversed acquisition and k_z oversampling into 3D multi-slab diffusion imaging to minimise distortion and boundary slice aliasing without increasing the scan time. However, achieving this goal is fundamentally challenging. Acquiring both blip-up and blip-down images without increasing the scan time requires an extra two-fold ($2\times$) acceleration for time compensation. The current 3D multi-slab dMRI acquisition typically uses 3D multi-shot EPI with each EPI readout covering a single k_z plane. Because 50% of sequence time is dedicated to diffusion preparation, shortening the readout with in-plane acceleration will not provide significant scan time reduction. Instead, the most effective way to reduce scan time is to accelerate along the slice direction, which determines the number of excitations required to sample along k_z . This is challenging for *in-vivo* diffusion MRI because each slab needs to be sufficiently thin (i.e., 10–20 mm) to ensure that the motion-induced phase variation within each slab can be accurately measured with a 2D navigator [4, 5, 7], which leads to very limited coil sensitivity variation along the slice direction. Therefore, direct under-sampling along the slice direction will result in a highly ill-posed reconstruction. Moreover, to additionally address the boundary slice aliasing, oversampling along k_z is needed, which will lead to an even higher acceleration factor and thus a more challenging reconstruction problem.

We propose an approach for 3D EPI that enables distortion- and boundary slice aliasing-corrected reconstruction similar to BUDA [14] but with no increase in scan time compared to the conventional 3D multi-slab acquisition [4, 5, 7] (Fig. 3.1). We acquire an equal number of shots with blip-up and blip-down phase encoding. The reconstruction proceeds in two stages. In the first stage, we reconstruct the highly under-sampled blip-up and blip-down images separately in order to estimate the field map [9, 63, 64]. In the second stage, we reconstruct a final image that is distortion corrected and has a more modest acceleration using all segments and accounting for the field map. In addition, we enlarge the FOV along the slice direction to reduce the boundary slice aliasing. This approach relies on the highly-accelerated reconstruction for stage 1 being sufficiently robust to accurately estimate the field map. To support this aim, we use a 3D EPI trajectory with blipped CAIPI [32] and partial Fourier along k_z . A high-SNR SPIRiT-based [23] reconstruction is also optimised and employed. Our framework is validated with *in-vivo* experiments on a 7T scanner and achieves whole-brain dMRI at 1.05 mm isotropic resolution. The diffusion analysis and tractography results reveal the higher anatomical fidelity and quantitative accuracy of the proposed method compared to the conventional 3D multi-slab imaging, demonstrating the potential of our method to facilitate high-resolution dMRI with improved image quality to benefit neuroscientific research.

3.2 Methods

3.2.1 Blip-Reversed Acquisition and Reconstruction Acquisition

Integrating 3D multi-slab dMRI with blip-reversed EPI requires the acquisition of two 3D EPI images with reversed phase encodings, which can be achieved at a cost of doubled scan time. One approach to shorten the scan time is to perform under-sampling along k_z . By acquiring half the segments with blip-up phase encoding and the other half with blip-down phase encoding, the total scan time of the integrated blip-reversed acquisition is identical to conventional 3D multi-slab EPI [4, 5, 7]. Importantly, 3D multi-slab diffusion MRI requires very thin slabs to facilitate the

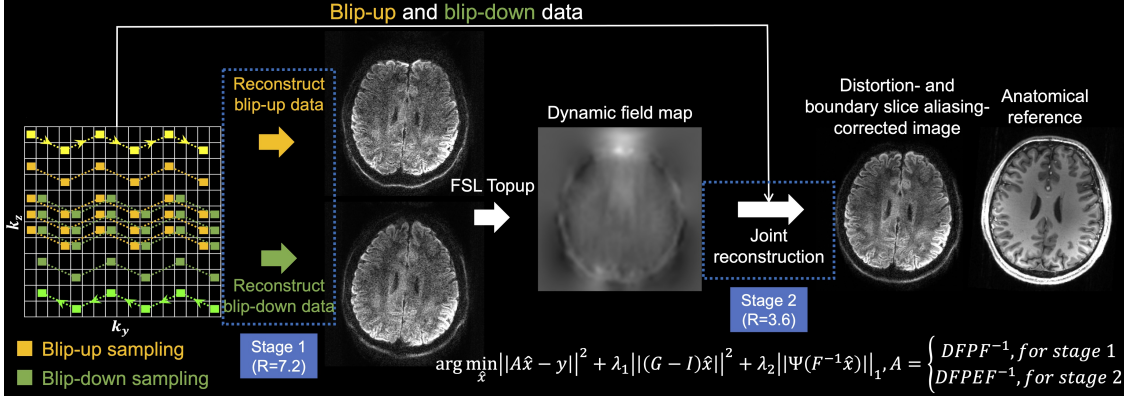


Figure 3.1: Proposed acquisition and reconstruction framework. The proposed sampling pattern with k_z blipped-CAIPI and complementary partial Fourier for blip-up (yellow) and blip-down (green) data and the two-stage reconstruction. The trajectory of one shot of the multi-shot sampling is marked in bright yellow and bright green for blip-up and blip-down sampling with arrows indicating the phase encoding direction. The forward operator in the reconstruction is $A = DFPP^{-1}$ for stage 1 reconstruction where F is Fourier Transform, P is the phase modulation representing motion-induced phase errors measured by 2D navigators, D is the k-space sampling operation and $A = DFPEF^{-1}$ where E is the distortion operation (captured by the field map) for stage 2. Example images are from a single volume diffusion MRI dataset (1.05 mm isotropic resolution) of a representative subject, with anatomical image listed for reference (acquired with MPRAGE at 0.86 mm isotropic resolution).

correction of motion-induced errors. This severely limits the variation of coil profiles along the slab direction, and as a result conventional rectangular under-sampling with integer reduction along k_y and k_z (e.g., Fig. 3.2c, iii) may suffer from a high g-factor penalty and residual aliasing.

Moreover, due to the limited FOV and non-rectangular RF profile along the slice direction, where transition bands and side lobes extend to adjacent slabs (Fig. 3.2a), slice aliasing happens at slab boundaries in 3D multi-slab imaging. The slice aliasing artifacts can be reduced by over-sampling along k_z with an extended FOV (Fig. 3.2b). To extend the FOV without increasing the number of shots (e.g., from 10 shots to 12 shots comparing Fig. 3.2c, i and Fig. 3.2c, ii), a higher acceleration factor along k_z is required, which leads to a more challenging reconstruction.

Here, we combined blip-reversed acquisitions with two sampling strategies to improve the under-sampled reconstruction, which is particularly important for the first stage of reconstruction in which images are separately reconstructed for the

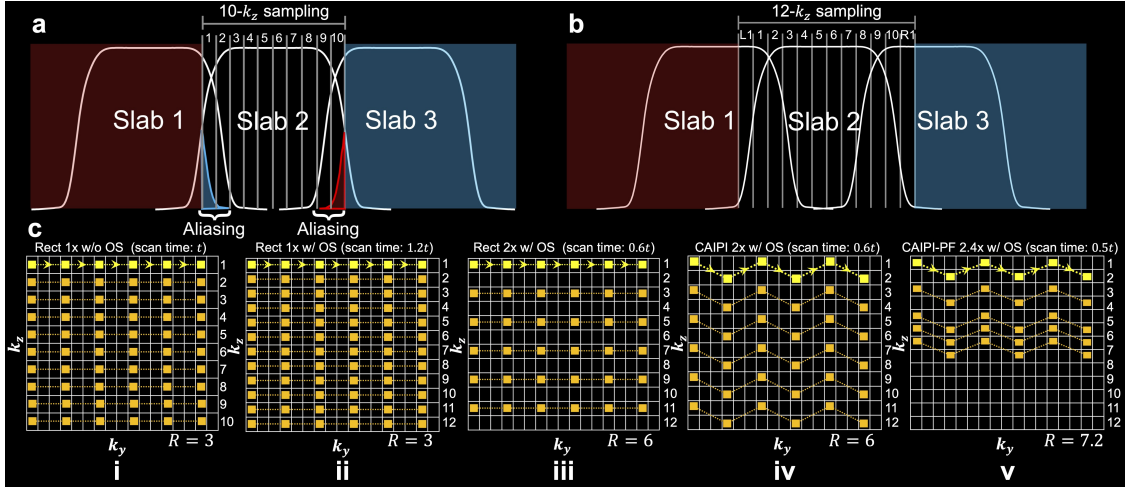


Figure 3.2: 3D multi-slab acquisition and sampling patterns for blip-up data. (a) Demonstration of slice aliasing at slab boundary caused by non-rectangular RF profile and limited FOV along the slice direction for 3D multi-slab acquisition ($10-k_z$ sampling). (b) Correction of slice aliasing by expanding the FOV for each slab through k_z oversampling ($12-k_z$ sampling) (L1 and R1 denote the 1 oversampled slice on the left and right side). (c) Comparison of various sampling patterns: rectangular sampling without oversampling along k_z (Rect $1 \times$ w/o OS) (i); rectangular sampling with 20% oversampling along k_z (Rect $1 \times$ w/OS) (ii); rectangular sampling with 20% oversampling and $2 \times$ acceleration along k_z (Rect $2 \times$ w/OS) (iii); CAIPI sampling with 20% oversampling and $2 \times$ acceleration along k_z (CAIPI $2 \times$ w/OS) (iv); CAIPI-PF sampling with 20% oversampling, partial Fourier and $2.4 \times$ acceleration along k_z (CAIPI-PF $2.4 \times$ w/ OS) (v). A $3 \times$ acceleration along k_y is applied in all sampling patterns. The trajectory of one EPI shot is marked in bright yellow with arrows indicating the phase encoding direction. The parameter t represents the scan time of the acquisition without oversampling or acceleration along k_z for one phase encoding direction (c, i). R is the total under-sampling factor for each sampling pattern.

blip-up and blip-down segments. First, k_z blipped-CAIPI [31, 32] was integrated for more effective use of coil sensitivity, where each shot covers the full extent of the phase encoding direction (k_y) with even spacing along k_y and includes “blips” to traverse closely spaced k_z planes (Fig. 3.2c, iv). We also employed partial Fourier along k_z (Fig. 3.2c, v), which reduces aliasing in the slice direction through a more densely sampled central k-space region. This achieves an even shorter scan time by reducing the number of shots (Fig. 3.2c, v, halved scan time compared to Fig. 3.2c, i). The proposed sampling approach enables simultaneous correction of distortions (by allowing joint acquisition of blip-up/down data) and boundary slice aliasing (by oversampling along k_z) without increasing the scan time. We refer to our sampling approach as “CAIPI-PF” hereafter.

For the joint acquisition of blip-reversed data, equal numbers of shots traversing k_y in opposite directions are acquired. The blip-up and blip-down sampling cover the complementary subsets of the k-space with a shift of Δk_y and complementary partial Fourier along k_z (Fig. 3.1, left), which reduces the noise amplification and prevents resolution loss in the stage 2 joint reconstruction.

Reconstruction

The cost function of the SPIRiT-based regularised reconstruction is:

$$\|A\hat{x} - y\|_2^2 + \lambda_1 \|(G - I)\hat{x}\|_2^2 + \lambda_2 \|\Psi(F^{-1}\hat{x})\|_1, \quad (3.1)$$

where \hat{x} is the multi-coil k-space data of the target aliasing- and distortion-corrected image, A is the forward operator, y is the acquired data, G is the SPIRiT kernel trained on coil calibration data, I is the identity matrix, Ψ is the wavelet operator, F^{-1} is the inverse Fourier Transform, and λ_1 and λ_2 are the parameters for the SPIRiT and sparsity regularisations (the SPIRiT weights for stage 1 and stage 2 are denoted as $\lambda_{1,1}$ and $\lambda_{1,2}$, respectively). The SPIRiT regularisation (λ_1) facilitates parallel imaging by enforcing calibration consistency between every k-space data point and its neighbours [23], while the sparsity regularisation (λ_2) is used to suppress the noise [26].

The forward operator A is constructed differently for the two stages:

$$A = \begin{cases} DFPF^{-1} & \text{for stage 1} \\ DFPEF^{-1} & \text{for stage 2} \end{cases} \quad (3.2)$$

where D is the down-sampling operator in k-space and P represents motion-induced phase errors captured by 2D navigators. The E operator is only applied in stage 2 and represents spatial distortion induced by field inhomogeneity estimated from the stage 1 reconstruction.

It is worth noting that unlike the original SPIRiT formulation [23] in which data consistency can be ensured implicitly by only estimating the missing k-space points, in our work the entire k-space needs to be estimated. This is because in our forward operator A the acquired k-space points are corrupted by the motion-induced phase

P and the distortion-induced displacement E , and are therefore different from those in the uncorrupted target k-space data \hat{x} . Hence, the SPIRiT constraint is explicitly added as a regularisation term and the data consistency term $\|A\hat{x} - y\|_2^2$ imposes the consistency constraint across all k-space points.

3.2.2 *In-Vivo* Experiments

A 3D multi-slab spin-echo diffusion MRI sequence [7] was modified to integrate blipped-CAIPI and k_z partial Fourier sampling. After the imaging echo, a second refocusing pulse was applied to acquire a low resolution 2D navigator to correct the motion-induced phase errors. The sequence diagram of the proposed acquisition is demonstrated in Fig. 3.3. Subjects were scanned on a Siemens 7T scanner (Siemens Magnetom, Erlangen, Germany) using a 32-channel receive coil. Written informed consent in accordance with local ethics was obtained from each subject.

Evaluation Protocols

To evaluate the impact of different sampling patterns and reconstruction parameters, fully sampled single-slab datasets using CAIPI sampling were acquired from a single subject using the following scan parameters: 1.22 mm isotropic resolution, 20 slices, TE1(imaging)/TE2(navigator)/TR=72/128/1800 ms, b=1000 s/mm², diffusion encoding along left–right direction, k_y under-sampling $R_y = 3$, echo spacing 0.76 ms. The 2D navigator acquired 64 phase encoding lines using the same phase-encoding direction, the same k_y under-sampling factor, and the same echo spacing as the imaging echo. An illustration of the navigator acquisition trajectory is showed in Fig. 3.4. To avoid slice aliasing, 1.2× oversampling along k_z was applied, which encoded a larger FOV of 29.3 mm along the slice direction with 24 slices. Blip-up and blip-down phase encoding datasets were acquired separately. Three scans with 0, 1, and 2 Δk_y shift were acquired and combined to produce fully sampled data. In addition, a dataset using conventional rectangular sampling without k_z oversampling was also acquired from the same subject to demonstrate slice aliasing.

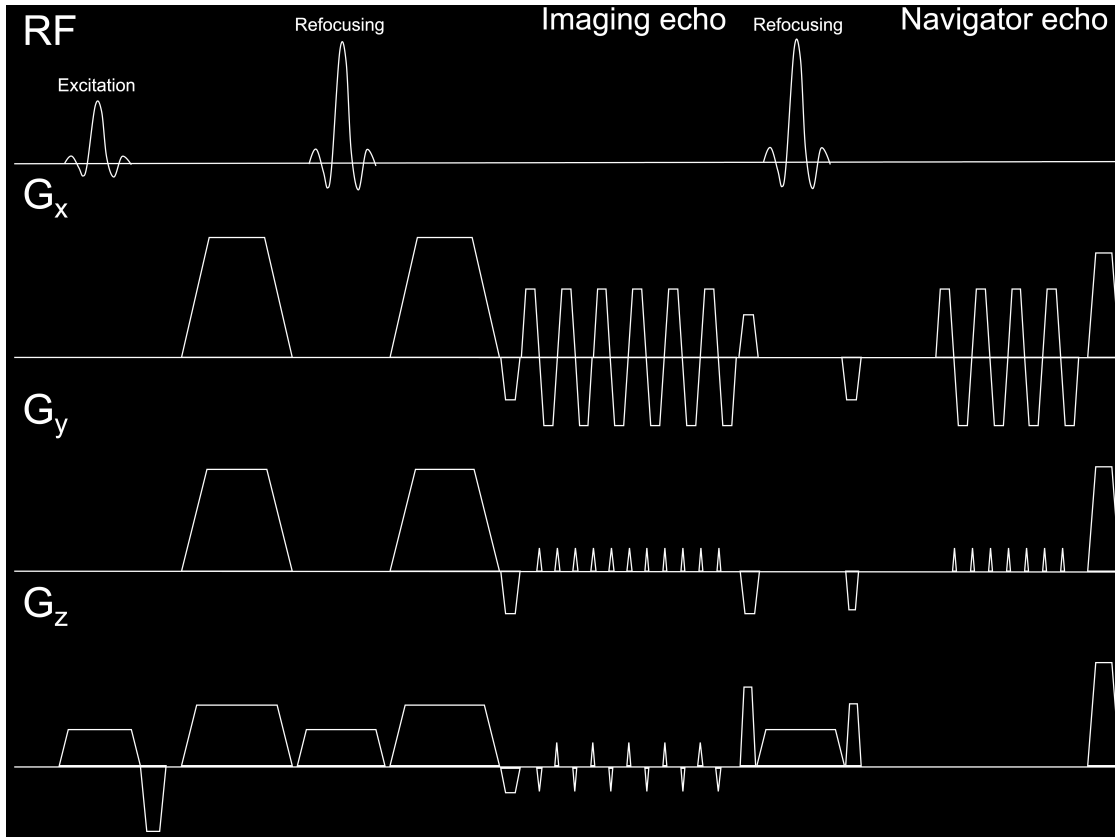


Figure 3.3: Sequence diagram of the proposed 3D multi-slab CAIPI-PF acquisition. The diagrams of the RF pulse, the readout gradient (G_x), the phase-encoding gradient (G_y), and the slice selection gradient (G_z) are displayed. The k_z -blipped CAIPI is achieved by adding slice selection blipped gradient. The 3D imaging echo is followed by a navigator echo used for 2D navigator acquisition.

Retrospective under-sampling was then performed to evaluate different sampling patterns, which were similar to those shown in Fig. 3.2 but with 20 or 24 k_z encoding planes (without or with k_z oversampling). For CAIPI-PF sampling (Fig. 3.2c, v), the shots #1, 3, 5, 7, 9, 11, 12, 13, 14, 15 were used for acquiring blip-up data and the shots #11, 12, 13, 14, 15, 16, 18, 20, 22, 24 were used for acquiring blip-down data (shot #n covers k_z plane n and n+1) (partial Fourier factor: $f_{PF} = 2/3$).

For whole-brain quantitative comparison of the proposed CAIPI-PF sampling (Fig. 3.2c, v) with the conventional rectangular sampling (Fig. 3.2c, i), three subjects were scanned using 1.22 mm isotropic resolution reference protocol, CAIPI-PF sampling protocol, and conventional sampling protocol. For all protocols, six slabs with 20 slices per slab were acquired. neighbouring slabs were overlapped




Figure 3.4: Illustration of 2D navigator acquisition trajectory. The 2D navigator acquires the central k_z plane of each slab after the imaging echo. Only the central part of the 2D k-space (i.e., 64 phase-encoding lines) is acquired with the same phase-encoding direction, the same k_y under-sampling factor ($R_y = 3$), and the same echo spacing as the imaging echo.

by 1 slice, resulting in 115 slices in the final reconstruction. The FOV was $220 \times 220 \times 140 \text{ mm}^3$ and voxel size was 1.22 mm isotropic. Interleaved slab acquisition was used to minimise cross talk between adjacent slabs. Diffusion-weighted images were acquired with $b=1000 \text{ s/mm}^2$ and 16 diffusion directions uniformly sampled on a sphere. The echo spacing was 0.78 ms, and $R_y = 3$ acceleration was applied along k_y phase encoding, resulting in an effective echo spacing of 0.26 ms. $\text{TE1}(\text{imaging})/\text{TE2}(\text{navigator})/\text{TR}=73/130/1800 \text{ ms}$. For the reference protocol, the blip-up and blip-down diffusion weighted datasets were separately acquired with rectangular sampling with no acceleration and 20% oversampling along k_z (as in Fig. 3.2c, ii), with a scan time of $\sim 25 \text{ min}$. For the CAIPI-PF imaging protocol, the diffusion weighted data were acquired using the CAIPI-PF sampling (i.e., Fig. 3.2c, v) ($f_{\text{PF}} = 2/3$) with 10 blip-up shots and 10

blip-down shots for each slab. The acquisitions of blip-up and blip-down shots were consecutive within each diffusion encoding direction. For the conventional imaging protocol, the diffusion weighted data were acquired using 20- k_z rectangular sampling (Fig. 3.2c, i). The scan times for the CAIPI-PF and conventional imaging protocols were both ~ 10.5 min. One set of $b=0$ image was acquired with k_z fully sampled (i.e., 24 blip-up shots and 24 blip-down shots for each slab) and used for the diffusion analyses of all quantitative comparison protocols, with a scan time of ~ 1.4 min.

High-Resolution Diffusion Protocols

Six subjects were scanned using a 1.05 mm isotropic resolution CAIPI-PF imaging protocol to demonstrate the robustness of the proposed method. Diffusion-weighted images were acquired with $b=1000$ s/mm² and 48 diffusion directions uniformly sampled on a sphere with interleaved 3 $b=0$ image volumes. The echo spacing was 0.82 ms, and $R_y = 3$ acceleration was applied along k_y phase encoding, resulting in an effective echo spacing of 0.27 ms. TE1(imaging)/TE2(navigator)/TR=82/150/1800 ms. The FOV was $220 \times 220 \times 121$ mm³ and voxel size was 1.05 mm isotropic. The diffusion weighted data were acquired using the same CAIPI-PF sampling as in the Evaluation Protocol (i.e., Fig. 3.2c, v). To reduce aliasing artifacts from CSF signal, $b=0$ images were acquired with k_z fully sampled (i.e., 24 blip-up shots and 24 blip-down shots for each slab). The total scan time was ~ 33 min (36 s per diffusion direction).

In one subject, a dataset using a conventional 3D multi-slab high-resolution protocol [7] was also acquired using 20- k_z rectangular sampling (Fig. 3.2c, i) to enable comparisons at 1.05 mm isotropic resolutions. Eight $b=0$ images, including six blip-up and two blip-down volumes, were interspersed into the diffusion-weighted image acquisition. Other scan parameters were the same as the CAIPI-PF protocol and the total scan time was ~ 34 min.

An MPRAGE image was also acquired for each subject as an anatomical reference (~ 5 min) at 0.86 mm isotropic resolution.

3.2.3 Reconstruction Details

The integrated blip-reversed 3D multi-slab EPI data were reconstructed using the proposed two-stage reconstruction, while the conventional 3D multi-slab EPI data and reference separate blip-up and blip-down data were reconstructed by the stage 1 reconstruction. The SPIRiT kernel used in the reconstruction was trained using gradient echo coil calibration data. The image reconstruction was conducted offline in MATLAB 2021a (Mathworks, Natick, MA, USA). Image processing were conducted using functions from FMRIB Software Library (FSL) [64] unless indicated otherwise. Equation 3.1 was optimised with preconditioned conjugate gradient method with variable splitting [65]. The k-space data were first Fourier transformed along k_x followed by reconstruction performed for each $k_y - k_z$ plane using a 5×5 SPIRiT kernel. The reconstructed 2D images were concatenated along the readout direction (x) to form the whole image volume. The 32-channel data were compressed to 8 to shorten the calculation time [66]. The field map was calculated using “topup” [9, 63]. On a 2.9 GHz Quad-Core Intel Core i7 CPU, the computation time for one $k_y - k_z$ plane of one slab is ~ 15 s for stage 1 reconstruction, and ~ 90 s for stage 2 reconstruction. The processing time for “topup” is ~ 5 min per slab.

In stage 1 reconstruction, the unacquired partial Fourier region was zero filled before inverse Fourier transform along k_z , resulting in smoothing along the slice direction. As the ΔB_0 field is spatially smooth and the slab is thin, the impact of slice smoothness on field map estimation with “topup” is expected to be not significant. No zero filling is necessary for the stage 2 reconstruction because it operates on the full k_z extent (complementary blip-up and blip-down segments).

The 2D navigator images were reconstructed with 2D GRAPPA [22] using a 2D calibration dataset, which was acquired separately for blip-up and blip-down data [62]. The phase images were extracted and used as an estimation of motion induced phase errors.

Based on the balance between the data consistency and the regularisation effect of SPIRiT, we empirically searched the appropriate weights at different orders of magnitude (i.e., $\lambda_{1,1}=0, 0.1, 1, 10$ and $\lambda_{1,2}=0, 1, 10, 100$ for stage 1




Figure 3.5: Impact of SPIRiT regularisation weight on reconstruction. Stage 1 blip-up (a) and stage 2 joint reconstruction (c) of CAIPI-PF data using different SPIRiT regularisation weights and their residuals compared to the reference (b, d) acquired with the evaluation protocol (1.22 mm isotropic resolution) are displayed. The optimal weight for stage 1 (i.e., $\lambda_{1,1} = 1$) was used when different weights were evaluated for stage 2. The normalised root mean squared errors (NRMSE) of the whole slab are listed to quantify the image similarity.

and stage 2, respectively). The weights that produce the lowest normalised root mean squared error (NRMSE) were selected (i.e., $\lambda_{1,1} = 1$ and $\lambda_{1,2} = 10$) and further fine-tuned to obtain the final weights for reconstruction (i.e., $\lambda_{1,1} = 1$ and $\lambda_{1,2} = 20$) (Fig. 3.5). These optimal weights were used for the reconstruction of the whole-brain data from the evaluation and high-resolution protocols. The sparsity regularisation parameter $\lambda_2 = 0.7 \times 10^{-3}$ was used in both stage 1 and stage 2 reconstructions. The estimated field maps were compared with reference field maps

derived from fully sampled blip-up and blip-down images, and voxel displacement errors were calculated by multiplying the field map difference (in Hz) with the readout duration (in seconds). NRMSE and mean voxel displacement errors were calculated within a brain mask. The reconstruction codes are openly available at <https://github.com/liziyu0929/distortion-free-3d-diffusion-mri>.

The g-factors for different blip-up sampling patterns demonstrated in Fig. 3.2 were evaluated, which were calculated using the pseudo-multiple replica method [37] with 100 repetitions of the Monte-Carlo simulation. In each Monte-Carlo repetition, independent and identically distributed complex Gaussian noise was added to the pre-whitened multi-channel k-space data. For the stage 1 CAIPI-PF sampling, the impact of k_z partial Fourier on g-factor was accounted by dividing the resultant g-factor with the square root of partial Fourier factor as suggested by Kettinger et al. [67] The sparsity constraint was not included in g-factor calculations to reduce non-linearity. The g-factors were calculated within a brain mask.

3.2.4 Post-Processing

Slab combination and correction for slab saturation artifacts were performed using “NPEN” [62]. The whole-brain images were corrected for Gibbs ringing using “mrdegibbs3D” (<https://github.com/jdtournier/mrdegibbs3D>) [68, 69]. For the reference data acquired with the evaluation protocol, the whole-brain blip-up and blip-down data were processed by “topup” to obtain the distortion corrected images for each diffusion direction. The distortion corrected images were then processed with “eddy” without “-topup” option. For conventional 3D multi-slab data acquired with the evaluation and high-resolution protocols, a whole-brain field map was estimated using blip-up and blip-down $b=0$ image volumes using “topup”, which was then input to “eddy” [70] along with all diffusion data to correct for susceptibility and eddy current induced distortions. The CAIPI-PF data acquired with the evaluation and high-resolution protocols were processed with “eddy” without “-topup” option, as field inhomogeneity-induced distortion has been corrected in the reconstruction.

To evaluate the efficacy of the proposed method in estimating diffusion direction-dependent dynamic field maps and reducing the eddy current induced geometric distortions, the high-resolution CAIPI-PF data and conventional 3D multi-slab data were also processed by “eddy” without eddy current correction, which was facilitated by specifying “—flm=movement” in the command line of “eddy”.

3.2.5 Diffusion Analyses

All diffusion analyses were performed in the native diffusion space. For the evaluation protocol, as the diffusion data from the different methods were acquired in the same scan and were all co-aligned with the same $b=0$ image, they were already in the same native diffusion space. Diffusion tensor imaging (DTI) was performed on the whole-brain diffusion data for quantitative comparison between different methods.

For the high-resolution diffusion protocols, one subject was scanned twice using the CAIPI-PF and conventional 3D multi-slab high-resolution protocol, whose results were co-registered to the anatomical space for comparison. These datasets were acquired in separate sessions due to the long scan times. The diffusion data from both samplings were co-registered to the MPRAGE image. Due to different head positions, distortions induced by gradient nonlinearity were different between the two scans. For a fair comparison, the MPRAGE image was acquired in a separate session along with an intermediate $b=0$ image volume with matched gradient nonlinearity distortions to minimise the impact of gradient nonlinearity and improve the co-registration accuracy. The co-registration steps include:

1. An intermediate $b=0$ image volume (denoted as I_1) was acquired in the same scan as the MPRAGE image (with matched gradient distortions) using conventional 2D EPI and corrected for distortion using “topup” to minimise the impact of gradient nonlinearity distortions and improve the co-registration accuracy.
2. The $b=0$ images of CAIPI-PF and conventional 3D multi-slab sampling (denoted as $b0_{\text{CAIPI-PF}}$ and $b0_{\text{conventional}}$) were co-registered to I_1 using “flirt” [71,

72] with default parameters to correct gradient nonlinearity distortions. The resultant images and transformations were denoted as $I_{2,\text{CAIPI-PF}}$, $I_{2,\text{conventional}}$ and $T_{2,\text{CAIPI-PF}}$, $T_{2,\text{conventional}}$.

3. $I_{2,\text{CAIPI-PF}}$ and $I_{2,\text{conventional}}$ were co-registered to the MPRAGE image using “epi_reg” [71, 72]. The resultant transformations were denoted as $T_{3,\text{CAIPI-PF}}$ and $T_{3,\text{conventional}}$.
4. $T_{2,\text{CAIPI-PF}}$ and $T_{2,\text{conventional}}$ were combined with $T_{3,\text{CAIPI-PF}}$ and $T_{3,\text{conventional}}$ using “convert_xfm” [71, 72]. The resultant transformations were denoted as $T_{4,\text{CAIPI-PF}}$ and $T_{4,\text{conventional}}$.
5. $b0_{\text{CAIPI-PF}}$ and $b0_{\text{conventional}}$ were co-registered to the MPRAGE image using “flirt” with “bbr” as the cost function [73] and $T_{4,\text{CAIPI-PF}}$, $T_{4,\text{conventional}}$ as initialisations. The resultant transformations were used to resample the results from the diffusion analyses into the anatomical space.

For the remaining subjects, T1w images were co-registered to the diffusion space. The $b=0$ images were co-registered to the T1w images using “epi_reg”. The resultant transformations were then inverted using “convert_xfm” and applied to the T1w images.

The diffusion tensor model fitting was performed using “dtifit” [63]. White matter tractography was performed using “autoPtx” [74–76], which includes the pre-processing stage that runs probabilistic model fit using “bedpostx” [75], and the tractography stage which runs the probabilistic tractography using “probtrackx” [75].

3.3 Results

Figure 3.6 shows the effectiveness of the proposed sampling in reducing the slice aliasing in the slab boundary slice. Without oversampling, the boundary slice suffers from aliasing (comparing Fig. 3.6, i with Fig. 3.6, iv) caused by non-rectangular RF profile and limited FOV (Fig. 3.2a). Simple oversampling along k_z increases the FOV (Fig. 3.2b) and corrects the aliasing (comparing Fig. 3.6, ii with Fig. 3.6,




Figure 3.6: Boundary slice aliasing correction. The boundary slices (a) and enlarged regions (b) of blip-up data from rectangular sampling without oversampling along k_z ($R_z/R_y=1/3$) (i), rectangular sampling with 20% oversampling and $2\times$ acceleration along k_z ($R_z/R_y=1/3$) (ii), the proposed CAIPI-PF sampling with 20% oversampling, partial Fourier and $2.4\times$ acceleration along k_z ($R_z/R_y=2.4/3$) (iii) acquired with the evaluation protocol (1.22 mm isotropic resolution), and the reference MPRAGE data (0.86 mm isotropic resolution) (iv) are displayed, with relative scan times listed for each method except for MPRAGE reference. The parameter t represents the scan time of a rectangular sampling without over-sampling or acceleration along k_z (as Fig. 3.2c, i).

iv), but requires longer scan time. Our proposed CAIPI-PF sampling produces boundary slice aliasing-corrected images with a much faster acquisition, reducing scan time from $1.2t$ to $0.5t$ (t represents the scan time of a rectangular sampling without over-sampling or acceleration along k_z , as Fig. 3.2c, i).

Results from stage 1 blip-up reconstruction for different sampling patterns are demonstrated in Fig. 3.7. The reference data were fully sampled with no acceleration along k_y or k_z . The results with fully k_z encoding and $R_y = 3$ under-sampling along k_y produces the most similar result to the reference and the lowest g-factor (NRMSE: 9.3%; g-factor: 1.12) with the longest acquisition time. Image from rectangular under-sampling with $R_y/R_z = 3/2$ suffers from strong aliasing due to the limited coil sensitivity variation along the thin slab, producing the highest error and g-factor (NRMSE=27.1%; g-factor: 1.75). Using CAIPI sampling alone can substantially improve the reconstruction and lower the NRMSE (16.4%) and g-factor (1.45) by more efficient use of coil sensitivity information. The partial Fourier strategy further reduces reconstruction errors with lower NRMSE (15.6%)




Figure 3.7: Stage 1 blip-up reconstruction. The blip-up data for a slice at slab centre (a) from fully sampled reference (i) and different under-sampling patterns: rectangular sampling with 20% oversampling along k_z (Rect $1\times$ w/OS) (ii); rectangular sampling with 20% oversampling and $2\times$ acceleration along k_z (Rect $2\times$ w/OS) (iii); CAIPI sampling with 20% oversampling and $2\times$ acceleration along k_z (CAIPI $2\times$ w/OS) (iv); CAIPI-PF sampling with 20% oversampling, partial Fourier and $2.4\times$ acceleration along k_z (CAIPI-PF $2.4\times$ w/OS) (v), their residuals with the fully sampled reference (b), and their g-factor maps (v) are displayed, with relative scan times listed for each method. The data were acquired with the evaluation protocol (1.22 mm isotropic resolution).

and g-factor (1.24) using an even shorter scan time. The blip-down reconstruction produces consistent results with the blip-up reconstruction (Fig. 3.8).

