

MR Shuffling: Accelerated Single-Scan Multi-Contrast Magnetic Resonance Imaging

Jonathan Tamir



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2018-124

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-124.html>

August 16, 2018

Copyright © 2018, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

MR Shuffling: Accelerated Single-Scan Multi-Contrast Magnetic Resonance Imaging

by

Jonathan I. Tamir

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Associate Professor Michael Lustig, Chair

Associate Professor Chunlei Liu

Assistant Professor Moriel Vandsburger

Summer 2018

MR Shuffling: Accelerated Single-Scan Multi-Contrast Magnetic Resonance Imaging

Copyright 2018
by
Jonathan I. Tamir

Abstract

MR Shuffling: Accelerated Single-Scan Multi-Contrast Magnetic Resonance Imaging

by

Jonathan I. Tamir

Doctor of Philosophy in Engineering – Electrical Engineering and Computer Sciences

University of California, Berkeley

Associate Professor Michael Lustig, Chair

Magnetic resonance imaging (MRI) is an attractive medical imaging modality as it is non-invasive and does not involve ionizing radiation. Routine clinical MRI exams obtain MR images corresponding to different soft tissue contrast by performing multiple scans. When two-dimensional (2D) imaging is used, these scans are often repeated in other scanning planes. As a result, the number of scans comprising an MRI exam leads to prohibitively long exam times as compared to other medical imaging modalities such as computed tomography. Many approaches have been designed to accelerate the MRI acquisition while maintaining diagnostic quality.

One approach is to collect multiple measurements while the MRI signal is evolving due to relaxation. This enables a reduction in scan time, as fewer acquisition windows are needed to collect the same number of measurements. However, when the temporal aspect of the acquisition is left unmodeled, artifacts are likely to appear in the reconstruction. Most often, these artifacts manifest as image blurring. The effect depends on the acquisition parameters as well as the tissue relaxation itself, resulting in spatially varying blurring. The severity of the artifacts is directly related to the level of acceleration, and thus presents a tradeoff with scan time. The effect is amplified when imaging in three dimensions, severely limiting scan efficiency. Volumetric variants would be used if not for the blurring, as they are able to reconstruct images at isotropic resolution and support multi-planar reformatting.

Another established acceleration technique, called parallel imaging, takes advantage of spatially sensitive receive coil arrays to collect multiple MRI measurements in parallel. Thus, the acquisition is shortened, and the reconstruction uses the spatial sensitivity information to recover the image. More recently, methods have been developed that leverage image structure such as sparsity and low rank to reduce the required number of samples for a well-posed reconstruction. Compressed sensing and its low rank extensions use these concepts to acquire incoherent measurements below the Nyquist rate. These techniques are especially suited to MRI, as incoherent measurements can be easily achieved through pseudo-random under-sampling. As the mechanisms behind parallel imaging

and compressed sensing are fundamentally different, they can be combined to achieve even higher acceleration.

This dissertation proposes accelerated MRI acquisition and reconstruction techniques that account for the temporal dynamics of the MR signal. The methods build off of parallel imaging and compressed sensing to reduce scan time and flexibly model the temporal relaxation behavior. By randomly shuffling the sampling in the acquisition stage and imposing low rank constraints in the reconstruction stage, intrinsic physical parameters are modeled and their dynamics are recovered as multiple images of varying tissue contrast. Additionally, blurring artifacts are significantly reduced, as the temporal dynamics are accounted for in the reconstruction.

This dissertation first introduces T2 Shuffling, a volumetric technique that reduces blurring and reconstructs multiple T2-weighted image contrasts from a single acquisition. The method is integrated into a clinical hospital environment and evaluated on patients. Next, this dissertation develops a fast and distributed reconstruction for T2 Shuffling that achieves clinically relevant processing time latency. Clinical validation results are shown comparing T2 Shuffling as a single-sequence alternative to conventional pediatric knee MRI. Based off the compelling results, a fast targeted knee MRI using T2 Shuffling is implemented, enabling same-day access to MRI at one-third the cost compared to the conventional exam. To date, over 2,400 T2 Shuffling patient scans have been performed.

Continuing the theme of accelerated multi-contrast imaging, this dissertation extends the temporal signal model with T1-T2 Shuffling. Building off of T2 Shuffling, the new method additionally samples multiple points along the saturation recovery curve by varying the repetition time durations during the scan. Since the signal dynamics are governed by both T1 recovery and T2 relaxation, the reconstruction captures information about both intrinsic tissue parameters. As a result, multiple target synthetic contrast images are reconstructed, all from a single scan. Approaches for selecting the sequence parameters are provided, and the method is evaluated on in vivo brain imaging of a volunteer.

Altogether, these methods comprise the theme of *MR Shuffling*, and may open new pathways toward fast clinical MRI.

To JP Costa

Contents

Contents	ii
List of Figures	v
List of Tables	vii
List of Algorithms	viii
1 Introduction	1
2 MR Imaging and Reconstruction	5
2.1 MR Physics	5
2.2 MR Imaging	7
2.2.1 Frequency Domain Formulation	7
2.2.2 Generalized Image Reconstruction	12
2.3 Image Contrast	13
3 T2 Shuffling	17
3.1 Introduction	17
3.2 Theory	19
3.2.1 Temporal Subspace	22
3.2.2 Echo Train Ordering	25
3.3 Methods	28
3.3.1 Subspace Selection	28
3.3.2 Data Acquisition	30
3.4 Results	35
3.4.1 Retrospective Under-sampling and Simulation Results	35
3.4.2 In Vivo Results	38
3.5 Discussion	41
3.6 Conclusion	43
4 Clinical Deployment	44
4.1 Introduction	44

4.2	T2 Shuffling Reconstruction	45
4.3	Optimizations	46
4.3.1	Improving Multi-Core Parallelism	47
4.3.2	Fusing Forward and Adjoint Linear Operators	47
4.3.3	Removing Unnecessary and/or Redundant Computations	49
4.3.4	Image Quality Improvements	50
4.3.5	Distributed Implementation	50
4.4	Experimental Setup	51
4.4.1	Scan Sequence and Reconstruction Parameters	51
4.4.2	Experimental Compute Cluster	52
4.4.3	Hospital Compute Cluster	53
4.5	Results	53
4.5.1	Single-Machine Optimization	53
4.5.2	Multi-Machine Scaling	55
4.5.3	Hospital Cluster Statistics	56
4.5.4	Multi-Scanner Simulation	56
4.6	Conclusion	59
5	Clinical Validation	60
5.1	Introduction	60
5.2	Comparison to 2D Multi-Planar Knee MRI	61
5.2.1	Materials and Methods	61
5.2.2	Results	65
5.2.3	Discussion	72
5.3	Targeted Rapid Knee MRI Exam	74
5.3.1	Materials and Methods	74
5.3.2	Results	77
5.3.3	Discussion	82
5.4	Conclusion	84
6	T1-T2 Shuffling	85
6.1	Introduction	85
6.2	Theory	86
6.2.1	Signal Equation	86
6.2.2	Low Rank Modeling	87
6.2.3	Synthetic Contrast	90
6.2.4	Optimized Selection of Repetition Time Durations	91
6.3	Methods	94
6.3.1	Signal Modeling	94
6.3.2	Acquisition	94
6.3.3	Reconstruction	95
6.3.4	In Vivo Experiments	96

6.4	Results	96
6.5	Discussion	100
6.6	Conclusion	103
7	Conclusions and Future Work	105
7.1	Summary	105
7.2	Suggestions for Future Work	106
7.2.1	Improved Low-Dimensional Modeling	107
7.2.2	Multi-Parameter Shuffling	107
7.2.3	Parameter Mapping	108
7.2.4	Optimal Spatio-Temporal Sampling	108
7.2.5	Learned Reconstructions	108
A	Low Rank Regularization	109
B	Subspace Noise Propagation	112
C	T2 Shuffling Transform Point Spread Function	114
D	Sampling Pattern and Echo Train Formation	115
	Bibliography	117

List of Figures

2.1	Magnetization precession and excitation.	6
2.2	Longitudinal recovery and transverse relaxation.	6
2.3	Fourier-based forward model relationship.	7
2.4	Multi-coil forward model operator.	8
2.5	Forward model extended to fast spin-echo dynamics.	9
2.6	Comparison of regular vs. pseudo-random under-sampling.	10
2.7	Blurring due to fast spin-echo.	11
2.8	Fast spin-echo simulations using the Extended Phase Graph algorithm.	14
2.9	Contrast generation using fast spin-echo.	15
2.10	Illustration of image and signal behavior for PD, T2, and T1 imaging.	16
3.1	Extended forward model for a fast spin-echo experiment.	21
3.2	Simulation of a fast spin-echo experiment with variable refocusing flip angles.	23
3.3	T2 Shuffling forward model.	24
3.4	Comparison of center-out and randomly shuffled view ordering schemes.	26
3.5	Transform point-spread function for center-out and randomly shuffled view ordering schemes.	27
3.6	Process for generating an anatomy-specific subspace.	29
3.7	Subspace model error due to B1-inhomogeneity.	30
3.8	Subspace model error vs. noise amplification.	31
3.9	T2 Shuffling sampling patterns.	32
3.10	Retrospective simulation results for T2 Shuffling with variable flip angle modulation	36
3.11	Retrospective simulation results for T2 Shuffling with acquired constant flip angle data	37
3.12	Time series T2 Shuffling reconstructions on retrospectively simulated data.	38
3.13	T2 Shuffling reconstruction comparison on pediatric patient scans.	39
3.14	Example T2 Shuffling reconstructions with marginal improvement.	40
3.15	Degrees of freedom from LLR reconstruction.	41
4.1	Fused forward and adjoint T2 Shuffling operators.	48
4.2	Distributed T2 Shuffling flow chart.	51

4.3	Hospital compute cluster configuration.	53
4.4	Run-time of PICS portion across multiple slices and cores.	55
4.5	Processing run-times scaled from 1 to 12 machines.	56
4.6	MRI scanner schedule for one day.	57
4.7	Latency and processing time simulation for hospital network.	58
5.1	T2 Shuffling vs. 2D FSE for 15-year-old boy with medial retinaculum avulsion.	67
5.2	T2 Shuffling vs. 2D FSE for 19-year-old with intra-articular bodies of the knee.	68
5.3	T2 Shuffling vs. 2D FSE for 14-year-old girl with intra-articular bodies of the knee.	68
5.4	T2 Shuffling vs. 2D FSE for 16-year-old girl with lateral meniscus tear.	69
5.5	Tear of the posterior horn of the lateral meniscus requiring multi-planar reformat for visualization.	70
5.6	T2 Shuffling vs. 2D FSE structure delineation.	71
5.7	Radial tear of the lateral meniscus well-visualized with T2 Shuffling.	72
5.8	Clinical indications for targeted knee MRI exams.	77
5.9	Patient with bone bruise based on targeted knee MRI exam.	78
5.10	Patient with bone bruise based on targeted knee MRI exam.	78
5.11	Number of days between targeted knee MRI order and completion.	79
5.12	Exam time for targeted knee MRIs.	80
5.13	T2 Shuffling reconstruction time in seconds before deployment of distributed reconstruction.	81
5.14	T2 Shuffling reconstruction time in seconds after deployment of distributed reconstruction.	82
6.1	Fast spin-echo simulation of longitudinal and transverse magnetization with fast recovery.	88
6.2	T1-T2 Shuffling tensor construction and unfolding.	89
6.3	Generating a temporal subspace for T1-T2 Shuffling.	90
6.4	Image contrast synthesis through linear combination of subspace coefficients.	91
6.5	Creating joint (T1, T2) distribution from conventional quantitative mapping.	97
6.6	Relative loss of TR schedules.	98
6.7	T1-T2 Shuffling sampling patterns for each TE and TR.	99
6.8	T1-T2 Shuffling linear combination contrast synthesis fits.	100
6.9	Synthetic contrast images at several contrasts and orientations using T1-T2 Shuffling.	101
6.10	Comparison of conventional 3D FSE imaging to T1-T2 Shuffling.	103
7.1	T2 Shuffling scans of different anatomy.	106
A.1	Block-wise locally low rank operator.	111
B.1	Noise variance due to subspace back-projection.	113

List of Tables

3.1	Common scan parameters used for T2 Shuffling experiments.	33
3.2	Scan parameters for each T2 Shuffling experiment.	34
4.1	T2 Shuffling sequence and reconstruction parameters.	51
4.2	T2 Shuffling end-to-end processing times.	54
4.3	Reconstruction run-time statistics for 35 pediatric knee scans over a one-month period.	57
5.1	T2 Shuffling reconstruction parameters for clinical evaluation.	62
5.2	MRI scan parameters used for clinical evaluation.	63
5.3	Anatomic structures evaluted for clinical study.	64
5.4	Structure delineation ratings criteria for clinical study.	65
5.5	Discrepancies and consensus outcome in clinical study	66
5.6	Targeted knee MRI scan parameters.	75
5.7	Targeted knee MRI reconstruction parameters.	76
6.1	Optimal choice of TR durations through exhaustive search.	93
6.2	T1-T2 Shuffling scan parameters.	96
6.3	Optimized TR schedules.	97

List of Algorithms

4.1	FISTA update steps for T2 Shuffling inverse problem.	46
A.1	Locally low rank singular value thresholding.	110
D.1	Randomly shuffled sampling pattern and echo train formation.	116

Acknowledgments

Whenever I found myself reading over a dissertation, I would skip to the acknowledgments and dream about one day writing that section. It is with great fortune that I have now landed in such a position.

First and foremost, I have the great opportunity to thank Miki Lustig, my advisor. Miki embodies all the traits needed for great mentorship and teaching. He regularly took his role beyond the call of duty, giving me the means to succeed in my research as well as my life at Berkeley in general. Miki introduced me to MRI over the phone six years ago. Since then, not only has he unconditionally supported my research needs, he has also actively participated. At the same time, Miki afforded me a great degree of independence to build collaborations and take ownership of my work. Miki never hesitates to help make things happen. There's a reason people gravitate toward Miki in so many settings, and his dry sense of humor is matched only by the British. When I tell people he is the best advisor in our department, I say it with the utmost confidence. Thanks Miki for a great 6+ years.

I have also been unreasonably fortunate to work so closely with Shreyas Vasanawala. Shreyas introduced me to the world of clinical MRI and gave me a wheel to steer. In many ways, he is a second advisor to me. Shreyas taught me how to collaborate with the most important component of this research: the users. He has taken my work into clinical practice, and for that I am forever grateful.

I want to thank my other committee members, Chunlei Liu and Moriel Vandsburger, for their mentorship. Chunlei taught me the importance of preparation the first day I met him, and Moriel helped me grow as a teacher. I thank Brian Hargreaves, the Nicest Person on the Planet. Brian helped elevate my knowledge of MRI and I value his support and advocacy for my work. Brian also happily served on my qualifying exam committee even though it involved multiple train and bus rides! Thanks to Peder Larson for hosting me at UCSF and to Steve Conolly for great career advice. Thanks to Kannan Ramchandran for teaching me to love teaching and to Laura Waller for teaching me how to teach. I also want to thank Shirley Salanio, the most helpful person in the department. Shirley cares so much about the well-being of the students; it's no surprise that she knew all our names after one week of grad school.

I want to thank the people involved in the many scanning sites that enabled my research: Stanford University and Lucile Packard Children's Hospital, the Berkeley Brain Imaging Center, GE Healthcare ASL West, UCSF QB3, Oakland Children's Hospital, and Philips Healthcare. I have had the opportunity to scan at so many places!

I was very lucky to land in a lab with such a welcoming and collaborative environment. For three years Martin Uecker answered every little question I had and helped guide my research. I thank him for teaching me about computational MRI and for letting me work so closely with him on BART. Anita Flynn taught me how to manage a project and how to build a coil. Anita's desire to change the world is positively addicting.

Thanks to Pat Virtue who helped me assimilate to Berkeley life on and off the field. Pat was the first person to introduce me to machine learning. Beyond that, Pat showed me how to be an effective communicator. Joe Corea is a man who needs no introduction in the department. I thank him for answering my hardware questions and for sharing with me his outrageous ideas. I also thank his sweet tooth for finding us great gelato and snow cones during our travels.

Wenwen Jiang is notorious for knowing all the best food spots, no matter the city or country. I thank her for helping me with early pulse sequence questions and bouncing recon ideas off of each other. I want to thank Frank Ong, who despite being one year my junior has passed all the milestones with me side-by-side. Frank possesses a tremendous intuition for signal processing and MRI, and he is always willing to share it. He is also a great historian of the fields.

I want to thank Karthik Gopalan, a true Renaissance Man, for making the last years of grad school so enjoyable. Karthik exposed me to the world of digital fabrication and inspired me to buy a 3D printer. When we weren't in the lab making phantoms, Karthik was teaching me how to do an oil change, how to use a table saw, and how to dunk (okay, not quite). Thanks to Michael Kellman for lively discussions about system design, image recon, and just about any other topic. Michael never hesitates to offer an ear for feedback while also brainstorming new ideas. He's also the fastest swimmer on Alcatraz.

Thanks to Zhiyong Zhang, the newest, oldest member in MikGroup. In the short time he has been here, Zhiyong has helped me appreciably improve my understanding of pulse sequence theory. Zhiyong will make everyone in the room laugh, and often has a secret trick to share. Thanks to Mariya Doneva for helping me better understand image reconstruction and providing me with feedback on my work.

I thank the former undergrads in the lab who helped me explore new research ideas with them. I thank Sid Iyer, who is proudly continuing the MR Shuffling legacy at MIT. I also thank Debbie Ma and Eilam Levitov for eagerly collaborating on projects with me. I thank Dara Bahri for discussing all matters technical with me; I always look forward to his unannounced pop-ins at Berkeley. Finally, I want to thank the other members of MikGroup who have helped me on this journey: Balthazar Lechene, Xucheng Zhu, Alan Dong, Michael Driscoll, Alina Pechacek, Teresa Ou, Michal Zarrouk, and Sana Vaziri.

I am very thankful for my second research home at Stanford. I thank Joseph Cheng and Tao Zhang for helping me acclimate into pediatric imaging. Thanks to Marc Alley for hooking me up with the tools to do real, actual recons on patients. Thanks to Shanshan Bao and Albert Roh for showing me the radiologist perspective and to Kendall O'Brien for showing me the technologist perspective. I also thank other Stanford collaborators, including Feiyu Chen, Chris Sandino, Enhao Gong, and Adam Bush.

I was very fortunate to collaborate with industry experts in MRI and beyond. I thank Peng Lai, Weitian Chen, and Valentina Taviani for their endless technical support and help with bringing these sequences to life. Thanks to Mike Anderson and Javier Tuko, who helped me dip my toes into super computing. Many thanks to Yuval Zur; spending a winter under his wing stands among my most memorable experiences in grad school.

I owe much of my accomplishments to my roots at UT Austin. Thanks to Ted Rappaport, Brian Evans, and Sriram Vishwanath – three absolutely amazing teachers. Thanks to Miki Szmuk for being my Israeli comrade and for jump-starting my career in imaging. I have fond memories of our late nights in the lab trying to make an airplane fly.

I will always have a fondness for my time in Berkeley because of the great friends I made. Rachel Traylor, thanks for teaching me how to garden, for bringing Hampton in to my life, and for giving me a shoulder to lean on. Austin Buchan, thanks for showing me the ropes when I moved here, for the entertaining ski and backpacking trips, and for the endless philosophizing over vegetarian sushi. Emily Marron and Dan Drew, thanks for expressing my inner goofiness and letting me third-wheel on your dates. Thanks Varun Mishra for including me on so many fun outdoor adventures. Thanks Thomas Rembert and the rest of JASONS for all those great years playing IM soccer. Thanks to Andrew Harris and Lee Kittams for being great housemates. A very special thanks to Kyu Seob Kim, who isn't actually from Berkeley but would follow and get me to go on a number of adventures. Kyu, you rock! Also, many thanks to Joey Greenspun and the PCB crew.

I am grateful for my close friends from Austin who routinely provided me an escape from California. I would like to thank Matt Swarthout for patiently waiting for me to come around. Matt is one of the most talented comedians I know. I thank Casey Lee and Kevan Tavasoli, two of my best friends since high school. We did a lot of growing up together. Thanks Stephane Mead for supporting me since the beginning.

I want to thank my extended family and friends in Israel for the many years of support from abroad. They never hesitate to provide me with housing and friendship, even when I am just passing through for a quick trip. I also want to thank my newest family members, the Abrams clan. Steve, Michael, Ruth, and Hannah show no bounds for compassion and acceptance in making me feel welcome.

I want to thank Marcelle Friedman, my girlfriend and partner. Marcelle has an endless ability to make those around her feel cared for and important, and her support during the good times and the bad have kept me going. Thanks for following me on my trips, for offering so many means of escape, for supporting my chicken strips habit, for accommodating me no-questions-asked, and for always being there for me. When the world goes sour, you will be the first I write to: Eyzeh Bullsclachten.

To my family I owe everything. I thank my mom Smadar, who cultivated my interests and shovavnik tendencies. And I thank my dad Dan, who provided me with the inspiration to pursue a PhD. I am so happy we are able to “speak the same language”. Thanks to the Kiddos: Hagar, Daphne, and Michelle. You are the best seesters I could have ever hoped to get. I value my friendship with you above all else and I am excited for the next chapters in our lives. And thanks to Cooper, the fastest dog in the South.

I would like to end this note with thanks to my best friend JP Costa, to whom I dedicate this dissertation. JP, you were stripped from this world early, but you will not be forgotten. You continue to inspire my enthusiasm for life.

Jon Tamir
August 10, 2018

Chapter 1

Introduction

Magnetic resonance imaging (MRI) takes advantage of spin angular momentum present in atoms with an odd number of protons and/or neutrons [1]. These MR-relevant nuclei, called *spins*, can be manipulated in the presence of different magnetic fields to produce a detectable signal. Quite serendipitously, hydrogen spins in water are abundantly found in biological tissue, making MRI especially suitable for non-invasive medical imaging. MRI is used to visualize soft tissue contrast, capture motion, track functional and behavioral dynamics, spectrally decompose metabolites, assess blood flow, measure tissue stiffness, and more. These different types of image contrasts play an important role in diagnostic imaging, and each is realized by carefully manipulating a large number of user controls during the scan. Thus, a conventional clinical MRI exam consists of multiple scans that aim to capture different snapshots of the body by varying the scan parameters. Unlike computed tomography (CT), MRI does not use ionizing radiation, which makes it an especially attractive choice for pediatric imaging.

