

Advances in Myelin Water Imaging and Stack-of-Spirals MRI: Image Reconstruction and Parameter Estimation

by

Naveen N. Murthy

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical and Computer Engineering)
in the University of Michigan
2024

Doctoral Committee:

Professor Jeffrey A. Fessler, Co-Chair
Research Associate Professor Jon-Fredrik Nielsen, Co-Chair
Associate Professor Raj Rao Nadakuditi
Professor Nicole Seiberlich

Naveen N. Murthy
nnmurthy@umich.edu
ORCID iD: 0000-0003-1365-3302

© Naveen N. Murthy 2024

ACKNOWLEDGMENTS

The PhD journey is a long and winding one, filled with highs and lows, and my personal journey has been made immeasurably better by all the people I met and interacted with along the way. I want to take this opportunity to acknowledge the part that everyone played in this journey of mine.

I am deeply thankful to my co-advisors, Jeffrey A. Fessler and Jon-Fredrik Nielsen, for their unwavering and invaluable support throughout my doctoral journey. Both of them have been incredibly supportive throughout my PhD, especially in a post-COVID world, and I am extremely grateful for that. Their words of encouragement through the years have meant a lot to me, and allowed me to keep progressing in my journey.

I first interacted with Jeff through the Image Processing course offered by him during my Master's degree at University of Michigan. Having thoroughly enjoyed the course, I reached out to Jeff again while working in the industry (after my Master's), expressing my intention to pursue a PhD. I am extremely grateful to him for welcoming me into his group and for serving as an exceptional mentor throughout my doctoral studies. Under Jeff's guidance, I have gained a wide range of experience working in various imaging problems, from quantitative MRI and MR image reconstruction (covered in this dissertation), to my early work with X-ray CT reconstruction [123, 109]. He allowed me to explore my research interests independently, while still providing me with the essential guidance and tools to make informed decisions. I have always admired several aspects of Jeff's advising style: his attention to detail, his thoroughness in thinking through a research problem, and his willingness to roll up his sleeves and work on the small details of any problem with me, whenever I was stuck.

I began working with Jon midway through my PhD, and he has been instrumental with all my projects in MRI, along with being a great resource for all my MRI related questions. He was very helpful with acquiring data for a lot of my projects, as well as explaining some of the practical issues to be considered for scanning. Jon has also been a great mentor, and has given me a lot of guidance during my PhD. I have learnt a lot from working with him, especially how to frame my work in a practical context, and how to make sure that I always keep the big picture in mind. I will miss our weekly meetings where nothing was off-limits, and I could ask any question, however fundamental they might be.

I am grateful to my committee members Nicole Seiberlich and Raj Rao Nadakuditi for their help and guidance in shaping this dissertation. Nicole's grant introduced me to the world of myelin water exchange, and her feedback has been extremely helpful for my work. As for Raj, I have had the pleasure of being a Graduate Student Instructor for him (twice), and I thoroughly enjoyed working with, and learning from him. Both their perspectives have been very important to me, and I thank them for being on my committee.

I have been fortunate to be part of two wonderful groups as part of my PhD journey: the Lab of Jeff (LOJ) as well as the fMRI group. I'm grateful to the members of both these groups (both past and present), and have enjoyed interacting and collaborating with them. Many labmates have taken their journeys around the same time as me, and I have had a lot of fun in their company (Anish, Cameron, Caroline, Claire, Melissa, Steven, David, Dinank, Javier, Rodrigo, Amaya, Mingjie, Mariama, etc.).

I would like to acknowledge the role of my collaborators throughout my PhD. My first project during my PhD was in image reconstruction for large-scale convex problems, and I enjoyed collaborating with Greg Ongie and Laura Balzano. For all my work with myelin water imaging, I had fun working with Steven Whitaker and brainstorming with him, since he worked on a related problem to mine. Scott Swanson has been an invaluable collaborator of mine; he has been extremely helpful with anything to do with estimating exchange (designing urea water phantoms, understanding the mechanisms of exchange, acquiring NMR data for me, etc.). I have also enjoyed working with Amaya Murguia for the urea water exchange project, as well as Dinank Gupta and Rodrigo Lobos for my work on stack-of-spirals MRI.

They say it takes a village to obtain a PhD, and I am extremely fortunate that I have been surrounded by an amazing group of friends and family who have supported me throughout my journey. Board game nights were a wonderful part of my PhD, and I have loved every single one of them (Thanks Vaishnav, Divya, Mrudula, Rinav, Sharmi, Madan, Saylee, Aakanksha!). Vaishnav was my roommate for the better part of a decade, and he is one of the most authentic and genuine people that I know; I am extremely glad that I have known him for so long and that he's a part of my life. Divya, Madan, and I have been friends since our days together in our undergrad, and our journeys together from Bangalore to Ann Arbor over more than a decade have been very special to me. I have known Rinav since our days playing football together in Ann Arbor, and I'm glad that we have become very close over the past few years. Mrudula is one of my closest friends, and I have loved our co-working sessions from coffee shops, as well as spontaneous conversations about nothing in particular. Anup and I had our PhD journeys overlap with each other, and I'm grateful for having been able to share my thoughts and journey with him. I have loved spending

time with Dinank and Sushmitha, especially during the latter parts of my PhD. I have also enjoyed random chats about life, PhD, and everything in between, with friends and family in Ann Arbor and elsewhere (Sridhar, Srini, Minisha, Udhaya, Aniket, Anish, etc.).

Finally, I owe everything in my life to my family. Appa and Amma, thank you for your constant and unwavering support from day 1; this would not have been possible without you. I love you both from the bottom of my heart. My sister, Navya, is my best friend, and has been an amazing pillar of support during this whole journey and beyond. My brother-in-law Harish, my nephew Samarth, and niece Shriya have provided me great respite from the rigors of the PhD, and have given me lots of joy and happiness. I would also love to acknowledge my grandmom, who strongly wished that I pursue a PhD even before I had any inkling of wanting to obtain one. I hope she's proud wherever she is.

TABLE OF CONTENTS

| | |
|--|-------------|
| Acknowledgments | ii |
| List of Figures | viii |
| List of Tables | xi |
| Abstract | xii |
| Chapter | |
| 1 Introduction | 1 |
| 2 Background | 4 |
| 2.1 MR Physics | 4 |
| 2.1.1 Bloch equation | 6 |
| 2.1.2 Bloch-McConnell equation | 10 |
| 2.1.3 Steady-state MR signal | 13 |
| 2.2 Non-Cartesian MRI | 14 |
| 2.3 Two pool exchanging model for myelinated tissue | 15 |
| 2.4 Parameter Estimation via Regression with Kernels (PERK) | 17 |
| 3 Quantifying Myelin Water Exchange using Optimized bSSFP Sequences | 21 |
| 3.1 Introduction | 21 |
| 3.2 Related work | 23 |
| 3.3 Methods | 24 |
| 3.3.1 bSSFP signal model | 24 |
| 3.3.2 CRB-based scan design | 27 |
| 3.3.3 Estimation of myelin water exchange | 30 |
| 3.4 Experiments | 31 |
| 3.4.1 RF phase cycling factors (ϕ) in bSSFP sequences | 31 |
| 3.4.2 Importance of optimizing RF phase cycling factors ϕ | 32 |
| 3.4.3 Optimized scan design | 33 |
| 3.4.4 Numerical simulations | 35 |
| 3.4.5 In vivo results | 38 |
| 3.4.6 Effect of off-resonance on exchange estimation | 40 |
| 3.4.7 Joint estimation of B_0 and B_1+ maps | 41 |

| | | |
|----------|---|-----------|
| 3.5 | Discussion | 43 |
| 3.6 | Conclusion | 45 |
| 4 | Validating Exchange Maps in a Urea Water System using Phase-Cycled bSSFP | 46 |
| 4.1 | Introduction | 46 |
| 4.2 | Background | 47 |
| 4.2.1 | Exchange estimation using inversion recovery based NMR (IR-NMR) spectroscopy | 48 |
| 4.2.2 | Estimating inter-compartmental exchange: Related work | 53 |
| 4.2.3 | Parameter estimation using phase-cycled bSSFP MRI | 55 |
| 4.3 | Materials and Methods | 56 |
| 4.3.1 | Urea phantom studies | 56 |
| 4.3.2 | Ground truth using IR-NMR spectroscopy | 57 |
| 4.3.3 | Estimating exchange: phase-cycled bSSFP imaging | 59 |
| 4.4 | Results | 64 |
| 4.4.1 | Ground truth using spectroscopic data | 64 |
| 4.4.2 | Parameter estimation using phase-cycled bSSFP acquisitions | 66 |
| 4.5 | Discussion | 68 |
| 4.6 | Conclusions | 70 |
| 4.7 | Appendix | 71 |
| 4.7.1 | Eigen analysis of longitudinal magnetization dynamics | 71 |
| 5 | Stack-of-Spirals MRI Reconstruction with Off-Resonance Correction | 72 |
| 5.1 | Introduction | 72 |
| 5.2 | Background | 74 |
| 5.2.1 | Stack-of-spirals MRI | 74 |
| 5.2.2 | Off-resonance correction in MRI | 75 |
| 5.3 | Modified NUFFT for stack-of-spirals | 76 |
| 5.3.1 | Stack of identical spirals (no off-resonance correction) | 77 |
| 5.3.2 | Stack of identical spirals (with off-resonance correction) | 79 |
| 5.3.3 | Rotated stack of spirals (no off-resonance correction) | 81 |
| 5.3.4 | Rotated stack-of-spirals (with off-resonance correction) | 84 |
| 5.4 | Experiments | 90 |
| 5.5 | Discussion | 93 |
| 5.6 | Conclusions | 94 |
| 6 | Harmonization of Acquisition and Reconstruction Methods for Stack-of-Spirals MRI | 96 |
| 6.1 | Introduction | 96 |
| 6.2 | Background | 98 |
| 6.2.1 | Software/toolboxes for MR acquisition | 98 |
| 6.2.2 | Software/toolboxes for MR reconstruction | 98 |
| 6.3 | SOSP3d toolbox | 99 |
| 6.4 | Experiments | 100 |
| 6.4.1 | SOSP3d-based reconstruction | 100 |

| | | |
|----------|---|------------|
| 6.4.2 | SOSP3d-based acquisition and reconstruction | 101 |
| 6.5 | Discussion | 102 |
| 6.6 | Conclusions | 103 |
| 7 | Future Work | 104 |
| 7.1 | Myelin water imaging | 104 |
| 7.2 | Validation using aqueous urea system | 106 |
| 7.3 | Stack-of-spirals MRI | 108 |
| | Bibliography | 110 |

LIST OF FIGURES

| | | |
|-----|---|----|
| 2.1 | Cartesian and non-Cartesian sampling trajectories in MRI: Cartesian (left), radial (middle), and spiral (right). (Figure adapted from [171].) | 14 |
| 2.2 | Structure of myelinated tissue (Figure taken from [128]). | 15 |
| 2.3 | Two pool exchanging pool for myelinated tissue. | 16 |
| 3.1 | Exchange maps obtained using mcDESPOT [29] and MRF-X [48]. | 23 |
| 3.2 | (a) bSSFP pulse sequence. (b) Effect of varying RF phase cycling factors (ϕ) on simulated bSSFP off-resonance profiles. (Note: α = flip angle, T_R = repetition time, T_E = echo time) | 25 |
| 3.3 | Simulated bSSFP off-resonance profiles for different values of the myelin-specific frequency shift $\Delta\omega_f$, as a function of phase accrued during one T_R | 26 |
| 3.4 | Simulated bSSFP off-resonance profiles for various flip angles α (top row) and for different RF phase cycling factors ϕ (bottom row), as a function of phase accrued during one T_R . (Note: $\phi = \pi$ for the top row, and $\alpha = 30^\circ$ for the bottom row.) | 26 |
| 3.5 | Comparison of bSSFP signal magnitudes for alternating ($\phi = \pi$) and non-alternating ($\phi = 0$) RF pulses. A typical WM voxel is simulated here, with a mean residence lifetime ($\tau_{f \rightarrow s}$) of 250ms and a myelin-specific frequency shift ($\Delta\omega_f$) of 7.5Hz. | 31 |
| 3.6 | Numerically optimized bSSFP scan design for estimating myelin water exchange, consisting of 40 bSSFP scans with flip angles (α) and RF phase cycling factors (ϕ) as shown. (Note: y-axes are in degrees. TR/TE = 20ms/4ms.) | 34 |
| 3.7 | Simulated bSSFP signal magnitudes obtained using the optimized scan design for a typical WM voxel (on resonance, and 5Hz additional frequency offset for myelin water), under three different exchange regimes. | 35 |
| 3.8 | Estimated white matter (WM) exchange maps from simulated test data based on a modified Brainweb phantom [23]. (Note: Scatter plot shows exchange estimates of a randomly selected subset of WM voxels.) | 35 |
| 3.9 | Comparison of scatter plots (using BrainWeb simulations) for the optimized scan design and for three different random sets of acquisition parameters for phase-cycled bSSFP scans, using PERK estimation. (Note: The scatter plots correspond to a different random subset of voxels than shown in Fig. 3.8.) | 37 |

| | | |
|------|---|----|
| 3.10 | Comparison of exchange estimates obtained using PERK (left) and NLLS (right) estimation methods, for white matter (using BrainWeb simulations), vs. ground truth exchange values. The solid blue line indicates the linear least-squares fit, while the dashed black line is the identity line. | 38 |
| 3.11 | bSSFP images from the optimized scan design. (Note: Every other acquisition is shown out of 40 acquisitions.) | 39 |
| 3.12 | Estimated exchange map $\tau_{f \rightarrow s}$ without (left) and with (right) CSF masking. . . | 39 |
| 3.13 | Parameter maps obtained using the optimized scan design, followed by PERK based estimation. (Note: The scan design was optimized specifically for estimating exchange, and not the other parameters of the two pool model.) | 40 |
| 3.14 | Effect of (a) bulk off-resonance and (b) myelin-specific frequency shift on predicted coefficient of variation for estimating myelin water exchange. (Note: We study a typical white matter voxel here.) | 41 |
| 3.15 | Coefficients of variation for scan designs optimized for jointly estimating B_0 and $B_1 +$ information along with exchange. | 42 |
| 4.1 | Inversion Recovery-based pulse sequence used for obtaining ground truth exchange values in the aqueous urea phantom. Note: $T_S =$ time for selective inversion, $T_M =$ mixing time, $T_E =$ echo time. | 48 |
| 4.2 | Phantom containing urea water solution. | 57 |
| 4.3 | NMR spectrum with urea and water peaks for a particular mixing time T_M . Area under the peaks yields the signal values for urea and water proton pools at T_M | 58 |
| 4.4 | Handcrafted set of bSSFP acquisitions (S2), showing flip angles and phase cycling factors of 69 different acquisitions. | 60 |
| 4.5 | bSSFP off-resonance profiles simulated (using (S1)) for a single voxel in aqueous urea, showing (a) magnitude and (b) complex signal behavior for a range of values for urea fraction f_u and mean residence time of urea τ_u . (Note: $\tau_u = 1/k_u$) | 62 |
| 4.6 | Nonlinear least squares fit to inversion recovery NMR (IR-NMR) data, using a bi-exponential model. (Urea protons are selectively inverted here.) | 65 |
| 4.7 | Magnitude images for all 69 bSSFP acquisitions in set (S2). | 66 |
| 4.8 | Estimated parameter maps of urea fraction f_u and mean residence time of urea τ_u | 67 |
| 4.9 | Estimated parameter maps of equilibrium magnetization M_0 , phase M_ϕ , chemical shift Δf_{cs} and bulk off-resonance Δf | 67 |
| 5.1 | Spiral sampling trajectories. (a) 2D variable density spiral (VDS). (b) Stack-of-spirals acquisition consisting of 2D VDS shots. | 74 |
| 5.2 | Off-resonance effects: Cartesian vs. non-Cartesian acquisitions. Spiral acquisition contains larger amount of blurring (red arrow), but with less geometric distortion (blue arrows) compared to EPI. Figure taken from [53]. | 75 |

| | | |
|-----|--|-----|
| 5.3 | Results of simulating different implementations of the NUFFT-based forward model for a stack-of-spirals acquisition. (a) (<i>Empirical</i>) wall clock time, (b) (<i>Predicted</i>) floating point operations (FLOPs) computed using Table 5.1, (c) NRMSD vs. number of components in the low-rank approximation of the off-resonance terms (L). | 91 |
| 5.4 | Comparison of density-compensated adjoint reconstructions for baseline 3D NUFFT and recommended implementations for stack-of-spirals MRI, and ground truth. (Note: All three sets of images were scaled individually between 0 and 1.) | 95 |
| 6.1 | Stack-of-spirals trajectories used for our experiments. (a) 2D variable density spiral (VDS). (b) Stack-of-spirals acquisition consisting of 2D VDS shots. . . . | 100 |
| 6.2 | Stack-of-spirals brain data acquired using TOPPE and reconstructed (with off-resonance correction) using our SOSp3d toolbox. (Undersampling factors: $R_{xy} = 3$, $R_z = 1$.) | 101 |
| 6.3 | Reconstructed images of the ball phantom both acquired as well as reconstructed with our SOSp3d toolbox. (Acquisition: stack-of-spirals with undersampling factors of $R_{xy} = 3$ and $R_z = 1$.) | 102 |

LIST OF TABLES

| | | |
|-----|---|----|
| 3.1 | Distributions of tissue parameters used for optimizing the bSSFP acquisition parameters, as well as the tissue distributions used to train the PERK estimator. (Note: bSSFP signal profiles are periodic with a period of $1/T_R$, and hence, we select the off-resonance distribution $\Delta\omega$ to span 50Hz, corresponding to a T_R of 20ms.) | 30 |
| 3.2 | Ablation study showing the importance of optimizing RF phase cycling factors $\{\phi_i\}$ for exchange estimation. Note: $\{T_{Ri}\}$ denotes repetition times and $\{\alpha_i\}$ denotes flip angles of the 40 bSSFP acquisitions. | 33 |
| 3.3 | Comparison of predicted coefficients of variation and empirical RMSE values in white matter (using BrainWeb simulations), for the optimized scan design and for three different random sets of bSSFP scans. | 36 |
| 4.1 | Lower and upper box constraints for parameters to be estimated in the aqueous urea system (to optimize the nonlinear least squares cost function in (4.20)). (Note: We assume that the T_1 values of both the urea and water pools are the same, and denote it by the parameter T_1 .) | 64 |
| 4.2 | Parameter estimates for the urea water phantom, using complex-valued bSSFP data. (Note: Means and standard deviations were computed for the ROI corresponding to the entire cross-section of the urea water phantom.) | 67 |
| 5.1 | Comparison of computational costs for three different implementations of the forward model for stack-of-spirals acquisitions. | 89 |

ABSTRACT

Myelin content plays a vital role in healthy functioning of the brain. Myelin water imaging (MWI) is a quantitative MRI technique that aims to image the myelin content in the brain. Characterizing and quantifying the exchange dynamics in myelin (between myelin water and non-myelin water) could lead to better understanding of the role of myelin loss in several neurological disorders. This dissertation introduces a method to estimate myelin water exchange in white matter, using MRI experiment design. We optimized the acquisition parameters of a set of phase-cycled bSSFP acquisitions using a cost function based on the Cramér-Rao bound, and obtained low coefficients of variation for estimating myelin water exchange in white matter ($\sim 13.5\%$).

Validating the obtained myelin water exchange maps *in vivo* is a challenging problem. This dissertation also explores the design and use of an aqueous urea system to validate our proposed bSSFP acquisitions for estimating exchange. To potentially translate our proposed acquisitions to clinically feasible settings, undersampling is an important consideration to bring down acquisition time. This dissertation studies an undersampling scheme based on stack-of-spirals MRI. We demonstrate that the 3D NUFFT-based forward model for stack-of-spirals MRI can be simplified and sped up by exploiting the Cartesian structure in the through-plane direction, even in the presence of off-resonance effects. We show that our recommended efficient implementation is at least $\sim 22\%$ faster than other existing 3D NUFFT-based implementations of the stack-of-spirals forward model, when run on a 20-core Intel® Xeon® processor.

CHAPTER 1

Introduction

Magnetic Resonance Imaging (MRI) provides a non-invasive way to study soft-tissue contrast in the body. Conventional MR images are generally qualitative in nature; various contrast-weighted images can be obtained by adjusting MR imaging acquisition parameters (e.g., T_1 -weighted and T_2 -weighted images). In contrast, quantitative MR imaging focuses on obtaining accurate quantitative maps of tissue parameters, such as T_1 , T_2 , and proton density, that potentially reflect the changes at the cellular level better.

Many classical quantitative approaches are based on the premise that the MR signal from a voxel arises from a single type of tissue. However, this assumption may not be appropriate when modeling the complex spin-dynamics of microstructural components in living tissues. A more suitable model would be a multi-compartment relaxation model, where the MR signal is assumed to originate from different tissue compartments. Quantifying the exchange dynamics between these various compartments has been suggested as a potential biomarker for several clinical applications [76, 21, 188, 4].

An example of a multi-compartment tissue model can be seen in myelin water imaging, where myelin content is studied in brain tissue [78]. E.g., for the two-compartment tissue model of myelin used in this work, the MR signal is assumed to arise from two water pools: a fast-relaxing tissue component corresponding to water trapped between myelin bilayers (called *myelin water*), and a slow-relaxing component consisting of intra and extra cellular water (called *non-myelin water*). The fraction of signal arising from the myelin water is defined as the myelin water fraction (MWF) and is taken to be a proxy for the myelin content in a voxel of brain tissue [2, 164]. Chapter 2 provides further background on the two pool exchanging model that we consider in this work, and how the Bloch-McConnell equations are used to model the evolution of spins in such an exchanging system [98].

The relationship between MWF and myelin content has been thought to depend on not just the relative water proton densities of the two compartments, but also the water exchange dynamics between them [187]. It has been thought that the presence of inter-compartmental water exchange could lead to conventional MWF methods underestimating

the true myelin content in the tissue [34, 51]. The ability to quantify aspects of the inter-compartmental exchange dynamics, including exchange rates and relative fractions, may provide researchers with a tool to study various disease processes that depend on water transport between compartments. For example, variations in water exchange may be associated with brain edema and demyelination processes [84].

Characterizing these tissue compartments and their interactions *in vivo* has been a long-standing research problem [60, 157]. However, this is a challenging problem since there are few baseline methods that can quantitatively estimate these water exchange dynamics *in vivo*. Another major issue is the lack of reproducibility of mean residence times (i.e., the inverse of the exchange rates); some studies have observed residence times of 160-310 ms *ex vivo* in myelin [145, 11], while other studies have measured values of 280-780 ms *in vivo* [70]. Chapter 3 focuses on obtaining precise estimates of myelin water exchange in white matter by optimizing the acquisition parameters of a set of MR sequences. In particular, we optimize the flip angles and RF phase cycling factors of a set of balanced Steady-State Free Precession (bSSFP) acquisitions, to minimize the variance of unbiased exchange estimates. This work was published as a conference abstract [110].

Validating the exchange maps *in vivo* obtained in Chapter 3 is challenging, since there is no established method to obtain ground truth exchange maps *in vivo* (cf. inversion recovery methods for T_1 estimation, multi-echo spin echo acquisitions for estimating myelin water fraction). Chapter 4 explores the design and use of an aqueous urea system to validate our proposed acquisitions for estimating exchange. We obtain ground truth values for urea water exchange using inversion recovery based NMR (IR-NMR) spectroscopic experiments [168]. We perform imaging experiments using phase-cycled bSSFP acquisitions and demonstrate their feasibility for estimating urea water exchange, by validating them against the ground truth exchange values.

The Cramér-Rao bound analysis in Chapter 3 assumes the acquisition of fully sampled data for each of the optimized bSSFP acquisitions. However, to make the acquisitions feasible for clinical settings, we require faster scan times, and undersampling is a popular approach to accelerate MRI scans. Chapter 5 studies an undersampling scheme based on stack-of-spirals MRI, which is a popular non-Cartesian trajectory for 3D imaging. We demonstrate that the forward model for stack-of-spirals MRI (based on the 3D non-uniform FFT (NUFFT)) can be efficiently implemented by exploiting the Cartesian structure in the through-plane direction, even in the presence of off-resonance effects.

There is a concerted push to increase harmonization and reproducibility of MRI studies across different experimental settings [105]. This is especially important for quantitative MRI, where researchers would like to study the variability of parameter estimates across

a range of experimental conditions. A critical requirement for harmonization is the development of open-source software and toolboxes that facilitate widespread adoption of MRI protocols. Chapter 6 introduces an open-source toolbox (called SOS3d) for acquiring and reconstructing stack-of-spirals MRI data. We demonstrate the feasibility of this toolbox for a few preliminary studies, e.g., for structural MRI scans.

The rest of this thesis (including our contributions in Chapters 3-6) is organized as follows:

- Chapter 2 contains relevant background information.
- Chapter 3 introduces a novel method for designing a set of bSSFP acquisitions to estimate myelin water exchange in white matter. In particular, optimizing the phase cycling factors of a set of bSSFP acquisitions using the Cramér-Rao bound (during scan design) is a novel contribution of this chapter.
- Chapter 4 introduces an aqueous urea system designed to validate our proposed techniques for estimating exchange.
- Chapter 5 develops an efficient implementation for the NUFFT based forward model for stack-of-spirals MRI, in the presence of off-resonance effects.
- (Preliminary work) Chapter 6 introduces an open-source toolbox for acquisition and reconstruction methods for stack-of-spirals MRI.
- Chapter 7 contains possible ideas for future work.

CHAPTER 2

Background

This chapter provides background information that is relevant to upcoming chapters. Subsequent chapters provide additional topic-specific background content. Section 2.1 provides a brief introduction to MRI principles that are pertinent to this thesis. Section 2.2 gives a brief description of non-Cartesian trajectories used in MRI. Section 2.3 introduces a two pool exchanging model for myelinated tissue, and describes the model parameters in detail. This thesis uses a parameter estimation approach called PERK (Parameter Estimation via Regression with Kernels), and Section 2.4 provides a brief description of PERK.

2.1 MR Physics

This section gives a very brief background about some fundamentals of MR physics that are relevant to this work. We refer readers to the following resources for more detailed descriptions and explanations of MR imaging [118, 85].

Magnetic Resonance Imaging (MRI) is an imaging modality based on the magnetization properties of hydrogen protons (^1H) in the body. Other nuclei satisfy the magnetic resonance properties required for MRI, but the hydrogen proton is the most abundant in the human body and is most convenient to image. Each hydrogen proton possesses a magnetic property called nuclear spin, which causes it to behave like a tiny rotating magnet. These spins¹ are typically oriented randomly in all directions. Hence, the vector sum of all these spins is zero; in other words, the net magnetization is zero in the absence of an external magnetic field.

In MRI, a static magnetic field, called the main magnetic field B_0 , is externally applied, causing the spins to align themselves either parallel to or anti-parallel to B_0 . Roughly half the spins can be found in the parallel state and half are in the anti-parallel state. To be more precise, under typical conditions, there is a tiny excess of protons that lie in the parallel

¹Throughout this work, we refer to protons or nuclei or nuclear spins or spins interchangeably.

state versus the anti-parallel state (roughly 7 spins out of 10^6 [118]). This tiny excess of protons results in a net magnetization vector $\mathbf{M} = [M_x, M_y, M_z]^T \in \mathbb{R}^3$ that is aligned parallel to B_0 , and this magnetization \mathbf{M} is at the heart of MR imaging.

Conventionally, the main field B_0 is assumed to be applied along the z -direction (referred to as the *longitudinal* direction), and the plane perpendicular to it (x - y plane) is referred to as the *transverse* plane. While only B_0 is applied, the net magnetization vector $\mathbf{M} = [0, 0, M_0]^T$ is aligned with the z -direction, where M_0 is the equilibrium magnetization due to the B_0 field. If the spins are not in equilibrium (i.e., when they are *excited*), the nuclear spins display a phenomenon called resonance, whereby they *precess* at a well-defined frequency around the B_0 field. The concept of precession is similar to a spinning top in a gravitational field, where the top is said to precess about the direction of the gravitational field.

The frequency of precession of the nuclear spins is governed by one of the most influential equations in MRI, the Larmor equation, given by

$$\omega_0 = \gamma B_0, \tag{2.1}$$

where ω_0 is the precessional frequency of a spin system experiencing an external static field B_0 . This is also referred to as the resonant frequency or Larmor frequency of the spin. γ is a known constant called the gyromagnetic ratio that is specific to each nuclear species. For instance, the gyromagnetic ratio for ^1H protons [118] is $\gamma/2\pi = 42.575 \text{ MHz/T}$.

Under the externally applied main field B_0 , the net magnetization is aligned with B_0 . If we momentarily *excite* or perturb the magnetization at its resonant frequency, the net magnetization \mathbf{M} is disturbed from its equilibrium state and tries to return to equilibrium along B_0 . In MR imaging, we use an RF pulse (oscillating at the resonant frequency) in the transverse plane to create a rotating magnetic field B_1 . The B_1 field is turned on for a very short period of time (on the order of a few μs or ms) and this excites some of the nuclear spins away from equilibrium. The net magnetization \mathbf{M} eventually returns to its equilibrium state along z (assuming B_0 is along z), and this is accompanied by three processes: free precession, longitudinal relaxation, and transverse relaxation. The longitudinal and transverse relaxations are governed by time constants T_1 and T_2 respectively that determine how long the magnetization \mathbf{M} takes to return to its equilibrium.

The signal of interest in MR imaging is the net magnetization in the transverse plane, i.e., $M_{xy} = M_x + iM_y$, where $\mathbf{M} = [M_x, M_y, M_z]^T \in \mathbb{R}^3$ and $i = \sqrt{-1}$ is the imaginary unit of complex numbers. We typically excite the spins using the B_1 field, and while the net magnetization is returning to equilibrium, we acquire the MR signal using RF receive coils. Depending on how long we wait after excitation (roughly referred to as the

echo time), we obtain signals that highlight different aspects/contrasts of the sample being imaged. Different biological tissues such as fat, muscle, water etc. have different T_1 and T_2 characteristics, and by manipulating the sequence timing parameters such as *echo time* (T_E) and *repetition time* (T_R), we can highlight the contrast between specific tissues.

The RF receive coils acquire the MR signal over the entire imaging volume, and it is impossible to distinguish signals from different spatial locations (post excitation), without the use of additional magnetic fields. In MR imaging, spatial localization is achieved by using linearly varying magnetic fields called *gradient* fields (G) in addition to the B_0 field [118]. The gradient fields cause variations in the effective magnetic fields seen by different spins, and this leads to a variation in their precessional frequencies (Larmor frequencies), thereby helping localize the signal spatially. We do not cover spatial localization in detail here; interested readers are referred to other resources for more details [85, 118].

A typical MR pulse sequence consists of a sequence of RF excitations², with periods of free precession and relaxation in between successive RF pulses. The MR signal is acquired during the periods of free precession and relaxation. The time period between successive RF pulses is called the *repetition time* (T_R). The time between each RF pulse and when the signal is acquired is roughly referred to as the *echo time* (T_E). We repeat the process of excitation and data acquisition until we have sufficient data to reconstruct the MR image.

2.1.1 Bloch equation

The time evolution of the net magnetization \mathbf{M} due to various processes such as excitation, free precession and relaxation, is an important part of MRI. The Bloch equation is a phenomenological equation that describes the evolution of \mathbf{M} in the sample being imaged.

Free precession

The precession of magnetization $\mathbf{M} = [M_x, M_y, M_z]^T \in \mathbb{R}^3$ around a magnetic field $\mathbf{B} = [B_x, B_y, B_z]^T \in \mathbb{R}^3$ is described by the following differential equation [13]:

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma \mathbf{B}, \quad (2.2)$$

where \times denotes the cross-product of two vectors. After expanding out the cross product, this can be written in matrix form as

²We refer to RF excitation and B_1 field interchangeably throughout this work.

$$\frac{d\mathbf{M}}{dt} = \begin{bmatrix} 0 & \gamma B_z & -\gamma B_y \\ -\gamma B_z & 0 & \gamma B_x \\ \gamma B_y & -\gamma B_x & 0 \end{bmatrix} \mathbf{M}. \quad (2.3)$$

The differential equation in (2.3) can be solved under different conditions, to describe various MR processes. For instance, under the influence of only the main magnetic field B_0 , i.e., $\mathbf{B} = [0, 0, B_0]^T \in \mathbb{R}^3$, the solution is given by [118]

$$\mathbf{M}(t) = \mathbf{R}_z(\omega_0 t) \mathbf{M}(0), \quad \mathbf{R}_z(\omega_0 t) = \begin{bmatrix} \cos \omega_0 t & \sin \omega_0 t & 0 \\ -\sin \omega_0 t & \cos \omega_0 t & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.4)$$

which describes the precession of \mathbf{M} around the z-axis. $\omega_0 = \gamma B_0$ is the Larmor frequency, while $\mathbf{M}(0)$ represents the initial conditions for \mathbf{M} . \mathbf{R}_z is the rotation matrix around the z-axis.

Longitudinal relaxation

As described before, when we excite the spins with an RF excitation pulse, they eventually return to equilibrium along the main magnetic field B_0 . One of the processes that accompanies the return to equilibrium is longitudinal relaxation. The longitudinal component M_z recovers to equilibrium in the following way [118]:

$$\frac{dM_z}{dt} = -\frac{M_z - M_0}{T_1}, \quad (2.5)$$

where M_0 is the equilibrium magnetization due to the B_0 field. Longitudinal relaxation is also called spin-lattice relaxation, since there is transfer of energy between nuclei and the surrounding lattice. T_1 is called the spin-lattice time constant, and characterizes the recovery of the longitudinal component of the net magnetization to equilibrium. The solution of the differential equation in (2.5) is

$$M_z(t) = M_z(0) e^{-t/T_1} + M_0 (1 - e^{-t/T_1}), \quad (2.6)$$

where $M_z(0)$ is the initial condition at time $t = 0$.

Transverse relaxation

Transverse relaxation describes the decay of the transverse components of \mathbf{M} , i.e., M_x and M_y , after excitation. This is described by [118]

$$\frac{dM_{xy}}{dt} = -\frac{M_{xy}}{T_2}, \quad (2.7)$$

where $M_{xy} = M_x + iM_y$ is the transverse component. T_2 is the time constant that characterizes the decay of the transverse magnetization after excitation. The solution of the above equation is

$$M_{xy}(t) = M_{xy}(0) e^{-t/T_2}, \quad (2.8)$$

where $M_{xy}(0)$ represents the initial conditions at time $t = 0$.

The Bloch equation [13, 118] combines the above equations into a single set of equations that describes the evolution of the net magnetization:

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma \mathbf{B} - \frac{M_x \mathbf{i} + M_y \mathbf{j}}{T_2} - \frac{(M_z - M_0) \mathbf{k}}{T_1}, \quad (2.9)$$

where $\mathbf{i}, \mathbf{j}, \mathbf{k}$ are the unit vectors along x, y, z directions respectively. In matrix form, this can be written as

$$\frac{d\mathbf{M}}{dt} = \begin{bmatrix} -\frac{1}{T_2} & \gamma B_z & -\gamma B_y \\ -\gamma B_z & -\frac{1}{T_2} & \gamma B_x \\ \gamma B_y & -\gamma B_x & -\frac{1}{T_1} \end{bmatrix} \mathbf{M} + \begin{bmatrix} 0 \\ 0 \\ \frac{M_0}{T_1} \end{bmatrix}. \quad (2.10)$$

The solutions of the Bloch equation under various conditions provide closed-form expressions for the magnetization under typical MR processes such as excitation, free precession and relaxation. Though we have assumed a perfectly uniform B_0 field throughout the imaging volume so far, in reality, there are variations in the resonant frequency due to factors such as main field inhomogeneities, chemical shift and susceptibility variations. These variations are called *off-resonance* and are modeled within the B_z term in (2.10).

Excitation

Using the Bloch equation, we have described the evolution of net magnetization \mathbf{M} during free precession and relaxation, assuming that the spins were already perturbed or *excited* in some way. Here, we dive deeper into how \mathbf{M} behaves during excitation. Modern MR

imaging systems typically generate a circularly polarized B_1 field³, that takes the following form [85]:

$$\mathbf{B}_1(t) = B_1^e(t) [\cos(\omega_{\text{rf}}t + \theta)\mathbf{i} - \sin(\omega_{\text{rf}}t + \theta)\mathbf{j}], \quad (2.11)$$

where $B_1^e(t)$ is the envelope of the RF pulse, ω_{rf} is the RF carrier frequency and θ is the initial phase of the pulse. In other words, the B_1 field has an x-component of $B_1^e(t) \cos(\omega_{\text{rf}}t + \theta)$ and a y-component of $-B_1^e(t) \sin(\omega_{\text{rf}}t + \theta)$ in the transverse plane. To excite a spin with Larmor frequency ω_0 , we typically set ω_{rf} to ω_0 (resonance condition). $B_1^e(t)$ plays an important role in determining how much the magnetization \mathbf{M} is tipped away from the B_0 field due to excitation. If the RF pulse is on for a duration τ , and if no gradients are played out during RF transmission, then the *tip-down* or *flip angle* is defined as

$$\alpha = \gamma \int_0^\tau B_1^e(t) dt, \quad (2.12)$$

where γ is the gyromagnetic ratio of the nuclear spin. A longer RF pulse duration or higher amplitude of the B_1 envelope function leads to a higher flip angle and typically results in the net magnetization being tipped away (excited) from equilibrium to a larger extent.

In the conventional coordinate system (called the *laboratory* frame), it is cumbersome to mathematically describe the evolution of \mathbf{M} during excitation. Hence, we use a transformed coordinate system, called the *rotating* frame, where the transverse plane is rotating at a chosen angular frequency. By choosing the Larmor frequency ω_0 for rotation, we demodulate the free precession due to the B_0 field, and this simplifies the solution for excitation. Solving the Bloch equation for the rotating frame (similar to (2.9)) for excitation results in [85]

$$\mathbf{M}_{\text{rot}}^+ = \mathbf{R}_{z'}(\theta)\mathbf{R}_{x'}(\alpha)\mathbf{R}_{z'}(-\theta)\mathbf{M}_{\text{rot}}^-, \quad (2.13)$$

where $\mathbf{M}_{\text{rot}}^-$ and $\mathbf{M}_{\text{rot}}^+$ refer to the net magnetization *pre* and *post* excitation respectively (in the rotating frame). Since the RF pulse is assumed to be of a very short duration (i.e., $\ll T_2$), we ignore the relaxation terms in the Bloch equation here. x', y', z' are the coordinates in the rotating frame⁴. $\mathbf{R}_{z'}(\theta)$ and $\mathbf{R}_{x'}(\alpha)$ are rotation matrices around the z' and x' axes respectively, defined as

³The same analysis holds for a linearly polarized B_1 field, which can be shown to decompose into two circularly polarized fields [118].

⁴The longitudinal components are the same in both the laboratory and rotating frames, i.e., $z' = z$.

$$\mathbf{R}_{z'}(\theta) \triangleq \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{R}_{x'}(\alpha) \triangleq \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & \sin \alpha \\ 0 & -\sin \alpha & \cos \alpha \end{bmatrix}. \quad (2.14)$$

2.1.2 Bloch-McConnell equation

So far, we have assumed that the sample being imaged is homogeneous within each voxel and a single-compartment tissue model is sufficient to model the MR signal (using the Bloch equation). However, in practice, it is more realistic to model the net magnetization in each voxel as arising from a multi-compartment tissue model, where the net magnetization in the voxel depends on the net magnetization of the individual tissue compartments. The spins in a specific compartment have their own properties (T_1 , T_2 etc.) depending on their chemical environments. Also, we need to account for exchange of magnetization between compartments (due to transfer of water for instance) in our analysis. The Bloch-McConnell equation provides a way to describe the evolution of magnetization in a multi-pool⁵ model with exchange [98].

Consider a two compartment tissue model with compartments denoted by a and b . In other words, the net magnetization in each voxel arises from spins in either compartment a or b . Let T_{1a}, T_{2a}, T_{1b} and T_{2b} be the relaxation time constants for compartments a and b respectively. Let f_a and $f_b = (1 - f_a)$ represent the fraction of spins in each voxel that correspond to compartments a and b respectively. The off-resonance experienced by nuclei in compartments a and b are denoted by $\Delta\omega_a$ and $\Delta\omega_b$ respectively. The exchange from compartment a to b is characterized by the mean residence time $\tau_{a \rightarrow b}$ (inverse of the exchange rate), and $\tau_{b \rightarrow a}$ denotes the exchange in the opposite direction from b to a . The Bloch-McConnell equation for the net magnetization $\mathbf{M} = [M_{a,x}, M_{a,y}, M_{a,z}, M_{b,x}, M_{b,y}, M_{b,z}]^T \in \mathbb{R}^6$ takes the form (in the *rotating frame*) [98, 179]:

$$\frac{d\mathbf{M}}{dt} = \begin{bmatrix} \mathbf{L}_a - \mathbf{K}_a & \mathbf{K}_b \\ \mathbf{K}_a & \mathbf{L}_b - \mathbf{K}_b \end{bmatrix} \mathbf{M} + \mathbf{c}, \quad (2.15)$$

where

$$\mathbf{c} = \left[0, 0, \frac{f_a M_0}{T_{1a}}, 0, 0, \frac{f_b M_0}{T_{1b}} \right]^T, \quad (2.16)$$

⁵We use the terms tissue compartments and pools interchangeably.

and M_0 is the equilibrium magnetization. L_a captures the compartment-specific characteristics of a , while K_a characterizes the exchange process from a to b . These are defined as:

$$L_a \triangleq \begin{bmatrix} -\frac{1}{T_{2a}} & \Delta\omega_a & -\gamma B_y \\ -\Delta\omega_a & -\frac{1}{T_{2a}} & \gamma B_x \\ \gamma B_y & -\gamma B_x & -\frac{1}{T_{1a}} \end{bmatrix}, \quad K_a \triangleq \begin{bmatrix} \frac{1}{\tau_{a \rightarrow b}} & 0 & 0 \\ 0 & \frac{1}{\tau_{a \rightarrow b}} & 0 \\ 0 & 0 & \frac{1}{\tau_{a \rightarrow b}} \end{bmatrix},$$

where B_x and B_y are the x and y components of the magnetic field in the transverse plane. L_b and K_b are defined in a similar manner as L_a and K_a . In the absence of any exchange, K_a and K_b contain all zeros (i.e., mean residence times are infinite), and (2.15) reduces to a set of *decoupled* differential Bloch equations specific to each compartment. The magnetization in each compartment can then be separately analyzed using the Bloch equation as described before.

Excitation

Assuming that the RF excitation pulse is of very short duration, we neglect the effects of free precession, relaxation and exchange during excitation. Consider a circularly polarized B_1 field with components $B_{1,x}$ and $B_{1,y}$ along the x and y directions respectively. Let α be the flip angle and θ be the initial phase of the RF pulse. Then, the Bloch-McConnell equation (Eq. (2.15)) reduces to

$$\frac{d\mathbf{M}}{dt} = \begin{bmatrix} 0 & 0 & -\gamma B_{1,y} & 0 & 0 & 0 \\ 0 & 0 & \gamma B_{1,x} & 0 & 0 & 0 \\ \gamma B_{1,y} & -\gamma B_{1,x} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\gamma B_{1,y} \\ 0 & 0 & 0 & 0 & 0 & \gamma B_{1,x} \\ 0 & 0 & 0 & \gamma B_{1,y} & -\gamma B_{1,x} & 0 \end{bmatrix} \mathbf{M}. \quad (2.17)$$

These decoupled differential equations can be solved separately for each compartment, and results in a solution that is similar to the Bloch equation solution in (2.13):

$$M^+ = \begin{bmatrix} \mathbf{R}_{z'}(\theta) \mathbf{R}_{x'}(\alpha) \mathbf{R}_{z'}(-\theta) & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{R}_{z'}(\theta) \mathbf{R}_{x'}(\alpha) \mathbf{R}_{z'}(-\theta) \end{bmatrix} M^-, \quad (2.18)$$

where $M^- \in \mathbb{R}^6$ and $M^+ \in \mathbb{R}^6$ represent the net magnetization *pre* and *post* excitation respectively. $R_{x'}$ and $R_{z'}$ are rotation matrices around the x' and z' directions in the rotating frame.