The estimated field maps and stage 2 joint reconstruction results for different joint blip-reversed acquisition strategies are shown in Figure 5. The data acquired with full k_z sampling still produces the most accurate field map and the best stage 2 reconstruction (mean displacement error: 0.17 pixels; NRMSE: 6.6%). However, the scan time of this acquisition strategy is $2.4\times$ the scan time of a conventional multi-slab acquisition. In conventional rectangular under-sampling, the reconstructed image exhibits obvious structural difference with the reference due to the inaccurate field map estimated from the aliased images from stage 1

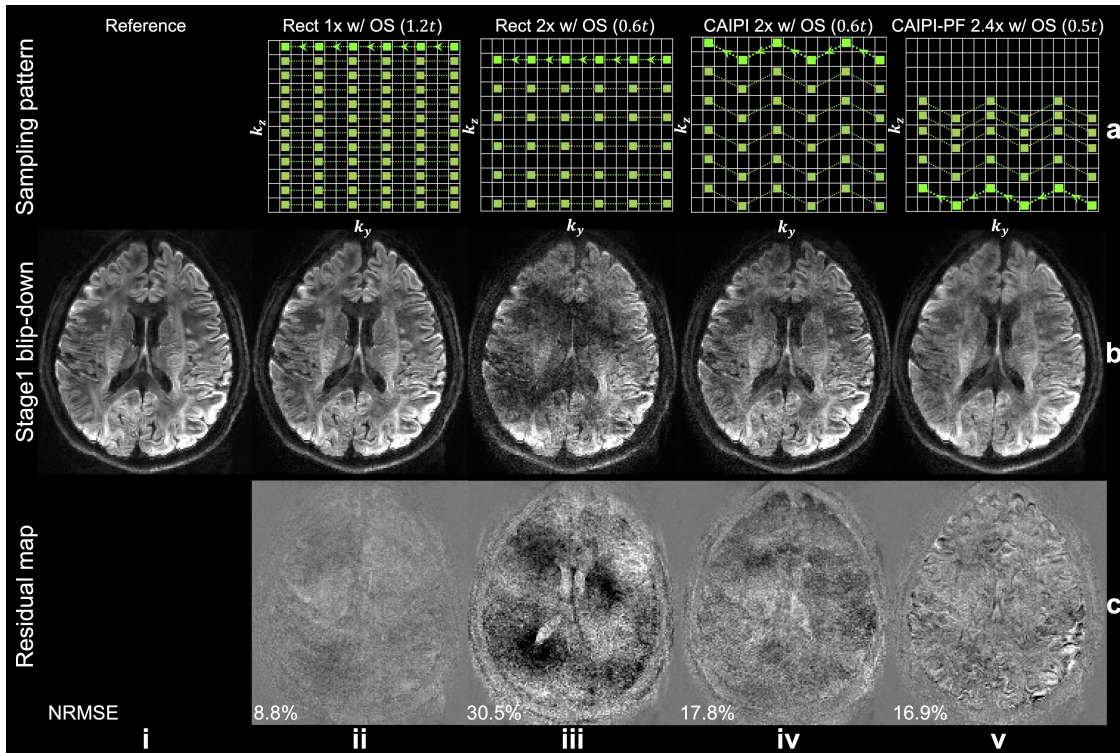


Figure 3.8: Stage 1 blip-down sampling and reconstruction. The sampling patterns (a), reconstruction results of a slab-centre slice (b) from fully sampled reference (i) and different sampling patterns (ii-v) and their residuals with the fully sampled reference (c) for the blip-down data acquired with the evaluation protocol (1.22 mm isotropic resolution) are displayed, with relative scan times listed for each method. The trajectory of one shot of the multi-shot sampling is marked in bright green with arrows.

reconstruction (mean displacement error: 1.09 pixels; NRMSE: 30.2%). The residual aliasing of regular CAIPI under-sampling in stage 1 reconstruction results in errors in the estimated field map, which leads to inaccuracies in joint reconstruction results (mean displacement error: 0.35 pixels; NRMSE: 14.2%). The proposed CAIPI-PF sampling provides an improved field map estimation and joint reconstruction with substantially better results due to a more densely sampled k-space centre (mean displacement error: 0.24 pixels; NRMSE: 12.6%) using the shortest scan time among all sampling patterns. The voxel displacement error map from CAIPI-PF also shows similar spatial pattern to the fully k_z -sampled data without obvious anatomical bias, indicating the CAIPI-PF sampling strategy does not introduce substantial errors in field map estimation.

Figure 3.10 demonstrates the whole-brain diffusion imaging results from the

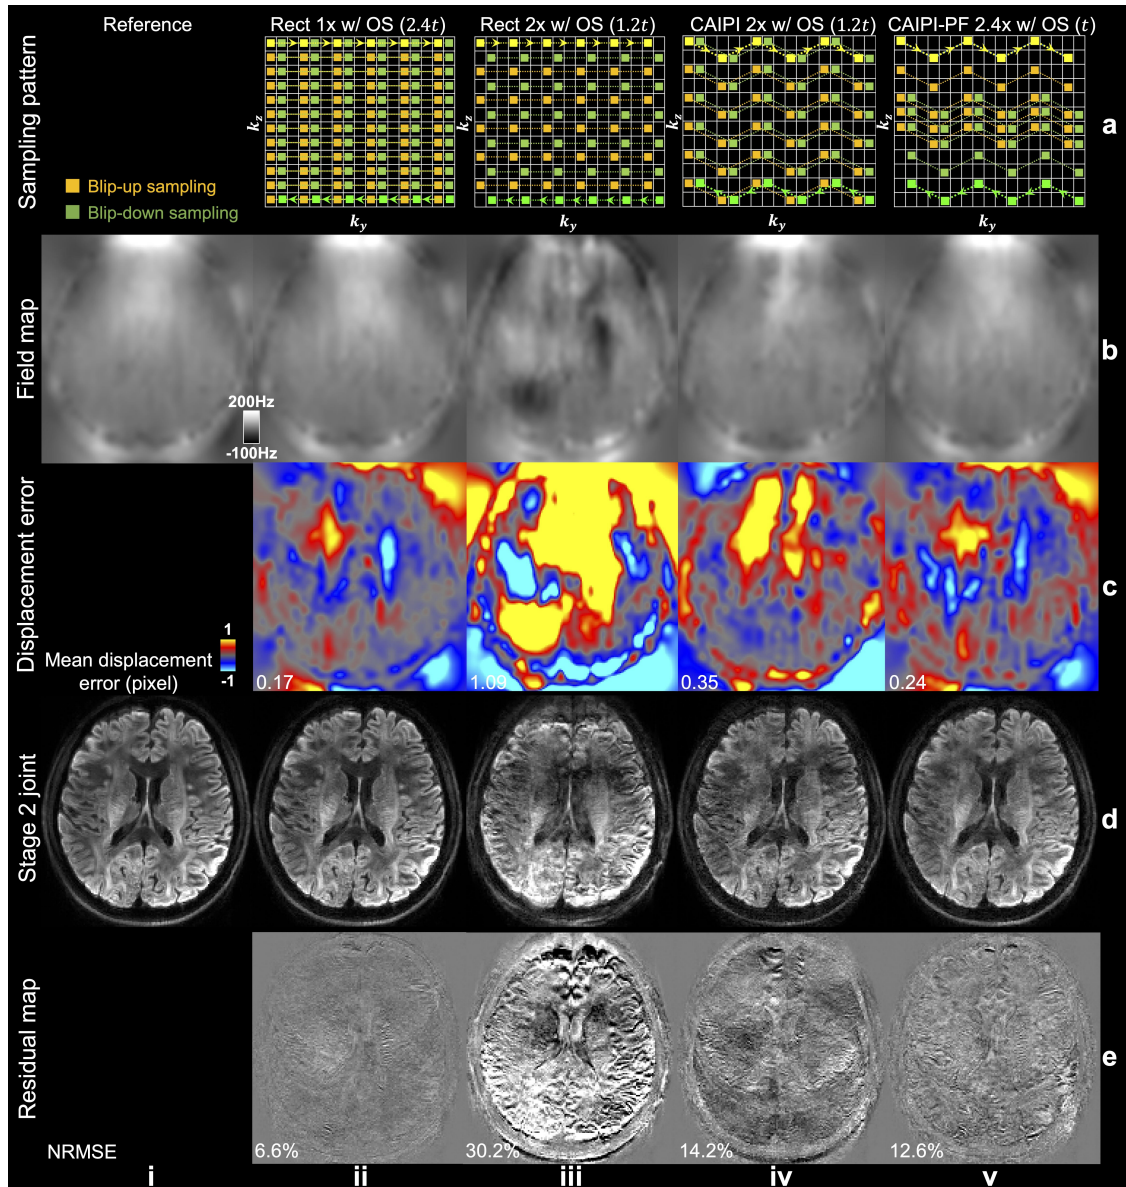


Figure 3.9: Field map estimation and joint reconstruction. The joint blip-up (yellow) and blip-down (green) sampling patterns (a), the corresponding estimated field maps from “topup” (b), the voxel displacement error maps (v), stage 2 joint reconstruction results (d), and residual maps (e) with the reference image acquired with the evaluation protocol (1.22 mm isotropic resolution) are displayed, with relative total scan times listed for each method. The trajectory of one shot of blip-up and blip-down sampling is marked in bright yellow and bright green, respectively.

evaluation protocols of one representative subject. Overall, the 16-direction DTI results exhibit high SNR thanks to the superior SNR efficiency provided by 3D multi-slab imaging. However, limited distortion correction performance is observed on the results from the conventional 3D multi-slab acquisition where diffusion

weighted images were acquired with only blip-up phase encoding, and distortion correction was applied using a field map derived from a pair of blip-up/down $b=0$ image, especially in the frontal region. The proposed CAIPI-PF protocol produces results visually similar to the reference results acquired with $2.4\times$ scan time. For quantitative comparison, the mean absolute differences of FA and MD compared to the reference results across the brain region (i.e., within a brain mask) were calculated for each subject. The group-level means (\pm standard deviation) of the mean absolute differences of the fractional anisotropy (FA) (calculated within FA masks where $FA > 0.05$) with the reference results are 0.0876 ± 0.00221 and 0.1004 ± 0.00170 for the proposed and conventional method, respectively. The mean absolute differences of mean diffusivity are $1.33 \times 10^{-4} \pm 6.49 \times 10^{-6} \text{mm}^2/\text{s}$ and $2.05 \times 10^{-4} \pm 2.29 \times 10^{-5} \text{mm}^2/\text{s}$ for the proposed and conventional method, respectively. The improved distortion and boundary slice-aliasing correction performance of the proposed method contribute to its better quantitative accuracy.

Results from the 1.05 mm whole-brain high-resolution protocols are showed in Fig. 3.11. Our proposed method produces single-volume diffusion-weighted image with high SNR and reduced artifacts and high-quality whole-brain DTI results. Compared to the conventional method, the proposed method achieves substantially improved anatomical fidelity, especially near the regions in the frontal lobe and pons where the field inhomogeneity is strong (Fig. 3.11a, yellow arrows). The conventional method also suffers from residual slice aliasing at slab boundaries (Fig. 3.11c, ii, the white arrow in the splenium of corpus callosum) due to the absence of over-sampling and limited overlapping between slabs (only one slice overlapping) to match the scan time with the proposed method.

The alignment of the diffusion data with the anatomical reference is showed in Fig. 3.12. When combining blip-up and blip-down data for distortion correction, the proposed method and the conventional method produce similar distortion correction on $b=0$ images. However, the distortion-corrected diffusion-weighted images from the proposed method exhibits consistently improved anatomical fidelity for different diffusion directions.




Figure 3.10: Comparisons of whole-brain diffusion images. The mean diffusion-weighted image (DWI) (a), the fractional anisotropy (FA) modulated primary eigenvector (V1) (b), and the mean diffusivity (v) of diffusion data acquired using blip-up and blip-down k_z -fully sampled (i.e., 24 blip-up and 24 blip-down shots for each diffusion direction) reference protocol (Reference, i), joint blip-reversed with CAIPI-PF protocol (Proposed, ii) and conventional 3D multi-slab protocol using only blip-up phase encoding (Conventional, iii) at 1.22 mm isotropic resolution from the same subject are displayed. The yellow arrows highlight the regions with strong distortions near frontal which cannot be fully corrected with the conventional method.

In addition, the proposed method can also correct eddy current induced distortions because the diffusion direction-dependent dynamic field map is incorporated into the joint reconstruction (Fig. 3.13). The co-registered images with and without eddy current correction are highly similar from the proposed reconstruction, with a low mean NRMSE across all diffusion directions (1.7%), indicating most eddy current induced distortions have already been corrected during the reconstruction. In comparison, images from the conventional method exhibit larger differences before and after the eddy current correction, with a much higher mean NRMSE for all diffusion directions (9.4%).

Tractography results show consistent improvement of the proposed method




Figure 3.11: High-resolution whole-brain diffusion images. The single volume diffusion-weighted images (DWI) along the diffusion direction $(-0.10, 0.88, 0.47)$ (a), the mean DWI (b), and the fractional anisotropy (FA) modulated primary eigenvector (V1) (c) of diffusion data acquired using joint blip-reversed with CAIPI-PF sampling (Proposed, i) and conventional 3D multi-slab sampling using only blip-up phase encoding (Conventional, ii) at 1.05 mm isotropic resolution from the same subject are displayed. The diffusion data are co-registered to the same anatomical data (0.86 mm isotropic resolution) for comparison.

compared to the conventional method. Improvements can be seen in frontal regions with residual distortion for the anterior thalamic radiation (Fig. 3.14a) and forceps minor (Fig. 3.14b), where the proposed method captures projections into the cortical grey matter more successfully. Similarly, the corticospinal tracts are represented as thicker white matter bundles in tractography based on the proposed reconstruction compared to the conventional method, whose thickness is considerably reduced near




Figure 3.12: Comparison of image anatomical fidelity. The axial (a, b) and sagittal (c, d) views of the MPRAGE reference (0.86 mm isotropic resolution, i), mean $b=0$ images (ii), mean diffusion-weighted images (DWI) (iii), and three DWI along different diffusion directions (iv-vi) acquired with CAIPI-PF sampling (Proposed, a, c) and conventional 3D multi-slab sampling (Conventional, b, d) at 1.05 mm isotropic resolution from the same subject are displayed. The white matter boundary segmented by FSL’s “fast” [77] is marked in red and overlaid. The diffusion data are co-registered to the MPRAGE reference for comparison.

	Proposed	Conventional	Intersection
Right anterior thalamic radiation	8955.4	7448.2	5205.8
Forceps minor	14281.6	9990.3	8017.7
Right corticospinal tract	11479.0	7679.7	6328.7

Table 3.1: Tract mask volumes. Tractography results of data from proposed CAIPI-PF (Proposed) and conventional 3D multi-slab (Conventional) sampling are binarised to obtain tract masks for different tracts (threshold: 0.3%). Tract mask volumes (in mm^3) are computed for tract masks of proposed and conventional data and their intersections.

the pons (red arrows, Fig. 3.14c). The tract mask (binarised with tract density threshold: 0.3%) volumes from the proposed method increase by 20.2%, 43.0%, and 49.5% compared to those from the conventional method for right anterior thalamic radiation, forces minor, and right corticospinal tracts, respectively. The detailed tract mask volumes can be found in Table 3.1.

These results generalise to other subjects (Fig. 3.15). The single volume




Figure 3.13: Correction of eddy current induced distortion. The single volume diffusion-weighted image along diffusion direction $(-0.16, -0.63, 0.76)$ acquired with CAIPI-PF (Proposed, i) and conventional 3D multi-slab sampling (Conventional, ii) at 1.05 mm isotropic resolution processed by FSL’s “eddy” with eddy current correction (a), without eddy current correction (b), and their difference maps (c) of the same subject are displayed. The normalised root mean squared errors (NRMSE) for diffusion-weighted images of all diffusion directions are listed to quantify the image similarity.

diffusion-weighted images show high-SNR and sharp textures. The DTI maps demonstrate high-quality and resolve fine structures with the high-resolution. The mean DWI results demonstrate high anatomical fidelity, with high structural similarity compared to the anatomical reference.

Tractography results for these subjects are shown in Fig. 3.16. The anterior thalamic radiation and forceps minor project successfully into the cortical grey matter




Figure 3.14: Comparison of tractography results. Tractography results of right anterior thalamic radiation (a), forces minor (b), right corticospinal tract (c) from the same subject using the proposed joint blip-reversed with CAIPI-PF (Proposed, i) and conventional 3D multi-slab sampling (Conventional, ii) with high-resolution protocols (1.05 mm isotropic) overlaid on their fractional anisotropy (FA) maps (tract density displaying range: 0.05%–5%), and their maximum intensity projection masks (iii) (binarised with tract density threshold: 0.3%) overlaid on the anatomical image. The red arrows highlight the regions with strong distortions near frontal (a, b, ii) and pons (c, ii), which cannot be fully corrected with the conventional method.

in regions of strong field inhomogeneity. The corticospinal tract is well reconstructed even in inferior regions such as the pons where B_0 field inhomogeneity is strong.

Diffusion imaging results of two other subjects are shown in Fig. 3.17.




Figure 3.15: Diffusion MRI results of multiple subjects using the proposed method. The single volume DWI along the diffusion direction $(-0.26, -0.81, -0.52)$ (a), the mean DWI (b), and the fractional anisotropy (FA) modulated primary eigenvector (V1) of diffusion tensor (c), and the anatomical images for reference (d) of three subjects are displayed. The diffusion images are acquired with the high-resolution protocol at 1.05 mm isotropic resolution using joint blip-up/down with CAIPI-PF acquisition. The anatomical images are acquired with MPRAGE at 0.86 mm isotropic resolution and co-registered to diffusion space for comparison.

3.4 Discussion

In this chapter, we demonstrate that high-resolution 3D multi-slab dMRI with minimal distortion and slice aliasing can be achieved by incorporating multiple sampling strategies into a joint reconstruction framework. For accurate distortion correction, both blip-up and blip-down shots for each diffusion direction were acquired for field mapping and distortion suppression. Over-sampling along k_z




Figure 3.16: Maximum intensity projections of fibre tracking results from three subjects. Three tracts including right anterior thalamic radiation (a), forceps minor (b), and right corticospinal tract (c) derived from data acquired with the high-resolution protocol from three subjects using the joint blip-reversed CAIPI-PF acquisition are displayed, overlaid on their fractional anisotropy (FA) maps. All tracts are visualized using the same track density threshold (0.05%–5%).

was used for an extended FOV to minimise slice aliasing at slab boundaries. Blip-reversed based distortion correction and boundary slice aliasing correction have been proposed previously [4, 12, 14]. One of the key challenges in this work is to acquire blip-reversed data with k_z over-sampling without requiring longer scan time than the conventional 3D multi-slab diffusion acquisition.

We address this challenge using two alterations in the sampling pattern. First,




Figure 3.17: Diffusion image results of the remaining subjects. The single volume diffusion-weighted images (DWI) along $(-0.26, -0.81, -0.52)$, the mean DWI (b), and the fractional anisotropy (FA) modulated primary eigenvector (V1) of diffusion tensor (c), and the anatomical images for reference (d) of two subjects are displayed. The diffusion images are acquired with CAIPI-PF sampling at 1.05 mm isotropic resolution. The anatomical images are acquired with MPRAGE at 0.86 mm isotropic resolution and co-registered to diffusion images for comparison.

blipped-CAIPI makes more efficient use of coil sensitivity and improves the conditioning of reconstruction. This is essential for 3D multi-slab imaging as each slab is very thin and variation of coil sensitivity across the slab is extremely limited. Without CAIPI the reconstruction is highly ill-posed and suffers from high noise amplifications (Fig. 3.7, iii). Blipped-CAIPI uses the in-plane coil profiles to encode information, improving the conditioning of the reconstruction and reducing the g-factor significantly. A relatively large slab thickness (~ 20 mm) is used in the

high-resolution protocol. Previous studies demonstrated the 2D navigator corrects most motion-induced phase errors even for slabs as thick as 30 mm [4] while the ideal range of slab thickness is 10–20 mm [4, 5, 7] for optimal phase error correction performance. Further evaluations are necessary to determine the optimal slab thickness in this highly accelerated scenario and achieve a well-balanced compromise between under-sampled reconstruction and phase correction performance.

Second, our k_z partial Fourier strategy enables a denser sampling of the k-space centre compared to regular under-sampling, which reduces the aliasing artifacts. The improvement is less obvious in stage 1 reconstruction (Fig. 3.7, iv, v), but is substantial in terms of the accuracy of the estimated field map and the image quality of stage 2 joint reconstruction (Fig. 3.9, iv, v). The resultant blurring along the slice direction in stage 1 reconstruction (residual structures in Fig. 3.7b, v) introduced by zero-filling of the partial Fourier region does not have an obvious impact on the accuracy of the estimated field map, presumably because field maps are intrinsically smooth, especially along the slice direction for thin slabs. In the joint reconstruction, such blurring is reduced thanks to the complementary patterns of blip-up and blip-down sampling. Furthermore, to extend the application of our proposed framework to higher spatial resolutions and/or higher b values that pose greater SNR challenges, one could incorporate partial Fourier encoding along the phase encoding direction in our framework.

The SPIRiT reconstruction used in this work is a model-based iterative k-space reconstruction approach [23] that outperforms conventional k-space reconstruction methods such as GRAPPA by making more efficient use of acquired data. Its iterative self-consistent formulation provides a noise averaging effect that reduces the overall noise [23, 78] and eliminates the need to train many kernels for our non-uniform sampling patterns as required by GRAPPA. Additionally, its model-based formulation allows straightforward incorporation of distortion and phase error correction into the reconstruction, which is essential for our work. Compared to image-space reconstruction methods such as SENSE [21], SPIRiT is advantageous in regions where accurate explicit sensitivity maps are challenging to obtain due to

low SNR or subject motion [34, 35]. Our experiments show the SPIRiT constraint improves SNR while preserving the sharp textures (Fig. 3.5). The g-factors of our reconstruction are also low, probably thanks to the implicit conditioning of noise for k-space-based reconstruction [37] and the noise averaging effect of SPIRiT regularisation [23]. Moreover, our regularised reconstruction framework offers considerable flexibility and can accommodate additional constraints to enhance the reconstruction performance. In our present study, we employed the sparsity constraint to suppress noise. The use of tailored constraints based on prior knowledge may further boost performance. Furthermore, our reconstruction can be highly parallelised because we used 2D SPIRiT (i.e., separately reconstructing each $k_y - k_z$ plane) in our framework. Using 3D SPIRiT may potentially further improve the reconstruction by leveraging data redundancy along the readout direction, but at a cost of longer computation time. Future work will be needed to explore the performance of 3D SPIRiT on the 3D multi-slab data.

The proposed framework’s motion robustness can be enhanced by integrating motion correction into the reconstruction process. Although the current *in-vivo* results do not exhibit significant head motion, presumably because the scan duration of 36 s for each diffusion direction is similar to the conventional method, it is important to note that head motion might still compromise the multi-shot joint reconstruction. To address this issue, a promising strategy is to incorporate the motion parameters estimated by “topup” into the forward model of the stage 2 joint reconstruction [79].

It is worth emphasising the importance of correcting the slab boundary slice aliasing in our framework. Our stage 2 joint reconstruction relies on an accurate field map to produce distortion-corrected reconstruction for each slab, which will be affected by boundary slice aliasing in stage 1 reconstruction. Therefore, removing boundary slice aliasing with over-sampling along k_z is essential for both aliasing-corrected reconstruction and accurate field map estimation. We utilise “NPEN” [62] to combine slabs and to reduce saturation artifacts. However, residual artifacts of slab combination still exist (Figs. 3.11, 3.15, 3.17), presumably due to the strong




Figure 3.18: The $b=0$ images of a representative subject. An axial slice of blip-up (a), blip-down (b), and “topup”-corrected (c) $b=0$ images of a representative subject scanned with CAIPI-PF sampling at 1.05 mm isotropic resolution are displayed.

B_0 field inhomogeneity at 7T that leads to slab profile distortion and shift, violating the assumption of “NPEN” that slab boundary artifacts are periodic along the slice dimension [62]. Nevertheless, these residual artifacts are much reduced in DTI fitting due to the normalisation with the $b=0$ data. These residual artifacts can be potentially reduced by improving the shimming and by incorporating B_0 information into the slab profile estimation.

The distortion correction efficacy of the proposed method is largely determined by the accuracy of field map estimation from stage 1 reconstruction. It is helpful to consider the “topup” correction using fully sampled blip-up and blip-down data as an upper bound for performance. In extreme cases when distortion is too strong for “topup” to correct even with fully sampled blip-up and blip-down data, the proposed method would have limited performance similar to “topup” correction. This is reflected in the results of one subject (Fig. 3.17, Subject 4), where residual distortions in the frontal region remain even after the joint reconstruction, probably due to the extremely high initial distortion level. Nevertheless, the level of residual distortion in the proposed method is highly similar to the “topup” corrected $b=0$ image with fully k_z sampling (Fig. 3.18), indicating that the main source of the residual

distortions is the extreme field inhomogeneity, likely due to poor shimming for that particular scan. A potential method to improve this is to use additional hardware to provide more homogeneous B_0 field (e.g., using local shimming coils [14]).

3.5 Conclusions

In this chapter, an acquisition and reconstruction framework are developed to minimise the distortion and boundary slice aliasing in high-resolution 3D multi-slab dMRI without increasing the scan time. The designed method achieves high-fidelity, robust reconstruction and produces high-SNR, high-quality diffusion images with superior anatomical fidelity compared to the conventional 3D multi-slab acquisition approach. The value of the method is demonstrated in improving the DTI fitting and tractography and can be further explored in more types of applications and by pushing to higher, submillimetre isotropic resolutions.

4

Self-Navigated 3D Diffusion MRI Using an Optimised CAIPI Sampling and Structured Low-Rank Reconstruction Estimated Navigator

Contents

4.1	Introduction	72
4.2	Methods	74
4.2.1	Theory	74
4.2.2	Experiments	83
4.3	Results	89
4.4	Discussion	101
4.5	Conclusion	105

In conventional 3D multi-slab diffusion MRI (dMRI), shot-to-shot phase variations caused by motion pose challenges due to the use of multi-shot k-space acquisition. Navigator acquisition after each imaging echo is typically employed to correct phase variations, which prolongs scan time and increases the specific absorption rate (SAR). The aim of this chapter is to develop a highly efficient, self-navigated method to correct for phase variations in 3D multi-slab dMRI without explicitly acquiring navigators. The sampling of each shot is carefully designed

to intersect with the central k_z plane of each slab, and the multi-shot sampling is optimised for self-navigation performance while retaining decent reconstruction quality. The central k_z intersections from all shots are jointly used to reconstruct a 2D phase map for each shot using a structured low-rank constrained reconstruction that leverages the redundancy in shot and coil dimensions. The phase maps are used to eliminate the shot-to-shot phase inconsistency in the final 3D multi-shot reconstruction. We demonstrate the method’s efficacy using retrospective simulations and prospectively acquired *in-vivo* experiments at 1.22 mm and 1.09 mm isotropic resolutions. Compared to conventional navigated 3D multi-slab imaging, the proposed self-navigated method achieves comparable image quality while shortening the scan time by 31.7% and improving the SNR efficiency by 15.5%. The proposed method produces comparable quality of diffusion tensor imaging (DTI) and white matter tractography to conventional navigated 3D multi-slab acquisition with a much shorter scan time. This chapter is based on a peer-reviewed journal paper we published in *IEEE Transactions on Medical Imaging* [80].

4.1 Introduction

High-resolution dMRI can provide detailed information about tissue microstructure and accurate representation of intricate fibre arrangements [2, 58]. However, the effectiveness of dMRI is limited by its inherently low SNR, which further decreases as the resolution increases.

3D multi-slab imaging is a promising approach for high-resolution dMRI due to its ability to achieve short TR=1-2 s [4–7, 54] for optimal SNR efficiency for spin-echo based dMRI. This technique divides the entire imaging volume into multiple thin slabs with 10-20 slices per slab, typically employing a 3D multi-shot echo-planar imaging (EPI) trajectory for high efficiency [4, 5, 7]. However, this method is sensitive to motion-induced shot-to-shot phase variations that degrade image quality when combining data from different shots [51, 81].

Conventional methods to correct the phase inconsistency in multi-shot diffusion MRI require navigators [4–7, 51, 54, 81]. The navigators are incorporated into

a model-based reconstruction to correct phase errors [51]. Previous studies have shown that the motion-induced phase within each 3D slab can be approximated by a 2D navigator if the slabs are sufficiently thin (i.e., <2 cm) [4–7, 54]. This 2D navigator necessitates an additional spin echo, leading to an extended TR and increasing the scan time by 25%-50% [4, 15]. Furthermore, the inclusion of another RF refocusing pulse increases the SAR.

Several studies have investigated the feasibility of navigator-free 2D multi-shot diffusion imaging. Multiplexed sensitivity-encoding (MUSE) reconstructs a 2D phase map from each under-sampled shot [52]. Despite the relatively high under-sampling factor per shot, clean phase maps can be reconstructed by exploiting the smoothness of motion-induced phase. Methods avoiding an explicit phase map have also been proposed [53, 82], which leverage the magnitude consistency between shots to create a structured Hankel matrix and estimate the missing data in each shot using structured low-rank (SLR) matrix completion [83].

Extending these methods for navigator-free 3D multi-slab diffusion imaging is intrinsically challenging. Specifically, to extend MUSE to 3D requires estimating a 2D phase map from each shot to capture in-plane phase variations. However, for conventional 3D Cartesian EPI (e.g., [4, 5, 7]), shots covering peripheral k_z planes encode high spatial frequency information, and estimating in-plane smooth phase variations is challenging. Additionally, in conventional 3D multi-slab diffusion imaging, each shot is effectively subjected to a high under-sampling factor with little redundancy between shots, which significantly hampers the feasibility of extending 2D SLR approach to 3D. Therefore, simplistically applying the 2D SLR approach to 3D multi-slab diffusion imaging might be impractical.

In this chapter, we present a novel acquisition and reconstruction framework for self-navigated 3D multi-slab dMRI. We propose a new 3D EPI sampling pattern that enables self-navigation with minimised k-space gaps in the shot-combined sampling and minimised overlapping between shots. A k-space based SLR constrained reconstruction is leveraged to jointly exploit data redundancy across shots and coils to reconstruct high-quality phase maps. Our multi-shot reconstruction incorporates

phase error correction to calculate the final image. The proposed method’s efficacy is evaluated through *in-vivo* experiments on a 7T scanner. The results demonstrate comparable image quality and significantly improved SNR efficiency compared to navigator-based methods. The resulting whole-brain DTI and tractography results highlight our proposed method’s potential to enable high-resolution 3D multi-slab diffusion imaging with improved time efficiency.

4.2 Methods

4.2.1 Theory

Review of 3D Multi-Slab Diffusion Imaging

3D multi-slab diffusion imaging uses multi-shot EPI to encode a 3D k-space for each slab. The acquired k-space signal for the j^{th} shot can be represented as:

$$y(\mathbf{k}_j) = \int \rho(\mathbf{r})\phi_{d_j}(\mathbf{r})e^{-i2\pi\mathbf{k}_j^T\mathbf{r}}d\mathbf{r} + n(\mathbf{k}_j), \quad (4.1)$$

where y is the diffusion k-space data, \mathbf{k}_j is the k-space sampling positions for the j^{th} shot, ρ is the phase error-free 3D image signal, \mathbf{r} is position in the image space, $\phi_{d_j} = e^{i\Psi_j}$, and Ψ_j is the non-diffusive motion induced phase in shot j , $j \in [1, N_{\text{shot}}]$, where N_{shot} denotes the total numbers of shots. In conventional 3D multi-slab diffusion imaging, each shot covers a k_z plane using a single EPI readout for efficient data acquisition (Fig. 4.1a), and N_{shot} equals to the number of encoded k_z planes (N_{k_z}). n is the additive Gaussian noise. The signal formulation described in Eq. 4.1 can be extended to the multi-coil case and represented in a matrix form as:

$$y_j = \mathcal{M}_j F \phi_{d_j} X + n_j, \quad (4.2)$$

where X is the matrix representation of the multi-coil phase error-free 3D image volume, \mathcal{M}_j is the k-space sampling matrix of shot j , and F is the Fourier Transform.

Non-diffusive motions during diffusion encoding introduced shot-dependent phase errors ϕ_d can lead to significant image corruptions if not corrected. Therefore, accurate information of phase errors ϕ_{d_j} is necessary for the reconstruction of X .

In conventional 3D multi-slab diffusion imaging, a 2D navigator is acquired at $k_z = 0$ for each shot using an extra refocusing RF pulse (i.e., the navigator samples the secondary spin echo) and the phase of the navigator is used as an estimation of motion induced phase errors (Fig. 4.1a). However, navigator acquisition suffers from several drawbacks, including prolonged scan time, increased SAR associated with a second spin echo, and reduced SNR efficiency.