Despite the benefits, long scans times remain a major limitation of MRI. A typical MRI scan might take 5 minutes; in comparison, a full-body CT scan takes a few seconds. To lower scan time, concessions are often made in the acquisition process. These tradeoffs can manifest as reduced imaging resolution, image blurring, and lower signal to noise ratio (SNR). In spite of these tradeoffs, several acceleration methods have been key enablers of MRI in a clinical environment.

By far the most established method for acceleration is the RARE technique, also referred to as fast spin-echo (FSE) [2]. FSE has served as the workhorse of clinical MRI since its introduction in the late 1980s. The FSE acceleration is a result of acquiring multiple MRI measurements (called *echoes*) while the signal is in a transient decay, and the acceleration factor is proportional to the *echo train length* (ETL). However, the temporal effects due to signal relaxation manifest as blurring in the image [3]. Thus, short ETLs are often necessary to limit the image blur, costing scan efficiency. To increase scan efficiency, clinical FSE is typically used in two-dimensional (2D) imaging with thick slices and large gaps. Due to voxel anisotropy, each desired image contrast may be acquired in more than one of the principal directions – axial, sagittal, and coronal. Volumetric (3D) variants of

FSE are desirable for clinical imaging because they provide isotropic resolution [4]. This would eliminate the need to scan at multiple imaging planes, as the single volume could be retrospectively reformatted into arbitrary oblique planes.

Other advances to MR imaging and reconstruction have led to markedly faster scanning. Parallel imaging [5–7] takes advantage of the spatial sensitivities of receive arrays to acquire measurements in parallel and thus reduce the required scan time. Parallel imaging has become ubiquitous in clinic practice, with state of the art array sizes approaching 128 channels. A more recent trend for faster scanning is to leverage image priors in order to reduce the required number of samples for reconstruction [8]. One such avenue that is especially suited to MRI is compressed sensing [9, 10]. In compressed sensing MRI, data are acquired below the Nyquist rate using an incoherent measurement scheme such as pseudo-random sampling. The reconstruction then leverages structure such as sparsity [11] and low rank [12] to recover the image as if it were sampled at the Nyquist rate. Although parallel imaging and compressed sensing have enabled near order-of-magnitude acceleration, MRI scan times are still significantly longer compared to CT.

This dissertation describes new accelerated MRI methods that aim to improve image quality. The methods extend the MR signal model to account for temporal dynamics that occur during the image acquisition. Traditionally, these dynamics are left unmodeled and are one cause of image artifacts such as blurring. The methods presented in this dissertation build off of parallel imaging and compressed sensing to constrain the temporal signal behaviors to low-dimensional subspaces [13]. As a result, blurring effects commonly seen with FSE are significantly reduced, making 3D FSE a viable diagnostic tool. In addition to reducing image blur, the temporal signal evolution behavior is captured in the subspace. The information in the subspace is used to generate multiple images of varying contrasts, all from a single scan. Since the sequence is volumetric, each image contrast can be retrospectively reformatted into arbitrary planes while maintaining isotropic sub-millimeter resolution. The approaches are tested in real clinical environments and compared to traditional schemes.

Altogether, the methods comprise the theme of *MR Shuffling*.

Outline

The remainder of this dissertation is organized as follows:

Chapter 2: MR Imaging and Reconstruction

This chapter provides an overview of MRI from a frequency domain perspective and introduces image reconstruction as a generalized inverse problem. Different effects of sampling are reviewed in the context of parallel imaging and compressed sensing. Image contrast formation due to signal relaxation is also reviewed.

Chapter 3: T2 Shuffling

This chapter introduces a technique based off of 3D FSE called T2 Shuffling. In T2 Shuffling, relaxation effects are included in the forward model and accounted for in the reconstruction. Blurring artifacts are mitigated by constraining the temporal signal relaxation to a low dimensional subspace in the reconstruction. Additionally, multiple images are acquired from a single scan, corresponding to different levels of T2 contrast. The chapter reviews the fundamentals of FSE and signal relaxation, presents the acquisition and reconstruction technique, and shows in vivo results on volunteers and patients.

Chapter 4: Clinical Deployment

Since T2 Shuffling effectively performs a 4D compressed sensing-based reconstruction, it requires a large amount of computation for the reconstruction. In order to use T2 Shuffling clinically, the reconstructed images should ideally be available to the scanner operator before the patient leaves the table. This chapter discusses the steps taken to successfully deploy T2 Shuffling in a clinical environment. This includes optimizing subroutines of the reconstruction algorithm as well as implementing a distributed processing that is capable of running on multiple high-performance computer nodes. End-to-end processing times are reduced to under two minutes and multi-scanner scalability is discussed.

Chapter 5: Clinical Validation

In this chapter, T2 Shuffling is evaluated as a single-sequence alternative to pediatric knee MRI. The conventional 2D clinical protocol is compared to T2 Shuffling on patients at the Lucile Packard Children's Hospital at Stanford, California. Statistical comparisons are made to assess the hypothesis that T2 Shuffling does not reduce diagnostic accuracy or image quality when compared to the conventional 2D protocol. Additionally, a fast, targeted knee MRI exam based on T2 Shuffling is deployed and assessed. By providing a shorter overall exam time, patients sub-selected for insurance pre-authorization are able to register for their exam on the same day as their orthopedic visit. As a result of the shorter exam times, the cost is three times lower than that of the conventional exam. To date, over 2,400 T2 Shuffling patient scans have been performed.

Chapter 6: T1-T2 Shuffling

This chapter extends the concepts in T2 Shuffling to additionally acquire T1 recovery information. A single 3D sequence is presented that is able to reconstruct multiple synthetic contrast images corresponding to proton-density, T1 weighted, and T2 weighted contrasts. The acquisition collects incoherent measurements at multiple echo times and repetition times, and the reconstruction constrains the temporal dynamics to a low dimensional subspace. The optimal choice for TR durations is formulated based on ideas from

decision theory. The sequence is evaluated on a volunteer and compared to conventional 3D approaches.

Chapter 7: Conclusions and Future Work

This chapter summarizes the methods presented in this dissertation and outlines directions for future work.

Chapter 2

MR Imaging and Reconstruction

The technology behind MRI builds off of numerous science and engineering disciplines. Although nuclear magnetic resonance is governed by quantum mechanics, the principles of MRI are readily understood through the lens of classical physics. Using the concepts in [1] as a foundation, this chapter builds the necessary machinery for understanding the MR acquisition and reconstruction methods developed in this dissertation. In particular, this chapter presents the frequency domain formulation for MR imaging and a generalized reconstruction framework. Additionally, methods for creating image contrast from relaxation are reviewed.

2.1 MR Physics

Spins placed in a large external magnetic field (called B_0) precess and produce a net magnetic moment parallel to B_0 , called the longitudinal or z -direction (Figure 2.1a). The precession frequency is $\omega_0 = \gamma B_0$, where the gyromagnetic ratio γ is a constant that varies by atomic species. The spin dynamics of the magnetization vector are manipulated through the use of radio frequency (RF) and gradient magnetic fields. RF pulses played at frequency ω_0 in the transverse ($x \leftrightarrow y$) direction act to rotate or *excite* the magnetization vector, effectively mixing the longitudinal and transverse components. The strength and duration of the RF pulse determines the degree of rotation experienced. For example, a 90° flip angle fully rotates the magnetization into the transverse plane (Figure 2.1b). After excitation, the precession causes a changing magnetic field in the transverse direction which is picked up by a receive coil through Faraday's Law.

Linearly varying magnetic field gradients are used to create a spatial dependence on precession frequency. As a result, magnetization is spatially resolved by linearly mapping frequency to position in the received signal. This leads to a Fourier relationship: MRI measures the spatial Fourier transform of the object being imaged. By convention, the spatial frequency domain is called *k-space*. After the frequency space is sufficiently covered, the image is reconstructed through an inverse Fourier transform.




Figure 2.1: Magnetization (blue arrow) is initially pointed in the longitudinal direction with precession frequency ω_0 . After 90° RF excitation, the magnetization rotates into the transverse plane.

Following an RF pulse, the magnetization returns to thermal equilibrium due to relaxation. The transverse component of the magnetization exponentially decays to zero and the longitudinal magnetization exponentially recovers to thermal equilibrium, as shown in Figure 2.2. After a 90° RF excitation pulse, the relaxation is given by

$$M_z(t) = M_0 \left(1 - e^{-\frac{t}{T_1}}\right), \quad (2.1)$$

$$M_{xy}(t) = M_0 e^{-\frac{t}{T_2}}, \quad (2.2)$$

where $M_z(t)$ and $M_{xy}(t)$ are the longitudinal and transverse magnetization components, respectively, and M_0 is the magnetization at thermal equilibrium. The longitudinal and transverse time constants, referred to as T_1 and T_2 , respectively, are intrinsic physical parameters and differ across tissue types.

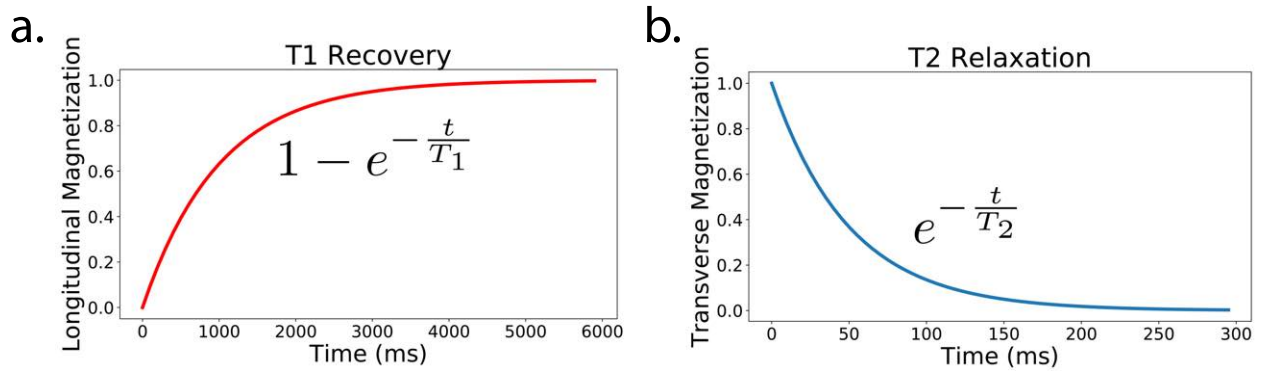


Figure 2.2: After 90° excitation, (a) the longitudinal magnetization will recover to thermal equilibrium with exponential time constant T_1 , and (b) the transverse magnetization will decay to zero with exponential time constant T_2 .

2.2 MR Imaging

2.2.1 Frequency Domain Formulation

The linear relationship between k-space and image space is formally described by the integral

$$s(t) = \int_{\mathbf{r}} M_{xy}(\mathbf{r}, t) e^{-j2\pi \mathbf{k}(t)^\top \mathbf{r}} d\mathbf{r} + w(t), \quad (2.3)$$

where $s(t) \in \mathbb{C}$ is the acquired signal at time t , $M_{xy}(\mathbf{r}, t)$ is the transverse component of the magnetization at position \mathbf{r} and time t , $\mathbf{k}(t)$ is the spatial encoding (k-space), and $w(t)$ is complex-valued white Gaussian noise. The symbol $(\cdot)^\top$ is used to denote the transpose operation. After change of variables and discretization, the MR signal equation is written as the linear operator

$$\mathbf{y} = \mathbf{E}\mathbf{x} + \mathbf{w}, \quad (2.4)$$

where \mathbf{x} is the image, \mathbf{y} are the acquired data in k-space, \mathbf{E} represents the spatial encoding, and \mathbf{w} is noise. When temporal magnetization effects are ignored, the forward model is

$$\mathbf{E} = \mathbf{P}\mathbf{F}, \quad (2.5)$$

where \mathbf{F} is the Fourier transform operator and \mathbf{P} is the operator that selects the acquired k-space samples. Figure 2.3 illustrates the forward model relationship. In the case of 3D Cartesian imaging, the acquired samples are represented as points on a 2D grid.




Figure 2.3: The discretized MRI forward model given by $\mathbf{y} = \mathbf{P}\mathbf{F}\mathbf{x}$. The MR image is Fourier transformed and measurements are acquired in k-space. In the case of 3D Cartesian imaging, the acquired samples are represented as points on a 2D grid.

Multiple Coils

The linear forward operator can be used to represent parallel imaging-based spatial encoding [5–7]. In the case of multiple receive coils, \mathbf{y} consists of multiple measurement vectors

$$\mathbf{y}_j = \mathbf{PFS}_j \mathbf{x} + \mathbf{w}_j, \quad j = 1, \dots, C \quad (2.6)$$

where \mathbf{S}_j and \mathbf{w}_j are the spatial sensitivity profile and noise at the j^{th} coil, respectively, and C is the number of coil elements in the array. The full forward model is illustrated in Figure 2.4 and described by

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_C \end{bmatrix} = \begin{bmatrix} \mathbf{PFS}_1 \\ \vdots \\ \mathbf{PFS}_C \end{bmatrix} \mathbf{x} + \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_C \end{bmatrix} = \mathbf{Ex} + \mathbf{w}. \quad (2.7)$$




Figure 2.4: Multi-coil forward model operator, \mathbf{PFSx} . The image is point-wise multiplied by the sensitivity profiles (\mathbf{S}) of the receive coil arrays ($C = 4$ channels). After Fourier transform (\mathbf{F}), the sampling operator \mathbf{P} is applied to each channel.

It is important to note that, because the received signals are acquired in parallel, the sampling operator \mathbf{P} is shared among all channels. It is common to represent (2.7) as $\mathbf{E}\mathbf{x} = \mathbf{P}\mathbf{F}\mathbf{S}\mathbf{x}$, where it is understood that the operators \mathbf{P} and \mathbf{F} operate on each channel independently. It is also convenient to assume that the noise \mathbf{w} is white and uncorrelated across channels. This is often accomplished through a noise pre-whitening step that precedes the reconstruction chain [14].

Multiple Time Points

Through a lifting procedure [15], the linear forward model can also be used to represent other imaging and encoding dimensions. For example, T2 relaxation dynamics in a fast spin-echo (FSE) experiment [2] can be represented as different echo time (TE) image states, $\mathbf{x} = [\mathbf{x}_1 \ \cdots \ \mathbf{x}_T]$, where \mathbf{x}_j is the image at the j^{th} TE and T is the echo train length (ETL). The forward model, illustrated in Figure 2.5, includes different encoding operators for time point i and coil j :

$$\mathbf{E}_{ij} = \mathbf{P}_i \mathbf{F} \mathbf{S}_j. \quad (2.8)$$

The same concept can be used to include additional dimensions such as motion states due to cardiac or respiratory motion [11, 16], spectral states for MR spectroscopy, dynamic contrast enhancement states [17, 18], and many others.




Figure 2.5: The forward model extended to T2 relaxation dynamics from FSE. Each image represents an echo time during the FSE echo train, $T = 4$ echoes in this example. The forward model includes the sensitivity map operator ($C = 2$ channels), the Fourier transform, and sampling. Each time point can be sampled differently, for instance using a center-out ordering of k-space.

Sampling

One way to analyze the effects of a specific sampling pattern operator is by passing an image through the forward and adjoint linear model, $\mathbf{E}^* \mathbf{E} \mathbf{x}$, where $(\cdot)^*$ is the adjoint (conjugate transpose) operation. This is equivalent to the point spread function (PSF) when the input is a point source at the center of the image. In single-channel Cartesian MRI, the PSF is the Fourier transform of the sampling pattern.

Many acceleration techniques rely on under-sampling k-space. Based on the sampling pattern, the effect on the image will differ. Figure 2.6 shows the result of sampling k-space below the Nyquist rate in two different ways. The left-most column shows the result of fully sampling k-space. In the middle column, the Nyquist rate is violated in the k_z direction. As a result, aliasing replicas occur indicating multiple coupled peaks in the PSF. In the right-most column, k-space is under-sampled using variable-density pseudo-random sampling. Because the sampling is under the Nyquist rate, aliasing still occurs. However, in this case the aliasing leads to artifacts that “look like noise,” and are incoherent with respect to the image.




Figure 2.6: Effects of sampling in k-space. (a). The result of sampling at the Nyquist rate. (b) Skipping every other k_z line leads to coherent aliasing. (c) Variable-density pseudo-random under-sampling leads to noise-like incoherent aliasing artifacts.

When neglecting relaxation, each sample is assumed to have the same signal level. In practice, however, the signal evolves while sampling. The effect can be mitigated by waiting a long time for longitudinal magnetization to fully recover between RF excitations.

This wait time, called the repetition time (TR), dictates the scan duration. However, this approach is not feasible for 3D imaging, as the scan time is equal to the number of phase encode points multiplied by the TR (e.g. 55 hours for 256^2 points when the TR is 3000 ms).

In this situation, the FSE sequence can be used to accelerate the acquisition by acquiring samples at multiple TEs during a single TR duration [2]. At each TE, a new line of k-space is acquired, affording an acceleration factor equal to the number of TEs (ETL). Since the TE of each line will differ based on the phase encode ordering, the signal level will vary in each k-space line due to T2 decay. When the temporal component of the acquisition is left unmodeled, the forward model amounts to modulating k-space phase encode lines with a T2-dependent relaxation. As Figure 2.7 shows, a center-out ordering of k-space lines is equivalent to multiplying k-space with a T2-dependent low-pass filter, effectively causing image blurring. Since relaxation depends on position, each voxel will undergo a different degree of blurring.

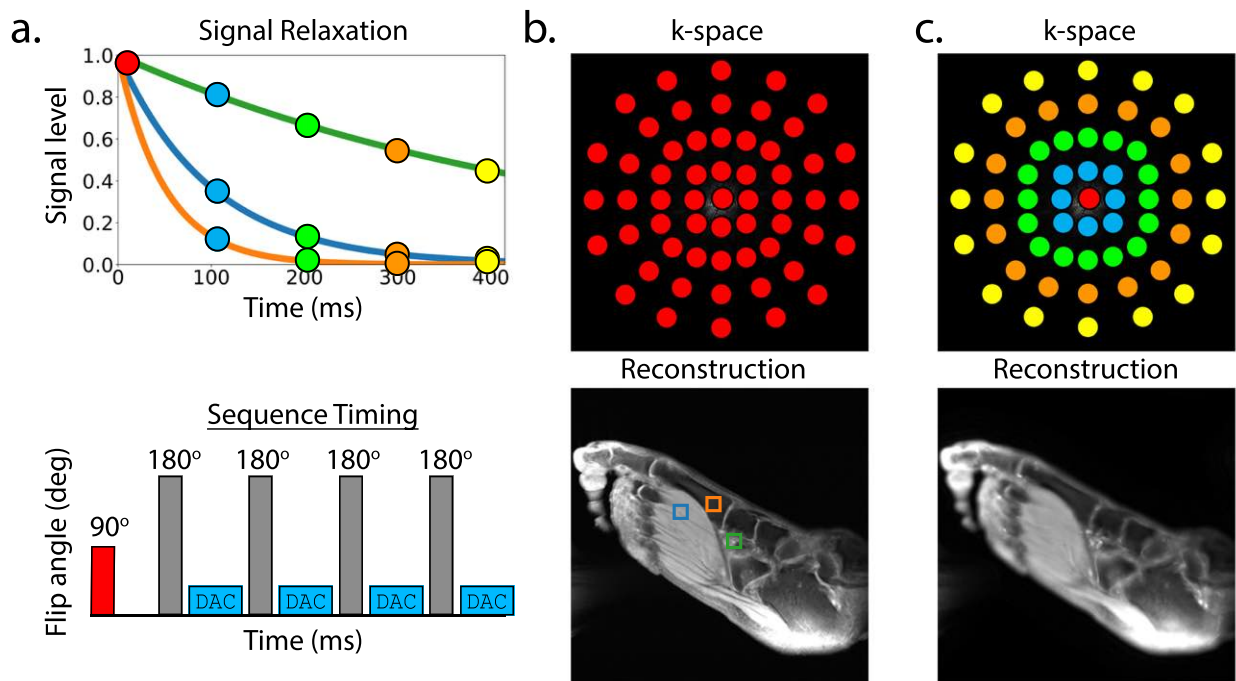


Figure 2.7: Effects of collecting samples at multiple TEs with FSE. (a) A 90° excitation pulse (red rectangle) followed by 180° refocusing pulses (gray rectangles) will lead to exponential decay due to T2 relaxation. (b) When a single sample is acquired in each TR duration (red circles), a sharp image can be reconstructed. To accelerate the acquisition, measurements can be acquired between each refocusing pulse, shown as digital-to-analog (DAC) blocks in (a). However, k-space will be modulated by T2 relaxation based on its acquisition time (multi-colored circles), leading to image blurring (c). Since T2 depends on position, each voxel will exhibit different amounts of blurring.

2.2.2 Generalized Image Reconstruction

Parallel imaging [5–7], compressed sensing [9, 10], and many other MRI reconstruction methods can be unified under a general inverse problem framework [19, 20]:

$$\arg \min_{\mathbf{x}} \quad \frac{\mu}{2} \|\mathbf{y} - \mathbf{E}\mathbf{x}\|_2^2 + \sum_{i=1}^M g_i(\mathbf{G}_i\mathbf{x}) \quad (2.9)$$

where the g_i are functionals applied to \mathbf{x} and the \mathbf{G}_i are linear transforms. The functionals can be used to represent constraints sets such as limited image support, real-valued and positivity, and other subspace constraints. Regularization can also be used to reduce noise amplification, promote image sparsity [11], or apply other priors on the images.

Even when all the g_i are convex, (2.9) is difficult to solve with general-purpose convex optimization methods because of the problem size. The ambient dimensions of the optimization variables are often in the tens of millions, necessitating matrix-free and Hessian-free approaches. Consequently, a number of efficient first-order iterative optimization algorithms that solve (2.9) have been developed [21–23]. These approaches are attractive because they support matrix-free methods and decompose each iteration into a linear data consistency update and point-wise proximal updates.

Compressed Sensing MRI

Compressed sensing MRI has seen great success in research settings, spawning thousands of technological developments and clinical applications. Recently, MRI vendors have released clinical packages that incorporate compressed sensing into their products [24–26]. The incoherence that results from sampling k-space in a pseudo-random fashion makes compressed sensing especially suited to MRI [11, 27]. Through a suitable transform (e.g. wavelets for images), a sparse or low-dimensional representation of the image can be used to dramatically reduce the required number of measurements for reconstruction. The transform PSF (TPSF) extends the PSF concept to the transform domain and is useful for analyzing the sampling behavior on the sparse coefficients themselves [11].