Free precession and Relaxation

After excitation, in the absence of the B_1 field, the Bloch-McConnell equation becomes

$$\frac{dM}{dt} = \mathbf{A}M + \mathbf{c}, \quad (2.19)$$

where

$$\mathbf{A} \triangleq \begin{bmatrix} \tilde{\mathbf{L}}_a - \mathbf{K}_a & \mathbf{K}_b \\ \mathbf{K}_a & \tilde{\mathbf{L}}_b - \mathbf{K}_b \end{bmatrix}, \quad \tilde{\mathbf{L}}_a \triangleq \begin{bmatrix} -\frac{1}{T_{2a}} & \Delta\omega_a & 0 \\ -\Delta\omega_a & -\frac{1}{T_{2a}} & 0 \\ 0 & 0 & -\frac{1}{T_{1a}} \end{bmatrix}. \quad (2.20)$$

$\tilde{\mathbf{L}}_b$ is defined similarly to $\tilde{\mathbf{L}}_a$, while \mathbf{K}_a , \mathbf{K}_b and \mathbf{c} are as defined before. The solution to this set of differential equations involves a matrix exponential, and is given by (assuming \mathbf{A} is invertible)

$$M(t) = e^{t\mathbf{A}}M(0) + (e^{t\mathbf{A}} - \mathbf{I}_6) \mathbf{A}^{-1}\mathbf{c}, \quad (2.21)$$

where $M(0)$ is the initial condition for the net magnetization and \mathbf{I}_6 is the 6×6 identity matrix. $e^{\mathbf{X}} = \sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{X}^k$ is the matrix exponential function. Let us define the compartmental equilibrium magnetization as $\mathbf{m}_0 = [0, 0, f_a M_0, 0, 0, f_b M_0]^T \in \mathbb{R}^6$. Expanding \mathbf{A} (using (2.20)), we obtain

$$\mathbf{A}\mathbf{m}_0 = \begin{bmatrix} 0 \\ 0 \\ \left(-\frac{1}{T_{1a}} - \frac{1}{\tau_{a \rightarrow b}}\right) f_a M_0 + \left(\frac{1}{\tau_{b \rightarrow a}}\right) f_b M_0 \\ 0 \\ 0 \\ \left(\frac{1}{\tau_{a \rightarrow b}}\right) f_a M_0 + \left(-\frac{1}{T_{1b}} - \frac{1}{\tau_{b \rightarrow a}}\right) f_b M_0 \end{bmatrix}. \quad (2.22)$$

Under the condition of chemical equilibrium⁶, which is assumed throughout this work, (2.22) further simplifies to $\mathbf{A}\mathbf{m}_0 = -\mathbf{c}$ (from (2.16)) [112]. Then, (2.21) simplifies as

$$\mathbf{M}(t) = e^{t\mathbf{A}}\mathbf{M}(0) + (\mathbf{I}_6 - e^{t\mathbf{A}})\mathbf{m}_0. \quad (2.23)$$

2.1.3 Steady-state MR signal

The Bloch equation in (2.9) and Bloch-McConnell equation in (2.15) provide ways to study the evolution of the net magnetization vector under various conditions. An MRI scan typically consists of applying different magnetic fields (B_0 , B_1 and gradient fields) and acquiring the MR signal at specific time points. Since the acquired MR signal is the net magnetization vector (in the *transverse* plane), the solutions to the Bloch and Bloch-McConnell equations (under various conditions such as free precession, relaxation, excitation etc.) play an important role in determining the acquired signal in MRI.

As described earlier, typical MRI acquisitions consist of a sequence of RF excitation pulses (separated by repetition time T_R), with periods of free precession and relaxation in between the RF pulses. If we set T_R sufficiently large (say $T_R \gg T_1$), then there is enough time between RF excitations for the transverse magnetization to decay to zero and for the longitudinal magnetization to completely recover to the equilibrium magnetization M_0 . However, if we use a rapid train of RF pulses, i.e., T_R is sufficiently small, then the net magnetization does not have enough time to fully recover and eventually reaches a *steady-state* that differs from M_0 . Such MRI sequences with rapidly applied RF excitations (typically with $T_R \ll T_2$) are called steady-state sequences [103]. In steady-state, the evolution of the net magnetization vector is identical in each T_R . We can compute the acquired MR signal in steady-state by modeling the magnetization in successive T_R periods as an affine system of equations, and then solving for the steady-state magnetization.

For instance, consider the balanced steady-state free precession (bSSFP) sequence that is a common steady-state sequence used in MRI. The bSSFP sequence consists of a rapid train of RF excitation pulses, with *balanced* gradient fields applied during each T_R . The term *balanced* refers to the fact that the net area of the applied gradients is designed to be zero in each T_R . Chapter 3 provides a more detailed look at the bSSFP signal model. If the phase of the RF pulses alternates between π and $-\pi$ (between successive excitations), the magnitude of the steady-state MR signal immediately after excitation, for an on-resonance spin, is given by [139]

⁶Two pools in chemical equilibrium satisfy the relation $\frac{f_a}{f_b} = \frac{\tau_{a \rightarrow b}}{\tau_{b \rightarrow a}}$.

$$s_{\text{bSSFP}} = M_0 \frac{(1 - E_1) \sin \alpha}{1 - (E_1 - E_2) \cos \alpha - E_1 E_2}, \quad (2.24)$$

where $E_1 = e^{-T_R/T_1}$, $E_2 = e^{-T_R/T_2}$, α is the flip angle of the RF excitation pulses and M_0 is the equilibrium magnetization. (2.24) assumes a single compartment tissue model. If we consider a two pool model with exchange, then the bSSFP steady-state signal does not have a simple analytical expression and has to be determined using the Bloch-McConnell equation. For a two pool model *without* exchange, each tissue compartment can be analyzed separately and the total steady-state signal in each voxel is a weighted sum of the signals from each compartment.

2.2 Non-Cartesian MRI

The simplest form of MR imaging typically acquires data in a Cartesian or a recti-linear trajectory. Fig. 2.1 (left) shows an example of a Cartesian sampling trajectory, where data is acquired line-by-line, and can be reconstructed in a straightforward manner using the fast Fourier transform (FFT). However, a recti-linear Cartesian acquisition is not the only possible sampling trajectory in MRI; we can acquire k-space data in a non-Cartesian manner as well. There are a lot of non-Cartesian trajectories that have been explored in MRI [171], e.g., radial [79, 43, 38], spiral [1, 100], rosette [121], PROPELLER [126], and stochastic trajectories [138]. Each of these non-Cartesian trajectories has its own unique set of benefits/properties, and are used in a wide range of applications in MRI. Fig. 2.1 shows examples of two non-Cartesian sampling trajectories: radial and spiral trajectories.

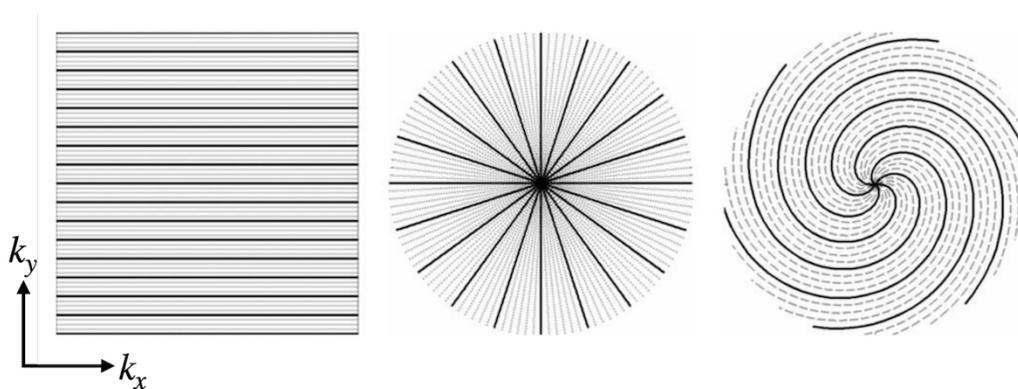


Figure 2.1: Cartesian and non-Cartesian sampling trajectories in MRI: Cartesian (left), radial (middle), and spiral (right). (Figure adapted from [171].)

Faster acquisition times are generally required for translating MR sequences to clinical applications, and non-Cartesian trajectories provide a way to reduce scan time through undersampled acquisitions. One of the most important benefits of non-Cartesian sampling is the potential for efficient coverage of k-space. Compared to Cartesian acquisitions, non-Cartesian trajectories typically make efficient use of the gradient hardware, and therefore lead to more rapid k-space coverage. Non-Cartesian trajectories are also potentially associated with other benefits such as fewer coherent undersampling artifacts [124], more robustness to motion [43, 87], potential to be used for self-navigation [89], etc. Ref. [171] provides an excellent review of non-Cartesian imaging, and is a good starting point for understanding more about the nuances of non-Cartesian MRI.

While non-Cartesian acquisitions have many benefits, it is also more complex to reconstruct non-Cartesian k-space data. Since the k-space data does not lie on a uniform Cartesian grid, we typically use reconstruction methods based on the non-uniform fast Fourier transform (NUFFT) [40, 176]. Also, since the acquisition k-space is sampled non-uniformly in non-Cartesian MRI, it is necessary to account for the non-uniform sampling density using a process called density compensation [127].

2.3 Two pool exchanging model for myelinated tissue

As described before, in MR imaging, the MR signal typically arises from hydrogen protons in the body and their interactions with various magnetic fields (a static main magnetic

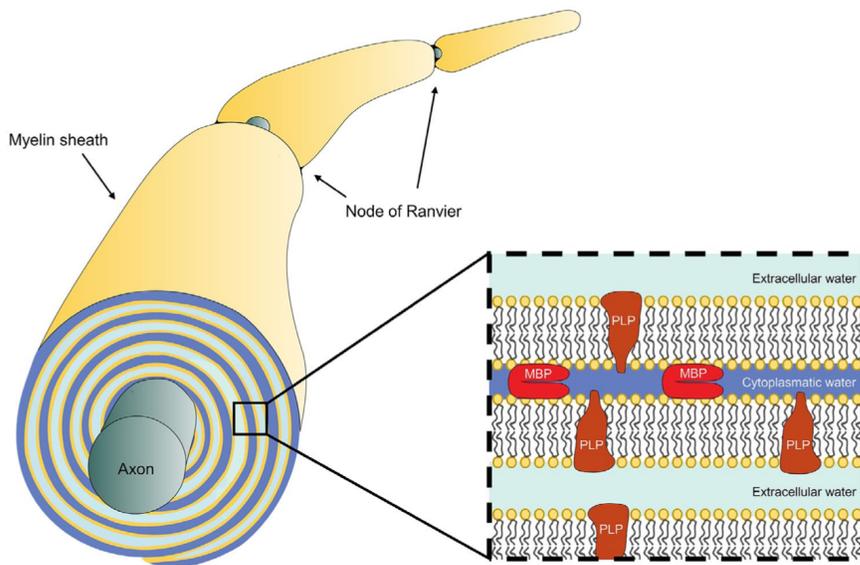


Figure 2.2: Structure of myelinated tissue (Figure taken from [128]).

field B_0 , a radiofrequency field B_1 and gradient fields G) [118]. However, individual water molecules can experience very different microscopic environments depending on their physical location, and this forms the basis of myelin water imaging [78]. Fig. 2.2 illustrates the myelination around axonal fibers. The water molecules trapped within this myelin sheath (called *myelin water*) experience a much different microscopic environment than other water molecules (denoted as *non-myelin water* or intra/extracellular water). There are also exchange dynamics at play here, with exchange of water molecules between these two environments, and our model must account for this process.

For the rest of this work, we model the myelinated tissue as a two pool exchanging model, as shown in Fig. 2.3. In other words, we assume that the MR signal in each voxel arises from two different compartments: a *myelin water* compartment representing the water trapped in the myelin bilayers, and a *non-myelin water* compartment that captures all other water, i.e., intra and extra-axonal water. The myelin water compartment has relaxation time constants denoted by T_{1f} and T_{2f} ; the letter ‘f’ signifies that the myelin water compartment is a fast-relaxing compartment. Similarly, the non-myelin water pool has relaxation time constants of T_{1s} and T_{2s} ; the ‘s’ signifies that the non-myelin water pool is a slow-relaxing compartment. In other words, the water molecules trapped in the myelin sheath undergo faster relaxation (i.e., have shorter time constants) than the non-myelin water.

The myelin water fraction is denoted by f_f ; this represents the fraction of the MR signal in a voxel that arises from protons in the myelin water pool. f_s represents the fraction of signal from the non-myelin water pool. Our model assumes that $f_f + f_s = 1$, so we have just a single fraction value to estimate (f_f).

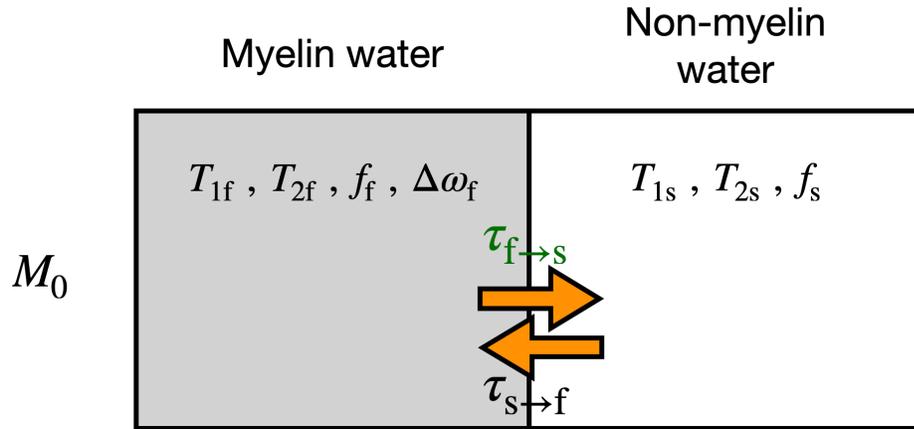


Figure 2.3: Two pool exchanging pool for myelinated tissue.

Studies have shown that myelin water experiences an additional frequency shift that is dependent on the orientation of myelinated axons wrt the main magnetic field B_0 [165, 36]. This is denoted by $\Delta\omega_f$ in our model. In other words, if $\Delta\omega$ is the bulk off-resonance, then the myelin water pool experiences a total off-resonance of $\Delta\omega + \Delta\omega_f$, while the non-myelin water experiences an off-resonance of just $\Delta\omega$.

We model the exchange dynamics between the two pools in the form of mean residence times (which are the inverse of exchange rates). $\tau_{f \rightarrow s}$ represents the mean residence time of myelin water; a low value of $\tau_{f \rightarrow s}$ implies a high exchange rate of water molecules from the myelin water pool to the non-myelin water pool and vice versa. Similarly, $\tau_{s \rightarrow f}$ denotes the mean residence time of non-myelin water, and this represents the exchange process in the opposite direction, i.e., from the non-myelin water pool to the myelin water pool. Additionally, this two pool exchanging model assumes the two compartments to be in chemical equilibrium, meaning $f_f \tau_{s \rightarrow f} = f_s \tau_{f \rightarrow s}$. In other words, $\tau_{f \rightarrow s}$ is the only exchange parameter to be estimated; once we obtain $\tau_{f \rightarrow s}$, the other exchange parameter $\tau_{s \rightarrow f}$ is determined using the aforementioned equation.

Parameters of the two pool model:

- Equilibrium magnetization M_0
- Myelin water fraction f_f
- Compartment-specific relaxation time constants $T_{1f}, T_{2f}, T_{1s}, T_{2s}$
- Mean residence time of myelin $\tau_{f \rightarrow s}$
- Frequency shift specific to myelin water $\Delta\omega_f$
- Bulk off-resonance $\Delta\omega$
- B_1+ scaling factor κ

2.4 Parameter Estimation via Regression with Kernels (PERK)

In quantitative MRI, parameter estimation is the process of obtaining parameter maps from the reconstructed images. Statistical approaches to parameter estimation include maximum likelihood (ML) [46] or maximum a posteriori (MAP) based estimation [183, 181].

These methods typically involve highly non-convex cost functions that are challenging to optimize. Another group of methods involve dictionary matching, where a dictionary is created consisting of several candidate signal vectors, and the acquired data is compared with the dictionary to obtain the estimate of the unknown parameters. This approach is quite commonly seen in fingerprinting-based approaches [91, 65]. This approach works well for smaller problems, but might not scale well with increase in the number of latent parameters to be estimated (as is the case with multi-compartment tissue models), since the size of the dictionary grows exponentially with increase in latent parameters. There has also been a lot of work exploring deep learning-based approaches for parameter estimation [177, 22, 58, 37, 158].

Parameter Estimation via Regression with Kernels (PERK) is a fast, dictionary-free approach for parameter estimation [114], based on kernel ridge regression. Standard linear regression techniques are simple and scalable, but are not suitable for many practical estimation problems which are non-linear, e.g., non-linear MR signal models in quantitative MRI. Kernel ridge regression gives us a way to perform non-linear regression in the input space, by first transforming or *lifting* the measurements into a higher-dimensional space and then performing linear regression in that *lifted* space. This transforms non-linear regression in the input space into linear regression in a feature space [148], defined by the use of an appropriate kernel function (e.g., Gaussian kernel).

Performing linear regression in such a high-dimensional feature space can be very computationally demanding due to the potentially large number of features. Fortunately, due to a popular trick called the kernel trick, kernel functions help us perform regression in the higher-dimensional space without explicitly computing the coordinates in the high-dimensional space, which reduces the computational burden of these techniques. Kernel methods have been quite popular in machine learning, e.g., classification [25] as well as regression [136]. Note that this machine learning form of kernel regression is fairly different from the form used in [151] for image processing.

PERK for Quantitative MRI:

In PERK, we first simulate many input-output measurement pairs based on prior distributions and a non-linear MR signal model. These are lifted to a higher-dimensional (possibly infinite) space using an appropriate kernel function, and then an affine regression function is learnt using the simulated training datapoints. To illustrate how PERK works, let us consider a typical non-linear MR signal model:

$$\mathbf{y} = \mathbf{s}(\mathbf{x}, \boldsymbol{\nu}; \mathbf{P}) + \boldsymbol{\epsilon}, \quad (2.25)$$

where $s(\mathbf{x}, \boldsymbol{\nu}; \mathbf{P})$ is the output of the MR signal model, $\mathbf{x} \in \mathbb{R}^L$ is the set of L latent parameters to be estimated, $\boldsymbol{\nu} \in \mathbb{R}^K$ is the set of K known parameters, \mathbf{P} is the collection of acquisition parameters for the sequence, and $\mathbf{y} \in \mathbb{C}^D$ is the set of D measurements. For instance, \mathbf{x} could contain tissue properties such as proton density and T_2 , while $\boldsymbol{\nu}$ could include B_0 and B_1+ information that are assumed to be known to us (e.g., acquired with separate pre-scans). $\boldsymbol{\epsilon} \sim \mathcal{CN}(\mathbf{0}_D, \boldsymbol{\Sigma})$ is zero-mean complex Gaussian noise with known covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{D \times D}$.

To generate a total of N input-output training pairs, we sample the space of $\mathbb{R}^L \times \mathbb{R}^K \times \mathbb{C}^D$ to generate N different realizations of the input parameters and noise $\{(\mathbf{x}_1, \boldsymbol{\nu}_1, \boldsymbol{\epsilon}_1), (\mathbf{x}_2, \boldsymbol{\nu}_2, \boldsymbol{\epsilon}_2), \dots, (\mathbf{x}_N, \boldsymbol{\nu}_N, \boldsymbol{\epsilon}_N)\}$ and the corresponding measurement vectors $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$ [112]. These N training datapoints are collected as $\{(\mathbf{q}_1, \mathbf{x}_1), (\mathbf{q}_2, \mathbf{x}_2), \dots, (\mathbf{q}_N, \mathbf{x}_N)\}$, where $\mathbf{q}_n := [|\mathbf{y}_n|^T, \boldsymbol{\nu}_n^T]^T \in \mathbb{R}^Q$ and $Q := D + K$. PERK seeks to learn a non-linear regression function (using an appropriate kernel function), that predicts the estimate \mathbf{x}_n based on the regressor \mathbf{q}_n for all training points $n = 1, 2, \dots, N$.

Let $\mathbf{q} \in \mathbb{R}^Q$ be the acquired test data. Then, the PERK estimator for the L latent parameters is given by [114]

$$\hat{\mathbf{x}}(\mathbf{q}) = \mathbf{X} \left(\frac{1}{N} \mathbf{1}_N + \mathbf{M}(\mathbf{M}\mathbf{K}\mathbf{M} + \rho N \mathbf{I}_N)^{-1} \mathbf{k}(\mathbf{q}) \right), \quad (2.26)$$

where $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_N] \in \mathbb{R}^{L \times N}$ contains the training estimates, $\mathbf{1}_N$ is a vector of N ones, \mathbf{I}_N is the $N \times N$ identity matrix, $\mathbf{M} := \mathbf{I}_N - \frac{1}{N} \mathbf{1}_N \mathbf{1}_N^T$ is a de-meaning operator and ρ is a regularization parameter. $\mathbf{K} \in \mathbb{R}^{N \times N}$ is the Gram matrix whose entries correspond to the pairwise kernel evaluations of all training data, i.e., $\mathbf{K}_{ij} = k(\mathbf{q}_i, \mathbf{q}_j)$, where $k(\cdot, \cdot)$ is the chosen kernel function. The term $\mathbf{k}(\mathbf{q})$ in (2.26) is defined as $\mathbf{k}(\mathbf{q}) := [k(\mathbf{q}, \mathbf{q}_1), k(\mathbf{q}, \mathbf{q}_2), \dots, k(\mathbf{q}, \mathbf{q}_N)]^T \in \mathbb{R}^N$.

For this work, we use a Gaussian kernel function defined as

$$k(\mathbf{q}_i, \mathbf{q}_j) = \exp \left(-\frac{1}{2} \|\mathbf{q}_i - \mathbf{q}_j\|_{\boldsymbol{\Lambda}^{-2}}^2 \right), \quad (2.27)$$

where $\boldsymbol{\Lambda} \in \mathbb{R}^{Q \times Q}$ is a diagonal matrix that controls the length scales of each latent parameter over which the estimator $\hat{\mathbf{x}}(\cdot)$ smooths and $\|\mathbf{z}\|_{\mathbf{W}} = \|\mathbf{W}^{1/2} \mathbf{z}\|_2$ is a weighted l^2 -norm with weights \mathbf{W} [112]. This can be interpreted as lifting the input data into an infinite dimensional space using a non-linear Gaussian kernel, before learning a linear regression function in that *lifted* space to predict the unknown latent parameters. The diagonal matrix

Λ represents the bandwidth matrix of the Gaussian kernel, and is chosen as [114]

$$\Lambda \triangleq \lambda \text{diag} \left([\mathbf{m}_y^T, \mathbf{m}_\nu^T]^T \right), \quad (2.28)$$

where $\mathbf{m}_y \in \mathbb{R}^D$ contains the sample means of all D measurements across voxels of magnitude test data and $\mathbf{m}_\nu \in \mathbb{R}^K$ contains sample means of the known parameter values. The scalar parameter $\lambda > 0$ provides an additional tuning knob to make PERK invariant to the scale of the test data; this is selected using holdout as described in [114]. $\text{diag}(\cdot)$ constructs a diagonal matrix using its entries.

PERK provides a scalable approach to multi-parametric estimation problems, and seems to be suited to our two-pool exchanging model, since we have many parameters to jointly estimate. Chapter 3 introduces methods to estimate myelin water exchange in white matter, and uses PERK as the estimator of choice.

CHAPTER 3

Quantifying Myelin Water Exchange using Optimized bSSFP Sequences

3.1 Introduction

In quantitative MRI, many classical approaches assume that the MR signal from a voxel arises from a single type of tissue. However, this assumption may not be appropriate when modeling the complex spin-dynamics of microstructural components in living tissues. A more suitable model would be a multi-compartment relaxation model, where the MR signal is assumed to originate from different tissue compartments. Quantifying the exchange dynamics between these various compartments has been suggested as a potential biomarker for several clinical applications [76, 21, 188, 4]. This work appeared as an abstract in ISMRM 2022 [110].

An example of a multi-compartment tissue model can be seen in myelin water imaging, where myelin content is studied in brain tissue [78, 2, 94, 81]. In myelin water imaging, the MR signal is typically assumed to arise from two water pools: myelin water and non-myelin water. A typical parameter of interest in myelin water imaging is the myelin water fraction (MWF) which is the fraction of MR signal that arises from myelin water [2, 164]. Myelin water fraction is thought to be useful for tracking demyelination [164], which is a hallmark of various neurological disorders, e.g., multiple sclerosis [93].

Previous work has shown that MWF is a proxy for myelin content¹ in tissue [147, 77]. However, some studies found variations in MWF between different rat spinal cord tracts, even when the underlying myelin content was similar [34, 51]. They suggested that the variations in MWF could be explained by variations in inter-compartmental water exchange, brought about by differences in axon diameter and myelin thickness. Hence, myelin water exchange potentially plays an important role in myelin water imaging, and in

¹In this work, we distinguish myelin water fraction (MWF) from myelin volume fraction (MVF) [164], which represents the myelin content in the voxel (by volume).

this work, we focus on quantifying the inter-compartmental exchange dynamics between myelin water and non-myelin water. Though there has been a lot of work focusing on the feasibility of myelin water fraction as a parameter of interest and a clinical biomarker [2, 94, 81, 166], this work specifically focuses on myelin water exchange, since we think that the exchange information could potentially be a useful biomarker in itself.

The question of how many tissue compartments to use for modeling a multi-compartment MR tissue model is context-dependent. For instance, in magnetization transfer (MT) experiments that study interactions between protons in free water and protons bound to macromolecules, a two pool model consisting of two exchanging compartments (one aqueous and one non-aqueous) is commonly used [143]. Manning et al., argue that, to get a complete picture of longitudinal relaxation in white matter, a two pool model is not sufficient [97]. They propose a four pool model considering T_1 relaxation in the myelin water, intra/extra cellular water and their associated non-aqueous pools [145]. Another example of a two pool model is one with two aqueous exchanging pools, corresponding to myelin water and non-myelin water. Chapter 2 reviewed the parameters of such a model for myelinated tissue. This model has been previously used for estimating myelin water fraction in the brain [166].

This chapter develops methods to quantify myelin water exchange between myelin water and non-myelin water. We use a similar two pool exchanging model as in [166], i.e., two aqueous pools corresponding to myelin water and non-myelin water respectively, and we perform MRI experiment design to help estimate myelin water exchange in white matter. We optimize the acquisition parameters (flip angles and phase cycling factors) of a set of balanced steady-state free precession (bSSFP) scans to estimate myelin water exchange with high precision. By optimizing a cost function based on the Cramér-Rao bound, we aim to design a set of bSSFP acquisitions that minimize the variance in the exchange estimates (for any unbiased estimator). A novel contribution of this work is to optimize a set of RF phase cycling factors of bSSFP acquisitions for quantitative MRI, using a Cramér-Rao bound-based cost function.

The rest of the chapter is organized as follows: Section 3.2 describes some related work for estimating inter-compartmental exchange. Section 3.3 describes the methods used for our work, including details of the bSSFP signal model, the Cramér-Rao bound based scan design, as well as how we estimate myelin water exchange using the PERK estimator (parameter estimation via regression with kernels) [114]. Section 3.4 provides details of our experimental setup and results, while Section 3.5 discusses our findings, along with a few ideas for future work. Section 3.6 talks about our conclusions.

3.2 Related work

Previous approaches have tried to characterize exchange dynamics from a few different directions. One approach is to use magnetization transfer and T_2 measurements in myelinated tissue, and fit the resulting data to a four pool model (myelin and non-myelin water, along with their associated macromolecular pools) [145, 11, 70]. These studies can be challenging due to the large number of free parameters to be jointly estimated. Another approach to studying inter-compartmental water exchange is through relaxation exchange spectroscopy (REXSY) [33, 32]. In REXSY, the presence of exchange during a mixing period is directly observed as off-diagonal components in a 2D $T_2 - T_2$ relaxation spectrum. REXSY could potentially be used as a baseline technique for estimating exchange, but it has prohibitively long scan times, which makes it infeasible for *in vivo* experiments.

Deoni et al. proposed a technique called mcDESPOT (multi-component driven equilibrium single pulse observation of T_1/T_2) for quantifying multi-component relaxation in the brain [29]. mcDESPOT uses a set of steady-state acquisitions (combination of SPGR and bSSFP acquisitions with varying flip angles) to encode multi-compartment T_1 and T_2 information. Though mcDESPOT primarily focuses on T_1 and T_2 estimation, the multi-parametric estimation approach in mcDESPOT also yields estimates of exchange rates in myelin water. Fig. 3.1 shows a sample map of mean residence time of myelin, obtained using mcDESPOT, taken from [29].

Typically, in quantitative imaging, separate acquisitions were previously used to estimate parameters of interest such as T_1 , T_2 , proton density, etc., leading to very long acquisition times. A relatively recent approach called MR Fingerprinting (MRF) proposed the use of a single acquisition scheme that simultaneously encodes information about multiple

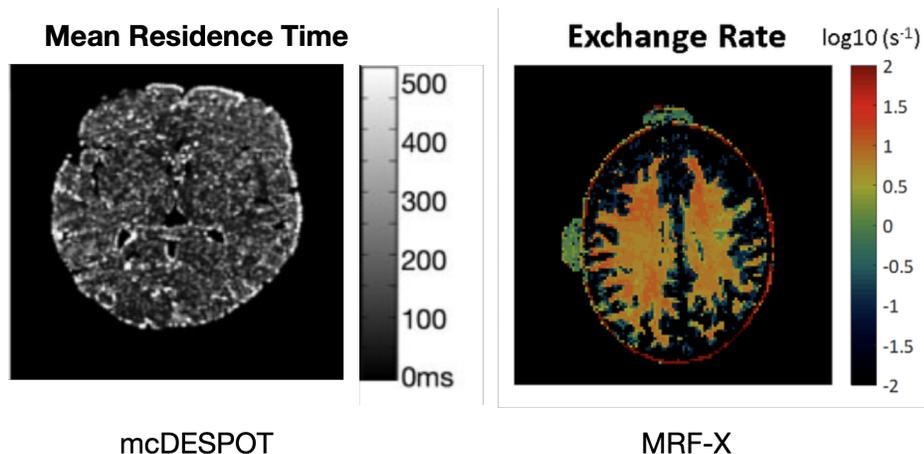


Figure 3.1: Exchange maps obtained using mcDESPOT [29] and MRF-X [48].

quantitative parameters [91]. Fingerprinting has been widely used for T_1 and T_2 quantification in the brain, and has demonstrated fast acquisition times along with good image resolution [5, 92]. Extending the original fingerprinting approach, Hamilton et al. proposed a technique called MRF-X (MR Fingerprinting with Chemical Exchange) to quantify the exchange dynamics in a multi-compartment model of exchanging pools [48]. Fig. 3.1 shows an exchange map derived using MRF-X, taken from [48].

One of the bottlenecks for exchange mapping *in vivo* is the lack of an established baseline method that could be used to generate ground truth (cf. inversion recovery approaches for T_1 mapping, multi-echo spin-echo methods for T_2 and MWF estimation etc.). This makes it hard to validate exchange maps obtained using a new approach. Existing approaches (such as mcDESPOT and MRF-X) produce exchange maps, but they were not validated against ground truth maps. This chapter focuses on developing methods for designing MRI acquisitions and estimating myelin water exchange, also without extensively validating the proposed techniques. Chapter 4 tackles the question of how to validate this proposed technique for estimating exchange; we do so by developing an aqueous urea system that can be used for obtaining ground truth exchange values using other established methods such as inversion recovery based NMR (IR-NMR) [168] spectroscopy experiments.

Another contribution of this work is to optimize the acquisition parameters of a set of bSSFP acquisitions with a focus on obtaining more precise exchange estimates. Both mcDESPOT and MRF-X use hand-crafted sequences where the acquisition parameters were not optimized for exchange estimation. Optimizing the acquisition parameters (especially flip angles and phase cycling factors) could potentially lead to sequences that are more sensitive to exchange, and hence result in better exchange estimates. In particular, optimizing a set of RF phase cycling factors of bSSFP acquisitions, for parameter estimation in quantitative MRI, is a novel contribution of this work.

3.3 Methods

3.3.1 bSSFP signal model

The bSSFP (balanced steady-state free precession) sequence consists of a rapid train of RF excitation pulses in the presence of balanced gradients, i.e., gradients whose net area is zero across a repetition period (T_R). Fig. 3.2 (a) shows a schematic of a bSSFP sequence². The RF excitation pulse is parametrized by flip angle (FA) α and RF phase cycling factor

²Schematic taken from YouTube lectures by Dr. Daniel Ennis and Dr. Brian Hargreaves.

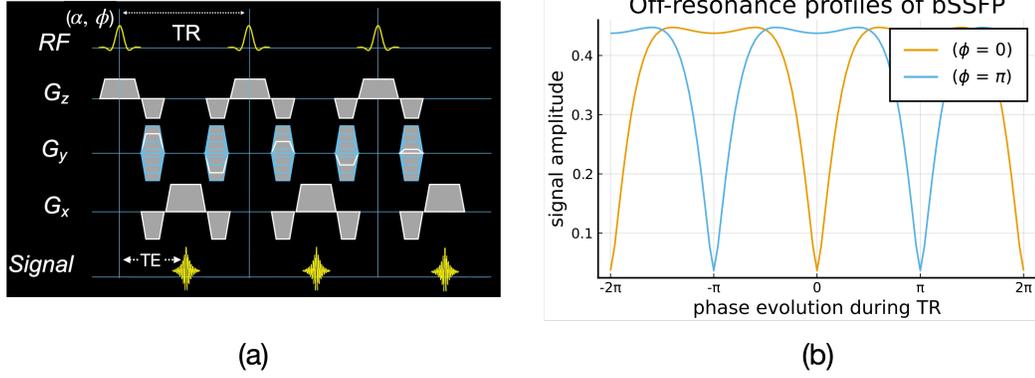


Figure 3.2: (a) bSSFP pulse sequence. (b) Effect of varying RF phase cycling factors (ϕ) on simulated bSSFP off-resonance profiles. (Note: α = flip angle, T_R = repetition time, T_E = echo time)

ϕ . The ϕ factor represents the phase increment between successive RF pulses; e.g., $\phi = 0$ corresponds to an RF pulse train where all pulses have the same phase, while $\phi = \pi$ corresponds to an alternating phase RF pulse train.

Since the gradients are balanced in a bSSFP sequence, the only source of phase accrual in every T_R is due to off-resonance precession, and not due to gradient dephasing. A consequence of this is that the bSSFP signal is quite sensitive to off-resonance, as illustrated in Fig. 3.2 (b). The orange line in the plot corresponds to the off-resonance profile (signal magnitude vs. phase accrued) when $\phi = 0$. The acquisition parameters were chosen to be $T_1 = 20T_R$, $T_2 = 15T_R$, $\alpha = 70^\circ$, to match the single pool MR system used in [139]. We can observe regions of high signal (referred to as the pass band), and regions with signal nulls (referred to as stop/transition band). MR spins that have a resonance frequency corresponding to these signal nulls have very low signal, and this results in dark banding artefacts in the image ([104] shows examples of these dark bands in the brain).

Fig. 3.2 (b) also shows the effect of varying the RF phase cycling factor ϕ . By using alternating phase RF pulses ($\phi = \pi$), the pass bands and transition bands shift by π , as shown using the blue line. For $\phi = 0$, on-resonance spins would fall into a signal null, while they lie within the pass band for an alternating RF pulse train ($\phi = \pi$), and hence result in much higher signal. Typical bSSFP acquisitions use RF pulses with alternating phase since it results in higher SNR for nearly on-resonance spins [56]. The location of the signal null manifests as a dark banding artefact in bSSFP images. In this work, we optimize a set of ϕ factors to better estimate exchange, and this shifts the dark bands accordingly (since ϕ determines the location of the signal null).

For a homogeneous sample with a single set of tissue properties, we obtain a symmetric bSSFP signal profile as shown in Fig. 3.2 (b), where the profiles are symmetric about

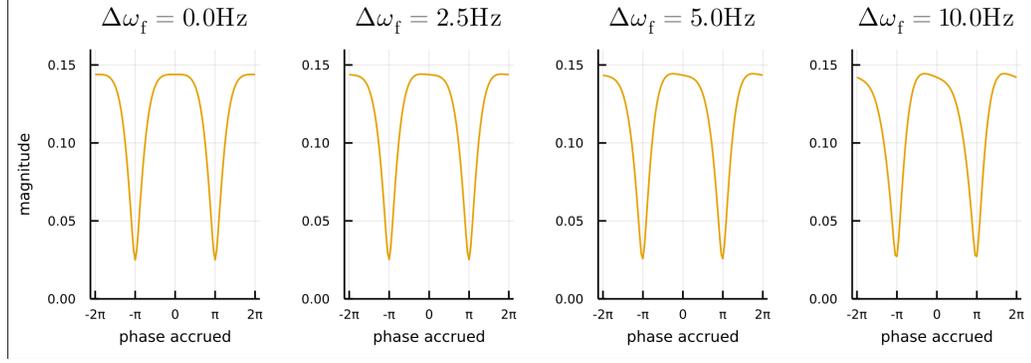


Figure 3.3: Simulated bSSFP off-resonance profiles for different values of the myelin-specific frequency shift $\Delta\omega_f$, as a function of phase accrued during one T_R .

the center frequency. However, bSSFP profiles measured *in vivo* have been shown to be asymmetric in nature [102]. This is thought to be due to the presence of tissue components with an asymmetric frequency distribution [101]. In our two-pool exchanging model, we model an additional frequency shift that is specific to myelin water ($\Delta\omega_f$) and is thought to arise from susceptibility effects due to the orientation of the fibers in white matter [36, 165]. Fig. 3.3 illustrates the asymmetries of bSSFP signal profiles when applied to a two pool

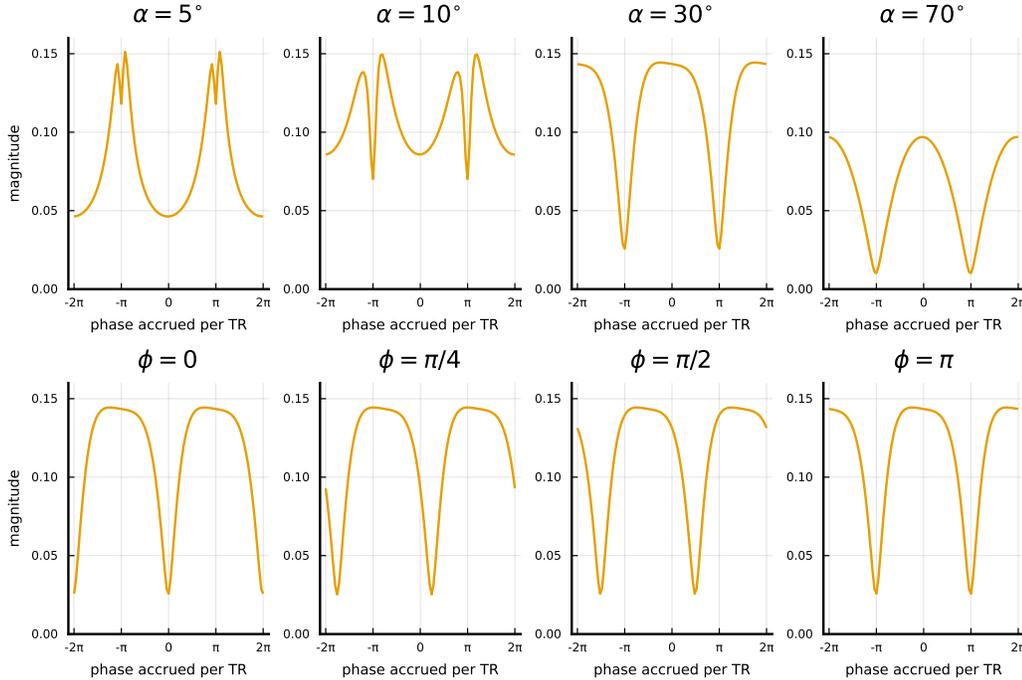


Figure 3.4: Simulated bSSFP off-resonance profiles for various flip angles α (top row) and for different RF phase cycling factors ϕ (bottom row), as a function of phase accrued during one T_R . (Note: $\phi = \pi$ for the top row, and $\alpha = 30^\circ$ for the bottom row.)

exchanging model. For these simulations, we considered a two pool exchanging model for a white matter voxel with the following tissue parameters: $f_f = 0.15$, $T_{1f} = 400\text{ms}$, $T_{1s} = 1000\text{ms}$, $T_{2f} = 20\text{ms}$, $T_{2s} = 80\text{ms}$, $\tau_{f \rightarrow s} = 100\text{ms}$. The acquisition parameters were chosen as $\alpha = 30^\circ$, $\phi = \pi$, and $T_R/T_E = 20\text{ms}/0\text{ms}$. The off-resonance profiles show increasing asymmetry as the myelin-specific frequency shift $\Delta\omega_f$ increases.

Fig. 3.4 shows the effects of varying the flip angles α (top row) and RF phase cycling factors ϕ (bottom row) on the bSSFP off-resonance profiles. We used the same tissue parameters as in Fig. 3.3, and fixed $\Delta\omega_f$ to 5Hz. Variations in flip angles modulate the shape of the off-resonance profiles, while varying the RF phase cycling factors *shifts* the off-resonance profiles [56, 189].

3.3.2 CRB-based scan design

This subsection describes our methods used for scan design. Our objective is to design a set of MR acquisitions to help us estimate myelin water exchange in white matter. We aim to do this by optimizing the acquisition parameters of a set of phase-cycled bSSFP acquisitions, using a cost function based on the Cramér-Rao bound. We provide more details about how we use the Cramér-Rao bound, and how it helps us encode information about myelin water exchange in our scan design.

The Cramér-Rao bound (CRB) provides a lower bound on the covariance of unbiased estimators [74]. It provides a way of evaluating how well the measurements can encode information about the parameters of interest. The CRB is typically applicable only for unbiased estimators. Most estimators used in common practice are biased, since regularized approaches generally result in biased estimates. However, minimizing the CRB can still be a useful tool for experiment design, since the CRB is estimator-agnostic (for unbiased estimators) and this design approach can ensure that the data acquisition process is sensitive to the parameter of interest [47, 182].