Review of SLR Reconstruction for 2D Navigator-Free Diffusion Imaging

The 2D SLR reconstruction (e.g., MUSSELS [53, 82, 84]) leverages the data redundancy across different shots to effectively restore missing data. Consequently, it enables accurate estimation of magnitude and phase information for each shot without the need for explicitly acquired phase maps. MUSSELS assumes different shots share the same underlying magnitude image m despite their different phases:

$$m = x_{j_1} \Phi_{j_1}^H = x_{j_2} \Phi_{j_2}^H, \quad j_1, j_2 \in [1, N_{\text{shot}}], \quad (4.3)$$

where x_{j_1} and x_{j_2} are the complex image for shot j_1 and j_2 , respectively, $\Phi = \phi_d \phi_c$ includes the diffusion-related phase ϕ_d and coil-related phase ϕ_c , and Φ^H denotes the conjugate of Φ . This leads to the establishment of an annihilation relation in both the image domain (Eq. 4.4) and the Fourier domain (Eq. 4.5):

$$x_{j_1} \Phi_{j_2} - x_{j_2} \Phi_{j_1} = 0, \quad j_1, j_2 \in [1, N_{\text{shot}}], \quad (4.4)$$

$$\hat{x}_{j_1} * \hat{\Phi}_{j_2} - \hat{x}_{j_2} * \hat{\Phi}_{j_1} = 0, \quad j_1, j_2 \in [1, N_{\text{shot}}], \quad (4.5)$$

where \hat{x} and $\hat{\Phi}$ denotes the Fourier Transform of x and Φ , respectively, and $*$ denotes the convolution operation. As Φ is typically smooth in image space, $\hat{\Phi}$ should be support limited in the Fourier domain. Utilising the block-Hankel matrix formulation $H_1(\hat{x})$ detailed in previous work [53], Eq. 4.5 can be expressed in a matrix form:

$$\begin{bmatrix} H_1(\hat{x}_1) & H_1(\hat{x}_2) \end{bmatrix} \begin{bmatrix} \text{vec}(\hat{\Phi}_{j_2}) \\ -\text{vec}(\hat{\Phi}_{j_1}) \end{bmatrix} = \mathbf{0}, \quad j_1, j_2 \in [1, N_{\text{shot}}]. \quad (4.6)$$

This relation holds for all pairs of shots as they are all assumed to share the same magnitude image, Thus, the structured matrix

$$H(\hat{x}) = \begin{bmatrix} H_1(\hat{x}_1) & H_1(\hat{x}_2) & \dots & H_1(\hat{x}_{N_{\text{shot}}}) \end{bmatrix} \quad (4.7)$$

exhibits a low-rank property due to the existence of a non-trivial null space P ($H(\hat{x})P = \mathbf{0}$):

$$P = \begin{bmatrix} \hat{\Phi}_2 & 0 & 0 & \hat{\Phi}_3 \\ -\hat{\Phi}_1 & \hat{\Phi}_3 & 0 & 0 \\ 0 & -\hat{\Phi}_2 & 0 & -\hat{\Phi}_1 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \hat{\Phi}_{N_{\text{shot}}} & 0 \\ 0 & 0 & -\hat{\Phi}_{N_{\text{shot}}-1} & 0 \end{bmatrix}. \quad (4.8)$$

The 2D SLR reconstruction enforces the low-rankness of $H(\hat{x})$ by minimising $\|H(\hat{x})\|_*$ during iterative reconstruction to facilitate high-fidelity navigator-free 2D diffusion imaging.

However, the extension of 2D SLR reconstruction to 3D multi-slab diffusion imaging presents notable challenges. While previous approaches, such as gSlider [14, 85], have effectively applied SLR to high-resolution multi-shot diffusion imaging, they essentially reconstructed RF-encoded 2D k-space. Our work focuses on leveraging SLR for 3D k-space reconstruction. A direct extension to 3D SLR (by substituting 2D images in Eq. 4.3 with 3D volumes) requires reconstructing an entire 3D volume from a single shot, which undergoes exceptionally high under-sampling factor for 3D multi-slab imaging. For instance, if a single slab is encoded with 10 shots and an acceleration factor of $R_y = 3$ is applied along the phase-encoding direction, the under-sampling factor for each shot would amount to 30. Furthermore, the construction of Hankel matrices using 3D volumes in Eq. 4.7 would entail substantial computational costs and feasibility concerns.

Self-Navigated 3D Multi-Slab Imaging Framework

Our method, illustrated in Fig. 4.1b, aims to integrate the SLR reconstruction into 3D multi-slab diffusion imaging to eliminate the need for navigator acquisition.




Figure 4.1: Comparison between the conventional navigated 3D multi-slab imaging (a) and the proposed self-navigated imaging framework (b). In the conventional approach (a), phase information is obtained from a separately acquired 2D navigator. In contrast, the proposed method (b) extracts shot phase directly from the imaging data using a novel acquisition and reconstruction framework, eliminating the need for acquiring a separate navigator.

This is achieved with an extended CAIPI [31] sampling trajectory where each shot covers a wide range of k_z planes and intersects with the central $k_z = 0$ plane (the intersections are called “self-navigation points” hereafter). The self-navigation points are used to reconstruct a fully sampled $k_z = 0$ plane for each shot, akin to a 2D navigator but without requiring additional scan time. The sampling trajectories of all shots are jointly optimised for robust self-navigation and multi-shot reconstruction. The limited self-navigation points result in high under-sampling. We developed a SPIRiT-based [23] SLR reconstruction that exploits the redundancy between coils and shots to produce a robust estimation of motion-induced phase error for each shot. This is then incorporated into the multi-shot reconstruction

to eliminate shot-to-shot phase inconsistencies as in conventional 3D multi-slab diffusion imaging reconstruction [4–7, 54] (Fig. 4.1).

K-Space Sampling Optimisation

The sampling coverage of each shot’s trajectory is expanded along k_z by inserting k_z blips between readout gradients, such that each shot traverses through the central $k_z = 0$ plane to provide the self-navigation points. The k_z width w of each shot was set to $w = \text{floor}(N_{k_z}/2) + 1$, the minimal value to ensure each shot intersects with the central $k_z = 0$ plane. Using this value can minimise the period of the sampling and therefore maximise the number of self-navigation points to benefit the 2D phase map reconstruction. A fundamental sampling configuration was established for an individual shot with a fixed k_y period and k_z width (Fig. 4.2a), based on which the sampling pattern was designed.

Drawing from this basic sampling, each shot could be acquired with a different pattern by modifying several parameters characterising the sampling. The k_z shift s_{k_z} can be altered to determine the sampling’s vertical position (Fig. 4.2b) $s_{k_z} \in [0, N_{k_z} - w]$. The periodic shift s_p controls the starting point of the periodic sampling (Fig. 4.2c). With a k_z width of w , the period of the sampling would be $2w - 2$. Therefore, the range of s_p is $[0, 2w - 3]$. Parallel imaging acceleration is usually applied along the phase encoding direction (k_y) to shorten TE and mitigate geometric distortion (not illustrated in Fig. 4.1 and Fig. 4.2a, b, c). Therefore, the sampling pattern can also be modified with a k_y shift s_{k_y} ($s_{k_y} \in [0, 1, \dots, R_y - 1]$), where R_y is the parallel imaging acceleration factor along k_y .

The sampling pattern can impact the image reconstruction quality in various ways. First, distinct shots might overlap, reducing sampling efficiency. Second, large gaps in k-space coverage could introduce substantial artifacts in the reconstructed images [86]. Lastly, the positioning of self-navigation points also affects phase correction performance. Ideally, each shot should incorporate some self-navigation points positioned close to the $k_y - k_z$ centre to capture sufficient low-frequency data for precise 2D phase map reconstruction. An optimal sampling scheme should

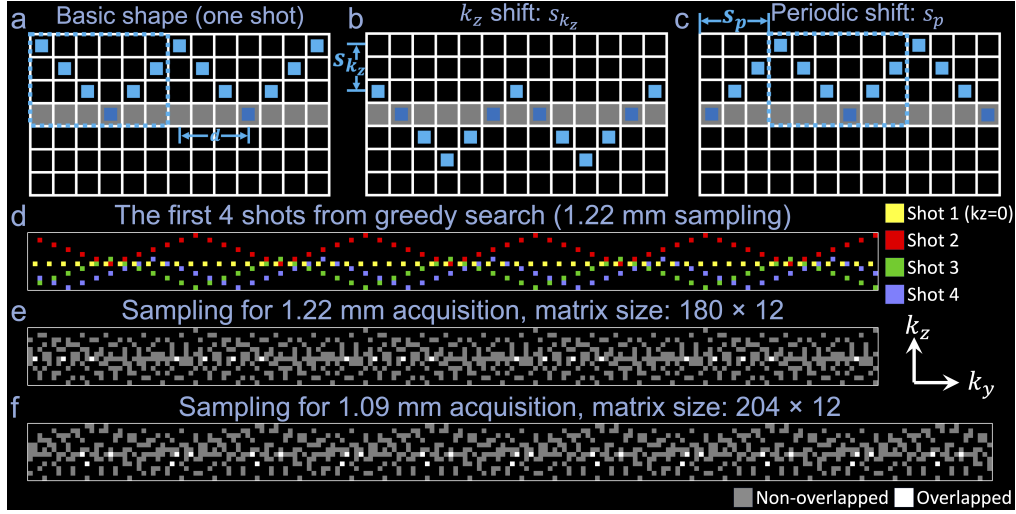


Figure 4.2: K-space sampling optimisation. The basic shape of one shot of the proposed extended CAIPI sampling (a), the illustration of sampling parameters: k_z shift s_{k_z} (b), and periodic shift s_p (c). Note the parallel imaging acceleration along the phase encoding direction (k_y) is not illustrated in (a-c) for simplicity, and acceleration factor=3 is used in (d-f). The intersections with central $k_z = 0$ plane (self-navigation points) are marked in dark blue (a-c). The parameter d which denotes the shortest distance between self-navigation points and the $k_y - k_z$ centre is also illustrated in (a). The dashed boxes indicate one period of the periodic sampling (a, c). The first 4 shots of resulting sampling patterns for 1.22 mm acquisition are plotted in yellow ($k_z = 0$ traversal shot), red, green, and blue, respectively (d). The overall sampling for 1.22 mm (matrix size: 180×12 , e) and 1.09 mm acquisition (matrix size: 204×12 , f) from the proposed sampling optimisation are demonstrated. The non-overlapped and overlapped points are marked in grey and white, respectively.

strive to minimise overlap and gaps in k-space coverage, while ensuring each shot includes some self-navigation points near the $k_y - k_z$ centre.

We developed a framework to optimise the sampling pattern through metrics of overlap, gaps, and self-navigation points. The metric o_i is the number of overlap points in the i -shot-combined sampling, where i is the optimisation step count and $i \in [1, N_{\text{shot}}]$:

$$o_i = T_o \left(\sum_{j=1}^i M_j(s_{k_{y_j}}, s_{k_{z_j}}, s_{p_j}) \right), \quad (4.9)$$

where $M_j(s_{k_{y_j}}, s_{k_{z_j}}, s_{p_j})$ is the sampling mask for shot j , specified by parameters $s_{k_{y_j}}, s_{k_{z_j}}, s_{p_j}$, respectively. T_o denotes an operation counting the number of elements greater than one within the combined sampling mask $\sum_{j=1}^i M_j$.

The metric g_i reflects the number of gaps larger than 3×3 in the i -shot-combined sampling:

$$g_i = T_g \left(\left(\sum_{j=1}^i M_j(s_{k_{y_j}}, s_{k_{z_j}}, s_{p_j}) \right) * \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \right), \quad (4.10)$$

where T_g counts the number of zeros in the convolution result.

The metric d is the shortest distance between self-navigation points and the $k_y - k_z$ centre for each shot (Fig. 4.2a). It is essential for each shot to have a small value of d to provide low-frequency information for accurate reconstruction of 2D phase maps. We establish a maximum allowable distance d_{\max} .

We optimise the sampling by solving the following problem:

$$\arg \min_{s_{k_{y_j}}, s_{k_{z_j}}, s_{p_j}} o_i + g_i + d_i, \quad s.t., \quad d_i \leq d_{\max}. \quad (4.11)$$

The optimisation is solved efficiently using greedy search in a shot-by-shot manner. Importantly, the first shot is designed to traverse the $k_z = 0$ plane without k_z blip to accurately capture the magnitude information of $k_z = 0$, which benefits the SLR reconstruction of the phase maps (see the following section).

SPIRiT-Based SLR Reconstruction for Phase Map Estimation

We address the extreme under-sampling of the phase maps by jointly leveraging the shared information across shots and coils using a SPIRiT [23, 87]-based SLR reconstruction.

The careful design of the sampling facilitates the adaptation of 2D SLR method for the reconstruction of the central $k_z = 0$ plane for each shot. Equations 4.3-4.8 can be readily applied to construct a low-rank block-Hankel matrix for $k_z = 0$ data, assuming x_j represents the $k_z = 0$ image for shot j . Notably, unlike the original 2D SLR formulation [53, 82, 84] where x_j is the coil-combined image, we utilise the multi-coil data to construct the block-Hankel matrix. Because coil sensitivity is smooth with limited k-space support, the multi-coil formulation promotes the low-rankness of the block-Hankel matrices, and therefore is believed to be beneficial




Figure 4.3: The block-Hankel matrix for the $k_z = 0$ plane reconstruction leverages the low-rank properties in both coil and shot dimensions. The sampled lines represent the self-navigation points in the $k_z = 0$ plane, reflecting the under-sampling without the parallel imaging acceleration. The actual sampling includes an additional parallel imaging acceleration factor=3, resulting in a threefold increase in the overall under-sampling factor.

for the parallel imaging reconstruction [28]. The multi-shot, multi-coil block-Hankel matrix of the $k_z = 0$ plane $H(\hat{x})$ is formulated as (Fig. 4.3):

$$H(\hat{x}) = \begin{bmatrix} H_{\text{shot}}(\hat{x}_1) & H_{\text{shot}}(\hat{x}_2) & \dots & H_{\text{shot}}(\hat{x}_{N_{\text{shot}}}) \end{bmatrix}, \quad (4.12)$$

$$H_{\text{shot}}(\hat{x}_u) = \begin{bmatrix} H_1(\hat{x}_{u,1}) & H_1(\hat{x}_{u,2}) & \dots & H_1(\hat{x}_{u,N_{\text{coil}}}) \end{bmatrix}, \quad (4.13)$$

where $H_{\text{shot}}(\hat{x}_u)$ is the multi-coil submatrix for shot u , $\hat{x}_{u,v}$ is the k-space data from the u^{th} shot and v^{th} coil, and N_{coil} denotes the total number of coils. The limited k-space support of the coil sensitivity resulting from its smoothness causes $H_{\text{shot}}(\hat{x}_u)$ to be rank-deficient [28], with $\text{rank}(H_{\text{shot}}(\hat{x}_u)) \leq (w_{\text{H}} + s - 1)^2$ where w_{H} is the kernel size for constructing the block-Hankel matrix, and s is the k-space support size of the coil sensitivity. This induces coil low-rankness in addition to shot low-rankness demonstrated by the annihilation relations in Equations 4.3-4.8. By minimising $\|H(\hat{x})\|_*$, we enforce low-rank and harness the redundancy in both the shot and coil dimensions. This constraint promotes similarity among the magnitude images derived from different shots. The use of one shot that traverses the $k_z = 0$ plane without k_z blips provides accurate magnitude information m' to improve the robustness and accelerate the convergence.

Our SPIRiT-based SLR reconstruction recovers the 2D phase maps for each shot from the self-navigation points by solving:

$$\arg \min_{\hat{x}} \|\mathcal{M}_{k_z 0} \hat{x} - y_{k_z 0}\|_2^2 + \lambda_1 \|(G_{k_z 0} - I)\hat{x}\|_2^2 + \lambda_2 \|H(\hat{x})\|_*, \quad (4.14)$$

where $\mathcal{M}_{k_z 0}$ selects self-navigation points, \hat{x} is the fully-sampled multi-coil k-space data for all shots at $k_z = 0$, $y_{k_z 0}$ is the acquired self-navigation points, $G_{k_z 0}$ is the SPIRiT kernel trained on the $k_z = 0$ calibration data (see below), I is the identity matrix, and λ_1, λ_2 are the hyperparameters for the SPIRiT and SLR constraints.

The SPIRiT kernel $G_{k_z 0}$ is trained on calibration data similarly to GRAPPA [22] and performs convolutions in k-space to reach self-consistency:

$$\hat{x} = G_{k_z 0} \hat{x}, \quad (4.15)$$

which is achieved by minimising $\|(G_{k_z 0} - I)\hat{x}\|_2^2$ in Eq. 4.14.

We reformulate Eq. 4.14 to solve it iteratively using alternating direction method of multipliers (ADMM) [88]:

$$\arg \min_{\hat{x}} \|\mathcal{M}_{k_z 0} \hat{x} - y_{k_z 0}\|_2^2 + \lambda_1 \|(G_{k_z 0} - I)\hat{x}\|_2^2 + \lambda_2 \|z\|_*, \quad s.t., z - H(\hat{x}) = 0. \quad (4.16)$$

In the k^{th} iteration, Eq. 4.16 is further split into the following three subproblems:

$$\hat{x}^k = \arg \min_{\hat{x}} \|\mathcal{M}_{k_z 0} \hat{x} - y_{k_z 0}\|_2^2 + \lambda_1 \|(G_{k_z 0} - I)\hat{x}\|_2^2 + \left(\frac{\beta}{2}\right) \|z^{k-1} - H(\hat{x}) + u^{k-1}\|_2^2, \quad (4.17)$$

$$z^k = \arg \min_z \lambda_2 \|z\|_* + \left(\frac{\beta}{2}\right) \|z - H(\hat{x}^k) + u^{k-1}\|_2^2, \quad (4.18)$$

$$u^k = u^{k-1} + z^k - H(\hat{x}^k), \quad (4.19)$$

where β is the dual update step length. As all shots share the same magnitude information, an explicit magnitude image m' for $k_z = 0$ is obtained from the $k_z = 0$ traversing shot, which is included in Eq. 4.17 as another constraint:

$$\begin{aligned} \hat{x}^k = \arg \min_{\hat{x}} & \|\mathcal{M}_{k_z 0} \hat{x} - y_{k_z 0}\|_2^2 + \lambda_1 \|(G_{k_z 0} - I)\hat{x}\|_2^2 + \\ & \lambda_3 \|\hat{x} - Fm'\Phi^{k-1}\|_2^2 + \left(\frac{\beta}{2}\right) \|z^{k-1} - H(\hat{x}) + u^{k-1}\|_2^2, \end{aligned} \quad (4.20)$$

where Φ^{k-1} is the image phase of \hat{x}^{k-1} . In practice, the constraint $\|\hat{x} - Fm'\Phi^{k-1}\|_2^2$ was found effective in improving the robustness and convergence speed of the reconstruction. We solve Eq. 4.18 by applying singular value decomposition (SVD) and hard thresholding similar to previous work [28, 53, 89]. Eq. 4.20 is solved using the conjugate gradient (CG) method.

A 2D phase map can then be extracted from the reconstructed $k_z = 0$ image \hat{x}_j for each shot:

$$\phi_{d_j} = \frac{S^H F^{-1} \hat{x}_j}{\|S^H F^{-1} \hat{x}_j\|}, \quad j \in [1, N_{\text{shot}}], \quad (4.21)$$

where S is the coil sensitivity maps for $k_z = 0$, which can be obtained from $k_z = 0$ calibration data. According to Eq. 4.2, the 2D phase maps can be incorporated into the forward model to address the phase variations in the multi-shot reconstruction:

$$\hat{X} = \arg \min_{\hat{X}} \sum_j \|\mathcal{M}_j F \phi_{d_j} F^{-1} \hat{X} - y_j\|_2^2 + \lambda_4 \|(G_{\text{slab}} - I)\hat{X}\|_2^2, \quad (4.22)$$

where \hat{X} is the phase error-corrected 3D k-space, G_{slab} is the SPIRiT kernel trained on coil calibration data for the whole slab, and λ_4 is the SPIRiT regularisation weight of the final phase-corrected reconstruction. The desired slab image can be obtained from the sum-of-squares of $F^{-1}\hat{X}$.

4.2.2 Experiments

Data Acquisition

A 3D multi-slab spin-echo diffusion MRI sequence [7] was modified to integrate the proposed CAIPI sampling for self-navigation (referred to as ‘‘Self-nav CAIPI’’ hereinafter). We used $N_{k_z} = 12$, $R_y = 3$, $d_{\text{max}} = 15$, such that $s_{k_z} \in [0, 5]$, $s_{k_y} \in [0, 2]$, $s_p \in [0, 11]$. N_{shot} was also set to 12 for consistency with the conventional 3D multi-slab imaging. In the i^{th} step of the sampling optimisation, we exhaustively

explored all combinations of $s_{k_y i}$, $s_{k_z i}$, s_{p_i} to identify lowest value for the cost function $o_i + g_i + d_i$. The resulting sampling patterns for 1.22 mm acquisition (matrix size: 180×12 for $k_y - k_z$ plane) and 1.09 mm acquisition (matrix size: 204×12 for $k_y - k_z$ plane) are demonstrated in Fig. 4.2d, e and Fig. 4.2f. There is little overlapping and no substantial k-space gaps. The samplings are denser near k-space centre with an overall random distribution.

Some experiments included an explicit navigator acquisition with 64 phase encoding lines and matching the imaging echo for phase encode direction, R_y , and echo spacing (ES).

Subjects were scanned on a Siemens 7T scanner (Siemens Magnetom, Erlangen, Germany) using a 32-channel receive coil. Written informed consent in accordance with local ethics was obtained from each subject.

Table 4.1 summarises six experimental protocols. Images were acquired at 1.22 and 1.09 mm isotropic resolutions (180×180 and 204×204 matrix, ES=0.76 and 0.82 ms, respectively). Unless otherwise specified, acquisitions used 13 slabs, 12 slices/slab, 2 slice slab overlap, partial Fourier (PF) factor 3/4 and $b=1000$ s/mm².

Experiment 1 Simulation Evaluation: To validate the proposed sampling and reconstruction, we acquired single-slab data (12 slices) from one subject at 1.22 mm (Table 4.1. 1A) and 1.09 mm (Table 4.1. 1B). Fully sampled reference data with $b=1000$ s/mm² were reconstructed from three scans (each with $R_y = 3$ and the conventional sampling, phase error-corrected using the 2D navigator) with 0, 1, and $2\Delta k_y$ shift and no PF. Realistic phase-corrupted multi-shot data were simulated by multiplying the reference data with one set of navigator phase maps followed by applying the shot sampling masks.

Using the fully sampled dataset, we conducted the following experiments to evaluate our sampling and reconstruction:

- Exp 1.1: we evaluated the conventional and Self-nav CAIPI samplings (Fig. 4.2e, f).

Exp.	Res. (mm ³)	TE1/TE2/ TR(ms) ^a	b-value (s/mm ²)	#b=0/ DWI ^b	Samp. ^c	Nav.	T _{acq} (min:sec)
1A	1.22	82/142/2000	1000	1/3	Conv.	Yes	01:36
1B	1.09	89/157/3500	1000	1/3	Conv.	Yes	02:48
2A	1.09	65/133/3500	1000	0/36	Conv.	Yes	25:12
2B	1.09	65/133/3500	1000	0/36	Self-nav	Yes	25:12
3A	1.09	65/133/3500	1000	0/12	Conv.	Yes	08:24
3B	1.09	64/-/2400	1000	0/12	Self-nav	No	05:46
4A	1.09	65/133/3500	1000	2/16	Conv.	Yes	12:36
4B	1.09	64/-/2400	1000	2/16	Self-nav	No	08:38
5A	1.09	72/140/3500	2000	0/1	Conv.	Yes	00:42
5B	1.09	72/140/3500	2000	0/1	Self-nav	Yes	00:42
5C	1.09	78/146/3500	3000	0/1	Conv.	Yes	00:42
5D	1.09	78/146/3500	3000	0/1	Self-nav	Yes	00:42
6	1.09	64/-/2400	1000	6/48	Self-nav	No	25:55

Table 4.1: Acquisition parameters. a. TE1 is the imaging echo time and TE2 is the navigator echo time. 3/4 partial Fourier along phase encoding direction was applied for Exp. 2-6. The TRs of Exp. 2-6 were the shortest achievable TRs due to SAR restriction. b. The number of DWIs refers to the number of scans with different encoding for each experiment. In Exp. 1, three Δk_y shift was acquired with the same diffusion encoding. In Exp. 2, 12 repetitions along three orthogonal diffusion encoding directions were obtained for SNR evaluation. In Exp. 3, 12 repetitions were acquired with a single diffusion encoding direction along readout. c. Conv.: conventional sampling; Self-nav: Self-nav CAIPI sampling. The Self-nav CAIPI samplings for 1.22 and 1.09 mm are illustrated in Fig. 4.2e and f respectively. All b=0 data were acquired with the conventional sampling.

- Exp 1.2: we simulated three sub-optimal samplings by solving Eq. 4.11 (i) without the overlap constraint o_i (i.e., $g_i + d_i$ is minimised subject to $d_i \leq d_{\max}$); (ii) without the k-space gap constraint g_i (i.e., $o_i + d_i$ is minimised subject to $d_i \leq d_{\max}$); (iii) without the self-navigation performance constraint d_i (i.e., $o_i + g_i$ is minimised), and evaluated the reconstructed images.
- Exp 1.3: we simulated Self-nav CAIPI sampling with and without the use of a non-blipped $k_z = 0$ shot and investigated the impact of including the magnitude constraint $\|\hat{x} - Fm'\Phi^{k-1}\|_2^2$ in Eq. 4.20.
- Exp 1.4: we compared our method with a previously proposed self-navigation method by Moeller et al. [15] (entitled ‘‘Self-nav Conventional’’), which used conventional sampling and extracted a motion-induced phase map for each shot from each k_z plane.

- Exp 1.5: we investigated the contribution of shot and coil dimension in our SLR reconstruction and implemented SLR only leveraging redundancy across shots (i.e., coil-combined SLR) or coils; we also selected 16-coil data from our 32-coil data to evaluate the performance of our method when fewer coils are available.

Experiment 2 SNR Evaluation with Matched TR: To quantitatively evaluate SNR, we acquired data from three subjects using conventional (Table 4.1. 2A) and Self-nav CAIPI sampling (Table 4.1. 2B) with matched TE and TR. To evaluate the accuracy of phase map estimation, navigators were also acquired. Diffusion weighting along the readout, phase, and slice directions was evaluated, each with twelve repetitions.

Experiment 3 SNR Efficiency Evaluation with Optimal TR: To quantitatively compare the SNR efficiency, we acquired data in four subjects with conventional (with navigators, Table 4.1. 3A) and Self-nav CAIPI (without navigators, Table 4.1. 3B) sampling, with different TR optimised for each scan. Diffusion weighting was along readout with twelve repetitions.

Experiment 4 DTI Comparison: To compare DTI results, we acquired data in two subjects with 16 diffusion encoding directions with the conventional sampling (with navigators, Table 4.1. 4A) and the Self-nav CAIPI sampling (without navigators, Table 4.1. 4B). The two $b=0$ images were acquired along opposite phase encoding directions (1 blip-up and 1 blip-down) using conventional rectangular sampling.

Experiment 5 High- b -value Protocol: To demonstrate higher b -values, we acquired data from one subject with the conventional sampling (Table 4.1. 5A, C) and the Self-nav CAIPI sampling (Table 4.1. 5B, D) (both with navigators) at $b=2000$ s/mm^2 (Table 4.1. 5A, B) and $b=3000$ s/mm^2 (Table 4.1. 5C, D).

Experiment 6 Tractography Protocol: To demonstrate tractography analysis, we acquired data in one subject with Self-nav CAIPI sampling and 48 diffusion directions along with 6 $b=0$ images (3 blip-up and 3 blip-down) (Table 4.1. 6).

Reconstruction Details

The image reconstruction was conducted in MATLAB 2021a (Mathworks, Natick, MA, USA). The SPIRiT kernels G_{k_z0} and G_{slab} were trained using a whole-brain gradient echo coil calibration scan (~ 1 min acquisition time). To train G_{k_z0} , the calibration data for each slab were Fourier transformed along the slice direction, and the central $k_z = 0$ plane data were used. The $k_z = 0$ coil sensitivity in Eq. 4.21 was estimated from the same data employed to train G_{k_z0} using ESPIRiT [24]. To obtain G_{slab} , the same calibration data was Fourier transformed along k_x , and G_{slab} was trained for each $k_y - k_z$ plane. All 32-coil k-space data were compressed to 8 coils [66].

For the 3D data reconstruction in Eq. 4.22, the k-space data were first Fourier transformed along k_x followed by reconstruction performed for each $k_y - k_z$ plane using the SPIRiT kernel G_{slab} . The reconstructed 2D images were concatenated along the readout direction to form the whole image volume. The kernel sizes for G_{k_z0} and G_{slab} were both set to 5×5 . The kernel size w_H for constructing $H(\hat{x})$ was set to 10 (i.e., 10×10 kernel). The hyperparameters $\lambda_1, \lambda_3, \lambda_4, \beta$ were empirically selected as 1, 1e-4, 1e-4, 10, respectively. The SVD thresholding for solving Eq. 4.18 was performed by retaining the largest N singular values and their corresponding vectors, where $N = w_H^2$ (i.e., $N = 100$ in this study). The ADMM iteration number was set to 50. The CG iteration numbers for solving Eqs. 4.20 and 4.22 were both set to 30.

The 2D navigator images were reconstructed with 2D GRAPPA. Because the phase variations are expected to be spatially smooth [52], the navigator images were filtered by a k-space Hamming window of size 32×32 to reduce noise. The phase images were extracted as an estimation of motion induced phase errors. The SLR estimated ϕ_d were also filtered using the same Hamming window for the reconstruction in Eq. 4.22.

In Exp 1.4, the implementation of ‘‘Self-nav Conventional’’ followed the steps detailed in [15], except a stronger Gaussian filter with $\sigma = 20$ was used in the final step to eliminate the high-frequency noise in the estimated phase map.

In Exp 1.5, for the coil-combined SLR, the multi-coil k-space data were combined using the $k_z = 0$ coil sensitivity map. The forward model, cost function, and construction of $H(\hat{x})$ followed the original 2D SLR configuration [53]. For the SLR only leveraging coil redundancy, the construction of $H(\hat{x})$ and the optimisation of Eq. 4.14 were conducted in a shot-by-shot manner. Both methods used the ADMM optimisation with the same parameters as mentioned previously. The 16-coil data were selected from the original 32-coil data using a greedy search, with the aim of finding a combination of 16 coils that closely matched the combined profile of all 32. At each step of the search, the mean absolute error between the sum-of-squares of the sensitivity maps from the selected coils and those from the full 32 coils was minimised.

The optimised sampling patterns and the codes for sampling optimisations and reconstructions are openly available at: https://github.com/liziyu0929/Self-nav_CAIP.

Image Analyses

Image post-processing was conducted using the FMRIB Software Library (FSL) [64] unless indicated otherwise. For Exp. 1, the reconstructed images from different methods were evaluated by calculating normalised root mean squared error (NRMSE) with the fully sampled ground truth data. For Exp. 2, 3, and 5, the whole-brain diffusion-weighted data were obtained by directly combining multiple slabs and averaging the overlapped slices. The SNR and SNR efficiency were calculated using the multi-repetition DWIs acquired in Exps. 2 and 3. To mitigate the impact of subject motion on SNR evaluation, all multi-repetition data were co-registered to the first repetition using “flirt” [72] with 6 degrees of freedom and trilinear interpolation. An interim mean image was computed by averaging all co-registered data, and the two volumes with the highest NRMSE with respect to this interim mean image were discarded to reduce motion contamination. The remaining ten-repetition data, denoted as S_{10} , were used for SNR calculation. The voxel-wise SNR map was calculated as $\text{mean}(S_{10}) ./ \text{std}(S_{10})$ ($./$ denotes the element-wise division). The SNR

efficiency (i.e., SNR per unit time) was calculated as $\text{SNR}/\sqrt{\text{TR}}$. The mean SNR and SNR efficiency were then calculated by taking the mean of the SNR and SNR efficiency maps within a brain mask, respectively.

For DTI and tractography experiments (Exp. 4 and 6), slab combination and correction for slab saturation artifacts were performed for the diffusion data using nonlinear inversion of slab profile encoding (NPEN) [62]. Whole-brain images were corrected for Gibbs ringing [69]. A whole-brain field map was estimated using blip-reversed $b=0$ image volumes using “topup” [9], which was then input to “eddy” [70] along with all diffusion data to correct for off-resonance distortions, eddy current effects, and subject motion. The diffusion analyses were conducted in the native diffusion space. For Exp. 4, the diffusion tensor model fitting was performed using “dtifit”. For Exp. 6, white matter tractography was performed using “autoPtx” [76] including probabilistic model fit “bedpostx” and probabilistic tractography “probtrackx”.

4.3 Results

Figure 4.4 demonstrates the image reconstruction results from the simulation data. For conventional sampling, the reconstructed images without phase error correction suffer from severe artifacts and exhibit high NRMSE (Fig. 4.4, iii). Employing navigator phase maps leads to significantly better image quality but requires additional scan time (Fig. 4.4, ii). In contrast, our proposed self-navigated approach avoids the necessity for navigators while achieving comparable efficacy in phase error correction with conventional navigated acquisition (Fig. 4.4, v). The optimised Self-nav CAIPI sampling produces similar image quality with navigator phase (Fig. 4.4, iv) and SLR-estimated phase (Fig. 4.4, v), demonstrating the robust phase variance estimation of the self-navigated approach. The residual maps exhibit random-noise-like error (Fig. 4.4b, d, iv, v) without anatomical structure, likely due to the pseudo-random sampling patterns.

The image reconstruction performance of the proposed method depends on the accuracy of the motion-induced phase variation maps. As show in Fig. 4.5, using




Figure 4.4: Image reconstruction results from the simulation experiment. Retrospectively under-sampled images from Simulation Evaluation (Exp. 1.1, slab-centre slices) are shown, reconstructed using: conventional sampling with ground truth phase maps (ii), conventional sampling without phase error correction (iii), self-navigated CAIPI sampling with ground truth phase maps (iv), and self-navigated CAIPI sampling with structured low-rank estimated phase maps (v), and at 1.22 mm (a) and 1.09 mm (c) isotropic resolutions, with residual maps (b and d) showing difference with the ground truth. The normalised root mean squared errors (NRMSE) of the entire slab are listed.

the proposed SPIRiT SLR reconstruction, the estimated phase maps are highly similar to the ground truth phase maps with consistent spatial patterns for all shots, confirming the accuracy of our method’s motion induced phase error estimations.

Figure 4.6 demonstrates the contribution of each constraint within our sampling optimisation framework. The absence of the overlap constraint (o_i) results in a sampling pattern with significant overlap (Fig. 4.6a), leading to higher NRMSE in




Figure 4.5: Phase map estimation from the simulation experiment. The ground truth phase maps (a) and structured low-rank (SLR) estimated phase maps (without filtering) from the proposed method (b) are shown for different shots of the simulated data from Simulation Evaluation (Exp. 1.1) at 1.22 mm isotropic resolutions. The phase of the shot traversing the entire $k_z = 0$ is subtracted from all shots to eliminate the phase offsets, effectively demonstrating the motion induced phase variations across shots (11 shots demonstrated).

the reconstructed images (Fig. 4.6d, e, iii). Omitting the k-space gap constraint (g_i) produces sampling patterns with large gaps (Fig. 4.6b), exacerbating residual artifacts in the reconstructed images (Fig. 4.6d, e, iv). Removing the self-navigation performance constraint (d_i) yields a uniform sampling pattern with little overlap (Fig. 4.6c). However, the reconstructed images exhibit slightly stronger aliasing and higher NRMSE (Fig. 4.6d, e, v) compared to Self-nav CAIPI, suggesting limitations in accurately estimating 2D phase maps due to insufficient low-frequency information contained in the self-navigation points of each shot.