Using the generalized image reconstruction framework, parallel imaging and compressed sensing can be formulated as the Lasso [28],

$$\arg \min_{\mathbf{x}} \quad \frac{1}{2} \|\mathbf{y} - \mathbf{E}\mathbf{x}\|_2^2 + \lambda \|\Psi\mathbf{x}\|_1, \quad (2.10)$$

where Ψ is a sparsifying transform and λ is a regularization term. In this dissertation, the combination of generalized parallel imaging and compressed sensing is referred to as PICS.

Low Rank Regularization

One generalization of compressed sensing is the recovery of low-rank matrices through nuclear norm minimization [8, 12]. Low rank structure is a useful model in dynamic

MRI because the images often exhibit global and local spatio-temporal correlation. This dissertation heavily uses low rank concepts to efficiently model relaxation effects during imaging. One way to apply low rank regularization is by including a nuclear norm functional

$$g_i(\mathbf{Z}) = \lambda \|\mathbf{Z}\|_*, \quad (2.11)$$

where the nuclear norm is defined as the sum of the singular values of the matrix \mathbf{Z} . Low rank regularization methods are elaborated on in Appendix A.

2.3 Image Contrast

Many types of image contrasts can be created with MRI. This dissertation focuses on contrast due to T1 recovery and T2 relaxation. After 90° excitation, the magnetization evolves over time as described by

$$M_{xy}(\mathbf{r}, t) = \rho(\mathbf{r}) f_t(\mathbf{r}), \quad (2.12)$$

where $\rho(\mathbf{r})$ is the intrinsic amount of magnetization at position \mathbf{r} (called proton-density, or PD), and $f_t(\mathbf{r})$ is a spatio-temporal signal evolution. When measurements are made at multiple TEs separated by an echo spacing T_s , it is sufficient to consider the magnetization at $t = kT_s$, where $k \in \{1, \dots, T\}$ indexes the TEs and T is the ETL. The extended phase graph (EPG) formalism [29, 30] is one approach to analyzing the effects of a series of RF excitation blocks, relaxation blocks, and gradient dephasing blocks on the magnetization signal evolution.

The EPG algorithm simulates these processes as instantaneous events given by linear operators. RF excitation mixes transverse and longitudinal magnetization; relaxation applies T1 recovery and T2 decay; gradients increase the cycles of phase across spins in a voxel. Figure 2.8 shows two examples of signal evolutions simulated with the EPG algorithm corresponding to the FSE pulse sequence [2]. The first uses a train of 180° RF refocusing flip angles, and the magnetization evolution at position \mathbf{r} and the k^{th} TE is given by

$$M_{xy}(\mathbf{r}, kT_s) = \rho(\mathbf{r}) e^{-\frac{kT_s}{T_2(\mathbf{r})}} \left(1 - e^{-\frac{\text{TR} - T \cdot T_s}{T_1(\mathbf{r})}} \right), \quad (2.13)$$

where TR is the repetition time duration. The second uses a train of variable refocusing flip angles and leads to a more complex spin evolution due to the continual mixing of the longitudinal and transverse magnetization components.

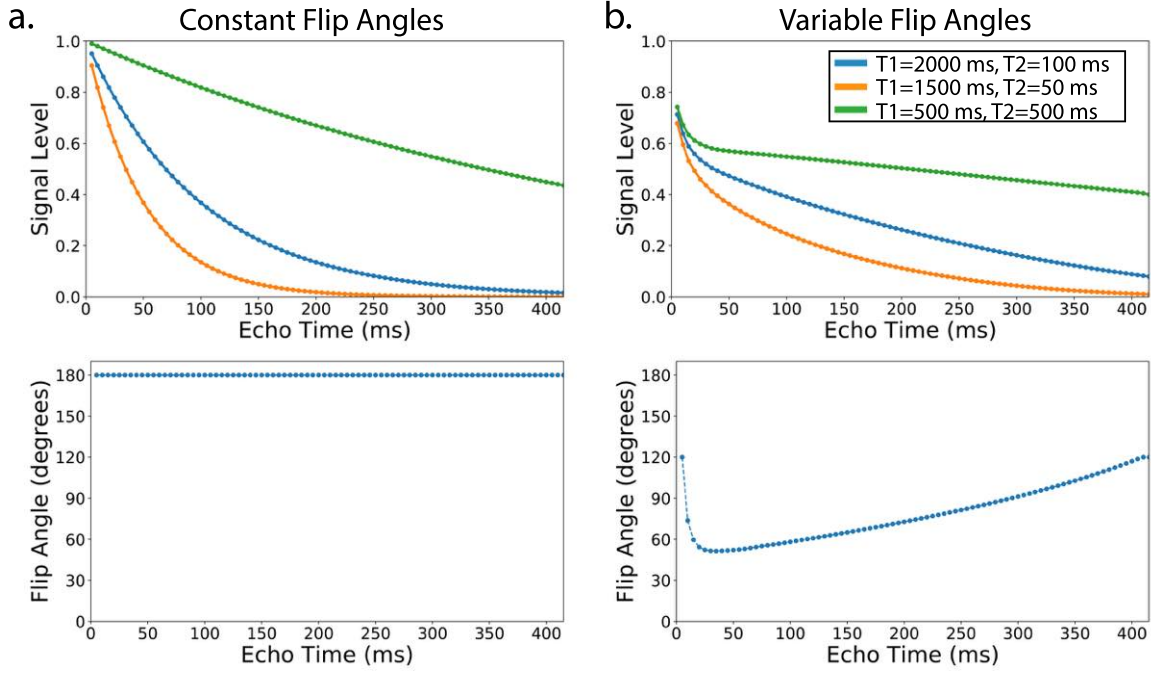


Figure 2.8: Fast spin-echo simulations using the Extended Phase Graph algorithm. Different combinations of T1 and T2 are plotted in different colors. (a) A series of refocusing RF pulses with 180° flip angles (bottom) leads to pure exponential decay (top). (b) When variable RF refocusing flip angles are used, the signal relaxation curves are a result of a mixture of T1 recovery and T2 decay.

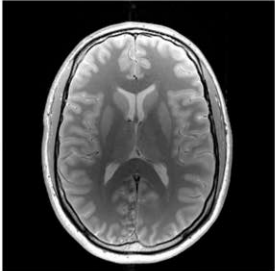
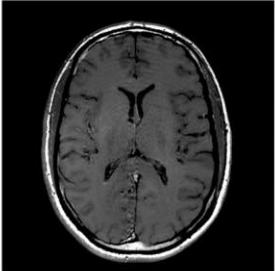
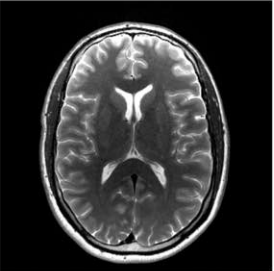
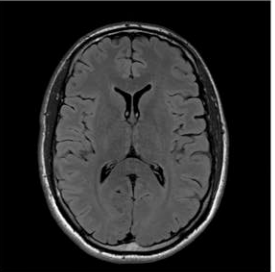
A common magnetization preparation technique in FSE is to precede the 90° excitation with a 180° inversion pulse. This inverts the longitudinal magnetization to the $-z$ direction, effectively doubling the dynamic range as the magnetization recovers back to thermal equilibrium [1]. After a suitable inversion time (TI) duration, the sequence continues as before with a 90° excitation. Based on the TI duration, the magnetization of different tissues will be nulled depending on their T1-dependent zero-crossing times. This is the basis for fluid-attenuated inverse recovery (FLAIR) contrast, in which the TI is chosen so that signal from fluid is nulled [31]. When 180° refocusing angles are used, the transverse magnetization at the k^{th} TE is given by

$$M_{xy}(\mathbf{r}, kT_s) = \rho(\mathbf{r}) \left(1 - 2e^{-\frac{\text{TI}}{T_1(\mathbf{r})}} e^{-\frac{\text{TR}-T \cdot T_s}{T_1(\mathbf{r})}} \right) e^{-\frac{kT_s}{T_2(\mathbf{r})}}, \quad (2.14)$$

where TI is the inversion time. It is possible to fit the signal dynamics in (2.13) and (2.14) to estimate T1 and T2 given images scanned at multiple time points (TEs, TRs, and TIs) [32–34]. In practice, different tissues comprise a single voxel, leading to a distribution of T1 and T2.

The FSE sequence flexibly supports different image contrasts based on PD, T1, and T2 by varying the three main pulse timing parameters: TE, TR, and TI. This concept is illustrated in Figure 2.9 for different types of image contrasts used in routine brain imaging

[35] at 3 Tesla. The sequence starts with an optional inversion recovery preparation module with duration TI , followed by 90° excitation and a train of refocusing pulses. When the refocusing pulses are each 180° , the magnetization is given by either (2.13) or (2.14).

Proton Density	T1 Weighted	T2 Weighted	T2 FLAIR
			
TE = 10 ms	TE = 10 ms	TE = 100 ms	TE = 100 ms
TR = 6000 ms	TR = 500 ms	TR = 6000 ms	TR = 9000 ms
TI = N/A	TI = N/A	TI = N/A	TI = 2250 ms

FSE Sequence Timing




Figure 2.9: Different image contrasts scanned at 3 Tesla (top) based on PD, T1, and T2 can be created using the FSE sequence (bottom). Inversion pulses are shown in orange, excitation pulses in red, and variable refocusing pulses in gray. TE imaging blocks are marked with blue DACs. PD, T1 weighted, and T2 weighted images are created by varying the TE and TR, given by (2.13). An inversion recovery preparation is included to produce images with nulled tissue such as T2 FLAIR, given by (2.14).

A PD weighted image is created using a short TE to limit T2 relaxation and a long TR to ensure full T1 recovery. The remaining contrast is then proportional to the intrinsic spin density of each voxel. A T2 weighted image also uses a long TR so that each repetition starts with full longitudinal magnetization, but uses a long TE to add contrast from T2 relaxation. A T1 weighted image is created using a short TE and a short TR. The short TE limits T2 relaxation, and the short TR limits the degree of longitudinal recovery before the next excitation. As a result, the initial magnetization is weighted by the partial T1 recovery. An alternative approach to T1 weighted imaging is to include the inversion recovery preparation pulse, called T1 FLAIR. T1 FLAIR provides additional tissue contrast due to

the extended dynamic range [36]. By combining the fluid-nulling inversion preparation with a long TE and a long TR, a T2 FLAIR image is created.

Figure 2.10 illustrates the magnetization evolution for PD, T2, and T1 imaging. The scan parameters lead to signal differences between tissue types, indicating different types of image contrast. PD imaging aims to minimize signal level differences, while T2 weighted images aim to maximize signal differences due to T2 relaxation. For T1 weighted imaging, it is useful to track the longitudinal magnetization immediately before it is excited into the transverse plane.

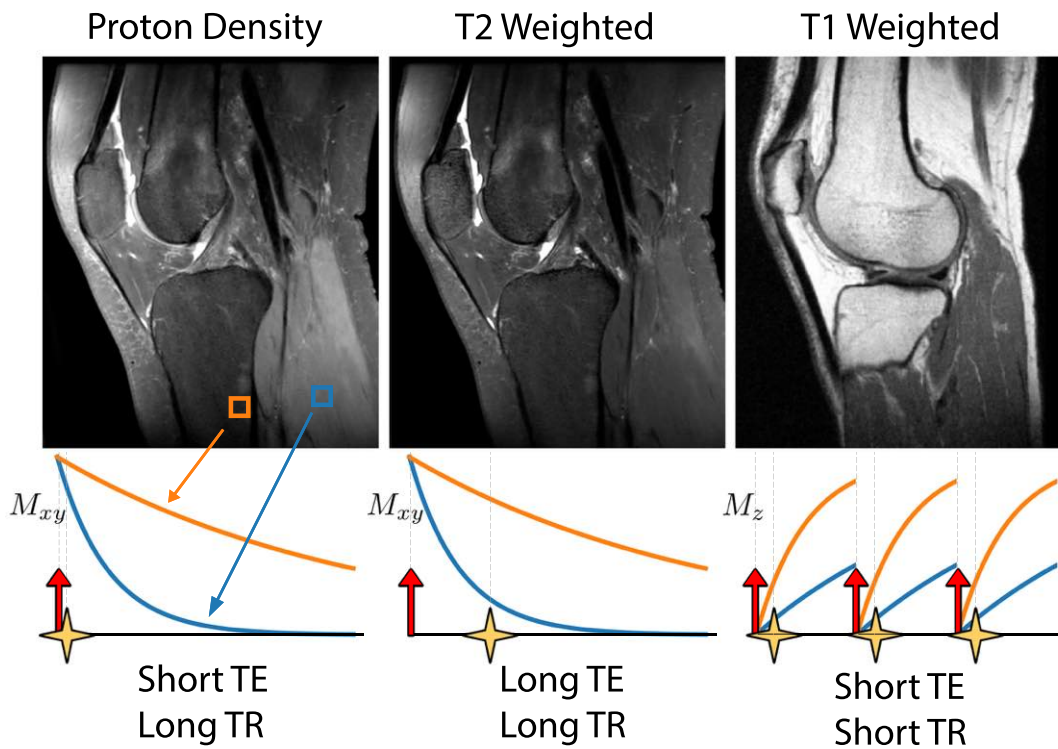


Figure 2.10: Proton-density, T2 weighted, and T1 weighted contrast generation by choosing different combinations of TE and TR. The RF excitation pulses are shown as red arrows and the TEs are marked with yellow stars.

Chapter 3

T2 Shuffling

3.1 Introduction

Fast spin-echo sequences are commonly used for MR imaging because of the ability to prescribe various image contrasts at target scan times in a flexible manner [3]. Most clinical musculoskeletal protocols separately acquire 2D FSE images at multiple orthogonal planes, and repeat the set of scans for different image contrasts. Altogether, the full exam time is lengthy and the 2D images suffer from relatively thick slices with large gaps, which can obscure pathology [37–39]. Volumetric (3D) FSE is an attractive alternative to 2D FSE because it provides isotropic resolution and the images can be reformatted in arbitrary orientations to evaluate pathology [3]. However, long echo trains are required to maintain scan efficiency, leading to blurring due to T2 decay [4, 38, 40]. Parallel imaging [5–7] in tandem with variable flip angle modulation e.g. using the CUBE (GE Healthcare, Waukesha, WI), SPACE (Siemens Healthineers, Erlangen, Germany), and VISTA (Philips Healthcare, Best Netherlands) pulse sequences, can reduce this effect by shortening echo trains and reducing the signal decay during the echo train, but blurring often persists, limiting diagnostic utility. Recent applications of compressed sensing to 3D FSE [39, 41, 42] have also shown promise in alleviating the tradeoff between scan time and blurring.

Cartesian 3D FSE scans acquire multiple k-space samples, i.e. an echo train, in each RF excitation. The speed of the acquisition depends primarily on the pulse TR, the ETL, and the parallel imaging acceleration factor. Conventional reconstruction neglects the echo train time progression between acquired phase encodes and produces a single image. Since the signal decays during the echo train, blurring artifacts corrupt the final image, leading to loss in apparent resolution.

A number of approaches to deblurring for 3D FSE exist. Several (i) first estimate T2 and (ii) next deconvolve the echo train signal evolution based on this estimate [43–45]. These methods are sensitive to errors in the T2 estimate, and can amplify noise in the deconvolution (high pass filtering) step. Other approaches exploit the Hermitian symmetry property of k-space to reduce the effect of T2 decay [46, 47], but are sensitive to phase er-

rors that invalidate the Hermitian symmetry. The above methods do not fundamentally change the signal decay behavior, and instead aim to correct it with post-processing.

Many approaches aim to directly recover tissue parameters from under-sampled data through a model-based, non-linear optimization [48–51]. These approaches often use non-Cartesian sampling to leverage incoherence in the acquired signal and improve the conditioning of the reconstruction. Nonetheless, the reconstructions must cope with added complexity in the non-linear forward model and are sensitive to model mismatch, e.g. due to partial voluming. Other works first impose subspace constraints through principal component analysis (PCA) [52–56] and dictionary learning [57, 58] to linearize the parametric forward model and enforce spatial sparsity on the basis images. The linearization is attractive because it lends itself to a convex formulation and PCA can be applied to training signals that account for stimulated echoes and imperfect slice profiles. Data-driven variants [59, 60] can also be used to build robustness to motion [61] and other non-ideal imaging considerations. The linear subspace constraint is robust to partial voluming and implicitly accounts for multi-compartmental models, as linear combinations of signal evolutions remain in the subspace. However, the approximation error from PCA and other basis expansions may lead to biased parameter estimates and the parametric fitting step must still explicitly account for imaging non-idealities such as partial voluming and RF-field inhomogeneity. Parametric estimation is further complicated by the subtle dependence of tissue relaxation on pulse sequence design [62]. More recently, pattern matching methods [63, 64] have shown promise at deriving tissue and system parameters. These methods employ pixel-wise matching of unknown signal evolutions to a database of simulated signals. However, they often neglect data fidelity in the reconstruction formulation. The addition of a data fidelity term has been shown to improve the reconstruction [65–68].

This chapter presents a subspace-constrained method, called T2 Shuffling, that simultaneously recovers many images at multiple contrasts with increased sharpness from a single acquisition that is about as long as a conventional 3D scan. T2 Shuffling reduces the blurring due to long echo trains and eliminates the need to perform separate scans at different target echo times (TEs) for each image contrast. The acquisition accounts for the echo train time progression, and the reconstruction combines recent advances in parallel imaging and compressed sensing in a software framework designed for a clinical workflow. Specifically, T2 Shuffling builds off of previous subspace [52–55] and low rank [18, 69–71] modeling methods to reduce the degrees of freedom in the reconstruction. In contrast to the subspace-based parametric methods, no parameter maps are explicitly estimated, and instead the linear subspace constraint is used as a flexible approach to de-blurring and recovering tissue dynamics. The method thus maintains the implicit robustness to imaging non-idealities present in other subspace-constrained approaches. Also, rather than regularizing in the parametric dimension, low-rank constraints are imposed directly on the subspace.

By accounting for signal decay, scan efficiency can be improved by acquiring more data in each TR without suffering from increased blurring. The wait time between the

end of the echo train and the start of the next TR has been previously used to jointly produce a proton-density and a T2-weighted image from a single 3D FSE scan [72]; here this observation is taken one step further – an image is recovered at each TE in the echo train. The authors in [73] and [74] have demonstrated a similar principle through the use of a radial acquisition with a view-sharing reconstruction for T1-weighted imaging. In contrast, here view sharing is replaced with an explicit subspace model that constrains the signal relaxation.

3.2 Theory

A model-based description for FSE acquisition is that the observed signal is a linear mixture of a proton-density image and a temporal signal evolution function:

$$y(t) = \int_{\mathbf{r}} \rho(\mathbf{r}) f_t(\mathbf{r}) e^{-j2\pi \mathbf{k}(t)^\top \mathbf{r}} d\mathbf{r}. \quad (3.1)$$

Here, $y(t)$ is the acquired signal at time t , $\rho(\mathbf{r})$ is the complex proton-density image at voxel \mathbf{r} , and $\mathbf{k}(t)$ is the vector-valued k-space trajectory. This model is easily extended to multiple receive coils in a straightforward manner. The temporal signal evolution function $f_t(\cdot)$ depends on several factors, including tissue relaxation and the refocusing flip angle schedule. For a constant 180 degree flip angle train, a simple and common model used for FSE imaging is the mono-exponential decay. In this model, each voxel position \mathbf{r} decays with $T_2(\mathbf{r})$ during the echo train and recovers with $T_1(\mathbf{r})$ after the echo train over a repetition time TR according to

$$f_t(\mathbf{r}) = e^{-\frac{t}{T_2(\mathbf{r})}} \left(1 - e^{-\frac{\text{TR}-T_s}{T_1(\mathbf{r})}} \right), \quad (3.2)$$

where T is the ETL and T_s is the echo spacing. In practice, each voxel will contain a distribution of relaxation values. In conventional FSE image reconstruction, the time progression is ignored, and an inverse Fourier transform is applied to (3.1), implicitly assuming that $f_t(\cdot)$ does not change over time. For example, it is often assumed that $f_t(\mathbf{r}) \approx f_{\text{TE}}(\mathbf{r})$ for all t , where TE is the (fixed) echo time when the center of k-space is acquired [3]. To appreciate the impact of this approximation, consider $f_t(\cdot)$ in terms of k-space position using the relationship $\mathbf{k}(t) = \frac{\gamma}{2\pi} \mathbf{G}(t - \text{TE})$, where \mathbf{G} is the gradient amplitude vector. For the special case of mono-exponential decay and neglecting the positional dependence of relaxation,

$$f_k = e^{-\frac{\mathbf{k}^\top (1/\mathbf{G})}{\frac{\gamma}{2\pi} T_2}} e^{-\frac{\text{TE}}{T_2}} \left(1 - e^{-\frac{\text{TR}-T_s}{T_1}} \right), \quad (3.3)$$

where $1/\mathbf{G}$ is the point-wise inverse of the gradient vector. As the k-space weighting in (3.3) is convolved in image space with the proton-density image, the approximation

impacts the apparent tissue contrast and leads to blurring due to T2 decay during the acquisition [4, 40]. Since relaxation depends on position, each voxel will undergo a different degree of blurring [45].

In this work, the temporal nature of the acquisition is modeled, and the Fourier relationship is maintained through a lifting procedure [15]. This does not assume a mono-exponential decay, and instead allows arbitrary signal evolutions formed by a distribution of T2 values. Consider an FSE experiment in which T echoes are acquired each TR with echo spacing T_s . Let

$$x_i(\mathbf{r}) = \rho(\mathbf{r})f_{iT_s}(\mathbf{r}) \quad (3.4)$$

be the magnetization at the i^{th} echo time and the \mathbf{r}^{th} voxel, and define the signal evolution at voxel \mathbf{r} as

$$\mathbf{x}(\mathbf{r}) = [x_1(\mathbf{r}) \quad x_2(\mathbf{r}) \quad \cdots \quad x_T(\mathbf{r})]^T \in \mathbb{C}^T. \quad (3.5)$$

In this notation, \mathbf{x}_i is the virtual echo time image at the i^{th} virtual echo time and \mathbf{x} is the collection of all echo time images. Thus, the forward model is extended to a time series of images:

$$\mathbf{y}_i = \mathbf{P}_i \mathbf{F} \mathbf{S} \mathbf{x}_i, \quad i = 1, \dots, T, \quad (3.6)$$

$$\mathbf{y} = \mathbf{P} \mathbf{F} \mathbf{S} \mathbf{x} := \mathbf{E} \mathbf{x}. \quad (3.7)$$

The encoding matrix \mathbf{E} operates on each virtual echo time image independently. Each image \mathbf{x}_i is point-wise multiplied by the coil sensitivity maps (\mathbf{S}) and Fourier transformed (\mathbf{F}). The data are then multiplied with the sampling mask of each virtual TE (\mathbf{P}), where the entries of the sampling pattern correspond to the acquired phase encodes at each virtual TE, and represented in k - t space by \mathbf{y} . Figure 3.1 depicts the forward model for the time-series of images. Figure 3.1b shows the sampling operator corresponding to a center-out phase encode ordering, which is commonly used for proton-density FSE [40]. In conventional 3D FSE, the acquired data are collapsed into a single k -space matrix, which in the lifted formulation is equivalent to summing along the time dimension. By representing the data in the lifted space, the model supports arbitrary sampling and is not restricted to a center-out view ordering. Figure 3.1c shows a *randomly shuffled* sampling procedure that leads to a better posed inverse problem [75], as discussed in the sequel.