We assume the acquisition of fully sampled k-space data in this chapter. This simplifies analysis since there are no aliasing artefacts, and parameter estimation can be done separately on a voxel-by-voxel basis. Under these conditions, the MR image data for a single voxel, for a single scan, is modeled as

$$y = s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{p}) + \epsilon, \quad (3.1)$$

where $s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{p}) \in \mathbb{C}$ is the output of the MR signal model, \mathbf{x} is the vector of unknown parameters to be estimated, $\boldsymbol{\nu}$ is the vector of known parameters, \mathbf{p} is the vector of acquisition parameters for a single scan and $\epsilon \sim \mathbb{C}\mathcal{N}(0, \sigma_\epsilon^2)$ is complex-valued Gaussian noise. E.g.,

for our work with estimating myelin water exchange, s represents the bSSFP signal model, $\mathbf{x} = [M_0, f_f, T_{1f}, T_{1s}, T_{2f}, T_{2s}, \tau_{f \rightarrow s}, \Delta\omega_f] \in \mathbb{R}^8$ collects the unknown tissue parameters to be estimated, $\boldsymbol{\nu} = [\Delta\omega, \kappa] \in \mathbb{R}^2$ contains B_0 and B_1+ parameters that are estimated through separate acquisitions, and \mathbf{p} contains the flip angle and RF phase cycling factor for a single bSSFP acquisition (to be optimized). The outputs of D scans are modeled together as

$$\mathbf{y} = \mathbf{s}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P}) + \boldsymbol{\epsilon}, \quad (3.2)$$

where $\mathbf{y} \in \mathbb{C}^D$ and $\mathbf{s}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P}) = [s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{p}_1), s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{p}_2), \dots, s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{p}_D)]^T \in \mathbb{C}^D$. $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_D]$ is the set of acquisition parameters for all D scans. E.g., for our work, \mathbf{P} collects the flip angles and RF phase cycling factors for all D bSSFP acquisitions. $\boldsymbol{\epsilon} \sim \mathcal{CN}(0, \boldsymbol{\Sigma})$ represents the noise vector. We assume that the noise in each acquisition is identical to, and independent of the other acquisitions; hence, $\boldsymbol{\Sigma} = \sigma_\epsilon^2 \mathbf{I}_D$, where σ_ϵ^2 is the noise variance of each acquisition, and \mathbf{I}_D is the $D \times D$ identity matrix.

For all our scan designs, we take the magnitude of the signal model. The noise in magnitude images in MR imaging is typically Rician distributed, but it has been shown that this can be approximated with Gaussian noise under fairly high SNR [45]. Following previous work, we model the noise in the magnitude images as Gaussian noise with variance σ^2 [166]; this variance is signal-dependent but we ignore that dependence for this work. Under these assumptions, the Fisher information matrix \mathbf{I} for the magnitude of the signal model in (3.2) is computed as [113]

$$\mathbf{I}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P}) = \frac{1}{\sigma^2} [\nabla_{\mathbf{x}} |s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P})|] [\nabla_{\mathbf{x}} |s(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P})|]^T, \quad (3.3)$$

where $\nabla_{\mathbf{x}}$ is the column gradient of the MR signal model output, with respect to the unknown parameters \mathbf{x} . The size of the Fisher information matrix is $L \times L$ where L is the number of unknown parameters in \mathbf{x} . The inverse Fisher matrix $\mathbf{I}^{-1}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P})$ gives the CRB, which is a theoretical lower bound for the covariance of unbiased estimators of \mathbf{x} [74]. In other words, if $\hat{\mathbf{x}}(\mathbf{y})$ is any unbiased estimator of \mathbf{x} , i.e., $\mathbb{E}[\hat{\mathbf{x}}(\mathbf{y})] = \mathbf{x}$ where $\mathbb{E}[\cdot]$ denotes expectation, then the covariance matrix is bounded by $\text{cov}(\hat{\mathbf{x}}(\mathbf{y})) \succeq \mathbf{I}^{-1}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P})$.

In particular, the variance of any unbiased estimator of the i th unknown parameter ($\hat{\mathbf{x}}_i(\mathbf{y})$) cannot be lower than the CRB of the i th parameter, i.e.,

$$\text{var}(\hat{\mathbf{x}}_i(\mathbf{y})) \geq \text{CRB}_i = [\mathbf{I}^{-1}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P})]_{(i,i)}, \quad (3.4)$$

where the CRB of the i th parameter is given by the i th diagonal element of the inverse Fisher matrix. The relation in (3.4) gives us a way to improve the precision in the estimates

of our desired parameter, by designing MR acquisitions that minimize the CRB of that parameter. The CRB has been previously used for MRI experiment design [47, 182, 113]. This chapter aims to design a set of phase-cycled bSSFP acquisitions for estimating myelin water exchange in white matter. We do this by minimizing the expectation of the weighted CRB over a search space of acquisition parameters \mathcal{P} [166]:

$$\hat{\mathbf{P}} = \underset{\mathbf{P} \in \mathcal{P}}{\operatorname{argmin}} \mathbb{E}_{\mathbf{x}, \boldsymbol{\nu}} [\operatorname{trace}(\mathbf{W}(\mathbf{I}^{-1}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{P})))], \quad (3.5)$$

where \mathbf{W} is a diagonal weighting matrix that is chosen to reflect our parameters of interest, $\mathbb{E}_{\mathbf{x}, \boldsymbol{\nu}}$ is an expectation w.r.t \mathbf{x} and $\boldsymbol{\nu}$ over a distribution of tissue types (such as white matter) and $\operatorname{trace}(\cdot)$ represents the trace operator of a matrix. For our work with estimating myelin water exchange, we have a single 1 along the diagonal of \mathbf{W} corresponding to the unknown parameter $\tau_{f \rightarrow s}$ and zeros corresponding to the other unknown parameters. We optimized (3.5) using a gradient-based method in the NLOpt³ package, implemented using the programming language Julia⁴.

The Cramér-Rao bound (CRB) gives us a sense of the variance in the parameter estimates (for unbiased estimators), but the CRB directly does not tell us much about how well the MR acquisitions encode information about our parameter of interest, since it does not account for the *mean values of the parameters*. Instead, we use a measure called **coefficient of variation**, which is the ratio of the standard deviation to the mean of the parameter. This is a dimensionless quantity, which measures the variability in parameter estimates relative to the mean value of the parameter. We use the coefficient of variation to examine how well a set of acquisitions encodes information about our desired parameter, and to evaluate different scan designs. The coefficient of variation is defined as:

$$\text{Coefficient of variation} = \frac{\sqrt{\text{Cramér-Rao bound}}}{\text{mean value of parameter}}. \quad (3.6)$$

(3.6) can be evaluated for any of the unknown parameters, and the coefficient of variation gives a sense of how well the optimized scan design encodes information about that parameter of interest. A lower value for the coefficient of variation is better, since it is related directly to the variance of unbiased parameter estimates.

3.3.3 Estimation of myelin water exchange

To optimize the acquisition parameters of our bSSFP sequences for better exchange estimation, we need to specify a set of tissue distributions over which to take the expectation (see (3.5)). Table 3.1 shows the distributions that we use for scan design. For most experiments in this work, we assume B_0 and B_{1+} information to be known, i.e., $\Delta\omega$ and κ are known parameters in ν , i.e., they are estimated using separate acquisitions. The other eight parameters $M_0, f_f, T_{1f}, T_{1s}, T_{2f}, T_{2s}, \tau_{f \rightarrow s}, \Delta\omega_f$ are unknown latent parameters (grouped in \mathbf{x}), and are jointly estimated.

We specifically focus on white matter (WM) since grey matter (GM) has very low myelin content and exchange estimates may be unreliable in GM. For scan design, we chose narrow distributions for many of the parameters, centered around reasonable values for white matter [166]. We chose broader ranges for mean residence time $\tau_{f \rightarrow s}$ and myelin-specific frequency shift $\Delta\omega_f$ since we want to optimize the bSSFP scans for a range of exchange rates and different orientations of the axonal fibers.

Once we have our optimized set of phase-cycled bSSFP acquisitions, we need to use a suitable estimator to estimate myelin water exchange. For our work, we use a fast dictionary-free method called Parameter Estimation via Regression with Kernels (PERK) [112, 114]. The PERK estimator first lifts the measurement data to a higher-dimensional (possibly infinite) space using a non-linear kernel, and then learns an affine regression function (in that lifted space) to predict the unknown tissue parameters. The PERK estimator

| Parameter | Distributions for scan design | PERK training ranges |
|----------------------------------|-------------------------------|----------------------|
| M_0 | Unif(0.769, 0.771) | Unif(0.75, 1.0) |
| f_f | Unif(0.149, 0.151) | Unif(0.03, 0.31) |
| T_{1f} (in ms) | Unif(399, 401) | Unif(300, 500) |
| T_{1s} (in ms) | Unif(831, 833) | Unif(800, 1350) |
| T_{2f} (in ms) | Unif(19.9, 20.1) | Unif(16, 24) |
| T_{2s} (in ms) | Unif(79.9, 80.1) | Unif(64, 96) |
| $\tau_{f \rightarrow s}$ (in ms) | Unif(50, 250) | Unif(50, 250) |
| $\Delta\omega_f$ (in Hz) | Unif(0, 10) | Unif(0, 10) |
| $\Delta\omega$ (in Hz) | Unif(-25, 25) | Unif(-25, 25) |
| κ | Unif(0.8, 1.2) | Unif(0.8, 1.2) |

Table 3.1: Distributions of tissue parameters used for optimizing the bSSFP acquisition parameters, as well as the tissue distributions used to train the PERK estimator. (Note: bSSFP signal profiles are periodic with a period of $1/T_R$, and hence, we select the off-resonance distribution $\Delta\omega$ to span 50Hz, corresponding to a T_R of 20ms.)

³<https://github.com/JuliaOpt/NLopt.jl>

⁴<https://julia-lang.org/>

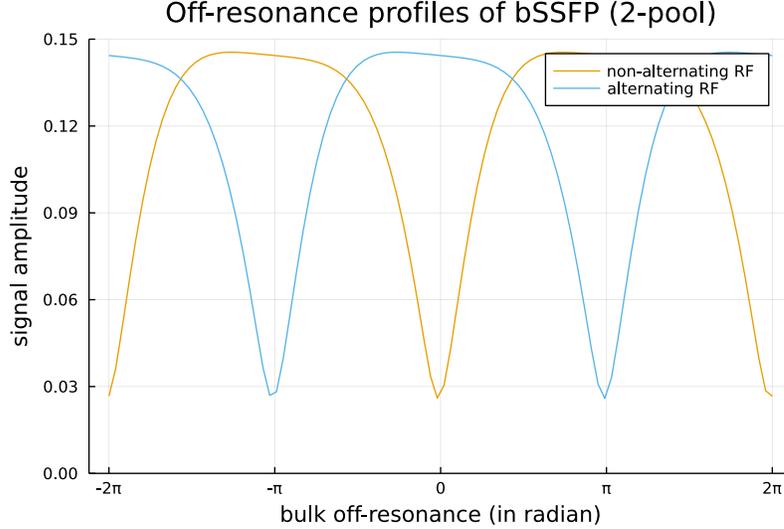


Figure 3.5: Comparison of bSSFP signal magnitudes for alternating ($\phi = \pi$) and non-alternating ($\phi = 0$) RF pulses. A typical WM voxel is simulated here, with a mean residence lifetime ($\tau_{f \rightarrow s}$) of 250ms and a myelin-specific frequency shift ($\Delta\omega_f$) of 7.5Hz.

needs to be trained using prior distributions of the tissue parameters, and this can then be used to make predictions from acquired data. Chapter 2 contains more details about the inner workings of PERK.

Though the bSSFP sequences are specifically optimized for white matter tissue characteristics (column 2 of Table 3.1), we train PERK over a wider range of tissue parameters since we want our estimator to be generalizable. Table 3.1 shows the tissue distributions used to train PERK (see column 3). We apply the trained PERK estimator to test data (simulated data or data acquired *in vivo*) to jointly estimate the exchange parameter $\tau_{f \rightarrow s}$ and the other unknown parameters.

3.4 Experiments

3.4.1 RF phase cycling factors (ϕ) in bSSFP sequences

In a typical bSSFP sequence, when all RF pulses are all at the same phase, signal nulls are obtained for on-resonance spins (see Fig. 3.5 for the case of constant phase RF pulses applied to a typical WM voxel). Alternating the phase of successive RF pulses shifts the passband by π so that we get higher signal for on-resonance spins [56, 189] (see the profile for alternating RF pulses in Fig. 3.5). The location of the signal nulls leads to the commonly seen banding artifacts in bSSFP images, since spins that are precessing at the stopband frequencies tend to have very low signal. By changing the RF phase increments between

successive pulses (referred to as RF phase cycling factors ϕ), the passband can be shifted to cover different regions of the spectrum.

A naive implementation of the bSSFP signal model for arbitrary ϕ would be tricky, since each successive RF excitation leads to a different linear system of matrices and it is hard to solve such a complex system of equations. It has been previously shown that, an alternating RF pulse train (i.e., $\phi = \pi$) can be treated exactly like a constant-phase RF pulse train ($\phi = 0$) but with resonance frequencies being offset by an amount corresponding to phase π [56]. We empirically extended this approach to an arbitrary ϕ which makes it much easier to solve for the steady-state magnetization for arbitrary RF phase cycling factors. In other words, we model an RF pulse train with arbitrary phase cycling factor ϕ , as a constant-phase RF pulse train but with the resonance frequencies being offset by $\phi/(2\pi T_R)$ [56]. This results in a simpler system of matrix equations describing the evolution of magnetization, and the steady-state bSSFP signal can be derived for any arbitrary phase cycling factor ϕ [50].

3.4.2 Importance of optimizing RF phase cycling factors ϕ

In this work, we optimize the acquisition parameters of a set of 40 bSSFP sequences to minimize the variance of unbiased exchange estimates in white matter. For phase-cycled bSSFP acquisitions, we could potentially optimize: flip angles (α), RF phase cycling factors (ϕ), repetition times (T_R) and echo times (T_E). The search space for acquisition parameters can be quite extensive and hence, in this set of experiments, we perform an ablation study to determine which acquisition parameters are most important in reducing the coefficient of variation for exchange. Our initial experiments suggested that ϕ factors play an important role in estimating myelin water exchange; varying the ϕ factors across the 40 bSSFP acquisitions seemed to result in lower coefficients of variation, i.e., more precise exchange estimates.

Table 3.2 contains the results of the ablation study. We compared four scan designs (T_E values were fixed at 4ms):

- (i) T_R , α and ϕ values were all optimized for 40 bSSFP acquisitions.
- (ii) Only α and ϕ values were optimized; T_R values were fixed to 20ms.
- (iii) Only T_R and α values were optimized; ϕ factors were fixed to 0° for half the scans, and 180° for the other half.
- (iv) Only α values were optimized; T_R values and ϕ factors were kept fixed as before.

| Experiment | Coeff. of variation |
|---|---------------------|
| optimizing $\{T_{Ri}\}$, $\{\alpha_i\}$ and $\{\phi_i\}$ | 0.13 |
| optimizing $\{\alpha_i\}$ and $\{\phi_i\}$ | 0.13 |
| optimizing $\{T_{Ri}\}$ and $\{\alpha_i\}$ | 0.58 |
| optimizing only $\{\alpha_i\}$ | 561.8 |

Table 3.2: Ablation study showing the importance of optimizing RF phase cycling factors $\{\phi_i\}$ for exchange estimation. Note: $\{T_{Ri}\}$ denotes repetition times and $\{\alpha_i\}$ denotes flip angles of the 40 bSSFP acquisitions.

Table 3.2 shows that designs (iii) and (iv), which do not optimize ϕ factors, result in higher coefficients of variation. In fact, for the last design, which optimizes only the flip angles α , the optimization failed due to extremely high Cramér-Rao bounds. On the other hand, the first two designs (where ϕ factors are optimized) resulted in much lower coefficients of variation. This suggests that variations in ϕ factors could be quite important for encoding information about exchange. Moreover, optimizing T_R values (along with α and ϕ values) did not reduce the coefficient of variation noticeably; both designs (i) and (ii) yielded very similar coefficients of variation. Hence, for the rest of this work, we focus on optimizing only the flip angles (α) and RF phase cycling factors (ϕ) for the 40 different bSSFP acquisitions, and keep the echo times and repetition times fixed.

3.4.3 Optimized scan design

For scan design, we focused on estimating exchange in white matter, because grey matter has very low myelin content. We optimized the flip angles (α) and RF phase cycling factors (ϕ) for 40 different bSSFP acquisitions, with a fixed $T_R/T_E = 20\text{ms}/4\text{ms}$ (chosen based on running Cramér-Rao bound experiments). During optimization, the flip angles were allowed to vary between 10° and 40° , while the phase cycling factors were optimized between $-\pi$ and $+\pi$. Fig. 3.6 shows the numerically optimized bSSFP scan design for estimating myelin water exchange in white matter. This scan design has a predicted coefficient of variation of **13.5%** for estimating exchange, at an average image SNR of 50dB.

It is interesting that, in the optimized scan design, roughly half of all scans have a flip angle of 10° while the other half has a flip angle of 40° . Starting from a randomly initialized scan design, it is also intriguing that the optimized set of 40 acquisitions has a very regular sweep of phase cycling factors ϕ at both the low and high flip angles. In fact, the numerically optimized set of bSSFP acquisitions in Fig. 3.6 could be mistaken for a *handcrafted*

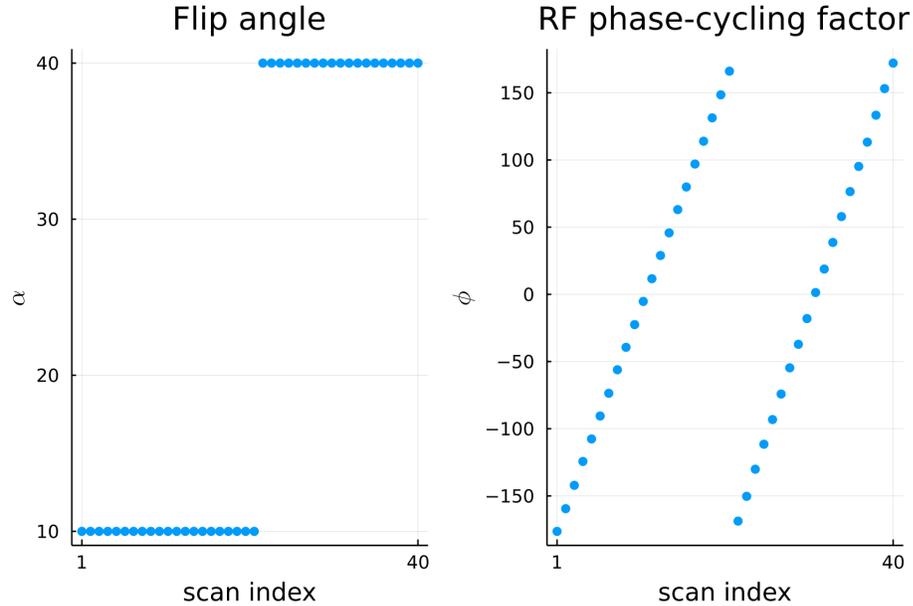


Figure 3.6: Numerically optimized bSSFP scan design for estimating myelin water exchange, consisting of 40 bSSFP scans with flip angles (α) and RF phase cycling factors (ϕ) as shown. (Note: y-axes are in degrees. TR/TE = 20ms/4ms.)

design, instead of one optimized using a cost function. This suggests that this regular variation in ϕ could potentially be important for exchange estimation, and we investigated the importance of optimizing ϕ in an earlier set of experiments (Section 3.4.2 contains an ablation study to examine the importance of optimizing ϕ factors). In Fig. 3.6, we *reordered* the 40 bSSFP acquisitions for visualization purposes; the optimization framework does not result in such an ordered set of α and ϕ values.

To study how sensitive the optimized scan design is to exchange, we performed simulations for a typical voxel in white matter using this optimized scan design, under three different exchange regimes (mean residence times of $\tau_{f \rightarrow s} = 50\text{ms}$, 150ms and 250ms). Fig. 3.7 shows the resulting bSSFP signal magnitudes for all 40 acquisitions. It might seem hard to distinguish between exchange regimes with *every* bSSFP acquisition; however, the low predicted coefficient of variation ($\sim 13.5\%$) suggests that, between them, the 40 scans potentially contain enough diversity and information to estimate exchange with good precision for the two-pool model considered here. Additionally, it is important to note that this is a joint parameter estimation problem, and some of the 40 acquisitions might potentially be helpful in estimating parameters other than exchange. Also, Fig. 3.7 shows the magnitudes for a single type of voxel; the plots could possibly look different if we consider voxels with different tissue parameters, i.e., a different subset of the 40 acquisitions might be sensitive to exchange.

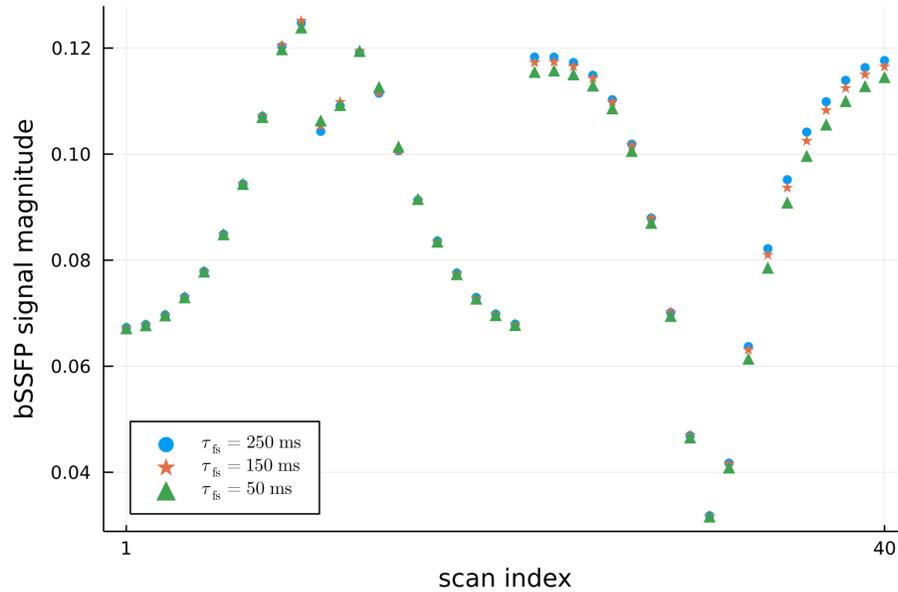


Figure 3.7: Simulated bSSFP signal magnitudes obtained using the optimized scan design for a typical WM voxel (on resonance, and 5Hz additional frequency offset for myelin water), under three different exchange regimes.

3.4.4 Numerical simulations

Directly evaluating the proposed approach on real data is hard, since there are no well-defined gold-standard methods for estimating myelin water exchange. Before acquiring data *in vivo*, a first step is to test the optimized scan design using numerical simulations. Using digital phantoms, the proposed approach can be validated under a wide range of

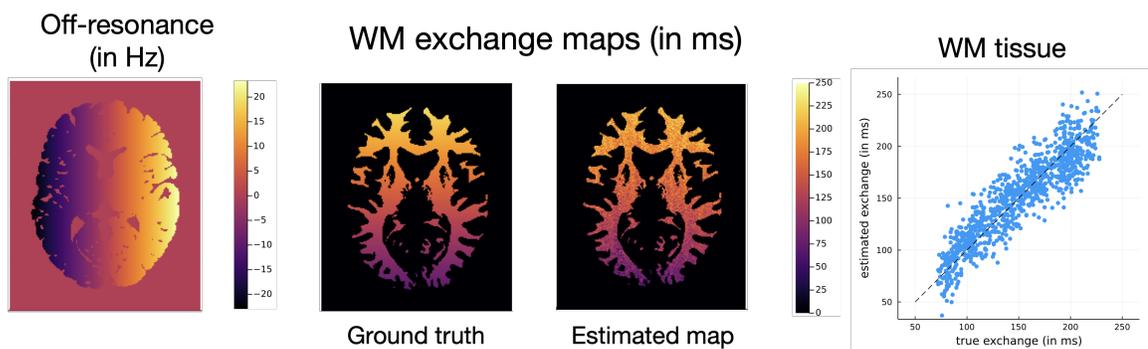


Figure 3.8: Estimated white matter (WM) exchange maps from simulated test data based on a modified Brainweb phantom [23]. (Note: Scatter plot shows exchange estimates of a randomly selected subset of WM voxels.)

conditions (different noise levels, off-resonance conditions, presence/absence of myelin-specific frequency shift, different tissue types etc.). For our experiments, we used a structurally realistic digital brain phantom (BrainWeb) [23], and modified it to reflect different experimental conditions such as varying off-resonance and mean residence time of myelin.

Fig. 3.8 shows the ground truth and estimated exchange maps using the proposed approach (i.e., optimized bSSFP scan design along with PERK estimation). To create ground truth maps of all tissue parameters, we assigned typical values from literature for white matter (WM) voxels [166]. The simulated bulk off-resonance varies linearly from left to right, while the mean residence time of myelin (i.e., inverse of exchange rate) varies along the anterior-posterior direction. Grey matter (GM) regions are masked out since they have low myelin content and exchange estimates in GM are potentially unreliable. The estimated myelin water exchange map for WM is quite close to the ground truth (RMSE of around 13.6%), and Fig. 3.8 illustrates this property using a scatter plot of exchange values in randomly selected WM voxels.

Optimized scan design vs. random baselines

How does our optimized scan design (in Section 3.4.3) compare against a baseline consisting of randomly chosen bSSFP acquisition parameters? Table 3.3 and Fig. 3.9 show the results of comparing the optimized scan design with three different random sets of bSSFP acquisitions. Each of the three random sets contain 40 different bSSFP scans, and the acquisition parameters were chosen as follows: T_R/T_E were chosen to be the same as the optimized scan design, while the flip angles were randomly chosen in the range $[10^\circ, 40^\circ]$ and the RF phase cycling factors were chosen randomly in the range $[-180^\circ, 180^\circ]$. Table 3.3 compares the predicted coefficients of variation (obtained using Cramér-Rao bound analysis), and the empirical RMSE values in white matter (obtained using BrainWeb simulations) for the optimized scan design vs. the three random sets of bSSFP scans. The optimized scan design (described in Section 3.4.3) has the best predicted coefficient of variation, as well as the lowest empirical RMSE in white matter (using BrainWeb simulations). These

| | Optimized design | Random set 1 | Random set 2 | Random set 3 |
|------------------------------------|-------------------------|--------------|--------------|--------------|
| Predicted coefficient of variation | 0.136 | 0.260 | 0.279 | 0.282 |
| Empirical RMSE (in BrainWeb) | 0.135 | 0.184 | 0.206 | 0.182 |

Table 3.3: Comparison of predicted coefficients of variation and empirical RMSE values in white matter (using BrainWeb simulations), for the optimized scan design and for three different random sets of bSSFP scans.

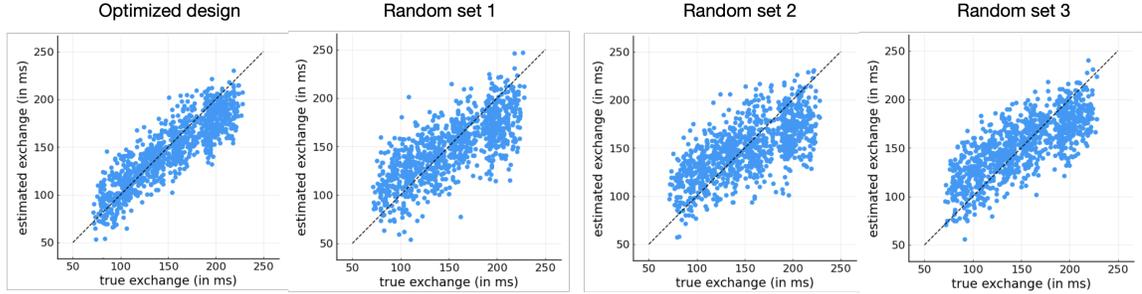


Figure 3.9: Comparison of scatter plots (using BrainWeb simulations) for the optimized scan design and for three different random sets of acquisition parameters for phase-cycled bSSFP scans, using PERK estimation. (Note: The scatter plots correspond to a different random subset of voxels than shown in Fig. 3.8.)

results illustrate the efficacy of our proposed scan design approach for estimating myelin water exchange, compared to random sets of scans (as baselines).

Fig. 3.9 shows the scatter plots of a randomly selected subset of voxels in white matter (from BrainWeb simulations) for all four sets of acquisitions. The true exchange values of mean residence time of myelin ($\tau_{f \rightarrow s}$) are plotted on the x -axis, while the estimated exchange values are plotted on the y -axis. We used PERK as the estimator for the results in both Table 3.3 and Fig. 3.9. From a visual inspection, the three random sets of bSSFP scans seem to have a wider spread in the estimated exchange values than the optimized scan design.

PERK vs. NLLS

This set of experiments compares the exchange estimates obtained using two estimation methods: PERK (Parameter Estimation via Regression with Kernels) and NLLS (Nonlinear Least Squares). We ran numerical simulations using a modified BrainWeb phantom, and estimated exchange values using both PERK and NLLS. Fig. 3.10 shows the estimated exchange values vs. ground truth, for a randomly selected subset of voxels in white matter (for both PERK and NLLS). Note that the PERK scatter plots in Figs. 3.8, 3.9 and 3.10 correspond to different random subsets of white matter voxels being plotted, as well as different random seeds used for training PERK.

In white matter, we obtained RMSE values of 13.5% with PERK estimation and 13.2% with NLLS estimation. The left plot in Fig. 3.10 shows the bias in PERK exchange estimates (since PERK is a Bayesian estimator); however, the speed of the PERK estimator (compared to the NLLS approach) makes it an attractive method for estimating exchange in this chapter. Running PERK estimation for the white matter voxels in a 2D slice of the

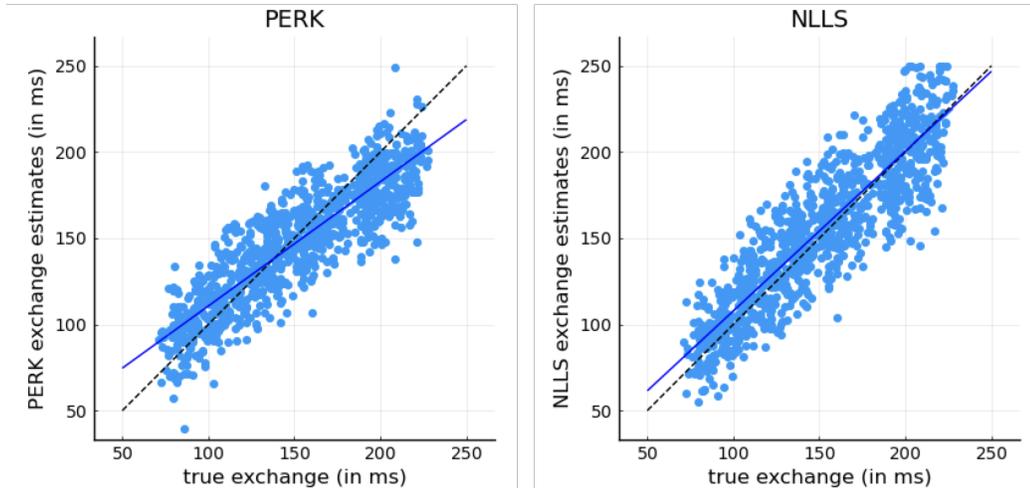


Figure 3.10: Comparison of exchange estimates obtained using PERK (left) and NLLS (right) estimation methods, for white matter (using BrainWeb simulations), vs. ground truth exchange values. The solid blue line indicates the linear least-squares fit, while the dashed black line is the identity line.

BrainWeb phantom of size 181×217 took around a minute. The NLLS estimation for the same slice had a compute time of around 8 hours. For the NLLS estimation, we ran 5 trials for each voxel and picked the trial with the lowest cost function for the exchange estimates. We parallelized the NLLS estimation across voxels using the multi-threading features of the Julia [9] programming language (with 25 threads). All experiments were performed on an Intel® Xeon® E5-2698 v4 processor (2.2GHz), using Julia (version 1.6.5). We use PERK estimation for our *in vivo* results in Section 3.4.5, since we typically have larger matrix sizes associated with *in vivo* data, and performing NLLS voxel-wise could be prohibitively slow.

3.4.5 In vivo results

We acquired *in vivo* data using the optimized set of phase-cycled bSSFP acquisitions on a GE 3T scanner. We used 3D acquisitions to avoid slice profile effects, and analyzed the central slice of the acquired data. The scan time for all 40 bSSFP acquisitions was 6 minutes and 40 seconds. We separately estimated B_1+ maps using Bloch-Siegert (BS) acquisitions [134]; the BS scans took 2 minutes and 23 seconds to acquire. Additionally, we acquired two Spoiled Gradient-Recalled Echo (SPGR) scans to estimate B_0 maps, which had a scan time of around 20 seconds. The two SPGR scans were acquired at echo times of 4 ms and 6.3 ms. The total scan time of our proposed protocol (implemented in TOPPE [116]) was

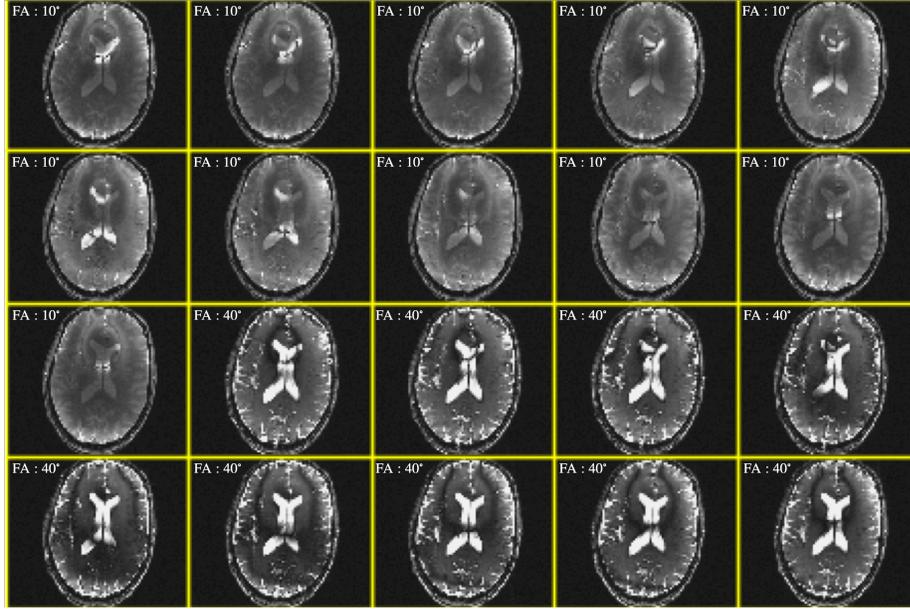


Figure 3.11: bSSFP images from the optimized scan design. (Note: Every other acquisition is shown out of 40 acquisitions.)

9 minutes and 23 seconds. All scans were acquired with a field-of-view (FOV) of $240 \times 200 \times 24 \text{ mm}^3$ and a matrix size of $192 \times 168 \times 8$.

To obtain the magnitude images required for exchange estimation, we took the inverse FFT of the fully sampled k-space data, followed by a square-root of sum-of-squares coil combination. Fig. 3.11 shows the images for every other bSSFP acquisition (out of 40 acquisitions), with the flip angles printed at the top-left of each image. The images are ordered in ascending order of RF phase cycling factors ϕ (as in Fig. 3.6). We observe that the locations of the dark bands (which are characteristic to bSSFP acquisitions) are shifted depending on the variations in ϕ .

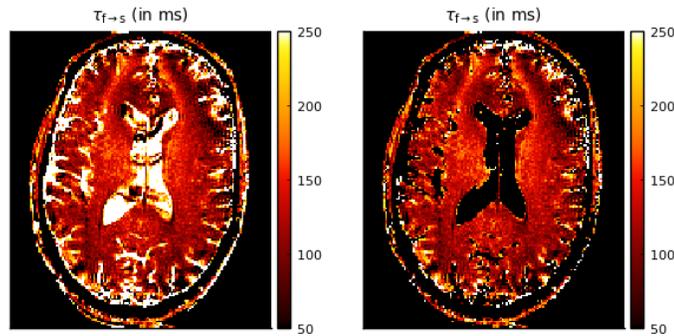


Figure 3.12: Estimated exchange map $\tau_{f \rightarrow s}$ without (left) and with (right) CSF masking.

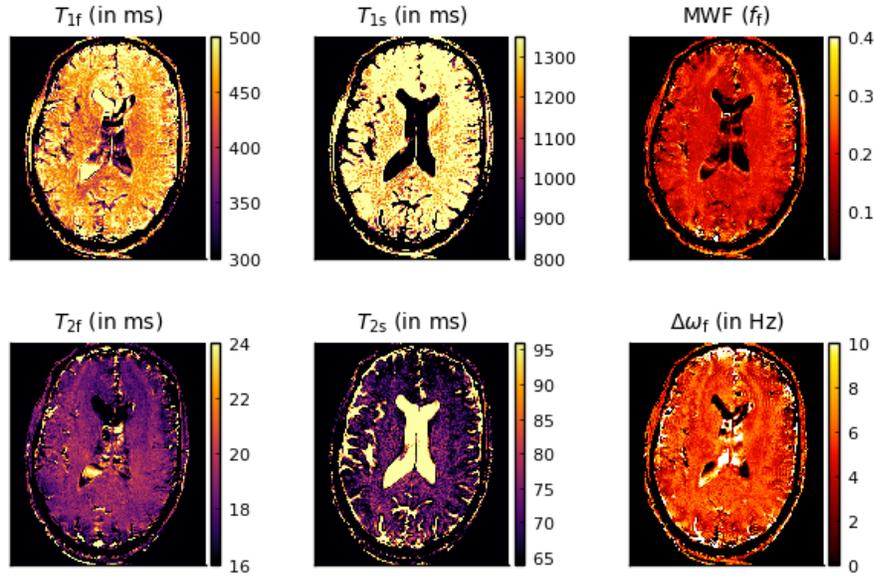


Figure 3.13: Parameter maps obtained using the optimized scan design, followed by PERK based estimation. (Note: The scan design was optimized specifically for estimating exchange, and not the other parameters of the two pool model.)

Fig. 3.12 shows the estimated map of the mean residence time of myelin water ($\tau_{f \rightarrow s}$), with and without masking of the CSF (cerebrospinal fluid) region. We can observe faint traces of the banding artefacts in the estimated exchange map, which suggests that the exchange map still contains some off-resonance effects. The values of the mean residence times seem to lie within reasonable estimates of exchange obtained in previous studies [3, 48]. However, as noted before, it is hard to interpret these exchange maps without validating the proposed approach through other methods (e.g., in vitro experiments using an aqueous urea system as described in Chapter 4).

Fig. 3.13 illustrates a few of the other estimated parameter maps. We did not mask the CSF region for these parameter maps; additionally, we set the colorbars to correspond to reasonable ranges for each parameter. Though the 40 bSSFP scans were optimized specifically for estimating exchange, it might still be informative to study the other parameter maps. We note that the precision in each map would potentially be better if we optimized the bSSFP acquisitions for that specific quantitative parameter.

3.4.6 Effect of off-resonance on exchange estimation

An important factor affecting the estimation of myelin water exchange could be off-resonance. Fig. 3.14(a) shows how exchange estimation is affected by bulk off-resonance.

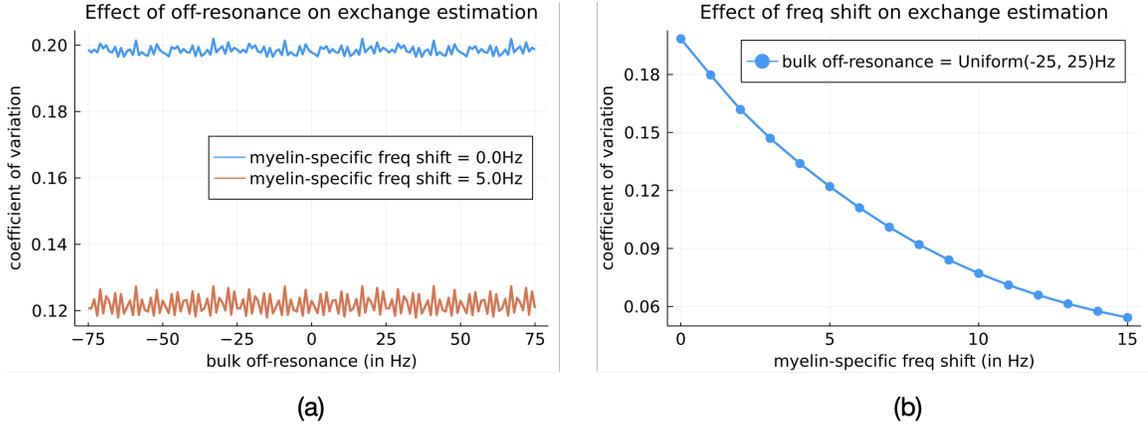


Figure 3.14: Effect of (a) bulk off-resonance and (b) myelin-specific frequency shift on predicted coefficient of variation for estimating myelin water exchange. (Note: We study a typical white matter voxel here.)

The predicted coefficient of variation (from the CRB) for estimating exchange remains relatively uniform with off-resonance. This result suggests that the proposed approach could potentially be robust to off-resonance and it would be interesting to explore this further. Since we use a fixed TR of ~ 20 ms, there is an inherent periodicity of ~ 50 Hz in the bSSFP off-resonance profiles, and it would be interesting to see if we can exploit this fact when setting the training distributions for PERK.

Fig. 3.14(b) shows how myelin water exchange estimation is affected by the myelin-specific frequency shift $\Delta\omega_f$. It has been shown that the orientation of nerve fibers in WM introduces susceptibility-related effects in myelin [165, 36]. This property is modeled as an additional frequency shift ($\Delta\omega_f$) that is specific to myelin water. The predicted coefficient of variation for exchange estimation worsens as the frequency shift goes to 0Hz. The orientation of WM fibers is out of our control when setting up experiments, so it is critical to design our acquisition parameters for a wide enough range of frequency shifts; we currently optimize the bSSFP scan design for a range of frequency shifts from 0 to 10Hz. A higher value of the myelin-specific shift seems to help in estimating exchange with higher precision.

3.4.7 Joint estimation of B_0 and B_1+ maps

For all scan designs and experiments above, we assumed B_0 and B_1+ information to be *known* parameters. In other words, we estimated B_0 and B_1+ maps using separate Spoiled Gradient-Recalled Echo (SPGR) and Bloch-Siegert (BS) acquisitions, respectively. In this subsection, we assume that B_0 and B_1+ maps are not known to us, and investigate whether

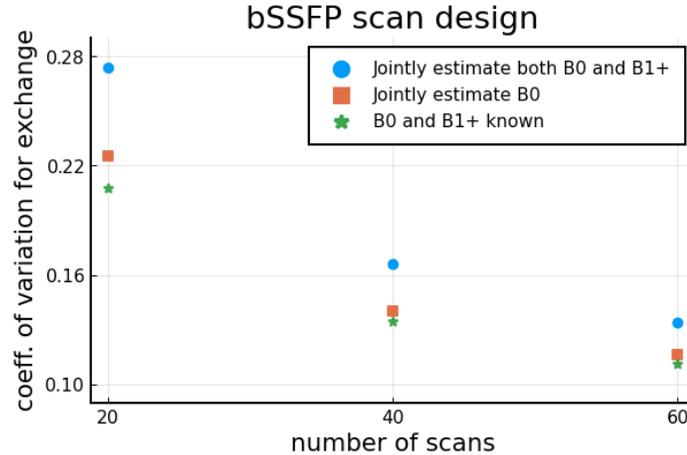


Figure 3.15: Coefficients of variation for scan designs optimized for jointly estimating B_0 and B_{1+} information along with exchange.

they can be jointly estimated along with exchange. If we obtain an optimized scan design with a low coefficient of variation, then we could potentially reduce acquisition time since we would not need to acquire SPGR and/or BS data separately.

Fig. 3.15 illustrates the coefficients of variation for scan designs that were optimized under three different conditions: (i) joint estimation of both B_0 and B_{1+} along with exchange, (ii) joint estimation of only B_0 along with exchange (assuming B_{1+} is known) and (iii) estimation of exchange assuming B_0 and B_{1+} are separately known. We ran three trials for each experimental condition and chose the scan design that yielded the lowest coefficient of variation among the three trials (to plot). We repeated the experiments, varying the number of bSSFP acquisitions (20, 40 and 60). We observe that the coefficient of variation for exchange is lowest when B_0 and B_{1+} maps are known, and this worsens when we move on to jointly estimating B_0 and B_{1+} along with exchange. It would be worth exploring further, using numerical simulations, how well these optimized scan designs are able to jointly estimate B_0 and B_{1+} maps for a modified BrainWeb phantom [23]. Fig. 3.15 also shows that, as expected, the coefficients of variation generally trend lower as we increase the number of bSSFP acquisitions. For a more apples-to-apples comparison, it could be worth repeating these experiments with more acquisitions for the case where we jointly estimate B_0 and B_{1+} (e.g., 22 acquisitions instead of 20), since we no longer need to acquire two spoiled gradient echo (SPGR) scans to estimate B_0 .