Figure 4.7 illustrates the impact of including the shot traversing the $k_z = 0$ plane. When the $k_z = 0$ traversing shot is replaced with another CAIPI shot obtained from greedy search in Eq. 4.11, the SLR reconstruction produces a less accurate phase map (Fig. 4.7b, iv) with a slower convergence (Fig. 4.7c). The lower phase accuracy




Figure 4.6: Evaluation of sampling optimisation. Samplings for 1.22 mm acquisition are obtained by optimising the cost function in Eq.4.11 without overlap constraint o_i (a), k-space gap constraint g_i (b), and self-navigation performance constraint d_i (c). Images (d) from Simulation Evaluation (Exp. 1.2, a slab-boundary slice) are under-sampled using the Self-nav CAIPI sampling (ii) and the three samplings in a-c (iii-v) and reconstructed with structured low-rank estimated phase maps, with residual maps (e) showing the difference with the ground truth.

affects the multi-shot reconstruction performance, leading to a higher reconstruction error (Fig. 4.7a, iv). Having the magnitude constraint $\|\hat{x} - Fm'\Phi^{k-1}\|_2^2$ in Eq. 4.20 is beneficial for further accelerating the convergence (Fig. 4.7c).

The proposed method achieves better self-navigation performance than the previous method proposed by Moeller et al. [15] (Fig. 4.8). Although their method provides improved image quality compared to the uncorrected image (Fig. 4.8, ii vs. Fig. 4.8, iv), its estimated phase map does not accurately capture the phase variation (Fig. 4.8c, ii), presumably due to the strong smoothing filter used to suppress the high-frequency noise from shots covering the peripheral k-space in conventional sampling. Self-nav CAIPI produces more accurate phase maps and

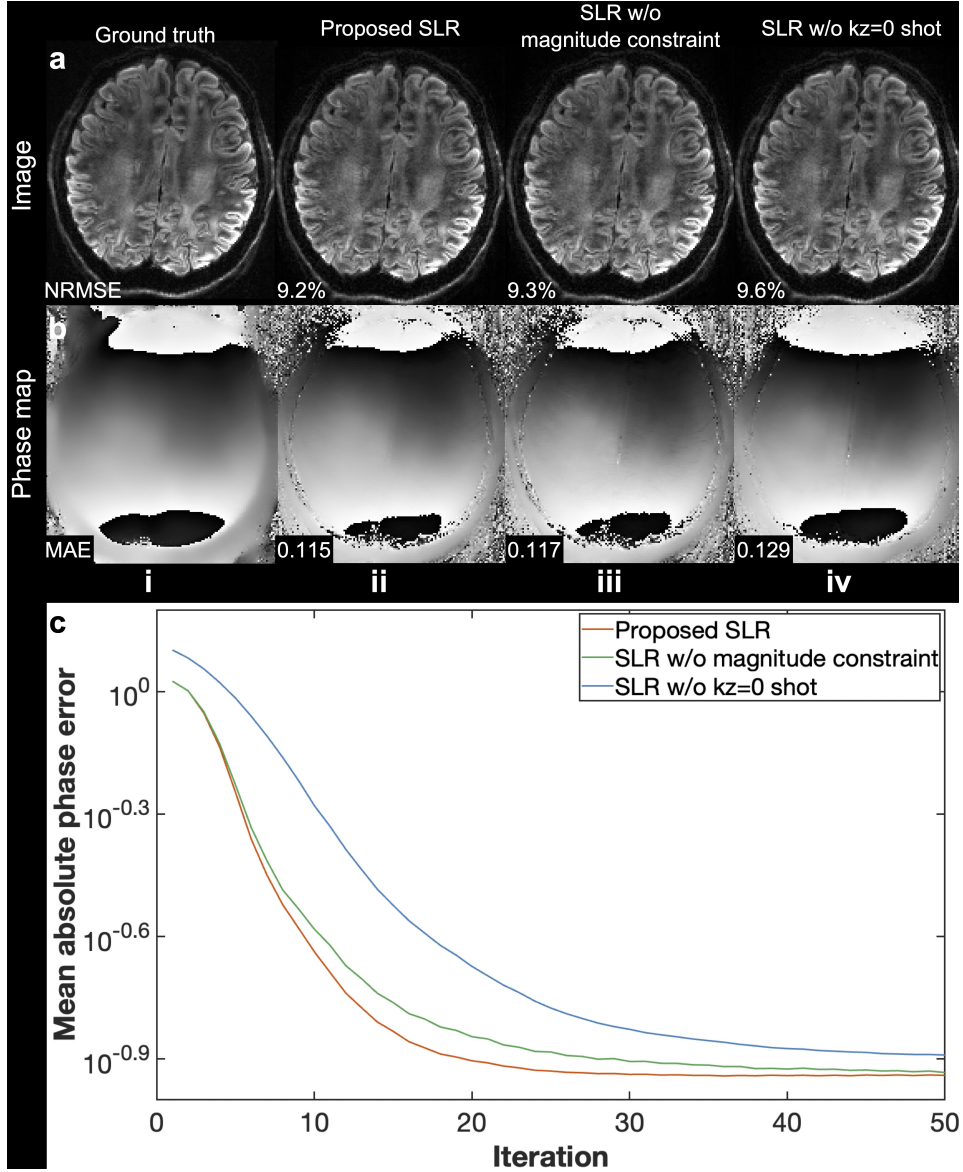


Figure 4.7: Impact of including the shot traversing the $k_z = 0$ plane. Retrospectively under-sampled images (a) and 2D phase maps (without filtering) (b) of one shot at 1.22 mm isotropic resolution from Simulation Evaluation (Exp. 1.3, a slab-centre slice) are shown. Reconstructions using the proposed structured low-rank (SLR) (ii), SLR without the magnitude constraint $\|\hat{x} - Fm'\Phi^{k-1}\|_2^2$ in Eq. 4.20 but with $k_z = 0$ shot data (iii) and SLR without the $k_z = 0$ shot (substituted with a greedy-searched CAIPI shot) (iv) are compared. The mean absolute error (MAE) of phase maps of all shots within the brain mask are provided. The phase MAE plot of different SLR strategies at different ADMM iterations is also shown to demonstrate the convergence of the optimisation (c).

image reconstruction with substantially lower NRMSE (Fig. 4.8, iii).

The inclusion of both shot and coil dimensions benefits the SLR reconstruction (Fig. 4.9). Without leveraging shared information across shots, the SLR fails to




Figure 4.8: Comparison of self-navigation methods. Retrospectively under-sampled images (a), residual maps to ground truth (b), and estimated 2D phase maps of one shot (c) at 1.22 mm isotropic resolution from Simulation Evaluation (Exp. 1.4, a slab-centre slice) from the method proposed by Moeller et al. [15] (entitled “Self-nav Conventional”, ii), our proposed method (Self-nav CAIPI, iii), and the conventional sampling without phase error correction (iv) are shown.

estimate a reliable 2D phase map due to the high under-sampling factor of the self-navigation points for each shot (see Fig. 4.2d), and the image reconstruction is contaminated by the phase error (Fig. 4.9, v). utilising the redundancy across shots substantially benefits the SLR (Fig. 4.9, iv), while exploiting both shot and coil dimensions further improves the phase estimation accuracy, leading to a better image reconstruction (Fig. 4.9, ii). Notably, our method demonstrates robustness when using only 16-coil data (Fig. 4.9, iii), highlighting its compatibility with receive RF coils consisting of fewer channels.




Figure 4.9: Comparison of different reconstruction configurations. Reconstruction of retrospectively under-sampled data (a), residual maps to the ground truth (b), and estimated 2D phase maps (without filtering) of one shot (c) at 1.22 mm isotropic resolution from Simulation Evaluation (Exp. 1.5, a slab-boundary slice) are shown. The 2D phase maps are reconstructed using the proposed structured low-rank (SLR) (ii), the proposed SLR with data from 16 coils (selected from 32 coils) (iii), SLR only leveraging redundancy across shots, i.e., coil-combined SLR (SLR-shots, iv), and SLR only leveraging redundancy across coils (SLR-coils, v). Note for SLR with data from 16 coils, the residual map and NRMSE are calculated based on the 16-coil fully sampled data.

Our method works well on prospectively acquired *in-vivo* data (Fig. 4.10). While the conventional method exhibits significant image artifacts when navigators are not acquired (Fig. 4.10a, ii, iii), our proposed self-navigated method achieves similar image reconstruction results using either navigator phase maps or SLR estimated phase maps (Fig. 4.10b). Notably, the navigator acquired phase map and estimated phase map exhibit high consistency (Fig. 4.10c).

The impact of the $k_z = 0$ traversal shot on reconstruction is notably evident in the *in-vivo* experiment (Fig. 4.11). Comparing reconstructions using the same amount of data from 11 shots, excluding the last CAIPI shot leads to a more accurate reconstruction (Fig. 4.11, ii) compared to exclusion of the $k_z = 0$ shot




Figure 4.10: Reconstruction of prospectively acquired data. *In-vivo* diffusion weighted (diffusion encoding along the slice selection direction) images at 1.09 mm isotropic resolution of a representative subject from Exp. 2 are shown. Reconstruction using conventional rectangular sampling with (a, i) and without navigator (a, ii), and using Self-nav CAIPI sampling with explicitly acquired (b, i) and SLR estimated navigator (b, ii) are compared. The difference maps (iii) between i and ii are also presented. Additionally, the explicitly acquired and SLR estimated navigator phase map of a representative shot of Self-nav CAIPI sampling, along with their difference, are displayed (c).

(Fig. 4.11, iii), indicating the importance of the $k_z = 0$ shot for precise 2D phase map estimation. Remarkably, removal of the last CAIPI shot does not significantly compromise image quality (Fig. 4.11, ii) thanks to the sampling optimisation based on the greedy search framework.

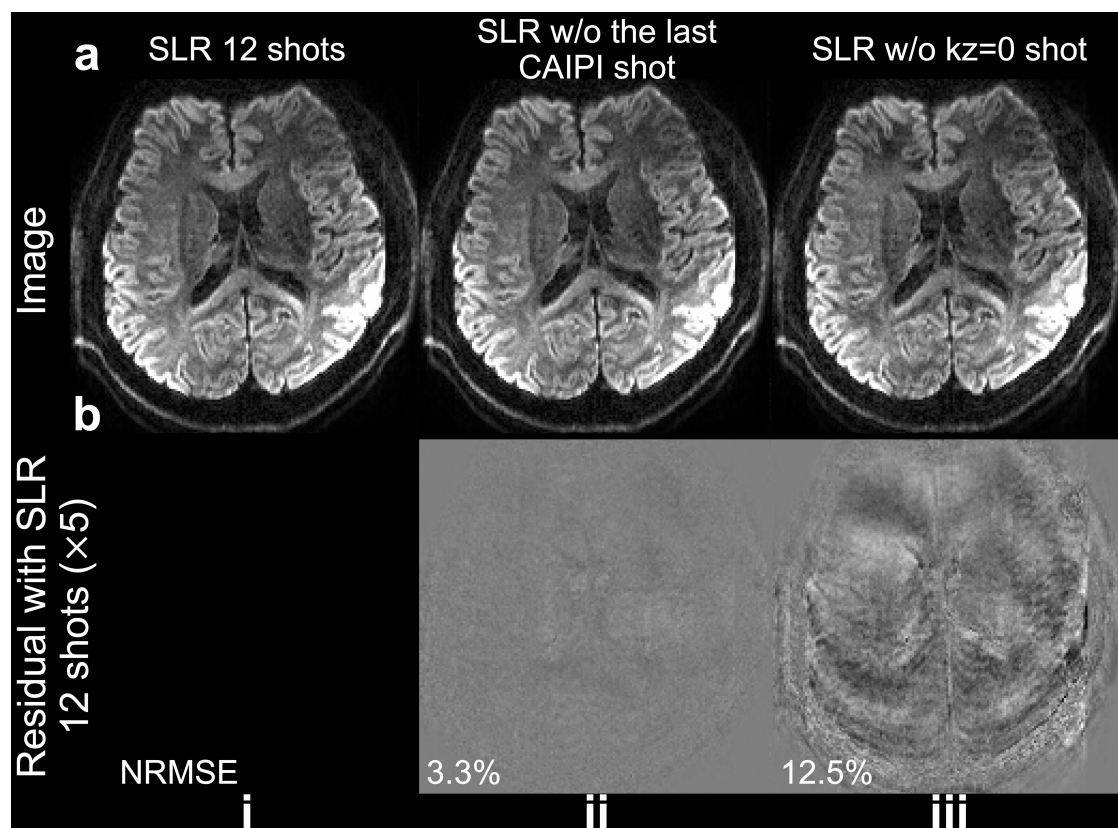


Figure 4.11: Impact of the $k_z = 0$ shot in the *in-vivo* experiment. *In-vivo* diffusion weighted (diffusion encoding along the slice selection direction) images acquired using Self-nav CAIPI sampling at 1.09 mm isotropic resolution of a representative subject from Exp. 2 are shown (a). Images reconstructed using the proposed structured low-rank (SLR) using data from all 12 shots (i), data excluding the last CAIPI shot (ii), and data excluding the $k_z = 0$ shot (iii) are compared, with residual maps (b) showing the difference with the 12-shot reconstruction.

Our proposed method exhibits robust performance along different diffusion encoding directions (Fig. 4.12). When acquired with the same TR, the self-navigated method (Fig. 4.12d) produces SNR values comparable to the conventional navigated method (Fig. 4.12a), even along the slice selection diffusion encoding direction (Fig. 4.12, iii) where motion-induced phase variance is most significant [51]. Notably, the SNR values obtained from our self-navigated method are substantially higher than those from conventional sampling without phase error correction and even marginally higher than those from navigated Self-nav CAIPI sampling. This could be attributed to the higher resolution of SLR reconstructed phase maps compared to navigators (64 phase encoding lines acquired) which may fail to capture high-frequency phase

changes in certain cases (i.e., larger motions). The group-level SNR values for three subjects (Exp. 2, Table 4.2) are consistent with the findings depicted in Fig. 4.12.

	Readout	Phase encoding	Slice selection
Conv. (w/ navigator)	11.66±1.23	11.33±0.37	11.65±0.38
Conv. (w/o navigator)	7.37±0.53	6.91±0.13	6.14±0.33
Self-nav. (w/ navigator)	11.46±0.72	12.46±0.45	10.57±0.73
Self-nav. (w/ SLR navigator)	11.95±0.60	13.04±0.16	11.43±0.81

Table 4.2: Group-level SNR values along different diffusion encoding directions. The group level SNR (mean \pm std) for different diffusion encoding directions for different acquisition and reconstruction methods (Conv.: conventional sampling; Self-nav.: Self-nav CAIPI sampling) of four subjects in Exp. 2 are compared. The highest SNR for each diffusion direction is marked in bold.

Removing the navigator acquisition can shorten TR, leading to an improved SNR efficiency. With the TR reduction, the SNR obtained from the self-navigated method is slightly lower compared to the conventional navigated method (Fig. 4.13, i), as would be expected due to the reduced T_1 recovery. For the four subjects scanned in Exp. 3, the group-level SNR values (mean \pm std) are 11.70 ± 0.74 and 11.18 ± 0.41 for the conventional and proposed method, respectively. In terms of SNR efficiency, the group-level results for the conventional and proposed methods are 6.25 ± 0.39 and 7.22 ± 0.26 , respectively. The self-navigated method improves SNR efficiency by 15.5% compared to the conventional method.

The self-navigated method enables high-quality DTI within a shorter scan time (Fig. 4.13, ii). Our method obtained the DTI maps within 8.6 minutes (Fig. 4.13b, ii), with comparable SNR to the conventional method requiring 12.6 minutes (Fig. 4.13a, ii). By eliminating the navigator acquisition, we save 31.7% of the total scan time without noticeable loss in image quality.

Figure 4.14 demonstrates the performance of our method at higher b-values, a scenario where the diffusion data have lower SNR and are more sensitive to motion-induced phase errors. The Self-nav CAIPI sampling with SLR estimated phase maps (Fig. 4.14, iii) produces similar results to conventional rectangular sampling (Fig. 4.14, i) and Self-nav CAIPI sampling with navigator phase maps (Fig. 4.14, ii). The SLR estimated phase maps (Fig. 4.14, v) are highly similar

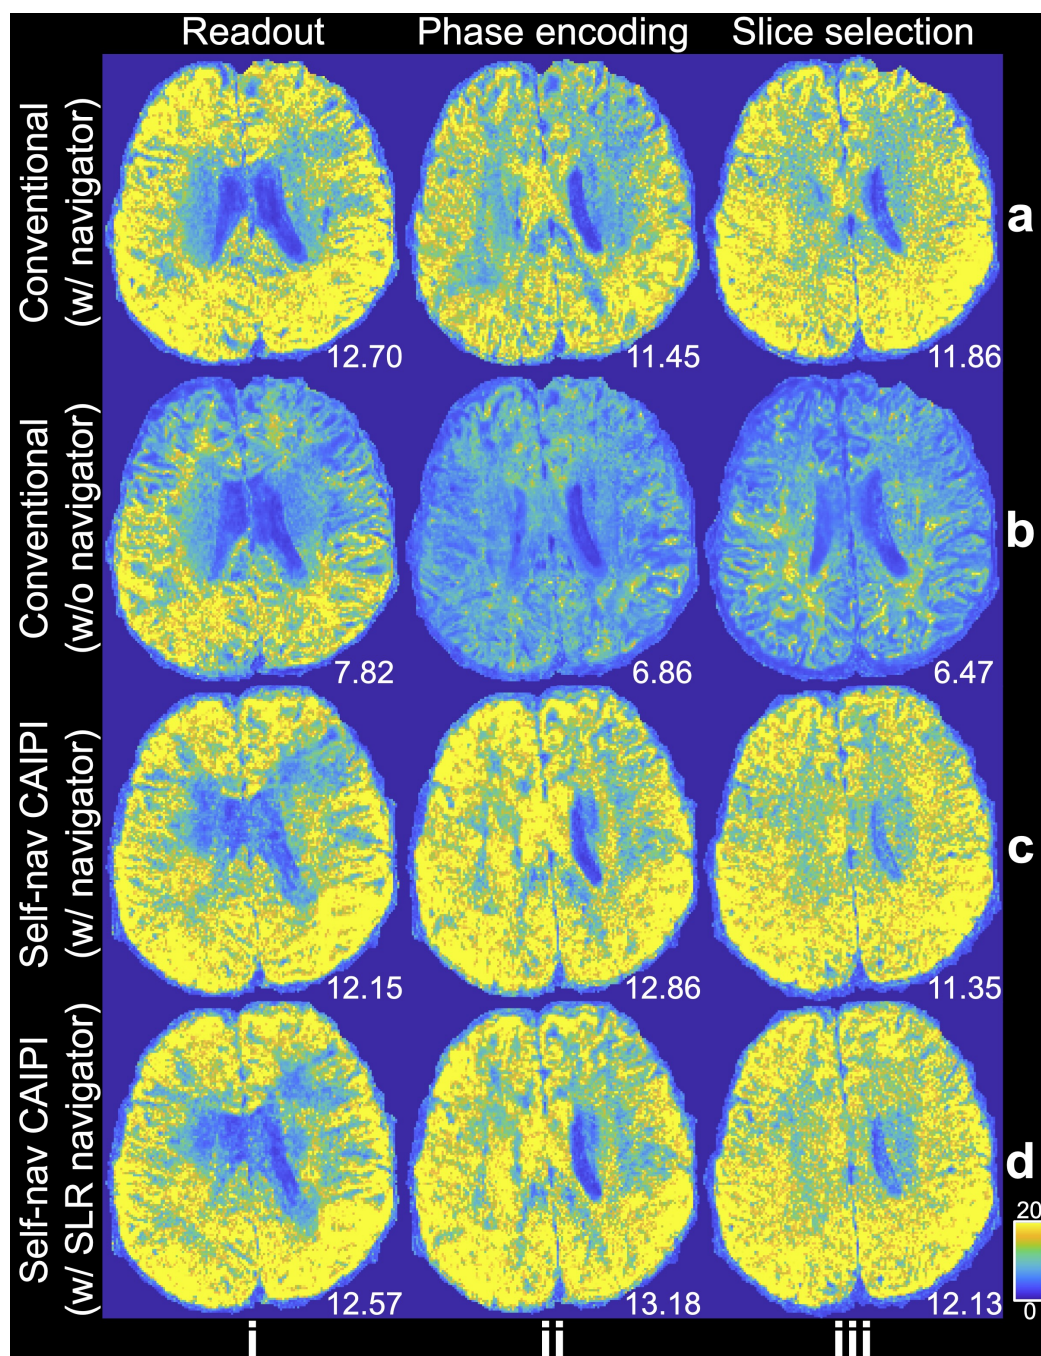


Figure 4.12: SNR maps of different diffusion encoding directions. The SNR maps for diffusion-weighted images (1.09 mm isotropic resolution) along readout (i), phase encoding (ii), and slice selection (iii) diffusion encoding directions from conventional and Self-nav CAIPI acquisition with the same TR of a representative subject from Exp. 2 are demonstrated. Four acquisition and reconstruction methods are compared: (a) conventional rectangular sampling and reconstruction with navigator phase maps. (b) conventional rectangular sampling and reconstruction without phase error correction. (c) Self-nav CAIPI acquisition and reconstruction with navigator phase maps. (d) Self-nav CAIPI acquisition and reconstruction with structured low-rank estimated phase maps. The mean SNR calculated within a brain mask for this subject is listed.




Figure 4.13: SNR efficiency and DTI comparisons. The comparison of SNR maps (Exp. 3, i) and DTI results (Exp. 4, ii, 1.09 mm isotropic resolution) of a representative subject is presented. The SNR maps and DTI results are obtained using the conventional navigated method (a) and the proposed self-navigated method (b). The whole-brain SNR values (calculated within a brain mask) for this subject are listed. Notably, the proposed method has a shorter TR by eliminating navigator acquisition.

to the navigator phase maps (Fig. 4.14, iv), even at high b -values (e.g., $b=3000$ s/mm²). This result demonstrates the robustness of the proposed sampling and reconstruction method at high b -values.

Figure 4.15 demonstrates whole-brain tractography results using the proposed self-navigated method, with 14 tracts illustrated in coronal and sagittal views. The data acquired with the proposed method support the delineation of major fibre bundles.

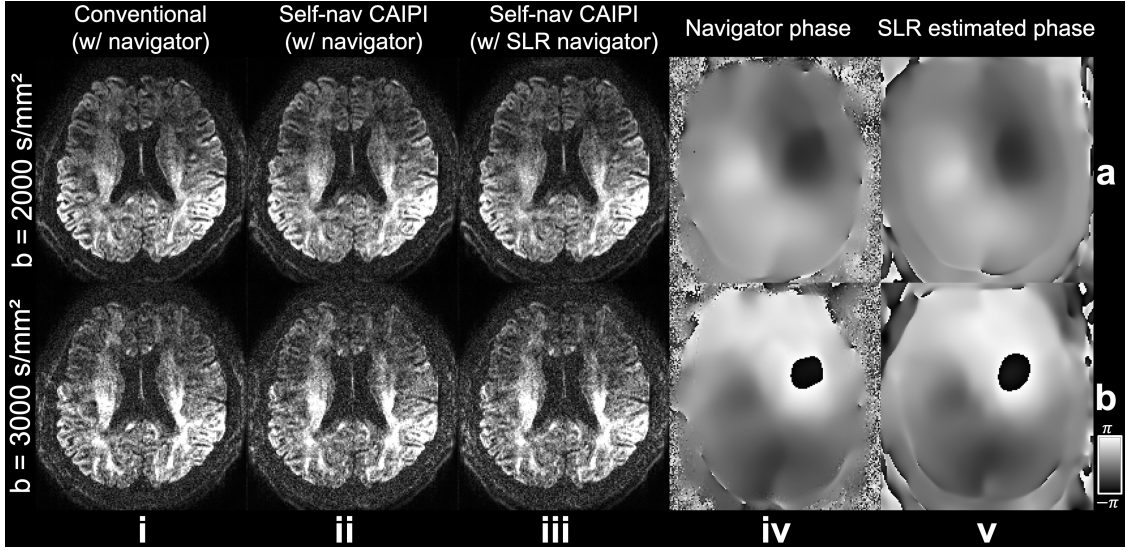


Figure 4.14: Results at high b -values. Diffusion-weighted images (1.09 mm isotropic resolution, diffusion encoding along readout) reconstructed using conventional sampling with navigator (i), Self-nav CAIPI sampling with navigator (ii), and Self-nav CAIPI sampling with structured low-rank estimated phase maps (iii) at $b=2000$ s/mm² (a) and 3000 s/mm² (b) from Exp. 5 are shown. The phase maps of one representative shot for Self-nav CAIPI sampling from the navigator (iv) and the proposed structured low-rank (SLR) reconstruction (v) are also demonstrated.

4.4 Discussion

In this chapter, we propose a novel acquisition and reconstruction framework that extends 2D SLR-based phase correction [53, 82, 84] to 3D multi-slab dMRI to eliminate the need for acquiring navigators. Each shot of our Self-nav CAIPI sampling intersects with the central $k_z = 0$ plane, providing self-navigation points. The sampling is optimised for overlapping, k-space gaps, and self-navigation performance using a greedy search algorithm. Self-nav CAIPI sampling allows us to leverage an SLR reconstruction method that exploits the redundancy across shots and coils to obtain a 2D phase map from self-navigation points for each shot. One shot traverses the entire $k_z = 0$ plane to provide accurate magnitude information, accelerating the convergence and improving the robustness of the reconstruction. *In-vivo* experiments from seven subjects validate the efficacy and robustness of our proposed method, which saves 31.7% of scan time by eliminating the navigator acquisition and achieves 15.5% higher SNR efficiency compared to conventional navigated 3D multi-slab imaging.

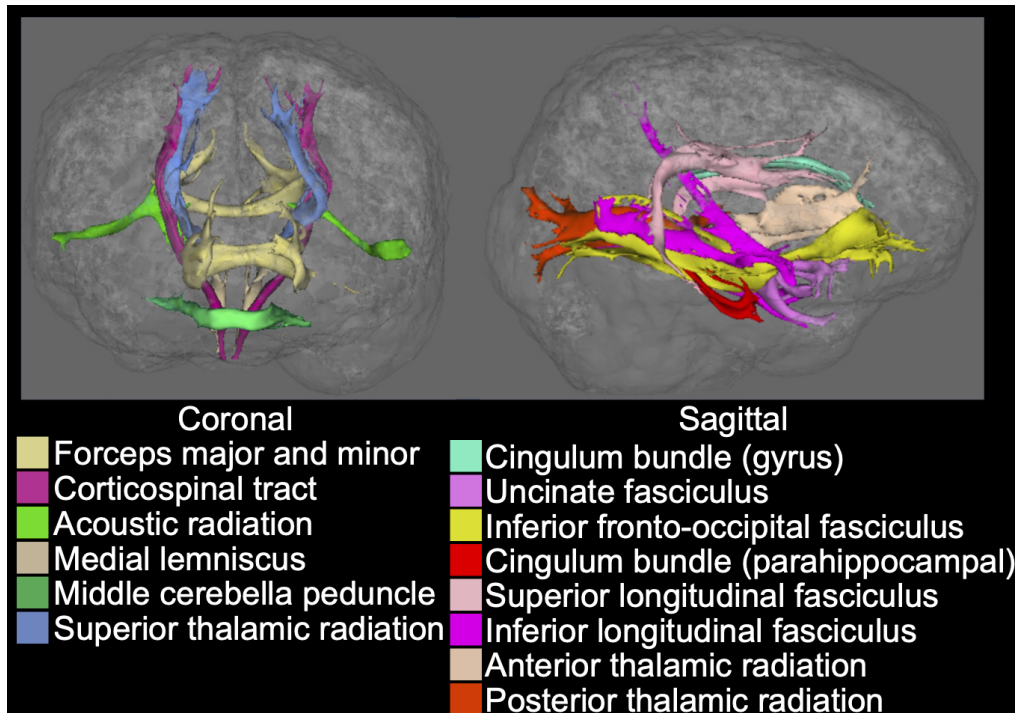


Figure 4.15: Tractography results. The binarised tractography (threshold: 0.01) results obtained from 48-direction diffusion data at 1.09 mm isotropic resolution using our proposed method from Exp. 6 are presented. Fourteen tracts are illustrated in coronal and sagittal views.

The SNR efficiency benefits of eliminating the navigator acquisition are more pronounced for high-field and high-resolution applications. High-field applications often face SAR constraints, necessitating longer TRs. In our baseline protocol, TR=3.5 s is the minimal achievable TR to comply with SAR limits. Eliminating the 180° RF pulse for the navigator echo effectively reduces SAR and the required TR. On the other hand, higher resolution necessitates the acquisition of more slabs, and therefore a longer TR. For instance, a previous study which achieved 0.85 mm isotropic resolution used TR=3 s at 3T [6]. In this case, eliminating navigator acquisition and shortening TR are expected to bring SNR efficiency benefits.

Our sampling design hinges on two critical considerations: optimising for the multi-shot reconstruction and self-navigation performance (i.e., the reconstruction of $k_z = 0$ plane). We developed an optimisation framework to determine the optimal sampling pattern using a shot-by-shot greedy search. Given the substantial parameter space (with $s_{k_z} \in [0, 5]$, $s_{k_y} \in [0, 2]$, $s_p \in [0, 11]$), conducting a global

search across 11 shots would involve evaluating $(6 \times 3 \times 12)^{11} \approx 4.78 \times 10^{25}$ sampling patterns. Our greedy search approach significantly reduces the search space to $6 \times 3 \times 12 \times 11 = 2376$ while achieving satisfactory reconstruction performance. Additionally, the uniform and random samplings may benefit compressed sensing reconstruction methods [86], indicating that adding tailored constraints in Eq. 4.22 may further enhance performance. Furthermore, our sampling strategy holds potential for additional acceleration by acquiring fewer shots (i.e., $N_{\text{shot}} < N_{k_z}$) without significant image quality loss (Fig. 4.11), as the greedy search-based sampling optimisation ensures efficient k-space coverage even with fewer shots. Recent studies have similarly demonstrated that CAIPI sampling optimisation can facilitate accelerated 3D multi-slab diffusion MRI [54, 90], supporting its benefits in under-sampled diffusion MRI reconstruction.

Our alternation in the sampling pattern is not expected to introduce significant additional geometric distortions or blurring compared to the conventional sampling. In our periodic sampling, the accumulations of B_0 induced phase and T_2^* decay along the slice direction is effectively counteracted during the k_z -blip-up and k_z -blip-down steps (Fig. 4.2d), thereby minimising additional geometric distortions and blurring along the slice direction. Consequently, our prospectively acquired *in-vivo* data using Self-nav CAIPI exhibit no significant additional distortions or blurring when compared to the conventional sampling (Figs. 4.10, 4.13, 4.14).

Due to the limited number of self-navigation points, the acceleration factor for $k_z = 0$ reconstruction is extremely high (R=18 or 36, Fig. 4.2d). The substantial under-sampling factor necessitates the utilisation of shared information from other shots for accurate reconstruction (Fig. 4.9, v). Incorporating the coil dimension exploits the inherent low-rankness between coils due to the smooth nature of coil sensitivity, thereby benefiting parallel imaging reconstruction (Fig. 4.9, iv) [28]. The impact is expected to be more pronounced for low-SNR *in-vivo* data. The shot traversing $k_z = 0$ provides an accurate magnitude estimation, reducing the unknowns from the complex image to only the phase image. This is proved beneficial for accelerating the convergence and improving the robustness of the reconstruction

(Figs. 4.7, 4.11). The optimised self-navigation performance (i.e., d_i in Eq. 4.11) ensures that each shot contains sufficient low-frequency information to produce a reliable phase map. These attributes enable robust reconstruction of 2D phase maps from limited data, even with low-SNR (e.g., high b-values, Fig. 4.14).

Our method outperforms the previous self-navigation method [15] (Fig. 4.8), which used conventional sampling and extracted a phase map from each k_z plane. We found our data necessitated a stronger filter than the original work to suppress the high-frequency noise, presumably due to the larger number of k_z planes per slab used in our work (12 k_z vs. 8 k_z). The estimated phase maps are over-smoothed and fail to accurately capture the phase variation (Fig. 4.8c, ii). This also highlights the intrinsic challenge to extract smooth and accurate phase maps from peripheral k_z planes that encode high-frequency information, which might limit the performance and robustness of the previous method.

The SLR estimated phase maps demonstrate slightly better reconstruction performance than navigator-acquired phase maps in prospective *in-vivo* experiments (Fig. 4.8, Table 4.2). However, in Exp. 1, the structured low-rank reconstruction exhibits slightly higher NRMSE compared to navigated reconstruction (Fig. 4.4). This is presumably because in Exp. 1 the navigator phase maps are ground truth phase maps that were retrospectively added to phase error-free data. In prospective acquisitions, navigators contain some measurement error due to the extended echo time of their acquisition and the associated SNR loss. It has been demonstrated that reducing the resolution along the readout direction and decreasing the acceleration factor along the phase encoding direction can effectively mitigate such SNR loss and measurement error [59, 91]. Furthermore, compared to the conventional method where the navigators are acquired with lower resolutions, our SLR reconstructed phase maps have the same resolution as the images, which might be helpful in capturing motion-induced phase variations more accurately.

The inclusion of both shot and coil dimensions contributes to the substantial size of the Hankel matrix $H(\hat{x})$ (Fig. 4.3), making the current reconstruction time-consuming. Each iteration of the ADMM takes approximately 8 minutes on

a 2.9 GHz Quad-Core Intel Core i7 CPU, resulting in a reconstruction time of ~ 7 hours for a 3D slab on a single CPU. The inclusion of the $k_z=0$ traversing shot and the magnitude constraint $\|\hat{x} - Fm'\Phi^{k-1}\|_2^2$ may help reduce the number of iterations without compromising the reconstructed phase map quality (Fig. 4.7c). Using the coil-combined SLR (i.e., “SLR – shots” in Fig. 4.9, iv) formulation can also effectively shorten the computation time of each iteration to 3 minutes with a slight performance compromise. Model-based deep learning reconstruction approaches [84] may also offer a promising avenue for accelerating the reconstruction process.