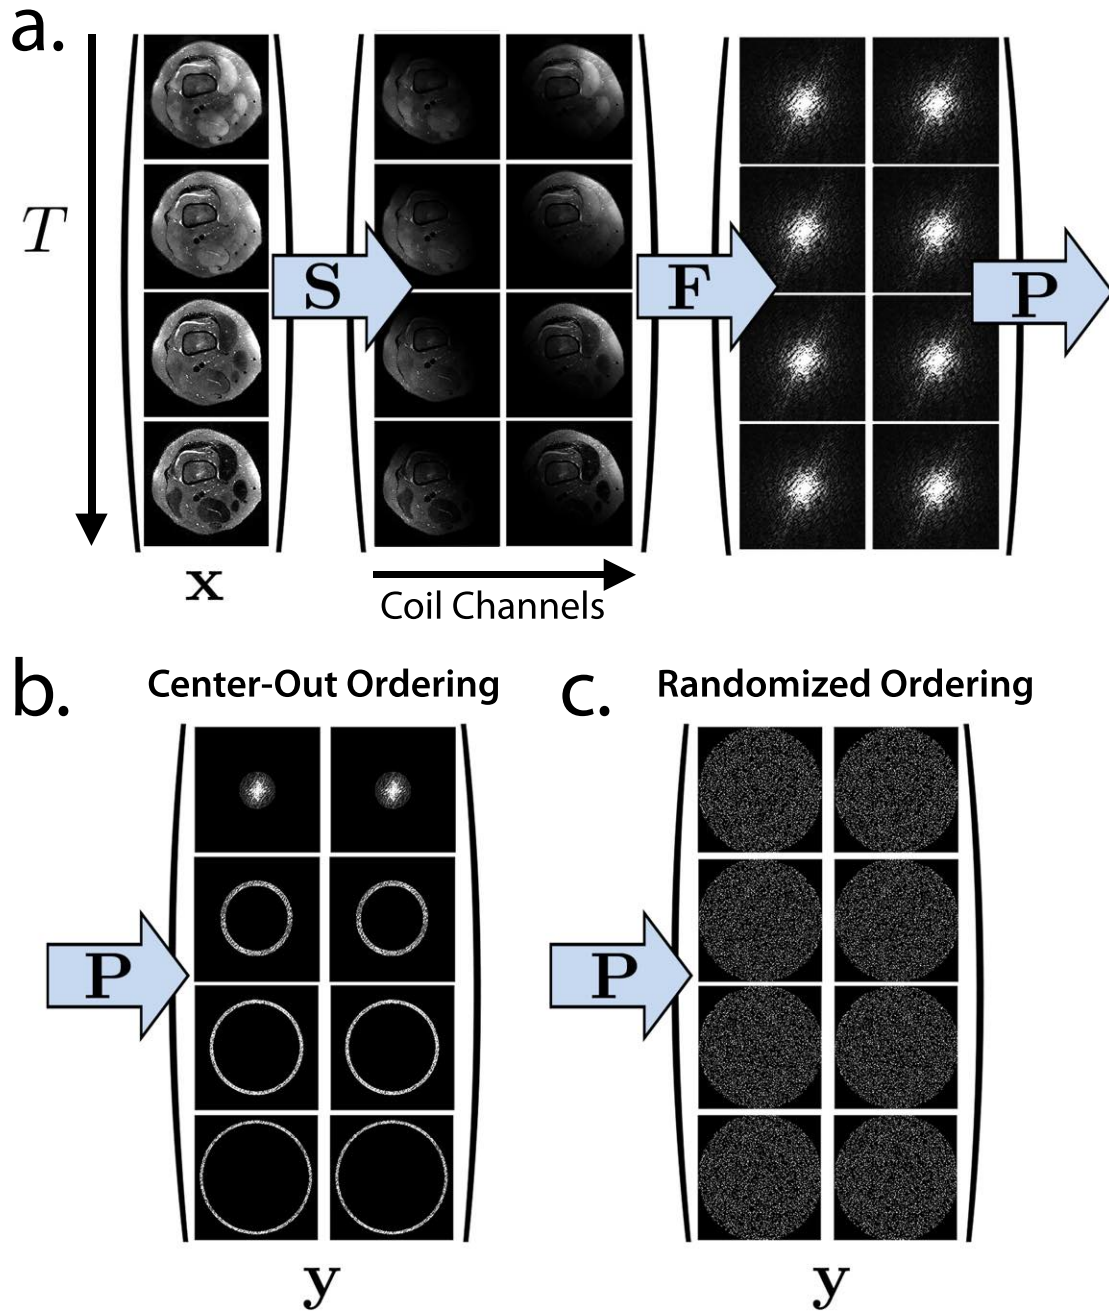


Figure 3.1: The forward model is extended to a time series of images at T virtual echo times. (a) Each virtual echo time image is point-wise multiplied by the sensitivity maps, Fourier transformed, and sampled. (b) A center-out view ordering used in conventional proton-density 3D FSE imaging. (c) A *randomly shuffled* view ordering suitable for T2 Shuffling.

Since the ambient dimension of the unknown images is LT , where L is the number of voxels, a naive well-posed reconstruction in the lifted space requires T times the number of observations compared to a standard 3D FSE reconstruction. To overcome this limitation, structure in the data is exploited to reduce dimensionality. This is achieved in two ways: first, the temporal signal evolution is modeled and represented by a low-dimensional subspace; second, the sampling problem is formulated based on ideas from compressed sensing [11]. These two approaches significantly reduce the dimensionality to a regime where the equivalent of a single 3D FSE acquisition can be used to reconstruct the full time series of images.

3.2.1 Temporal Subspace

The signal evolutions observed in an FSE experiment are functions of both the tissue (e.g. T1, T2) and system parameters (e.g. flip angles in the echo train). Other factors such as B_1 inhomogeneity will also influence the signal evolution [64]. Using these parameters, signal evolutions can be generated with Bloch simulation and the Extended Phase Graph (EPG) algorithm [4, 29, 64]. In the case of a constant flip angle schedule, there also exist analytical solutions [50, 76]. Figure 3.2 shows example signal evolutions using the variable flip angle modulation scheme described in [40]. The signal evolutions corresponding to different tissues follow similar trends despite differing in relaxation parameters. It is this correlation that makes parameter mapping a difficult problem. However, the correlation implies that the signal evolutions of different tissues form a low-dimensional subspace [13].

Consider an ensemble of N spins sampled at T echo times, $\mathbf{X} \in \mathbb{C}^{T \times N}$. Each column in \mathbf{X} represents the signal evolution of a spin with a particular (T_1, T_2) pair. Let $\Phi \in \mathbb{C}^{T \times T}$ be an orthonormal temporal basis, i.e.

$$\mathbf{X} = \Phi \Phi^* \mathbf{X}. \quad (3.8)$$

The goal is to design $\Phi = [\varphi_1 \cdots \varphi_T]$ and a K -dimensional subspace, $\Phi_K = \text{span}\{\varphi_1, \dots, \varphi_K\}$, such that

$$\|\mathbf{X} - \Phi_K \Phi_K^* \mathbf{X}\| < \epsilon, \quad (3.9)$$

where ϵ is a modeling error tolerance. The choice of norm in (3.9) will affect the chosen subspace. If the Frobenius norm is used, then this metric captures the sum-of-squares error in the fit. If a max-column norm is used, then the tolerance can be interpreted as the worst-case error on any signal evolution. By choosing the ensemble \mathbf{X} to match the distribution of (T_1, T_2) values in the tissue of interest, a suitable subspace that minimizes the Frobenius norm can be generated using PCA [53, 54, 77].

In conventional 3D FSE, variable refocusing flip angles are used to reduce the rate of decay for a particular range of T2 values [3, 4, 40]. This leads to less intrinsic blurring in the final image. The choice of refocusing flip angles for the echo trains has a significant




Figure 3.2: (a) Simulated signal evolutions during an echo train using variable refocusing flip angles shown in (b) ($T_1 = 1000$ ms). For $T_2 = 500$ ms, the signal is nearly constant after the initial 20 ms and will lead to minimal blurring in the image. For $T_2 = 100$ ms and below, the signal decays throughout the echo train. Despite the different temporal profiles, the signal curves are highly correlated.

impact on the signal evolution, and thus the temporal subspace design. Since the model in (3.7) accounts for the temporal signal, it is no longer necessary to design the flip angle train to reduce signal decay. Instead, other metrics such as signal to noise ratio (SNR) and contrast to noise ratio (CNR) could be used as flip angle design criteria. For this work, the flip angle modulation scheme was not changed from that described in [40].

Using the subspace Φ_K ,

$$\mathbf{y} = \mathbf{E}\mathbf{x} \quad (3.10)$$

$$= \mathbf{E}\Phi\Phi^*\mathbf{x} \quad (3.11)$$

$$\approx \mathbf{E}\Phi_K\Phi_K^*\mathbf{x} \quad (3.12)$$

$$:= \mathbf{E}\Phi_K\boldsymbol{\alpha}, \quad (3.13)$$

where $\alpha := \Phi_K^* \mathbf{x}$ are the K temporal basis coefficients that describe the time series of images. Neglecting model error,

$$\mathbf{x} = \Phi_K \alpha. \quad (3.14)$$

The reconstruction problem can now be posed in terms of α :

$$\min_{\alpha} \frac{1}{2} \|\mathbf{y} - \mathbf{E} \Phi_K \alpha\|_2^2 + \lambda \sum_{\mathbf{r}} \|R_{\mathbf{r}}(\alpha)\|_*. \quad (3.15)$$

After solving for α , the unique time series is computed by projecting through the basis using (3.14). The subspace constraint significantly reduces the dimensionality of the unknown variable in the optimization problem. Rather than resolving each image voxel at $T \approx 80$ time points, it is sufficient to represent each voxel with $K \approx 4$ coefficients [53, 55]. The full T2 Shuffling forward model is shown in Figure 3.3. The subspace coefficient images are explicitly used to generate the time series of images through back-projection.

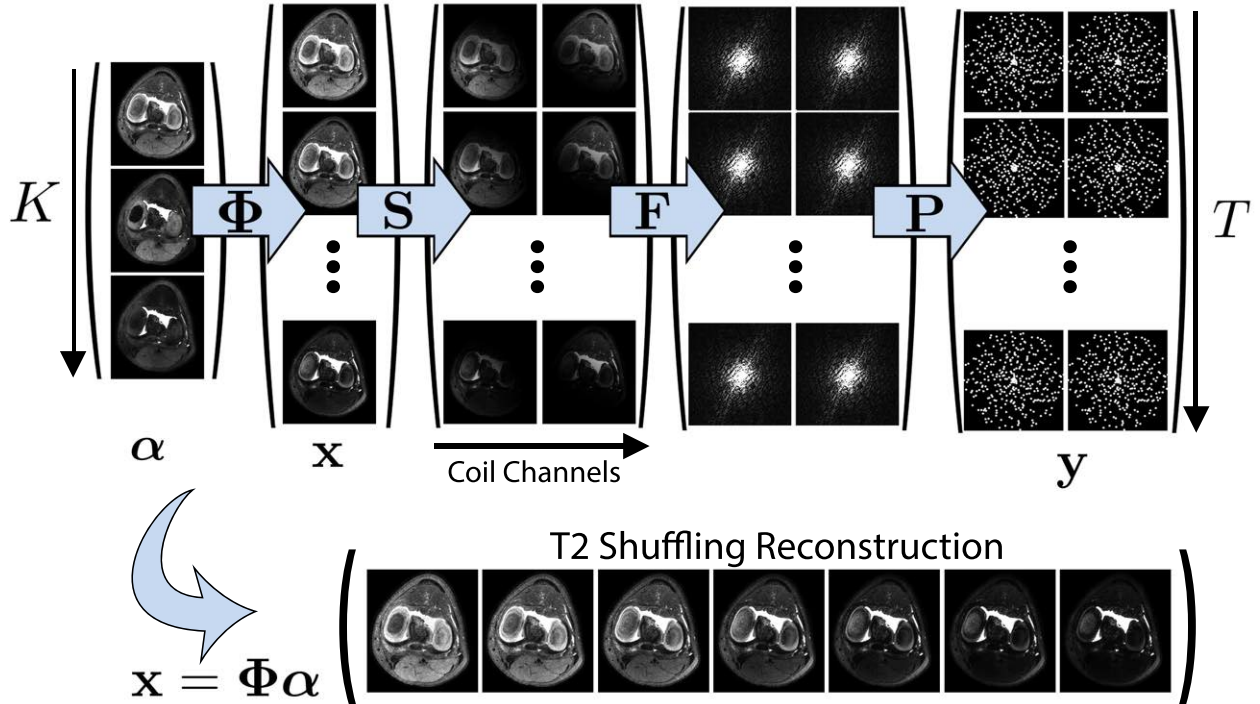


Figure 3.3: T2 Shuffling forward model with explicit subspace constraint. The time series of images are generated by back-projecting through the low-dimensional subspace with K subspace coefficients ($K = 3$ in this example).

A locally low rank (LLR) regularization functional $\lambda \sum_{\mathbf{r}} \|R_{\mathbf{r}}(\alpha)\|_*$ is used to further reduce sample complexity [18, 69, 70, 78, 79]. LLR regularization is implemented in a


manner similar to the approach described in [69]. The operator R_r extracts a block from each temporal coefficient image centered around voxel \mathbf{r} and reshapes each block into a column of a matrix. The nuclear norm is applied to each matrix with regularization parameter λ and the result is summed. LLR regularization exploits spatial correlations in the temporal image coefficients, providing substantial dimensionality reduction beyond the capabilities of joint wavelet regularization or finite differences [55, 77, 80, 81]. In effect, LLR constrains local image patches into a smaller space *within* the subspace Φ_K , which compresses the representation to fewer than K coefficients per voxel. To avoid blocking artifacts [82], the temporal image coefficients are randomly shifted at each iteration prior to LLR regularization [83]. Low rank regularization is described further in Appendix A.

As discussed in Appendix B, a consequence of the subspace constraint (3.14) for a fully sampled, linear reconstruction is that the noise in the system does not depend on the ETL, even if the signal evolutions fully decay. This is because noise outside the subspace is removed through the projection operation, i.e. the final noise standard deviation is $\sigma\sqrt{K}$, where σ is the noise standard deviation of a conventional 3D FSE acquisition that neglects the time progression [84]. Thus a linear reconstruction leads to a fixed \sqrt{K} reduction in SNR, as indicated by (B.11). Although an under-sampled imaging experiment is not analyzed in this work, the one-dimensional result is used as a heuristic motivation for increasing the ETL to collect more data without incurring any additional blurring in the resulting reconstruction. In addition, the use of LLR regularization in the reconstruction formulation (3.15) provides additional denoising.

3.2.2 Echo Train Ordering

In conventional proton-density 3D FSE acquisition, phase encodes along an echo train are ordered according to radial distance so that the k-space center is sampled first and the periphery is sampled last during each echo train. By using a center-out ordering, the k-space center is acquired early during the signal evolution, providing high SNR and producing apparent proton-density contrast. Figure 3.4a depicts the echo train ordering for center-out sampling. Because the signal decays during the echo train, k-space is apodized. The effect is equivalent to space-variant low-pass filtering in image space and leads to a blurring point spread function (PSF). Refocusing flip angle modulation reduces this blurring for tissues with a particular range of T2 values, but it does not fully compensate for this effect. Furthermore, tissues with short T2 remain affected and are blurred out.

a. Center-Out Ordering



b. Randomized Ordering

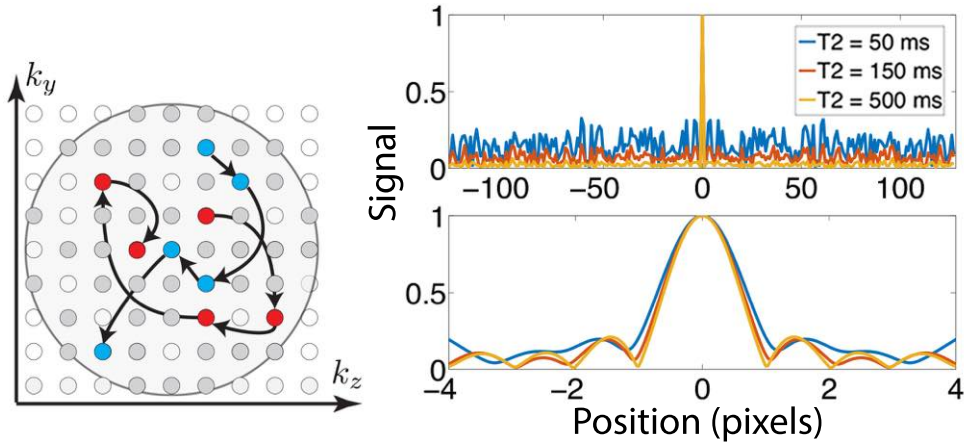


Figure 3.4: Comparison of center-out and randomly shuffled view ordering schemes for an exponential decay. (a) The echo trains in a center-out ordering start at central k -space and move outward in a radial fashion. The resulting normalized PSF for species with high T_2 is sharp. For low T_2 the normalized PSF has non-negligible width, indicating intrinsic image blur. (b) The echo trains for a randomly shuffled ordering traverse k -space in a random fashion, and phase encodes may be repeatedly re-sampled. The PSF has a distinct sharp peak, indicating reduced blur, and incoherent artifacts. The peak remains sharp for low T_2 values at the cost of larger incoherent artifacts.

Since the inverse problem defined in (3.15) uses ideas from compressed sensing to reduce sample complexity through the use of LLR regularization, the data space should be sampled incoherently with respect to the unknown [11]. The forward model (3.7) represents the data in k - t space. In this view, a center-out ordering acquires low-frequency k -space samples at early echo times and high-frequency k -space samples at later echo times,

which is not incoherent with respect to the temporal coefficients. One realizable incoherent sampling approach that fits the signal distribution uses randomized sampling both in k-space and in time [75, 85]. Figure 3.4b depicts an alternative view ordering scheme, in which echo trains are formed by randomly selecting phase encodes throughout k-space. This randomly shuffled ordering induces a PSF with a sharp peak, indicating reduced blur, and spreads artifacts incoherently. However, the peak-to-sidelobe ratio decreases as T2 decreases, highlighting the difficulty of recovering images with short T2.

Transform PSF




Figure 3.5: The system input and TPSF due to the first temporal coefficient (other coefficients show similar behavior). Center-out ordering (middle) blurs and spreads interference coherently, while randomized ordering (bottom) spreads interference incoherently.

The effect of randomly shuffled echo trains can also be observed in the subspace domain. Figure 3.5 shows a system input containing a uniform patch in the first temporal coefficient and the resulting transform PSF (TPSF). A uniform patch is used as it represents a single coefficient (point) in the LLR domain. The center-out ordering (top) leads to blurring and coherent interference. The coherent interference couples the coefficients, which demonstrates the deconvolution is ill-posed under the center-out sampling scheme. The randomized ordering (bottom) spreads interference incoherently through space and time,

and does not couple the coefficients. The LLR soft threshold operation on the randomized TPSF will reduce the incoherent artifacts. The TPSF calculation is described in Appendix C.

3.3 Methods

3.3.1 Subspace Selection

The subspace selection affects both the compression power and the modeling error. In this work the subspace was pre-computed by simulating an ensemble of signal evolutions that match the distribution of T2 values in the anatomy of interest. In a clinical environment, this subspace could be computed offline once and stored for future exams from each imaging protocol (e.g. knee, foot). Figure 3.6a shows a T2 map of a human foot and the corresponding histogram of T2 values. A total of 256 T2 values were sampled from this distribution and EPG-simulated with a variable refocusing flip angle train to produce the ensemble of signal evolutions shown in Figure 3.6b. Because of the weak dependence of T1 on the signal evolution [4, 40], EPG simulations were repeated with T1 values of 500, 700, 1000, and 1800 ms, for a total of $N = 1024$ signal evolutions. The temporal subspace was generated by taking the first $K = 4$ principal components of the ensemble of signal evolutions through PCA. Four principal components are sufficient to represent each signal evolution with greater than 99% accuracy. Figure 3.6c shows the subspace curves for $K = 4$ and an ETL of 78, and their use in forming signal evolutions corresponding to different tissue relaxation through (3.14). Since the subspace flexibly represents arbitrary signal evolutions, it need not precisely represent all relaxation values. In particular, if a voxel contains multiple tissue species, and each respective tissue's signal evolution lives in the subspace Φ_K , then the sum of the signal evolutions will also live in the subspace [53]. For the FSE signal model, the subspace property also allows for mismatches between the simulated and actual refocusing flip angle trains, e.g. due to B_1 inhomogeneity. Figure 3.7 shows a simulation of the impact of B_1 inhomogeneity on the subspace model error. To simulate B_1 homogeneity, EPG simulations were repeated by scaling the RF excitation and refocusing flip angles by up to 40% and the normalized model error was computed according to

$$\frac{\|\tilde{\mathbf{x}} - \Phi_K \Phi_K^* \tilde{\mathbf{x}}\|_2}{\|\tilde{\mathbf{x}}\|_2}, \quad (3.16)$$

where $\tilde{\mathbf{x}}$ is the signal evolution generated with B_1 inhomogeneity and Φ_K was generated from the purely homogeneous B_1 field. As expected, the model error is lowest at 100% scaling, with a maximal error of 0.5%. The worst-case model error increases to 3% at high B_1 inhomogeneity and remains at a tolerable level for most signal evolutions.