3.5 Discussion

This chapter focuses on estimating myelin water exchange in white matter in the brain. The proposed approach can be divided into two parts: (i) scan design, where we optimize the acquisition parameters of a set of MR acquisitions to help estimate myelin water exchange, and (ii) parameter estimation, where we use the PERK estimator to estimate myelin water exchange using our optimized scan design. For scan design, we optimize the flip angles and phase cycling factors of 40 different bSSFP acquisitions by minimizing a cost function based on the Cramér-Rao bound. From Fig. 3.15, the choice of 40 acquisitions seems to be a reasonable choice since we obtain a very low coefficient of variation for estimating myelin water exchange ($\sim 13\%$), along with a reasonable acquisition time (compared to higher number of acquisitions such as 60).

Starting from a randomly initialized scan design, it is intriguing that the optimized set of bSSFP acquisitions has a very regular sweep of RF phase cycling factors ϕ at both the low and high flip angles (see Fig. 3.6). This variation in ϕ appears to be quite important for exchange estimation; the predicted coefficient of variation was worse when ϕ was not varied as much. It could be interesting to explore whether optimizing the phase cycling factors might prove helpful in other quantitative MRI techniques that use bSSFP acquisitions. E.g., the mcDESPOT approach to multi-parametric estimation also uses a set of bSSFP scans with a fixed set of ϕ factors (along with SPGR scans) [29, 28]; a future direction could be to optimize this set of ϕ values to minimize the variance of the desired parameter(s) of interest (e.g., myelin water fraction).

In our experiments, we demonstrate the potential effectiveness of the optimized bSSFP scan design using numerical simulations (Section 3.4.4) and *in vivo* data (Section 3.4.5). However, validating the proposed approach *in vivo* is quite tricky since there is no established baseline method for *in vivo* exchange mapping (cf. inversion recovery acquisitions for T_1 estimation or multi-echo spin echo acquisitions for myelin water fraction). One possibility is to compare the exchange maps obtained using our optimized scan designs, with previous methods such as mcDESPOT [29] or MRF-X [48]. However, both those studies looked at exchange maps without necessarily validating against other established baselines/methods. These methods might be useful for comparison but would not help systematically validate our proposed approach for estimating exchange.

A future direction would be to design a phantom/system to validate our proposed techniques, by obtaining ground truth exchange estimates using other established methods. Spectroscopy-based methods provide a way of obtaining ground-truth values for exchange rates [97, 33, 32, 6, 168]. Relaxation exchange spectroscopy (REXSY) has been previously

used to characterize exchange in aqueous urea systems [33] as well as excised rat and frog samples [32]. Inversion recovery based NMR (IR-NMR) spectroscopy is another method that could be used to obtain ground truth exchange rates in small samples [168]. Chapter 4 designs an aqueous urea system to analyze exchange, and obtain ground truth exchange values using IR-NMR spectroscopy. We then use the ground truth exchange values to validate our proposed acquisition scheme using phase-cycled bSSFP protocols.

For our uncertainty analysis using the Cramér-Rao bound, we assume the acquisition of fully sampled data for every bSSFP scan. This assumption simplifies the process of parameter estimation since we can perform parameter fitting voxel-by-voxel separately. However, many MR sequences of interest in quantitative MRI lie in the highly under-sampled regime to accelerate data acquisition. This complicates uncertainty analysis since systematic errors (due to aliasing artefacts) tend to dominate over random errors in the highly undersampled regime [68]. A future direction would be to model these undersampling artefacts [71, 68] (instead of just bundling all sources of error into random Gaussian noise). For instance, in [71], aliasing artefacts are modeled through an additional Gaussian noise term, whose magnitude is signal-dependent. In [68], undersampling artefacts are modeled directly from first principles using computer simulations, and incorporated into a cost function for sequence design in magnetic resonance fingerprinting.

For translation to clinical applications, the long scan times of the fully sampled acquisitions discussed in this chapter could be a potential roadblock. A common approach to bring down the acquisition time is through the use of undersampling. Non-Cartesian acquisition schemes such as spiral trajectories are quite popular in undersampled acquisitions due to their efficient k-space coverage and robustness to motion. A future direction could be to explore undersampled spiral acquisitions to estimate myelin water exchange using phase-cycled bSSFP. For the repetition times discussed in this chapter ($T_R = 20\text{ms}$), we would potentially need to use spiral shots with longer readouts and this increases sensitivity to off-resonance effects, along with the need for more computationally intensive reconstruction approaches. Chapter 5 analyzes the off-resonance effects for spiral trajectories, and derives fast implementations of the NUFFT-based forward model for stack-of-spirals MRI. This would potentially be useful for estimating myelin water exchange in a 3D volume using undersampled bSSFP acquisitions with our numerically optimized scan design.

For all our analysis and experiments, we acquire only steady-state data during our bSSFP acquisitions. It seems intuitive that, to better estimate exchange values, we might need data in the transient regime. One option to sample the transient regime might be to collect data while we wait for the bSSFP acquisitions to reach steady-state [137, 170].

Another option could be to use fingerprinting-like sequences (instead of steady-state sequences) to drive the system into a persistent transient state [65, 48] and potentially encode more information about exchange. A future direction could be to perform a Cramér-Rao bound analysis assuming the acquisition of both transient and steady-state data.

3.6 Conclusion

This chapter optimized the acquisition parameters (flip angles and RF phase cycling factors) of a set of bSSFP acquisitions, to help estimate myelin water exchange in white matter (WM). We showed, using numerical simulations, that the optimized scan design resulted in WM exchange maps with low RMSE. Through additional experiments, we showed that optimizing a set of RF phase cycling factors ϕ led to improved exchange estimates; it would be interesting to investigate whether this would be beneficial for estimating other quantitative parameters such as myelin water fraction. We also obtained exchange maps *in vivo* but it is hard to interpret and validate them. The next step would be to validate the proposed approach using ground truth exchange values obtained using other established methods; Chapter 4 attempts to validate exchange maps in an aqueous urea system using phase-cycled bSSFP acquisitions (through *in vitro* experiments).

CHAPTER 4

Validating Exchange Maps in a Urea Water System using Phase-Cycled bSSFP

4.1 Introduction

Chapter 3 optimized bSSFP acquisitions for estimating myelin water exchange in white matter and demonstrated low coefficients of variation for estimating exchange rates using a modified BrainWeb [23] digital phantom. However, it is much harder to show the effectiveness of the proposed bSSFP acquisitions for estimating exchange *in vivo*, since it is difficult to obtain ground truth exchange maps *in vivo*. In this chapter, the objective is to design a two-compartment physical phantom with exchange, which could be used to validate the optimized bSSFP acquisitions. We aim to obtain ground truth exchange values using other established techniques and use this ground truth to validate the exchange maps obtained using our proposed acquisitions.

It is challenging to design a phantom that exactly mimics the exchange dynamics occurring in the myelin sheath in our brain, but we set out to design something that can act as a reasonable proxy. One of our collaborators, Dr. Scott Swanson, designed an aqueous urea system for us to use for validating our proposed methods for estimating exchange. We provide more details of the urea water phantom in Section 4.3.1. Such an aqueous urea system has been previously used to study chemical exchange [83, 159, 41, 33, 144]. Aqueous urea contains two compartments/pools (water protons and urea protons) that are chemically shifted by about 1.0 ppm (around 298Hz at 7T). This is not identical to the myelin system in our brain, since the *myelin water* pool possesses a much smaller frequency shift than the 1.0 ppm shift present in the aqueous urea system [165, 36]. Additionally, the source of the frequency shift in myelin water is thought to be the orientation of the nerve fibers in white matter [165], instead of a chemical shift as seen in the urea phantom. Despite these differences, the urea water phantom is a reasonable model for studying methods for quantifying exchange, and is a good first step to eventually demonstrating the effectiveness of our

proposed approach *in vivo*. Aqueous urea has previously been found to be a good system to study exchange, and Section 4.3.1 provides more details for why it is a good candidate to analyze exchange dynamics.

This chapter can be divided into two parts: the first part describes how we obtained ground truth values for exchange rates using inversion recovery based NMR (IR-NMR) [168] spectroscopy experiments, and the second part explains how we performed imaging experiments using our phase-cycled bSSFP acquisitions and validated the exchange maps against the IR-NMR ground truth data. For obtaining ground truth exchange values using IR-NMR, we selectively invert one of the two pools (urea protons or water protons), and sample the recovery of the longitudinal magnetization at different time points using NMR measurements, i.e., non-imaging MR measurements. In this selectively inverted urea water system, T_1 recovery is influenced by the exchange dynamics between the urea and water proton pools. Through parameter fitting of the NMR measurements, we obtain ground truth values for the mean residence time of the urea protons (i.e., the inverse of the exchange rate of protons going from urea to water). We compare these ground truth exchange values against the parameter estimates obtained using our phase-cycled bSSFP acquisitions (through imaging experiments), and demonstrate the effectiveness of our proposed approach for estimating exchange rates in an aqueous urea system.

The rest of the chapter is organized as follows: Section 4.2 contains relevant background for the rest of this chapter, while Section 4.3 describes the methods used for estimating exchange in our aqueous urea system. We provide details of our experiments in Section 4.4, while Sections 4.5 and 4.6 contain some discussions and conclusions of this chapter, along with several ideas for future work.

4.2 Background

This section provides relevant background information for the methods used in this chapter. Section 4.2.1 provides detailed background about the inversion recovery based NMR (IR-NMR) method for estimating exchange. We use this technique to obtain ground truth values for urea water exchange rates in this work. Section 4.2.2 gives a brief description of related work for estimating inter-compartmental exchange, such as relaxation exchange spectroscopy (REXS) and chemical exchange saturation transfer (CEST). Our proposed technique for estimating exchange involves phase-cycled bSSFP acquisitions, and Section 4.2.3 provides some background as well as some prior work with using phase-cycled bSSFP for quantitative MRI.

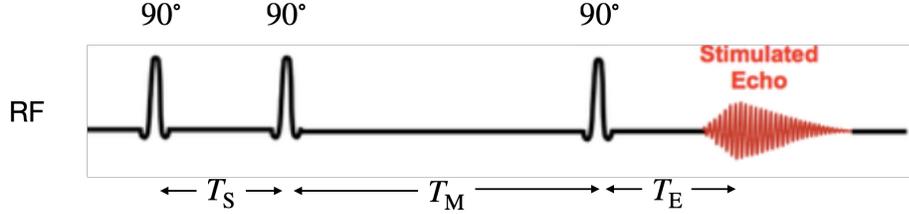


Figure 4.1: Inversion Recovery-based pulse sequence used for obtaining ground truth exchange values in the aqueous urea phantom. Note: T_S = time for selective inversion, T_M = mixing time, T_E = echo time.

4.2.1 Exchange estimation using inversion recovery based NMR (IR-NMR) spectroscopy

To obtain ground truth values for exchange rates in the urea water phantom, we use an inversion recovery based NMR (IR-NMR) spectroscopy experiment [168]. Fig. 4.1 shows the pulse sequence used for our IR experiments. This is not an imaging experiment, so we do not use imaging gradients and hence, NMR data is acquired from the whole sample. The first two 90° pulses are used to selectively invert one of the two pools (urea or water proton pools). The third 90° pulse is applied after a mixing time T_M (allowing the two pools to mix/interact) and the free induction decay (FID) signal is read out. We repeat the experiment for different values of T_M , thereby sampling the longitudinal magnetization recovery curve at different time points. The T_1 recovery of this selectively inverted system is influenced by the exchange dynamics between the urea protons and water protons resulting in a bi-exponential behavior; this forms the principle on which this exchange estimation method is based. We apply parameter fitting techniques to this bi-exponential curve and extract ground truth values for the exchange rates in the aqueous urea system.

To estimate exchange values, we could selectively invert either urea or water protons; without loss of generality, let us assume that the urea proton pool is selectively inverted. In our acquisition module, we choose our central frequency so that the urea protons are on resonance. The time for selective inversion T_S is chosen so that the protons in the water pool (which are off-resonance) precess through a phase of π between the two 90° pulses, i.e., they are out-of-phase with the urea protons. For our aqueous urea system, we set

$$T_S = \left(\frac{1}{2}\right) \left(\frac{1}{\Delta f_s}\right), \quad (4.1)$$

where Δf_s is the chemical shift in the urea water system ($\sim 298\text{Hz}$ at 7T). The water protons (which are off-resonance) complete one cycle of precession during a period of $1/\Delta f_s$, and come back in phase with the urea protons (which are on-resonance). We set

T_S to half this period, since we want the water protons to be out-of-phase with the urea protons, i.e., T_S is set to $\sim 1.68\text{ms}$ for our experiments.

When we set T_S to the value given in (4.1), the water protons are 180° out-of-phase with the urea protons. Then, the second 90° pulse selectively inverts the urea protons (i.e., urea magnetization lies along the negative z -axis), while the water protons are tipped back along the positive z -axis. Magnetization is now *stored* along the longitudinal axis (z -axis) and does not undergo T_2 decay (assuming ideal 90° RF pulses, exact values of T_S , and an infinitely narrow frequency distribution). T_1 recovery takes place during the period between the second and third 90° pulses, termed as mixing time (T_M). Initially, the evolution of longitudinal magnetization is influenced by the exchange dynamics between the urea and water protons (assuming that exchange occurs at faster timescales than simple T_1 recovery), before reaching an equilibrium corresponding to the T_1 recovery of both pools.

To model the evolution of longitudinal magnetization after applying the pulse sequence in Fig. 4.1, we use a two pool exchanging model. The longitudinal relaxation rates for the two pools are denoted by R_w and R_u (subscripts w and u denote the water protons and urea protons respectively). The exchange rate of magnetization out of the water pool is denoted by k_w , while the rate of exchange out of the urea pool is denoted by k_u . The mean residence times, denoted by τ_w and τ_u , are defined as the average time spent by a spin in the water and urea proton pools respectively; these are given by the inverse of the exchange rates.

The longitudinal magnetization dynamics in a two pool exchanging model are characterized by the Bloch-McConnell equation [98] as

$$\begin{aligned}\frac{dM_w(t)}{dt} &= (M_{0,w} - M_w(t))R_w - k_w M_w(t) + k_u M_u(t) \\ \frac{dM_u(t)}{dt} &= (M_{0,u} - M_u(t))R_u - k_u M_u(t) + k_w M_w(t),\end{aligned}\quad (4.2)$$

where $M_w(t)$ and $M_u(t)$ represent the longitudinal magnetization (at time t) in the water and urea proton pools respectively, while $M_{0,w}$ and $M_{0,u}$ represent the longitudinal magnetization in the two pools at thermal equilibrium. The longitudinal magnetization can also be written in the form of a fractional saturation level as

$$S_w(t) = 1 - \frac{M_w(t)}{M_{0,w}},\quad (4.3)$$

where $S_w(t)$ represents the fractional saturation level of the longitudinal magnetization of water protons. This takes a value of 2 for perfect inversion (i.e., $M_w(t) = -M_{0,w}$), 1 for full saturation (i.e., $M_w(t) = 0$) and 0 for full longitudinal recovery (i.e., $M_w(t) = M_{0,w}$) [162].

Similarly, $S_u(t) = 1 - \frac{M_u(t)}{M_{0,u}}$ represents the fractional saturation level of urea protons. It is important to perform the IR-NMR experiment for a long enough mixing time (T_M) so that we get a good estimate of the equilibrium magnetization values $M_{0,w}$ and $M_{0,u}$.

The longitudinal dynamics in (4.2) can be equivalently written in terms of $S_w(t)$ as

$$\begin{aligned}
\frac{dS_w(t)}{dt} &= \left(\frac{-1}{M_{0,w}} \right) \frac{dM_w(t)}{dt} \\
&= \frac{-1}{M_{0,w}} \left\{ (M_{0,w} - M_w(t))R_w - k_w M_w(t) + k_u M_u(t) \right\} \\
&= -R_w S_w(t) + k_w \frac{M_w(t)}{M_{0,w}} - k_u \frac{M_u(t)}{M_{0,w}} \\
&= -R_w S_w(t) + k_w \frac{M_w(t)}{M_{0,w}} - k_w \frac{M_u(t)}{M_{0,u}} \\
&= -R_w S_w(t) - k_w S_w(t) + k_w S_u(t),
\end{aligned} \tag{4.4}$$

where we have related the exchange rates between the water and urea pools as

$$\frac{k_u}{k_w} = \frac{M_{0,w}}{M_{0,u}} = \frac{1 - f_u}{f_u}. \tag{4.5}$$

Here, f_u is the urea pool fraction while $1 - f_u$ is the water pool fraction.

A similar analysis of longitudinal magnetization in urea protons results in

$$\frac{dS_u(t)}{dt} = -R_u S_u(t) - k_u S_u(t) + k_u S_w(t). \tag{4.6}$$

(4.4) and (4.6) can be combined and written in the form of a matrix ordinary differential equation (ODE) system as

$$\frac{d\mathbf{s}(t)}{dt} = \mathbf{A} \mathbf{s}(t), \tag{4.7}$$

where

$$\mathbf{A} = \begin{bmatrix} -R_w - k_w & k_w \\ k_u & -R_u - k_u \end{bmatrix} \quad \text{and} \quad \mathbf{s}(t) = \begin{bmatrix} S_w(t) \\ S_u(t) \end{bmatrix}. \tag{4.8}$$

The matrix ODE system in (4.7) has an analytical solution given by¹

$$\mathbf{s}(t) = e^{t\mathbf{A}} \mathbf{s}(0), \tag{4.9}$$

¹https://en.wikipedia.org/wiki/Matrix_differential_equation

where $e^{\mathbf{X}} = \sum_{k=0}^{\infty} \mathbf{X}^k/k!$ is called the matrix exponential of \mathbf{X} . Computing the matrix exponential of \mathbf{A} is easier when \mathbf{A} possesses certain properties. E.g., if \mathbf{A} is diagonalizable², then the matrix exponential in (4.9) can be computed in an efficient manner, by using the property that the matrix exponential of a diagonal matrix is the same as the element-wise exponentiation of its diagonal elements. More generally, the matrix exponential of a matrix can be also be computed using the Jordan canonical form³ of the matrix.

If we are able to diagonalize \mathbf{A} (given in (4.8)), i.e., if we compute the eigenvectors and eigenvalues of \mathbf{A} , then we can plug that into (4.9) and obtain a parametric representation describing the longitudinal magnetization recovery in our aqueous urea system. We derive the eigenvalues and eigenvectors of \mathbf{A} in the next section (Section 4.2.1.1), and demonstrate that the longitudinal magnetization recovery in the IR-NMR method can be described using bi-exponential curves.

4.2.1.1 Eigen analysis

Eigenvalues:

The eigenvalues of \mathbf{A} in (4.7) are computed as

$$\begin{aligned} & \left| \begin{bmatrix} -R_w - k_w - \lambda & k_w \\ k_u & -R_u - k_u - \lambda \end{bmatrix} \right| = 0 \\ & (-R_w - k_w - \lambda)(-R_u - k_u - \lambda) - k_w k_u = 0 \\ & \lambda^2 + \lambda(R_w + R_u + k_w + k_u) + (R_w R_u + R_u k_w + R_w k_u) = 0 \\ & \lambda = \frac{-(R_w + R_u + k_w + k_u) \pm \sqrt{(R_w + R_u + k_w + k_u)^2 - 4(R_w R_u + R_u k_w + R_w k_u)}}{2} \\ & \lambda = \frac{-(R_w + R_u + k_w + k_u) \pm \sqrt{(R_w - R_u + k_w - k_u)^2 + 4k_w k_u}}{2}, \end{aligned} \quad (4.10)$$

where $|\mathbf{X}|$ denotes the determinant of \mathbf{X} . To demonstrate that \mathbf{A} is diagonalizable, it is sufficient to show that the eigenvalues of \mathbf{A} are distinct. If both exchange rates k_w and k_u are positive, then the term under the square root in (4.10) is positive, and hence, the eigenvalues are distinct. Otherwise, if the exchange rates k_w and k_u are both 0, then the

²The matrix \mathbf{X} is diagonalizable if \mathbf{X} can be written as $\mathbf{X} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$, where $\mathbf{\Lambda}$ is a diagonal matrix containing the eigenvalues of \mathbf{X} and \mathbf{V} contains the linearly independent eigenvectors of \mathbf{X} .

³https://en.wikipedia.org/wiki/Matrix_exponential#Using_the_Jordan_canonical_form

eigenvalues are distinct as long as $R_w \neq R_u$. Under these conditions, the eigenvalues are distinct, and hence, \mathbf{A} is indeed diagonalizable.

We define the fast (λ_f) and slow (λ_s) relaxation rates for this longitudinal magnetization system as⁴

$$\lambda_{f,s} = \frac{(R_w + R_u + k_w + k_u) \pm \sqrt{(R_w - R_u + k_w - k_u)^2 + 4k_w k_u}}{2}. \quad (4.11)$$

Eigenvectors:

Let $\mathbf{v}_f = [v_{f,w} \ v_{f,u}]^T \in \mathbb{R}^2$ and $\mathbf{v}_s = [v_{s,w} \ v_{s,u}]^T \in \mathbb{R}^2$ be the eigenvectors corresponding to the fast and slow relaxation rates respectively. The eigenvector corresponding to the eigenvalue $-\lambda_f$ can be determined using $\mathbf{A}\mathbf{v}_f = (-\lambda_f)\mathbf{v}_f$, i.e.,

$$\begin{bmatrix} -R_w - k_w + \lambda_f & k_w \\ k_u & -R_u - k_u + \lambda_f \end{bmatrix} \begin{bmatrix} v_{f,w} \\ v_{f,u} \end{bmatrix} = \mathbf{0}.$$

Solving for \mathbf{v}_f and performing a similar analysis for \mathbf{v}_s , the eigenvectors can be written as (up to a scale factor)

$$\mathbf{v}_f = \begin{bmatrix} v_{f,w} \\ v_{f,u} \end{bmatrix} = \begin{bmatrix} k_w \\ R_w + k_w - \lambda_f \end{bmatrix}, \quad \mathbf{v}_s = \begin{bmatrix} v_{s,w} \\ v_{s,u} \end{bmatrix} = \begin{bmatrix} k_w \\ R_w + k_w - \lambda_s \end{bmatrix}. \quad (4.12)$$

The eigenvalues in (4.11) and eigenvectors in (4.12) can be written in matrix form as

$$\mathbf{\Lambda} = \begin{bmatrix} -\lambda_f & 0 \\ 0 & -\lambda_s \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} v_{f,w} & v_{s,w} \\ v_{f,u} & v_{s,u} \end{bmatrix}. \quad (4.13)$$

Now that we have the eigenvalues and eigenvectors of \mathbf{A} in (4.13), we can substitute $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ in (4.9) to obtain

$$\begin{aligned} \mathbf{s}(t) &= e^{t\mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}} \mathbf{s}(0) \\ &= \mathbf{V} e^{t\mathbf{\Lambda}} \mathbf{V}^{-1} \mathbf{s}(0) \\ &= \begin{bmatrix} v_{f,w} & v_{s,w} \\ v_{f,u} & v_{s,u} \end{bmatrix} \begin{bmatrix} e^{-\lambda_f t} & 0 \\ 0 & e^{-\lambda_s t} \end{bmatrix} \left(\frac{1}{|\mathbf{V}|} \right) \begin{bmatrix} v_{s,u} & -v_{s,w} \\ -v_{f,u} & v_{f,w} \end{bmatrix} \begin{bmatrix} S_w(0) \\ S_u(0) \end{bmatrix} \end{aligned}$$

⁴ λ_f and λ_s are defined by taking the negative of the eigenvalues given in (4.10).

$$\begin{aligned}
&= \left(\frac{1}{|\mathbf{V}|} \right) \begin{bmatrix} v_{f,w} & v_{s,w} \\ v_{f,u} & v_{s,u} \end{bmatrix} \begin{bmatrix} (v_{s,u}S_w(0) - v_{s,w}S_u(0)) e^{-\lambda_f t} \\ (-v_{f,u}S_w(0) + v_{f,w}S_u(0)) e^{-\lambda_s t} \end{bmatrix} \\
&= \left(\frac{1}{|\mathbf{V}|} \right) \begin{bmatrix} \{v_{f,w}v_{s,u}S_w(0) - v_{f,w}v_{s,w}S_u(0)\} e^{-\lambda_f t} + \\ \{-v_{s,w}v_{f,u}S_w(0) + v_{s,w}v_{f,w}S_u(0)\} e^{-\lambda_s t} \\ \{v_{f,u}v_{s,u}S_w(0) - v_{f,u}v_{s,w}S_u(0)\} e^{-\lambda_f t} + \\ \{-v_{s,u}v_{f,u}S_w(0) + v_{s,u}v_{f,w}S_u(0)\} e^{-\lambda_s t} \end{bmatrix} \\
\begin{bmatrix} S_w(t) \\ S_u(t) \end{bmatrix} &= \begin{bmatrix} c_{f,w}e^{-\lambda_f t} + c_{s,w}e^{-\lambda_s t} \\ c_{f,u}e^{-\lambda_f t} + c_{s,u}e^{-\lambda_s t} \end{bmatrix}. \tag{4.14}
\end{aligned}$$

From (4.14), the longitudinal magnetization recovery is characterized by bi-exponential behavior in the IR-NMR experiments. The coefficients for the bi-exponential fit for the water pool and urea pool are given by (see Appendix in Section 4.7 for more details)

$$\begin{aligned}
c_{f,w} &= \frac{(R_w + k_w - \lambda_s)S_w(0) - k_w S_u(0)}{\lambda_f - \lambda_s} \\
c_{s,w} &= \frac{-(R_w + k_w - \lambda_f)S_w(0) + k_w S_u(0)}{\lambda_f - \lambda_s} \\
c_{f,u} &= \frac{(R_w + k_w - \lambda_f) \{ (R_w + k_w - \lambda_s)S_w(0) - k_w S_u(0) \}}{k_w(\lambda_f - \lambda_s)} \\
c_{s,u} &= \frac{-(R_w + k_w - \lambda_s) \{ (R_w + k_w - \lambda_f)S_w(0) - k_w S_u(0) \}}{k_w(\lambda_f - \lambda_s)}. \tag{4.15}
\end{aligned}$$

4.2.2 Estimating inter-compartmental exchange: Related work

4.2.2.1 Relaxation exchange spectroscopy (REXSY)

Another technique to measure inter-compartmental exchange is relaxation exchange spectroscopy (REXSY) [82, 163]. REXSY consists of two Carr-Purcell-Meiboom-Gill (CPMG)⁵ sequences separated by a mixing period where the magnetization is stored along the longitudinal axis. By arraying the number of refocusing pulses in both CPMG pulse trains, and analyzing the resulting $T_2 - T_2$ decay data, we obtain a $T_2 - T_2$ spectrum. Spins

⁵CPMG is a spin-echo technique where the refocusing 180° pulses are applied along appropriately chosen axes so as to reduce the dephasing due to off-resonance inhomogeneities (i.e., B_0 effects) as well as imperfect 180° pulses (i.e., B_1+ effects). In CPMG sequences, we typically apply a 90° pulse along $+x$ -direction, followed by 180° pulses along $+y$ -direction.

that reside in the same pool during both CPMG trains show up as peaks along the diagonal of the $T_2 - T_2$ spectrum, while spins that undergo exchange during the mixing period appear as off-diagonal peaks.

At a given mixing time, the $T_2 - T_2$ spectrum helps us obtain a snapshot of the exchange dynamics in the system. This experiment needs to be repeated for various mixing times to estimate the inter-compartmental exchange rates, which makes it a very long experiment for estimating exchange. To speed up the acquisition, Dortch et al. proposed an inversion recovery (IR) magnetization preparation which reduces it to a shorter experiment [33], but still too long for *in vivo* whole-brain imaging.

REXSY-based methods have previously been used to study exchange in various systems. They have been used to study exchange dynamics in porous media [163, 99] and hydrated elastin [149]. They have also been used to analyze exchange dynamics in aqueous urea systems [33] (similar to the one used in this chapter). Additionally, REXSY-based methods have been applied to intact specimens as well, specifically in the study of inter-compartmental water exchange in freshly excised rat optic nerves and frog sciatic nerves [32].

4.2.2.2 Chemical exchange saturation transfer (CEST)

Chemical exchange saturation transfer (CEST) is a relatively recent MR technique that enables imaging certain compounds at concentrations that are too low to impact the contrast of standard MR imaging [173]. CEST relies on selective saturation of the chemical species of interest, and the indirect observation of this saturation effect on water protons. For this to be possible, the target compound must resonate at a different frequency than bulk water (termed as *chemical shift*), and must be capable of exchanging protons with the surrounding water protons (referred to as *chemical exchange*). In CEST, we selectively saturate the chemical species of interest using RF irradiation [186]. Due to continuous exchange of excited protons from the target compound with non-excited water protons via chemical exchange, the saturation of the target chemical species is transferred to water over time, which can be conveniently detected using standard MR imaging sequences [173].

CEST MRI is typically associated with long acquisition times, and this is one of the factors that potentially makes it harder for widespread use clinically. However, the long acquisition times are not problematic when it comes to *in vitro* experiments. CEST imaging has previously been used to study chemical exchange in systems such as aqueous urea solutions [144] as well as choline, glucose and glycogen solutions [180]. This potentially makes CEST MRI a good choice to analyze our designed urea water phantoms, and to obtain ground truth estimates of exchange rates. Though REXSY and CEST are potentially

useful baseline techniques to obtain ground truth exchange values, we chose to use the inversion recovery NMR (IR-NMR) method in this work, since it is easier to design a pulse sequence for (i.e., it is already a standard sequence on the scanner and requires minimal work setting it up), and is a good starting point for analyzing exchange.

4.2.3 Parameter estimation using phase-cycled bSSFP MRI

Balanced steady state free precession (bSSFP) is a popular pulse sequence due to its rapid acquisition time as well as high signal-to-noise ratio, along with its T_2/T_1 contrast [10]. However, it is highly susceptible to off-resonance effects which manifest themselves as dark banding artifacts in the images. One approach to reduce these banding artefacts is to reduce the T_R of the sequence, which decreases the phase accrual in each voxel and hence minimizes the banding artifacts. However, a shorter T_R potentially leads to lower maximum spatial resolution and might not be desirable in all cases. An alternate approach is to use RF phase cycling, where the phase of the excitation RF pulses is linearly incremented along the pulse train. By acquiring multiple images with different RF phase increments (called *RF phase cycling factors*), it has been shown that the resulting bSSFP images can be combined to minimize the off-resonance artifacts [7].

Along with off-resonance correction, phase-cycled bSSFP acquisitions have also been used for parameter estimation in quantitative MRI. DESPOT2 is one such technique that was initially introduced for T_2 estimation using multiple bSSFP acquisitions [30], and was later improved to account for off-resonance effects using phase-cycling [27]. In another work [12], the authors used phase-cycled bSSFP acquisitions to estimate the unknown off-resonance by linearizing the nonlinear bSSFP signal model, and studied the feasibility of jointly estimating T_1 and T_2 from phase-cycled bSSFP images. Another group of methods uses elliptical models for the complex-valued phase-cycled bSSFP signal [174, 57], and employs parameter fitting techniques for multi-parameter estimation (T_1 , T_2 , off-resonance etc.) [140, 141, 75]. E.g., PLANET is a method that first fits an ellipse to the acquired set of phase-cycled bSSFP images, and then estimates parameters such as T_1 and T_2 voxel-by-voxel, using the geometric properties of an ellipse [140]. There are also purely data-driven approaches for multi-parameter estimation using phase-cycled bSSFP; in [54], the authors use artificial neural networks to estimate T_1 and T_2 maps. There has also been work exploring the benefits of varying RF phase in MR fingerprinting using bSSFP acquisitions [160].

The previous paragraph described some methods that use phase-cycled bSSFP acquisitions for parameter estimation assuming a *single* pool/compartment of MR spins within

each voxel. There has also been relevant work developing similar methods for estimating parameters assuming the presence of *multiple* tissue components in each voxel. One group of methods analyzes asymmetries observed in the bSSFP frequency profiles (in brain tissue, muscle etc.) [101, 102] to estimate parameters in a multi-compartment setting. E.g., SPARCQ is a novel approach to quantify fat fraction in tissue, using multiple bSSFP acquisitions [133]. SPARCQ uses phase cycling to obtain bSSFP frequency profiles, and uses multi-compartment dictionary matching to generate quantitative maps of water and fat fractions. Similar methods have also been explored to analyze chemical exchange processes in a multi-compartment setting (e.g., using *in vitro* experiments in CEST imaging [180, 55]).

4.3 Materials and Methods

4.3.1 Urea phantom studies

Aqueous urea has been used as a system to study chemical exchange before [83, 159, 41, 33]. It contains two proton pools (urea and water) that are chemically shifted by around 1.0 ppm (~ 298 Hz at 7T). Aqueous urea has a few properties that make it useful to analyze exchange [33]: (i) it is biexponential in nature (urea protons have a shorter T_2 than water), (ii) urea has a high solubility in water (thus, we can create samples where 30% or more of the signal comes from urea protons), (iii) both urea and water relaxation rates can individually be manipulated using contrast reagents [59], (iv) exchange rates can be adjusted by changing temperature and pH [159], and (v) urea fraction can be known from the stoichiometry of the solution. All of these factors have contributed to the use of aqueous urea as a potential gold standard for evaluating methods to estimate exchange. This influenced our decision to design a urea water phantom for validating our exchange estimation techniques based on phase-cycled bSSFP.

4.3.1.1 Phantom preparation

We prepared the aqueous urea solution to contain a ratio of 10% / 90% for urea/water protons. The relaxation rates of urea and water proton were adjusted by adding 200 μM Gd-DTPA (Magnvist[®]; Berlex, Inc.) and 85 μM $MnCl_2$ as reagents. Fig. 4.2 shows the aqueous urea phantom used for our experiments.



Figure 4.2: Phantom containing urea water solution.

4.3.1.2 Scanner

For all our experiments in this chapter, we acquired data at bore temperature ($\sim 20^\circ\text{C}$), using an Agilent 7/310 (7.0T, 310mm bore) (Walnut Creek, CA) NMR/MRI scanner. This system is actively shielded, cryo-cooled and offers a 4-channel, phased array for reception and an additional modular transmitter channel for up to four transmit channels. Three gradient inserts with inner diameters 115 mm, 205 mm, and 305 mm, and strengths 60 mT/m, 120 mT/m, and 210 mT/m are available.

4.3.2 Ground truth using IR-NMR spectroscopy

Section 4.2.1 provides background information for the inversion recovery NMR (IR-NMR) experiment that we use to obtain ground truth exchange estimates. This section summarizes the main equations governing longitudinal magnetization recovery, and describes how we use them to obtain ground truth exchange estimates. The inversion recovery experiment gives us longitudinal magnetization data for both water and urea proton pools. From Section 4.2.1, we can model the longitudinal magnetization recovery using bi-exponential curves as

$$\begin{aligned} S_w(t) &= c_{f,w}e^{-\lambda_f t} + c_{s,w}e^{-\lambda_s t} \\ S_u(t) &= c_{f,u}e^{-\lambda_f t} + c_{s,u}e^{-\lambda_s t}, \end{aligned} \quad (4.16)$$

where

$$S_w(t) = 1 - \frac{M_w(t)}{M_{0,w}}$$

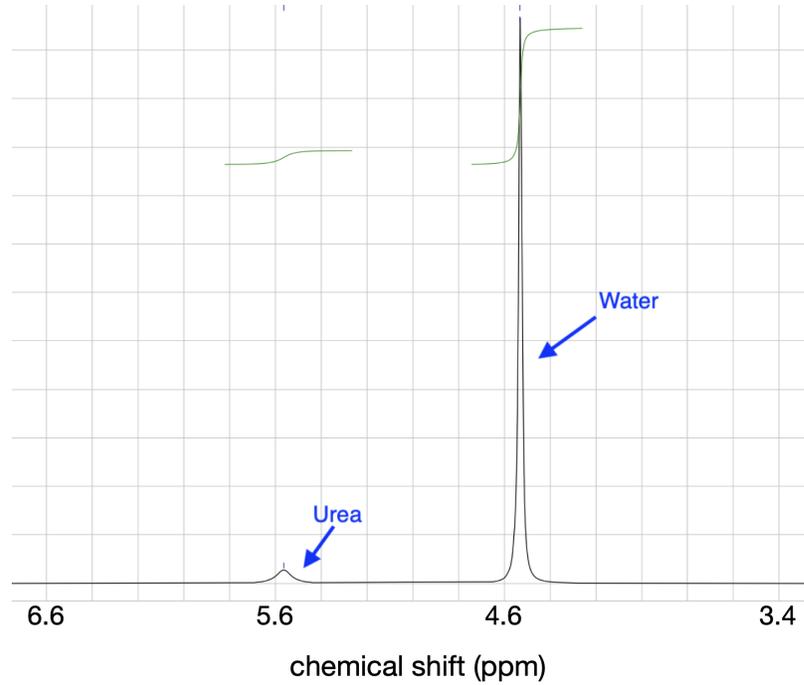


Figure 4.3: NMR spectrum with urea and water peaks for a particular mixing time T_M . Area under the peaks yields the signal values for urea and water proton pools at T_M .

$$S_u(t) = 1 - \frac{M_u(t)}{M_{0,u}} \quad (4.17)$$

are the fractional saturation levels of water and urea proton pools respectively. The fast/slow relaxation rates and coefficients are given in (4.11) and (4.15) respectively; Section 4.2.1 provides more details on how they are computed.

For our experiments to estimate the ground truth, we use an inversion recovery NMR (IR-NMR) experiment. As described earlier in Section 4.2, this consists of three 90° pulses, and the pulse sequence is shown in Fig. 4.1. The first two 90° pulses are timed appropriately to selectively invert one of the two pools (urea or water pools), while the third 90° pulse is applied after a *mixing time* T_M , allowing the two pools to exchange. In our experiments, we collected IR-NMR data at 21 T_M values, exponentially arrayed between 10ms and 10s. For each mixing time, we acquired FID (free induction decay) data consisting of 10,000 time points. We then took a 1-dimensional FFT of this to obtain a spectrum, and computed the areas under the urea and water peaks. This gives us the acquired longitudinal magnetization for each of the two pools (urea and water) as a function of mixing time. Fig. 4.3 shows an example NMR spectrum with urea and water peaks, at a given mixing time T_M .

We included a mixing time of 10s at the end to ensure full longitudinal recovery and to make sure that we have good estimates of the equilibrium magnetization values in urea

and water, i.e., $M_{0,w}$ and $M_{0,u}$. These values are used to convert the acquired longitudinal magnetization signal into fractional saturation levels, also called *reduced magnetization* (in a similar manner to (4.17)).

We denote the acquired IR-NMR data (after expressing it as a fractional saturation level) as $\mathbf{y} \in \mathbb{R}^{2N}$, where $N = 21$ is the number of mixing times. We have $2N$ measurements, since this is an NMR experiment, and we have N measurements each, for the urea and water proton pools. We use nonlinear least squares (NLLS) fitting to fit the acquired data using bi-exponential curves as

$$\hat{\mathbf{x}} = \underset{\mathbf{x} \neq \mathbf{0}}{\operatorname{argmin}} \|\mathbf{y} - f(\mathbf{x})\|^2, \quad (4.18)$$

where $\hat{\mathbf{x}} = [R_w, R_u, k_w, k_u]$ collects the unknown parameters to be estimated, and $f(\mathbf{x}) \in \mathbb{R}^{2N}$ captures the bi-exponential signal model corresponding to \mathbf{x} (see (4.16)), evaluated at the desired mixing times. We performed a Cramer-Rao Bound (CRB) analysis and found that we obtained lower bounds when we assume the longitudinal relaxation rates of both pools to be the same. Hence, for the fitting of our NMR experiments in (4.18), we assume $R_w = R_u$, which reduces the number of unknowns to 3, i.e., R_w , k_w and k_u . For optimizing the NLLS cost function in (4.18), we use a box-constrained version of the BFGS method using the Optim.jl package in Julia [106]. The BFGS (Broydon-Fletcher-Goldfarb-Shanno) method is a very popular quasi-Newton algorithm used in numerical optimization [120].

The value of k_u provides the ground truth estimate of the exchange rate from the urea pool to the water pool, and this is used for validation of our proposed bSSFP acquisitions for estimating exchange. Additionally, we compute the urea pool fraction f_u from these estimated exchange rates as follows [162]:

$$f_u = \frac{k_w}{k_w + k_u}. \quad (4.19)$$

4.3.3 Estimating exchange: phase-cycled bSSFP imaging

Section 4.3.2 described how we compute ground truth estimates of exchange parameters using inversion recovery NMR experiments. This section contains details of our imaging experiments using phase-cycled bSSFP acquisitions, and describes how we validate the obtained exchange estimates against the ground truth values.

4.3.3.1 Hand-crafted phase-cycled bSSFP acquisitions

Chapter 3 studied the problem of estimating myelin water exchange in the brain, and optimized scan designs to help estimate exchange in white matter. Since our objective in this dissertation is to eventually estimate myelin water exchange *in vivo*, where scan time is an important consideration, we optimized our scans to help encode information about exchange in *as few acquisitions as possible*. This chapter studies exchange in an aqueous urea system using a urea water phantom, where the requirements for scan time are not as stringent as *in vivo* scans. Thus, we did not explore *optimized* scan designs for estimating urea water exchange, and instead, investigated the use of *handcrafted* acquisitions, with (potentially) sufficient diversity in flip angles and RF phase cycling factors. We aim to study how well these *handcrafted* sets of bSSFP acquisitions encode information about exchange in our simple urea water phantom, and help us estimate urea water exchange.

For our handcrafted bSSFP acquisitions in this chapter, we consider two sets of acquisitions (S1) and (S2) described below:

- (S1) - Set of 91 bSSFP acquisitions, each with a fixed flip angle of 30° . The phase cycling factors are linearly spaced between -180° and 180° with 4° increments, resulting in 91 different acquisitions.
- (S2) - Set of 69 bSSFP acquisitions. We have 3 different phase cycling factors (0° , 50° and 180°), and for every phase factor, we have 23 different flip angles from 1° to 80° (see Fig. 4.4 for more details about how the flip angles are spaced). To account for SAR limitations at higher flip angles, we finely sampled flip angles at lower angles (up to 40°) and coarsely sampled the higher range between 40° and 80° .

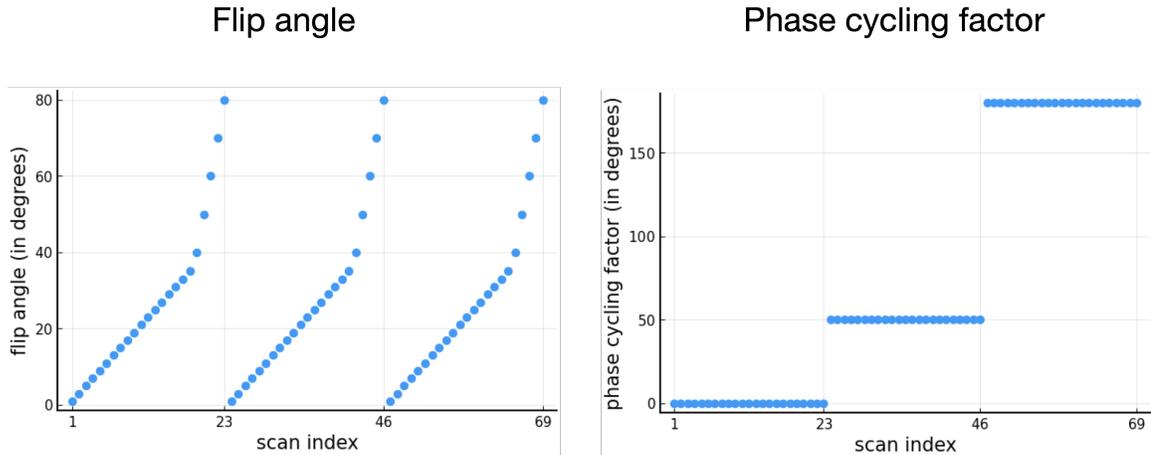


Figure 4.4: Handcrafted set of bSSFP acquisitions (S2), showing flip angles and phase cycling factors of 69 different acquisitions.