It is also worth noting that our method primarily addresses phase error induced by physiological motions such as respiration and cardiac pulsations. Additional considerations are necessary to address larger-scale motions. For example, inter-shot bulk motions violate the assumption of the SLR that the underlying $k_z = 0$ magnitude images are consistent across different shots. Recent advances [92] which integrate rigid motion compensation into SLR hold promises in addressing such motions without acquiring navigators. However, for large-scale motion particularly along the slice direction, the 2D phase maps might be no longer sufficient for accurately capturing the motion. In this case, data rejection and reacquisition [46] or prospective motion correction methods [93] should be considered.

4.5 Conclusion

We present a novel acquisition and reconstruction framework that eliminates the requirement for navigator acquisition in high-resolution 3D multi-slab dMRI. It effectively shortens the TR and reduces SAR, enhancing scan efficiency and safety, and permitting the use of better RF pulses. The effectiveness of our method is evident in high-fidelity fast DTI and tractography and can be explored in more applications.

5

Submillimetre Diffusion MRI Using an In-Plane Segmented 3D Multi-Slab Acquisition and Denoiser-Regularised Reconstruction

Contents

5.1	Introduction	108
5.2	Methods	111
5.2.1	Submillimetre 3D Diffusion Imaging Framework	111
5.2.2	In-Plane Segmented 3D Multi-Slab Acquisition	111
5.2.3	Denoiser-Regularised Reconstruction	114
5.2.4	<i>In-Vivo</i> Experiments	115
5.2.5	Reconstruction Details	117
5.2.6	Diffusion Analyses	119
5.3	Results	120
5.4	Discussion	132
5.5	Conclusion	137

High-resolution diffusion MRI (dMRI) provides valuable insights into brain microstructure, particularly at submillimetre resolutions, where it enables more precise delineations of curved and crossing white matter pathways. However, achieving high-quality submillimetre dMRI *in-vivo* poses significant challenges due to the intrinsically low SNR, along with the long echo spacing, readout time, and

TE required for the large matrix size, leading to significant image distortion, T_2^* blurring, and T_2 signal decay. In this chapter, we propose a novel acquisition and reconstruction framework to overcome these challenges. Based on numerical simulations, we introduce an in-plane segmented 3D multi-slab acquisition that leverages the optimal SNR efficiency of 3D multi-slab imaging while reducing echo spacing, readout times, and TE using in-plane segmentation. This approach minimises distortion, improves image sharpness, and enhances SNR. Additionally, we develop a denoiser-regularised reconstruction to suppress noise while maintaining data fidelity, which reconstructs high-SNR images without introducing substantial blurring or bias. Comprehensive *in-vivo* experiments demonstrate that our method consistently produces high-quality dMRI data at 0.65 mm and 0.53 mm isotropic resolutions on a 3T clinical scanner. The submillimetre dMRI datasets reveal richer microstructural details, reduce gyral bias, and improve U-fibre mapping compared to prospectively acquired 1.22 mm diffusion data. Our method demonstrates robustness at 7T and generates high-SNR 0.61 mm diffusion datasets, showing excellent agreement with previous *post-mortem* studies at the same scanner. Implemented using the open-source, scanner-agnostic framework Pulseq, our approach may facilitate broader adoption across different scanner platforms to benefit a wider range of applications. These results underscore the potential of our method to advance medical image analysis and neuroscientific research of the human brain. The preprint of this chapter is available on *bioRxiv* [94].

5.1 Introduction

High-resolution diffusion MRI (dMRI) offers a powerful tool for investigating detailed brain microstructure. The benefits of submillimetre dMRI have been demonstrated in *post-mortem* studies [2, 95], where it provides more precise delineations of curved and crossing white matter pathways (e.g., transverse pontine fibres) compared to conventional resolutions (e.g., 2 mm). Additionally, submillimetre dMRI may help address a known challenge in dMRI fibre tracking (tractography) called “gyral bias”, where tracked fibres tend to terminate at gyral crowns instead of accurately following

the gyral walls [96, 97]. Higher spatial resolutions are also advantageous for the identification of short cortical association fibres, commonly known as U-fibres, which connect cortical regions between adjacent gyri [98]. This is particularly important for studying the brain’s structural connectivity [99]. Furthermore, high-resolution dMRI holds great promise for precisely delineating small but crucial subcortical structures, which can serve as targets for deep brain stimulation in treating conditions such as Parkinson’s disease, essential tremor, and dystonia [100, 101].

Despite these appealing advantages, achieving high-resolution *in-vivo* dMRI, especially at submillimetre levels, presents considerable challenges. For the conventional 2D single-shot EPI, a major hurdle is the intrinsically limited SNR due to the small voxel size and the long TE associated with the large number of phase encoding lines. The long TR necessary to cover a large number of slices reduces SNR efficiency, which is optimised with TR=1-2 s for spin echo-based dMRI [5]. Additionally, the large imaging matrix size also results in extended echo spacing and readout times, leading to significant image distortion that compromises image anatomical fidelity, and T_2^* blurring that affects effective resolution [102]. Moreover, the limitations in gradient strength for delivering slice-selective RF pulses to resolve thin slices constrain the achievable through-plane resolution.

To overcome these challenges, previous studies have proposed super-resolution-based approaches that have shown potential in enabling submillimetre dMRI. One notable technique is gSlider [14, 47, 85], which uses 2D acquisitions to excite a thin slab multiple times with distinct RF pulses with complementary slice profiles, allowing thin slices to be resolved from the slab (typically 5 slices per slab). Coupled with the simultaneous multi-slice (SMS) technique, gSlider-SMS can effectively shorten the TR and improve SNR efficiency. More recently, super-resolution-based submillimetre dMRI by acquiring multiple thick-slice volumes with rotated field-of-view (FOV) to resolve a high-resolution volume has also been proposed [50]. However, a potential drawback of these super-resolution methods is the blurring effect that can result from the use of conditioning techniques (e.g., Tikhonov regularisation) when solving the inverse problem. Furthermore, the TRs of these

methods, while shorter than conventional 2D EPI, still typically exceed 3 seconds and remain outside the optimal range for SNR efficiency.

An alternative and promising approach for achieving high-resolution, SNR-efficient dMRI is 3D multi-slab imaging, which is compatible with shorter TRs of 1-2 seconds [4, 5, 7, 54, 80]. This technique divides the whole brain into several slabs, each with a thickness of less than 2 cm to ensure that motion-induced phase variations can be effectively captured by a 2D navigator [4]. Within each slab, 3D EPI is typically employed for encoding, offering superior SNR and image sharpness due to the use of orthogonal Fourier bases. However, conventional 3D multi-slab imaging often uses a single shot to cover one k_z plane, resulting in extended echo spacing, readout time, and TE, which can lead to significant distortions, T_2^* blurring, and SNR penalties, particularly at submillimetre resolutions.

In this chapter, we present a novel acquisition and reconstruction framework designed to achieve submillimetre *in-vivo* dMRI. We leverage 3D multi-slab imaging for superior SNR efficiency, with in-plane segmented EPI used to shorten echo spacing, readout times, and TE, thereby reducing distortion, T_2^* blurring, and improving SNR. A denoiser-regularised reconstruction approach is proposed to effectively suppress noise during reconstruction while maintaining consistency with the acquired data to avoid biases. *In-vivo* experiments at 3T demonstrate that our method produces high-quality submillimetre dMRI at 0.65 mm and 0.53 mm isotropic resolutions. The efficacy of our approach is further validated at a higher field strength of 7T, where it consistently delivers high-quality data. Comprehensive diffusion analyses of the acquired submillimetre datasets reveal significantly richer microstructural details, underscoring our method’s potential to address the “gyral bias” and improve U-fibre tracking compared to conventional 1.22 mm resolution data. These findings hold considerable promise for advancing neuroscientific research into the human brain.

5.2 Methods

5.2.1 Submillimetre 3D Diffusion Imaging Framework

Our goal is to develop a 3D multi-slab EPI framework combined with denoiser-regularised reconstruction to achieve high-fidelity, high-SNR submillimetre *in-vivo* dMRI. Conventional 3D multi-slab EPI, where each k_z plane is acquired in a single shot, is constrained by SNR penalties, T_2^* blurring, and geometric distortions, particularly for high-resolution imaging with large matrix sizes. To overcome these issues, we introduce in-plane segmentation to cover each k_z using multiple segments, which reduces echo spacing, readout duration, and TE, effectively mitigating these challenges. A numerical simulation framework is employed to quantify the impact of in-plane segmentation on SNR and effective resolution, guiding the determination of the number of in-plane segmentations.

For reconstruction, we leverage the SNR enhancement provided by image denoising methods. Rather than using the denoised image directly as the final output, we integrate it as a regularisation term within a model-based reconstruction framework. This approach ensures noise suppression while enforcing data consistency, thereby improving SNR without introducing significant blurring or bias and preserving overall image fidelity.

5.2.2 In-Plane Segmented 3D Multi-Slab Acquisition

In conventional 3D multi-slab diffusion imaging, each shot typically covers a k_z plane using a single EPI readout for efficient data acquisition [4, 7, 59]. However, for high-resolution EPI with a large matrix size, this single-shot method can result in extended echo spacing, readout durations, and TE, leading to significant image distortion, T_2^* blurring, and reduced SNR.

In-plane segmented multi-shot acquisitions hold promise in mitigating these issues. Previous studies have proposed readout-segmented 3D multi-slab EPI [5, 62], but these methods have not demonstrated submillimetre-level resolution, likely due to limitations in maximum slew rate. An alternative approach is phase encoding

segmented EPI, where each shot undergoes a high under-sampling factor along the phase encoding direction (k_y). This effectively shortens the effective echo spacing, readout duration, and TE. The effective echo spacing and associated image distortion decreases inversely proportional to the number of in-plane segments (N_{seg}). For the effective resolution and SNR, we quantitatively assess how they are affected by N_{seg} at 0.6 mm and 1 mm resolutions, assuming the data are fully sampled with effective acceleration factor $R_{\text{eff}} = N_{\text{seg}}/N_{\text{acq}} = 1$, where R_{eff} is the effective acceleration factor and N_{acq} is the number of acquired segments. We extend the single-shot simulation framework by Feizollah and Tardif [102] for multi-segment acquisitions, modifying the parameters to reflect realistic 3D multi-slab dMRI acquisitions: TR=2.5 s, maximum gradient amplitude $G_{\text{max}}=80$ mT/m, RF excitation duration=6 ms, RF refocusing duration=10 ms, b-value=1000 s/mm², and bandwidths of 992 Hz and 1384 Hz for 0.6 mm and 1 mm acquisitions, respectively. The T_1 , T_2 , and T_2^* values are set to 800/79.6/53.2 ms at 3T and 1200/47/26.8 ms at 7T for white matter [102].

The effective resolution is quantified by the full-width-half-maximum (FWHM) of the point spread functions simulated under various acquisition parameters. For the partial Fourier (PF) sampling, conjugate symmetric filling and zero-padding are investigated. The SNR is quantified based on:

$$\text{SNR} \propto B_0(\Delta x \Delta y \Delta z) e^{-TE/T_2} (1 - e^{-TR/T_1}), \quad (5.1)$$

where B_0 is the main magnetic field strength, Δx , Δy , and Δz are the effective resolution along readout, phase-encoding, and slice-selection directions, respectively, and $(\Delta x \Delta y \Delta z)$ is the effective voxel size. While Δy is determined by FWHM, Δx and Δz are assumed to match the nominal resolution. The TE is simulated under different acquisition conditions.

As detailed in the Results section, the simulations indicate that a large N_{seg} is necessary for improving effective resolution and SNR. Based on our simulation results, we select N_{seg} of 6 for 3T and 8 for 7T acquisitions of submillimetre diffusion data, by considering the balance between scan time, effective resolution, and SNR.




Figure 5.1: Impact of the sampling order of in-plane segmented 3D EPI. The “ $k_z - k_y$ ” (a, all k_z planes for one k_y segment are acquired before proceeding to the next segment) and “ $k_y - k_z$ ” (b, all k_y segments for one k_z plane are acquired before proceeding to the next k_z) sampling orders are illustrated with a simplified 4-segment example. Representative $b=0$ images (0.65 mm isotropic resolution, 6-segment acquisition) are demonstrated below the corresponding sampling orders.

For the sampling order of in-planed segmented 3D EPI, we choose to acquire all in-plane k_y segments of a given k_z in immediate succession before moving to the next k_z (“ $k_y - k_z$ ” in Fig. 5.1b). Compared to the alternative order where all k_z planes for one k_y segment are acquired before proceeding to the next segment (“ $k_z - k_y$ ” in Fig. 5.1a), our chosen sampling order reduces the inconsistencies between in-plane segments and is believed to be beneficial for improving the motion robustness and reducing the artifacts [103, 104], especially for the $b=0$ image acquisition due to the pronounced cerebrospinal fluid (CSF) signal aliasing introduced by inter-segment inconsistency (Fig. 5.1).

5.2.3 Denoiser-Regularised Reconstruction

The accuracy of high-resolution dMRI reconstruction is often limited by the inherently low SNR. Image denoising is a common strategy to enhance SNR, with various methods available both for general natural images (e.g., Non-Local Means [105], BM3D [106], BM4D [107]) and specifically for dMRI (e.g., MPPCA [108], NORDIC [109]). However, these methods typically function as standalone post-processing steps, which can introduce biases and blurring during dMRI denoising [110].

Our approach integrates denoising within a SPIRiT-based [23] multi-shot reconstruction framework, aiming to suppress noise while minimising biases and blurring by enforcing data consistency:

$$\arg \min_{\hat{x}} \sum_i^{N_{\text{shot}}} \|D_i F P_i^H F^{-1} \hat{x} - y_i\|_2^2 + \lambda_1 \|(G - I)\hat{x}\|_2^2 + \lambda_2 \|F^{-1} \hat{x} - \Phi(F^{-1} \hat{x})\|_2^2, \quad (5.2)$$

where \hat{x} is the desired fully-sample, high-SNR, multi-coil k-space data to be reconstructed, N_{shot} is the total number of shots, D_i is the shot-sampling mask for the i^{th} shot, F is the Fourier Transform, P_i^H is the conjugate of 2D navigator phase for addressing motion-induced phase variance [4, 51] for the i^{th} shot, y_i is the acquired data for the i^{th} shot, G is the SPIRiT kernel trained on coil calibration data, I is the identity matrix, Φ is the denoiser, and λ_1, λ_2 are the weights for SPIRiT and denoiser regularisations, respectively.

Directly solving Eq. 5.2 is challenging due to the computational complexity of differentiating Φ , especially for modern sophisticated denoisers. We leverage the “plug-and-play” approach [111], iteratively alternating image denoising with forward-model-based reconstruction. At the k^{th} iteration:

$$\hat{x}^k = \arg \min_{\hat{x}} \sum_i^{N_{\text{shot}}} \|D_i F P_i^H F^{-1} \hat{x} - y_i\|_2^2 + \lambda_1 \|(G - I)\hat{x}\|_2^2 + \lambda_2 \|F^{-1} \hat{x} - \Phi(F^{-1} \hat{x}^{k-1})\|_2^2, \quad (5.3)$$

which enables the incorporation of advanced denoisers while enforcing the data consistency by minimising $\sum_i^{N_{\text{shot}}} \|D_i F P_i^H F^{-1} \hat{x} - y_i\|_2^2$ and the consistency with the

calibration data by minimising the SPIRiT constraint $\|(G - I)\hat{x}\|_2^2$. The proposed denoiser-regularised SPIRiT reconstruction is referred to as “DnSPIRiT” hereafter.

5.2.4 *In-Vivo* Experiments

We adapted a 3D multi-slab spin-echo diffusion MRI sequence [7] to incorporate in-plane segmented acquisitions following the sampling order illustrated in Fig. 5.1a. The sequence was implemented on Siemens platform as well as an open-source, scanner agnostic framework “Pulseq” [112] to promote broader accessibility and applications. *In-vivo* experiments at 3T (Siemens Prisma, Erlangen, Germany) and 7T (Siemens Magnetom, Erlangen, Germany) were conducted, both with 32-channel receive coils. Written informed consents in accordance with local ethics were obtained from subjects before scanning. Submillimetre dMRI data at 0.65 mm, 0.61 mm, and 0.53 mm isotropic resolutions were acquired using the following protocols, with key parameters listed in Table 5.1:

Res. (mm ³)	B_0	Matrix size	PF	TE1/TE2/TR(ms) ^a	$N_{\text{seg}}/N_{\text{acq}}$ ^b	ES _{eff} (ms) ^c	T_{acq} per dir.
0.65	3T	336×336×164	1	102/195/2500	6/6	0.21	6 min
0.53	3T	414×414×186	3/4	80/201/2000	6/3	0.22	2.7 min
0.61	7T	360×360×174	1	87/153/2600	8/8	0.15	7.6 min

Table 5.1: Submillimetre dMRI acquisition parameters. a. TE1 is the imaging echo time and TE2 is the navigator echo time. b. For the 3T 0.53 mm Protocol, the 3 acquired k_y segments are evenly spaced. c. ES_{eff} refers to the effective echo spacing (i.e., echo spacing/ N_{seg}).

3T 0.65 mm Protocol: A 6-direction, fully-sampled dataset at 0.65 mm isotropic resolution was acquired at 3T to validate the efficacy of the proposed acquisition and reconstruction framework. Nine slabs were acquired, each containing 20 slices, with a 20% oversampling along the k_z direction to mitigate slab boundary aliasing, resulting in 24 slices per slab. Adjacent slabs were overlapped by 2 slices, leading to a total of 164 slices in the final reconstruction. Slab excitation was achieved using a Shinnar-Le-Roux (SLR) pulse [113] with a time-bandwidth product of 8 for excitation and 12 for refocusing. The pulse durations were 6 ms for excitation and 9 ms for refocusing, with flip angles set at 90° for excitation and 160° for

refocusing to reduce slab boundary saturation artifacts. The navigator was acquired for each shot with a lower parallel imaging acceleration factor of $R_{\text{nav}} = 4$ for more robust reconstruction. Since the motion-induced phase is typically smooth [52], the distortion difference between the image and the navigator is not expected to significantly affect the reconstruction quality. Diffusion-weighted images were acquired along anterior-posterior phase-encoding direction with $b=1000 \text{ s/mm}^2$ and 6 diffusion directions with 2 $b=0$ image volumes along opposite phase-encoding directions (i.e., anterior-posterior and posterior-anterior) for distortion correction. The total scan time was 48 minutes.

3T 0.53 mm Protocol: To enable more advanced diffusion analyses at an ultrahigh resolution, a 20-direction dataset at 0.53 mm isotropic resolution with $R_{\text{eff}} = 2$ was acquired at 3T. For this protocol, the number of slab was 8 with 25 slices per slab, with 8% oversampling along k_z applied (27 acquired slices per slab). neighbouring slabs were overlapped by 2 slice, resulting in 186 slices, providing a ~ 100 mm coverage. An SLR pulse with matched parameters with the 3T 0.65 mm Protocol was applied for slab excitation. Partial Fourier (PF) with a factor of 3/4 was applied along the phase encoding direction to further shorten the TE. The navigator was acquired with $R_{\text{nav}} = 4$. Diffusion-weighted images were acquired along anterior-posterior phase-encoding direction with $b=1000 \text{ s/mm}^2$ and 20 diffusion directions. Two fully-sampled $b=0$ image volumes along opposite phase-encoding directions were acquired. The total scan time was 65 minutes.

7T 0.61 mm Protocol: To demonstrate the robustness of our method at higher field strength, a 6-direction dataset at 0.61 mm isotropic resolution was acquired at 7T. Nine slabs, each containing 20 slices, were used, with 10% oversampling along k_z (22 slices per slab). Adjacent slabs were overlapped by 2 slices, resulting in 174 slices in total. A sinc pulse with excitation and refocusing duration=6 and 8 ms was used for slab excitation. The navigator was acquired with $R_{\text{nav}} = 6$. Diffusion-weighted images were acquired along anterior-posterior phase-encoding direction with $b=1000 \text{ s/mm}^2$ and 6 diffusion directions and 2 $b=0$ images. The total scan time was 61 minutes.

For comparison, dMRI data at conventional resolutions of the same subjects were acquired. At 3T, a dMRI dataset at 1.22 mm isotropic resolution was obtained using conventional 3D multi-slab EPI with the following parameters: matrix size= $180 \times 180 \times 91$, $R=3$, $PF=3/4$, $TE1/TE2/TR=70/150/2700$ ms, effective echo spacing= 0.27 ms, number of slabs= 10 , slices per slab= 10 , k_z oversampling= 20% , number of overlapped slices between slabs= 1 , T_{acq} per volume= 32.4 s, 24 diffusion encoding directions ($b=1000$ s/mm²) and 4 $b=0$ images (with 2 images acquired along the opposite phase encoding direction), total scan time= 15.1 minutes. At 7T, another dMRI dataset at 1.05mm isotropic resolution was acquired using the 3D multi-slab acquisition described in Chapter 3 [54] with matrix size= $210 \times 210 \times 115$, 48 diffusion encoding directions ($b=1000$ s/mm²) and 6 $b=0$ images, $R=3$, $TE1/TE2/TR = 82/150/1800$ ms, $T_{\text{acq}}=36$ s per volume and 33 minutes in total. Additionally, a T_1 -weighted (T1w) anatomical image at 0.7 mm isotropic resolution was acquired at 3T using magnetization-prepared rapid gradient-echo imaging (MPRAGE) [114] following the protocol from the Human Connectome Project (HCP) [115].

5.2.5 Reconstruction Details

The image reconstruction was conducted in MATLAB 2021a (Mathworks, Natick, MA, USA). The SPIRiT kernels G in Eq. 5.3 was trained using a whole-brain gradient echo coil calibration scan (~ 2 min acquisition time). All 32-coil k-space data were compressed to 8 coils [66].

The 2D navigator images were reconstructed with 2D GRAPPA [22] and subsequently filtered using a k-space Hamming window of size 32×32 to reduce noise. Phase images were then extracted as estimates of motion-induced phase errors (i.e., P_i^H in Eq. 5.3). The reconstruction in Eq. 5.3 was performed for each $k_y - k_z$ plane after the k-space data being first Fourier transformed along k_x . The reconstructed 2D images were concatenated along the readout direction to form the whole image volume. The kernel size for G was set to 5×5 . Equation 5.3 was solved using a

conjugate gradient (CG) method. The first iteration of Eq. 5.3 essentially performed a standard SPIRiT reconstruction with motion-induced phase error correction [80].

For the 6-direction 3T 0.65 mm and 7T 0.61 mm protocols, BM4D [107] was employed as the denoiser, with the number of iterations (N_{iter}) set to 5 (i.e., $k \in [1, 5]$) and hyperparameters $\lambda_1 = 10$ and $\lambda_2 = 2$. During each iteration, the sum-of-squares of $F^{-1}x$ was denoised by BM4D using adaptive noise level estimation with Rician noise distribution, while other parameters were kept at their default values. The denoised magnitude image was then multiplied by coil sensitivities, estimated from the coil calibration data using ESPIRiT [24], to generate multi-coil complex data for the next iteration.

For the 20-direction 3T 0.53 mm protocol, NORDIC [109] was used as the denoiser, leveraging shared information across diffusion directions. In this case, N_{iter} was set to 2, with $\lambda_1 = 10$ and $\lambda_2 = 3$. In each iteration, the multi-coil $F^{-1}x$ was coil-combined using the sensitivity map and then denoised by NORDIC (complex denoising), with the kernel size for g-factor map estimation set to $20 \times 20 \times 5$, while other parameters remained at their defaults. The denoised complex image was then multiplied by the sensitivity map to generate multi-coil data for the subsequent iteration. The PF data were reconstructed using P-LORAKS [116].

To evaluate the performance of our reconstruction method on under-sampled data, fully sampled data from the 3T 0.65 mm protocol were retrospectively under-sampled by selecting 2 and 3 evenly spaced segments from the original 6 segments, corresponding to acceleration factors of $R_{\text{eff}}=3$ and 2, respectively. The normalised root mean squared error (NRMSE) between the under-sampled and fully-sampled data reconstructed using SPIRiT and DnSPIRiT was calculated within a brain mask to quantify their similarities.

The 1.22 mm data acquired at 3T were reconstructed using a standard SPIRiT with motion-induced phase error correction [80]. A joint SPIRiT reconstruction with distortion and phase error correction detailed in Chapter 3 [54] was used for reconstructing the 1.05 mm data acquired at 7T.

5.2.6 Diffusion Analyses

Image post-processing was conducted using the FMRIB Software Library (FSL) [64] unless indicated otherwise. Slab combination and correction for slab saturation artifacts were performed for the diffusion data using nonlinear inversion of slab profile encoding (NPEN) [62]. The images were corrected for Gibbs ringing [69]. A whole-brain field map was estimated using blip-reversed $b=0$ image volumes using “topup” [9]. For the 6-direction diffusion data from 3T 0.65 mm Protocol and 7T 0.61 mm Protocol, data from different diffusion directions were aligned using “eddy_correct”, and “applytopup” were used to address off-resonance distortions. For the 20-direction data from 3T 0.53 mm Protocol, the diffusion data and the field map estimated by “topup” were input to “eddy” [9] to correct for off-resonance distortions, eddy current effects, and subject motion. The 1.22 mm and 1.05 mm data were also processed by NPEN, Gibbs ringing correction, “topup”, and “eddy”, followed by “flirt” [71] for head position alignment with high-resolution data but without upsampling. The T1w image was corrected for bias field using “fast” [77] and co-registered to the diffusion space by applying the inverse transform obtained from “epi_reg” [71, 72]. Co-registered T1w image was then processed by FreeSurfer’s “recon-all” [117] to produce whole-brain segmentations. All diffusion analyses were performed in the native diffusion space.

Diffusion tensor model fitting was performed on all diffusion data using “dtifit” [63]. For the 20-direction 0.53 mm data and the 24-direction 1.22 mm data, fibre orientation distributions (FOD) were estimated voxel-wise using spherical deconvolution [118, 119]. A response function was first estimated with maximum harmonic degrees of 4 using MRtrix3’s “dwi2response” function and the “fa” algorithm [120], followed by FOD estimation via MRtrix3’s “dwi2fod” function with the “csd” algorithm [119].

Tractography was performed to reconstruct three representative white matter bundles from the 0.53 mm and 1.22 mm datasets to evaluate the data quality. Specifically, the acoustic radiation, cingulate gyrus part of cingulum, and corticospinal tract were reconstructed with the seed, waypoint, target, and exclusion

masks from “autoPtx” [74, 75]. Tractography was executed with MRtrix3’s “tckgen” function using the “iFOD2” algorithm [121], with the cutoff value set to 0.3. For each tract, a matched number of streamlines were generated from the 0.53 mm and 1.22 mm resolution data.

To assess the “gyral bias” problem, masks were manually drawn to delineate several representative gyri on 0.53 mm and 1.22 mm data, and tractography was performed within the gyri using these masks as seeds and region-of-interests (ROIs) using “tckgen” with “iFOD2”.

For U-fibre comparison between the 0.53 mm and 1.22 mm data, tractography was performed to reconstruct the short association fibres on both datasets. Figure 5.2 shows the seed, waypoint, and ROI masks for the tractography. Specifically, the seed mask was set to cortical grey matter, obtained by subtracting the white matter from the cortical ribbon mask segmented from the T1w image by FreeSurfer’s “recon-all” (Fig. 5.2a). The waypoint mask was the white matter mask (Fig. 5.2b). The ROI mask was the cortical ribbon mask dilated using MRtrix3’s “maskfilter” with “-npass” set to four (Fig. 5.2c), excluding non-relevant deep white matter tracts. The tractography was performed using “tckgen” with “iFOD2” and a minimum track length of 10 mm. On the 0.53 mm data, one seed was generated from each voxel and the cutoff value was set to 0.4. On the 1.22 mm data, two seeding mechanisms were evaluated – one seed per voxel and 12 seeds per voxel, with the latter compensating the voxel number difference due to resolution and demonstrating the impact of an increased seed number. A higher cutoff value of 0.7 was applied to reduce false positive tracts that primarily traversed in grey matter. The 0.53 mm data (one seed per voxel) and 1.22 mm data (12 seeds per voxel) produced comparable streamline numbers (1.75M and 1.84M, respectively), while the 1.22 mm data with one seed per voxel yielded 0.15M tracts.

5.3 Results

Simulation results for effective resolution and SNR across different in-plane segmentation numbers (N_{seg}) for 0.6 mm and 1 mm EPI-based dMRI acquisitions at 3T

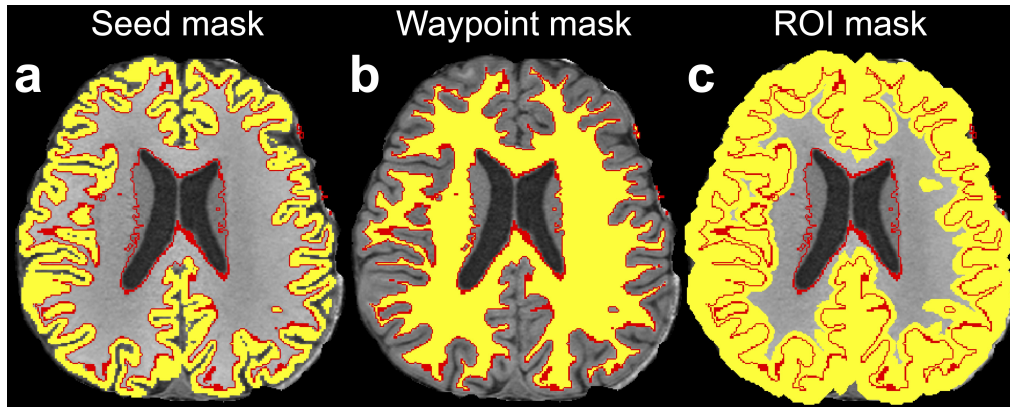


Figure 5.2: Masks for tracking the short association fibres. The seed (a), waypoint (b), and region-of-interest (ROI) (c) masks (yellow) for tracking the short association fibres are displayed on an axial slice of the T1w image, with the white matter boundary marked in red and overlaid.

are shown in Fig. 5.3. In all cases, T_2^* blurring reduces the effective resolution, but increasing N_{seg} helps mitigate this effect (Fig. 5.3a, c). PF acquisitions can improve SNR by shortening the TE, though they introduce additional blurring, particularly when zero-padding (ZP) is used. In practice, phase errors lead to imperfections in conjugate symmetry (CS) in PF acquisitions, making the simple PF-CS simulation model less feasible and further degrading effective resolution. These findings emphasise the importance of using large N_{seg} to achieve higher effective resolution, especially for submillimetre imaging. Additionally, SNR improvements are also observed with increasing N_{seg} due to the shorter TE (Fig. 5.3b, d). Notably, the SNR for 0.6 mm acquisitions is considerably lower than for 1 mm acquisitions, with over a fivefold difference (Fig. 5.3b, d), highlighting the substantial SNR challenges in high-resolution dMRI and the critical need for higher N_{seg} to compensate for SNR loss.

Simulations at 7T follow similar trends (Fig. 5.4). The shorter T_2 at 7T intensifies T_2^* blurring, requiring a higher N_{seg} to reduce the readout time and achieve a resolution comparable to 3T. Moreover, the shorter T_2 at 7T demands a larger N_{seg} to shorten the TE and take full advantage of the SNR benefit from the higher field strength. Based on these simulations, we selected $N_{\text{seg}}=6$ for the 3T protocols and $N_{\text{seg}}=8$ for the 7T protocol to balance scan time, effective resolution, and SNR. The simulations suggest that our three high-resolution protocols – 3T

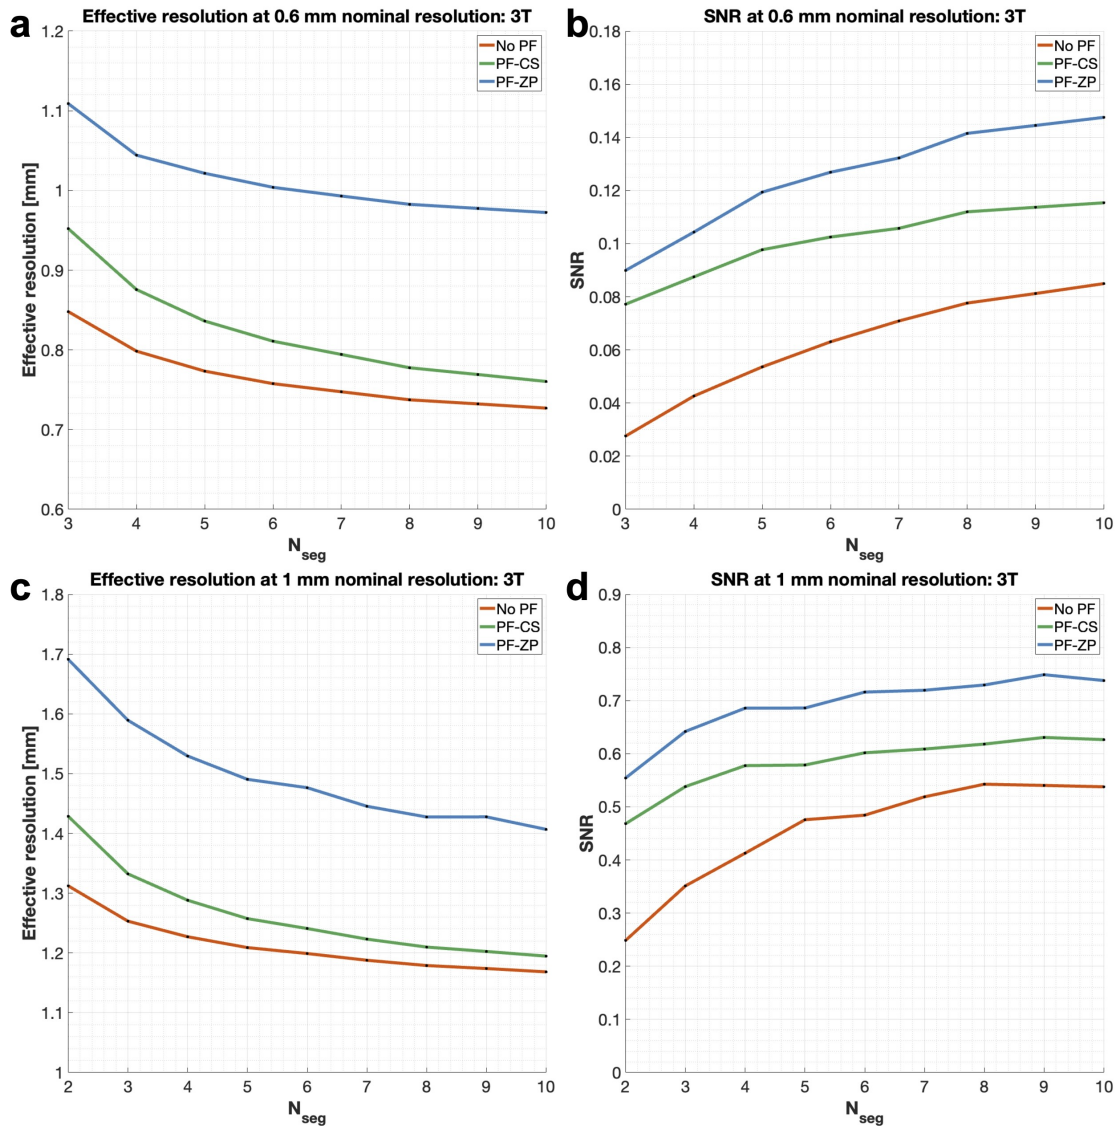


Figure 5.3: Simulation of effective resolution and SNR for high-resolution diffusion-weighted EPI at 3T. No partial Fourier (No PF, red), 6/8 PF with conjugate symmetric filling (PF-CS, green), and 6/8 PF with zero padding (PF-ZP, blue) samplings are investigated at 1 mm (a, b) and 0.6 mm (c, d) for white matter at 3T ($TR = 2.5\text{s}$, $b\text{-value} = 1000\text{ s/mm}^2$, bandwidths = 992/1384 Hz/pixel for 0.6/1 mm acquisitions, $T_1/T_2/T_2^* = 800/79.6/53.2\text{ ms}$) for different acceleration factors along phase encoding direction (N_{seg}). The SNR is calculated based on the simulated effective voxel sizes, TE, and TR.