The subspace size presents a tradeoff between bias and noise. As K increases, the model error decreases and the noise variance increases. Figure 3.8 shows the tradeoff for

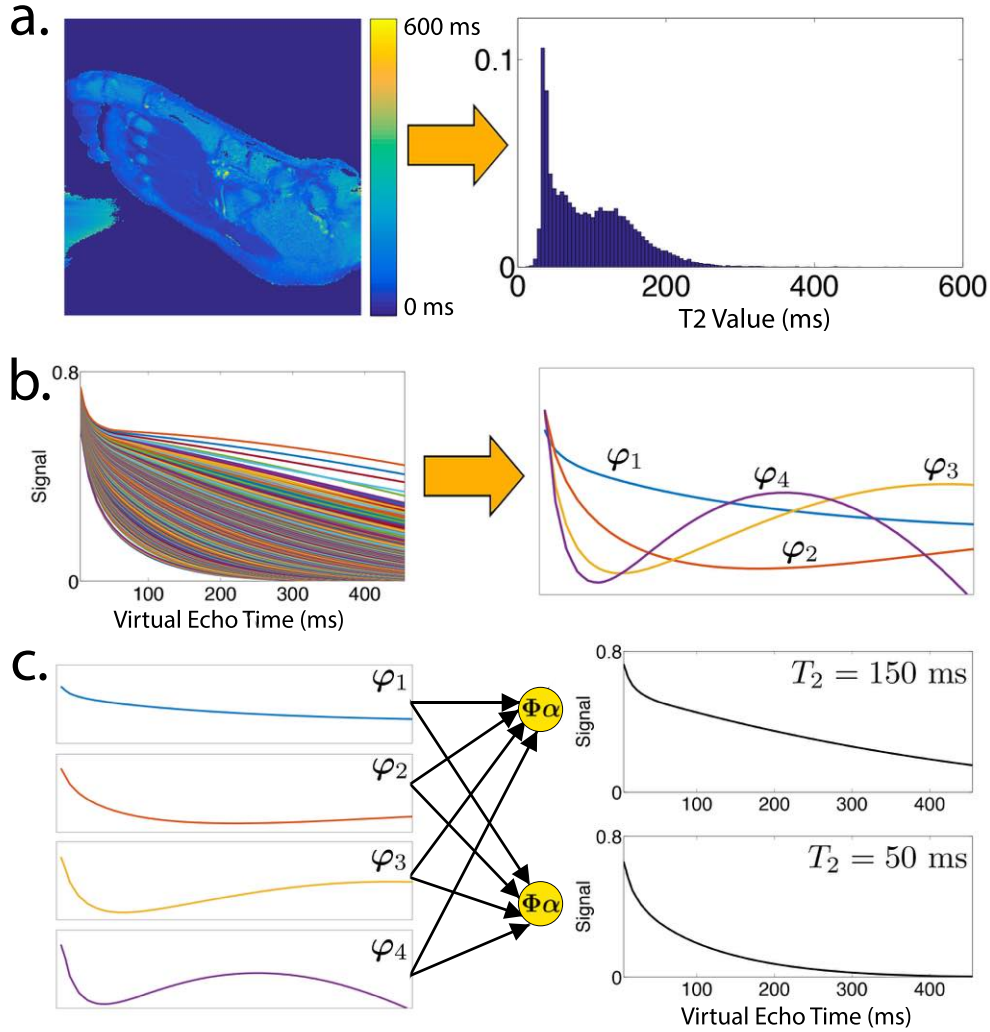


Figure 3.6: (a) T2 map and histogram of the foot derived from a 2D multi-echo scan. (b) Ensemble of EPG-simulated signal evolutions using values from the T2 histogram and the first four principal components that comprise the temporal subspace. (c) Temporal subspace curves are linearly combined to form signal evolutions.

short- T_2 and long- T_2 signal evolutions with $T_1 = 1000$ ms. The top plot shows normalized model error given by (3.16), and is less than 0.5% across all simulated signal evolutions for $K = 4$. The bottom plot shows the overall root mean-squared error (RMSE) derived from (B.10). The RMSE was calculated with $\sigma = 0.7$, which corresponds to an SNR of about 15 dB for $T_2 = 14$ ms.




Figure 3.7: Normalized subspace model error ($K = 4$ coefficients) in the presence of B_1 inhomogeneity for a range of EPG-simulated signal evolutions. B_1 inhomogeneity was simulated by scaling the RF excitation and refocusing flip angles during the echo train by a fixed percentage. The subspace was generated from a 100% homogeneous B_1 field.

3.3.2 Data Acquisition

The CUBE 3D FSE pulse sequence (GE Healthcare, Waukesha, WI, USA) was previously modified to support Poisson disc sampling [39]. In this work, the sequence was extended to re-sample and store phase encodes at multiple echo times, enabling the prescription of arbitrary sampling patterns and echo trains. In a conventional 3D FSE scan, the operator first prescribes imaging parameters such as resolution, TR, and TE to meet a target image contrast, and next chooses a suitable spatial acceleration factor and ETL to meet a target scan time. Because T2 Shuffling resolves the full tissue dynamics at any virtual echo time during the echo train, it is no longer necessary to set a target TE. Instead, the operator chooses a target resolution, TR, and scan time. The TR and scan time provide an upper limit on N_{trains} , the number of echo trains, independent of ETL:

$$N_{\text{trains}} = \left\lfloor \frac{\text{Scan Time}}{\text{TR}} \right\rfloor, \quad (3.17)$$

where $\lfloor \cdot \rfloor$ is the floor function. The ETL is now a free parameter, and can be experimentally chosen for each protocol to account for SNR, T1 weighting, and specific absorption

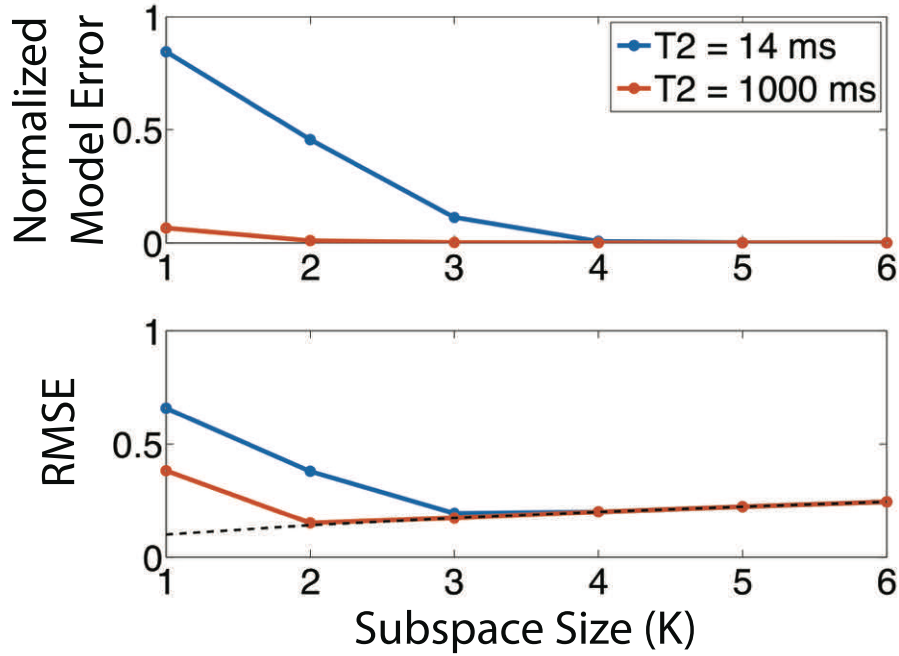


Figure 3.8: Normalized subspace model error (top, (3.16)) and overall RMSE (bottom, $\sqrt{(B.10)}$) vs. subspace size for short- T_2 and long- T_2 signal evolutions with $T_1 = 1000$ ms. The dashed black line shows the RMSE due to noise for $\sigma = 0.07$. With $K = 4$ coefficients, a low model error is uniformly achieved with a reasonable noise penalty.

rate (SAR) limits.

Under this view, and because phase encodes can be re-sampled, conventional parallel imaging acceleration factor is no longer a scan parameter. Instead, the *relative acceleration* is defined as the number of phase encode measurements acquired with respect to a fully sampled circular k-space coverage. Although this definition does not account for the ETL, it serves as a metric for the effective amount of data collected. In addition, since K temporal coefficient images are reconstructed with T2 Shuffling, the *apparent acceleration* is defined as K times the relative acceleration. The apparent acceleration represents the amount of data collected with respect to the number of degrees of freedom in the forward model.

One approach to incoherent sampling is to generate a unique variable-density Poisson disc mask with N_{trains} samples for each echo time. An echo train is then formed by choosing one sample from each mask. However, naively chaining together samples from each mask could result in large gradient switching during the echo train and lead to eddy current artifacts. To overcome this limitation, an alternative sampling approach is described in Appendix D that reduces the distance between points in each echo train. Figure 3.9 depicts the different sampling patterns across the echo times. Due to the variable refocusing

flip angle scheme, the signal evolutions start in a transient decay before reaching pseudo steady state [40, 45]. As a result, traditional center-out 3D FSE discard up to 15 initial echoes in each echo train. Since T2 Shuffling models the signal decay, data are acquired during the transient state, increasing scan efficiency. The first two echoes in the echo train are used to fully cover the center of k-space. This provides a time-consistent calibration region for estimating ESPIRiT coil sensitivity maps. All remaining echoes are used for reconstruction.

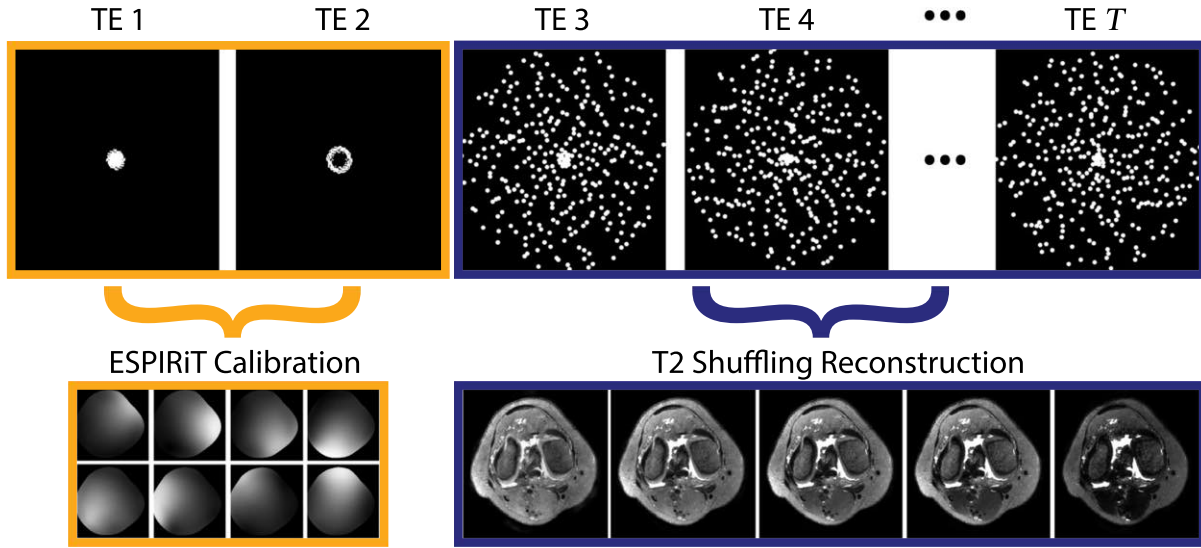


Figure 3.9: Sampling patterns across each echo time for T2 Shuffling. The first echoes are acquired during the transient decay in a center-out fashion and used for coil sensitivity estimation. The remaining echoes follow a variable-density Poisson disc sampling pattern and are reconstructed into a time series of images.

2D Retrospective Under-Sampling and Anatomical Simulation

To test the forward model in (3.7) and the proposed sampling procedure, a fully sampled 2D multi-echo sequence was used to scan a sagittal slice of a volunteer's foot at equally spaced echo times at 3T (TR/TE = 3060/10 ms, 32 echoes, 32 coils, 180 degree refocusing flip angles) under an approved IRB study. To reduce computation time, the data were coil-compressed to 8 virtual channels [86] and used to estimate ESPIRiT coil sensitivity maps. A T2 map was estimated by discarding the first echo time and performing a mono-exponential fit (Figure 3.6a). From the T2 map, synthetic multi-echo data sets were generated through EPG simulation with variable refocusing flip angles and a fixed value of $T_1 = 1000$ ms. The synthetic data were then multiplied by the ESPIRiT maps and Fourier transformed to produce k - t space. Figure 3.10 shows the first TE of the fully sampled, simulated data with an ETL of $T = 80$.

To simulate the effects of a 3D FSE acquisition on a single slice in the readout direction, the coil-compressed, multi-echo data were retrospectively under-sampled using a center-out view ordering. The data were also retrospectively under-sampled using the randomly shuffled view ordering scheme. The two under-sampled data sets were each reconstructed with T2 Shuffling to produce a time series of images. The data were also summed along the time dimension to simulate conventional 3D FSE with a spatial acceleration factor of 2 (40 in k - t space). These were then reconstructed using L1-ESPIRiT, a parallel imaging and compressed sensing method that does not compensate for signal decay [87]. The retrospective experiments were also repeated on the original multi-echo data (no flip angle modulation, 32 time points) to investigate the performance of T2 Shuffling for constant flip angle schemes (approximately exponential decay).

In Vivo Data

Adult volunteers and pediatric patients were scanned under an approved IRB study with informed consent/assent using the modified CUBE pulse sequence for knee imaging with an 8-channel knee coil array at 3T. Table 3.1 lists the common scan and imaging parameters used for the experiments, and Table 3.2 shows the scan parameters for each individual scan type. The first scan (Scan 1a) matched the current pediatric knee protocol used at the Lucille Packard Children’s Hospital (Stanford, California) with a center-out view ordering and variable-density Poisson disc sampling. The second scan (Scan 2a) used a randomly shuffled view ordering and sampling as discussed in Appendix D. The first scan was reconstructed into a single proton-density image using L1-ESPIRiT. The second scan was reconstructed into 78 virtual echo time images using T2 Shuffling. Both scans were 6 minutes and 30 seconds in duration. A second set of scans with scan times of 7 minutes and 30 seconds was also used to reduce the ETL for the center-out ordering (Scan 1b) and to reduce the acceleration for the randomly shuffled ordering (Scan 2b). T2 Shuffling was compared to L1-ESPIRiT because compressed sensing has shown great promise toward improving 3D FSE [39, 41, 42].

Table 3.1: Common scan and imaging parameters used for in vivo experiments.

	All Scans
TR (ms)	1400
Echo Spacing (ms)	5.5
Receive Bandwidth (kHz)	± 62.5
Fat Saturation Efficiency	0.85
Acquisition Matrix	$288 \times 260 \times 240$
Reconstruction Matrix	$512 \times 512 \times 472$
Acquisition Voxel (mm^3)	$0.6 \times 0.6 \times 0.7$
Reconstruction Voxel (mm^3)	$0.3 \times 0.3 \times 0.4$
Reformat Slice Thickness (mm)	Axial: 0.3, Sagittal: 1.5, Coronal: 2.0

Table 3.2: Scan and imaging parameters used for each scan type.

	Scan 1a	Scan 2a	Scan 1b	Scan 2b
View Ordering	Center-out	Randomly Shuffled	Center-out	Randomly Shuffled
ETL	33	80	28	82
N_{trains}	278	278	321	321
Initial Discarded Echoes	5	2	5	2
TE_{eff} (ms)	23	-	23	-
Relative Acceleration ^a	6.3	2.3	6.6	1.9
Apparent Acceleration ^b	-	9.2	-	7.6
Image Contrast	Proton	Proton, T2 (78 images)	Proton	Proton, T2 (80 images)
Scan Time	6 min. 30 sec.	6 min. 30 sec.	7 min. 30 sec.	7 min. 30 sec.

^aThe acceleration with respect to a fully sampled circular k-space coverage.

^bThe relative acceleration multiplied by the number of temporal coefficient images.

Reconstruction Implementation

The T2 Shuffling reconstruction was implemented in C and CUDA using the Berkeley Advanced Reconstruction Toolbox (BART) [19, 88]. MATLAB demonstration code is available for download at <http://eecs.berkeley.edu/~mlustig/Software.html>.

The reconstruction problem (3.15) is convex and can be solved using iterative algorithms such as FISTA [21] and ADMM [22]. In both cases, the algorithms require repeated application of the forward and adjoint operators described by the normal equations,

$$\Phi_K^* \mathbf{S}^* \mathbf{F}^* \mathbf{P} \mathbf{F} \mathbf{S} \Phi_K \alpha = \Phi_K^* \mathbf{S}^* \mathbf{F}^* \mathbf{P} \mathbf{y}. \quad (3.18)$$

In this naive formulation, the run-time and memory requirements scale with the ETL. For an ETL of $T = 80$, $K = 4$ temporal coefficients, and 8 coils, the forward operator amounts to an 80×4 matrix multiplication per voxel, followed by 80 sensitivity map multiplications and 640 FFTs. The result is multiplied by the sampling mask and similar operations are applied in reverse order for the adjoint. Although some operations can be parallelized, they pose a large memory and computation bottleneck on the reconstruction.

To reduce computation, a similar approach to Step 1 in [89] is used. Since the temporal basis Φ_K operates across time and the sensitivity map multiplication \mathbf{S} and Fourier transform \mathbf{F} operate across space, the operators commute:

$$\mathbf{F} \mathbf{S} \Phi_K = \Phi_K \mathbf{F} \mathbf{S}. \quad (3.19)$$

Thus, the normal equations can be rewritten more compactly:

$$\mathbf{S}^* \mathbf{F}^* \Phi_K^* \mathbf{P} \Phi_K \mathbf{F} \mathbf{S} \alpha = \mathbf{S}^* \mathbf{F}^* \Phi_K^* \mathbf{P} \mathbf{y} \quad (3.20)$$

$$\implies \mathbf{S}^* \mathbf{F}^* \Psi_K \mathbf{F} \mathbf{S} \alpha = \mathbf{S}^* \mathbf{F}^* \Phi_K^* \mathbf{P} \mathbf{y}. \quad (3.21)$$

Here, $\Psi_K := \Phi_K^* \mathbf{P} \Phi_K$ is a $K \times K$ sampling kernel per voxel and can be pre-computed prior to reconstruction. For $K = 4$, the forward operation now only requires the application of 4 sensitivity map multiplications and 32 FFTs, followed by a 4×4 matrix

multiplication per voxel. The adjoint operation sees a similar reduction in computation. The data y and time series of images x are no longer explicitly stored, making the reconstruction's memory and run-time requirements independent of the ETL. The complexity is roughly $K = 4$ times greater than that of an L1-ESPIRiT reconstruction.

As the 4D reconstruction involves a large amount of data and computation, processing times remain a bottleneck for clinical adoption. To overcome this bottleneck, a fast, distributed reconstruction was designed and implemented, and is described in Chapter 4.

3.4 Results

3.4.1 Retrospective Under-sampling and Simulation Results

The first experiment, carried out on the anatomical simulation with variable refocusing flip angles, is shown in Figure 3.10. The left image shows the first TE of the simulated, fully sampled data and serves as the gold-standard proton-density contrast. The middle-left shows the L1-ESPIRiT reconstruction with center-out ordering and represents a conventional compressed sensing approach. The blurring due to T2 decay is evident, reaffirming the limitation of current FSE. After randomly shuffling the echo train ordering (middle image), the L1-ESPIRiT reconstruction shows incoherent artifacts, which are explained by the PSF in Figure 3.4b. There is also mixed contrast due to the k-space modulation induced by the random sampling. Despite the image artifacts, detailed structure remains, indicating the potential to “denoise” with T2 Shuffling. The first virtual echo time from the T2 Shuffling reconstruction applied to center-out ordering is shown in the middle-right image. The reconstruction does not successfully recover the time series of images because the sampling is coherent in k - t space (ill-posed deconvolution). The T2 Shuffling reconstruction in tandem with randomized view ordering (right image) is sharp, and the first virtual echo time image shows similar quality and contrast to that of the fully sampled image. There is a clear delineation of fine structure that is not afforded by the conventional 3D FSE pipeline.




Figure 3.10: Proton-density reconstructions on retrospectively under-sampled data from variable flip angle modulation simulation. From left to right: first TE of the fully sampled slice; L1-ESPIRiT reconstruction using a center-out echo train ordering; L1-ESPIRiT reconstruction using a randomly shuffled echo train ordering; first virtual TE of T₂ Shuffling reconstruction with center-out echo train ordering; first virtual TE of T₂ Shuffling reconstruction with randomly shuffled echo train ordering.

Figure 3.11 shows a similar trend on the simulations that were based on the acquisition data. Since the multi-echo sequence used a constant flip angle schedule, the blurring from center-out ordering is even more pronounced. The T₂ Shuffling reconstruction with randomized view ordering is comparable to the fully sampled slice despite the approximate exponential signal decay.




Figure 3.11: Proton-density reconstructions on simulation derived from retrospectively under-sampled acquisition data with constant flip angles. From left to right: first TE of the fully sampled slice; L1-ESPIRiT reconstruction using a center-out echo train ordering; L1-ESPIRiT reconstruction using a randomly shuffled echo train ordering; first virtual TE of T2 Shuffling reconstruction with center-out echo train ordering; first virtual TE of T2 Shuffling reconstruction with randomly shuffled echo train ordering.

The T2 Shuffling reconstruction recovers the full time series of virtual echo time images and provides various degrees of T2 contrast, all from the same data. Figure 3.12a shows a later virtual echo time image (approximately 100 ms) for both the fully sampled simulation and acquisition slices, and well as the T2 Shuffling reconstructions (with the randomly shuffled ordering). As suggested by [4], the effective contrast-equivalent TE (TE_{eff}) is shown for the variable flip angle modulation. The reconstructions closely match the fully sampled slices. Figure 3.12b shows representative temporal signal evolutions for the cases considered in Figure 3.12a. The signal curves for the variable flip angles show close agreement, indicating low model error and successful reconstruction. Note the deviation from the true signal evolution in the case of the constant flip angle; this is a result of the non-exponential behavior that was not modeled in the basis and indicates the importance of matching the temporal subspace to the pulse sequence design.