Both (S1) and (S2) were handcrafted with different intentions in mind. Set (S1) has a fixed flip angle and linearly incremented phase cycling factors. Such a linearly spaced phase-cycled bSSFP scheme samples the bSSFP frequency profiles of our aqueous urea system, and can be very helpful in visualizing the asymmetries in both magnitude and complex-valued data of our phase-cycled bSSFP acquisitions [55, 133].

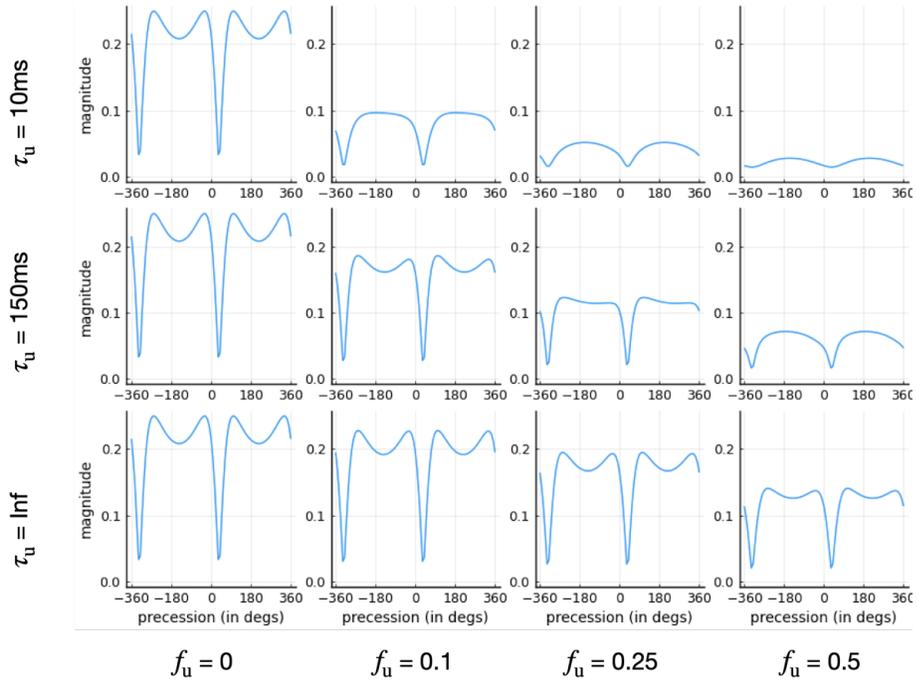
Fig. 4.5 shows the results of simulating the set of bSSFP acquisitions in (S1) for a single voxel in our aqueous urea system, both in the magnitude as well as complex domains. To study how sensitive the bSSFP signal model is to the exchange dynamics in the aqueous urea system, we vary the urea fraction f_u as [0, 0.1, 0.25, 0.5], and the mean residence time of urea τ_u as [10, 150, Inf] ms. This corresponds to urea exchange rates of [100, 6.67, 0] s⁻¹, representing three exchange regimes: fast, slow and zero exchange. For these simulations, we set the T_1 and T_2 values to be [T_{1u} , T_{1w}] = [700, 800] ms, and [T_{2u} , T_{2w}] = [10, 200] ms, where the subscripts u and w represent the urea and water pools respectively. We set the equilibrium magnetization to be $M_0 = 1.0$, where M_0 is related to the equilibrium magnetization of the urea and water pools as $M_{0,u} = f_u M_0$ and $M_{0,w} = (1 - f_u) M_0$. We model a chemical shift of 298 Hz between the urea and water proton pools (~ 1.0 ppm at 7T), and a bulk off-resonance of 15 Hz. We did not model any B_1 + inhomogeneities for these simulations, i.e., $\kappa = 1$. We set $T_R/T_E = 6.6\text{ms} / 3.3\text{ms}$ for the simulations in Fig. 4.5. Also, we extended the phase cycling factors in (S1) from -360° to 360° for this set of simulations, so that we could study two periods of the bSSFP signal profiles.

To actually encode information about the exchange processes in the aqueous urea system, we potentially need more variations in our acquisition parameters than (S1) provides. To account for this, we handcrafted the acquisitions in set (S2) to have variability in both flip angles and phase cycling factors. Fig. 4.4 shows the flip angles and phase cycling factors used in (S2).

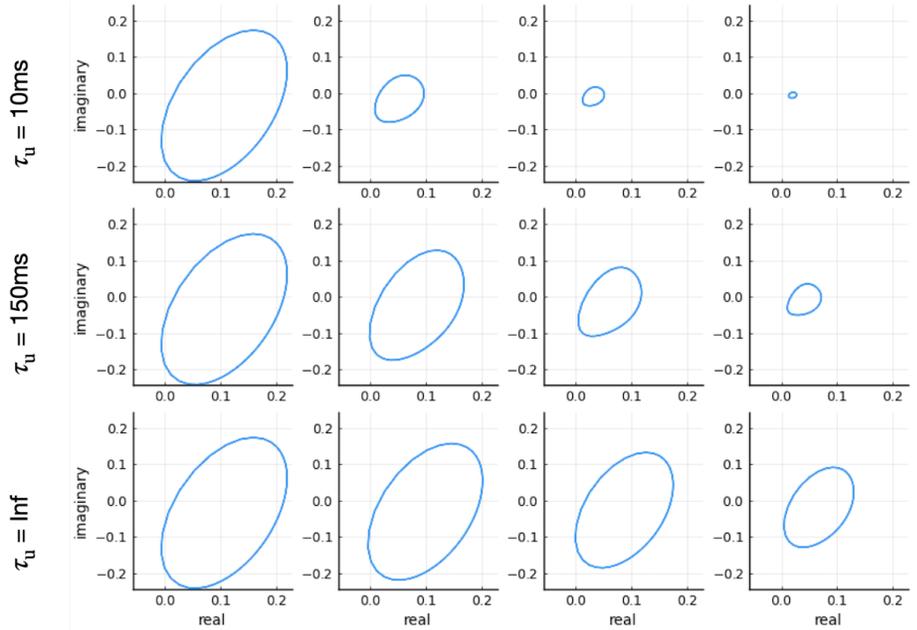
4.3.3.2 MR acquisition details

For our bSSFP imaging experiments, we acquired data on the same 7T Agilent (Varian) scanner used for our NMR experiments. We collected fully sampled Cartesian data for our handcrafted set of bSSFP acquisitions in (S2) (see Fig. 4.4 for more details of the acquisition parameters). We performed 2D slice-selective acquisitions for all 69 scans in (S2), with a matrix size of 64×64 , and set $T_R/T_E = 8\text{ms} / 4\text{ms}$.

We modified the bSSFP sequence on the Varian scanner to be able to manipulate the RF phase cycling factors, and passed in our flip angles and phase cycling factors in (S2) as parameter arrays. To ensure that steady-state is reached, we ran a total of 375 dummy T_{RS}



(a)



(b)

Figure 4.5: bSSFP off-resonance profiles simulated (using (S1)) for a single voxel in aqueous urea, showing (a) magnitude and (b) complex signal behavior for a range of values for urea fraction f_u and mean residence time of urea τ_u . (Note: $\tau_u = 1/k_u$)

at the beginning (corresponding to a time of 3s). At present, we run dummy acquisitions

only at the beginning of the parameter array corresponding to (S2), i.e., only before the first bSSFP acquisition. To potentially ensure that steady-state is reached for *every* acquisition, we would need to further modify the sequence so that the dummy acquisitions are run before *every* acquisition in (S2).

4.3.3.3 Nonlinear least squares (NLLS) data fitting

Chapter 3 compared the performance of two estimators: PERK (Parameter Estimation via Regression with Kernels) and NLLS, in the context of estimating myelin water exchange. Through numerical simulations, we found that NLLS produced less biased exchange estimates than PERK (see Fig. 3.10), albeit with a higher computational burden. For a phantom study such as the one investigated in this chapter, compute time is potentially not a major factor, and thus, we focused on NLLS-based estimation (due to its lower bias), in spite of its higher compute time. The rest of this section describes our NLLS fitting approach for our imaging experiments using phase-cycled bSSFP acquisitions.

We used a conventional inverse FFT reconstruction for our fully sampled Cartesian bSSFP acquisitions. Additionally, since we obtain fully sampled data, we do not need to model any undersampling artefacts, and we can perform parameter estimation *independently for each voxel*. For each voxel, we represent the complex-valued reconstructed data as $\mathbf{y} \in \mathbb{C}^N$, where $N = 69$ for the set of bSSFP acquisitions in (S2). We use nonlinear least squares (NLLS) fitting to estimate the unknown parameters as

$$\hat{\mathbf{x}} = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{s}(\mathbf{x}; \mathbf{P})\|^2, \quad (4.20)$$

where $\hat{\mathbf{x}} = [M_0, M_\phi, f_u, T_{1u}, T_{1w}, T_{2u}, T_{2w}, \tau_u, \Delta f_{cs}, \Delta f, \kappa]$ represents the unknown parameters to be estimated for the given voxel. T_{1u}, T_{1w}, T_{2u} and T_{2w} represent the T_1 and T_2 values for the urea and water pools (subscripts u and w respectively). Similar to our NMR experiments, we assume that the T_1 values of both the urea and water pools are the same (denoted by T_1), and that reduces the number of estimated parameters by one. The urea fraction and mean residence time of urea are given by f_u and τ_u respectively. Δf_{cs} stands for the chemical shift between the two pools, while Δf stands for bulk off-resonance (B_0 inhomogeneities). The equilibrium magnetization is given by a real-valued scalar M_0 , while M_ϕ models an additional phase term since we are dealing with complex-valued data. B_1+ effects are modeled with the scaling factor κ . $\mathbf{s}(\mathbf{x}; \mathbf{P}) \in \mathbb{C}^N$ represents the bSSFP signal model that takes the parameters in \mathbf{x} as input, and outputs bSSFP signal values corresponding to the set of acquisitions in (S2). \mathbf{P} collects the acquisition parameters of all scans in (S2).

| Parameter | Lower constraint | Upper constraint |
|-------------------------|------------------|------------------|
| M_0 | 0.1 | 10 |
| M_ϕ | $-\pi$ | π |
| f_u | 0.04 | 0.40 |
| T_1 (in ms) | 500 | 1000 |
| T_{2u} (in ms) | 5 | 30 |
| T_{2w} (in ms) | 100 | 400 |
| τ_u (in ms) | 30 | 150 |
| Δf_{cs} (in Hz) | 250 | 350 |
| Δf (in Hz) | -25 | 15 |
| κ | 0.9 | 1.1 |

Table 4.1: Lower and upper box constraints for parameters to be estimated in the aqueous urea system (to optimize the nonlinear least squares cost function in (4.20)). (Note: We assume that the T_1 values of both the urea and water pools are the same, and denote it by the parameter T_1 .)

To optimize the NLLS cost function in (4.20), we use a box-constrained version of the BFGS method [120] using the Optim.jl package in Julia [106]. Table 4.1 gives the lower and upper ends of the box constraints for all parameters. For each voxel, we optimize the cost function in (4.20) using 5 different initializations (randomly initialized within the box constraints in Table 4.1), and pick the minimizer that results in the lowest cost function. We combine the results from every voxel in the region of interest (ROI) and obtain parameter maps.

4.4 Results

4.4.1 Ground truth using spectroscopic data

To obtain ground truth estimates of exchange rates, we performed inversion recovery NMR (IR-NMR) experiments as described in Section 4.3.2. The central frequency is chosen so that the urea protons are on resonance; hence, the urea proton pool is selectively inverted in our NMR experiments. For each mixing time, we acquired FID (free induction decay) data consisting of 10,000 time points. We then performed a 1-dimensional FFT to obtain a spectrum, and computed the areas under the water and urea peaks, resulting in signal magnitudes for water and urea pools at that particular mixing time. Repeating this for every mixing time provided us with two curves for the acquired data: one corresponding to the urea signal vs. mixing time, and the other corresponding to the water signal vs. mixing time.

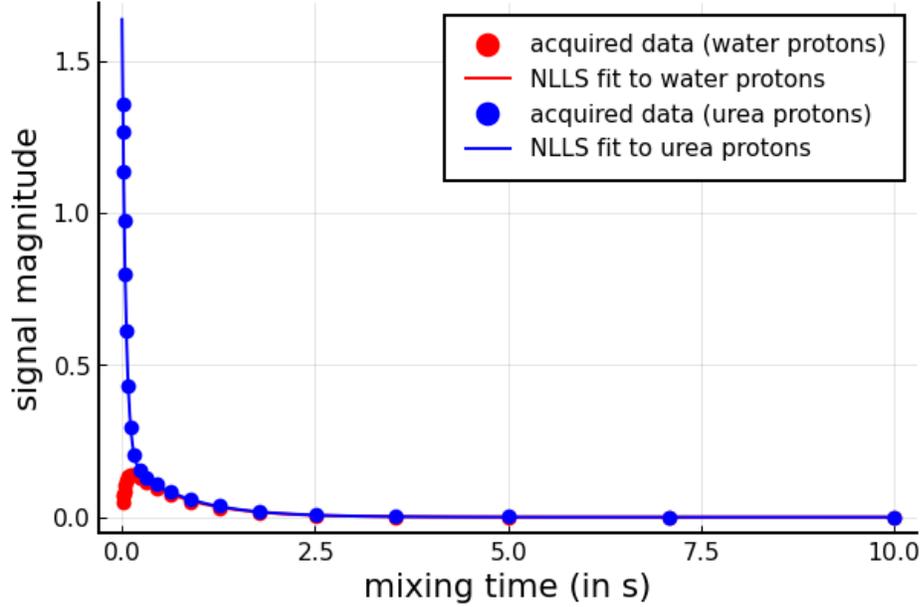


Figure 4.6: Nonlinear least squares fit to inversion recovery NMR (IR-NMR) data, using a bi-exponential model. (Urea protons are selectively inverted here.)

The acquired data is normalized using the equilibrium magnetization values, resulting in *reduced magnetization* data as described in Section 4.3.2. Fig. 4.6 shows the reduced magnetization signal values for both the urea and water pools, as a function of mixing time. The initial reduced magnetization data would be 2 for perfect inversion, 1 for full saturation and approach 0 for recovery to equilibrium.

As we see from the acquired data in Fig. 4.6, the urea pool is selectively inverted; hence, the urea signal starts from a value closer to 2. The water protons are tipped back along the positive z -axis by the third 90° pulse in our IR-NMR pulse sequence (see 4.1 for more details), and the water signal starts close to 0 in Fig. 4.6. As both pools recover to equilibrium, the exchange dynamics between the two pools results in the bi-exponential behavior that we see in the acquired data.

To obtain ground truth values for exchange, we performed nonlinear least squares (NLLS) fitting (with bi-exponential curves) as described in Section 4.3.2, for the cost function shown in (4.18). Fig. 4.6 shows the resulting NLLS fit; we obtained a very good fit with the acquired data. The final values for the optimization variables were: $[R_w, R_u, k_w, k_u] = [1.31\text{s}^{-1}, 1.31\text{s}^{-1}, 1.93\text{s}^{-1}, 17.9\text{s}^{-1}]$. Taking the inverse, we obtain the T_1 values and the mean residence times of the two pools as $[T_{1w}, T_{1u}, \tau_w, \tau_u] = [764\text{ms}, 764\text{ms}, 517\text{ms}, 55.9\text{ms}]$. The urea fraction is computed as $f_u = k_w / (k_w + k_u) = 0.097$, which is close to the fraction (0.10) expected from the stoichiometry of the aqueous urea solution. Hence, the ground truth values for the urea

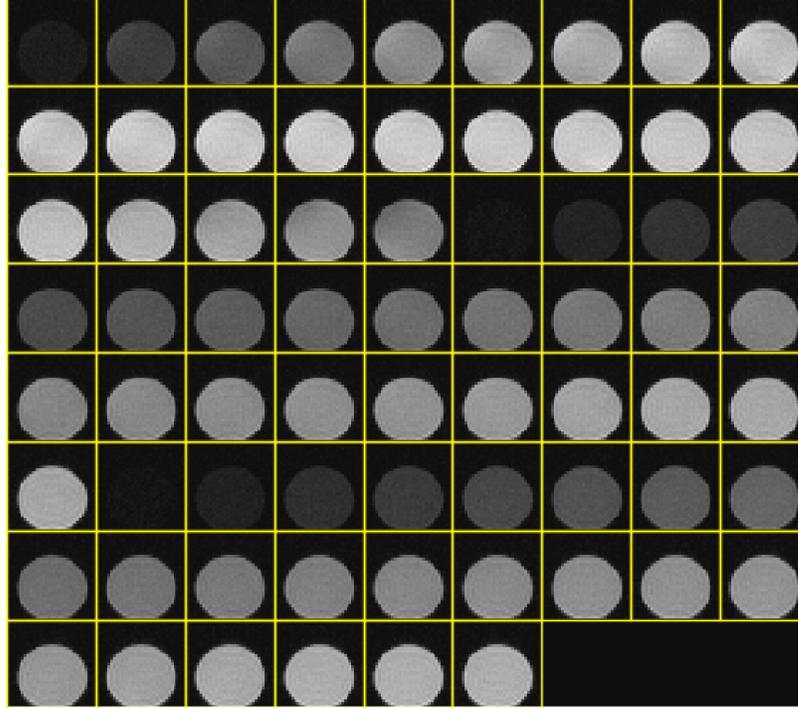


Figure 4.7: Magnitude images for all 69 bSSFP acquisitions in set (S2).

fraction and mean residence time of urea are **0.097** and **55.9 ms** respectively. We will use these ground truth values to validate the proposed approach based on phase-cycled bSSFP acquisitions in the next set of imaging experiments.

4.4.2 Parameter estimation using phase-cycled bSSFP acquisitions

This section describes our imaging experiments to estimate urea water exchange, using phase-cycled bSSFP acquisitions. Before performing our imaging experiments, we shimmed the magnetic field in our urea phantom, since it is desirable to obtain as uniform a magnetic field as possible. For our imaging experiments, we acquired 2D slice-selective fully sampled data for all 69 bSSFP acquisitions in set (S2) (see Fig. 4.4 for more details of the flip angles and phase cycling factors in (S2)). Fig. 4.7 shows the magnitude images for all acquisitions. We do not notice any significant dark bands that are typically associated with bSSFP images, suggesting that the B_0 field is well-shimmed.

For each voxel, we perform nonlinear least squares (NLLS) fitting on the complex-valued data and obtain estimates for the urea fraction f_u and mean residence time of urea τ_u , along with other parameters (see Section 4.3.3.3 for more details of the parameter fitting approach). Fig. 4.8 shows the estimated parameter maps of fraction f_u and mean residence time τ_u (in ms). We thresholded the M_0 parameter map to obtain a mask corresponding

| | M_0 | f_u | T_1 | T_{2u} | T_{2w} | τ_u | Δf_{cs} | Δf | κ |
|-----------|---------------|------------------|-------------|---------------|-----------------|----------------|-----------------|-----------------|-------------|
| bSSFP | 4.9 ± 0.2 | 0.115 ± 0.01 | 500 ± 0 | 7.5 ± 0.6 | 398.2 ± 9.7 | 50.0 ± 5.2 | 288.3 ± 3.3 | -13.8 ± 1.1 | 0.9 ± 0 |
| Reference | - | 0.097 | - | - | - | 55.9 | - | - | - |

Table 4.2: Parameter estimates for the urea water phantom, using complex-valued bSSFP data. (Note: Means and standard deviations were computed for the ROI corresponding to the entire cross-section of the urea water phantom.)

to our ROI. We use this mask to compute our ROI statistics and to display the estimated parameter maps. Fig. 4.9 shows a few of the other estimated parameter maps: equilibrium magnetization M_0 , phase M_ϕ , chemical shift Δf_{cs} and bulk off-resonance Δf .

Table 4.2 shows the mean and standard deviations for all parameters in our ROI. We also show the reference values for the urea fraction f_u and mean residence time τ_u obtained

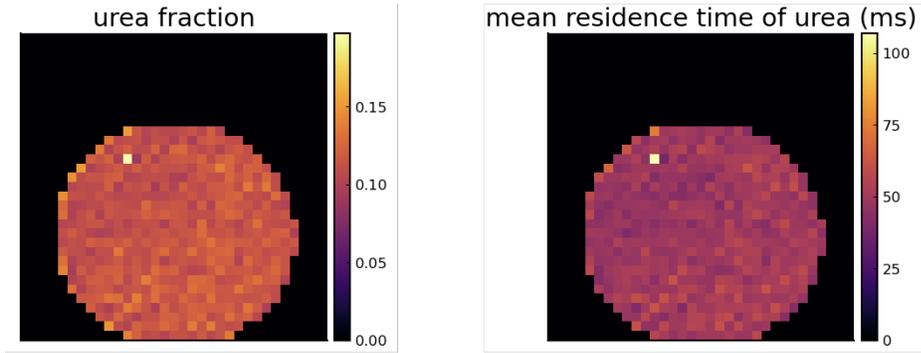


Figure 4.8: Estimated parameter maps of urea fraction f_u and mean residence time of urea τ_u .

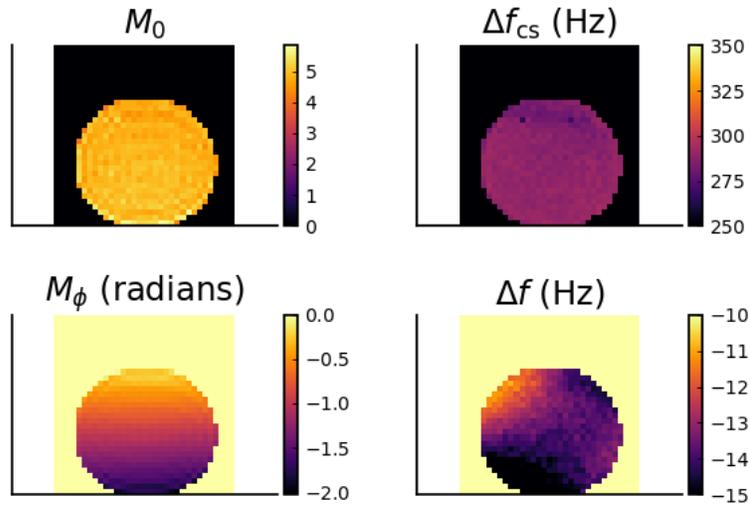


Figure 4.9: Estimated parameter maps of equilibrium magnetization M_0 , phase M_ϕ , chemical shift Δf_{cs} and bulk off-resonance Δf .

using our inversion recovery based NMR (IR-NMR) experiments (see Section 4.4.1 for more details about how we compute these reference values). The estimated urea fraction and mean residence times were in reasonable agreement with the reference values (within 19% and 11% respectively).

4.5 Discussion

In Chapter 3, we designed MRI experiments (using phase-cycled bSSFP acquisitions) for estimating myelin water exchange, but it was difficult to validate the obtained exchange maps *in vivo*, since we do not have an established gold standard technique to estimate exchange *in vivo*. Thus, we designed a urea water phantom to help validate our proposed methodologies for estimating exchange. Though the aqueous urea system is not identical to the myelin system that we hope to replicate, we feel that it is a reasonable proxy and a good starting point to validate our proposed scan designs and estimation methods.

To obtain ground truth values for urea water exchange, we used inversion recovery based NMR spectroscopy (IR-NMR) experiments [168]. We could selectively invert either the urea or the water protons to analyze our aqueous urea system; in this chapter, we have selectively inverted the urea protons, and obtained ground truth values for urea fraction and the urea exchange rate. It would be interesting to repeat our experiments by selectively inverting the *water* pool. This would still give us two curves (as in Fig. 4.6), but the water signal would begin near 2 (due to selective inversion of water), while the urea signal would begin near 0. Fitting this data using nonlinear least squares (NLLS) fitting should theoretically result in similar fraction and exchange rates, and it would be interesting to compare them with our ground truth values obtained by selectively inverting urea.

For estimating the ground truth exchange rates from the acquired IR-NMR spectroscopic data, we perform NLLS fitting on the *reduced magnetization* curves. The *reduced magnetization* vectors are obtained after normalizing the acquired longitudinal magnetization with respect to the equilibrium magnetization (as described in Section 4.3.2). For this approach to work, we need to sample reasonably large values of *mixing time* T_M to ensure that we wait long enough to reach equilibrium. For our experiments, we sample mixing times of up to 10s which should be sufficient to attain equilibrium. It is important to ensure that the equilibrium magnetization values be as accurate as possible. An alternative approach might be to fit the longitudinal magnetization curves directly instead of normalizing them to *reduced magnetization* vectors.

The NLLS fitting of the IR-NMR data yielded ground truth values of 0.097 for the urea fraction, and around 55.9ms for the mean residence time of urea. As a sanity check, the

reference value for the urea fraction (9.7%) is close to the urea fraction known from the stoichiometry of the solution [33], which is 10% for our phantom. This gives us some confidence in our ground truth exchange rates as well, since the urea fraction is reasonably close to what we expect from the contents of the urea solution. One direction for future work could be to run other established techniques to estimate exchange such as REXSY and CEST, and to compare the ground truth exchange rates with the value obtained here using inversion recovery NMR experiments (IR-NMR).

After obtaining ground truth values from the NMR data, we performed imaging experiments to validate our proposed method of estimating urea water exchange using a set of phase-cycled bSSFP acquisitions. Chapter 3 *optimized* the acquisition parameters of phase-cycled bSSFP acquisitions (using the Cramér-Rao bound) for estimating myelin water exchange. This chapter uses a *handcrafted* set of phase-cycled bSSFP acquisitions (see details of set (S2) in Fig. 4.4) to help estimate urea water exchange. As a direction of future work, it would be interesting to *optimize* bSSFP scans for estimating urea water exchange, and compare their performance against our *handcrafted* acquisitions.

This chapter uses nonlinear least squares (NLLS) based estimation for the imaging experiments (with phase-cycled bSSFP acquisitions) to estimate urea water exchange. Chapter 3 compared NLLS with PERK (Parameter Estimation via Regression with Kernels) for estimating myelin water exchange using numerical simulations, and found that NLLS was less biased than PERK, albeit with a higher computational runtime. We used NLLS in this chapter since it has lower bias, and compute time is not a major consideration for a phantom study, as compared to *in vivo* settings. As a future research direction, it would be desirable to repeat these imaging experiments with PERK as the estimator, and compare the performance of PERK vs NLLS in the context of estimating urea water exchange.

Using the bSSFP acquisitions in (S2), we estimated all parameters of our two pool exchanging urea water model and Table 4.2 shows the parameter estimates in our region of interest. The values of the urea fraction (0.115 ± 0.01) and mean residence time of urea ($50.0 \text{ ms} \pm 5.2 \text{ ms}$) were in reasonable agreement with the reference values (0.097 and 55.9 ms respectively). An extension of this work might be to repeat this experiment with aqueous urea samples at different *pH* values (resulting in different exchange rates [159]), and test whether the proposed method of estimating urea water exchange is generalizable across different samples and conditions.

From Table 4.2, we do not seem to attain good fits for a few of the parameters, e.g., T_1 and κ , which seem to have hit the boundaries of the box constraints (see Table 4.1 for a list of the constraints for optimization). This could be potentially attributed to a couple of sources of model mismatch. We acquired 2D data but the signal model in our fitting

method does not consider slice profile effects. Also, we acquire dummy TRs at the very beginning of the set of bSSFP acquisitions (S2), instead of taking dummy scans before *every* acquisition in (S2). A future direction would be to take 3D acquisitions, so that slice profile effects are no longer an issue, as well as acquire dummy scans before *every* bSSFP acquisition. This would potentially result in better parameter fits for our bSSFP imaging experiments. Also, for our inversion recovery NMR experiments to obtain ground truth values, it would be interesting to explore the effects of imperfect 90° pulses for selective inversion.

Finally, instead of using the *handcrafted* set of acquisitions in (S2), we could *optimize* them using the Cramér-Rao bound (as in Chapter 3), or in an end-to-end manner [167]. A direction for future research could be to compare the optimized scan design for urea water exchange, with the optimized set of acquisitions for myelin water exchange in Chapter 3. It would be interesting to see how similar or dissimilar the optimized scan designs are, for estimating exchange in different exchanging systems (e.g., myelin water system vs. aqueous urea system).

4.6 Conclusions

We designed an aqueous urea system to help us validate our phase-cycled bSSFP scan designs for estimating exchange. We estimated urea water exchange rates using imaging experiments with phase-cycled bSSFP acquisitions, and validated them against ground truth exchange rates obtained using inversion recovery-based NMR (IR-NMR) spectroscopy experiments. From the bSSFP data, the urea fraction f_u and the mean residence time of urea τ_u were found to be in reasonable agreement with the reference values obtained using NMR spectroscopy. Future work includes optimizing the bSSFP acquisitions to better estimate urea water exchange, as well as testing the generalizability of the proposed approach across various samples with different exchange rates.

Though the urea water phantom was designed initially with an aim of helping us specifically model myelin water exchange, it could be potentially useful for other exchange related experiments as well. We can perform time-consuming experiments such as REXSY (relaxation exchange spectroscopy) or CEST to obtain ground truth estimates of exchange, and use that to validate any proposed techniques to estimate exchange, before proceeding to *in vivo* data. A urea water phantom with multiple vials containing different concentrations of urea, as well as at different *pH* values (resulting in different exchange rates [159]), would be very useful as a gold standard for estimating fractions and exchange rates in a two-pool exchanging model.

4.7 Appendix

4.7.1 Eigen analysis of longitudinal magnetization dynamics

For a two pool exchanging model, the coefficients for the water protons in (4.14) are derived as

$$\begin{aligned}
 c_{f,w} &= \frac{(v_{f,w}v_{s,u}S_w(0) - v_{f,w}v_{s,w}S_u(0))}{|\mathbf{V}|} \\
 &= \frac{k_w(R_w + k_w - \lambda_s)S_w(0) - k_w(k_w)S_u(0)}{k_w(R_w + k_w - \lambda_s) - k_w(R_w + k_w - \lambda_f)} \\
 &= \frac{(R_w + k_w - \lambda_s)S_w(0) - k_w S_u(0)}{\lambda_f - \lambda_s},
 \end{aligned}$$

and

$$\begin{aligned}
 c_{s,w} &= \frac{(-v_{s,w}v_{f,u}S_w(0) + v_{s,w}v_{f,w}S_u(0))}{|\mathbf{V}|} \\
 &= \frac{-k_w(R_w + k_w - \lambda_f)S_w(0) + k_w k_w S_u(0)}{k_w(R_w + k_w - \lambda_s) - k_w(R_w + k_w - \lambda_f)} \\
 &= \frac{-(R_w + k_w - \lambda_f)S_w(0) + k_w S_u(0)}{\lambda_f - \lambda_s}.
 \end{aligned}$$

The coefficients for the urea protons in (4.14) are similarly derived as

$$\begin{aligned}
 c_{f,u} &= \frac{(v_{f,u}v_{s,u}S_w(0) - v_{f,u}v_{s,w}S_u(0))}{|\mathbf{V}|} \\
 &= \frac{(R_w + k_w - \lambda_f)(R_w + k_w - \lambda_s)S_w(0) - (R_w + k_w - \lambda_f)k_w S_u(0)}{k_w(R_w + k_w - \lambda_s) - k_w(R_w + k_w - \lambda_f)} \\
 &= \frac{(R_w + k_w - \lambda_f) \{(R_w + k_w - \lambda_s)S_w(0) - k_w S_u(0)\}}{k_w(\lambda_f - \lambda_s)},
 \end{aligned}$$

and

$$\begin{aligned}
 c_{s,u} &= \frac{(-v_{s,u}v_{f,u}S_w(0) + v_{s,u}v_{f,w}S_u(0))}{|\mathbf{V}|} \\
 &= \frac{-(R_w + k_w - \lambda_s)(R_w + k_w - \lambda_f)S_w(0) + (R_w + k_w - \lambda_s)k_w S_u(0)}{k_w(\lambda_f - \lambda_s)} \\
 &= \frac{-(R_w + k_w - \lambda_s) \{(R_w + k_w - \lambda_f)S_w(0) - k_w S_u(0)\}}{k_w(\lambda_f - \lambda_s)}.
 \end{aligned}$$

CHAPTER 5

Stack-of-Spirals MRI Reconstruction with Off-Resonance Correction

5.1 Introduction

Chapter 3 worked on MRI experiment design to estimate myelin water exchange in white matter, while Chapter 4 analyzed the exchange dynamics in aqueous urea as a means to validate our proposed methods for exchange estimation. In both chapters, we considered fully sampled acquisitions for our analyses. This simplifies the scan design and estimation problems, since we do not need to model undersampling artefacts and can perform parameter estimation *voxel-by-voxel*. However, fully sampled acquisitions are typically associated with longer scan times, which can lead to issues such as motion artefacts, patient discomfort, lower throughput etc.

The long scan times associated with fully sampled acquisitions are typically not problematic when we perform *in vitro* experiments since the sample is generally fixed in place, and we do not need to worry about long acquisition times. E.g., for our *in vitro* experiments with the urea water phantom in Chapter 4, acquisition time is not a huge factor, since the sample is fixed in the scanner. In fact, even a time-consuming technique such as REXSY (Relaxation Exchange Spectroscopy) could potentially be used to obtain ground truth exchange values for such *in vitro* experiments [33]. However, to translate these techniques for *in vivo* applications and to make them clinically viable, we need to ensure that the acquisition times are reasonably low. E.g., for the *in vivo* scans for estimating myelin water exchange in Chapter 3, the scan time for 40 fully sampled bSSFP acquisitions was ~ 7 minutes (for a $192 \times 168 \times 8$ volume). This could be prohibitively long in clinical settings, especially when we consider larger matrix sizes for whole-brain coverage, and we would need to reduce the acquisition time to make it clinically feasible.

A common approach to speed up MRI acquisitions is to acquire a subset of k-space data (instead of fully sampling all of k-space), which is called *undersampling*. There has been

a lot of work on developing acquisition and reconstruction methods for accelerating MRI scans through undersampling of k-space [130, 90]. To maximize the acceleration factors attained through undersampling, we often consider sampling trajectories that efficiently sample k-space. Non-Cartesian sampling trajectories provide a way to achieve efficient coverage of k-space, and in particular, spiral trajectories provide an efficient way to sample k-space in as few shots as possible (sometimes as few as a single shot [73]) [42].

Spiral trajectories have a lot of benefits such as efficient k-space sampling, robustness to motion (due to oversampling of the center of k-space), faster scan times etc. They have been used in quite a few quantitative MRI applications [65, 122, 115, 172, 185]. However, spiral trajectories are also associated with a few challenges such as sensitivity to off-resonance effects [14, 15] and the need for more complex reconstruction strategies [129]. During reconstruction, spiral trajectories require the non-uniform Fast Fourier Transform (NUFFT) [35, 40, 176] due to the non-Cartesian nature of the sampling, which is more computationally demanding than Cartesian acquisitions that can be reconstructed using the standard Fast Fourier Transform (FFT).

Some of the above challenges have been dealt with by incorporating both static and dynamic field inhomogeneities in the signal model [169], along with efficient iterative reconstruction algorithms for off-resonance correction [150, 39]. Off-resonance effects play a bigger role in the case of spiral shots with longer readouts. E.g., in Chapter 3, we have T_{RS} of around 20ms and if we use spiral trajectories to achieve undersampling, off-resonance effects would be significant at these readout durations and would potentially require correction.

For three-dimensional imaging, a common way to extend the benefits of 2D spiral trajectories is to acquire a stack of 2D spiral shots, referred to as a stack-of-spirals acquisition [64]. In a typical stack-of-spirals acquisition, we have phase encoding gradients along the through-plane direction, and for each phase encoding step, we acquire 2D spiral shots in-plane. The phase encoding direction along k_z in 3D is often referred to as *partition* encoding to distinguish it from the k_y encoding direction. Stack-of-spirals acquisitions have been used in a variety of MRI studies [185, 86, 142, 125, 146, 61].

For reconstructing stack-of-spirals data, the 3D NUFFT is a key block of the forward model. This is typically the computational bottleneck for iterative reconstruction, especially when modeling off-resonance effects (which is necessary for spirals of longer readout durations). This chapter focuses on efficient implementations for the NUFFT in the forward model for stack-of-spirals acquisitions, when off-resonance effects are considered. Section 5.2 provides some background for stack-of-spirals data and off-resonance correction in MRI. Section 5.3 analyses various configurations of the stack-of-spirals acquisition,

and derives efficient ways to implement the NUFFT in the forward model. Section 5.4 provides details of experiments, while Section 5.5 discusses the obtained results (along with a few ideas for future work). Section 5.6 talks about our conclusions from this chapter.

5.2 Background

5.2.1 Stack-of-spirals MRI

A stack-of-spirals acquisition is a popular 3D sampling trajectory used in non-Cartesian MRI [64]. As the name suggests, it consists of a stack of 2D spiral shots, acquired at specific k_z encodes in the through-plane direction. This involves the use of spiral gradients in-plane (k_x and k_y directions), and a Cartesian phase-encoding gradient along k_z . This trajectory differs from a more general 3D spiral sampling scheme, where each spiral shot may traverse the whole of 3D k -space, instead of being constrained to a specific k_z encoding plane. The stack could contain spiral shots of any kind, e.g., Archimedian spirals, variable density spirals (VDS) etc. [100, 42, 18]. Fig. 5.1 shows an example of a 2D VDS shot, and a stack-of-spirals acquisition scheme consisting of VDS shots.

The specific structure of a stack-of-spirals trajectory makes it potentially easier to design the sequence and acquire data, compared to a general 3D non-Cartesian trajectory, since we need to design just phase encoding gradient steps along k_z , instead of more complex gradient waveforms. Reconstructing stack-of-spirals data is also potentially easier

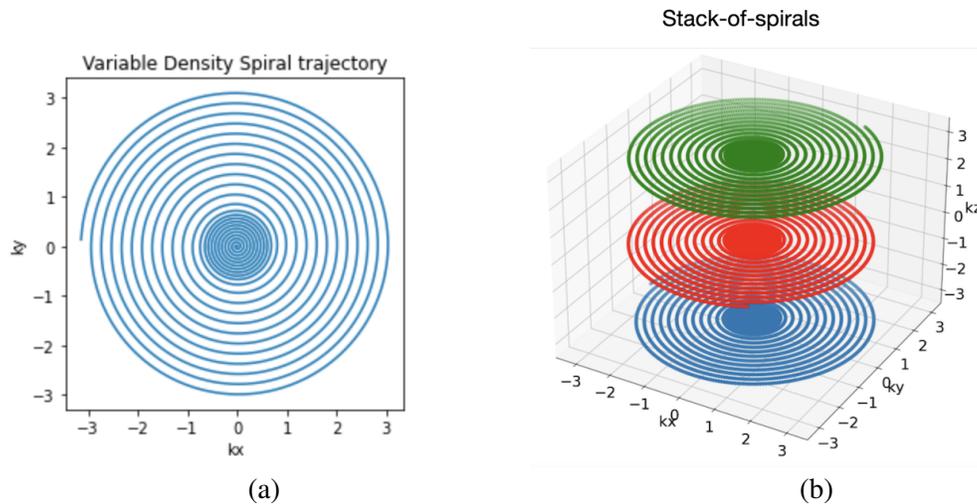


Figure 5.1: Spiral sampling trajectories. (a) 2D variable density spiral (VDS). (b) Stack-of-spirals acquisition consisting of 2D VDS shots.

than general 3D spiral or other non-Cartesian trajectories, since we can exploit the Cartesian structure along k_z to speed up reconstruction. General 3D trajectories require a 3D gridding operation (which is part of the NUFFT), and this is a computationally expensive operation. For stack-of-spirals data, it can be shown that the 3D gridding operation can be replaced by 2D gridding followed by a Fourier transform along the k_z direction, and there has been work that exploits this structure to avoid 3D gridding [155]. However, a lot of these works do not consider the effects of modeling off-resonance, where this simple method of 2D gridding followed by a 1D FFT along k_z breaks down. This chapter analyzes the stack-of-spirals scheme under the effects of off-resonance, and demonstrates that, under certain assumptions (*low-rankness*) of the off-resonance terms, we can still derive efficient implementations of the forward model.

5.2.2 Off-resonance correction in MRI

In MRI scanners, inhomogeneities in the main magnetic field (B_0 field) lead to non-uniform precessional frequencies of spins, called *off-resonance*, and this can lead to artifacts in the reconstructed images. Ref. [53] is a great review paper that describes, in detail, the sources of off-resonance as well as strategies to mitigate off-resonance effects. Some common sources of B_0 inhomogeneity include: magnet inhomogeneity, magnetic susceptibility effects (e.g., near air-filled sinuses), chemical shift (e.g., fat-water chemical shift), and metal implants. Off-resonance can cause artifacts in images such as: geometric distortion that warps the anatomy, loss of signal in areas with severe off-resonance, and blurring. The

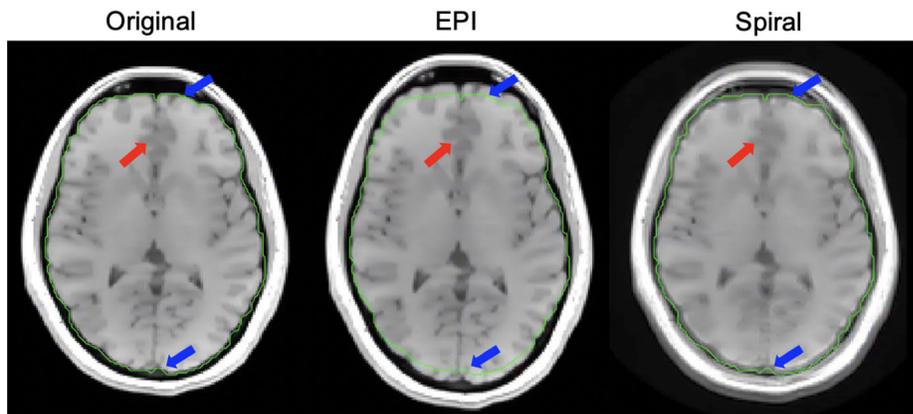


Figure 5.2: Off-resonance effects: Cartesian vs. non-Cartesian acquisitions. Spiral acquisition contains larger amount of blurring (red arrow), but with less geometric distortion (blue arrows) compared to EPI. Figure taken from [53].

artifacts can range from mild to severe based on various factors such as the type of acquisition, anatomy, application etc. Fig. 5.2 compares the off-resonance artifacts for spiral imaging vs. echo-planar imaging (EPI) and Cartesian imaging.

Off-resonance effects are more pronounced in acquisitions with longer readouts because of higher phase accrual due to off-resonance. In several applications, long readout times are preferred due to factors such as the need for efficient k-space coverage (e.g., single shot spiral imaging [52]), high echo times (e.g., BOLD fMRI [18]), and the need for high temporal resolution in dynamic and functional imaging [73]. In such acquisitions, off-resonance artifacts effects are typically significant and the image artifacts could be severe if not corrected for. Modeling off-resonance effects in iterative reconstruction methods is computationally expensive, and there has been work on developing efficient iterative algorithms for off-resonance correction [150, 39]. Additionally, for the specific case of spiral trajectories with long readouts, there have been works that study the effect of B_0 inhomogeneities [178, 52].