0.65 mm ($N_{\text{seg}}=6$, no PF), 3T 0.53 mm ($N_{\text{seg}}=6$, with PF=3/4), and 7T 0.61 mm ($N_{\text{seg}}=8$, no PF) – achieve effective resolutions of approximately 0.8 mm.

Figure 5.5 demonstrates the image reconstruction results from 3T 0.65mm Protocol. Our carefully designed acquisition scheme produces ultrahigh-resolution diffusion images with superior sharpness and adequate SNR even when using the

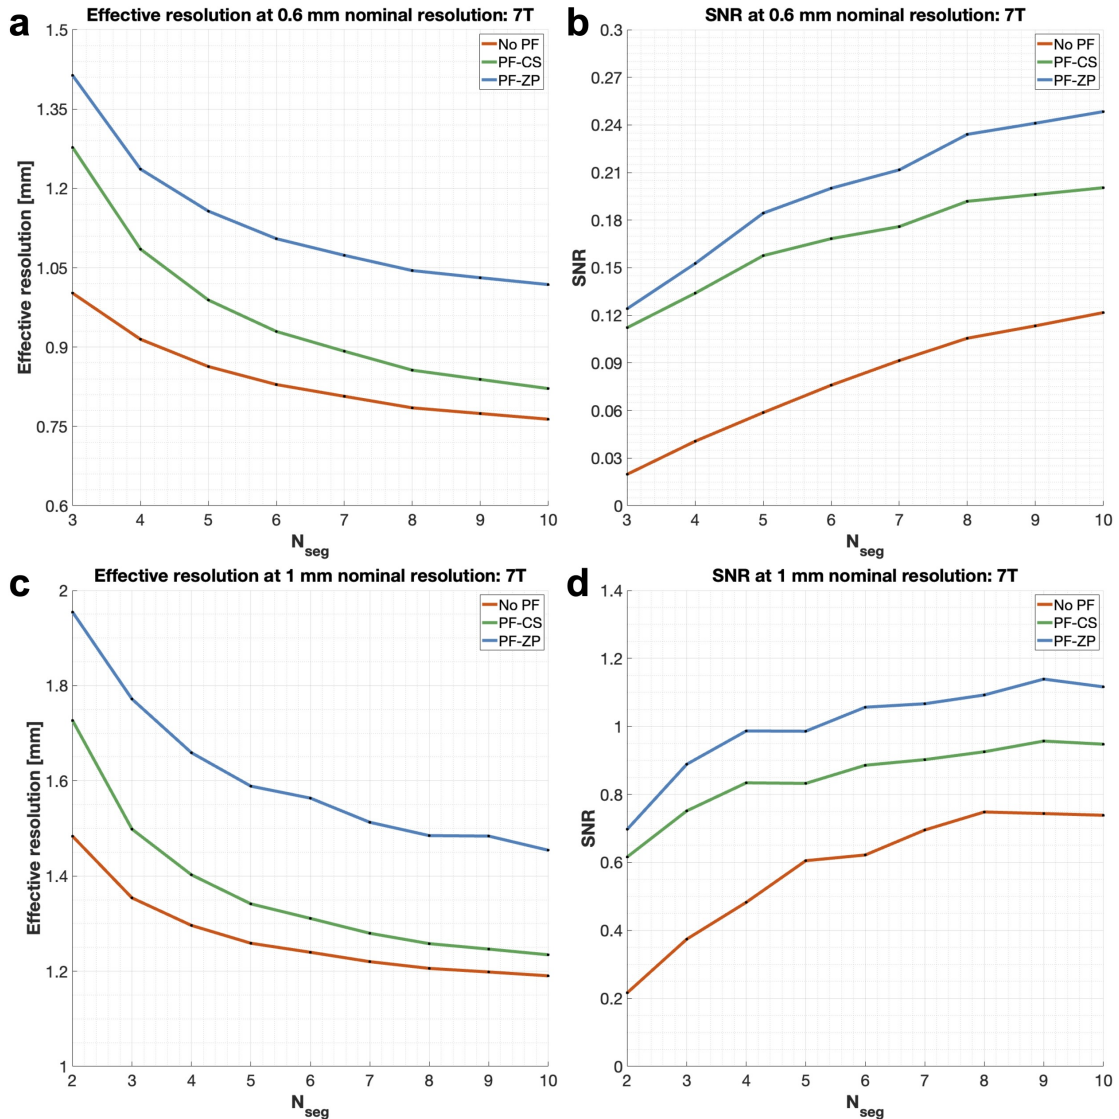


Figure 5.4: Simulation of effective resolution and SNR for high-resolution diffusion-weighted EPI at 7T. No partial Fourier (No PF, red), 6/8 PF with conjugate symmetric filling (PF-CS, green), and 6/8 PF with zero padding (PF-ZP, blue) samplings are investigated at 1 mm (a, b) and 0.6 mm (c, d) for white matter at 7T ($TR=3$ s, $b\text{-value}=1000$ s/mm², bandwidths=992/1384 Hz/pixel for 0.6/1 mm acquisitions, $T_1/T_2/T_2^* = 1200/47/26.8$ ms) for different acceleration factors along phase encoding direction (N_{seg}).

standard SPIRiT reconstruction (Fig. 5.5a, b). The proposed DnSPIRiT further enhances SNR with negligible biases or blurring (Fig. 5.5c, d). The difference maps between SPIRiT and DnSPIRiT primarily display noise without noticeable anatomical structures or biases (Fig. 5.5e).

The benefit of DnSPIRiT becomes more pronounced in under-sampled recon-




Figure 5.5: Denoiser-regularised reconstruction results. *In-vivo* diffusion-weighted data ($b = 1000 \text{ s/mm}^2$) from 3T 0.65 mm Protocol along 3 representative diffusion encoding directions (i-iii, the direction is displayed above each image) are reconstructed using SPIRiT (a) and Denoiser-regularised SPIRiT (DnSPIRiT, c), with an enlarged region showing the image detail (b, d). Their difference is also shown to demonstrate the removed noise (e).




Figure 5.6: Retrospective under-sampled reconstruction. Retrospective under-sampling are applied to the fully sampled data (i) from 3T 0.65 mm Protocol by selecting 3 segments ($R_{\text{eff}} = 2$, ii) and 2 segments ($R_{\text{eff}} = 3$, iii) from the total of 6 segments. The diffusion-weighted images along $(-0.45, -0.86, -0.23)$ reconstructed using SPIRiT (a) and DnSPIRiT (c) and their difference with fully sampled reference (b, d) are shown to demonstrate the under-sampled reconstruction fidelity. The normalised root mean squared errors (NRMSE) between the under-sampled reconstructed image and the fully sampled reference of the whole slab are calculated within a brain mask to quantify their similarity.

structions (Fig. 5.6). While the standard SPIRiT reconstruction exhibits significant noise for under-sampled data (Fig. 5.6a, b), DnSPIRiT effectively suppresses the noise (Fig. 5.6c, d), resulting in substantially lower NRMSE with the fully sampled reference. However, when only 2 segments are used ($R_{\text{eff}} = 3$, Fig. 5.6, iii), anatomical errors begin to emerge even with DnSPIRiT, presumably due to the extremely high noise level introduced by $\sqrt{R_{\text{eff}}}$ and g-factor penalty. We therefore used $R_{\text{eff}} = 2$ in 3T 0.53 mm Protocol to balance the acquisition time and reconstruction accuracy. Notably, the prospective under-sampled acquisition is expected to be more robust against motion artifacts compared to retrospective under-sampling due to shortened scan time.

The whole-brain diffusion data from 3T 0.65 mm Protocol resolve rich information in brain microstructure (Fig. 5.7). The diffusion data exhibit sharp details and high-SNR even with only 6 diffusion encoding directions, thanks to our carefully optimised sampling and reconstruction strategies. The short effective echo spacing (0.21 ms), achieved through in-plane segmentation, significantly enhances the anatomical fidelity (Fig. 4b). For comparison, the effective echo spacing for HCP 3T diffusion imaging at 1.25 mm isotropic resolution was 0.78 ms [122]. Small structures such as fornix, external capsule, and anterior commissure are clearly visible on the 0.65 mm data with excellent anatomical fidelity (Fig. 5.7d). Moreover, the results align well with prior *post-mortem* data at 0.5 mm resolution [95] (Fig. 5.7d).

Compared to the conventional 1.22 mm resolution diffusion data, data from 3T 0.65 mm Protocol reveals substantially finer microstructure details (Fig. 5.8). This includes the improved visualisation of the striations through the internal capsule (Fig. 5.8, i), the interdigitating transverse pontine fibres (Fig. 5.8, ii), and the clearer delineation of the cingulum bundle (Fig. 5.8, iii).

Figure 5.9 shows the whole-brain diffusion data from 3T 0.53 mm Protocol. Across both subjects, our proposed acquisition and reconstruction framework consistently achieves high image quality with adequate SNR (Fig. 5.9a, c) and superior anatomical fidelity (Fig. 5.9b) at 0.53 mm isotropic resolution.

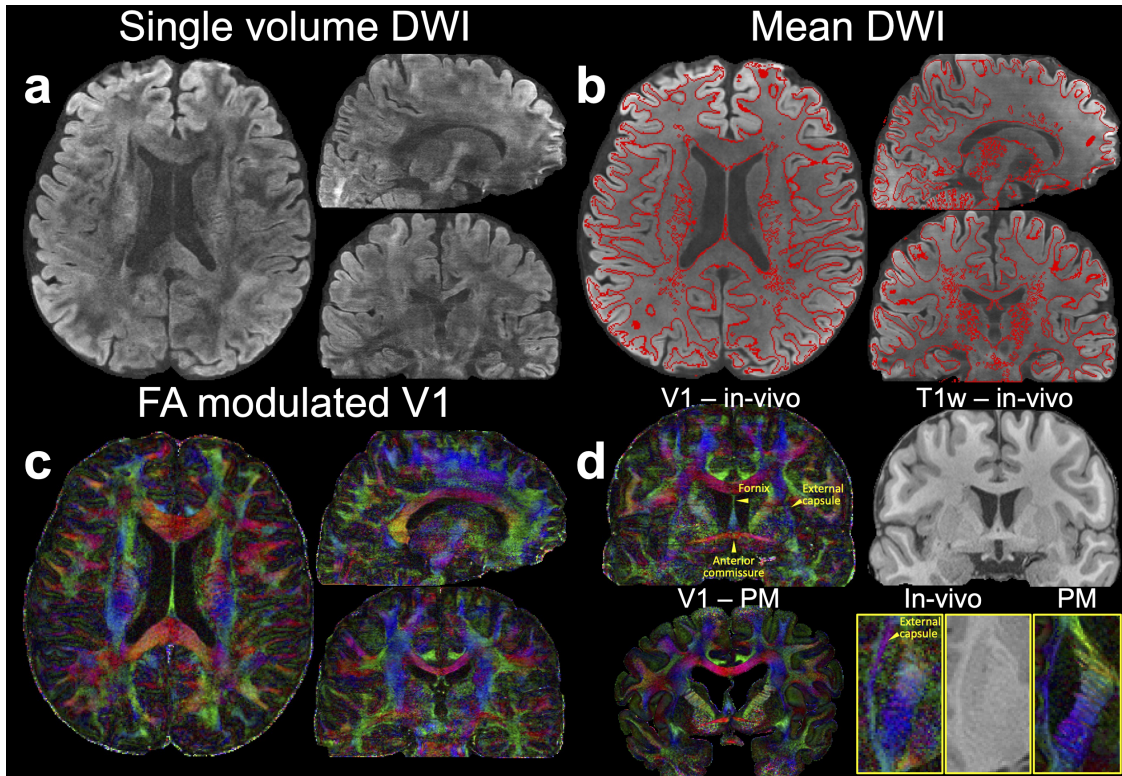


Figure 5.7: Whole-brain data from 3T 0.65 mm Protocol. Whole-brain *in-vivo* diffusion data ($b=1000$ s/mm²) including diffusion-weighted image (DWI) along $(-0.45, 0.83, -0.32)$ (a), 6-direction mean DWI with grey-white matter boundary derived from the T1w image using FSL’s “fast” overlaid (b), FA modulated V1 (c) at 0.65 mm isotropic resolution are shown. A representative sagittal view of the FA modulated V1 (V1 – *in-vivo*) alongside the T1-weighted image (T1w – *in-vivo*), highlighting key structures including the fornix, external capsule, and anterior commissure are also presented. For comparison, *post-mortem* data (V1 – PM) from previous studies [95] are shown. Enlarged views of the external capsule from an axial plane are provided for both *in-vivo* and *post-mortem* data (d).

The superior quality of the 0.53 mm data is further demonstrated by the successful reconstruction of three representative white matter tracts (Fig. 5.10). Specifically, our data enable the tracking of the acoustic radiation, which crosses several major brain fibre systems with complex fibre crossings, making it difficult to reconstruct from *in-vivo* data [123]. The tracking results for the acoustic radiation on the 0.53 mm data are consistent with those from the 1.22 mm data (Fig. 5.10, i). The cingulate gyrus portion of the cingulum, which predominantly follows an anterior-posterior trajectory, is also successfully reconstructed in both datasets. The tract appears more extensive in the retrosplenial region in the 0.53 mm data, likely

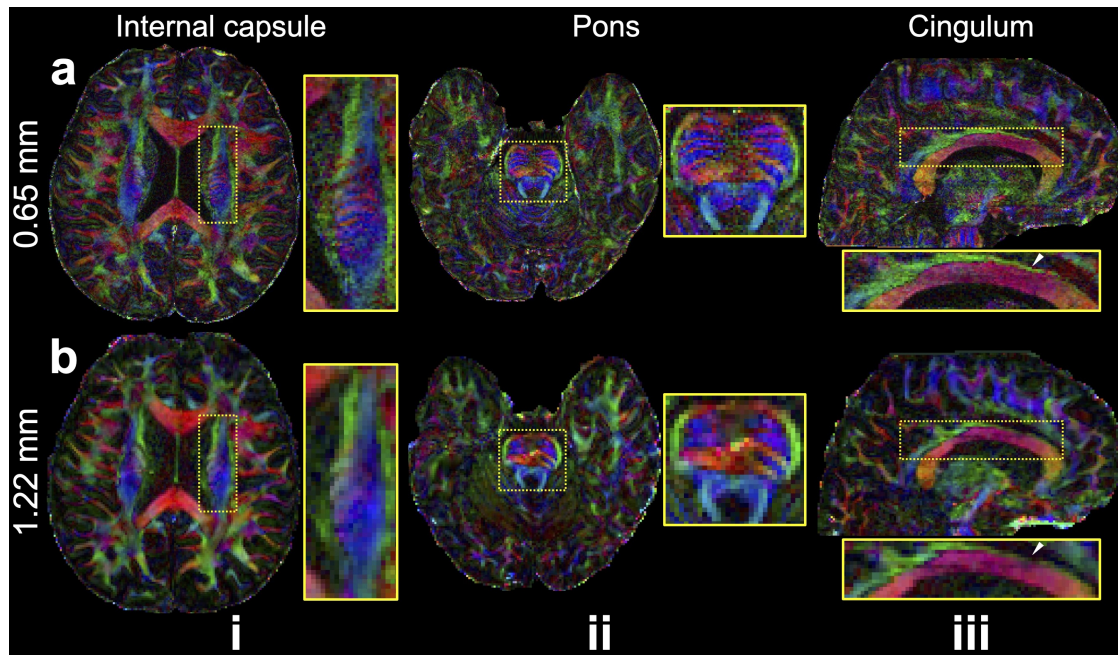


Figure 5.8: Comparison of 3T 0.65 mm and 1.22 mm isotropic resolution DTI. Two axial slices showing the internal capsule (i) and pons (ii), and a sagittal slice showing the cingulum bundle (iii) of FA modulated V1 of 0.65 mm (a) and 1.22 mm (b) isotropic resolution *in-vivo* diffusion data acquired at 3T are demonstrated. White arrows highlight the improved delineation of the cingulum bundle in the 0.65 mm data.

due to the higher resolution (Fig. 5.10, ii, yellow arrows). For the corticospinal tract, which primarily runs along the slice selection direction, the 0.53 mm data reveal clearer branching (Fig. 5.10, iii, yellow arrows at the top) and a thicker appearance near the pons (Fig. 5.10, iii, yellow arrows near the bottom). The clearer branching is likely due to the improved delineation of crossing fibres thanks to the higher resolution. The thicker tract near the pons is attributed to reduced distortion and better anatomical fidelity in the 0.53 mm data compared to the 1.22 mm data. Notably, the tractography utilised masks from FSL’s “autoPtx”, which is designed for conventional resolution fibre tracking (1-2 mm). It is expected that the advantage of submillimetre resolution data would be even more pronounced with tailored tractography methods designed for such spatial resolution.

The 0.53 mm data also effectively reduce the gyral bias problem, particularly in small gyri (Fig. 5.11). In the 1.22 mm data, gyral bias is evident, with streamlines often terminating at the gyral crowns (Fig. 5.11, ii, iv). In contrast, the 0.53 mm




Figure 5.9: Whole-brain data from 3T 0.53 mm Protocol. Whole-brain *in-vivo* diffusion data of two subjects ($b=1000 \text{ s/mm}^2$) including diffusion-weighted image (DWI) along $(0.5, -0.86, 0.02)$ (a), 20-direction mean DWI with grey-white matter boundary derived from the T1w image using FSL’s “fast” overlaid (b), FA modulated V1 (c) at 0.53 mm isotropic resolution are shown.

data show the expected fanning pattern, with a greater number of streamlines extending to the gyral walls (Fig. 5.11, i, iii). This more even distribution of streamline termination points between gyral crowns and walls in the 0.53 mm data indicates a significant reduction in gyral bias with higher spatial resolution, consistent with previous *post-mortem* findings [95]. Notably, we found the gyral bias problem in our data is less pronounced in larger gyri, where the expected fanning pattern is also visible in the 1.22 mm data 5.12.

Compared to the 1.22 mm data, the 0.53 mm data also provide an improved

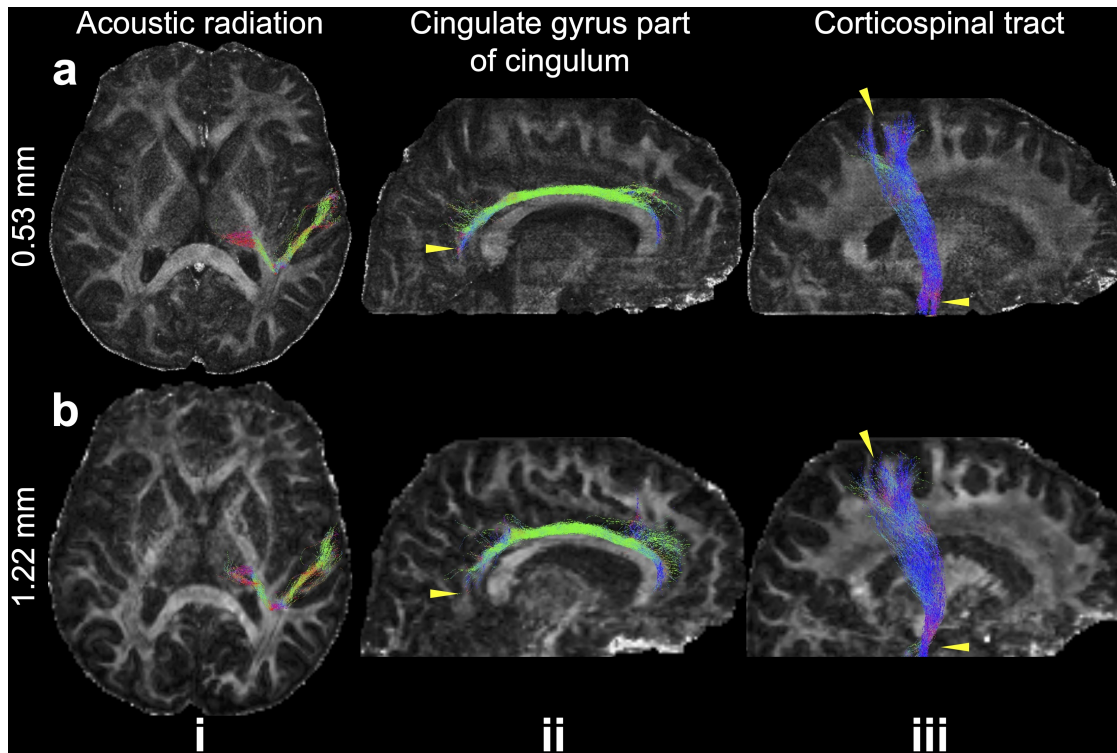


Figure 5.10: Example white matter tracts from the 0.53 mm and 1.22 mm data. The maximum intensity projections of three representative white matter tracts including acoustic radiation (i), cingulate gyrus part of cingulum (ii), and corticospinal tract (iii) are overlaid on the fractional anisotropy maps of 0.53 mm (a) and 1.22 mm (b) datasets. The yellow arrows indicate the region where the tractography on 0.53 mm data shows improvement compared to that on 1.22 mm data.

mapping of U-fibres (Fig. 5.13). Our tractography setup successfully identified short association fibres in both datasets (Fig. 5.13a, b). When using one seed per voxel for both, the 0.53 mm data reveal denser whole-brain short association fibres, likely due to the substantially increased voxel number. The distribution of short association fibres in the 0.53 mm data is also more continuous along the grey-white matter boundary. These findings are consistent with previous studies [98]. Remarkably, the 0.53 mm data is able to resolve U-fibres even at sharp turnings along the grey-white matter boundary (Fig. 5.13c-e, i, ii), whereas the 1.22 mm data struggle to capture these fibres (Fig. 5.13c-e, iii, iv), even with increased seed numbers to compensate the voxel number difference (Fig. 5.13c-e, v). The tractography on 0.53 mm data with one seed per voxel and that on 1.22 mm data with 12 seeds per voxel produced comparable numbers of streamlines (1.75M vs.

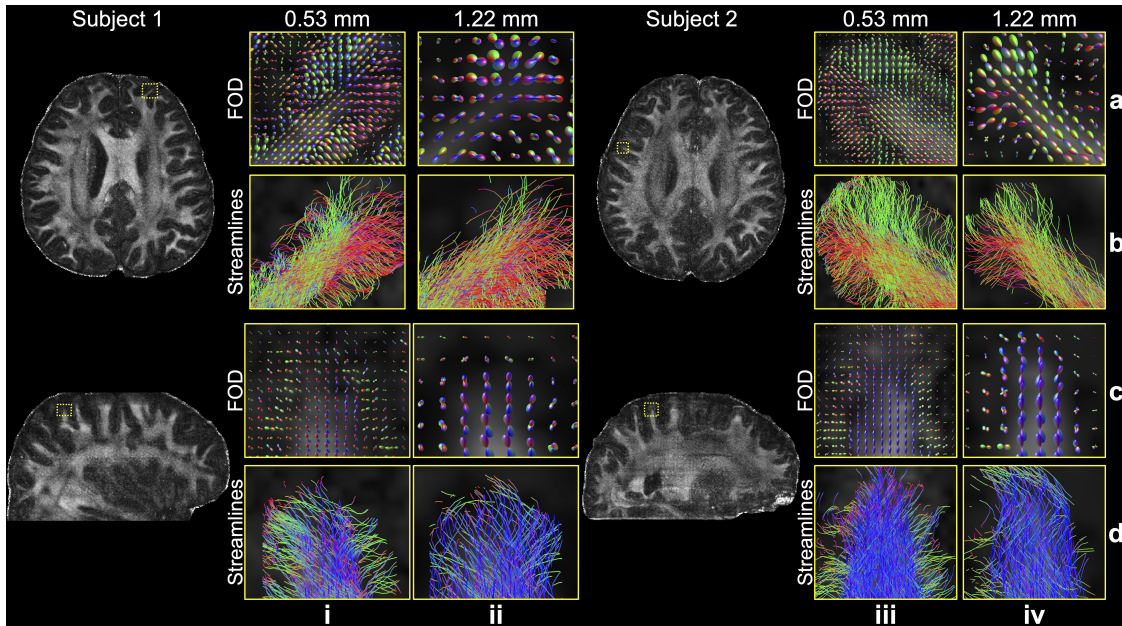


Figure 5.11: Gyral bias comparison between 0.53 mm and 1.22 mm data. The fibre orientation distributions (FOD) (a, c) and tractography streamlines (b, d) for representative gyri from the 0.53 mm (i, iii) and 1.22 mm (ii, iv) data of two subjects are shown to demonstrate the reduced gyral bias on high-resolution data. The fractional anisotropy maps at 0.53 mm resolution are also displayed to indicate the anatomical position of these gyri.

1.84M). This confirms the improvement on the 0.53 mm data originates from the better resolved fibre turning curvature due to the higher spatial resolution, rather than simply an increase in seed or streamline numbers.

Our acquisition and reconstruction framework shows robustness at higher field strength of 7T (Fig. 5.14). The diffusion data from the 7T 0.61 mm Protocol exhibit excellent SNR even with only six diffusion encoding directions, thanks to the SNR benefit at 7T (Fig. 5.4). Compared to the 1.05 mm isotropic resolution data, the submillimetre dataset reveals significantly more detailed microstructure, including better resolution of fanning patterns in the gyri and clearer delineation of small structures (Fig. 5.14, white arrows). Notably, our *in-vivo* results for resolving transverse pontine fibres are in excellent agreement with previous *post-mortem* data acquired with much longer scan times on the same 7T scanner [95].




Figure 5.12: Fibre distributions at a large gyrus. The fibre orientation distributions (FOD) (a) and tractography streamlines (b) for a large gyrus from the 0.53 mm (i) and 1.22 mm (ii) data are shown. At this large gyrus, the gyral bias problem is less pronounced even on 1.22 mm data. The fractional anisotropy map at 0.53 mm resolution is also displayed to indicate the anatomical position of the gyrus.

5.4 Discussion

This chapter introduces an acquisition and reconstruction framework to achieve high-quality submillimetre dMRI. Our acquisition takes advantage of the optimal SNR efficiency provided by 3D multi-slab EPI, with in-plane segmentations integrated to shorten the effective echo spacing, readout, and TE for reduced distortion, T_2^* blurring, and SNR penalty. The sampling order is designed for improved robustness against subject motion. A denoiser-regularised reconstruction approach is proposed to suppress noise while maintaining consistency with the acquired data to reduce bias or blurring. Our *in-vivo* experiments at 3T demonstrate that our method consistently produces high-quality diffusion data at 0.65 mm and 0.53 mm isotropic resolutions. The analyses of these submillimetre datasets reveal enhanced microstructural detail, mitigating the “gyral bias” problem and improving the precision of U-fibre tracking compared to data acquired at conventional resolutions of 1.22 mm from the same subjects at the same scanner. The framework also proves effective at 7T, delivering similarly robust results. These findings suggest that the proposed framework represents a promising advance in high-resolution dMRI,

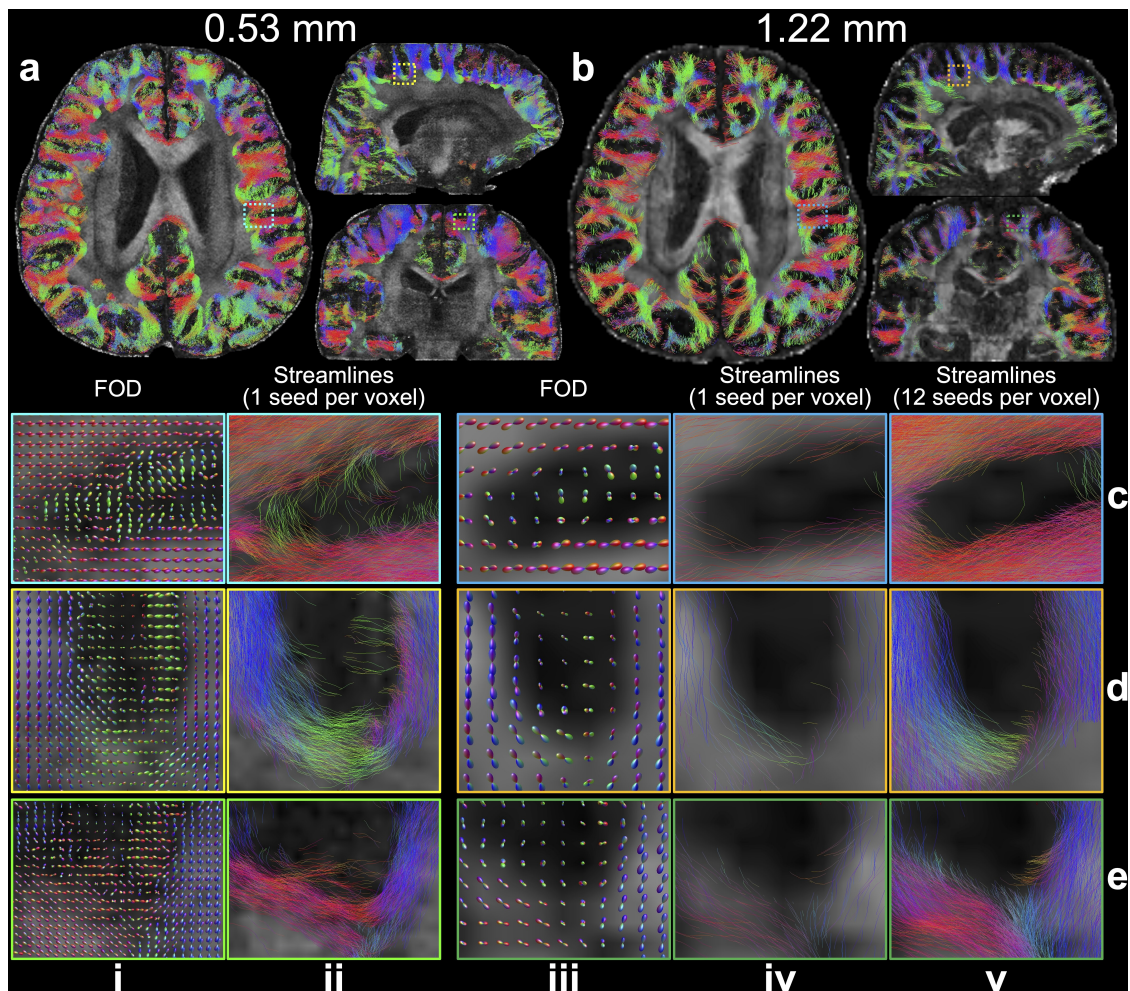


Figure 5.13: U-fibres comparison between 0.53 mm and 1.22 mm data. The whole-brain short association fibres tracked with one seed per voxel for 0.53 mm (a) and 1.22 mm (b) data are overlaid on fractional anisotropy maps. Representative enlarged regions (c-e) show the fibre orientation distributions (FOD) (i, iii), the tracked streamlines with one seed per voxel (ii, iv) for the 0.53 mm (i, ii) and 1.22 mm (iii, iv) data, and the tracked streamlines with 12 seeds per voxel for the 1.22 mm data (v) to compensate the difference in voxel numbers due to resolution.

offering the potential for more accurate characterisation of brain microstructure and expanded applications in neuroscientific research.