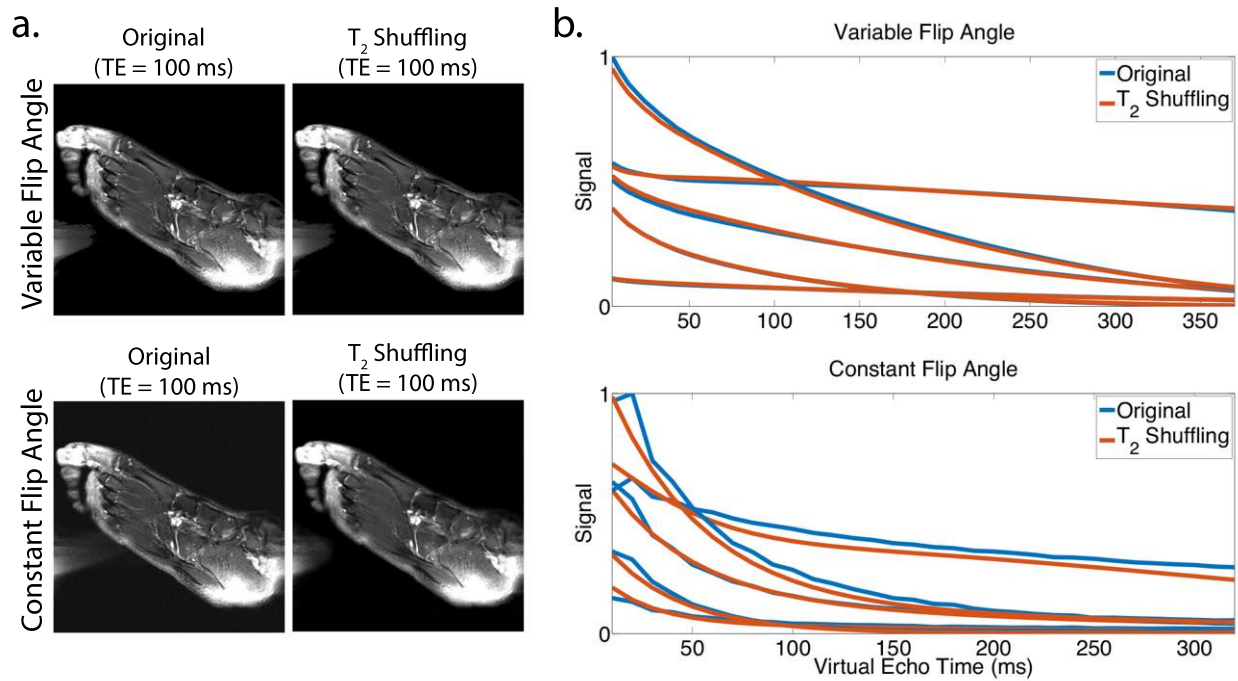


Figure 3.12: (a) Comparison of the 100 ms TE_{eff} image for the fully sampled slices and the T2 Shuffling reconstructions (with randomly shuffled ordering) for both the simulation and the acquisition data. (b) Representative signal evolution curves from the fully sampled slices and the reconstructions in (a).

3.4.2 In Vivo Results

Figure 3.13a depicts axial reformatted images of modified Poisson disc CUBE with L1-ESPIRiT vs. T2 Shuffling reconstructions on a pediatric patient with bone marrow edema (Table 3.2, Scans 1a and 2a, respectively). The proton-density T2 Shuffling image shows increased delineation of the patellar tendon and other fine structures as compared to the L1-ESPIRiT reconstruction. The later virtual echo time images clearly show the edema due to the higher T2 value of the fluid relative to the surrounding tissue. Figure 3.13b shows sagittal and coronal reformatted T2 Shuffling images from a knee scan of a pediatric patient with a discoid meniscus (Table 3.2, Scan 2b). The early virtual echo time image points to intrasubstance degeneration, but the later virtual echo time clearly shows the tear and fluid.




Figure 3.13: (a) Axial reformatted reconstructions of L1-ESPIRiT using Poisson disc CUBE and T2 Shuffling at three virtual echo times. The solid yellow arrows depict increasingly improved contrast of the bone marrow edema. The dashed white arrows (zoomed-in region) show enhanced delineation of the patellar tendon in the T2 Shuffling images. Acquisition parameters are listed in Table 3.2, Scan 1a and Scan 2a. (b) T2 Shuffling reconstruction of a pediatric patient scan with a discoid meniscus. The coronal reformat at 50 ms TE_{eff} shows the discoid meniscus (dashed white arrow). The sagittal reformat at 20 ms TE_{eff} shows potential intrasubstance degeneration. The same slice at 90 ms TE_{eff} clearly depicts the meniscal tear and fluid (solid yellow arrow). Acquisition parameters are listed in Table 3.2, Scan 2b.

In some cases, T2 Shuffling only afforded marginal qualitative improvement. Figure 3.14 presents such a case for a pediatric patient, comparing sagittal reformats of Poisson disc CUBE with L1-ESPIRiT and T2 Shuffling at two virtual echo times (Table 3.2,

Scan 1a and 2a, respectively). Although there are some differences between the proton density images, it is not immediately clear if they are due to deblurring or due to residual aliasing artifacts. The solid yellow arrow shows increased apparent resolution in the muscle, while the dashed white arrow shows ringing artifacts in the bone only present in the T2 Shuffling images.




Figure 3.14: Sagittal reformatted reconstructions of (a) L1-ESPIRiT using Poisson disc CUBE and (b-c) T2 Shuffling at two virtual echo times. The solid yellow arrow shows increased apparent resolution in the muscle, while the dashed white arrow shows ringing artifacts in the bone signal of the T2 Shuffling images. Acquisition parameters are listed in Table 3.2, Scan 1a and Scan 2a.

Figure 3.15a shows an axial reformatted slice of the first virtual echo time image using T2 Shuffling for a volunteer. The image depicts the trabecular bone structure in the distal femur. This area is challenging in particular for parameter mapping methods because of the multi-compartmental nature of the bone structure. Since T2 Shuffling does not explicitly estimate the T2 value and instead flexibly represents the voxels by a linear subspace, the modeling error is low and the structure is preserved. The dimensionality reduction offered by LLR is depicted in Figure 3.15b; each voxel neighborhood is constrained to a smaller subspace *within* the global subspace. The local rank is determined by taking non-overlapping patches of the reconstructed temporal coefficient images, reshaping each patch into a matrix, and computing the number of non-zero singular values.

Regions in the image containing a single tissue species are captured by a single coefficient. The regions at tissue boundaries require more coefficients to represent due to the diversity in T2. The correlation between tissues is naturally exploited by the subspace constraint; Figure 3.15c shows the result of applying k-means to the coefficient images. With little effort, the image is segmented into broad tissue classes based on the temporal coefficients of each voxel.




Figure 3.15: (a) Axial reformatted slice of the T2 Shuffling reconstruction at the first virtual echo time. (b) Number of degrees of freedom in each voxel after LLR reconstruction. (c) K-means clustering applied to the coefficient images after reconstruction. Each color represents a broad tissue class.

3.5 Discussion

Conventional FSE imaging methods must balance the tradeoffs between scan time, image sharpness, and image contrast. T2 Shuffling is able to break this dependency and resolve a time series of sharp images from a single acquisition. Instead, T2 Shuffling inherits similar benefits and challenges present with compressed sensing: the limited acquisition time manifests as incoherent artifacts in the images. Short T2 species impart larger artifacts, and these artifacts spread through the entire image volume. Thus, successful recovery is tied to the sampling rate and distribution in k - t space. From this viewpoint, a center-out echo train ordering does not adequately sample the temporal signal evolution curve and leads to an ill-conditioned inverse problem. However, it is possible to sample k - t space incoherently by randomly shuffling the echo train ordering. The modification is straightforward and can be added to existing 3D FSE pulse sequences with little change. The signal modulation from randomized echo train orderings is no longer smooth, as phase encodes are randomly shuffled. Nonetheless, eddy current effects can be mitigated by assigning nearby phase encodes in the same echo train, as discussed in Appendix D.

The ability to reconstruct the time series of images from the limited data is due to the high degree of sparsity. T2 Shuffling builds upon similar work [13, 52–55, 80, 90–92] to

produce a faithful low-dimensional representation of the temporal signal evolutions. This temporal basis can be computed offline once and stored for future exams. The temporal subspace constraint alone does not provide a sufficiently compact representation, as $K \approx 4$ observations are required at each voxel. However, the use of LLR regularization further reduces the sample complexity by about a factor of 2. In addition, the added denoising offered by LLR enables the use of a larger subspace (to reduce model error) without a large increase in sample complexity. By constraining patches to be locally low rank, the spatio-temporal images are represented as a union of subspaces. Thus, it may be possible to improve upon this representation through subspace clustering methods [93]. Unlike LLR, subspace clustering can group voxels that are not spatially localized.

A second consequence of the subspace constraint is that the intrinsic dimension of the image set does not grow with increasing ETLs. For feasible TRs, there remains observable signal throughout the TR due to tissue with high T_2 . This residual signal informs the entire reconstruction to help reduce incoherent artifacts throughout the image volume. Nonetheless, the ETL cannot occupy the entire TR in order to retain sufficient time for longitudinal recovery. For a particular imaging metric, it may be possible to choose an optimal flip angle schedule and ETL given SAR constraints. In [94] and [95], the optimal TEs are chosen to maximize SNR and parameter estimation admissibility given a multi-exponential decay model. A similar approach could be used for the flip angle and ETL selection.

To use T2 Shuffling in a clinical workflow, it is necessary to maintain a fast reconstruction. The use of BART as a CPU- and GPU-accelerated iterative reconstruction platform has afforded a distributed reconstruction with 90 second end-to-end processing times. Others have noted it is not feasible to save and display all images along the signal relaxation curve [73, 96]. Further clinical validation is needed to motivate the selection and visualization of the data. At Lucile Packard Children’s Hospital at Stanford, CA, three virtual echo time images are ultimately transmitted to the PACS, corresponding to proton-density, intermediate T2-weighted, and T2-weighted images. However, since the signal evolutions are fully derived from the temporal coefficients through (3.14), it is feasible to save the temporal coefficients only and compute arbitrary virtual echo time images in real-time. This could be used with advanced image viewers to intuitively navigate the 4D data set in space and in time in a manner similar to the capabilities of the SyMRI software (SyntheticMR AB, Linköping, Sweden).

Despite the demonstrated benefits of T2 Shuffling, the linear relaxation through PCA comes at the cost of lower apparent SNR. In addition, if the temporal subspace does not suitably model the signal formation, e.g. due to motion, the model mismatch will lead to poor data fidelity and degrade the reconstruction. The bias vs. noise tradeoff presented in Appendix B only considered fully sampled signal evolutions. Even though the analysis does not directly extend to under-sampled acquisitions, it was used as a rule-of-thumb for the T2 Shuffling acquisition and reconstruction. Because the acquisition is highly accelerated, the reconstruction relies on a strong sparsity prior. In the absence of sparsity, longer scan times are necessary to fully suppress the incoherent aliasing artifacts. These

drawbacks must be fully understood and addressed before wider clinical adoption.

Although this work focused on volumetric imaging, similar tradeoffs exist for 2D FSE. The signal decay that leads to blurring is partially mitigated by reducing ETLs and interleaving slices in multi-shot 2D FSE [3]. Variable refocusing flip angles have also been used to reduce blurring in single-shot 2D FSE [97]. The ideas behind T2 Shuffling could also be used in 2D FSE to further reduce blurring due to long ETLs as well as resolve multiple image contrasts.

3.6 Conclusion

In conclusion, this chapter presented T2 Shuffling, a novel acquisition and reconstruction method for volumetric FSE imaging. Conventional 3D FSE trades off scan time and image quality, e.g. shorter scans incur more blurring. Additionally, current clinical methods separately acquire proton-density and T2-weighted image contrasts. T2 Shuffling addresses both issues by simultaneously reconstructing multi-contrast images with increased sharpness from a single 3D FSE acquisition. Each image in the time series exhibits T2 contrast corresponding to a particular virtual echo time. Based on the experiments, the use of multiple image contrasts at isotropic resolution has a demonstrated diagnostic benefit, as the signal evolution of tissue microstructure provides additional information on pathology.

Chapter 4

Clinical Deployment

4.1 Introduction

Many exciting advances to MRI acquisition and reconstruction methods are enabling markedly faster scanning. However, to realize the benefits clinically, the methods must integrate into the existing hospital system. This includes ease of use for technologists and low reconstruction latency. A number of clinical applications have made great leaps toward this goal, showing promise in the wild [17, 18, 24, 25, 98–100]. T2 Shuffling and other parallel imaging and compressed sensing (PICS) approaches require iterative solvers for reconstruction [11]. Therefore, the reduction in scan time is offset by the increased computational complexity in the reconstruction. As the acquisition and reconstruction methods grow in complexity, the computational requirements continue to grow, leading to reconstructions that can take tens of minutes to several hours [98, 101, 102]. In order for these approaches to gain clinical traction, the image reconstruction pipeline should maintain a low latency so that the images are available at the scanner before the patient leaves the table. Based on initial experience, clinical integration requires latencies of at most a few minutes [103].

Motivated by the need for fast PICS reconstructions, a number of general-purpose software packages have been built as a mechanism to translate the work to the clinic [104–110]. As a result, several applications of PICS have seen success in a clinical research setting [24–26]. Initial efforts focused on accelerating 3D reconstructions using CPUs and GPUs, achieving reconstruction times on the order of one minute [103, 111–113]. Extensions to cloud-based processing [114] have enabled reconstructions that take advantage of scalable compute. More recently, the focus has shifted to 2D and 3D dynamic imaging [99, 100, 115, 116]. The latter often involves significantly higher memory/computation requirements.

This chapter describes the steps taken to integrate the T2 Shuffling sequence into a clinical protocol. The reconstruction pipeline and its algorithmic implementation are described. Previously, the end-to-end reconstruction time for T2 Shuffling was 20 minutes

to one hour on a single non-optimized shared-memory system [117, 118]. To reduce computation, a distributed reconstruction that uses multiple high-performance computers is proposed and subsequently deployed into a clinical setting. The main optimizations to the computational components are outlined and the scalability of the distributed computing is investigated. After applying machine-level optimization, the reconstruction time is reduced to eight minutes on a single shared-memory system, and to under 90 seconds on a four-node distributed system. The fast processing frees up time for additional signal processing steps that improve image quality robustness [14, 119, 120] while maintaining a suitable latency.

4.2 T2 Shuffling Reconstruction

T2 Shuffling reconstructs subspace coefficient images that can be back-projected to a time series of virtual echo time images. However, in a clinical environment, it is not feasible or necessary to view all the time points. For pediatric knee imaging, it is sufficient to visualize the result at three virtual echo times, corresponding to proton-density weighting, intermediate weighting, and T2 weighting [118]. The reconstructed DICOMS for these three time points are transmitted to the scanner, where a technologist is able to apply multi-planar reformatting and slice averaging. The reformats are then transferred to the Picture Archiving and Communication System (PACS), where radiologists view the result. As the reconstruction requires highly computational, specialized algorithms, the data processing is done on dedicated machines separate from the scanner suite. This requires a mechanism for streaming the scan data from the scanner to the machines, in a process called offline reconstruction. In order to provide quick feedback to the scanner operator, a fast linear reconstruction is initially performed and sent to the scanner.

The T2 Shuffling reconstruction involves solving the inverse problem given by (3.15). As the Cartesian acquisition is fully sampled in the readout direction, it is possible to decouple the slices in the readout direction through an inverse Fourier transform [103]. Then, (3.15) can be independently solved for each slice and rejoined into a volume after. After applying the operator commutativity (3.19), the inverse problem is

$$\arg \min_{\alpha} \frac{1}{2} \|y - \mathbf{P}\Phi_K \mathbf{F}\mathbf{S}\alpha\|_2^2 + \lambda \sum_{\mathbf{r}} \|\mathbf{R}_{\mathbf{r}}\alpha\|_* . \quad (4.1)$$

The inverse problem is a special case of the more general convex inverse problem given by (2.9), where $\mathbf{E} := \mathbf{P}\Phi_K \mathbf{F}\mathbf{S}$ and $g(\mathbf{z}) := \lambda \sum_{\mathbf{r}} \|\mathbf{R}_{\mathbf{r}}\mathbf{z}\|_*$. Thus, a number of iterative algorithms can be used. Two suitable first-order iterative methods are FISTA [21] and ADMM [22]. Both approaches have benefits and tradeoffs but perform roughly the same in terms of convergence speed and computational complexity. Taking advantage of the compressed normal operator (3.21), the FISTA update steps for (4.1) are shown in Algorithm 4.1.

Algorithm 4.1 FISTA update steps for T2 Shuffling inverse problem (4.1).

Inputs: y –data, μ –step size, N –maximum iterations

Outputs: $\alpha^{(N)}$ –subspace coefficient images after N iterations

- 1: $\mathbf{b} \leftarrow \mathbf{S}^* \mathbf{F}^* \Phi_K^* \mathbf{P} \mathbf{y}, \mathbf{z} \leftarrow \mathbf{b}, \alpha^{(n)} \leftarrow \mathbf{0}, t_1 = 1$
 - 2: **for** $n = 1 : N$ **do**
 - 3: $\alpha^{(n)} \leftarrow \text{LSVT}(\mathbf{z} - \mu(\mathbf{S}^* \mathbf{F}^* \Psi_K \mathbf{F} \mathbf{S} \mathbf{z} - \mathbf{b}), \mu\lambda)$
 - 4: $t_{n+1} \leftarrow \frac{1 + \sqrt{1 + 4t_n^2}}{2}$
 - 5: $\mathbf{z} \leftarrow \alpha^{(n)} + \frac{t_n - 1}{t_{n+1}}(\alpha^{(n)} - \alpha^{(n-1)})$
 - 6: **end for**
-

After solving (4.1), a 3D spatial wavelet soft-threshold is applied to each basis coefficient image for additional spatial denoising. The time-series of images \mathbf{x} are then recovered through (3.14). T2 Shuffling can also incorporate partial Fourier sampling [47, 121], in which approximate conjugate symmetry is used to reduce the number of acquired measurements. In this case, the reconstruction is processed by a Homodyne filter [122] that accounts for low-frequency phase \mathbf{x} .

Although the primary computational component of the T2 Shuffling reconstruction is the PICS-based inverse problem, a number of other steps comprise the full reconstruction pipeline, shown in Table 4.2. This includes processing the vendor-formatted raw data, computing the basis based on the scan parameters, estimating coil sensitivity maps, applying gradient non-linearity corrections, converting the result to DICOM images, and transferring the images to the PACS. The full pipeline must maintain a low latency so that the images are available at the scanner console before the patient leaves the table. The pre-processing and post-processing overhead is a non-negligible portion of the overall latency; these aspects are often not accounted for when comparing algorithm run-times.

4.3 Optimizations

In this section, strategies for speeding up the baseline implementation of T2 shuffling are described, as well as improvements to image quality which are enabled by the speedups. Several serial bottlenecks are identified and parallelized in the BART toolbox [88] (e.g. FFT phase modulation, normalization, and basis generation). Additionally, inter-operator fusion is exploited to improve cache behavior and eliminate costly transpose operations. Some computations are also completely removed from the pipeline using pre-computation and/or code refactoring. Finally, the computation is distributed across multiple machines by exploiting parallelism across image slices in the PICS portion, and exploiting parallelism across time points during post-processing. This improved perfor-

mance enables time for additional signal processing steps to the pipeline, which improve image quality robustness in a patient setting.

4.3.1 Improving Multi-Core Parallelism

Multi-dimensional operations There are several opportunities to take advantage of efficient parallelism in BART. From the baseline implementation, the `parallel_md_nary` operation, which applies an arbitrary function to many elements in a multi-dimensional array, was modified to be a flat parallel loop instead of a nested for loop. By adding a configuration capability for the loop chunk size, it was also possible to separately tune the parallelism for different operations such as coil compression and data re-sorting. An additional parallelism bottleneck involves FFT phase modulation which enables the use of un-centered FFT packages (e.g. FFTW) As each data element is multiplied by a phase term that is relative to its position in the multi-dimensional array, the operation requires several modulo divisions to compute the correct index. However, when the dimensions of the array are even, for consecutive elements the phase alternates between $+1$ and -1 , or $+j$ and $-j$. Thus, the phase modulation can be parallelized over blocks of consecutive elements. For each block, the position and phase is computed for the initial element only, and the phase is alternated for the remaining elements.

Basis generation As the basis is computed by simulating an ensemble of signal time courses, it is embarrassingly parallel across the different signal simulations. Thus, each signal evolution can be simulated on a single CPU core, and the full simulation can be parallelized across all the cores of the machine. In each simulation, the signal value is propagated through the time points by repeatedly applying matrix operations that represent the acquisition dynamics [30]. To further reduce the simulation time, matrix terms that do not change during the simulation are pre-computed.

Echo time processing The echo time processing is reduced by operating directly on the coil-compressed data, reducing coil-by-coil processing (e.g. Homodyne filter, FFT). Additionally, all three time points are processed in parallel, allocating one third of the machine's resources to each TE.

4.3.2 Fusing Forward and Adjoint Linear Operators

Although the basis in T2 Shuffling in general could be complex-valued, an FSE experiment with CPMG conditions will lead to a constant phase across the signal evolution [34]. In other words, the basis need not represent phase. As a result, the space-time kernel Ψ is real-valued and symmetric. Thus, the storage and computation were modified to include only the real part and load half of the matrix, reducing memory access by a total factor of four.

Each call to the normal operator involves the application of the forward 2D Fourier transform followed by inverse 2D Fourier transform, i.e. $\mathbf{F}^*\Psi\mathbf{F}$. This operation is shown visually in Figure 4.1. The 2D Fourier transforms are applied typically by doing a set of 1D FFTs on the rows, followed by a matrix transposition, then a set of 1D FFTs again on the rows, followed by a matrix transposition to return the data to its original format. The lines in the figure correspond to the matrix ordering after each operation, i.e row or column. Since the application of Ψ is an independent computation across spatial positions, it is agnostic to the ordering. Therefore, the argument can be left in transposed form in between the forward and inverse Fourier transform computations, thus avoiding two of four transpose operations. This optimization is shown as the shortcut path in the figure. Additional cache reuse in traffic reduction to and from memory is accomplished by fus-




Figure 4.1: Graphical depiction of the fused forward and adjoint T2 Shuffling operators. Two transpose operations are avoided by leveraging the fact that the Ψ operator is agnostic to the image data ordering.

ing and re-ordering loops in the forward and adjoint operators. The sensitivity maps and FFT computations are fused and applied to each image row to improve cache reuse. The application of the sensitivity maps is moved to the outer loop; that is, all subspace coefficient images are multiplied by a single sensitivity map before moving to the next coil, reducing intermediate storage by a factor of the number of coils.