5.3 Modified NUFFT for stack-of-spirals

This section analyses the forward models for four different cases: stack of identical spirals with and without off-resonance correction, and rotated stack-of-spirals with and without off-resonance correction. For each of the four cases, we derive efficient implementations for the 3D NUFFT-based forward model, by exploiting the structure inherent to that specific stack-of-spirals configuration. We progressively build up to the most general case, which is a stack of rotated spirals, with off-resonance effects being modeled.

Notation

- N_{xy} : number of voxels in-plane ($N_{xy} = N_x N_y$)
- N_z : number of slices
- N_s : number of spiral partitions (we assume 1 shot per partition)
- N_r : number of samples in each spiral readout

In all our analyses, we assume the general case where $N_s \leq N_z$, i.e., the number of spiral partitions is at most the number of slices reconstructed. However, we do assume that *the spiral partitions lie on a Cartesian grid corresponding to the N_z slices*, i.e., we do not need to perform any interpolation along k_z for our stack-of-spirals forward model. Petrov

et al. provide an example of such a stack-of-spirals acquisition, where the central k_z phase encoding planes are fully sampled, while the remaining k_z planes are randomly sampled at every time point in an fMRI study [125].

5.3.1 Stack of identical spirals (no off-resonance correction)

We begin by considering a stack-of-spirals where every spiral shot is identical. Also, we do not model off-resonance effects in this case. For a 3D stack-of-spirals dataset acquired using C coils, the signal model can be written as

$$\mathbf{y}_c = \mathbf{A}_c \mathbf{x} + \boldsymbol{\epsilon}_c, \quad c = 1, \dots, C, \quad (5.1)$$

where $\mathbf{y}_c \in \mathbb{C}^{N_s N_r}$ denotes the complex-valued k-space measurements for the c^{th} coil, $\mathbf{x} \in \mathbb{C}^{N_z N_{xy}}$ is the transverse magnetization image to be reconstructed, \mathbf{A}_c represents the forward model for the c^{th} coil while $\boldsymbol{\epsilon}_c$ denotes complex Gaussian noise for the c^{th} coil.

Disregarding off-resonance effects, the system matrix can be written as

$$\begin{aligned} \mathbf{A}_c &= \tilde{\mathbf{F}}_3 \mathbf{S}_c \\ &= \mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3 \mathbf{S}_c, \end{aligned} \quad (5.2)$$

where $\tilde{\mathbf{F}}_3 \in \mathbb{C}^{N_s N_r \times N_z N_{xy}}$ denotes the three-dimensional Non-Uniform Fast Fourier Transform (NUFFT) operation [35, 40, 176] and each $\mathbf{S}_c \in \mathbb{C}^{N_z N_{xy} \times N_z N_{xy}}$ represents a coil sensitivity map. The NUFFT consists of three steps: zero-padding, oversampled FFT, and interpolation of Cartesian samples onto a non-Cartesian grid. The NUFFT operation in (5.2) can be broken down as $\tilde{\mathbf{F}}_3 = \mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3$, where \mathbf{Z}_3 denotes a zero-padding operation, \mathbf{F}_3 represents an oversampled FFT, and \mathbf{K}_3 denotes the interpolation of Cartesian k-space values onto a non-Cartesian grid (for our spiral trajectories in this case). Note that \mathbf{F}_3 is separable because of the properties of the discrete Fourier transform (DFT), i.e., \mathbf{F}_3 can be written separably in terms of the FFT operations along z as well as the 2D FFT operation in the x - y plane. We use this property to further simplify the 3D NUFFT below.

Our goal is to exploit the Cartesian structure along k_z that is inherent to stack-of-spirals acquisitions, and simplify the 3D NUFFT in (5.2). Since we acquire a stack of identical spirals, we make the following simplifications in our NUFFT block:

$$\begin{aligned} \mathbf{Z}_3 &= \mathbf{I}_{N_s} \otimes \mathbf{Z}_2 \\ \mathbf{F}_3 &= \mathbf{P} \mathbf{F}_z \otimes \mathbf{F}_{x'y'} \\ \mathbf{K}_3 &= \mathbf{I}_{N_z} \otimes \mathbf{K}_2, \end{aligned} \quad (5.3)$$

where $\mathbf{Z}_2 \in \mathbb{R}^{N_{x'y'} \times N_{xy}}$ denotes zero-padding in-plane so as to help us take the FFT on a finer grid in the next step. $N_{x'y'}$ denotes the number of voxels in-plane after zero-padding (a typical oversampling factor is 1.25 or 2 in each dimension). \mathbf{I}_{N_z} denotes the $N_z \times N_z$ identity matrix. $\mathbf{F}_z \in \mathbb{C}^{N_z \times N_z}$ computes the FFT along the z direction and $\mathbf{F}_{x'y'} \in \mathbb{C}^{N_{x'y'} \times N_{x'y'}}$ denotes the oversampled 2D FFT in-plane. $\mathbf{K}_2 \in \mathbb{R}^{N_r \times N_{x'y'}}$ represents 2D interpolation onto non-Cartesian samples. In other words, \mathbf{K}_2 denotes the in-plane interpolation of the oversampled Cartesian k-space samples onto our desired spiral trajectory samples. $\mathbf{P} \in \mathbb{R}^{N_s \times N_z}$ denotes a selection matrix that picks out only the required k_z planes corresponding to the acquired spiral partitions. Each row in \mathbf{P} contains a single 1 corresponding to the relevant k_z plane, and the rest are zeroes. For the case when $N_s = N_z$, then $\mathbf{P} = \mathbf{I}_{N_z}$.

In (5.3), we used the separability of \mathbf{F}_3 due to the properties of the DFT, i.e.,

$$\begin{aligned} \mathbf{F}_3 &= (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{x'y'}})(\mathbf{I}_{N_z} \otimes \mathbf{F}_{x'y'}) \\ &= \mathbf{P}\mathbf{F}_z \otimes \mathbf{F}_{x'y'}. \end{aligned} \quad (5.4)$$

Using the Kronecker products from (5.3) and using the mixed-product property of Kronecker products¹, (5.2) simplifies as

$$\begin{aligned} \mathbf{A}_c &= \mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3 \mathbf{S}_c \\ &= (\mathbf{I}_{N_s} \otimes \mathbf{K}_2)(\mathbf{P}\mathbf{F}_z \otimes \mathbf{F}_{x'y'})(\mathbf{I}_{N_z} \otimes \mathbf{Z}_2) \mathbf{S}_c \\ &= (\mathbf{I}_{N_s} \otimes \mathbf{K}_2)(\mathbf{I}_{N_s} \otimes \mathbf{F}_{x'y'})(\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{x'y'}})(\mathbf{I}_{N_z} \otimes \mathbf{Z}_2) \mathbf{S}_c \\ &= \underbrace{(\mathbf{I}_{N_s} \otimes \mathbf{K}_2)(\mathbf{I}_{N_s} \otimes \mathbf{F}_{x'y'})(\mathbf{I}_{N_s} \otimes \mathbf{Z}_2)}_{\text{stack of } N_s \text{ 2D NUFFT operations}} (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \mathbf{S}_c \\ &= (\mathbf{I}_{N_s} \otimes \mathbf{K}_2 \mathbf{F}_{x'y'} \mathbf{Z}_2)(\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \mathbf{S}_c \\ &= (\mathbf{I}_{N_s} \otimes \tilde{\mathbf{F}}_2)(\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \mathbf{S}_c, \end{aligned} \quad (5.5)$$

where $\tilde{\mathbf{F}}_2 = \mathbf{K}_2 \mathbf{F}_{x'y'} \mathbf{Z}_2$ is a 2D NUFFT operation performed in-plane. In (5.5), the term $(\mathbf{I}_{N_s} \otimes \tilde{\mathbf{F}}_2)$ corresponds to a stack of 2D NUFFT operations applied shot-by-shot for the N_s spiral shots. The term $(\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}})$ corresponds to taking the 1-dimensional FFT along the z -direction, and then selecting the desired k_z planes. Thus, we have simplified the 3D NUFFT in (5.2) in terms of a 1D FFT along the z -direction, followed by a stack of 2D NUFFT operations performed in-plane. This potentially provides a significant benefit when $N_s < N_z$, since we require fewer 2D NUFFT operations than before.

¹If $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{p \times q}$, then $\mathbf{A} \otimes \mathbf{B} = (\mathbf{A} \otimes \mathbf{I}_p)(\mathbf{I}_n \otimes \mathbf{B}) = (\mathbf{I}_m \otimes \mathbf{B})(\mathbf{A} \otimes \mathbf{I}_q)$. See https://en.wikipedia.org/wiki/Kronecker_product.

In (5.5), we manipulate the equations to push the $(\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}})$ term all the way to the right. An alternate approach is to push it all the way to the left of the expression. This comes in handy for the specific case when $N_s = N_z$, i.e., the number of slices is the same as the number of spiral partitions acquired. When $N_s = N_z$, we have $\mathbf{P} = \mathbf{I}_{N_z}$ and the system matrix can be written as

$$\mathbf{A}_c = (\mathbf{F}_z \otimes \mathbf{I}_{N_r})(\mathbf{I}_{N_z} \otimes \tilde{\mathbf{F}}_2)\mathbf{S}_c. \quad (5.6)$$

The term $(\mathbf{F}_z \otimes \mathbf{I}_{N_r})$ is invertible, and its inverse is given by $(\mathbf{F}_z^{-1} \otimes \mathbf{I}_{N_r})$, which corresponds to the 1D inverse FFT along the z -direction. This implies that we could take the 1D inverse FFT (along z) of the acquired stack-of-spirals data *once at the beginning*, and then the reconstruction simplifies to a stack of 2D NUFFT operations (parallelizable across slices). This simplification applies only for the specific case when $N_s = N_z$.

5.3.2 Stack of identical spirals (with off-resonance correction)

This section continues to analyze a stack of *identical* spirals, but with off-resonance effects being modeled. As discussed earlier, spiral trajectories are sensitive to off-resonance effects due to their longer readouts [14], and this results in blurring in the reconstructed images, if not corrected for. Correcting off-resonance effects during general MR image reconstruction has been investigated before, and there are efficient reconstruction algorithms to deal with B_0 effects by including them in the signal model [150, 39]. In this section, we investigate whether we can obtain a simplified forward model for a stack-of-spirals acquisition, while accounting for off-resonance effects.

The MR signal equation at time t , when considering off-resonance effects, is given by

$$s(t) = \int_V \underbrace{f(\mathbf{r}) e^{-i2\pi(\mathbf{k}(t) \cdot \mathbf{r})}}_{\text{non-uniform FT}} \underbrace{e^{-i \Delta\omega(\mathbf{r}) t}}_{\text{off-resonance term}} d\mathbf{r}, \quad (5.7)$$

where $\mathbf{r} \in \mathbb{R}^3$ represents spatial location, V is the imaging volume, $f(\mathbf{r})$ denotes the transverse magnetization to be reconstructed, $\mathbf{k}(t) \in \mathbb{R}^3$ is the location in k-space at readout time t . $\Delta\omega(\mathbf{r})$ is the off-resonance at location \mathbf{r} (in radians). From (5.7), the signal equation can be written as the Hadamard product of two terms: the first term can be captured by a non-uniform Fourier transform, while the second term captures the phase accrued due to off-resonance effects.

If we discretize the latent image f as well as the off-resonance map $\Delta\omega$ [39] in (5.7), we can write the system matrix for the c th coil as,

$$\mathbf{A}_c = (\tilde{\mathbf{F}}_3 \odot \boldsymbol{\Omega}) \mathbf{S}_c, \quad (5.8)$$

where $\boldsymbol{\Omega} \in \mathbb{C}^{N_s N_r \times N_z N_{xy}}$ captures the phase accrual due to field inhomogeneity, i.e., $\Omega_{ij} = e^{-i\Delta\omega_j t_i}$ is the phase accrued by spins in voxel j at readout time t_i . $\Delta\omega \in \mathbb{R}^{N_z N_{xy}}$ denotes the fieldmap (in units of radians) that is separately estimated. The symbol \odot denotes the element-wise product or the Hadamard product. The Hadamard product arises from the MR signal equation in (5.7), due to the product of the non-uniform FT term with the off-resonance phase term.

Reusing the result in (5.5), we simplify (5.8) as

$$\begin{aligned} \mathbf{A}_c &= (\tilde{\mathbf{F}}_3 \odot \boldsymbol{\Omega}) \mathbf{S}_c \\ &= (\mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3 \odot \boldsymbol{\Omega}) \mathbf{S}_c \\ &= (((\mathbf{I}_{N_s} \otimes \tilde{\mathbf{F}}_2)(\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}})) \odot \boldsymbol{\Omega}) \mathbf{S}_c \\ &= ((\mathbf{P}\mathbf{F}_z \otimes \tilde{\mathbf{F}}_2) \odot \boldsymbol{\Omega}) \mathbf{S}_c. \end{aligned} \quad (5.9)$$

For a stack-of-spirals acquisition, the vector of readout times \mathbf{t} has a special repetitive structure: every spiral shot has the same readout times. The time restarts at zero after each excitation pulse, because spins begin in phase then. This results in a periodic structure in the off-resonance phase matrix $\boldsymbol{\Omega}$ given by

$$\boldsymbol{\Omega} = \begin{bmatrix} \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \\ \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \\ \vdots & \vdots & \vdots & \vdots \\ \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \end{bmatrix} = \mathbf{1}_{N_s} \otimes \begin{bmatrix} \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \end{bmatrix}, \quad (5.10)$$

where $\Omega_l \in \mathbb{C}^{N_r \times N_{xy}}$ represents the phase accrued for all voxels in slice l for all readout times. The phase accrued in voxel j of slice l at readout time t_i is given by $[\Omega_l]_{ij} = e^{-i[\Delta\omega_l]_j t_i}$, where $\Delta\omega_l \in \mathbb{R}^{N_{xy}}$ is the off-resonance map for slice l . $\mathbf{1}_{N_s}$ represents a vector of N_s ones.

Expanding the matrices in (5.9) and using the notation $[\mathbf{P}\mathbf{F}_z]_{i,j} = \phi_{i,j}$ for the 1-dimensional FFT along z , we obtain

$$\mathbf{A}_c = ((\mathbf{P}\mathbf{F}_z \otimes \tilde{\mathbf{F}}_2) \odot \boldsymbol{\Omega}) \mathbf{S}_c$$

$$\begin{aligned}
&= \left(\begin{bmatrix} \phi_{1,1} \tilde{\mathbf{F}}_2 & \phi_{1,2} \tilde{\mathbf{F}}_2 & \dots & \phi_{1,N_z} \tilde{\mathbf{F}}_2 \\ \phi_{2,1} \tilde{\mathbf{F}}_2 & \phi_{2,2} \tilde{\mathbf{F}}_2 & \dots & \phi_{2,N_z} \tilde{\mathbf{F}}_2 \\ \vdots & \vdots & & \vdots \\ \phi_{N_s,1} \tilde{\mathbf{F}}_2 & \phi_{N_s,2} \tilde{\mathbf{F}}_2 & \dots & \phi_{N_s,N_z} \tilde{\mathbf{F}}_2 \end{bmatrix} \odot \begin{bmatrix} \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \\ \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \\ \vdots & \vdots & \vdots & \vdots \\ \Omega_1 & \Omega_2 & \dots & \Omega_{N_z} \end{bmatrix} \right) \mathbf{S}_c \\
&= \begin{bmatrix} \phi_{1,1}(\tilde{\mathbf{F}}_2 \odot \Omega_1) & \dots & \phi_{1,N_z}(\tilde{\mathbf{F}}_2 \odot \Omega_{N_z}) \\ \vdots & \ddots & \vdots \\ \phi_{N_s,1}(\tilde{\mathbf{F}}_2 \odot \Omega_1) & \dots & \phi_{N_s,N_z}(\tilde{\mathbf{F}}_2 \odot \Omega_{N_z}) \end{bmatrix} \mathbf{S}_c \\
&= \begin{bmatrix} \phi_{1,1} \mathbf{I}_{N_r} & \dots & \phi_{1,N_z} \mathbf{I}_{N_r} \\ \vdots & \vdots & \vdots \\ \phi_{N_s,1} \mathbf{I}_{N_r} & \dots & \phi_{N_s,N_z} \mathbf{I}_{N_r} \end{bmatrix} \underbrace{\begin{bmatrix} (\tilde{\mathbf{F}}_2 \odot \Omega_1) & & & \\ & (\tilde{\mathbf{F}}_2 \odot \Omega_2) & & \\ & & \ddots & \\ & & & (\tilde{\mathbf{F}}_2 \odot \Omega_{N_z}) \end{bmatrix}}_{\text{Stack of } N_z \text{ 2D NUFFTs with off-resonance correction}} \mathbf{S}_c \\
&= (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_r}) \mathbf{B} \mathbf{S}_c, \tag{5.11}
\end{aligned}$$

where \mathbf{B} is a block-diagonal matrix that corresponds to a stack of N_z 2D NUFFT operations with off-resonance correction. The structure of \mathbf{B} lends itself to an embarrassingly parallel² slice-by-slice implementation. The term $(\tilde{\mathbf{F}}_2 \odot \Omega_l)$ corresponds to a 2D NUFFT operation for slice l (with off-resonance correction). We have exploited the periodicity in the off-resonance matrix Ω , owing to the fact that the readout times repeat for every single-shot spiral. To summarize, for the specific case of identical stack-of-spirals (with off-resonance effects), the derivation above simplifies the forward model from a full 3D NUFFT to a stack of 2D NUFFT operations (with off-resonance correction) followed by a 1D FFT along z .

5.3.3 Rotated stack of spirals (no off-resonance correction)

This section considers a rotated stack-of-spirals acquisition [26], i.e., the spiral interleaves in each k_z -plane are not identical, but are rotated by a pre-determined amount instead. Also, we do not consider off-resonance effects in this subsection. This section analyzes the forward model from two different perspectives: (i) the matrix form of the forward model, and (ii) the MRI signal equation (integral form).

²An embarrassingly parallel problem is one which can be divided into smaller tasks that have no dependencies on each other, and so can be solved independently of each other.

5.3.3.1 Perspective: Matrix form of forward model

The forward model for a rotated stack of spirals is similar to the forward model for a stack of identical spirals (refer to (5.2) in Section 5.3.1), with one prominent difference. For the case of a rotated stack-of-spirals, the interpolation matrix \mathbf{K}_3 takes the form of

$$\mathbf{K}_3 = \begin{bmatrix} \mathbf{K}_{2,1} & & & \\ & \mathbf{K}_{2,2} & & \\ & & \ddots & \\ & & & \mathbf{K}_{2,N_s} \end{bmatrix}, \quad (5.12)$$

where $\mathbf{K}_{2,m} \in \mathbb{R}^{N_r \times N_{x'y'}}$ is the interpolation matrix corresponding to the m th spiral trajectory. N_r is the number of samples in the spiral readout, while $N_{x'y'}$ is the number of voxels (in-plane) on the oversampled Cartesian grid prior to interpolation. N_s is the number of spiral shots.

Starting from the same forward model as in Section 5.3.1 and using the expression for \mathbf{K}_3 from (5.12), we obtain

$$\begin{aligned} \mathbf{A}_c &= \mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3 \mathbf{S}_c \\ &= \mathbf{K}_3 (\mathbf{P}\mathbf{F}_z \otimes \mathbf{F}_{x'y'}) (\mathbf{I}_{N_z} \otimes \mathbf{Z}_2) \mathbf{S}_c \\ &= \mathbf{K}_3 (\mathbf{I}_{N_s} \otimes \mathbf{F}_{x'y'}) (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{x'y'}}) (\mathbf{I}_{N_z} \otimes \mathbf{Z}_2) \mathbf{S}_c \\ &= \mathbf{K}_3 (\mathbf{I}_{N_s} \otimes \mathbf{F}_{x'y'}) (\mathbf{I}_{N_s} \otimes \mathbf{Z}_2) (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \mathbf{S}_c \\ &= \begin{bmatrix} \mathbf{K}_{2,1} \mathbf{F}_{x'y'} \mathbf{Z}_2 & & & \\ & \mathbf{K}_{2,2} \mathbf{F}_{x'y'} \mathbf{Z}_2 & & \\ & & \ddots & \\ & & & \mathbf{K}_{2,N_s} \mathbf{F}_{x'y'} \mathbf{Z}_2 \end{bmatrix} (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \mathbf{S}_c \\ &= \underbrace{\begin{bmatrix} \tilde{\mathbf{F}}_{2,1} & & & \\ & \tilde{\mathbf{F}}_{2,2} & & \\ & & \ddots & \\ & & & \tilde{\mathbf{F}}_{2,N_s} \end{bmatrix}}_{\text{Stack of } N_s \text{ 2D NUFFT operations}} (\mathbf{P}\mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \mathbf{S}_c, \end{aligned} \quad (5.13)$$

where we used the separability of \mathbf{F}_3 into \mathbf{F}_z and $\mathbf{F}_{x'y'}$ (due to property of the discrete FT) in step 2, and the mixed-product property of Kronecker products to go from step 2 to step 5. $\tilde{\mathbf{F}}_{2,m}$ denotes the 2D NUFFT operation corresponding to the m th spiral shot.

Similar to our results in Section 5.3.1 (see (5.5)), we have simplified the 3D NUFFT operation into a composition of a 1D FFT along z , followed by a stack of 2D NUFFT operations performed in each of the N_s spiral partitions. This potentially provides significant computational speedups for the case $N_s < N_z$, since we need to perform only N_s 2D NUFFT operations instead of N_z of them.

5.3.3.2 Alternate perspective: MRI signal equation

This section analyses the rotated stack-of-spirals data from a different perspective: the MRI signal equation. This perspective is potentially more informative when we consider off-resonance effects later. The MRI signal equation for the m th spiral shot in the stack-of-spirals (not considering off-resonance effects and coil sensitivity information) is given by

$$s_m(t_i) = \int_V f(\mathbf{r}) e^{-i2\pi(\mathbf{k}_m(t_i) \cdot \mathbf{r})} d\mathbf{r}, \quad m = 1, 2, \dots, N_s, \quad (5.14)$$

where $\mathbf{r} \in \mathbb{R}^3$ represents spatial location, V is the imaging volume, $f(\mathbf{r})$ denotes the transverse magnetization to be reconstructed, $\mathbf{k}_m(t_i) \in \mathbb{R}^3$ is the location in k-space for the m th spiral at readout time t_i . This considers the most general case of arbitrarily rotated spiral shots in every k_z -plane. Additionally, we assume that the length of the spiral readout N_r and the readout times $\{t_i\}_{i=1}^{N_r}$ are identical for each of the N_s spiral shots. Expanding the spatial location vector \mathbf{r} as (x, y, z) and further simplifying, we obtain

$$\begin{aligned} s_m(t_i) &= \int \int \int f(x, y, z) e^{-i2\pi(\mathbf{k}_m(t_i) \cdot (x, y, z))} dx dy dz \\ &= \int \int e^{-i2\pi((k_{m,x}(t_i), k_{m,y}(t_i)) \cdot (x, y))} \left(\int f(x, y, z) e^{-i2\pi(k_{m,z}(t_i) z)} dz \right) dx dy \\ &= \int \int e^{-i2\pi((k_{m,x}(t_i), k_{m,y}(t_i)) \cdot (x, y))} \underbrace{\left(\int f(x, y, z) e^{-i2\pi(\frac{k_m}{N_z} z)} dz \right)}_{\text{1D FT (along } z)} dx dy, \end{aligned} \quad (5.15)$$

where $\mathbf{k}_m(t_i) = (k_{m,x}(t_i), k_{m,y}(t_i), k_{m,z}(t_i)) \in \mathbb{R}^3$, and we have used the Cartesian structure present in the stack-of-spirals data to write $k_{m,z}(t_i) = k_m$. Also, $\{k_m\}_{m=1}^{N_s} \subseteq \{-N_z/2, \dots, N_z/2 - 1\}(\frac{1}{N_z \Delta z})$, since we assume that all N_s spiral partitions lie on the Cartesian k_z grid corresponding to the N_z slices. Here, Δz denotes the resolution along z (in units of length). In this case, the inner integral in (5.15) simplifies as the 1-dimensional Fourier Transform (FT) of $f(x, y, z)$ along z , evaluated at the desired values $\{k_m\}$, i.e.,

$$s_m(t_i) = \int \int F_z(x, y; k_m) e^{-i2\pi((k_{m,x}(t_i), k_{m,y}(t_i)) \cdot (x,y))} dx dy, \quad (5.16)$$

where $F_z(x, y; k_m)$ represents the result of taking the 1-dimensional FT of $f(x, y, z)$ along z and extracting the k_z -plane corresponding to k_m . The double integral in (5.16) is the 2-dimensional non-uniform FT of $F_z(x, y; k_m)$ for the m th spiral shot. If we appropriately discretize the transverse magnetization $f(x, y, z)$ and write (5.16) in a block matrix formulation (along with coil sensitivity information), the forward model can be written as a composition of functions:

- Weighting by sensitivity maps
- 1D FFT along z , evaluated at $\{k_m\}_{m=1}^{N_s}$
- Stack of N_s 2D NUFFT operations

This is the same as the result we obtained in (5.13), which we derived by writing the forward model in the form of matrices.

5.3.4 Rotated stack-of-spirals (with off-resonance correction)

This section considers a rotated stack-of-spirals with off-resonance effects being modeled. This is the most general configuration of a stack-of-spirals acquisition. We analyze the forward model under both perspectives: the matrix form, as well as the integral form of the MRI signal equation (basic principles). The matrix form seems to be challenging to simplify (Section 5.3.4.1), while the integral form gives us more insight into deriving an efficient implementation of the forward model (Section 5.3.4.2).

5.3.4.1 Attempt 1: Matrix form of forward model

When off-resonance effects are modeled, the system matrix is written as

$$\begin{aligned} \mathbf{A}_c &= (\mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3 \odot \mathbf{\Omega}) \mathbf{S}_c \\ &= \left\{ \left\{ \left[\begin{array}{ccc} \tilde{\mathbf{F}}_{2,1} & & \\ & \tilde{\mathbf{F}}_{2,2} & \\ & & \ddots \\ & & & \tilde{\mathbf{F}}_{2,N_s} \end{array} \right] (\mathbf{P} \mathbf{F}_z \otimes \mathbf{I}_{N_{xy}}) \right\} \odot \mathbf{\Omega} \right\} \mathbf{S}_c, \end{aligned} \quad (5.17)$$

where we have started with the forward model from Section 5.3.2, and we have used the result from (5.13) for simplifying the 3D NUFFT terms, i.e., $\mathbf{K}_3 \mathbf{F}_3 \mathbf{Z}_3$. $\mathbf{\Omega} \in \mathbb{C}^{N_s N_r \times N_z N_{xy}}$

captures the phase accrual due to field inhomogeneity, i.e., $\Omega_{ij} = e^{-i\Delta\omega_j t_i}$ is the phase accrued in voxel j at readout time t_i . For our specific case where we assume that the readout times are the same for every spiral shot, the structure of Ω was shown before in (5.10). \mathbf{F}_z denotes the 1D FFT along z , while $\tilde{\mathbf{F}}_{2,m}$ denotes the 2D NUFFT operation corresponding to the m th spiral shot.

Expanding \mathbf{F}_z using the notation $[\mathbf{P}\mathbf{F}_z]_{i,j} = \phi_{i,j}$ and exploiting the periodic structure of Ω shown in (5.10), we obtain

$$\begin{aligned}
\mathbf{A}_c &= \left(\begin{bmatrix} \tilde{\mathbf{F}}_{2,1} & & & \\ & \tilde{\mathbf{F}}_{2,2} & & \\ & & \ddots & \\ & & & \tilde{\mathbf{F}}_{2,N_s} \end{bmatrix} \begin{bmatrix} \phi_{1,1}\mathbf{I}_{N_{xy}} & \phi_{1,2}\mathbf{I}_{N_{xy}} & \cdots & \phi_{1,N_z}\mathbf{I}_{N_{xy}} \\ \phi_{2,1}\mathbf{I}_{N_{xy}} & \phi_{2,2}\mathbf{I}_{N_{xy}} & \cdots & \phi_{2,N_z}\mathbf{I}_{N_{xy}} \\ \vdots & \vdots & & \vdots \\ \phi_{N_s,1}\mathbf{I}_{N_{xy}} & \phi_{N_s,2}\mathbf{I}_{N_{xy}} & \cdots & \phi_{N_s,N_z}\mathbf{I}_{N_{xy}} \end{bmatrix} \right. \\
&\quad \left. \odot \begin{bmatrix} \Omega_1 & \Omega_2 & \cdots & \Omega_{N_z} \\ \Omega_1 & \Omega_2 & \cdots & \Omega_{N_z} \\ \vdots & \vdots & \vdots & \vdots \\ \Omega_1 & \Omega_2 & \cdots & \Omega_{N_z} \end{bmatrix} \right) \mathbf{S}_c \\
&= \left(\begin{bmatrix} \phi_{1,1}\tilde{\mathbf{F}}_{2,1} & \phi_{1,2}\tilde{\mathbf{F}}_{2,1} & \cdots & \phi_{1,N_z}\tilde{\mathbf{F}}_{2,1} \\ \phi_{2,1}\tilde{\mathbf{F}}_{2,2} & \phi_{2,2}\tilde{\mathbf{F}}_{2,2} & \cdots & \phi_{2,N_z}\tilde{\mathbf{F}}_{2,2} \\ \vdots & \vdots & & \vdots \\ \phi_{N_s,1}\tilde{\mathbf{F}}_{2,N_s} & \phi_{N_s,2}\tilde{\mathbf{F}}_{2,N_s} & \cdots & \phi_{N_s,N_z}\tilde{\mathbf{F}}_{2,N_s} \end{bmatrix} \odot \begin{bmatrix} \Omega_1 & \Omega_2 & \cdots & \Omega_{N_z} \\ \Omega_1 & \Omega_2 & \cdots & \Omega_{N_z} \\ \vdots & \vdots & \vdots & \vdots \\ \Omega_1 & \Omega_2 & \cdots & \Omega_{N_z} \end{bmatrix} \right) \mathbf{S}_c \\
&= \begin{bmatrix} \phi_{1,1}(\tilde{\mathbf{F}}_{2,1} \odot \Omega_1) & \phi_{1,2}(\tilde{\mathbf{F}}_{2,1} \odot \Omega_2) & \cdots & \phi_{1,N_z}(\tilde{\mathbf{F}}_{2,1} \odot \Omega_{N_z}) \\ \phi_{2,1}(\tilde{\mathbf{F}}_{2,2} \odot \Omega_1) & \phi_{2,2}(\tilde{\mathbf{F}}_{2,2} \odot \Omega_2) & \cdots & \phi_{2,N_z}(\tilde{\mathbf{F}}_{2,2} \odot \Omega_{N_z}) \\ \vdots & \vdots & & \vdots \\ \phi_{N_s,1}(\tilde{\mathbf{F}}_{2,N_s} \odot \Omega_1) & \phi_{N_s,2}(\tilde{\mathbf{F}}_{2,N_s} \odot \Omega_2) & \cdots & \phi_{N_s,N_z}(\tilde{\mathbf{F}}_{2,N_s} \odot \Omega_{N_z}) \end{bmatrix} \mathbf{S}_c.
\end{aligned} \tag{5.18}$$

In hopes of reducing computation, a goal is to disentangle the FFT along the z -dimension (comprising of $\phi_{i,j}$ terms) from the 2D NUFFT operations with off-resonance correction (the bracketed terms with the Hadamard product). That would result in an efficient implementation of the forward model, since the 2D NUFFT operations could then be performed slice-by-slice in parallel. However, there does not seem to be a way to do this in

(5.18). In the following section, we start from first principles (the basic MRI signal equation) and explore whether that helps us obtain any insights (using approximations where required).

5.3.4.2 Attempt 2: MRI signal equation

To account for off-resonance effects, we write the MRI signal equation for the m th spiral shot as (see (5.15))

$$\begin{aligned}
s_m(t_i) &= \int \int \int f(x, y, z) e^{-i2\pi(\mathbf{k}_m(t_i) \cdot (x, y, z))} \underbrace{e^{-i\Delta\omega(x, y, z)t_i}}_{\text{Off-resonance term}} dx dy dz \\
&= \int \int e^{-i2\pi(\mathbf{k}_{m,2}(t_i) \cdot (x, y))} \left(\int f(x, y, z) e^{-i\Delta\omega(x, y, z)t_i} e^{-i2\pi(k_{m,z}(t_i) z)} dz \right) dx dy \\
&= \int \int e^{-i2\pi(\mathbf{k}_{m,2}(t_i) \cdot (x, y))} \left(\int f(x, y, z) e^{-i\Delta\omega(x, y, z)t_i} e^{-i2\pi(\frac{k_m}{N_z} z)} dz \right) dx dy,
\end{aligned} \tag{5.19}$$

where $\Delta\omega(x, y, z)$ is a separately acquired off-resonance map. $\mathbf{k}_{m,2}(t_i)$ is the 2-dimensional k-space location along the k_x and k_y axes for the m th spiral shot at readout time t_i . Unlike (5.15) in Section 5.3.3.2, we cannot simplify the inner integral into the form of a 1D FT along z , owing to the presence of the additional $e^{-i\Delta\omega(x, y, z)t_i}$ term. It seems hard to simplify further without making some approximations for the off-resonance terms.

We consider a low-rank approximation to the off-resonance terms (for each spiral shot), given by

$$e^{-i\Delta\omega(x, y, z)t_i} \approx \sum_{l=1}^L b_l(t_i) c_l(x, y, z), \tag{5.20}$$

where $\{b_l(\cdot)\}_{l=1}^L$ and $\{c_l(\cdot)\}_{l=1}^L$ are L basis functions underlying the low-rank approximation of the non-Fourier exponential terms [39]. In matrix form (after discretization of the object), we can write this as

$$[e^{-i\Delta\omega_j t_i}] \approx \mathbf{BC}, \tag{5.21}$$

where $\mathbf{B} \in \mathbb{C}^{N_r \times L}$ and $\mathbf{C} \in \mathbb{C}^{L \times N_v}$ capture the basis functions in matrix form. N_r is the number of samples in each spiral shot, while $N_v = N_z N_{xy}$ is the number of voxels being reconstructed. Using this low-rank approximation in (5.19) and rearranging, we obtain

$$\begin{aligned}
s_m(t_i) &\approx \int \int e^{-i2\pi(\mathbf{k}_{m,2}(t_i) \cdot (x,y))} \\
&\quad \left(\int f(x, y, z) \left(\sum_{l=1}^L b_l(t_i) c_l(x, y, z) \right) e^{-i2\pi\left(\frac{k_m}{N_z} z\right)} dz \right) dx dy \\
&= \sum_{l=1}^L b_l(t_i) \left[\int \int e^{-i2\pi(\mathbf{k}_{m,2}(t_i) \cdot (x,y))} \right. \\
&\quad \left. \underbrace{\left(\int f(x, y, z) c_l(x, y, z) e^{-i2\pi\left(\frac{k_m}{N_z} z\right)} dz \right)}_{\text{1D FT (along } z \text{) with a weighting function } \triangleq F_z(x,y;k_m,l)} dx dy \right] \\
&= \sum_{l=1}^L b_l(t_i) \underbrace{\left[\int \int F_z(x, y; k_m, l) e^{-i2\pi(\mathbf{k}_{m,2}(t_i) \cdot (x,y))} dx dy \right]}_{\text{2D NUFFT operation}}, \tag{5.22}
\end{aligned}$$

where $F_z(x, y; k_m, l)$ represents the result of taking the 1D Fourier Transform (FT) of the term $f(x, y, z) c_l(x, y, z)$ along z , evaluated at the k_z -plane corresponding to k_m .

After discretizing the integrals in (5.22), the order of operations for a practical implementation of the forward model can be written as (for $l = 1, 2, \dots, L$):

- Scaling of input volume $f(x, y, z)$ by the l th spatial basis function, i.e., $c_l(x, y, z)$, and coil sensitivity maps
- 1D FFT (along z) of this scaled volume, evaluated at the desired planes $\{k_m\}$
- 2D NUFFT operations performed for each spiral shot in parallel
- Scaling with temporal basis functions, i.e., $b_l(t_i)$, and sum over all L terms

We began with 3D NUFFT operations in the forward model, and simplified it to a composition of 1D FFTs (along z) and 2D NUFFT operations which is parallelizable across spiral shots. This needs to be repeated L times, since our approximation in (5.20) consists of L basis functions. However, the computations for each of the L terms can be independently done, providing further opportunities for speedup. In the next section, we perform a big-O analysis of (5.22) to investigate how much of a computational speedup it could potentially offer, as compared to a conventional implementation based on the full 3D NUFFT.

Big-O analysis

This section analyses the computational overhead of three different implementations of the stack-of-spirals forward model:

- Recommended formulation in (5.22) (1D FFT along z followed by 2D NUFFT)
- Full 3D NUFFT
- Modified 3D NUFFT (exploiting Cartesian structure along k_z)

Our recommended formulation in (5.22) involves taking a 1D FFT along the z direction, followed by N_s 2D NUFFTs (i.e., one NUFFT per spiral shot), for every basis function, i.e., L times. For all our big-O analyses, we assume that the oversampling factor (within the NUFFT) is 2 and that the k-space interpolator width is J along each dimension. Under these conditions, the computational cost of our recommended approach in (5.22) is

$$\begin{aligned}
 & O(L \{ \underbrace{N_{xy}N_z \log N_z}_{\text{Cost of 1D FFT along } z} + N_s \underbrace{(2^2 N_{xy} \log(2^2 N_{xy}) + J^2 N_r)}_{\text{Cost of 2D NUFFT}} \}) \\
 &= O(L \{ N_{xy}N_z \log N_z + 4N_s N_{xy} \log(2^2 N_{xy}) + J^2 N_s N_r \}) \\
 &= O(L \{ N_z N_{xy} \log N_z + 4N_s N_{xy} \log N_{xy} + (8 \log 2) N_s N_{xy} + J^2 N_s N_r \}). \quad (5.23)
 \end{aligned}$$

For the conventional approach that does not exploit the stack-of-spirals structure and takes full 3D NUFFTs of the whole image volume, the computational cost is [39]

$$\begin{aligned}
 & O(L \{ \underbrace{2^3 N_z N_{xy} \log(2^3 N_z N_{xy}) + J^3 N_s N_r}_{\text{Cost of 3D NUFFT}} \}) \\
 &= O(L \{ 8 N_z N_{xy} \log(2^3 N_z N_{xy}) + J^3 N_s N_r \}) \\
 &= O(L \{ 8 N_z N_{xy} \log N_z + 8 N_z N_{xy} \log N_{xy} + (24 \log 2) N_z N_{xy} + J^3 N_s N_r \}). \quad (5.24)
 \end{aligned}$$

The implementation in (5.24) considers the full 3D NUFFT without exploiting the stack-of-spirals structure in any way. For a more apples-to-apples comparison with our recommended approach in (5.23), we consider a modified 3D NUFFT approach where we exploit the Cartesian structure along z that is present in stack-of-spirals acquisitions. E.g., we do not need to oversample along the z dimension in the 3D NUFFT, since the k_z partitions already lie on a Cartesian grid. Also, instead of a 3D interpolation step, we can perform 2D interpolation in-plane. Incorporating these modifications, the computational

| Approach | Computational overhead |
|--|--|
| Full 3D NUFFT (5.24) | $O(L \{8 N_z N_{xy} \log N_z + 8 N_z N_{xy} \log N_{xy} + (24 \log 2) N_z N_{xy} + J^3 N_s N_r\})$ |
| Modified 3D NUFFT (5.25) | $O(L \{4 N_z N_{xy} \log N_z + 4 N_z N_{xy} \log N_{xy} + (8 \log 2) N_z N_{xy} + J^2 N_s N_r\})$ |
| Recommended implementation (5.23) | $O(L \{N_z N_{xy} \log N_z + 4 N_s N_{xy} \log N_{xy} + (8 \log 2) N_s N_{xy} + J^2 N_s N_r\})$ |

Table 5.1: Comparison of computational costs for three different implementations of the forward model for stack-of-spirals acquisitions.

cost is

$$\begin{aligned}
& O(L \{ \underbrace{2^2 N_z N_{xy} \log(2^2 N_z N_{xy})}_{\substack{\text{Cost of modified 3D NUFFT} \\ \text{(exploiting Cartesian structure along } z)}} + J^2 N_s N_r \}) \\
&= O(L \{4 N_z N_{xy} \log(2^2 N_z N_{xy}) + J^2 N_s N_r\}) \\
&= O(L \{4 N_z N_{xy} \log N_z + 4 N_z N_{xy} \log N_{xy} + (8 \log 2) N_z N_{xy} + J^2 N_s N_r\}).
\end{aligned} \tag{5.25}$$

The modified NUFFT approach in (5.25) could potentially be implemented using existing image reconstruction toolboxes. For instance, in MIRTorch, which is a PyTorch-based image reconstruction toolbox [161], there are keyword arguments for oversampling factors and k-space interpolation widths in the forward model, that knowledgeable users can select, to exploit the Cartesian structure in stack-of-spirals data. E.g., the number of points to be used for interpolation in the 3D NUFFT can be passed in as an argument to the forward model. By default, this is set to a non-zero value in all three directions (typically 6). For stack-of-spirals acquisitions, one can obtain a reconstruction speedup by modifying this argument to ensure that no interpolation is performed along z .

Table 5.1 summarizes the computational costs of the three approaches in (5.23), (5.24) and (5.25). Compared to the full 3D NUFFT, our recommended implementation significantly reduces computational cost, especially as we consider higher number of interpolation points J (since our recommended implementation is $O(J^2)$ vs. the full 3D NUFFT which is $O(J^3)$). The modified 3D NUFFT implementation provides a more apples-to-apples comparison with our recommended implementation. The computational cost of our

recommended implementation is slightly lower than the modified 3D NUFFT, and the savings become more pronounced when N_s is lower than N_z .

5.4 Experiments

Table 5.1 shows the reduction in computational costs obtained for the recommended NUFFT implementation for stack-of-spirals (2D NUFFTs + 1D FFT along k_z). To see if the theoretical speed-up of the recommended approach translated to a computational speed-up in practice, we performed simulations to obtain the wall clock times of all three implementations of the forward model for stack-of-spirals. We simulated a stack-of-spirals acquisition with a field of view of [22cm, 22cm, 10cm] and a matrix size of [92, 92, 42]. This imaging setup corresponds to an isotropic resolution of 2.4mm.

Fig. 5.1 shows the stack-of-spirals trajectories that we used for our experiments. For our spiral trajectory, we used a variable density spiral (VDS) with a fully sampled center and 3-fold undersampling at the edges of k -space. We used an acceleration factor of 3 in-plane; every k_z encode has a single spiral shot, and successive shots are rotated by 120° , for a total of 42 shots. We used off-resonance maps estimated from a separate 3D GRE acquisition of a human brain. For all our experiments, we used a PyTorch-based toolbox called MIRTorch [161], which is an open-source toolbox for image reconstruction. MIRTorch uses a package called torchkbnufft [107] for non-uniform FFT operations under the hood. We performed all experiments on a 20-core Intel® Xeon® E5-2698 v4 processor (2.2GHz). We found that the full 3D NUFFT implementation used $\sim 70\%$ of total CPU time (obtained using the *top* command in Linux, under the %CPU column), while the other two implementations used an average of $\sim 50\%$.