The optimisation of our sampling strategies considers the effective resolution, SNR, distortion, scan time, and motion robustness. Compared to previous studies that achieve submillimetre resolutions using super-resolution methods [47, 50, 85], our method using 3D multi-slab EPI offers several key advantages: (i) superior SNR efficiency enabled by short TR and orthogonal encoding bases from Fourier encoding, and (ii) high-fidelity voxel shapes produced directly by Fourier encoding,




Figure 5.14: Comparisons of 7T 0.61 mm and 1.05 mm diffusion data. The sagittal, coronal, and axial slices of fractional anisotropy (FA) modulated V1 and their enlarged regions of 0.61 mm and 1.05 mm isotropic resolution *in-vivo* diffusion data acquired at 7T are demonstrated, with white arrows highlighting the more detailed microstructure resolved by the high resolution at 0.61 mm. The transverse pontine fibres demonstrate similar patterns compared to previous *post-mortem* study [95] at 0.5 mm isotropic resolution from the same scanner.

avoiding the need for Tikhonov regularisation used in super-resolution reconstruction. However, for EPI-based sampling, higher resolution imaging requires longer echo spacing, readout, and TE, and therefore more severe distortion, T_2^* blurring, and T_2 signal decay. For example, for $N_{\text{seg}}=3$ sampling without PF at 3T, image displacement (assuming a 50 Hz B_0 offset), image blurring ((effective resolution-nominal resolution)/nominal resolution) and T_2 signal decay (e^{-TE/T_2}) for 0.6 mm imaging are 4.4 mm, 41.3%, and 0.103, respectively, compared to 3.2 mm, 25.3%, and 0.293 for 1 mm imaging (Fig. 5.3). This requires a larger N_{seg} for submillimetre imaging to improve effective resolution, SNR, and minimise distortion, though at the cost of longer scan times. In the 3T 0.65 mm Protocol, we acquired fully-sampled data with $N_{\text{seg}}=6$ without using PF, which achieved superior SNR, image sharpness, and anatomical fidelity (Fig. 5.7). To enable more advanced diffusion analyses with

more diffusion encoding directions, we explored the potential for acceleration by acquiring fewer segments (Fig. 5.6) and achieved $R_{\text{eff}}=2$ in the 3T 0.53 mm Protocol without significantly compromising the image quality, allowing for the acquisition of 20 diffusion encoding directions in a one-hour session. Notably, the TR of 2 s in this protocol is within the SNR-efficient range. Additionally, we applied a motion-robust sampling order for 3D EPI (Fig. 5.1), which particularly improves the $b=0$ images where motion artifacts are more noticeable due to strong CSF signal aliasing.

Our denoiser-regularised reconstruction DnSPIRiT effectively suppresses noise while maintaining image fidelity by enforcing consistency with the acquired data. Essentially, the denoised image serves as a regularisation to suppress noise during reconstruction, instead of directly being used as the final output. The experiment shows our method effectively reduces the noise without introducing significant blurring or bias (Fig. 5.5), supported by the fact that the removed noise generally exhibits random distribution without noticeable anatomical structures (Fig. 5.5e). DnSPIRiT proves even more beneficial in under-sampled reconstruction where the noise level is higher (Fig. 5.6). The plug-and-play nature of DnSPIRiT allows for the integration of advanced denoisers without adding significant computational burden in directly solving Eq. 5.2, which is crucial for high-resolution dMRI reconstruction due to the large data size. We demonstrated its compatibility with BM4D for 6-direction data denoising (3T 0.65 mm Protocol, 7T 0.61 mm Protocol) and NORDIC for 20-direction data denoising (3T 0.53 mm Protocol). Incorporating recently developed deep learning-based dMRI denoisers [124–126] that outperform conventional methods is expected to further improve the reconstruction.

The comprehensive *in-vivo* experiments comparing submillimetre data with prospectively acquired conventional-resolution data (1.22 mm and 1.05 mm) validate the benefits of higher spatial resolution for diffusion analyses. In previous studies [6, 98, 127], the comparisons between the high-resolution and low-resolution data were based on retrospective under-sampling. One limitation of those comparisons is that the increased image blurring and SNR penalty associated with high-resolution imaging (Figs. 5.3, 5.4) can propagate into the under-sampled low-resolution data,

leading to lower image quality than what can be practically achieved. Instead, we prospectively acquired data at 1.22 mm and 1.05 mm with practical imaging protocols from the same scanners as the submillimetre resolution data. Using such approach, the datasets at conventional resolutions also demonstrate excellent data quality. For instance, the expected fibre fanning pattern can be resolved at large gyri from the 1.22 mm data (Fig. 5.12). These prospective acquisitions provide a more realistic comparison, revealing that submillimetre data still significantly outperform 1.22 mm data in reducing gyral bias at small gyri and mapping U-fibres in sharp grey-white matter boundary turnings. Moreover, our results' strong agreement with previous *post-mortem* findings [95] on the same 7T scanner (Fig. 5.14) further reinforces the reliability of these comparisons.

Our implementation of the sequence using Pulseseq enhances the accessibility of our method. The scanner-agnostic, open-source nature of Pulseseq facilitates the straightforward translation of our framework to different scanner platforms. Demonstrating the ability to produce high-quality data on a clinical scanner (Siemens 3T Prisma) highlights the potential for broader application of our methods across various clinical and research settings. We further demonstrated the method's compatibility a higher field strength at 7T (Siemens 7T Magnetom), leveraging the SNR advantage (Fig. 10). Going further, the state-of-the-art high-performance scanners with stronger gradients, such as the Siemens 3T Connectome 2.0 [128], GE 3T Ultra-High Performance, and Siemens 7T NexGen [129], offer promising opportunities to further reduce readout times and TE, improving image quality and potentially enabling even higher spatial resolutions.

Robust acquisition of submillimetre dMRI data opens up new possibilities for medical image analysis and neuroscientific research. Our study demonstrates its potential to address "gyral bias" and improve U-fibre mapping, leading to more accurate reconstruction of intra- and inter-gyri connections. Future work could leverage the rich detail of ultrahigh-resolution data through more quantitative analyses. One potential direction is integrating cortical curvature from T1w images to enable precise whole-brain gyri segmentation to improve whole-brain

U-fibre mapping at submillimetre resolution. Submillimetre dMRI may also aid in mapping small but crucial subcortical structures targeted for deep brain stimulation [100, 101]. Public subcortical atlases (e.g., <https://www.lead-dbs.org/helpsupport/knowledge-base/atlassesresources/atlasses-2/>) offer useful resources for defining seeding and targeting masks for tractography to map white matter circuits in these regions.

One limitation of our method is the relatively long scan and reconstruction times. The scan time per volume (2.7 minutes in the 3T 0.53 mm Protocol) restricts the number of diffusion encoding directions. Further acceleration could be achieved by leveraging k-q space joint reconstruction, which shares information across diffusion directions [130]. Additionally, recent advances in eliminating navigator acquisition in 3D multi-slab EPI [80] could be incorporated into our framework to further enhance scan efficiency. On the reconstruction side, processing submillimetre data remains time-intensive due to the large matrix size (~12 hours per slab on a single CPU for data from the 3T 0.53 mm Protocol). Model-based deep learning reconstruction methods [131] present a promising avenue to significantly reduce computational time and potentially further improve SNR.

5.5 Conclusion

In this chapter, a novel acquisition and reconstruction framework is proposed to achieve high-quality submillimetre dMRI 0.5-0.6 mm isotropic resolution for *in-vivo* human brains. Comprehensive *in-vivo* experiments demonstrate the superior SNR, image sharpness, and anatomical fidelity of the submillimetre data, which reveal substantially more microstructure details, reduce the gyral bias, and improve the U-fibre mapping compared to data at conventional resolutions (1-1.2 mm). The method also shows robustness at 7T, where data show excellent agreement with previous *post-mortem* data at 0.5 mm resolution acquired from the same scanner. With the scanner-agnostic, open-source implementation using Pulseseq, our approach holds promise in benefiting a wide range of medical image analyses and neuroscientific research.

6

Summary and Future Directions

Contents

6.1	Thesis Summary	139
6.2	Future Directions	141
6.3	Final Remarks	149

6.1 Thesis Summary

Improving the SNR and resolution has long been a critical challenge in dMRI acquisition and reconstruction. 3D multi-slab acquisitions offer a promising solution, as they combine the benefits of 3D Fourier encoding with compatibility for optimal TR in spin echo-based dMRI, resulting in high-resolution, SNR-efficient imaging. This thesis focuses on developing acquisition and reconstruction techniques to improve the image fidelity and time efficiency of 3D multi-slab dMRI.

Conventional 3D multi-slab EPI-based dMRI is prone to image distortion and slab boundary aliasing, which degrade anatomical accuracy and overall image quality. In Chapter 3, we integrate advanced sampling strategies, including blip-reversed acquisition, CAIPI sampling, and k_z partial Fourier, with a joint reconstruction of blip-reversed data. This method enables the correction of distortion and slab

boundary aliasing without additional time cost. *In-vivo* experiments reveal that our method facilitates more accurate diffusion quantitative measurement and tractography compared to conventional 3D multi-slab dMRI with matched scan times.

Another key challenge of 3D multi-slab dMRI is its sensitivity to motion due to its multi-shot nature, which typically requires navigator acquisition after each imaging echo to address motion-induced shot-to-shot phase inconsistencies. However, acquiring navigators increases the SAR and prolongs the TR, compromising the SNR efficiency especially for high-field and high-resolution imaging. We introduce a self-navigated acquisition and reconstruction framework in Chapter 4 that enables robust extraction of motion-induced phase variations directly from the imaging data. Our method therefore eliminates the requirement for navigator acquisition in 3D multi-slab dMRI, which leads to 15.5% SNR efficiency increase and 31.2% scan time reduction without sacrificing image quality.

Chapter 5 investigates the potential of 3D multi-slab acquisitions to achieve submillimetre resolution in dMRI, a capability highly sought after for numerous neuroscientific applications. Achieving such high resolution in *in-vivo* settings is particularly challenging due to the intrinsically limited SNR and T_2^* decay-induced blurring. We leverage an in-plane segmented 3D multi-slab EPI acquisition, combined with a denoiser-regularised reconstruction to overcome these difficulties. Numerical simulations and *in-vivo* experiments at 3T and 7T support our method's ability to achieve high effective resolution, SNR, and anatomical fidelity. This facilitates high-quality submillimetre dMRI (0.5-0.6 mm isotropic resolution) which provides substantially richer microstructure details and more accurate mapping of intra- and inter-gyri fibre connections compared to conventional 3D multi-slab dMRI at 1-1.2 mm resolutions. Furthermore, our sequence is implemented using the open-source, scanner-agnostic platform "Pulseq", which may facilitate a broader adoption of this technique.

In conclusion, this thesis introduces novel acquisition and reconstruction methods for 3D multi-slab dMRI, offering significant improvements in data quality and time efficiency. These advances hold the promise to expand the use of dMRI

and push the boundaries of its SNR efficiency and resolution, thereby benefiting neuroscience explorations.

6.2 Future Directions

Self-Navigated, Distortion-Corrected 3D Multi-Slab Diffusion MRI

Building on the advances presented in Chapters 3 and 4, a 3D multi-slab dMRI framework can be developed to simultaneously correct both B_0 -induced geometric distortions and motion-induced phase variations without requiring additional scan time. This technique offers the potential to enable SNR-efficient dMRI with enhanced anatomical fidelity, particularly benefiting high-field imaging (e.g., 7T), where strong B_0 field inhomogeneities and SAR constraints are common challenges. As demonstrated in Chapter 3, blip-reversed sampling enables high-quality distortion correction even in the presence of significant B_0 field inhomogeneities. Additionally, the elimination of navigator echoes through self-navigation, as shown in Chapter 4, effectively reduces SAR and shortens the minimum allowable TR to deliver the optimal SNR efficiency.

The development of this self-navigated, blip-reversed 3D multi-slab dMRI framework involves several key considerations for acquisition and reconstruction. First, an integrated sampling optimisation is essential to balance distortion correction, self-navigation, and image reconstruction performance. In Chapter 3, the optimisation focuses on robust reconstruction of blip-up and blip-down images from highly under-sampled data, which is critical for accurate field map estimation and distortion correction. It also aims to achieve high-quality joint reconstruction that combines all blip-up and blip-down data and the estimated field map. In Chapter 4, we employ a greedy search-based sampling optimisation to ensure reliable self-navigation while maintaining satisfying multi-shot reconstruction performance. All these factors need to be comprehensively considered to implement the self-navigated, blip-reversed framework effectively.

On the reconstruction side, the structured low-rank (SLR) method for estimating motion-induced phase maps must be adapted for blip-reversed sampling. In Chapter 4, phase maps are reconstructed by solving:

$$\arg \min_{\hat{x}} \|\mathcal{M}_{k_z0}\hat{x} - y_{k_z0}\|_2^2 + \lambda_1 \|(G_{k_z0} - I)\hat{x}\|_2^2 + \lambda_2 \|H(\hat{x})\|_*, \quad (6.1)$$

where \mathcal{M}_{k_z0} selects the self-navigation points, \hat{x} is the fully-sampled multi-coil k-space data for all shots at $k_z = 0$, y_{k_z0} is the acquired self-navigation points, G_{k_z0} is the SPIRiT kernel trained on the $k_z = 0$ calibration data, and I is the identity matrix. This method assumes consistent magnitude across all shots, despite phase differences. However, blip-reversed sampling violates this assumption due to distortion-induced magnitude inconsistencies between blip-up and blip-down images. If phase maps for blip-up and blip-down shots are reconstructed separately, shared information across shots is reduced, potentially degrading performance.

An alternative strategy is to incorporate a distortion operator into the forward model:

$$\arg \min_{\tilde{x}} \|\mathcal{M}_{k_z0} F E F^{-1} \tilde{x} - y_{k_z0}\|_2^2 + \lambda_1 \|(G_{k_z0} - I)\tilde{x}\|_2^2 + \lambda_2 \|H(\tilde{x})\|_*, \quad (6.2)$$

where \tilde{x} is the distortion-free, fully-sampled k-space data for all shots at $k_z = 0$, F is the Fourier transform, and E is the spatial distortion operator, acting inversely for blip-up and blip-down shots. This ensures magnitude consistency of \tilde{x} across all blip-up and blip-down shots, enabling accurate reconstruction of the motion-induced phase maps. In practice, the slab-averaged field map from the $b=0$ images can be used to generate E . With the motion-induced phase maps, a distortion-corrected image can then be obtained using the two-stage reconstruction framework from Chapter 3.

We conducted a simulation study to validate the feasibility of this method using the fully-sampled blip-up/down data acquired in the Evaluation Protocol in Chapter 3. For the sampling, we applied the 1.22 mm self-navigated scheme established in Chapter 4 (Fig. 4.2e, 12-shot sampling) and selected the six shots with the

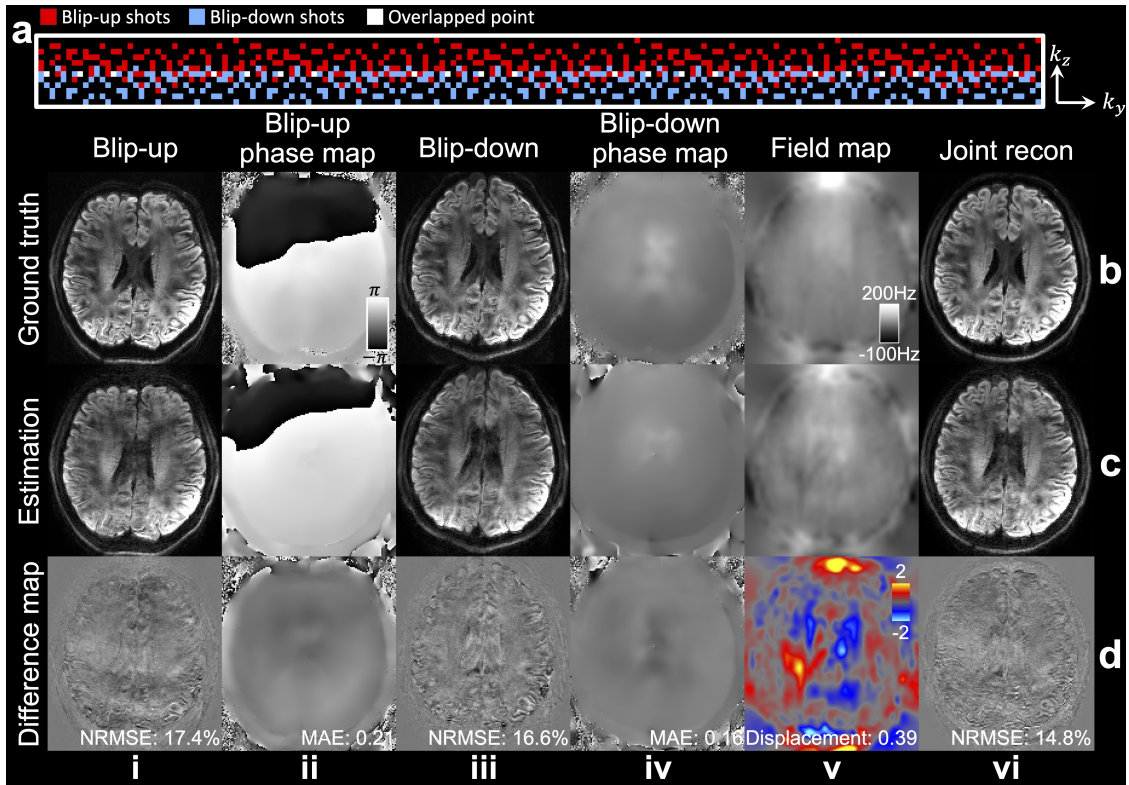


Figure 6.1: Simulation study for self-navigated, distortion-corrected 3D multi-slab dMRI. The sampling for self-navigated, distortion-corrected imaging at 1.22 mm resolution (matrix size: 180×12) is shown in (a). The blip-up/down image and motion-induced phase maps (i-iv), the field map estimated from the blip-reversed data (v), and the joint reconstruction results (vi) from the fully-sampled ground truth data at 1.22 mm isotropic resolution (7T, $b = 1000 \text{ s/mm}^2$) acquired from the Evaluation Protocol in Chapter 3 (b), the retrospective under-sampled reconstructions (c), and their differences (d) are demonstrated. The differences in reconstructed images, motion-induced phase maps, and field map are quantified using normalised root mean squared error (NRMSE), mean absolute error (MAE), and mean displacement (in pixels), respectively, for the whole-slab data within a brain mask.

largest k_z shifts (i.e., sampling the upper part of the k-space) as the blip-up shots and the remaining as the blip-down shots (Fig. 6.1a). This design mimics the complementary partial Fourier distribution of blip-up/down shots used in Chapter 3. Motion-induced phase maps for the blip-up/down shots were reconstructed using Eq. 6.2, and these phase maps were incorporated into the Chapter 3 framework, where the individual blip-up/down images are first reconstructed to calculate a field map, and then a joint reconstruction of all blip-up/down images with the field map is performed to obtain the final distortion-corrected image.

The simulation study supports the feasibility of our method. By accounting for

distortion differences, motion-induced phase variations are reliably estimated for both blip-up and blip-down data (Fig. 6.1c, ii, iv), enabling successful reconstruction of the blip-up (i) and blip-down (iii) images. While individual reconstructions of the blip-up/down images show some blurring (Fig. 6.1c, d, i, iii), presumably due to the k_z partial Fourier sampling and high under-sampling factor, the estimated field map closely resembles the ground truth field map from fully sampled data (Fig. 6.1b-d, v). The final joint reconstruction (Fig. 6.1b-d, vi) demonstrates reduced blurring and artifacts, as the overall sampling effectively covered the full range of k_z .

The performance of this method can be further enhanced through refined sampling strategies and the incorporation of tailored constraints into reconstruction. One promising avenue for sampling optimisation is modifying the greedy search framework from Chapter 4 to achieve a balanced optimisation for both individual blip-up/down reconstructions and the joint reconstruction. The pseudo-random sampling with a high under-sampling factor is likely to benefit from compressed sensing constraints [86]. Additionally, conducting further *in-vivo* experiments, particularly at high field strengths (e.g., 7T), would provide valuable evidence of the method's ability to improve anatomical fidelity and SNR efficiency for high-resolution dMRI.

Fast SNR-Efficient 3D Multi-Slab Diffusion MRI with Self-Navigation and Improved Slab-Boundary Artifacts Correction

The self-navigated imaging framework introduced in Chapter 4 eliminates the requirement for navigator acquisitions in 3D multi-slab dMRI, enabling faster and more SNR-efficient diffusion imaging. This potential has been demonstrated in *in-vivo* experiments at 7T (Chapter 4), where the minimum allowable TR constrained by SAR was reduced from 3.5 s to 2.4 s for 1.09 mm resolution imaging. At standard field strengths like 3T, where SAR constraints are less restrictive, the TR could be further shortened. For example, assuming a TE of 75 ms for 1 mm resolution imaging at 3T, the TR could be reduced to less than 1.4 s for a 14-slab acquisition (i.e., less than 100 ms per slab) if SAR constraints can be ignored. Additionally, as demonstrated in Fig. 4.11, the greedy search-based sampling optimisation from

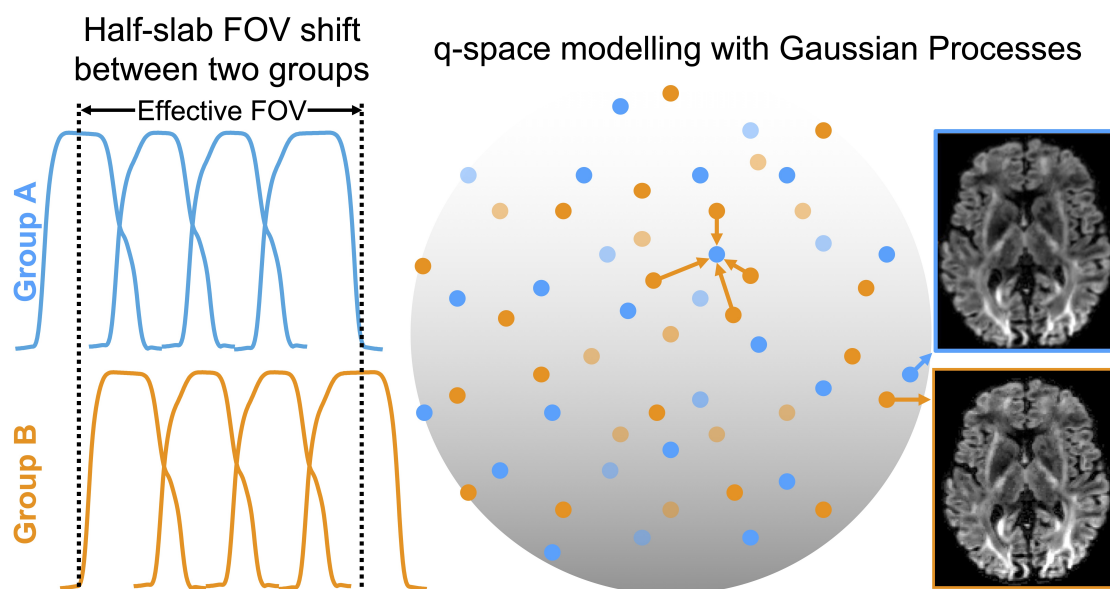


Figure 6.2: FOV shifting and q-space modelling for improved slab boundary artifacts correction. Diffusion directions are split into two groups (blue and yellow), each uniformly distributed in q-space. Each q-space point can be estimated from neighbouring points in the other group. A half-slab FOV shift between groups creates complementary slab profiles, benefiting accurate slab profile estimation.

Chapter 4 allows for a reduction in the number of shots without significantly compromising image quality. These advances suggest that a high-resolution (e.g., 1 mm isotropic) imaging protocol with a TR of less than 1.5 s and fewer than 10 shots (assuming 12 k_z sampling as in Chapter 4) is feasible, leading to a short scan time of under 15 seconds per volume. Importantly, this TR duration falls within the SNR-efficient range (i.e., TR=1-2 s).

A significant challenge in 3D multi-slab imaging is the presence of slab boundary artifacts, which become more pronounced as TR shortens [62]. In Chapters 3-5, “NPEN” [62] was employed to mitigate these artifacts by jointly estimating the slab profile and the corrected image through a nonlinear optimisation. However, with extremely short TRs, further improvements are needed for effective artifact reduction. Due to the elimination of navigator acquisition, sharper RF pulses can be designed and applied within SAR limits, reducing slab boundary artifacts.

Additionally, the performance of slab boundary artifact correction can be potentially improved using FOV shifting combined with q-space modelling across diffusion encoding directions (Fig. 6.2). In this proposed method, diffusion encoding

directions are divided into two groups, each uniformly distributed in q-space. A half-slab FOV shift is applied between the two groups along the slice direction, allowing them to capture complementary slab profile information. Using Gaussian Processes [130, 132], q-space points for each group can be estimated by leveraging neighbouring points from the other group. This strategy produces two sets of images: one native image and another estimated image that has identical contrast but encodes complementary slab profile information. By combining the complementary slab profiles from both images, the method potentially allows for more accurate slab profile estimation, leading to improved correction of slab boundary artifacts even in the context of very short TRs. Furthermore, as suggested in [130], Gaussian Process-based estimation may also benefit the under-sampled reconstruction of each q-space point, especially for slab-boundary slices where SNR is extremely limited.

One limitation of this method is the need to excite an additional slab due to the reduced effective FOV caused by FOV shifting (Fig. 6.2). However, since the self-navigated method reduces imaging time for each slab to less than 100 ms, this additional requirement is not expected to significantly extend the overall scan time. Another consideration is that robust q-space modelling typically requires a large number of diffusion encoding directions (e.g., 64). Given that each direction takes only ~ 15 s to acquire, this requirement can still be accommodated within a standard imaging protocol.

Self-Navigated Submillimetre 3D Multi-Slab Diffusion MRI

To integrate the self-navigation technique from Chapter 4 into the submillimetre dMRI framework described in Chapter 5 is another promising future direction. In submillimetre dMRI, navigator acquisitions suffer from long echo times due to the large matrix size (Table 5.1), which compromises scan efficiency, reduces navigator SNR, and potentially affects phase error correction. Therefore, the self-navigation method offers the potential to improve phase error correction performance while significantly shortening scan time and enhancing SNR efficiency.

However, implementing self-navigation in the in-plane segmented 3D multi-slab EPI used in Chapter 5 poses challenges, particularly due to high under-sampling in 2D phase map reconstruction. In Chapter 4, with $R_y = 3$, only around 10 k_y lines per shot were available to reconstruct a 180×180 phase map for 1.22 mm imaging. As matrix size and R_y increase in submillimetre imaging, the under-sampling factor becomes even more severe. Reducing the in-plane segmentation factor or grouping shots for structured low-rank reconstruction [89] could mitigate this issue by lowering the effective acceleration factor.

Another challenge is designing an effective sampling pattern when the number of shots is large. Ensuring each shot intersects the $k_z = 0$ plane concentrates points near the k-space centre (Fig. 4.2e, f), which is more pronounced in submillimetre imaging (e.g., 144 shots per slab for the 3T 0.65 mm protocol). Further considerations, such as incorporating point spread function analysis into sampling optimisation, is needed to maintain reconstruction quality and minimise blurring.

Deep Learning Reconstruction for Submillimetre 3D Diffusion MRI

Deep learning-based methods hold great promise to improve image quality and reduce computation time for submillimetre 3D multi-slab dMRI. Deep learning reconstruction methods have demonstrated the ability to achieve higher SNR than conventional model-based MRI reconstructions [133]. Additionally, specialized deep learning-based denoisers for dMRI, like Patch2Self [124], DeepDTI [125], and SDnDTI [126], could be integrated into our denoiser-regularised reconstruction framework, which are expected to outperform the conventional denoisers used in Chapter 5. Furthermore, reconstructing submillimetre data is currently time-intensive, taking approximately 12 hours per slab on a single CPU for data from the 3T 0.53 mm protocol. The significantly faster reconstruction times offered by trained deep learning models present a valuable opportunity to improve computational efficiency in this scenario.

The practical challenges of obtaining high-SNR reference data for submillimetre dMRI require additional considerations in deep learning approaches. Self-supervised

denoising methods, such as SDnDTI and Patch2Self, are particularly promising as they enhance SNR without the need for high-SNR reference data. For reconstruction, unrolled networks [134] have shown robust performance even with limited training data (e.g., as few as four subjects for MoDL [131]), making them suitable candidates for our application.

Analysis Methods Development for Submillimetre Diffusion MRI

The robust acquisition of high-quality submillimetre dMRI data opens exciting possibilities for neuroscience research. In Chapter 5, preliminary results demonstrated the potential of submillimetre dMRI to address “gyral bias” and improve U-fibre mapping. More quantitative and comprehensive analyses could further leverage the rich information provided by ultrahigh-resolution data.

One promising research direction is the accurate reconstruction of whole-brain intra- and inter-gyri connections at submillimetre resolutions, supporting detailed studies of brain connectivity. Current frameworks [135, 136] are optimised for conventional resolutions (1-2 mm) and may not fully exploit the finer details available in submillimetre data. By incorporating curvature information from cortical surfaces derived from T_1 -weighted images, it may be possible to enable the segmentation of whole-brain gyri, facilitating submillimetre connectivity studies.

Submillimetre dMRI may also help delineate small but critical subcortical structures, such as those targeted for deep brain stimulation in treating conditions like Parkinson’s disease, essential tremor, and dystonia [137]. Publicly available subcortical atlases (e.g., <https://www.lead-dbs.org/helpsupport/knowledge-base/atlassesresources/atlasses-2/>) provide valuable resources for defining seeding and targeting masks for tractography, enabling the reconstruction of white matter circuits in subcortical regions. A prior study [100] successfully reconstructed these circuits using 0.76 mm dMRI data acquired in 18 hours [127]. Our methods from Chapter 5 hold promise for achieving similar results within clinically feasible scan times.

6.3 Final Remarks

Over the past few decades, advances in diffusion acquisition and reconstruction methods have paved the way for numerous groundbreaking neuroimaging studies, such as the Human Connectome Project (HCP) and UK Biobank (UKB), significantly deepening our understanding of the human brain. These advances were primarily based on 2D simultaneous multislice acquisition, which faces inherent limitations in achieving higher resolution and SNR. The development of 3D multi-slab imaging has further improved SNR efficiency and enabled higher spatial resolutions for dMRI, opening new possibilities in the “post HCP and UKB era” [3, 138]. This thesis addresses several key challenges in 3D multi-slab dMRI with the aim to further advance imaging fidelity and efficiency. It will be exciting to see how 3D multi-slab imaging and the methods developed here might contribute to next-generation neuroimaging studies, bringing us closer to unravelling the complexities of the human brain.

References

- [1] Benedetta Bodini and Olga Ciccarelli. “Chapter 11 - Diffusion MRI in Neurological Disorders”. In: *Diffusion MRI (Second Edition)*. Ed. by Heidi Johansen-Berg and Timothy E.J. Behrens. Second Edition. San Diego: Academic Press, 2014, pp. 241–255. URL: <https://www.sciencedirect.com/science/article/pii/B9780123964601000111>.
- [2] Karla L. Miller et al. “Diffusion imaging of whole, post-mortem human brains on a clinical MRI scanner”. In: *NeuroImage* 57.1 (July 2011), pp. 167–181. URL: <http://dx.doi.org/10.1016/j.neuroimage.2011.03.070>.
- [3] Wenchuan Wu and Karla L. Miller. “Image formation in diffusion MRI: A review of recent technical developments”. In: *Journal of Magnetic Resonance Imaging* 46.3 (Feb. 2017), pp. 646–662. URL: <http://dx.doi.org/10.1002/jmri.25664>.
- [4] Mathias Engström and Stefan Skare. “Diffusion-weighted 3D multislabs echo planar imaging for high signal-to-noise ratio efficiency and isotropic image resolution”. In: *Magnetic Resonance in Medicine* 70.6 (Jan. 2013), pp. 1507–1514. URL: <http://dx.doi.org/10.1002/mrm.24594>.
- [5] Robert Frost et al. “3D Multi-slab diffusion-weighted readout-segmented EPI with real-time cardiac-reordered k-space acquisition”. In: *Magnetic Resonance in Medicine* 72.6 (Dec. 2013), pp. 1565–1579. URL: <http://dx.doi.org/10.1002/mrm.25062>.
- [6] Hing-Chiu Chang et al. “Human brain diffusion tensor imaging at submillimeter isotropic resolution on a 3 Tesla clinical MRI scanner”. In: *NeuroImage* 118 (Sept. 2015), pp. 667–675. URL: <http://dx.doi.org/10.1016/j.neuroimage.2015.06.016>.
- [7] Wenchuan Wu et al. “High-resolution diffusion MRI at 7T using a three-dimensional multi-slab acquisition”. In: *NeuroImage* 143 (Dec. 2016), pp. 1–14. URL: <http://dx.doi.org/10.1016/j.neuroimage.2016.08.054>.
- [8] Mathias Engström et al. “On the signal-to-noise ratio efficiency and slab-banding artifacts in three-dimensional multislabs diffusion-weighted echo-planar imaging”. In: *Magnetic Resonance in Medicine* 73.2 (Mar. 2014), pp. 718–725. URL: <http://dx.doi.org/10.1002/mrm.25182>.
- [9] Jesper L.R. 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 (Oct. 2003), pp. 870–888. URL: [http://dx.doi.org/10.1016/S1053-8119\(03\)00336-7](http://dx.doi.org/10.1016/S1053-8119(03)00336-7).