Finally, the new optimized normal operator is parallelized across image panels, which are groups of roughly 10 – 20 consecutive image rows assigned to a single CPU thread. This creates two levels of coarse-grained parallelism: (i) the volume split across slices, and (ii) across panels within a slice. The second level of parallelism wasn't initially necessary due to the amount of parallelism available across slices, however it became useful for two reasons. First, by increasing the number of cores applied to each problem, intermediate

results more easily fit in on-chip memory. The benefit of this effect is shown in more detail in Section 4.5. Secondly, this enables scaling to more cores than the number of slices.

4.3.3 Removing Unnecessary and/or Redundant Computations

Fast linear reconstruction The original pipeline included a fast linear reconstruction to provide quick feedback to the technologist. As the overall latency of the full reconstruction became nearer to that of the linear reconstruction due to the optimizations, the linear reconstruction was removed altogether.

Basis and space-time kernel caching The basis operator Φ is determined by the target scan anatomy (e.g. knee), the refocusing RF pulse flip angle schedule, and the associated echo spacing between the pulses during the data acquisition. As it does not depend on the particular patient scan, it can be stored and re-used for future scans that have the same acquisition parameters. To accomplish this, a SHA-1 hash of the list of flip angles, echo spacing, and target anatomy is calculated. Before computing the basis, the hash is searched for in a look-up table to determine if the basis has been previously computed. If it is not found, the basis matrix is computed and stored in the look-up table. In a similar fashion, the space-time kernel is pre-computed and stored a separate look-up table. The SHA-1 hash of Ψ consists of the SHA-1 of the basis and of the sampling pattern. Since the basis and the space-time kernel are 7.1 KB and 7.7 MB, respectively, the load times were negligible.

Single-pass locally lower rank regularizer The LLR operates on non-overlapping patches of size $b \times b$ from the coefficient images. In each iteration of the solver, the coefficient images are first circular-shifted randomly in both dimensions by a small amount less than the patch size [83]. Then the patches are extracted from the K coefficient images and rearranged into a 2D matrix per patch with one column per coefficient image and b^2 rows. Common factors are extracted across coefficient images using the singular value decomposition (SVD), and any singular values below a certain threshold are set to zero before regenerating the coefficient images back from the factors.

The LLR computation is embarrassingly parallel across patches, and this can easily be exploited using OpenMP. However, it is also possible to reduce the number of passes through the data, and hence reduce the amount of data read and written to memory. In the baseline BART implementation, (i) a temporary image is allocated which is padded to allow for random circular shifting, (ii) the shifted coefficient images are copied into the padded buffer, and (iii) the buffer is reshaped into a batch of $b^2 \times K$ contiguous matrices to be processed by a parallel batch SVD. When the batch SVD is complete, these steps are repeated in reverse order to bring the data back to its original location in the coefficient images. Operations (ii), (iii), and their reverse versions each require a full pass through the coefficient images which increases memory traffic at the expense of performance. This

procedure was modified to work in a single pass through the coefficient images in order to reduce memory traffic and unnecessary copies. Each thread allocates a temporary buffer of size $b^2 \times K$ ahead of time. In a single pass over the image, data are extracted according to the parameters of the circular shift, without explicitly performing the shift. The SVD threshold operation is performed in the thread-local buffer, and finally the data are scattered back to the original coefficient image when all the work for that patch is complete.

4.3.4 Image Quality Improvements

The reduction in processing time enabled additional signal processing steps to be added with the aim of improving robustness across a diverse set of pediatric patients. Specifically, Auto-ESPIRiT [120], a parameter-free ESPIRiT implementation in which the internal parameters are chosen based on Stein’s Unbiased Risk Estimate (SURE), was incorporated. Auto-ESPIRiT produces more consistent sensitivity maps in cases with a high degree of data inconsistency, e.g. due to motion or noise. To use Auto-ESPIRiT, an accurate noise estimate was needed. To this end, channel noise pre-whitening [14, 119] was incorporated as a step in pre-processing. As the SURE-based parameter estimate can be computed for each slice independently, it was performed immediately before the PICS reconstruction.

4.3.5 Distributed Implementation

Figure 4.2 depicts the distributed reconstruction pipeline implementation for a T2 Shuffling scan with 288 slices. The PICS and post-processing portions are distributed over multiple machines in the cluster using the Message Passing Interface (MPI), while the pre-processing happens on a single machine.

The PICS portion is divided across slices with approximately equal numbers of slices sent to all machines. In order to do this, the pre-processed sensitivity maps, k-space data, sampling pattern, basis matrix, and other files are sent to each machine using Python and `mpi4py`. Once the data are distributed, the PICS computation proceeds in a hybrid fashion with MPI parallelism across slices and OpenMP parallelism within each slice. The post-processing portion is distributed over only three machines using `mpi4py`, as there are currently three time points being post-processed. The results from the distributed PICS portion are transmitted across the machines using the `Allgather` primitive.

Pre-processing is left to run on a single machine so as to not over-complicate the implementation. This makes the distributed framework useful for other reconstruction methods with a similar parallelism profile but different implementation details. Such a framework would be complicated by excessive use of distributed primitives within areas such as pre-processing that should stay easily customizable by other MRI reconstruction experts.




Figure 4.2: Distributed Implementation including Broadcast and Allreduce MPI primitives.

4.4 Experimental Setup

With institutional review board approval and informed consent/assent, pediatric patients presenting with indications of knee pain were scanned using the T2 Shuffling sequence on a 3 Tesla MRI scanner (GE Healthcare, Waukesha WI) at Lucile Packard Children’s Hospital from April to May, 2018. A representative case from the scans was used to carry out the design and analysis.

4.4.1 Scan Sequence and Reconstruction Parameters

The acquisition and reconstruction dimensions and sizes are shown in Table 4.1.

Table 4.1: Sequence and reconstruction parameters.

Description	Size
Spatial matrix size	$288 \times 260 \times 240$
Receive channels	16
Coil compressed channels	7
Echo train length	80
Soft-SENSE maps	2
Basis coefficient images	4
LLR patch size	12×12
LLR regularization parameter	0.0125
FISTA iterations	250
ESPIRiT calibration size	$288 \times 24 \times 24$

Acquisition T2 Shuffling scans were performed sagittally at 0.6 mm isotropic resolution with a 16-channel GEM-Flex receive coil array. The pulse repetition time and echo spacing were 1200 ms and 6 ms, respectively, and the echo train length was 83. The data from the first three echoes covered the central portion of k-space and were used to perform ESPIRiT calibration [118]. The remaining 80 echoes were used in the iterative reconstruction. Each TE was under-sampled with a variable-density Poisson disc sampling pattern, for an acceleration factor of 139 per time point. Collectively, this represented a relative acceleration of 1.7 and apparent acceleration of 6.9, as defined in Table 3.2. A partial Fourier acceleration of 0.65 was used in the slice direction [121]. The acquisition 3D array was $288 \times 260 \times 240$, and the total scan time was 7 minutes and 2 seconds.

Reconstruction The distributed reconstruction was controlled through a Python script using `mpi4py` for distributed communication. Each underlying component of the reconstruction was implemented using BART. In the reconstruction, two soft-SENSE maps [87] were used and 4 subspace coefficient images were reconstructed. The data were coil compressed from 16 channels to 7 virtual channels [78]. The fully sampled readout direction was inverse Fourier transformed, and each slice was independently solved with 250 iterations of FISTA. The data were normalized based on the maximum signal value of the raw data and the LLR regularization parameter was held fixed at 0.0125 across all scans. Homodyne filtering was applied to each reconstructed virtual coil image and coil-combined with root sum-of-squares. Gradient non-linearity correction and DICOM generation was performed using the Orchestra reconstruction library (GE Healthcare) and the Ox-BART post-processing code, available at <http://github.com/mrirecon/ox-bart>. The ESPIRiT calibration was performed in two steps using a parameter-free approach [120] to improve robustness to patient scan variability. The first step used the 3D calibration data to estimate the signal subspace [87]. The second step separately estimated the sensitivity maps for each slice immediately before calling the iterative solver.

4.4.2 Experimental Compute Cluster

Each compute node had two (dual-socket) Intel® Xeon® Platinum 8180 processors @ 2.5 GHz. with 28 cores and 38.5 MB total L3 cache per socket. Each node also contained 192 GB DDR4 RAM, arranged as 12 DIMMs \times 16 GB per DIMM in order to fully utilize the system's 12 available memory channels.

Experiments used a RAM-mapped temporary storage space for data between multiple invocations of the BART library. For the timing experiments, the input data were pre-loaded into the temporary storage and the transfer time was not counted in the run-time results. Similarly, the output DICOM files were written to the temporary storage in order to remove the consideration of network transfer time from the experiments. The reasoning for this is that transfer times varied widely depending on other jobs running in the cluster, which was a shared resource.

4.4.3 Hospital Compute Cluster

A four-node system was deployed locally at the hospital, with each node matching the specifications and compilations used in the experimental compute cluster. The compute cluster was placed in a room adjacent to the scanner and connected to the hospital network and MRI scanners with a 1 Gigabit network link. The cluster itself used a 10 Gigabit switch for communication between compute nodes during reconstruction. Figure 4.3 shows the details of the system as it is deployed in the hospital.




Figure 4.3: Configuration deployed in the hospital for fast image reconstruction.

4.5 Results

4.5.1 Single-Machine Optimization

Table 4.2 quantifies the benefits of applying the optimizations described in Section 4.3, averaged over five instances of the same reconstruction. The main improvements to the pre-processing step are a result of the changes to FFT phase modulation and to parallelizing and caching the basis. Accounting for data slicing and joining, the full PICS reconstruction is sped up from 275 seconds to 109 seconds on one machine. The post-processing is reduced from 70 to 18 seconds, benefiting from parallel processing. Overall, the optimizations led to a 3x speedup. Due to the processing gain, reconstruction quality improvements were added at a cost of 30 seconds on a single machine.

Figure 4.4 shows the effect of increasing the number of cores per problem while running many independent PICS slice reconstructions in parallel. All cores on the machine are fully occupied for each point on the graph. In the case of one core per problem, there are 56 independent problems being solved simultaneously. For two cores per problem there are 28 independent simultaneous problems, and so on. As the number of cores per problem increases, it becomes more likely that the working set derived from processing the slices fits in L2 cache, which is 1 MB/core in the configuration. The for-

Table 4.2: Run-time breakdown comparison for T2 Shuffling baseline implementation, with optimizations, with optimizations and improved quality, and with optimizations and improved quality on the 4-node cluster. Time shown in seconds.

	Baseline	Optimizations Only	Optimizations and Improved Quality	Optimizations and Improved Quality 4-node
Pre-processing				
Conversion to BART Format	2	1	1	1
Noise Whitening	-	-	2	1
Re-sort and Average	8	4	4	4
Coil Compression	7	4	4	4
Scale Factor	4	0	-	-
ESPIRiT Maps	6	6	6	6
Create Basis	38	-	-	-
Re-sort and Project Data	4	1	1	1
Broadcast	-	-	-	1
PICS Reconstruction				
Linear Recon.	20	-	-	-
Linear Recon. Process DICOM	23	-	-	-
Slice and Detailed Recon.	271	-	-	-
Slice and ESPIRiT Maps	-	13	24	9
No Slice and Detailed Recon.	-	94	106	28
Join Reconstructed Slices	-	2	6	1
Post-processing				
Allgather	-	-	-	3
Process All TEs	70	19	18	9
Other	6	6	7	15
Total	458	150	180	85

ward/adjoint operation per-iteration run-time, LLR per-iteration run-time, and the full solver per-iteration run-time are shown separately. The forward/adjoint run-time and full solver run-times improve in performance with increasing cores up to four due to beneficial caching effects from decreasing the number of simultaneous problems running on the system. Beyond four cores, the overall throughput decreases due to increasing overheads of parallelism as well as the impact of serial portions elsewhere in the solver. As a result, four cores per problem was ultimately chosen.




Figure 4.4: PICS run-time for a batch of 288 total slices, varying number of cores assigned per slice. All cores are utilized at each point. Improved performance from 1 core per problem to 4 is due to reduced working set size. Run-time is normalized by number of slices, machines, and iterations.

4.5.2 Multi-Machine Scaling

Figure 4.5 shows the run-time of T2 Shuffling when executed on one up to 12 machines. The PICS portion is the only part that is distributed, so therefore this portion scales with the number of machines added. However, the computation is currently limited to one machine for pre-processing and three machines for post-processing, so these components quickly become the bottleneck as the number of machines increases. The pre and post-processing steps were not distributed further so as to hide the distributed computing details from future users of the software who are more productive in the shared-memory BART environment. The four-node configuration, which was deployed in the hospital, provides a total compute time of roughly 85 seconds.

Since the collectives are implemented in Python using `mpi4py`, most of the data transfer run-time is being spent loading data into Python from disk and preparing it for the collective operation, not in the communication itself. This is consistent with the experimental results, in which the total time spent in the distributed collective operations Broadcast and Allgather stays roughly constant as the number of machines is increased from two to 12. There is a decrease in the post-processing run-time from one to three nodes, as it is distributed across three machines. There is also a sharp improvement in PICS run-time between 10 and 11 nodes. This is the result of load balancing. At this scale, each machine processes 14 slices in parallel. At 10 nodes, this requires three steps to fully pass through 288 slices. At 11 nodes, there are no slices left over after two steps, so the performance

improves sharply due to the improved overall load balance.




Figure 4.5: Run-time of a T2 Shuffling reconstruction on configurations from 1 to 12 machines. A 60-second total run-time is achieved with 11 machines.

4.5.3 Hospital Cluster Statistics

Over the one-month deployment period, 35 pediatric knee T2 Shuffling scans were performed. The mean, standard deviation, and maximum times for the main steps of the processing pipeline across these scans are shown in Table 4.3. The increase in run-time relative to that reported in Table 4.2 is attributed to latency associated with reading the raw data and writing the DICOM images across the 1 Gigabit network. There was an additional overhead of 10-20 seconds while writing to PACS, which is common to all offline reconstruction implementations.

4.5.4 Multi-Scanner Simulation

In the initial deployment of the hospital compute cluster, only four scanners were connected, and the cluster was only used for T2 Shuffling knee exams. This section investigates the possible tradeoffs involved with serving many exams from many scanners using compute clusters, if the current distributed system and reconstruction method were used for all exams from these scanners. This serves as an upper bound on the amount of computation over the course of the day.

Table 4.3: Reconstruction run-time statistics in seconds for 35 pediatric knee scans over a one-month period.

	Mean (s)	Std. Deviation (s)	Max (s)
Pre-processing	27.97	2.70	34.40
Broadcast	2.26	0.16	2.73
Reconstruction	37.24	4.14	52.60
Allgather	4.47	0.31	5.16
Post-processing	18.27	1.59	23.19
Cleanup	1.13	0.07	1.34
Total	91.85	7.02	108.68

Figure 4.6 shows the MRI scanner schedule from one day out of the one-month study at Lucile Packard Children’s Hospital. Each row represents an MRI scanner, and each entry shows the duration of a scan. The dotted box shows a zoomed-in portion from the schedule to better visualize the individual scans. With the exception of one scanner used in the Emergency Department and reserved for emergencies, there is high utilization across the scanners.




Figure 4.6: MRI scanner schedule from one day out of the one-month study. Each row represents an MRI scanner, and each entry shows the duration of a scan. The dotted box shows a zoomed-in portion from the schedule to better visualize the individual scans.

Using this data, a simulation is performed for the case where one or more compute clusters are reconstructing images from the scanners. To do so, the image sizes are used

to estimate the processing times if they were all to employ PICS-based acquisitions and reconstructions. The cost of the distributed primitives in the cluster were estimated by scaling the measured values from the experiments by the total size of the volume compared to the experimental data. For simplicity, it was assumed that the pre-processing and post-processing times scaled as $O(N_x N_y N_z \log(N_x N_y N_z))$, and that the PICS time scaled as $O(N_x N_y N_z \log(N_y N_z))$; that is, the processing is dominated by the FFT factors and linearly scales with the number of concurrent slices. Constants for each stage were fit to match the processing time for the T2 Shuffling knee dataset (Table 4.2).

The `simpy` discrete event simulation library was used to model contention on the cluster. The model has one process representing the hospital, which creates a new reconstruction process for each scan reconstruction task as they arrive. Each scan reconstruction task tries to access a priority resource representing the clusters. Earlier scans are given higher priority. Once a cluster resource is acquired, the reconstruction process holds the resource for a number of time units equal to our run-time estimates for that particular scan. The time spent waiting and reconstructing are logged for each task throughout the simulated day.




Figure 4.7: Simulated average total reconstruction time and average wait time running reconstruction tasks for all scanners in a large hospital over the course of one day, for different cluster configurations

Figure 4.7 shows the result of the simulation. The average wait time is the average time that each reconstruction task spends waiting for the cluster resource. Average total time is the average compute time plus the wait time. The average wait time is reduced

significantly when using two clusters rather than one, but there are diminishing returns to adding more clusters after this. The plot also informs how to best configure the clusters for a fixed number of machines. Given 8 machines, it makes the most sense, in terms of average and max total reconstruction time, to configure these as two clusters with four nodes per cluster. For this configuration, clinically feasible run-times ($< 2\text{min}$) are achieved with average wait time near zero, and slightly higher max time performance of roughly 200 seconds due to some scans which have very large matrix sizes.

4.6 Conclusion

Developments to PICS-based acquisitions have shown great promise at reducing MRI scan times, but the reconstruction times have remained a barrier to clinical adoption. This chapter explored a distributed reconstruction architecture for T2 Shuffling that achieved clinically feasible latency times. The system was deployed at Lucile Packard Children's Hospital and used for pediatric knee imaging. The distributed reconstruction approach is well-suited for other PICS acquisitions in which the slices can be independently reconstructed, and it presents an avenue for other reconstruction formulations to benefit from the distributed optimizations.

The optimization of the pipeline led to an overall improvement in performance by $5.3\times$ after speedups and adding additional operations to improve image quality. A $3\times$ speedup came from single-node improvements and a $2.1\times$ speedup came from distributed parallelization on a 4-node cluster. By enabling clinically feasible reconstruction times, a larger number of clinical protocols could potentially leverage PICS acquisitions, leading to shorter overall exam times. Thus, scanner utilization should improve as more patients can be accommodated into the same schedule. Dedicated compute clusters could support this increased utilization without the need for transferring data outside the hospital network.

Chapter 5

Clinical Validation

5.1 Introduction

MRI is the predominant imaging method for evaluating musculoskeletal (MSK) pathology based on its ability to provide superior soft tissue contrast and lack of ionizing radiation. In children and adolescents, the knee is the most common joint assessed with MRI [123]. In very young pediatric patients, knee MRI is often used to evaluate developmental conditions and non-traumatic pain [124]. However, the often lengthy exam times continue to be challenging for the pediatric population, compounded by other factors such as patient anxiety and discomfort. Efforts to optimize pediatric MRI protocols aim to maximize information acquired and minimize scan time, often working against an unpredictable deadline when cooperation of the pediatric patient is exceeded [125, 126]. Sedation of the pediatric patient increases the likelihood of a high quality scan, but carries associated risks, including increasing concern for anesthesia associated neurocognitive impairment [127]. Moreover, the need for an MRI exam under sedation requires a complex multidisciplinary approach [128] and can decrease the overall efficiency of a radiology department [129]. A secondary effect of the long exams is high cost, limiting the value of the MRI exam in the context of clinical care. Nevertheless, the basic conventional 2D knee MRI protocol still requires multiple acquisitions in axial, coronal, and sagittal planes due to voxel anisotropy, and also includes sequences with different contrasts to optimize assessment of various structures and pathology [125, 130]. A single-sequence MRI protocol can potentially alleviate some of the difficulties encountered in pediatric MRI.

Volumetric variants of FSE [3, 40] have been described to simplify and accelerate joint MRI based on potential reformattability into multiple planes [131–133]. However in clinical practice, standard 2D imaging still predominates largely due to blurring of the volumetric images. This is likely due to long echo trains required for scan efficiency of volumetric imaging, which leads to image blurring due to T2 decay [4, 40]. This work evaluates T2 Shuffling [117] as an alternative volumetric FSE acquisition due to the re-

duction in blurring that results from modeling the sampled echo time progression. This roughly seven-minute acquisition effectively supports a four-dimensional reconstruction, producing images corresponding to each echo time within the echo train, and thus with image contrast varying from proton density (PD) to heavy T2 weighting. Therefore, a single acquisition yields multiple sets of sagittal images with isotropic resolution, which can be reformatted to axial, coronal, and arbitrary oblique planes. This technique has the potential to significantly streamline joint MRI, which may be especially advantageous in pediatric imaging.

This chapter first presents a clinical study to test the hypothesis that T2 Shuffling can suffice as a single-sequence rapid pediatric knee protocol. In Section 5.2, T2 Shuffling is compared to the routine knee MRI protocol for image quality and diagnostic performance. Building on the evidence that missing clinically relevant pathology is unlikely, Section 5.3 evaluates a targeted knee MRI exam based on T2 Shuffling with significantly shorter exam times. As the table time is significantly reduced compared to the conventional protocol, patients are scheduled for the targeted exam between existing slots. The targeted knee MRI is billed with a reduced charge modifier code, lowering the cost by a factor of three compared to the standard lower extremity joint MRI. As additional patients can now be accommodated into the MRI schedule, the cost of the MRI is simply professional costs of the radiologist and minimal costs associated with reporting infrastructure. The magnet, technologists, and receptionists are already present. A significant additional benefit is same-day scheduling, which reduces lost productivity of parents and children at their work and school.

5.2 Comparison to 2D Multi-Planar Knee MRI

T2 Shuffling modifies the 3D FSE sequence to randomize the echo train view ordering and re-sample positions in k-space at multiple echo times. The revised acquisition requires a modified compressed sensing iterative reconstruction and results in sharp images at each echo time. The overall methodology is described in Chapter 3.

5.2.1 Materials and Methods

Data Acquisition

The echo train phase encode ordering was chosen such that at each echo time, the sampling follows an under-sampled variable-density Poisson disc distribution, which is known to be suitable for compressed sensing and parallel imaging [11, 134]. An exception to the rule was the first two echoes, which are typically discarded in product sequences. The data in these echoes covered the central portion of k-space (18-by-18) and were used for parallel imaging auto-calibration of the ESPIRiT maps [87]. Figure 3.9 depicts the sampling and sensitivity map calibration.