Fig. 5.3 shows the timing and accuracy results of our simulations. We compared three different implementations of the forward model for stack-of-spirals MRI: (i) Full 3D NUFFT, (ii) modified 3D NUFFT call with interpolation performed only in the $k_x - k_y$ plane, and (iii) our recommended implementation (2D NUFFT followed by a 1D FFT along k_z). For our timing experiments, we randomly initialized an input image, and simulated each of our forward model implementations 5 times (for the same random input). We then took the mean of all 5 runs; this was done to hopefully get a consistent estimate of wall clock times. Fig. 5.3a shows the wall clock times vs. the number of bases used in the low-rank approximation for off-resonance correction (denoted by L). The timing plot shows a clear reduction in wall clock time with our recommended implementation of the forward model. The modified implementation of the 3D NUFFT reduces the wall clock time by $\sim 68\%$ compared to the full 3D NUFFT, while our recommended implementation (with 2D

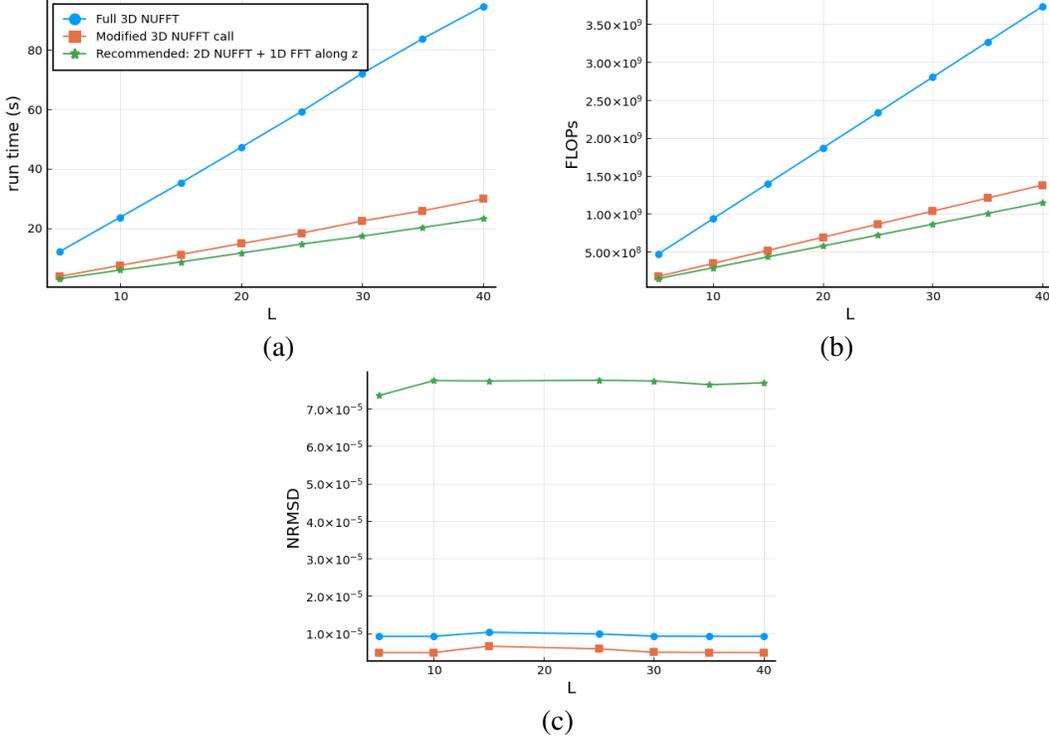


Figure 5.3: Results of simulating different implementations of the NUFFT-based forward model for a stack-of-spirals acquisition. (a) (*Empirical*) wall clock time, (b) (*Predicted*) floating point operations (FLOPs) computed using Table 5.1, (c) NRMSD vs. number of components in the low-rank approximation of the off-resonance terms (L).

NUFFTs + 1D FFT along k_z) further reduces the wall clock time by $\sim 22\%$, at the higher values of L . Fig. 5.3b shows the (predicted) floating point operations (FLOPs) per coil vs. L for all three implementations; these are computed from Table 5.1 with $N_x = N_y = 90$, $N_z = N_s = 42$, and $J = 6$.

To ensure that the recommended approach does not affect the accuracy of the result, we studied the differences between our recommended approach and a baseline implementation. For our baseline method, we used the full 3D NUFFT call in MIRTorch [161] with more neighbors used for interpolation (11 neighbors used in all 3 directions vs. 6 used in the default call). We then computed the normalized root mean square difference as

$$\text{NRMSD} = \frac{\|\mathbf{y} - \mathbf{y}^*\|_2}{\|\mathbf{y}^*\|_2}, \quad (5.26)$$

where \mathbf{y} is the result of the given forward model implementation under consideration, and \mathbf{y}^* is the result of the baseline implementation. We repeated this for all three implementations of the forward model and across a range of values of L ; Fig. 5.3c shows the results. To ensure an apples-to-apples comparison between all three implementations in Fig. 5.3c, we

fixed the bases used in the low-rank approximation for off-resonance correction (i.e., matrices B and C in (5.21) are kept fixed). In other words, we are only interested in numerical differences caused by the NUFFT implementation, and not by the low-rank approximation used in off-resonance correction. For the recommended implementation, we obtain a very low value of NRMSD, on the order of $8e-5$. All operations (including the non-uniform FFT) were performed using 32 bit floating point for the real and imaginary parts, i.e., we used the `torch.complex64` datatype in PyTorch.

Images: Density-compensated adjoint reconstruction

The preceding sections analyzed the timing and accuracy of the forward model for stack-of-spirals MRI, comparing our recommended implementation (1D FFT along z followed by a stack of 2D NUFFT operations) against conventional 3D NUFFT-based implementations. It is also desirable to see how the recommended implementation performs with reconstructing images. To analyze this, we implemented the adjoint of the recommended implementation, which corresponds to a stack of 2D adjoint NUFFT operations, followed by a 1D inverse FFT along the k_z direction. Using a simple digital phantom, we compared the adjoint reconstructions of our recommended implementation for stack-of-spirals MRI vs. the baseline implementation based on the conventional 3D NUFFT.

Fig. 5.4 shows the results of this comparison. We created a simple digital phantom, and simulated a stack-of-spirals acquisition with a field of view of [22cm, 22cm, 10cm] and a matrix size of [92, 92, 42]. Fig. 5.4a shows the ground truth of our digital phantom. In previous experiments, we demonstrated that the recommended implementation is very close to the baseline implementations (on the order of $8e-5$; see Fig. 5.3c), for a wide range of values for L . Thus, in this set of experiments, we do not consider off-resonance effects and only look at the adjoint reconstructed images in the absence of off-resonance. For our stack-of-spirals trajectory, we used the same setup as described before: we used a variable density spiral with a fully sampled center and 3-fold undersampling at the edges of k -space. Each k_z -encode contained a single spiral shot, and successive shots were rotated by 120° .

Figs. 5.4b and 5.4c show the adjoint reconstructions (with density compensation [127]) for the baseline 3D NUFFT-based implementation and the recommended implementation respectively. The baseline implementation involved the 3D NUFFT with $J = 11$ interpolation neighbors in all three directions. The reconstructed magnitude images are scaled between 0 and 1 by normalizing with the maximum value over the entire volume. The recommended implementation produces visually similar images to the baseline reference implementation. Fig. 5.4 shows 15 slices in the middle of the digital phantom.

5.5 Discussion

Undersampling in MRI is a major factor in obtaining faster acquisitions and higher temporal resolutions in several MRI applications (quantitative, dynamic, functional imaging etc.). Spiral trajectories are a popular choice for undersampling because of their efficient k-space coverage and robustness to motion, among other factors. However, depending on the application, we often require spiral readouts with longer durations, and this makes them sensitive to off-resonance effects. Correcting for these off-resonance effects makes the reconstruction problem more complex, and typically increases reconstruction time. This chapter aims to reduce reconstruction time specifically for stack-of-spirals MRI, in the presence of off-resonance.

In the absence of off-resonance effects, and if we consider a stack of identical spirals, then it is sufficient to take a 1D FFT of the k-space data at each sample location, along the k_z direction once at the beginning, and then just work with 2D NUFFTs for each k_z encode. However, the forward model becomes more complex when we consider non-identical shots in every partition, and if we would like to model off-resonance effects due to field inhomogeneities. This chapter considers the most general case of a stack of rotated spirals, as well as with off-resonance effects being modeled. Our contribution is that we mathematically analyze the forward model for this most general case of the stack-of-spirals, and show that we can exploit the Cartesian structure inherent in stack-of-spirals to obtain a computational speedup in the forward model. We show that we can replace the 3D NUFFT operation with a stack of 2D NUFFT operations followed by a 1D FFT performed along the z direction. We demonstrate the computational benefits of the recommended approach using timing and accuracy plots in Fig. 5.3.

The timing plot in Fig. 5.3a shows a clear increasing linear trend for the three different implementations of the forward model as we increase L , i.e., as we increase the number of bases in the low-rank approximation to the off-resonance terms. Also, the gap between the three lines increases as we move to higher values of L . For the highest value of L in the plot, i.e., $L = 40$, the modified NUFFT implementation reduces wall clock time by about **68%** compared to the full 3D NUFFT. In other words, if we just replace a full 3D NUFFT call (which takes 6 neighbors in all 3 directions for interpolation) with a modified NUFFT call (with 6 neighbors only in x and y directions, and no interpolation along z), then we obtain a huge speedup. Additionally, our recommended approach further reduces the wall clock time by around **22%**, which is a significant speedup, especially at large values of L and at higher spatial resolutions.

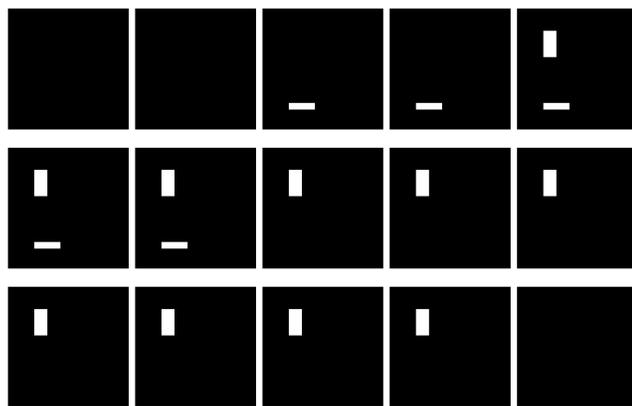
The accuracy plot in Fig. 5.3c demonstrates very low values of normalized root mean square difference (NRMSD) for the recommended implementation (on the order of $8e-5$). However, it is interesting that the accuracy curves for both the full 3D NUFFT as well as the modified 3D NUFFT are lower than the recommended approach. This suggests that our choice of the baseline implementation for the computation of the NRMSD curves might need to be carefully examined. For our baseline implementation, we used the full 3D NUFFT call with 11 neighbors used for interpolation along all three directions, instead of the default 6 neighbors. It would be interesting to explore a different baseline method, based on an exact computation, though that would be very slow to evaluate.

For all of our experiments, we used a 20-core Intel® Xeon® processor. An extension of this work would be to harness the power of graphical processing units (GPU) and potentially obtain further computational speedups. Our recommended approach contains a series of 2D NUFFT operations which could be performed slice-by-slice, and this parallelism could be further exploited.

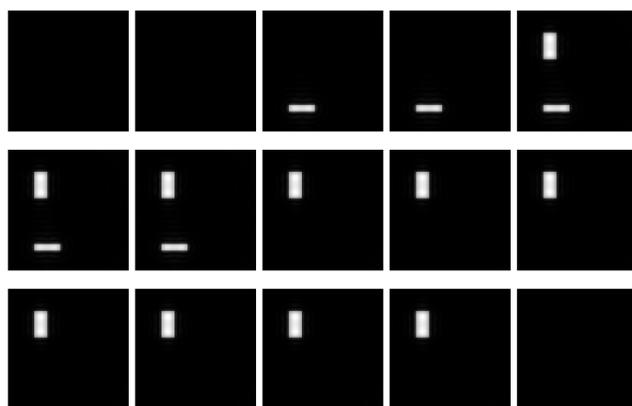
For speeding up iterative reconstruction with off-resonance correction, there has been work exploiting the Toeplitz-like structure of the Gram matrix of the forward model (using suitable approximations) [39]. A future direction could be to explore whether we could incorporate this into our recommended implementation for stack-of-spirals MRI to obtain further speedups.

5.6 Conclusions

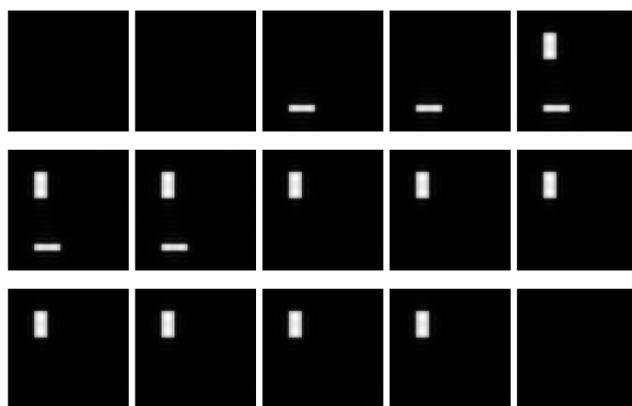
We analyzed the forward model for stack-of-spirals MRI and provided a way to exploit the Cartesian structure inherent in stack-of-spirals acquisitions (along the k_z direction) to obtain a computational speedup (even in the presence of off-resonance). We showed that the 3D NUFFT in the forward model for stack-of-spirals MRI can be replaced with a series of 2D NUFFT operations in-plane and a 1D FFT along the through-plane direction. We analyzed the computational complexity of our recommended approach and showed that it has a lower computational cost than existing implementations. We simulated the forward model for stack-of-spirals MRI, and demonstrated that our recommended implementation leads to faster wall clock times than the conventional 3D NUFFT-based implementations.



(a) Ground truth



(b) Baseline 3D NUFFFT implementation



(c) Recommended implementation

Figure 5.4: Comparison of density-compensated adjoint reconstructions for baseline 3D NUFFFT and recommended implementations for stack-of-spirals MRI, and ground truth. (Note: All three sets of images were scaled individually between 0 and 1.)

CHAPTER 6

Harmonization of Acquisition and Reconstruction Methods for Stack-of-Spirals MRI

6.1 Introduction

This chapter includes preliminary work on creating vendor-agnostic workflows for acquiring and reconstructing stack-of-spirals MRI data. We develop and introduce a toolbox called SOSP3d, which is an open-source package containing both acquisition and reconstruction methods for stack-of-spirals MRI, under one umbrella.

Reproducibility and repeatability of MRI studies has been a well-researched topic over the years, across a wide variety of different applications in MRI [105, 119]. To promote reproducibility in the MR community, there have been a lot of challenges, tasks and datasets released by research groups [108, 44, 8, 96]. The need for reproducibility becomes even more important in the case of quantitative MRI (qMRI), where we would like the estimated parameter maps to be invariant to different experimental conditions [72]. Variability in experimental conditions could arise due to several sources such as number of channels used for parallel imaging (e.g., 16-channel or 32-channel coils), sensitivity of the coils, imaging gradient non-linearity, magnetic field inhomogeneity, differences in toolboxes and algorithms used to reconstruct the data, different software versions on the scanner, as well as other scanner related factors [184, 152, 69]. These variations make it challenging to *harmonize* MRI protocols across different vendors (GE, Siemens, Philips etc.) as well as across different scanners from the same vendor.

An important milestone on the road to achieving reproducibility of results in MRI is *harmonization* of MRI protocols and workflows across different vendors and settings. We define *harmonization* of MRI as designing vendor-agnostic MRI acquisition and reconstruction protocols. Studies have shown that vendor-neutral MR workflows result in

reduced inter-vendor variability of qMRI parameters such as T_1 , magnetization transfer ratio (MTR), and magnetization transfer saturation-index (MTsat) [72]. There are both open-source and proprietary solutions for developing vendor-neutral MR frameworks [24, 153, 135].

This chapter introduces a vendor-agnostic toolbox for the acquisition and reconstruction of stack-of-spirals MRI data, called SOSP3d [111]. Chapter 5 provided background on stack-of-spirals MRI and derived efficient implementations for its forward model, in the presence of off-resonance. This chapter provides a vendor-neutral workflow for stack-of-spirals MRI from start to end, right from acquisition of stack-of-spirals data to reconstructing the data. For developing acquisition scripts for stack-of-spirals MRI, our toolbox uses a Python-based open-source package for MRI pulse sequence design called PyPulseq [131]. For reconstructing the acquired data, we use a PyTorch-based image reconstruction toolbox called MIRTorch [161].

Our SOSP3d toolbox is written in a combination of Python and MATLAB; MATLAB is primarily used to read the acquired stack-of-spirals data (e.g., using TOPPE [116] for GE scanners). Our toolbox is aimed at MRI users who would like to use stack-of-spirals MRI for their own applications, and requires only basic proficiency in Python. Also, we provide detailed documentation in the GitHub repository of our toolbox¹, starting from the very basics of setting up a Python environment with the relevant packages. By combining both acquisition and reconstruction methods for stack-of-spirals MRI under one umbrella, and hopefully providing sufficient documentation for practitioners who are not seasoned Python users, we hope that our SOSP3d toolbox would be useful for researchers to create vendor-agnostic workflows for various applications, e.g., fMRI studies using stack-of-spirals acquisitions.

The rest of the chapter is organized as follows: Section 6.2 contains relevant background for this chapter. Section 6.3 describes the SOSP3d toolbox, while Section 6.4 provides details of our experiments using the toolbox. Sections 6.5 and 6.6 contain our discussions and conclusions respectively, along with a few future research directions.

¹<https://github.com/HarmonizedMRI/SOSP3d>

6.2 Background

6.2.1 Software/toolboxes for MR acquisition

For MR pulse sequence design, there are various open-source pulse sequence development platforms from multiple research groups [67, 95, 80, 116, 131]. These tools provide useful cross-platform alternatives to vendor-supplied pulse sequence design and editing tools. ODIN is an open-source cross-platform development environment (implemented in C++) for NMR experiments. SequenceTree is an open-source graphical environment for pulse sequence programming [95]. RTHawk [135] is another vendor-neutral solution for real-time MRI, but it is a proprietary platform, which could potentially limit its use in academic research settings.

TOPPE [116] and Pulseq [80] are two relatively recent pulse sequence prototyping platforms. Both TOPPE and Pulseq share the same design philosophy: to clearly separate the sequence specification from scanner implementation [116]. E.g., using Pulseq, pulse sequence programming can be done in a hardware-independent manner (using MATLAB or JEMRIS GUI), and the resulting file is then executed using a hardware-specific interpreter for the specific vendor of choice (Siemens, GE, Bruker etc.) [80]. With a push towards open-source languages, there has been recent work with programming Pulseq sequences using Python, called PyPulseq [131], or using an open-source visual interface called Pulseq-GPI [132]. There is also recent work developing an open-source MRI simulation framework in Julia, called KomaMRI.jl [17].

6.2.2 Software/toolboxes for MR reconstruction

Image reconstruction is another key part of the MRI pipeline, and scanner-independent image reconstruction platforms are essential for improving reproducibility in MRI studies. Michigan Image Reconstruction Toolbox (MIRT)² is a collection of open-source algorithms for image reconstruction, written in MATLAB. Specifically for MR image reconstruction, MIRT has a lot of relevant algorithms including the non-uniform FFT (NUFFT) as well as algorithms for off-resonance correction. SigPy³ is another Python-based package useful for iterative algorithms for signal processing and image reconstruction. Additionally, both MIRT and SigPy also have modules for MRI pulse design. Gadgetron [49] is another open-source framework proposed for image reconstruction, along with an extension for distributed MRI [175]. BART⁴ is another free and open-source image reconstruction

²<https://github.com/JeffFessler/mirt>

³<https://github.com/mikgroup/sigpy>

⁴<https://mrirecon.github.io/bart/index.html>

toolbox [154], providing support for both MATLAB and Python; it has a suite of iterative algorithms as well as algorithms for sensitivity map estimation.

MIRTorCh [161] is another open-source differentiable image reconstruction toolbox based on PyTorch. MIRTorCh is inspired by MIRT, and contains a suite of linear operators and algorithms for various imaging applications such as MRI and SPECT. It has a modular structure, which facilitates easy development of novel image reconstruction algorithms. With native PyTorch support, MIRTorCh supports the implementation of modern model-based deep learning methods, since it is fully auto-differentiable. Also, MIRTorCh provides fast iterative algorithms for image reconstruction, supported by GPU implementations.

6.3 SOSP3d toolbox

SOSP3d is our open-source toolbox for implementing stack-of-spirals MRI in a vendor-agnostic manner, from start to finish. We provide tools for implementing both acquisition and reconstruction, and help facilitate rapid prototyping of entire workflows based on stack-of-spirals MRI. This toolbox is written in a combination of Python and MATLAB. We use Python for all of our acquisition and reconstruction code, while MATLAB is currently used for the vendor-specific parts such as reading in the data acquired on the scanner (e.g., reading GE data using TOPPE).

For reconstruction, we use code for iterative algorithms from MIRTorCh. The aim is to create a toolbox for users who are not necessarily proficient in Python or maybe not very well-versed with iterative reconstruction in general. To that end, we implemented wrappers around simple iterative algorithms such as CG-SENSE [129] (implemented using MIRTorCh), with default settings that should hopefully work right out of the box for many applications. More knowledgeable users can adjust other settings such as regularization strength, NUFFT arguments, number of iterations etc. to optimize for better image quality or faster reconstruction. Our reconstruction algorithms expect the stack-of-spirals data in a vendor-independent format, either as .h5 or .mat files. The vendor-specific k-space data from the scanner needs to be converted into either of these two formats. We also provide an option for off-resonance correction in CG-SENSE; there are options to estimate fieldmaps using either MATLAB (MIRT) or Julia (MRIFieldmaps.jl⁵).

For acquisition, we base our stack-of-spirals pulse sequence design on the PyPulseq [131] framework. For generation of gradients for the spiral trajectories, as well as other

⁵<https://github.com/MagneticResonanceImaging/MRIFieldmaps.jl>

parts of the RF pulse design itself, we use the `mri.rf` module⁶ in the SigPy package. The output of pulse sequence design is a `.seq` file (which is a vendor-independent open file format) that can be directly run on Siemens scanners. To run on GE scanners, we require a GE-specific interpreter that has been recently developed for arbitrary Pulseseq sequences [117].

6.4 Experiments

6.4.1 SOSP3d-based reconstruction

This subsection describes experiments that test only the reconstruction part of our SOSP3d toolbox, and not the acquisition module. We used TOPPE to write the stack-of-spirals acquisition code. Fig. 6.1 shows the variable density spirals (VDS) used in our stack-of-spirals acquisition. Each of our spiral shots has a fully sampled center and is 3-fold undersampled at the outer edges of k -space. For this set of experiments, we acquired a field-of-view (FOV) of $22.08 \times 22.08 \times 10.08 \text{ cm}^3$, with a matrix size of $92 \times 92 \times 42$ on a GE 3T scanner. This corresponds to an isotropic resolution of 2.4 mm. For estimation of fieldmaps, we acquired two gradient echo (GRE) scans with an echo time shift of ~ 2.27 ms.

For our stack-of-spirals scan, we acquired fully sampled data, i.e., 3 spiral shots per k_z -encode, for a total of 126 shots. We retrospectively undersampled the data to process

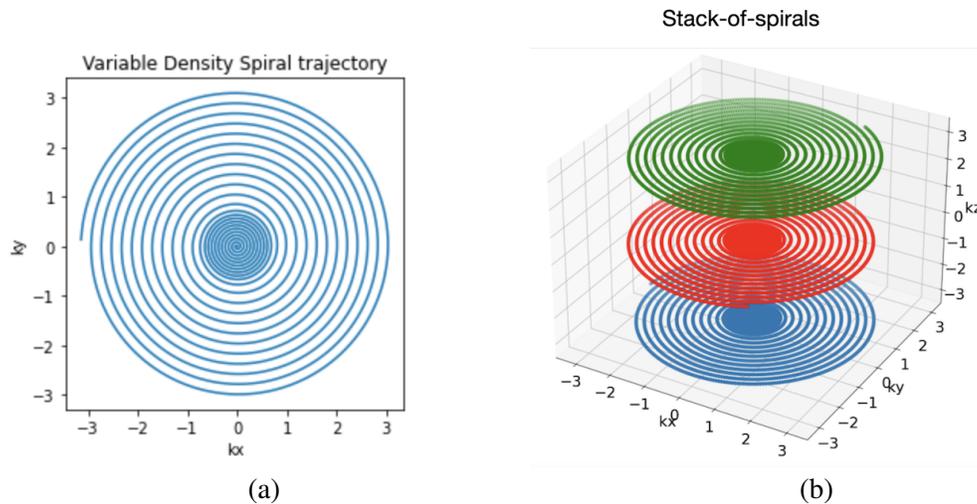


Figure 6.1: Stack-of-spirals trajectories used for our experiments. (a) 2D variable density spiral (VDS). (b) Stack-of-spirals acquisition consisting of 2D VDS shots.

⁶https://sigpy.readthedocs.io/en/latest/mri_rf.html

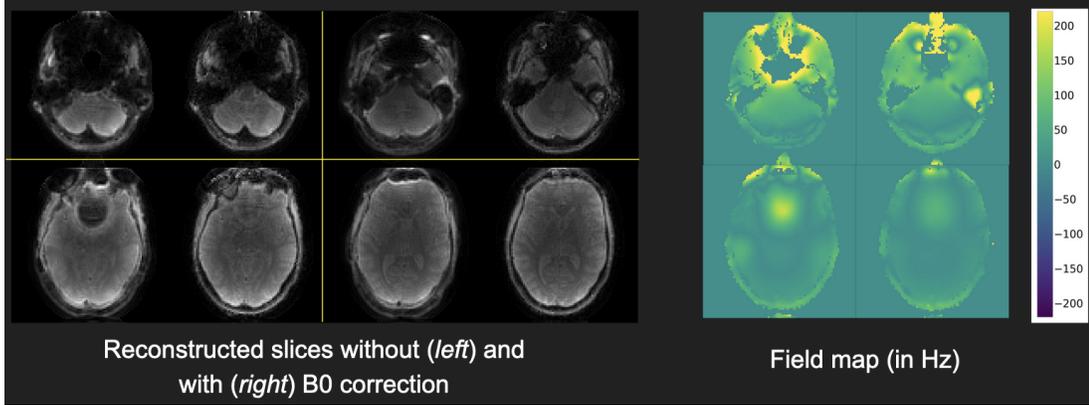


Figure 6.2: Stack-of-spirals brain data acquired using TOPPE and reconstructed (with off-resonance correction) using our SOSP3d toolbox. (Undersampling factors: $R_{xy} = 3$, $R_z = 1$.)

just 1 spiral shot per k_z -encode (for a total of 42 shots). To obtain a rotated stack-of-spirals dataset, we rotated successive spiral shots by 120° . This corresponds to an undersampling factor of $R_{xy} = 3$ in-plane, and $R_z = 1$ in the through-plane direction.

We reconstructed the undersampled stack-of-spirals data using our SOSP3d toolbox. We ran 10 iterations of CG-SENSE [129] with regularization using a quadratic roughness penalty, along with off-resonance correction [39]. Fig. 6.2 shows the field maps, as well as the reconstructed images (with and without off-resonance correction).

6.4.2 SOSP3d-based acquisition and reconstruction

This subsection describes our experiments to test both the acquisition and reconstruction modules of our SOSP3d toolbox, using the ball phantom. In contrast with the previous set of experiments in Section 6.4.1, we developed our acquisition scripts using PyPulseq (as part of our toolbox), instead of using TOPPE. We used the same stack-of-spirals setup with variable density spirals (VDS) as shown in Fig. 6.1. The VDS shots are fully sampled in the center and are 3-fold undersampled at the outer edges of k-space. We acquired a field of view of $21.6 \times 21.6 \times 14.4 \text{ cm}^3$, with a matrix size of $90 \times 90 \times 60$, corresponding to an isotropic resolution of 2.4 mm.

For our stack-of-spirals acquisition, we acquired a prospectively undersampled dataset with 1 spiral shot per k_z -encode, for a total of 60 shots. Shots in successive k_z -encodes were rotated by 120° . We reconstructed the acquired data using the reconstruction module of our toolbox. We ran 20 iterations of conjugate gradient with SENSE (CG-SENSE), without off-resonance correction. For regularization, we used a quadratic roughness penalty. Fig. 6.3 shows the reconstructed images of the ball phantom.

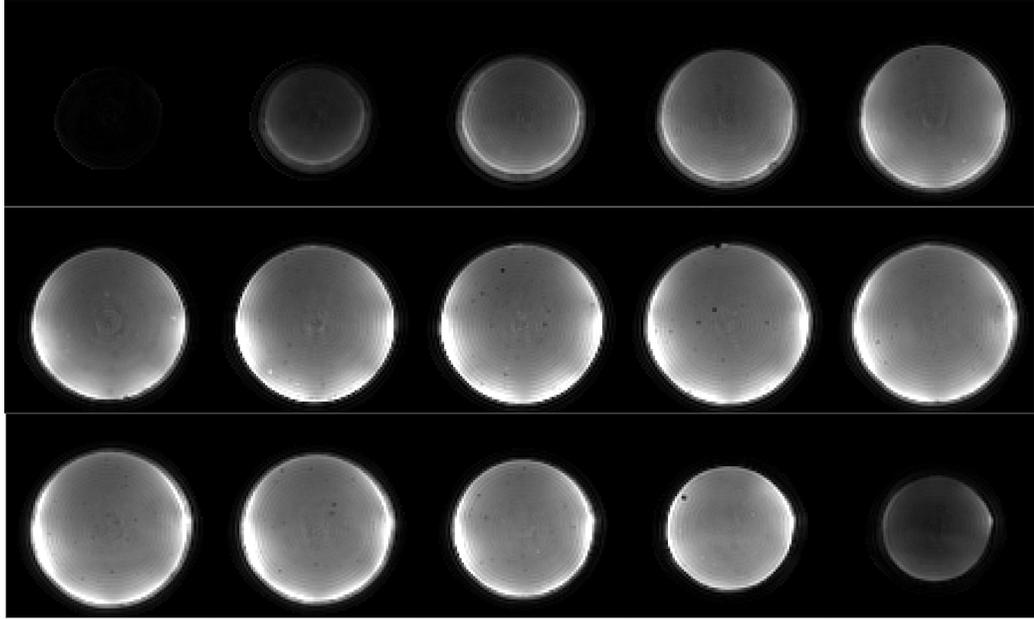


Figure 6.3: Reconstructed images of the ball phantom both acquired as well as reconstructed with our SOSP3d toolbox. (Acquisition: stack-of-spirals with undersampling factors of $R_{xy} = 3$ and $R_z = 1$.)

6.5 Discussion

This chapter develops tools for harmonizing acquisition and reconstruction methods for stack-of-spirals MRI; in particular, we introduced a toolbox called SOSP3d that contains both acquisition as well as reconstruction modules for stack-of-spirals under one umbrella. Section 6.4 described our experiments to test the acquisition and reconstruction modules of our SOSP3d toolbox. We tested our PyPulseseq-based acquisition scripts using a ball phantom; Fig. 6.3 shows the reconstructed images. The next step would be to acquire a brain dataset using a volunteer, and to verify that the stack-of-spirals acquisition module works well with an example of a human brain.

For non-Cartesian trajectories such as spiral trajectories, it is especially important to have knowledge of the actual trajectory that was acquired on the scanner; this is required for the non-uniform FFT (NUFFT) based reconstruction methods. However, the actual trajectory might be different from the desired trajectory due to hardware imperfections [171]. For our experiments, during reconstruction, we modeled a gradient delay so as to minimize the differences between the desired trajectory and the measured trajectory; this refinement reduced artifacts in the reconstructed images. A future direction of research could be to measure the actual spiral trajectory played out on the scanner, e.g., using a field camera [31].

The objective of our toolbox is to make it easy to design cross-vendor workflows for any applications involving stack-of-spirals MRI. With that in mind, a future direction of research would be to test our toolbox in a multi-vendor, multi-scanner setting. It would be interesting to explore the reproducibility of our experiments across different scanners from the same vendor, as well as different vendors (e.g., GE and Siemens).

To achieve our goal of repeatability and reproducibility of studies across scanners and vendors, it might be important to build quality assurance into our toolbox. There have been many studies evaluating MRI quality assurance of scanners using a standard phantom (e.g., the American College of Radiology (ACR) phantom) [19, 63]. A future extension of this work might be to incorporate quality metrics into our toolbox, e.g., geometric accuracy, high-contrast resolution, slice thickness accuracy, slice position accuracy, image intensity uniformity, percent signal ghosting, and low-contrast object detectability [19]. As an initial sanity check in multi-site experiments, we could run an initial scan of the ACR phantom, and evaluate the desired image quality metrics as quality assurance.

Another extension of our work in this chapter would be to use our toolbox for functional MRI (fMRI) applications based on stack-of-spirals acquisitions. We could acquire stack-of-spirals data for fMRI purposes, and process them to obtain activation maps. Another direction for future research could be to compare the activation maps obtained using our toolbox with maps obtained using other popular techniques such as simultaneous multi-slice echo planar imaging (SMS-EPI) [20].

Our SOSP3d toolbox has been developed using a combination of Python and MATLAB. We currently use MATLAB mainly for interfacing with the acquired data on the scanner, e.g., reading the P-files acquired using a GE scanner. Future work could focus on developing code in either Python or Julia to interface with the acquired data. This would help us align with the push towards using open-source languages for all of our code.

6.6 Conclusions

Reproducibility of MRI studies is an extremely important goal to achieve, and harmonization of MRI tools in a cross-vendor manner is a key milestone on the path to achieving this. This chapter develops tools to harmonize acquisition and reconstruction methods for stack-of-spirals MRI. We introduce a toolbox called SOSP3d that can be used to create vendor-agnostic workflows for applications involving stack-of-spirals MRI. We demonstrate the feasibility of acquiring and reconstructing structural data for stack-of-spirals MRI using our toolbox. As an extension of this work, we could design functional MRI (fMRI) workflows using our toolbox, and process them to obtain activation maps.

CHAPTER 7

Future Work

7.1 Myelin water imaging

CRB-based scan designs to estimate both fraction and exchange

Chapter 3 described our methods to estimate myelin water exchange in white matter. We optimized our scan design using a cost function based on the Cramér-Rao bound (see (3.5) for more details). The cost function contains a diagonal weighting matrix \mathbf{W} that determines the relative importance of estimating the different unknown parameters. For our work, we set \mathbf{W} to all zeros, except for a single 1 corresponding to myelin water exchange (i.e., mean residence time of myelin $\tau_{f \rightarrow s}$). Previous work has explored scan designs for estimating the myelin water fraction, by setting the corresponding entry (f_f) to 1 in \mathbf{W} [166]. An interesting future research direction would be to focus on scan designs for estimating multiple parameters; e.g., we could optimize the Cramér-Rao bound based cost function in (3.5) to help estimate *both* myelin water fraction as well as myelin water exchange. This would potentially help us obtain both fraction and exchange information in a single set of acquisitions, thereby reducing acquisition time.

End-to-end scan design for exchange

Optimizing MR acquisitions for quantitative MRI using the Cramér-Rao bound (CRB) comes with a caveat; the CRB helps minimize the variance in parameter estimates but is only directly applicable to *unbiased* estimators. The optimized scan designs might be sub-optimal when paired with a specific estimator, since most estimators in practice are *biased* in nature. An alternative approach to MR experiment design is to use an end-to-end design that directly minimizes the *mean square error* of our parameter(s) of interest, e.g., myelin water exchange, or cost functions based on other combinations of bias and variance.

The CRB-based cost function in (3.5) in Chapter 3 can be alternatively formulated in an end-to-end manner (using the mean square error) as

$$\hat{\mathbf{P}} = \underset{\mathbf{P} \in \mathcal{P}}{\operatorname{argmin}} \Psi(\mathbf{P}), \quad \Psi(\mathbf{P}) = \mathbb{E}_{\mathbf{x}, \boldsymbol{\nu}} [(x_e - \hat{x}_e(\mathbf{P}; \boldsymbol{\nu}))^2], \quad (7.1)$$

where $\Psi(\mathbf{P})$ is the MSE-based end-to-end cost function to be optimized, $\hat{\mathbf{P}}$ is the optimized scan design and \mathcal{P} is the search space of acquisition parameters. As described in Chapter 3, \mathbf{x} contains the unknown parameters, and $\boldsymbol{\nu}$ contains the known information (e.g., separately estimated B_0 maps). x_e is the true value of the exchange parameter (e.g., mean residence time of myelin), while $\hat{x}_e(\mathbf{P}; \boldsymbol{\nu})$ is the estimated value using our estimator of choice, e.g., PERK [114]. The end-to-end approach in (7.1) is tailored to a specific estimator, and potentially leads to scan designs that work well with our estimator of choice. Such an end-to-end optimization scheme has been previously used to design MR acquisitions for estimating myelin water fraction, when paired with PERK as the estimator [167]. Estimating myelin water exchange through this end-to-end training would be an interesting future direction to explore.

Undersampled acquisitions for estimating exchange

For our Cramér-Rao bound analyses, as well as our experiments in Chapter 3, we assumed that every bSSFP scan had full k-space sampling. However, to achieve clinically feasible times, we would potentially need to incorporate undersampling in our acquisitions. This complicates uncertainty analysis since systematic errors (due to aliasing artefacts) tend to dominate over random errors in the highly undersampled regime [68]. A future direction of research would model these undersampling artifacts in our uncertainty analysis, instead of grouping all sources of noise into random Gaussian noise. We could model the aliasing artifacts directly from first principles using computer simulations [68]. An alternate approach might be to model the undersampling effects through an additional Gaussian noise term, whose magnitude is signal-dependent [71].

Since we acquired fully sampled bSSFP data for estimating exchange, Chapters 3 and 4 used the simple inverse FFT method for reconstruction. However, to reconstruct non-Cartesian undersampled acquisitions, we would likely require non-uniform FFT (NUFFT) based iterative algorithms. Additionally, since there is potentially a lot of correlation along the contrast dimension (e.g., 40 bSSFP acquisitions to estimate myelin water exchange in Chapter 3), an extension of this work is to explore the use of regularizers along the contrast dimension, e.g., group sparsity [62].

Acquisition of transient data

Chapters 3 and 4 assumed that we acquire only steady-state data for our bSSFP scans. It seems intuitive that, to better estimate exchange values, we might benefit from recording data in the transient regime. A direction for future work could be to sample the transient regime by collecting data while we wait for the bSSFP acquisitions to reach steady-state [137, 170]. Another option could be to use fingerprinting-like sequences (instead of steady-state sequences) to drive the system into a persistent transient state [65, 48] and potentially encode more information about exchange. An extension of our Cramér-Rao bound analysis would potentially model the acquisition of both transient and steady-state data.

Combination of acquisitions for estimating exchange

This dissertation proposed the use of phase-cycled bSSFP acquisitions to estimate myelin water exchange, and we optimized the flip angles and RF phase cycling factors of a set of bSSFP scans to help estimate exchange. An area of future research would explore a combination of MR sequences instead of only bSSFP acquisitions. Estimating exchange is part of a complex multi-parametric estimation problem, and different types of MR acquisitions might help estimate different parameters of the joint estimation problem. For instance, we could design a composite set of MR sequences including bSSFP, SPGR (spoiled gradient-echo), and MESE (multi-echo spin echo) acquisitions. The SPGR acquisitions could potentially provide us with B_0 information, while the MESE scans might help with estimating myelin water fraction.

7.2 Validation using aqueous urea system

Test generalizability to different exchange rates

Exchange rates in an aqueous urea system can be manipulated by altering pH and temperature [159]. A future direction of work could be to design phantoms with different concentrations of urea, as well as at different pH values, resulting in different urea fraction values and urea water exchange rates. Such phantoms would be a good test of the generalizability of the proposed exchange estimation methods to different urea fractions as well as exchange rates.

REXSY/CEST as a baseline method for exchange

Relaxation exchange spectroscopy (REXSY) [33] and chemical exchange saturation transfer (CEST) [173] have previously been used to study exchange in exchanging systems, but are not feasible *in vivo* due to their prohibitively long acquisition times. However, they could potentially serve as baseline techniques for our aqueous urea system, since scan time is not a limiting factor for *in vitro* experiments. Chapter 4 uses inversion recovery NMR (IR-NMR) [168] spectroscopic experiments to obtain ground truth exchange values; a future direction of research could be to use REXSY or CEST to generate ground truth data and compare them with the estimates obtained from IR-NMR data. REXSY and CEST involve different methods and mechanisms of characterizing exchange, and we hope to get more insight into quantifying urea water exchange through these different techniques.

Suite of acquisitions to characterize urea water phantom

Chapter 4 specifically focused on estimating urea water exchange in our aqueous urea system. It might potentially be helpful to run separate acquisitions to obtain estimates of other parameters such as the T_1 values, T_2 values, and chemical shift; we could use these estimates to inform our nonlinear least squares fitting process. An important future direction would be to characterize our designed aqueous urea system in a more complete manner. This becomes even more important when we have a phantom with multiple urea solutions (with different concentrations and pH values), and it would be useful to characterize each sample separately. For instance, using a dedicated NMR spectrometer such as the Agilent 700 MHz NMR spectrometer, we could run a battery/suite of acquisitions to obtain relevant information about our sample. For instance, inversion recovery (IR) acquisitions would help us obtain compartment-specific T_1 values for both urea and water proton pools. Similarly, fitting data to GRE acquisitions would give us T_2 values for both pools. It would also be interesting to get estimates of the chemical shift in our aqueous urea sample through techniques such as chemical shift imaging [16].

Characterizing our urea water phantoms in this manner also would help us with our proposed techniques for estimating exchange using phase-cycled bSSFP; we could potentially use these acquisitions to guide our estimator. E.g., this suite of acquisitions could help inform our training ranges for PERK, or to set our bounds for non-linear least squares (NLLS) fitting.

Stability of aqueous urea system over time

The rate of exchange in the aqueous urea system is highly dependent on factors such as pH and temperature [159]. Studies have shown that one potential issue with the aqueous urea model is pH instability, which arises due to decomposition of urea into ammonium and cyanate ions [88]. Previous work has shown that buffering the urea solution (e.g., using phosphate buffer) stabilizes the pH value but the buffer itself can increase exchange, if added in relatively high concentrations [33]. For our urea water phantom, it would be interesting to explore the stability of exchange rates over time. We could conduct experiments to estimate urea fraction and urea water exchange rates, and repeat them over a course of time to test the repeatability and reproducibility of those measurements in our designed urea water system.

7.3 Stack-of-spirals MRI

Efficient implementation on GPU

Chapter 5 analyzed and derived an efficient implementation of the forward model for stack-of-spirals MRI. We demonstrated the computational speedup of the recommended implementation over conventional 3D non-uniform FFT (NUFFT) based implementations. The next step would be to implement this using graphical processing units (GPU), using software packages such as `torchkbnufft`¹ [107] and `MIRTorCh`² [161].

BOLD fMRI studies with PRESTO gradients

Chapter 6 described our attempts to harmonize acquisition and reconstruction methods for stack-of-spirals MRI using our `SOSP3d`³ toolbox. We demonstrated the use of our toolbox to acquire and reconstruct structural stack-of-spirals scans. An interesting future direction would be to run functional MRI (fMRI) studies or quantitative MRI (qMRI) scans using the toolbox.

BOLD (blood-oxygen-level-dependent) fMRI protocols require reasonably high temporal resolution (typically $\sim 0.8s$ per frame). However, we require fairly long echo times (~ 30 ms) to see the relevant T_2^* contrast for BOLD fMRI⁴, and this could be a limiting factor in achieving the desired temporal resolution. The PRESTO technique is an approach

¹<https://github.com/mmuckley/torchkbnufft>

²<https://github.com/guanhuaw/MIRTorCh>

³<https://github.com/HarmonizedMRI/SOSP3d>

⁴<https://mriquestions.com/bold-pulse-sequences.html>

to overcome this, and allows for echo times longer than the sequence repetition time [156]. There has been work combining the PRESTO technique with stack-of-spirals acquisitions for fMRI applications [66]. A future direction would be to incorporate the PRESTO technique in the acquisition module of our SOSP3d toolbox, and run fMRI studies with it. It would also be interesting to compare the performance of stack-of-spirals fMRI with other popular fMRI acquisition protocols such as simultaneous multi-slice echo planar imaging (SMS-EPI) [20].