- [10] Dominic Holland, Joshua M. Kuperman, and Anders M. Dale. “Efficient correction of inhomogeneous static magnetic field-induced distortion in Echo Planar Imaging”. In: *NeuroImage* 50.1 (Mar. 2010), pp. 175–183. URL: <http://dx.doi.org/10.1016/j.neuroimage.2009.11.044>.
- [11] M. Okan Irfanoglu et al. “DR-BUDDI (Diffeomorphic Registration for Blip-Up blip-Down Diffusion Imaging) method for correcting echo planar imaging distortions”. In: *NeuroImage* 106 (Feb. 2015), pp. 284–299. URL: <http://dx.doi.org/10.1016/j.neuroimage.2014.11.042>.
- [12] Benjamin Zahneisen et al. “Extended hybrid-space SENSE for EPI: Off-resonance and eddy current corrected joint interleaved blip-up/down reconstruction”. In: *NeuroImage* 153 (June 2017), pp. 97–108. URL: <http://dx.doi.org/10.1016/j.neuroimage.2017.03.052>.
- [13] Iain P. Bruce, Christopher Petty, and Allen W. Song. “Simultaneous and inherent correction of B0 and eddy-current induced distortions in high-resolution diffusion MRI using reversed polarity gradients and multiplexed sensitivity encoding (RPG-MUSE)”. In: *NeuroImage* 183 (Dec. 2018), pp. 985–993. URL: <http://dx.doi.org/10.1016/j.neuroimage.2018.09.055>.
- [14] Congyu Liao et al. “Distortion-free, high-isotropic-resolution diffusion MRI with gSlider BUDA-EPI and multicoil dynamic B0 shimming”. In: *Magnetic Resonance in Medicine* 86.2 (Mar. 2021), pp. 791–803. URL: <http://dx.doi.org/10.1002/mrm.28748>.
- [15] Steen Moeller et al. “Self-navigation for 3D multishot EPI with data-reference”. In: *Magnetic Resonance in Medicine* 84.4 (Mar. 2020), pp. 1747–1762. URL: <http://dx.doi.org/10.1002/mrm.28231>.
- [16] Zhipei Liang and Paul C. Lauterbur. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. New York: IEEE Press, 2000.
- [17] Matt A. Bernstein, Kevin F. King, and Xiaohong Joe Zhou. *Handbook of MRI Pulse Sequences*. Amsterdam: Elsevier Academic Press, 2004.
- [18] Sean Deoni. “Chapter 1 - Biophysical and Physiological Principles of T1 and T2”. In: *Quantitative Magnetic Resonance Imaging*. Ed. by Nicole Seiberlich et al. Vol. 1. Advances in Magnetic Resonance Technology and Applications. Academic Press, 2020, pp. 3–17. URL: <https://www.sciencedirect.com/science/article/pii/B9780128170571000032>.
- [19] Greg J. Stanisz et al. “T1, T2 relaxation and magnetization transfer in tissue at 3T”. In: *Magnetic Resonance in Medicine* 54.3 (Aug. 2005), pp. 507–512. URL: <http://dx.doi.org/10.1002/mrm.20605>.
- [20] Allen D. Elster. *Slice-selective excitation*. 2024. URL: <https://mriquestions.com/slice-selective-excitation.html>.
- [21] Klaas P. Pruessmann et al. “SENSE: Sensitivity encoding for fast MRI”. In: *Magnetic Resonance in Medicine* 42.5 (Nov. 1999), pp. 952–962. URL: [http://dx.doi.org/10.1002/\(SICI\)1522-2594\(199911\)42:5%3C952::AID-MRM16%3E3.0.CO;2-S](http://dx.doi.org/10.1002/(SICI)1522-2594(199911)42:5%3C952::AID-MRM16%3E3.0.CO;2-S).
- [22] Mark A. Griswold et al. “Generalized autocalibrating partially parallel acquisitions (GRAPPA)”. In: *Magnetic Resonance in Medicine* 47.6 (June 2002), pp. 1202–1210. URL: <http://dx.doi.org/10.1002/mrm.10171>.

- [23] Michael Lustig and John M. Pauly. “SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space”. In: *Magnetic Resonance in Medicine* 64.2 (June 2010), pp. 457–471. URL: <http://dx.doi.org/10.1002/mrm.22428>.
- [24] Martin Uecker et al. “ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA”. In: *Magnetic Resonance in Medicine* 71.3 (May 2013), pp. 990–1001. URL: <http://dx.doi.org/10.1002/mrm.24751>.
- [25] Klaas P. Pruessmann et al. “Advances in sensitivity encoding with arbitrary k-space trajectories”. In: *Magnetic Resonance in Medicine* 46.4 (Oct. 2001), pp. 638–651. URL: <http://dx.doi.org/10.1002/mrm.1241>.
- [26] Michael Lustig, David Donoho, and John M. Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic Resonance in Medicine* 58.6 (Oct. 2007), pp. 1182–1195. URL: <http://dx.doi.org/10.1002/mrm.21391>.
- [27] Dong Liang et al. “Accelerating SENSE using compressed sensing”. In: *Magnetic Resonance in Medicine* 62.6 (Sept. 2009), pp. 1574–1584. URL: <http://dx.doi.org/10.1002/mrm.22161>.
- [28] Peter J. Shin et al. “Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion: Calibrationless Parallel Imaging”. In: *Magnetic Resonance in Medicine* 72.4 (Nov. 2013), pp. 959–970. URL: <http://dx.doi.org/10.1002/mrm.24997>.
- [29] Kyong Hwan Jin, Dongwook Lee, and Jong Chul Ye. “A General Framework for Compressed Sensing and Parallel MRI Using Annihilating Filter Based Low-Rank Hankel Matrix”. In: *IEEE Transactions on Computational Imaging* 2.4 (Dec. 2016), pp. 480–495. URL: <http://dx.doi.org/10.1109/TCI.2016.2601296>.
- [30] Hammad Omer and Robert Dickinson. “Regularization in parallel MR image reconstruction”. In: *Concepts in Magnetic Resonance Part A* 38A.2 (Mar. 2011), pp. 52–60. URL: <http://dx.doi.org/10.1002/cmr.a.20206>.
- [31] Felix A. Breuer et al. “Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA)”. In: *Magnetic Resonance in Medicine* 55.3 (Jan. 2006), pp. 549–556. URL: <http://dx.doi.org/10.1002/mrm.20787>.
- [32] Kawin Setsompop et al. “Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty”. In: *Magnetic Resonance in Medicine* 67.5 (Aug. 2011), pp. 1210–1224. URL: <http://dx.doi.org/10.1002/mrm.23097>.
- [33] Berkin Bilgic et al. “Wave-CAIPI for highly accelerated 3D imaging”. In: *Magnetic Resonance in Medicine* 73.6 (July 2014), pp. 2152–2162. URL: <http://dx.doi.org/10.1002/mrm.25347>.
- [34] David J Larkman and Rita G Nunes. “Parallel magnetic resonance imaging”. In: *Physics in Medicine and Biology* 52.7 (Mar. 2007), R15–R55. URL: <http://dx.doi.org/10.1088/0031-9155/52/7/R01>.

- [35] Jesse Hamilton, Dominique Franson, and Nicole Seiberlich. “Recent advances in parallel imaging for MRI”. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 101 (Aug. 2017), pp. 71–95. URL: <http://dx.doi.org/10.1016/j.pnmrs.2017.04.002>.
- [36] Felix A. Breuer et al. “General formulation for quantitative G-factor calculation in GRAPPA reconstructions”. In: *Magnetic Resonance in Medicine* 62.3 (July 2009), pp. 739–746. URL: <http://dx.doi.org/10.1002/mrm.22066>.
- [37] Philip M. Robson et al. “Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions”. In: *Magnetic Resonance in Medicine* 60.4 (Sept. 2008), pp. 895–907. URL: <http://dx.doi.org/10.1002/mrm.21728>.
- [38] Heidi Johansen-Berg and Timothy E.J. Behrens. *Diffusion MRI: From Quantitative Measurement to In vivo Neuroanatomy (Second Edition)*. Amsterdam: Academic Press, 2014.
- [39] Vinit Baliyan et al. “Diffusion weighted imaging: Technique and applications”. In: *World Journal of Radiology* 8.9 (2016), p. 785. URL: <http://dx.doi.org/10.4329/wjr.v8.i9.785>.
- [40] James G. Pipe, Victoria G. Farthing, and Kirsten P. Forbes. “Multishot diffusion-weighted FSE using PROPELLER MRI”. In: *Magnetic Resonance in Medicine* 47.1 (Dec. 2001), pp. 42–52. URL: <http://dx.doi.org/10.1002/mrm.10014>.
- [41] Chunlei Liu et al. “Self-navigated interleaved spiral (SNAILS): Application to high-resolution diffusion tensor imaging”. In: *Magnetic Resonance in Medicine* 52.6 (Nov. 2004), pp. 1388–1396. URL: <http://dx.doi.org/10.1002/mrm.20288>.
- [42] Tie-Qiang Li, Dong-Hyun Kim, and Michael E. Moseley. “High-resolution diffusion-weighted imaging with interleaved variable-density spiral acquisitions”. In: *Journal of Magnetic Resonance Imaging* 21.4 (Mar. 2005), pp. 468–475. URL: <http://dx.doi.org/10.1002/jmri.20287>.
- [43] David A. Feinberg et al. “Multiplexed Echo Planar Imaging for Sub-Second Whole Brain fMRI and Fast Diffusion Imaging”. In: *PLoS ONE* 5.12 (Dec. 2010). Ed. by Pedro Antonio Valdes-Sosa, e15710. URL: <http://dx.doi.org/10.1371/journal.pone.0015710>.
- [44] Markus Barth et al. “Simultaneous multislice (SMS) imaging techniques”. In: *Magnetic Resonance in Medicine* 75.1 (Aug. 2015), pp. 63–81. URL: <http://dx.doi.org/10.1002/mrm.25897>.
- [45] Stamatiou N. Sotiropoulos et al. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project”. In: *NeuroImage* 80 (Oct. 2013), pp. 125–143. URL: <http://dx.doi.org/10.1016/j.neuroimage.2013.05.057>.
- [46] 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: *Magnetic Resonance in Medicine* 62.2 (May 2009), pp. 468–475. URL: <http://dx.doi.org/10.1002/mrm.22024>.

- [47] Kawin Setsompop et al. “High-resolution in vivo diffusion imaging of the human brain with generalized slice dithered enhanced resolution: Simultaneous multislice (gSlider-SMS)”. In: *Magnetic Resonance in Medicine* 79.1 (Mar. 2017), pp. 141–151. URL: <http://dx.doi.org/10.1002/mrm.26653>.
- [48] Benoit Scherrer, Ali Gholipour, and Simon K. Warfield. “Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions”. In: *Medical Image Analysis* 16.7 (Oct. 2012), pp. 1465–1476. URL: <http://dx.doi.org/10.1016/j.media.2012.05.003>.
- [49] Geraline Vis et al. “Accuracy and precision in super-resolution MRI: Enabling spherical tensor diffusion encoding at ultra-high b-values and high resolution”. In: *NeuroImage* 245 (Dec. 2021), p. 118673. URL: <http://dx.doi.org/10.1016/j.neuroimage.2021.118673>.
- [50] Zijng Dong et al. “Romer-EPTI: rotating-view motion-robust super-resolution EPTI for SNR-efficient distortion-free in-vivo mesoscale dMRI and microstructure imaging”. In: *bioRxiv* (Jan. 2024). URL: <http://dx.doi.org/10.1101/2024.01.26.577343>.
- [51] Karla L. Miller and John M. Pauly. “Nonlinear phase correction for navigated diffusion imaging”. In: *Magnetic Resonance in Medicine* 50.2 (July 2003), pp. 343–353. URL: <http://dx.doi.org/10.1002/mrm.10531>.
- [52] 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 (May 2013), pp. 41–47. URL: <http://dx.doi.org/10.1016/j.neuroimage.2013.01.038>.
- [53] Merry Mani et al. “Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS): Annihilating Filter K-Space Formulation for Multi-Shot DWI Recovery”. In: *Magnetic Resonance in Medicine* 78.2 (Aug. 2016), pp. 494–507. URL: <http://dx.doi.org/10.1002/mrm.26382>.
- [54] Ziyu Li et al. “Sampling strategies and integrated reconstruction for reducing distortion and boundary slice aliasing in high-resolution 3D diffusion MRI”. In: *Magnetic Resonance in Medicine* 90.4 (June 2023), pp. 1484–1501. URL: <http://dx.doi.org/10.1002/mrm.29741>.
- [55] P.J. Basser, J. Mattiello, and D. LeBihan. “MR diffusion tensor spectroscopy and imaging”. In: *Biophysical Journal* 66.1 (Jan. 1994), pp. 259–267. URL: [http://dx.doi.org/10.1016/S0006-3495\(94\)80775-1](http://dx.doi.org/10.1016/S0006-3495(94)80775-1).
- [56] Jens H. Jensen et al. “Diffusional kurtosis imaging: The quantification of non-gaussian water diffusion by means of magnetic resonance imaging”. In: *Magnetic Resonance in Medicine* 53.6 (May 2005), pp. 1432–1440. URL: <http://dx.doi.org/10.1002/mrm.20508>.
- [57] Alard Roebroeck et al. “High-resolution diffusion tensor imaging and tractography of the human optic chiasm at 9.4 T”. In: *NeuroImage* 39.1 (Jan. 2008), pp. 157–168. URL: <http://dx.doi.org/10.1016/j.neuroimage.2007.08.015>.
- [58] Jennifer A. McNab et al. “Surface based analysis of diffusion orientation for identifying architectonic domains in the in vivo human cortex”. In: *NeuroImage* 69 (Apr. 2013), pp. 87–100. URL: <http://dx.doi.org/10.1016/j.neuroimage.2012.11.065>.

- [59] Erpeng Dai, Simin Liu, and Hua Guo. “High-resolution whole-brain diffusion MRI at 3T using simultaneous multi-slab (SMSlab) acquisition”. In: *NeuroImage* 237 (Aug. 2021), p. 118099. URL: <http://dx.doi.org/10.1016/j.neuroimage.2021.118099>.
- [60] Peter Jezzard and Robert S. Balaban. “Correction for geometric distortion in echo planar images from B0 field variations”. In: *Magnetic Resonance in Medicine* 34.1 (July 1995), pp. 65–73. URL: <http://dx.doi.org/10.1002/mrm.1910340111>.
- [61] M. Okan Irfanoglu et al. “Evaluating corrections for Eddy-currents and other EPI distortions in diffusion MRI: methodology and a dataset for benchmarking”. In: *Magnetic Resonance in Medicine* 81.4 (Nov. 2018), pp. 2774–2787. URL: <http://dx.doi.org/10.1002/mrm.27577>.
- [62] Wenchuan Wu et al. “Reducing slab boundary artifacts in three-dimensional multislabs diffusion MRI using nonlinear inversion for slab profile encoding (NPEN)”. In: *Magnetic Resonance in Medicine* 76.4 (Oct. 2015), pp. 1183–1195. URL: <http://dx.doi.org/10.1002/mrm.26027>.
- [63] Stephen M. Smith et al. “Advances in functional and structural MR image analysis and implementation as FSL”. In: *NeuroImage* 23 (Jan. 2004), S208–S219. URL: <http://dx.doi.org/10.1016/j.neuroimage.2004.07.051>.
- [64] Mark Jenkinson et al. “FSL”. In: *NeuroImage* 62.2 (Aug. 2012), pp. 782–790. URL: <http://dx.doi.org/10.1016/j.neuroimage.2011.09.015>.
- [65] Feng Huang et al. “A rapid and robust numerical algorithm for sensitivity encoding with sparsity constraints: Self-feeding sparse SENSE”. In: *Magnetic Resonance in Medicine* 64.4 (June 2010), pp. 1078–1088. URL: <http://dx.doi.org/10.1002/mrm.22504>.
- [66] Tao Zhang et al. “Coil compression for accelerated imaging with Cartesian sampling”. In: *Magnetic Resonance in Medicine* 69.2 (Apr. 2012), pp. 571–582. URL: <http://dx.doi.org/10.1002/mrm.24267>.
- [67] Adam O. Kettinger et al. “Full utilization of conjugate symmetry: combining virtual conjugate coil reconstruction with partial Fourier imaging for g-factor reduction in accelerated MRI”. In: *Magnetic Resonance in Medicine* 82.3 (May 2019), pp. 1073–1090. URL: <http://dx.doi.org/10.1002/mrm.27799>.
- [68] Elias Kellner et al. “Gibbs-ringing artifact removal based on local subvoxel-shifts”. In: *Magnetic Resonance in Medicine* 76.5 (Nov. 2015), pp. 1574–1581. URL: <http://dx.doi.org/10.1002/mrm.26054>.
- [69] Thea Bautista et al. “Removal of Gibbs ringing artefacts for 3D acquisitions using subvoxel shifts”. In: *Proc Int Soc Magn Reson Med*. Vol. 29. 2021, p. 3535.
- [70] Jesper L.R. Andersson and Stamatios N. Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *NeuroImage* 125 (Jan. 2016), pp. 1063–1078. URL: <http://dx.doi.org/10.1016/j.neuroimage.2015.10.019>.
- [71] Mark Jenkinson and Stephen Smith. “A global optimisation method for robust affine registration of brain images”. In: *Medical Image Analysis* 5.2 (June 2001), pp. 143–156. URL: [http://dx.doi.org/10.1016/S1361-8415\(01\)00036-6](http://dx.doi.org/10.1016/S1361-8415(01)00036-6).

- [72] Mark Jenkinson et al. “Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images”. In: *NeuroImage* 17.2 (Oct. 2002), pp. 825–841. URL: <http://dx.doi.org/10.1006/nimg.2002.1132>.
- [73] Douglas N. Greve and Bruce Fischl. “Accurate and robust brain image alignment using boundary-based registration”. In: *NeuroImage* 48.1 (Oct. 2009), pp. 63–72. URL: <http://dx.doi.org/10.1016/j.neuroimage.2009.06.060>.
- [74] T.E.J. Behrens et al. “Characterization and propagation of uncertainty in diffusion-weighted MR imaging”. In: *Magnetic Resonance in Medicine* 50.5 (Oct. 2003), pp. 1077–1088. URL: <http://dx.doi.org/10.1002/mrm.10609>.
- [75] T.E.J. Behrens et al. “Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?” In: *NeuroImage* 34.1 (Jan. 2007), pp. 144–155. URL: <http://dx.doi.org/10.1016/j.neuroimage.2006.09.018>.
- [76] Marius de Groot et al. “Improving alignment in Tract-based spatial statistics: Evaluation and optimization of image registration”. In: *NeuroImage* 76 (Aug. 2013), pp. 400–411. URL: <http://dx.doi.org/10.1016/j.neuroimage.2013.03.015>.
- [77] Y. Zhang, M. Brady, and S. 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. URL: <http://dx.doi.org/10.1109/42.906424>.
- [78] Zijong Dong et al. “Interleaved EPI diffusion imaging using SPIRiT-based reconstruction with virtual coil compression”. In: *Magnetic Resonance in Medicine* 79.3 (June 2017), pp. 1525–1531. URL: <http://dx.doi.org/10.1002/mrm.26768>.
- [79] Xiaozhi Cao et al. “MOCO-BUDA: motion-corrected blip-up/down acquisition with joint reconstruction for motion-robust and distortion-free diffusion MRI of the brain”. In: *Proc Int Soc Magn Reson Med*. Vol. 28. 2020, p. 465.
- [80] Ziyu Li et al. “Self-navigated 3D diffusion MRI using an optimized CAIPI sampling and structured low-rank reconstruction estimated navigator”. In: *IEEE Transactions on Medical Imaging* (2024), pp. 1–1. URL: <http://dx.doi.org/10.1109/TMI.2024.3454994>.
- [81] Kim Butts et al. “Diffusion-weighted interleaved echo-planar imaging with a pair of orthogonal navigator echoes”. In: *Magnetic Resonance in Medicine* 35.5 (May 1996), pp. 763–770. URL: <http://dx.doi.org/10.1002/mrm.1910350518>.
- [82] Merry Mani et al. “Improved MUSSELS reconstruction for high-resolution multi-shot diffusion weighted imaging”. In: *Magnetic Resonance in Medicine* 83.6 (Dec. 2019), pp. 2253–2263. URL: <http://dx.doi.org/10.1002/mrm.28090>.
- [83] Ivan Markovskiy. “Structured low-rank approximation and its applications”. In: *Automatica* 44.4 (Apr. 2008), pp. 891–909. URL: <http://dx.doi.org/10.1016/j.automatica.2007.09.011>.
- [84] Hemant K. Aggarwal, Merry P. Mani, and Mathews Jacob. “MoDL-MUSSELS: Model-Based Deep Learning for Multishot Sensitivity-Encoded Diffusion MRI”. In: *IEEE Transactions on Medical Imaging* 39.4 (Apr. 2020), pp. 1268–1277. URL: <http://dx.doi.org/10.1109/TMI.2019.2946501>.

- [85] Congyu Liao et al. “High-fidelity mesoscale in-vivo diffusion MRI through gSlider-BUDA and circular EPI with S-LORAKS reconstruction”. In: *NeuroImage* 275 (July 2023), p. 120168. URL: <http://dx.doi.org/10.1016/j.neuroimage.2023.120168>.
- [86] Matthias Seeger et al. “Optimization of k-space trajectories for compressed sensing by Bayesian experimental design”. In: *Magnetic Resonance in Medicine* 63.1 (Oct. 2009), pp. 116–126. URL: <http://dx.doi.org/10.1002/mrm.22180>.
- [87] M. Murphy et al. “Fast ℓ_1 -SPIRiT Compressed Sensing Parallel Imaging MRI: Scalable Parallel Implementation and Clinically Feasible Runtime”. In: *IEEE Transactions on Medical Imaging* 31.6 (June 2012), pp. 1250–1262. URL: <http://dx.doi.org/10.1109/TMI.2012.2188039>.
- [88] Stephen Boyd. “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers”. In: *Foundations and Trends® in Machine Learning* 3.1 (2010), pp. 1–122. URL: <http://dx.doi.org/10.1561/22000000016>.
- [89] Xi Chen, Wenchuan Wu, and Mark Chiew. “Improving robustness of 3D multi-shot EPI by structured low-rank reconstruction of segmented CAIPI sampling for fMRI at 7T”. In: *NeuroImage* 267 (Feb. 2023), p. 119827. URL: <http://dx.doi.org/10.1016/j.neuroimage.2022.119827>.
- [90] Chu-Yu Lee and Merry Mani. “2D CAIPI accelerated 3D multi-slab diffusion weighted EPI combined with qModeL reconstruction for fast high resolution microstructure imaging”. In: *Magnetic Resonance Imaging* 111 (Sept. 2024), pp. 57–66. URL: <http://dx.doi.org/10.1016/j.mri.2024.04.003>.
- [91] Erpeng Dai et al. “The effects of navigator distortion and noise level on interleaved EPI DWI reconstruction: a comparison between image- and k-space-based method”. In: *Magnetic Resonance in Medicine* 80.5 (Mar. 2018), pp. 2024–2032. URL: <http://dx.doi.org/10.1002/mrm.27190>.
- [92] Xi Chen, Wenchuan Wu, and Mark Chiew. “Motion compensated structured low-rank reconstruction for 3D multi-shot EPI”. In: *Magnetic Resonance in Medicine* 91.6 (Feb. 2024), pp. 2443–2458. URL: <http://dx.doi.org/10.1002/mrm.30019>.
- [93] Julian Maclaren et al. “Prospective motion correction in brain imaging: A review”. In: *Magnetic Resonance in Medicine* 69.3 (May 2012), pp. 621–636. URL: <http://dx.doi.org/10.1002/mrm.24314>.
- [94] Ziyu Li et al. “Submillimeter diffusion MRI using an in-plane segmented 3D multi-slab acquisition and denoiser-regularized reconstruction”. In: *bioRxiv* (Oct. 2024). URL: <http://dx.doi.org/10.1101/2024.10.10.617536>.
- [95] Benjamin C Tendler et al. “The Digital Brain Bank, an open access platform for post-mortem imaging datasets”. In: *eLife* 11 (Mar. 2022). URL: <http://dx.doi.org/10.7554/eLife.73153>.
- [96] Michiel Cottaar et al. “Modelling white matter in gyral blades as a continuous vector field”. In: *NeuroImage* 227 (Feb. 2021), p. 117693. URL: <http://dx.doi.org/10.1016/j.neuroimage.2020.117693>.

- [97] Silei Zhu et al. “Imaging the structural connectome with hybrid diffusion MRI-microscopy tractography”. In: *bioRxiv* (Jan. 2024). URL: <http://dx.doi.org/10.1101/2024.01.08.574641>.
- [98] Allen W. Song et al. “Improved Delineation of Short Cortical Association Fibers and Gray/White Matter Boundary Using Whole-Brain Three-Dimensional Diffusion Tensor Imaging at Submillimeter Spatial Resolution”. In: *Brain Connectivity* 4.9 (Nov. 2014), pp. 636–640. URL: <http://dx.doi.org/10.1089/brain.2014.0270>.
- [99] Minhui Ouyang et al. “Short-range connections in the developmental connectome during typical and atypical brain maturation”. In: *Neuroscience Biobehavioral Reviews* 83 (Dec. 2017), pp. 109–122. URL: <http://dx.doi.org/10.1016/j.neubiorev.2017.10.007>.
- [100] Chiara Maffei et al. “Submillimeter dMRI protocol optimization for accurate in-vivo reconstruction of deep-brain circuitry”. In: *Proc Int Soc Magn Reson Med*. Vol. 30. 2022, p. 160.
- [101] Kadharbatcha S. Saleem et al. “High-resolution mapping and digital atlas of subcortical regions in the macaque monkey based on matched MAP-MRI and histology”. In: *NeuroImage* 245 (Dec. 2021), p. 118759. URL: <http://dx.doi.org/10.1016/j.neuroimage.2021.118759>.
- [102] Sajjad Feizollah and Christine L. Tardif. “High-resolution diffusion-weighted imaging at 7 Tesla: Single-shot readout trajectories and their impact on signal-to-noise ratio, spatial resolution and accuracy”. In: *NeuroImage* 274 (July 2023), p. 120159. URL: <http://dx.doi.org/10.1016/j.neuroimage.2023.120159>.
- [103] 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 (Mar. 2015), pp. 665–679. URL: <http://dx.doi.org/10.1002/mrm.25628>.
- [104] Dimo Ivanov et al. “Robust ACS acquisition for 3D echo planar imaging”. In: *Proc Int Soc Magn Reson Med*. Vol. 23. 2015, p. 2059.
- [105] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. “Non-Local Means Denoising”. In: *Image Processing On Line* 1 (Sept. 2011), pp. 208–212. URL: http://dx.doi.org/10.5201/ipol.2011.bcm_nlm.
- [106] Kostadin Dabov et al. “Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering”. In: *IEEE Transactions on Image Processing* 16.8 (Aug. 2007), pp. 2080–2095. URL: <http://dx.doi.org/10.1109/TIP.2007.901238>.
- [107] M. Maggioni et al. “Nonlocal Transform-Domain Filter for Volumetric Data Denoising and Reconstruction”. In: *IEEE Transactions on Image Processing* 22.1 (Jan. 2013), pp. 119–133. URL: <http://dx.doi.org/10.1109/TIP.2012.2210725>.
- [108] Jelle Veraart et al. “Denoising of diffusion MRI using random matrix theory”. In: *NeuroImage* 142 (Nov. 2016), pp. 394–406. URL: <http://dx.doi.org/10.1016/j.neuroimage.2016.08.016>.

- [109] Steen Moeller et al. “NOise reduction with DIstribution Corrected (NORDIC) PCA in dMRI with complex-valued parameter-free locally low-rank processing”. In: *NeuroImage* 226 (Feb. 2021), p. 117539. URL: <http://dx.doi.org/10.1016/j.neuroimage.2020.117539>.
- [110] Jose Pedro Manzano Patron et al. “Denoising diffusion MRI: Considerations and implications for analysis”. In: *Imaging Neuroscience* 2 (Jan. 2024), pp. 1–29. URL: http://dx.doi.org/10.1162/imag_a_00060.
- [111] Rizwan Ahmad et al. “Plug-and-Play Methods for Magnetic Resonance Imaging: Using Denoisers for Image Recovery”. In: *IEEE Signal Processing Magazine* 37.1 (Jan. 2020), pp. 105–116. URL: <http://dx.doi.org/10.1109/MSP.2019.2949470>.
- [112] Kelvin J. Layton et al. “Pulseseq: A rapid and hardware-independent pulse sequence prototyping framework: Rapid Hardware-Independent Pulse Sequence Prototyping”. In: *Magnetic Resonance in Medicine* 77.4 (June 2016), pp. 1544–1552. URL: <http://dx.doi.org/10.1002/mrm.26235>.
- [113] J. 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 (Mar. 1991), pp. 53–65. URL: <http://dx.doi.org/10.1109/42.75611>.
- [114] John P. Mugler and James R. Brookeman. “Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE)”. In: *Magnetic Resonance in Medicine* 15.1 (July 1990), pp. 152–157. URL: <http://dx.doi.org/10.1002/mrm.1910150117>.
- [115] Matthew F Glasser et al. “The Human Connectome Project’s neuroimaging approach”. In: *Nature Neuroscience* 19.9 (Aug. 2016), pp. 1175–1187. URL: <http://dx.doi.org/10.1038/nn.4361>.
- [116] Justin P. Haldar and Jingwei Zhuo. “P-LORAKS: Low-rank modeling of local k-space neighborhoods with parallel imaging data”. In: *Magnetic Resonance in Medicine* 75.4 (May 2015), pp. 1499–1514. URL: <http://dx.doi.org/10.1002/mrm.25717>.
- [117] Bruce Fischl. “FreeSurfer”. In: *NeuroImage* 62.2 (Aug. 2012), pp. 774–781. URL: <http://dx.doi.org/10.1016/j.neuroimage.2012.01.021>.
- [118] J.-Donald Tournier et al. “Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution”. In: *NeuroImage* 23.3 (Nov. 2004), pp. 1176–1185. URL: <http://dx.doi.org/10.1016/j.neuroimage.2004.07.037>.
- [119] J-Donald Tournier, Fernando Calamante, and Alan Connelly. “Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution”. In: *NeuroImage* 35.4 (May 2007), pp. 1459–1472. URL: <http://dx.doi.org/10.1016/j.neuroimage.2007.02.016>.
- [120] J-Donald Tournier et al. “MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation”. In: *NeuroImage* 202 (Nov. 2019), p. 116137. URL: <http://dx.doi.org/10.1016/j.neuroimage.2019.116137>.

- [121] J Donald Tournier, Fernando Calamante, Alan Connelly, et al. “Improved probabilistic streamlines tractography by 2nd order integration over fibre orientation distributions”. In: *Proceedings of the international society for magnetic resonance in medicine*. Vol. 18. 2010, p. 1670.
- [122] Kamil Uğurbil et al. “Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project”. In: *NeuroImage* 80 (Oct. 2013), pp. 80–104. URL: <http://dx.doi.org/10.1016/j.neuroimage.2013.05.012>.
- [123] Chiara Maffei et al. “Topography of the human acoustic radiation as revealed by ex vivo fibers micro-dissection and in vivo diffusion-based tractography”. In: *Brain Structure and Function* 223.1 (Sept. 2017), pp. 449–459. URL: <http://dx.doi.org/10.1007/s00429-017-1471-6>.
- [124] Shreyas Fadnavis, Joshua Batson, and Eleftherios Garyfallidis. “Patch2Self: Denoising Diffusion MRI with Self-Supervised Learning”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle et al. Vol. 33. Curran Associates, Inc., 2020, pp. 16293–16303. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/bc047286b224b7bfa73d4cb02de1238d-Paper.pdf.
- [125] Qiyuan Tian et al. “DeepDTI: High-fidelity six-direction diffusion tensor imaging using deep learning”. In: *NeuroImage* 219 (Oct. 2020), p. 117017. URL: <http://dx.doi.org/10.1016/j.neuroimage.2020.117017>.
- [126] Qiyuan Tian et al. “SDnDTI: Self-supervised deep learning-based denoising for diffusion tensor MRI”. In: *NeuroImage* 253 (June 2022), p. 119033. URL: <http://dx.doi.org/10.1016/j.neuroimage.2022.119033>.
- [127] Fuyixue Wang et al. “In vivo human whole-brain Connectom diffusion MRI dataset at 760 μm isotropic resolution”. In: *Scientific Data* 8.1 (Apr. 2021). URL: <http://dx.doi.org/10.1038/s41597-021-00904-z>.
- [128] Susie Y. Huang et al. “Connectome 2.0: Developing the next-generation ultra-high gradient strength human MRI scanner for bridging studies of the micro-, meso- and macro-connectome”. In: *NeuroImage* 243 (Nov. 2021), p. 118530. URL: <http://dx.doi.org/10.1016/j.neuroimage.2021.118530>.
- [129] David A. Feinberg et al. “Next-generation MRI scanner designed for ultra-high-resolution human brain imaging at 7 Tesla”. In: *Nature Methods* 20.12 (Nov. 2023), pp. 2048–2057. URL: <http://dx.doi.org/10.1038/s41592-023-02068-7>.
- [130] Wenchuan Wu et al. “Diffusion Acceleration with Gaussian process Estimated Reconstruction (DAGER)”. In: *Magnetic Resonance in Medicine* 82.1 (Mar. 2019), pp. 107–125. URL: <http://dx.doi.org/10.1002/mrm.27699>.
- [131] Hemant K. Aggarwal, Merry P. Mani, and Mathews Jacob. “MoDL: Model-Based Deep Learning Architecture for Inverse Problems”. In: *IEEE Transactions on Medical Imaging* 38.2 (Feb. 2019), pp. 394–405. URL: <http://dx.doi.org/10.1109/TMI.2018.2865356>.

- [132] Carl Edward Rasmussen. “Gaussian Processes in Machine Learning”. In: *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures*. Ed. by Olivier Bousquet, Ulrike von Luxburg, and Gunnar Rätsch. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 63–71. URL: https://doi.org/10.1007/978-3-540-28650-9_4.
- [133] Florian Knoll et al. “Deep-Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A Survey of the Current Approaches, Trends, and Issues”. In: *IEEE Signal Processing Magazine* 37.1 (Jan. 2020), pp. 128–140. URL: <http://dx.doi.org/10.1109/MSP.2019.2950640>.
- [134] Dong Liang et al. *Deep MRI Reconstruction: Unrolled Optimization Algorithms Meet Neural Networks*. 2019. arXiv: 1907.11711 [eess.IV]. URL: <https://arxiv.org/abs/1907.11711>.
- [135] Dmitri Shastin et al. “Surface-based tracking for short association fibre tractography”. In: *NeuroImage* 260 (Oct. 2022), p. 119423. URL: <http://dx.doi.org/10.1016/j.neuroimage.2022.119423>.
- [136] Kurt G. Schilling et al. “Short superficial white matter and aging: A longitudinal multi-site study of 1293 subjects and 2711 sessions”. In: *Aging Brain* 3 (2023), p. 100067. URL: <http://dx.doi.org/10.1016/j.nbas.2023.100067>.
- [137] Siobhan Ewert et al. “Toward defining deep brain stimulation targets in MNI space: A subcortical atlas based on multimodal MRI, histology and structural connectivity”. In: *NeuroImage* 170 (Apr. 2018), pp. 271–282. URL: <http://dx.doi.org/10.1016/j.neuroimage.2017.05.015>.
- [138] Steen Moeller et al. “Diffusion Imaging in the Post HCP Era”. In: *Journal of Magnetic Resonance Imaging* 54.1 (June 2020), pp. 36–57. URL: <http://dx.doi.org/10.1002/jmri.27247>.