Explore complex regularizers for stack-of-spirals fMRI

Chapter 6 demonstrates the use of our SOSP3d toolbox for harmonizing acquisition and reconstruction methods for stack-of-spirals MRI. For all of our experiments with structural scans, we used a simple quadratic roughness penalty as the regularizer. One of our areas for future work is to use our toolbox for fMRI studies using stack-of-spirals acquisitions. To achieve the necessary temporal resolution for fMRI studies, we would potentially need to undersample even more in our stack-of-spirals acquisitions, and a simple quadratic roughness regularizer might not be sufficient. An extension of this work could explore more complex regularizers along the temporal dimension too, e.g., low-rank + sparsity based approaches [125]. Such regularizers would also be important for qMRI studies where good spatial resolution is desired.

BIBLIOGRAPHY

- [1] C. B. Ahn, J. H. Kim, and Z. H. Cho. High-speed spiral-scan echo planar NMR imaging-I. *IEEE Trans. Med. Im.*, 5(1):2–7, 1986.
- [2] E. Alonso-Ortiz, I. R. Levesque, and G. B. Pike. MRI-based myelin water imaging: A technical review. *Mag. Res. Med.*, 73(1):70–81, January 2015.
- [3] E. C. A. Araujo, Y. Fromes, and P. G. Carlier. New insights on human skeletal muscle tissue compartments revealed by in vivo T2 NMR relaxometry. *Biophys. J.*, 106(10):2267–2274, 2014.
- [4] Ishu Arpan, Sean C Forbes, Donovan J Lott, Claudia R Senesac, Michael J Daniels, William T Triplett, Jasjit K Deol, H Lee Sweeney, Glenn A Walter, and Krista Vandendorpe. T2 mapping provides multiple approaches for the characterization of muscle involvement in neuromuscular diseases: a cross-sectional study of lower leg muscles in 5–15-year-old boys with Duchenne muscular dystrophy. *NMR in Biomedicine*, 26(3):320–328, 2013.
- [5] Chaitra Badve, Alice Yu, Matthew Rogers, Dan Ma, Yiyang Liu, Mark Schluchter, Jeffrey Sunshine, Mark Griswold, and Vikas Gulani. Simultaneous T1 and T2 brain relaxometry in asymptomatic volunteers using magnetic resonance fingerprinting. *Tomography*, 1(2):136–144, 2015.
- [6] R. Bai, D. Benjamini, J. Cheng, and P. J. Basser. Fast, accurate 2D-MR relaxation exchange spectroscopy (REXSY): Beyond compressed sensing. *J. Chem. Phys.*, 145(15):154202, October 2016.
- [7] N. K. Bangerter, B. A. Hargreaves, S. S. Vasanawala, J. M. Pauly, G. E. Gold, and D. G. Nishimura. Analysis of multiple-acquisition SSFP. *Mag. Res. Med.*, 51(5):1038–1047, 2004.
- [8] Y. Beauferris, J. Teuwen, D. Karkalousos, N. Moriakov, M. Caan, L. Rodrigues, A. Lopes, H. Pedrini, L. Rittner, M. Dannecker, et al. Multi-channel MR reconstruction (MC-MRRec) challenge—Comparing accelerated MR reconstruction models and assessing their generalizability to datasets collected with different coils. *arXiv preprint arXiv:2011.07952*, 2020.
- [9] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.

- [10] O. Bieri and K. Scheffler. Fundamentals of balanced steady state free precession MRI. *J. of Mag. Res. Imag.*, 38(1):2–11, 2013.
- [11] TA Bjarnason, IM Vavasour, CLL Chia, and AL MacKay. Characterization of the NMR behavior of white matter in bovine brain. *Mag. Res. Med.*, 54(5):1072–1081, 2005.
- [12] M. Björk, R. R. Ingle, E. Gudmundson, P. Stoica, D. G. Nishimura, and J. K. Barral. Parameter estimation approach to banding artifact reduction in balanced steady-state free precession. *Mag. Res. Med.*, 72(3):880–892, 2014.
- [13] F. Bloch. Nuclear induction. *Phys. Rev.*, 70(7-8):460–74, October 1946.
- [14] K. T. Block and J. Frahm. Spiral imaging: a critical appraisal. *J. of Mag. Res. Im.*, 21(6):657–668, 2005.
- [15] P. Börnert, H. Schomberg, B. Aldefeld, and J. Groen. Improvements in spiral MR imaging. *Mag. Res. Mat. in Phys., Bio. and Med.*, 9:29–41, 1999.
- [16] H. F. Brink, M. D. Buschmann, and B. R. Rosen. NMR chemical shift imaging. *Computerized medical imaging and graphics*, 13(1):93–104, 1989.
- [17] C. Castillo-Passi, R. Coronado, G. Varela-Mattatall, C. Alberola-López, R. Botnar, and P. Irarrazaval. KomaMRI.jl: an open-source framework for general MRI simulations with GPU acceleration. *Mag. Res. Med.*, 90(1):329–342, 2023.
- [18] C. Chang and G. H. Glover. Variable-density spiral-in/out functional magnetic resonance imaging. *Mag. Res. Med.*, 65(5):1287–1296, 2011.
- [19] C-C. Chen, Y-L. Wan, Y-Y. Wai, and H-L. Liu. Quality assurance of clinical MRI scanners using ACR MRI phantom: preliminary results. *Journal of digital imaging*, 17(4):279–284, 2004.
- [20] L. Chen, A. T. Vu, J. Xu, S. Moeller, K. Ugurbil, E. Yacoub, and D. A. Feinberg. Evaluation of highly accelerated simultaneous multi-slice EPI for fMRI. *Neuroimage*, 104:452–459, 2015.
- [21] Otavio R Coelho-Filho, François-Pierre Mongeon, Richard Mitchell, Heitor Moreno Jr, Wilson Nadruz Jr, Raymond Kwong, and Michael Jerosch-Herold. Role of transcytolemmal water-exchange in magnetic resonance measurements of diffuse myocardial fibrosis in hypertensive heart disease. *Circulation: Cardiovascular Imaging*, 6(1):134–141, 2013.
- [22] O. Cohen, B. Zhu, and M. S. Rosen. MR fingerprinting Deep RecOnstruction Network (DRONE). *Mag. Res. Med.*, 80(3):885–94, September 2018.
- [23] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Trans. Med. Imag.*, 17(3):463–8, June 1998.

- [24] C. Cordes, S. Konstandin, D. Porter, and M. Günther. Portable and platform-independent MR pulse sequence programs. *Mag. Res. Med.*, 83(4):1277–1290, 2020.
- [25] C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20:273–297, 1995.
- [26] W. Deng, B. Zahneisen, and V. A. Stenger. Rotated stack-of-spirals partial acquisition for rapid volumetric parallel MRI. *Mag. Res. Med.*, 76(1):127–135, 2016.
- [27] S. C. L. Deoni. Transverse relaxation time (T2) mapping in the brain with off-resonance correction using phase-cycled steady-state free precession imaging. *J. of Mag. Res. Imag.*, 30(2):411–417, 2009.
- [28] S. C. L. Deoni. Correction of main and transmit magnetic field (B0 and B1) inhomogeneity effects in multicomponent-driven equilibrium single-pulse observation of T1 and T2. *Mag. Res. Med.*, 65(4):1021–35, April 2011.
- [29] S. C. L. Deoni, B. K. Rutt, T. Arun, C. Pierpaoli, and D. K. Jones. Gleaning multi-component T1 and T2 information from steady-state imaging data. *Mag. Res. Med.*, 60(6):1372–87, December 2008.
- [30] S. C. L. Deoni, B. K. Rutt, and T. M. Peters. Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Mag. Res. Med.*, 49(3):515–526, 2003.
- [31] B. E. Dietrich, D. O. Brunner, B. J. Wilm, C. Barmet, S. Gross, L. Kasper, M. Haerberlin, T. Schmid, S. J. Vannesjo, and K. P. Pruessmann. A field camera for MR sequence monitoring and system analysis. *Mag. Res. Med.*, 75(4):1831–1840, 2016.
- [32] R. D. Dortch, K. D. Harkins, M. R. Juttukonda, J. C. Gore, and M. D. Does. Characterizing inter-compartmental water exchange in myelinated tissue using relaxation exchange spectroscopy. *Mag. Res. Med.*, 70(5):1450–9, November 2013.
- [33] RD Dortch, RA Horch, and MD Does. Development, simulation, and validation of NMR relaxation-based exchange measurements. *J. of Chemical Phys.*, 131(16):164502, 2009.
- [34] Adrienne N Dula, Daniel F Gochberg, Holly L Valentine, William M Valentine, and Mark D Does. Multiexponential T2, magnetization transfer, and quantitative histology in white matter tracts of rat spinal cord. *Mag. Res. Med.*, 63(4):902–909, 2010.
- [35] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM Jour. on Sci. Comp.*, 14(6):1368–1393, 1993.
- [36] Jeff H Duyn. Frequency shifts in the myelin water compartment. *Mag. Res. Med.*, 71(6):1953, 2014.

- [37] Zhenghan Fang, Yong Chen, Weili Lin, and Dinggang Shen. Quantification of relaxation times in MR fingerprinting using deep learning. In *Proc. Intl. Soc. Mag. Res. Med.*, volume 25. NIH Public Access, 2017.
- [38] L. Feng. Golden-angle radial MRI: basics, advances, and applications. *J. Mag. Res. Im.*, 56(1):45–62, 2022.
- [39] J. A. Fessler, S. Lee, V. T. Olafsson, H. R. Shi, and D. C. Noll. Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity. *IEEE Trans. on Sig. Proc.*, 53(9):3393–3402, 2005.
- [40] J. A. Fessler and B. P. Sutton. Nonuniform fast Fourier transforms using min-max interpolation. *IEEE Trans. On Sig. Proc.*, 51(2):560–574, 2003.
- [41] EG Finer, F Franks, and MJ Tait. Nuclear magnetic resonance studies of aqueous urea solutions. *J. Am. Chem. Soc.*, 94(13):4424–4429, 1972.
- [42] G. H. Glover. Simple analytic spiral k-space algorithm. *Mag. Res. Med.*, 42(2):412–415, 1999.
- [43] G. H. Glover and J. M. Pauly. Projection reconstruction techniques for reduction of motion effects in MRI. *Mag. Res. Med.*, 28(2):275–289, 1992.
- [44] W. A. Grissom, K. Setsompop, S. A. Hurley, J. Tsao, J. V. Velikina, and A. A. Samsonov. Advancing RF pulse design using an open-competition format: report from the 2015 ISMRM challenge. *Mag. Res. Med.*, 78(4):1352–1361, 2017.
- [45] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Mag. Res. Med.*, 34(6):910–4, December 1995.
- [46] J. P. Haldar, J. Anderson, and S. W. Sun. Maximum likelihood estimation of T1 relaxation parameters using VARPRO. In *Proc. Intl. Soc. Mag. Res. Med.*, page 41, 2007.
- [47] Justin P Haldar and Daeun Kim. OEDIPUS: An experiment design framework for sparsity-constrained MRI. *IEEE Trans. Med. Imag.*, 38(7):1545–1558, 2019.
- [48] J. I. Hamilton, A. Deshmane, M. Griswold, and N. Seiberlich. MR fingerprinting with chemical exchange (MRF-X) for in vivo multi-compartment relaxation and exchange rate mapping. In *Proc. Intl. Soc. Mag. Res. Med.*, page 0431, 2016.
- [49] M. S. Hansen and T. S. Sørensen. Gadgetron: an open source framework for medical image reconstruction. *Mag. Res. Med.*, 69(6):1768–1776, 2013.
- [50] B. A. Hargreaves, S. S. Vasanawala, J. M. Pauly, and D. G. Nishimura. Characterization and reduction of the transient response in steady-state MR imaging. *Mag. Res. Med.*, 46(1):149–58, July 2001.

- [51] K. D. Harkins, A. N. Dula, and M. D. Does. Effect of intercompartmental water exchange on the apparent myelin water fraction in multiexponential T2 measurements of rat spinal cord. *Mag. Res. Med.*, 67(3):793–800, 2012.
- [52] T. B. Harshbarger and D. B. Twieg. Iterative reconstruction of single-shot spiral MRI with off resonance. *IEEE Trans. on Med. Im.*, 18(3):196–205, 1999.
- [53] M. W. Haskell, J-F. Nielsen, and D. C. Noll. Off-resonance artifact correction for MRI: A review. *NMR in Biomed.*, 36(5):e4867, 2023.
- [54] R. Heule, J. Bause, O. Pusterla, and K. Scheffler. Multi-parametric artificial neural network fitting of phase-cycled balanced steady-state free precession data. *Mag. Res. Med.*, 84(6):2981–2993, 2020.
- [55] R. Heule, A. Deshmane, M. Zaiss, K. Herz, P. Ehses, and K. Scheffler. Structure or exchange? On the feasibility of chemical exchange detection with balanced steady-state free precession in tissue—an in vitro study. *NMR in Biomedicine*, 33(3):e4200, 2020.
- [56] W. S. Hinshaw. Image formation by nuclear magnetic resonance: The sensitive point method. *J. Appl. Phys.*, 47(8):3709, August 1976.
- [57] M. N. Hoff, J. B. Andre, and Q-S. Xiang. Combined geometric and algebraic solutions for removal of bSSFP banding artifacts with performance comparisons. *Mag. Res. Med.*, 77(2):644–654, 2017.
- [58] Elisabeth Hoppe, Gregor Körzdörfer, Tobias Würfl, Jens Wetzl, Felix Lugauer, Josef Pfeuffer, and Andreas K Maier. Deep learning for magnetic resonance fingerprinting: A new approach for predicting quantitative parameter values from time series. *GMDS*, 243:202–206, 2017.
- [59] R. A. Horch and M. D. Does. Aqueous urea as a model system for bi-exponential relaxation. *Mag. Res. Mat. in Phys., Biol. and Med.*, 20:51–56, 2007.
- [60] Charles Randall House. *Water transport in cells and tissues*. E. Arnold, 1974.
- [61] Y. Hu and G. H. Glover. Three-dimensional spiral technique for high-resolution functional MRI. *Mag. Res. Med.*, 58(5):947–951, 2007.
- [62] J. Huang, C. Chen, and L. Axel. Fast multi-contrast MRI reconstruction. *Mag. Res. Im.*, 32(10):1344–52, December 2014.
- [63] T. M. Ihalainen, N. T. Lönnroth, J. I. Peltonen, J. K. Uusi-Simola, M. H. Timonen, L. J. Kuusela, S. E. Savolainen, and O. E. Sipilä. MRI quality assurance using the ACR phantom in a multi-unit imaging center. *Acta oncologica*, 50(6):966–972, 2011.
- [64] P. Irrazabal and D. G. Nishimura. Fast three dimensional magnetic resonance imaging. *Mag. Res. Med.*, 33(5):656–662, 1995.

- [65] Y. Jiang, D. Ma, N. Seiberlich, V. Gulani, and M. A. Griswold. MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout. *Mag. Res. Med.*, 74(6):1621–31, December 2015.
- [66] M. M. Jimeno, S. Geethanath, J-F. Nielsen, and D. C. Noll. Cross-vendor implementation of a Stack-of-spirals PRESTO BOLD fMRI sequence using TOPPE and Pulseq. In *Proc. Intl. Soc. Mag. Res. Med.*, page 3649, 2020.
- [67] T. H. Jochimsen and M. Von Mengershausen. ODIN—object-oriented development interface for NMR. *J. Mag. Res.*, 170(1):67–78, 2004.
- [68] S. P. Jordan, S. Hu, I. Rozada, D. F. McGivney, R. Boyacioglu, D. C. Jacob, S. Huang, M. Beverland, H. G. Katzgraber, M. Troyer, M. A. Griswold, and D. Ma. Automated design of pulse sequences for magnetic resonance fingerprinting using physics-inspired optimization. *Proc. Natl. Acad. Sci.*, 118(40):e2020516118, October 2021.
- [69] J. Jovicich, M. Marizzoni, B. Bosch, D. Bartrés-Faz, J. Arnold, J. Benninghoff, J. Wiltfang, L. Roccatagliata, A. Picco, F. Nobili, et al. Multisite longitudinal reliability of tract-based spatial statistics in diffusion tensor imaging of healthy elderly subjects. *NeuroImage*, 101:390–403, 2014.
- [70] S. Kalantari, C. Laule, T. A. Bjarnason, I. M. Vavasour, and A. L. MacKay. Insight into in vivo magnetization exchange in human white matter regions. *Mag. Res. Med.*, 66(4):1142–51, 2011.
- [71] D. Kara, M. Fan, J. Hamilton, M. Griswold, N. Seiberlich, and R. Brown. Parameter map error due to normal noise and aliasing artifacts in MR fingerprinting. *Mag. Res. Med.*, 81(5):3108–23, May 2019.
- [72] Agah Karakuzu, Labonny Biswas, Julien Cohen-Adad, and Nikola Stikov. Vendor-neutral sequences and fully transparent workflows improve inter-vendor reproducibility of quantitative MRI. *Magnetic Resonance in Medicine*, 88(3):1212–1228, 2022.
- [73] L. Kasper, M. Engel, J. Heinzle, M. Mueller-Schrader, N. N. Graedel, J. Reber, T. Schmid, C. Barmet, B. J. Wilm, K. E. Stephan, et al. Advances in spiral fMRI: a high-resolution study with single-shot acquisition. *NeuroImage*, 246:118738, 2022.
- [74] S. M. Kay. *Fundamentals of statistical signal processing: Estimation theory*. Prentice-Hall, New York, 1993.
- [75] K. Keskin, U. Yılmaz, and T. Çukur. Constrained ellipse fitting for efficient parameter mapping with phase-cycled bSSFP MRI. *IEEE Trans. on Med. Imag.*, 41(1):14–26, 2021.
- [76] Sunghoon Kim, Harry Quon, Laurie A Loevner, Mark A Rosen, Lawrence Dougherty, Alex M Kilger, Jerry D Glickson, and Harish Poptani. Transcytolemmal water exchange in pharmacokinetic analysis of dynamic contrast-enhanced MRI

- data in squamous cell carcinoma of the head and neck. *J. Mag. Res. Im.*, 26(6):1607–1617, 2007.
- [77] Cornelia Laule, E Leung, D KB Li, AL Traboulsee, DW Paty, AL MacKay, and G RW Moore. Myelin water imaging in multiple sclerosis: quantitative correlations with histopathology. *Multiple Sclerosis Journal*, 12(6):747–753, 2006.
- [78] Cornelia Laule and GR Wayne Moore. Myelin water imaging to detect demyelination and remyelination and its validation in pathology. *Brain Pathology*, 28(5):750–764, 2018.
- [79] P. C. Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 242(5394):190–191, 1973.
- [80] K. J. Layton, S. Kroboth, F. Jia, S. Littin, H. Yu, J. Leupold, J-F. Nielsen, T. Stöcker, and M. Zaitsev. Pulseq: a rapid and hardware-independent pulse sequence prototyping framework. *Mag. Res. Med.*, 77(4):1544–1552, 2017.
- [81] J. Lee, J-W. Hyun, J. Lee, E-J. Choi, H-G. Shin, K. Min, Y. Nam, H. J. Kim, and S-H. Oh. So you want to image myelin using MRI: an overview and practical guide for myelin water imaging. *J. Mag. Res. Im.*, 53(2):360–373, 2021.
- [82] J. H. Lee, C. Labadie, C. S. Springer Jr, and G. S. Harbison. Two-dimensional inverse Laplace transform NMR: altered relaxation times allow detection of exchange correlation. *J. of Am. Chem. Soc.*, 115(17):7761–7764, 1993.
- [83] Jing Huei Lee, Christian Labadie, Charles S Springer Jr, and Gerard S Harbison. Two-dimensional inverse Laplace transform NMR: altered relaxation times allow detection of exchange correlation. *J. Am. Chem. Soc.*, 115(17):7761–7764, 1993.
- [84] Vanda A Lennon, Thomas J Kryzer, Sean J Pittock, AS Verkman, and Shannon R Hinson. IgG marker of optic-spinal multiple sclerosis binds to the aquaporin-4 water channel. *J. of Experimental Med.*, 202(4):473–477, 2005.
- [85] Z-P. Liang and P. C. Lauterber. *Principles of magnetic resonance imaging*. IEEE, New York, 2000.
- [86] C. Liao, B. Bilgic, M. K. Manhard, B. Zhao, X. Cao, J. Zhong, L. L. Wald, and K. Setsompop. 3D MR fingerprinting with accelerated stack-of-spirals and hybrid sliding-window and GRAPPA reconstruction. *NeuroImage*, 162:13–22, 2017.
- [87] J-R. Liao, J. M. Pauly, T. J. Brosnan, and N. J. Pelc. Reduction of motion artifacts in cine MRI using variable-density spiral trajectories. *Mag. Res. Med.*, 37(4):569–575, 1997.
- [88] M-F. Lin, C. Williams, M. V. Murray, G. Conn, and P. A. Ropp. Ion chromatographic quantification of cyanate in urea solutions: estimation of the efficiency of cyanate scavengers for use in recombinant protein manufacturing. *J. Chromatography B*, 803(2):353–362, 2004.

- [89] C. Liu, R. Bammer, D-H. Kim, and M. E. Moseley. Self-navigated interleaved spiral (SNAILS): application to high-resolution diffusion tensor imaging. *Mag. Res. Med.*, 52(6):1388–1396, 2004.
- [90] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Mag. Res. Med.*, 58(6):1182–1195, 2007.
- [91] D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A. Griswold. Magnetic resonance fingerprinting. *Nature*, 495:187–93, March 2013.
- [92] Dan Ma, Yun Jiang, Yong Chen, Debra McGivney, Bhairav Mehta, Vikas Gulani, and Mark Griswold. Fast 3D magnetic resonance fingerprinting for a whole-brain coverage. *Mag. Res. Med.*, 79(4):2190–2197, 2018.
- [93] A. Mackay, K. Whittall, J. Adler, D. Li, D. Paty, and D. Graeb. In vivo visualization of myelin water in brain by magnetic resonance. *Mag. Res. Med.*, 31(6):673–7, June 1994.
- [94] A. L. MacKay and C. Laule. Magnetic resonance of myelin water: an in vivo marker for myelin. *Brain plasticity*, 2(1):71–91, 2016.
- [95] J. F. Magland, C. Li, M. C. Langham, and F. W. Wehrli. Pulse sequence programming in a dynamic visual environment: SequenceTree. *Mag. Res. Med.*, 75(1):257–265, 2016.
- [96] O. Maier, S. H. Baete, A. Fyrdahl, K. Hammernik, S. Harrevelt, L. Kasper, A. Karakuzu, M. Loecher, F. Patzig, Y. Tian, et al. CG-SENSE revisited: Results from the first ISMRM reproducibility challenge. *Mag. Res. Med.*, 85(4):1821–1839, 2021.
- [97] A. P. Manning, A. L. MacKay, and C. A. Michal. Understanding aqueous and non-aqueous proton T1 relaxation in brain. *J. Mag. Res.*, 323:106909, February 2021.
- [98] H. M. McConnell. Reaction rates by nuclear magnetic resonance. *J. of Chemical Phys.*, 28(3):430–31, March 1958.
- [99] P. J. McDonald, J-P. Korb, J. Mitchell, and L. Monteilhet. Surface relaxation and chemical exchange in hydrating cement pastes: a two-dimensional NMR relaxation study. *Physical Review E*, 72(1):011409, 2005.
- [100] C. H. Meyer, B. S. Hu, D. G. Nishimura, and A. Macovski. Fast spiral coronary artery imaging. *Mag. Res. Med.*, 28(2):202–213, 1992.
- [101] K. L. Miller. Asymmetries of the balanced SSFP profile. Part I: theory and observation. *Mag. Res. Med.*, 63(2):385–395, 2010.
- [102] K. L. Miller, S. M. Smith, and P. Jezzard. Asymmetries of the balanced SSFP profile. Part II: White matter. *Mag. Res. Med.*, 63(2):396–406, February 2010.

- [103] K. L. Miller, R. H. N. Stijssen, N. Stikov, and T. W. Okell. Steady-state MRI: methods for neuroimaging. *Imaging in Medicine*, 3(1):93, 2011.
- [104] Karla L Miller. fMRI using balanced steady-state free precession (SSFP). *NeuroImage*, 62(2):713–719, 2012.
- [105] H. Mirzaalian, L. Ning, P. Savadjiev, O. Pasternak, S. Bouix, O. Michailovich, G. Grant, C. E. Marx, R. A. Morey, L. A. Flashman, et al. Inter-site and inter-scanner diffusion MRI data harmonization. *NeuroImage*, 135:311–323, 2016.
- [106] P. K. Mogensen and A. N. Riseth. Optim: A mathematical optimization package for Julia. *Journal of Open Source Software*, 3(24):615, 2018.
- [107] M. J. Muckley, R. Stern, T. Murrell, and F. Knoll. TorchKbNufft: A high-level, hardware-agnostic non-uniform fast Fourier transform. In *ISMRM Workshop on Data Sampling & Image Reconstruction*, 2020. Source code available at <https://github.com/mmuckley/torchkbnufft>.
- [108] Matthew J Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, et al. Results of the 2020 fastMRI challenge for machine learning MR image reconstruction. *IEEE Trans. Med. Im.*, 40(9):2306–2317, 2021.
- [109] N. Murthy and J. A. Fessler. Block axial checkerboarding: A distributed algorithm for helical X-Ray CT reconstruction. In *Proc. IEEE Intl. Symp. Biomed. Imag.*, pages 191–4, 2020.
- [110] N. Murthy, J-F. Nielsen, S. T. Whitaker, M. W. Haskell, S. D. Swanson, N. Seiberlich, and J. A. Fessler. Quantifying exchange using optimized bSSFP sequences. In *Proc. Intl. Soc. Mag. Res. Med.*, page 2068, 2022.
- [111] N. Murthy and D. Gupta and J. A. Fessler and J-F. Nielsen. SOSP3d: Toolbox for harmonization of stack-of-spirals MRI, 2023. Available at <https://github.com/HarmonizedMRI/SOSP3d>.
- [112] G. Nataraj. *Advances in quantitative MRI: acquisition, estimation, and application*. PhD thesis, Univ. of Michigan, Ann Arbor, MI, 48109-2122, Ann Arbor, MI, 2018.
- [113] G. Nataraj, J-F. Nielsen, and J. A. Fessler. Optimizing MR scan design for model-based T1, T2 estimation from steady-state sequences. *IEEE Trans. Med. Imag.*, 36(2):467–77, February 2017.
- [114] G. Nataraj, J-F. Nielsen, C. D. Scott, and J. A. Fessler. Dictionary-free MRI PERK: Parameter estimation via regression with kernels. *IEEE Trans. Med. Imag.*, 37(9):2103–14, September 2018.
- [115] Thanh D Nguyen, Kofi Deh, Elizabeth Monohan, Sneha Pandya, Pascal Spince-maille, Ashish Raj, Yi Wang, and Susan A Gauthier. Feasibility and reproducibility of whole brain myelin water mapping in 4 minutes using fast acquisition with spiral

- trajectory and adiabatic T2prep (FAST-T2) at 3T. *Magnetic resonance in medicine*, 76(2):456–465, 2016.
- [116] J-F. Nielsen and D. C. Noll. TOPPE: A framework for rapid prototyping of MR pulse sequences. *Mag. Res. Med.*, 79(6):3128–34, June 2018.
- [117] J-F. Nielsen, K. Wang, R. F. Schulte, A. Zhu, A. Ajala, S. Huang, D. Gupta, S. Peltier, M. Egan, R. Fung, M. Zaitsev, and D. C. Noll. An accurate interpreter of Pulseq files for GE scanners. In *Proc. Intl. Soc. Mag. Res. Med.*, 2024.
- [118] D. G. Nishimura. Principles of magnetic resonance imaging, 1996. Unpublished textbook.
- [119] S. Noble, D. Scheinost, and R. T. Constable. A decade of test-retest reliability of functional connectivity: A systematic review and meta-analysis. *NeuroImage*, 203:116157, 2019.
- [120] J. Nocedal and S. J. Wright. *Numerical optimization*. Springer, New York, 2006. 2nd edition.
- [121] D. C. Noll. Multishot rosette trajectories for spectrally selective MR imaging. *IEEE Trans. Med. Im.*, 16(4):372–377, 1997.
- [122] J. Oh, E. T. Han, M. C. Lee, S. J. Nelson, and D. Pelletier. Multislice brain myelin water fractions at 3T in multiple sclerosis. *NeuroImaging*, 17(2):156–163, 2007.
- [123] G. Ongie, N. Murthy, L. Balzano, and J. A. Fessler. A memory-efficient algorithm for large-scale sparsity regularized image reconstruction. In *Proc. 5th Intl. Mtg. on Image Formation in X-ray CT*, pages 20–3, 2018.
- [124] D. C. Peters, P. Rohatgi, R. M. Botnar, S. B. Yeon, K. V. Kissinger, and W. J. Manning. Characterizing radial undersampling artifacts for cardiac applications. *Mag. Res. Med.*, 55(2):396–403, 2006.
- [125] A. Y. Petrov, M. Herbst, and V. A. Stenger. Improving temporal resolution in fMRI using a 3D spiral acquisition and low rank plus sparse (L+ S) reconstruction. *NeuroImage*, 157:660–674, 2017.
- [126] J. G. Pipe. Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging. *Mag. Res. Med.*, 42(5):963–969, 1999.
- [127] J. G. Pipe and P. Menon. Sampling density compensation in MRI: rationale and an iterative numerical solution. *Mag. Res. Med.*, 41(1):179–186, 1999.
- [128] G. F. Piredda, T. Hilbert, J-P. Thiran, and T. Kober. Probing myelin content of the human brain with MRI: A review. *Mag. Res. Med.*, 85(2):627–52, February 2021.
- [129] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Mag. Res. Med.*, 46(4):638–651, 2001.

- [130] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Mag. Res. Med.*, 42(5):952–962, 1999.
- [131] K. S. Ravi, S. Geethanath, and J. T. Vaughan. PyPulseq: A python package for mri pulse sequence design. *Journal of Open Source Software*, 4(42):1725, 2019.
- [132] K. S. Ravi, S. Potdar, P. Poojar, A. K. Reddy, S. Kroboth, J-F. Nielsen, M. Zaitsev, R. Venkatesan, and S. Geethanath. Pulseq-Graphical Programming Interface: Open source visual environment for prototyping pulse sequences and integrated magnetic resonance imaging algorithm development. *Mag. Res. Im.*, 52:9–15, 2018.
- [133] G. M. C. Rossi, A. L. C. Mackowiak, B. C. Açikgöz, K. Pierzchała, T. Kober, T. Hilbert, and J. A. M. Bastiaansen. SPARCQ: A new approach for fat fraction mapping using asymmetries in the phase-cycled balanced SSFP signal profile. *Mag. Res. Med.*, 90(6):2348–2361, 2023.
- [134] L. I. Sacolick, F. Wiesinger, I. Hancu, and M. W. Vogel. B1 mapping by Bloch-Siegert shift. *Mag. Res. Med.*, 63(5):1315–22, May 2010.
- [135] J. M. Santos, G. A. Wright, and J. M. Pauly. Flexible real-time magnetic resonance imaging framework. In *The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 1, pages 1048–1051. IEEE, 2004.
- [136] C. Saunders, A. Gammerman, and V. Vovk. Ridge regression learning algorithm in dual variables. In *Proc. Intl. Conf. Mach. Learn.*, pages 515–21, 1998.
- [137] K. Scheffler. On the transient phase of balanced SSFP sequences. *Mag. Res. Med.*, 49(4):781–3, April 2003.
- [138] K. Scheffler and J. Hennig. Frequency resolved single-shot MR imaging using stochastic k-space trajectories. *Mag. Res. Med.*, 35(4):569–576, 1996.
- [139] K. Scheffler and Juergen Hennig. Is TrueFISP a gradient-echo or a spin-echo sequence? *Mag. Res. Med.*, 49(2):395–7, February 2003.
- [140] Y. Shcherbakova, C. A. T. van den Berg, C. T. W. Moonen, and L. W. Bartels. PLANET: an ellipse fitting approach for simultaneous T1 and T2 mapping using phase-cycled balanced steady-state free precession. *Mag. Res. Med.*, 79(2):711–722, 2018.
- [141] Y. Shcherbakova, C. A. T. van den Berg, C. T. W. Moonen, and L. W. Bartels. On the accuracy and precision of PLANET for multiparametric MRI using phase-cycled bSSFP imaging. *Mag. Res. Med.*, 81(3):1534–1552, 2019.
- [142] T. Shin, K. S. Nayak, J. M. Santos, D. G. Nishimura, B. S. Hu, and M. V. McConnell. Three-dimensional first-pass myocardial perfusion MRI using a stack-of-spirals acquisition. *Mag. Res. Med.*, 69(3):839–844, 2013.

- [143] J. G. Sled. Modelling and interpretation of magnetization transfer imaging in the brain. *NeuroImage*, 182:128–35, November 2018.
- [144] J. Stabinska, P. Neudecker, A. Ljimini, H-J. Wittsack, R. S. Lanzman, and A. Müller-Lutz. Proton exchange in aqueous urea solutions measured by water-exchange (WEX) NMR spectroscopy and chemical exchange saturation transfer (CEST) imaging in vitro. *Mag. Res. Med.*, 82(3):935–947, 2019.
- [145] G. J. Stanisz, A. Kecojevic, M. J. Bronskill, and R. M. Henkelman. Characterizing white matter with magnetization transfer and T2. *Mag. Res. Med.*, 42(6):1128–36, December 1999.
- [146] V. A. Stenger, F. E. Boada, and D. C. Noll. Three-dimensional tailored RF pulses for the reduction of susceptibility artifacts in T_2^* -weighted functional MRI. *Mag. Res. Med.*, 44(4):525–531, 2000.
- [147] Wendy A Stewart, Alex L Mackay, Kenneth P Whittall, GR Wayne Moore, and Donald W Paty. Spin-spin relaxation in experimental allergic encephalomyelitis. analysis of CPMG data using a non-linear least squares method and linear inverse theory. *Mag. Res. Med.*, 29(6):767–775, 1993.
- [148] F. Stulp and O. Sigaud. Many regression algorithms, one unified model: A review. *Neural Networks*, 69:60–79, 2015.
- [149] C. Sun and G. S. Boutis. Measurement of the exchange rate of waters of hydration in elastin by 2D T2–T2 correlation nuclear magnetic resonance spectroscopy. *New journal of physics*, 13(2):025026, 2011.
- [150] B. P. Sutton, D. C. Noll, and J. A. Fessler. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE Trans. on Med. Im.*, 22(2):178–188, 2003.
- [151] H. Takeda, S. Farsiu, and P. Milanfar. Kernel regression for image processing and reconstruction. *IEEE Trans. on Im. Proc.*, 16(2):349–366, 2007.
- [152] S. J. Teipel, S. Reuter, B. Stieltjes, J. Acosta-Cabronero, U. Ernemann, A. Fellgiebel, M. Filippi, G. Frisoni, F. Henschel, F. Jessen, et al. Multicenter stability of diffusion tensor imaging measures: a European clinical and physical phantom study. *Psychiatry Research: Neuroimaging*, 194(3):363–371, 2011.
- [153] G. Tong, A. S. Gaspar, E. Qian, K. S. Ravi, J. T. Vaughan Jr, R. G. Nunes, and S. Geethanath. A framework for validating open-source pulse sequences. *Mag. Res. Im.*, 87:7–18, 2022.
- [154] M. Uecker, J. I. Tamir, F. Ong, and M. Lustig. The BART toolbox for computational magnetic resonance imaging. In *Proc Intl Soc Magn Reson Med*, volume 24, page 1, 2016.

- [155] G. Valvano, N. Martini, L. Landini, and M. F. Santarelli. Variable density randomized stack of spirals (VDR-SoS) for compressive sensing MRI. *Mag. Res. Med.*, 76(1):59–69, 2016.
- [156] P. van Gelderen, J. H. Duyn, N. F. Ramsey, G. Liu, and C. T. W. Moonen. The PRESTO technique for fMRI. *NeuroImage*, 62(2):676–681, 2012.
- [157] AS Verkman. Water permeability measurement in living cells and complex tissues. *Journal of Membrane Biology*, 173(2):73–87, 2000.
- [158] P. Virtue, S. X. Yu, and M. Lustig. Better than real: Complex-valued neural nets for MRI fingerprinting. In *Proc. IEEE Intl. Conf. on Image Processing*, pages 3953–7, 2017.
- [159] Robert L Vold, ES Daniel, and SO Chan. Magnetic resonance measurements of proton exchange in aqueous urea. *J. Am. Chem. Soc.*, 92(23):6771–6776, 1970.
- [160] C. Y. Wang, S. Coppo, B. B. Mehta, N. Seiberlich, X. Yu, and M. A. Griswold. Magnetic resonance fingerprinting with quadratic RF phase for measurement of T2* simultaneously with δf , T1, and T2. *Mag. Res. Med.*, 81(3):1849–1862, 2019.
- [161] G. Wang, N. Shah, K. Zhu, D. C. Noll, and J. A. Fessler. MIRTorch: A PyTorch-powered differentiable toolbox for fast image reconstruction and scan protocol optimization. In *Proc. Intl. Soc. Magn. Reson. Med. (ISMRM)*, page 4982, 2022.
- [162] Y. Wang, P. van Gelderen, J. A. de Zwart, and J. H. Duyn. B0-field dependence of MRI T1 relaxation in human brain. *NeuroImage*, 213:116700, 2020.
- [163] K. E. Washburn and P. T. Callaghan. Tracking pore to pore exchange using relaxation exchange spectroscopy. *Phys. Rev. Letters*, 97(17):175502, 2006.
- [164] Kathryn L West, Nathaniel D Kelm, Robert P Carson, Daniel F Gochberg, Kevin C Ess, and Mark D Does. Myelin volume fraction imaging with MRI. *NeuroImage*, 182:511–521, 2018.
- [165] S. Wharton and R. Bowtell. Fiber orientation-dependent white matter contrast in gradient echo MRI. *Proc. Natl. Acad. Sci.*, 109(45):18559–64, November 2012.
- [166] S. T. Whitaker, G. Nataraj, J-F. Nielsen, and J. A. Fessler. Myelin water fraction estimation using small-tip fast recovery MRI. *Mag. Res. Med.*, 84(4):1977–90, October 2020.
- [167] S. T. Whitaker, J-F. Nielsen, and J. A. Fessler. End-to-End Scan Parameter Optimization for Improved Myelin Water Imaging. In *Proc. Intl. Soc. Mag. Res. Med.*, page 1273, 2022.
- [168] T. J. Williams, A. D. Kershaw, V. Li, and X. Wu. An inversion recovery NMR kinetics experiment. *J. of Chem. Ed.*, 88(5):665–669, 2011.

- [169] B. J. Wilm, C. Barmet, M. Pavan, and K. P. Pruessmann. Higher order reconstruction for MRI in the presence of spatiotemporal field perturbations. *Mag. Res. Med.*, 65(6):1690–1701, 2011.
- [170] Pauline W Worters and Brian A Hargreaves. Balanced SSFP transient imaging using variable flip angles for a predefined signal profile. *Mag. Res. Med.*, 64(5):1404–1412, 2010.
- [171] K. L. Wright, J. I. Hamilton, M. A. Griswold, V. Gulani, and N. Seiberlich. Non-Cartesian parallel imaging reconstruction. *J. Mag. Res. Im.*, 40(5):1022–1040, 2014.
- [172] B. Wu, W. Li, A. V. Avram, S-M. Gho, and C. Liu. Fast and tissue-optimized mapping of magnetic susceptibility and T2* with multi-echo and multi-shot spirals. *NeuroImage*, 59(1):297–305, 2012.
- [173] B. Wu, G. Warnock, M. Zaiss, C. Lin, M. Chen, Z. Zhou, L. Mu, D. Nanz, R. Tuura, and G. Delso. An overview of CEST MRI for non-MR physicists. *EJNMMI physics*, 3(1):1–21, 2016.
- [174] Q-S. Xiang and M. N. Hoff. Banding artifact removal for bSSFP imaging with an elliptical signal model. *Mag. Res. Med.*, 71(3):927–933, 2014.
- [175] H. Xue, S. Inati, T. S. Sørensen, P. Kellman, and M. S. Hansen. Distributed MRI reconstruction using gadgetron-based cloud computing. *Mag. Res. Med.*, 73(3):1015–1025, 2015.
- [176] Z. Yang and M. Jacob. Mean square optimal NUFFT approximation for efficient non-Cartesian MRI reconstruction. *J. Mag. Res.*, 242:126–135, 2014.
- [177] Jaeyeon Yoon, Enhao Gong, Itthi Chatnuntaweck, Berkin Bilgic, Jingu Lee, Woojin Jung, Jingyu Ko, Hosan Jung, Kawin Setsompop, Greg Zaharchuk, et al. Quantitative susceptibility mapping using deep neural network: QSMnet. *Neuroimage*, 179:199–206, 2018.
- [178] E. Yudilevich and H. Stark. Spiral sampling in magnetic resonance imaging - the effect of inhomogeneities. *IEEE Trans. On Med. Im.*, 6(4):337–345, 1987.
- [179] Moritz Zaiss and Peter Bachert. Chemical exchange saturation transfer (CEST) and MR Z-spectroscopy in vivo: a review of theoretical approaches and methods. *Phys. Med. Biol.*, 58(22):R221, 2013.
- [180] S. Zhang, Z. Liu, A. Grant, J. Keupp, R. E. Lenkinski, and E. Vinogradov. Balanced Steady-State Free Precession (bSSFP) from an effective field perspective: Application to the detection of chemical exchange (bSSFPX). *J. of Mag. Res.*, 275:55–67, 2017.
- [181] T. Zhang, J. M. Pauly, and I. R. Levesque. Accelerating parameter mapping with a locally low rank constraint. *Mag. Res. Med.*, 73(2):655–61, 2015.

- [182] B. Zhao, J. P. Haldar, C. Liao, D. Ma, Y. Jiang, M. A. Griswold, K. Setsompop, and L. L. Wald. Optimal experiment design for magnetic resonance fingerprinting: -Rao bound meets spin dynamics. *IEEE Trans. Med. Imag.*, 38(3):844–61, March 2019.
- [183] B. Zhao, W. Lu, T. K. Hitchens, F. Lam, C. Ho, and Z-P. Liang. Accelerated MR parameter mapping with low-rank and sparsity constraints. *Mag. Res. Med.*, 74(2):489–98, August 2015.
- [184] T. Zhu, R. Hu, X. Qiu, M. Taylor, Y. Tso, C. Yiannoutsos, B. Navia, S. Mori, S. Ekholm, G. Schifitto, et al. Quantification of accuracy and precision of multi-center DTI measurements: a diffusion phantom and human brain study. *NeuroImage*, 56(3):1398–1411, 2011.
- [185] R. Zi, D. Zhu, and Q. Qin. Quantitative T2 mapping using accelerated 3D stack-of-spiral gradient echo readout. *Mag. Res. Im.*, 73:138–147, 2020.
- [186] P. C. M. Van Zijl and N. N. Yadav. Chemical exchange saturation transfer (CEST): what is in a name and what isn't? *Mag. Res. Med.*, 65(4):927–948, 2011.
- [187] JR Zimmerman and W E. Brittin. Nuclear magnetic resonance studies in multiple phase systems: lifetime of a water molecule in an adsorbing phase on silica gel. *J. of Chemical Phys.*, 61(10):1328–1333, 1957.
- [188] Berislav V Zlokovic. The blood-brain barrier in health and chronic neurodegenerative disorders. *Neuron*, 57(2):178–201, 2008.
- [189] Y. Zur, S. Stokar, and P. Bendel. An analysis of fast imaging sequences with steady-state transverse magnetization refocusing. *Mag. Res. Med.*, 6(2):175–93, February 1988.