Reconstruction

Prior to the study, a T2 map of an adult volunteer's knee was acquired and used to simulate a distribution of signal evolutions with variable refocusing flip angles using the Extended Phase Graph Algorithm [4, 29]. A basis representing the relaxation curves subspace was formed through principal component analysis [53]. The iterative reconstruction differed from (3.15) in that both LLR regularization and ℓ_1 -wavelet regularization were employed, as follows:

$$\min_{\alpha} \frac{1}{2} \|\mathbf{y} - \mathbf{E}\Phi_K \alpha\|_2^2 + \lambda_1 \sum_r \|R_r(\alpha)\|_* + \lambda_2 \|\Psi \mathbf{x}\|_1. \quad (5.1)$$

Spatial wavelet regularization was added as an additional sparsity constraint as it appeared to qualitatively improve the SNR in the smaller subspace coefficient images, which tended to have significantly lower signal amplitudes (i.e. SNR) compared to the first subspace coefficient image. A non-distributed reconstruction based on BART [88] was used. The ADMM algorithm [20, 22] was used to solve (5.1) in order to handle the multiple regularization terms. The regularization parameters λ_1 and λ_2 were experimentally tuned and kept constant for all reconstructions. To reduce reconstruction time, a software coil compression [86] was applied to compress the 8 channels of data to 6 virtual channels. As the full reconstruction is computationally intensive due to the multiple regularization terms, for the purpose of clinical workflow, an initial linear reconstruction with no regularization was solved and transmitted to the scanner console within 10 minutes to verify that the scan was successful. In order to manage diagnosis, only three of the 80 reconstructed images at effective echo times of 27 ms, 63 ms, and 98 ms were transmitted to the PACS. The reconstruction parameters for this study are shown in Table 5.1.

Table 5.1: T2 Shuffling Reconstruction Parameters.

	Full Reconstruction	Fast Linear Reconstruction
Iterations	800	100
LLR Regularization Parameter	0.00065	—
L1-wavelet Regularization Parameter	0.00021	—
Subspace Size	4	1
Approx. Reconstruction Time (min)	60	10

The reconstructions were performed on a dedicated machine with dual Intel Xeon E5-2600 CPUs and 128 GB RAM.

Patient Recruitment and MRI protocol

This prospective study was approved by the institutional review board and complied with the Health Insurance Portability and Accountability Act. Written informed consent/assent was obtained from all subjects. Twenty-eight consecutive pediatric patients

(mean age 14 years, 11 male, 17 female) satisfying the inclusion criterion of referral to the Lucile Packard Stanford Children’s Hospital for clinical knee MRI exams were recruited. Two patients had requests for bilateral knee MRIs, resulting in 30 scans from May to August, 2015. The exclusion criterion was any clinical issue that required gadolinium-based contrast enhanced imaging. Subjects underwent the institution’s standard clinical knee MRI protocol followed by the T2 Shuffling sequence. All scans were performed on a clinical 3T MRI scanner (MR750, GE Healthcare, Waukesha, WI, USA) with an 8-channel knee coil. T2 Shuffling scan parameters include TR of 1400 ms, ETL of 82, bandwidth of 62.5 kHz, 16 cm field of view (FOV), 288x288 matrix, 0.6 mm isotropic resolution, 85% fat suppression efficiency, 240 sagittal source slices reconstructed to 472 slices, with total scan time of approximately 7 minutes. The standard clinical 2D FSE knee protocol consisted of coronal T1-weighted, sagittal PD, fat-suppressed axial PD, and fat-suppressed sagittal and coronal T2-weighted images. The scan parameters include 14 cm FOV, 2.5 mm slice thickness, and 0.5 mm slice gap. Other parameters vary by sequence and are summarized in Table 5.2. Scan time for 2D sequences totaled approximately 20 minutes.

Table 5.2: Conventional knee MRI and T2 Shuffling scan parameters.

	Axial PD FS	Coronal T2 FS	Coronal T1	Sagittal PD	Sagittal T2 FS	Sagittal T2 Shuffling FS
TR (ms)	3000	3700	600	3000	3000	1400
TE (ms)	30	56	18	9	60	—
Echo spacing (ms)	9.6	11.5	9.6	9.7	10.5	—
Echo train length	6	8	3	8	6	82
Bandwidth (kHz)	50	41.6	50	41.6	31.25	62.5
Field of view (cm)	14	14	14	14	14	16
Matrix size	416 × 256	416 × 256	416 × 256	416 × 256	384 × 320	288 × 260
Slices	38	32	28	32	29	240
Slice thickness (mm)	2.5	2.5	2.5	2.5	2.5	0.6
Slice gap (mm)	0.5	0.5	0.5	0.5	0.5	—
Scan time (min)	2:30	4:00	2:30	2:00	2:30	7:00

Image Analysis

The T2 Shuffling images were reconstructed at three effective echo times corresponding to PD, intermediate, and T2-weighted contrasts and reviewed on a multi-planar reformat capable workstation (Osirix). Though the source images were acquired at 0.6 mm slice thickness in the sagittal plane, they were reviewed at a slice thickness format of 2.5 mm in sagittal, axial, and coronal planes in a similar fashion as the conventional 2D images. To evaluate the diagnostic performance of T2 Shuffling, two radiologists (S. Bao and S.S. Vasanawala, 5 and 15 years experience respectively), blinded to clinical history, independently interpreted each case, forming a structured diagnostic report based solely on the

T2 Shuffling images. To avoid recall bias, this was performed prior to image quality evaluation.

Then, each reader independently evaluated T2 Shuffling image quality. The quality of delineation of 9 anatomic structures were rated on a five-point scale from non-diagnostic to outstanding quality, detailed in Table 5.3. Then, T2 Shuffling was compared to conventional 2D and the delineation of the 9 anatomic structures, as well as the appearance of bone marrow and joint fluid at different contrasts, were rated on a relative scale (Table 5.4). To validate the apparent image contrast of the T2 Shuffling images, for five cases signal intensities of muscle, cartilage, meniscus, bone marrow, and joint fluid were measured on T2 Shuffling reconstructions at 10 effective echo times. Regions of interests were placed over the semimembranosus muscle, posterior articular cartilage of the medial femoral condyle, posterior horn of the medial meniscus, marrow space of the medial femoral condyle, and anterior joint fluid.

Table 5.3: Anatomic structures evaluated.

Structures evaluated
Anterior cruciate ligament (ACL)
Posterior cruciate ligament (PCL)
Medial meniscus
Lateral meniscus
Extensor mechanism
Medial collateral ligament (MCL)
Lateral collateral ligament (LCL)
Retinaculum
Cartilage
Fluid
Marrow

Statistical Analysis

Using the clinical diagnostic report based on conventional 2D sequences as the standard of reference, the sensitivity, specificity, and accuracy of the T2 Shuffling structured reports were calculated based on the number of significant concordant findings. Significant findings included meniscal pathology, ligament and tendon injury, cartilage defect, focal bony and soft tissue lesions, bone marrow edema/contusion, and intra-articular bodies. Findings present on both T2 Shuffling structured reports and diagnostic reports were classified as concordant. Findings present on diagnostic reports and not present on T2 Shuffling structured reports were classified as false negatives. Findings detected on T2 Shuffling but not identified in diagnostic reports were classified as false positives. The

Table 5.4: Structures delineation ratings criteria.

T2 shuffling structure delineation rating scale		
1	—	Non-diagnostic - cannot see structure
2	—	Limited - can see structure but not evaluate for pathology
3	—	Diagnostic - can evaluate structure with some confidence
4	—	Good - can evaluate structure with high confidence
5	—	Outstanding - best quality of delineation
T2 shuffling and 2D comparison scale		
-2	—	Conventional 2D more delineation
-1	—	Conventional 2D preferred
0	—	Same
1	—	T2 shuffling preferred
2	—	T2 shuffling more delineation

proportion of concordant findings between the T2 Shuffling structured interpretations and diagnostic reports was calculated with 95% confidence interval (CI).

After this initial analysis, a consensus review of both conventional 2D and T2 Shuffling images was conducted as well as chart review. The rationale for consensus review was that diagnostic reports were imperfect as the sole reference standard, and as arthroscopy is less prevalent in the pediatric population than in adults [135], a consensus review would better establish a gold standard, as well as to determine the detection of pathologic findings on T2 Shuffling images. The consensus data were used to re-calculate sensitivity, specificity, accuracy, and 95% CI for concordance. A 95% CI for proportion of diagnostic or better rating was calculated for T2 Shuffling structure delineation. A Wilcoxon rank-sum test assessed the null hypothesis that the relative quality of T2 Shuffling structure delineation compared to conventional 2D is unchanged. Intraclass correlation coefficients were calculated to evaluate inter-observer agreement.

5.2.2 Results

Of the 30 cases reviewed, there were 10 cases without clinically significant findings. There were 9 findings of meniscal pathology, including one case of discoid lateral meniscus and one of post-surgical change of meniscal saucerization. There were 10 instances of ligamentous or tendon abnormalities, 3 cartilage defects, 5 focal bony or soft tissue lesions, including osteochondroma and osteonecrosis, 3 joint bodies, and 13 cases with findings of bone marrow edema, contusion, or fracture. There were 9 cases with more than one clinically significant finding.

For reader 1, there was 81% (95% CI of 70-92%) concordance between T2 Shuffling structured reports and diagnostic reports, with a sensitivity of 77%, specificity of 100%, and accuracy of 81%. For reader 2, there was 87% (95% CI of 78-96%) concordance, with

a sensitivity of 84%, specificity of 91%, and accuracy of 85%. There were 10 discrepancies between T2 Shuffling structured reports and diagnostic reports by either reader and 4 discrepancies by both readers. On consensus review of the discrepancies, three discrepant findings, a lateral meniscus posterior horn tear, a focal patellar cartilage defect, and discoid lateral meniscus, were less well visualized on T2 Shuffling compared to conventional 2D, even in retrospect. Of the remaining discrepancies, all missed findings were equally visualized on T2 Shuffling and 2D sequences on consensus review. The discrepant findings and consensus results are detailed in Table 5.5. Four findings were missed by both readers: one case of medial retinaculum avulsion; one lateral meniscus posterior horn tear, and two cases of joint bodies. Of those missed findings, only the case of the lateral meniscus tear was considered to be less delineated on T2 Shuffling on review while the other three were equally delineated.

Table 5.5: Discrepancies and consensus outcome.

Discrepancy	Reader 1	Reader 2	Consensus Outcome
Tibial plateau fracture	+		Similar delineation on T2 Shuffling and 2D
MCL sprain (false positive)		+	Similar delineation on T2 Shuffling and 2D
Medial retinaculum avulsion	+	+	Similar delineation on T2 Shuffling and 2D (Figure 5.1)
Patellar cartilage defect	+		Less well seen on T2 Shuffling due to complex signal in hemorrhagic joint fluid
Medial retinaculum sprain		+	Similar delineation on T2 Shuffling and 2D
Lateral meniscus posterior horn tear	+	+	Similar delineation on oblique sagittal reformats (Figure 5.5)
Lateral meniscus central tear	+		Similar delineation on T2 Shuffling and 2D (Figure 5.4)
Post saucerization of lateral meniscus	+		Similar delineation on T2 Shuffling and 2D
Infrapatellar ligament sprain		+	Similar delineation on T2 Shuffling and 2D
Discoid lateral meniscus		+	Similar delineation but may be less obvious on T2 Shuffling
Sliding-Larsen Johansson syndrome	+		Similar delineation on T2 Shuffling and 2D
Joint bodies (3)	+++	++	Similar delineation on T2 Shuffling and 2D (Figure 5.2 and Figure 5.3)

+ indicates discrepancy by the reader.

Figure 5.1 shows T2 Shuffling and conventional 2D images of a medial retinaculum avulsion which was not initially described in the T2 Shuffling structured report but seen with similar delineation on T2 Shuffling and conventional 2D sequences. Figure 5.2 and Figure 5.3 show two cases of joint bodies not prospectively identified on T2 Shuffling, but easily visualized upon consensus review. A subtle tear of the lateral meniscus is shown in Figure 5.4. While this was not prospectively identified on the T2 Shuffling structured report, it is apparent that this can be visualized to a similar degree on T2 Shuffling and 2D images.




Figure 5.1: A 15-year-old boy with medial retinaculum avulsion. (a) T2 Shuffling proton density (PD) axial reformatted and (b) T2 Shuffling T2-weighted (T2W) coronal reformatted MR images of the knee show a medial retinaculum avulsion and adjacent bone marrow edema (arrows) with similar delineation compared to (c) conventional 2D fast spin-echo (FSE) fat-suppressed (FS) PD axial and (d) FS T2W coronal images.




Figure 5.2: A 19-year-old boy with a small joint body (arrows) in the inferolateral joint space with deep infrapatellar bursitis, shown on T2 Shuffling (top row) in PD axial reformatted, inter-mediate coronal reformatted, and T2W sagittal source images (left to right), with similar delineation compared to conventional 2D FSE (bottom row) FS PD axial, FS T2W coronal, and FS T2W sagittal images (left to right). Metal artifact from ACL graft is noted in the distal femur.




Figure 5.3: A 14-year-old girl with a joint body (arrowheads) in the superolateral joint space with complex signal joint effusion. There is similar delineation between T2 Shuffling (top row) PD axial reformatted, intermediate sagittal source, and T2W coronal reformatted images (left to right), and conventional 2D FSE (bottom row) FS PD axial, FS T2W sagittal, and FS T2W coronal images (left to right).




Figure 5.4: A 16-year-old girl with lateral meniscus tear. Subtle tear of the lateral meniscus (arrows) is seen equally on (a-c) T2 Shuffling PD, intermediate, and T2W coronal reformatted images and (d) conventional 2D FSE FS T2W coronal knee MRI.

Figure 5.5 shows a tear of the posterior horn of the lateral meniscus that was not as well seen on the T2 Shuffling sagittal source images compared to conventional 2D, due to difference in the prescribed axis of imaging – oriented to the magnet for T2 Shuffling (thus oblique to the knee) and oriented to the knee on 2D FSE. In this case, the patient's leg was externally rotated. When the T2 Shuffling sagittal images were reformatted relative to the orientation of the knee, the tear could be easily seen (Figure 5c). A case of MCL injury was initially categorized as a false positive as it was identified on T2 Shuffling but not described on the clinical diagnostic report. However, on consensus review, it was apparent that abnormal signal along the deep fibers of the MCL was present both on the T2 Shuffling and 2D images, and thus likely a false negative of the diagnostic report. Using the consensus data, there was 94% concordance between T2 Shuffling and conventional 2D (95% CI of 88-100%), with a sensitivity of 93%, specificity of 100%, and accuracy of 94%.

Figure 5.6a shows the distribution of T2 Shuffling structure delineation ratings for both readers. Quality of delineation of the 9 anatomic structures was at least diagnostic in all cases for both readers (95% CI of a diagnostic or better rating is between 95-100%), except in one case where delineation of the medial collateral ligament, lateral collateral ligament, and retinaculum was limited (however, 95% CI is 90-100% for a diagnostic quality delineation for these three structures). The most frequent rating given for all anatomic structures was outstanding, 17-27 cases out of 30 for reader 1 and 22-28 out of 30 cases for reader 2. Inter-observer agreement was moderate or strong (ICC 0.53-0.79) for all structures except for the posterior collateral ligament, medial meniscus, and cartilage, where there was fair agreement (ICC 0.33-0.43).




Figure 5.5: Tear of the posterior horn of the lateral meniscus in a 15-year-old girl. Lateral meniscus tear (arrow) is well seen on (a) the 2D FSE FS T2W sagittal image but not well visualized on (b) the T2 Shuffling T2W sagittal source image at the same level, likely due to the orientation of the prescribed axis of image acquisition shown on the inset generated from the axial reformatted image of the same dataset. (c) Reformatted T2 Shuffling T2W sagittal image reoriented relative to the knee (see c inset) shows the lateral meniscus tear (arrow).

There was no significant difference in relative quality of structure delineation between T2 Shuffling and conventional 2D except for the retinaculum (Wilcoxon rank-sum test, $p < 0.05$), where 2D was preferred for both readers. Figure 5.6b shows the spectrum of ratings for both readers, with the predominant score being that of no difference between T2 Shuffling and conventional 2D. Additionally, while there is no statistically significant difference, the menisci may be delineated to better advantage on T2 Shuffling, as shown in Figure 5.7, where a radial tear of the lateral meniscus is more conspicuous in the axial plane with T2 Shuffling compared to conventional 2D.

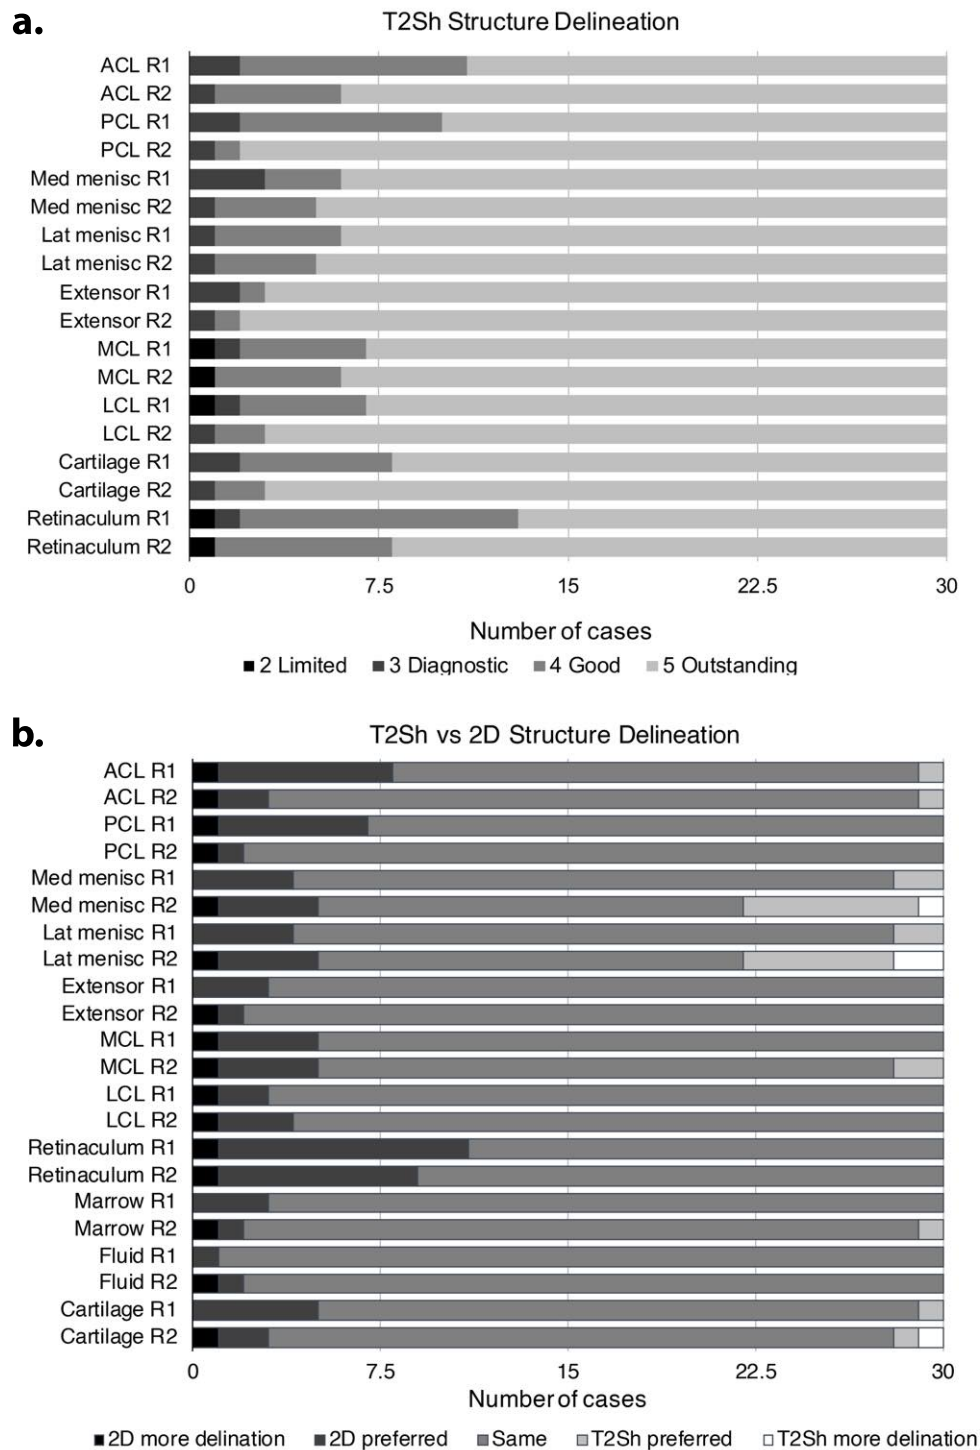


Figure 5.6: Frequency of scores for each anatomic structure for (a) T2 Shuffling structure delineation and (b) relative structure delineation between T2 Shuffling and conventional 2D. Note the higher frequency of higher ratings in both instances. For (a), there were no cases with a rating of 1 (non-diagnostic); therefore, this category was not included.




Figure 5.7: A 16-year-old girl with a radial tear of the lateral meniscus. (a-c) T2Sh T2W sagittal, intermediate coronal reformatted, and PD axial reformatted images with similar delineation of anatomic structures compared to (d-f) 2D FSE FS T2W sagittal, FS T2W coronal, and FS PD axial images. The radial tear of the lateral meniscus (arrows) is well visualized on the T2Sh PD axial reformatted image (c).

5.2.3 Discussion

There have been many studies on the clinical implementation of volumetric FSE methods for knee imaging [132, 136–138], including incorporating sparsity based methods [139–141]. This work seeks to further improve image quality, with the additional capability to provide reconstructions at multiple effective echo times with one acquisition, suggesting the feasibility of a single-sequence pediatric knee MRI protocol based on T2 shuffling. This technique accounts for T2 decay during the long echo trains of volumetric fast spin-echo and thereby mitigates T2 decay-related blurring that has historically limited adoption of the volumetric approach in joint imaging.

The consensus data show high diagnostic accuracy for T2 Shuffling. That most of the discrepancies on retrospective review were equally well seen on T2 Shuffling and conventional 2D, as evidenced in the above figures, suggests that these are likely attributable to readers' experience levels with this approach to imaging and also highlight the value of clinical history and comparison studies. For example, one case of joint body had prior MRIs and radiographs showing this finding, which were not available to the blinded readers. Additionally, most of the discrepant findings were isolated to individual readers, with only four findings missed by both readers. This may also suggest that factors other than those related to the T2 Shuffling technique and image quality have contributed to the assessment of diagnostic performance. Moreover, the pathologies that posed diagnostic challenges for T2 Shuffling in the study, namely lateral meniscus tears and intra-articular bodies, have been shown to be challenging even with conventional 2D knee MRI [142, 143].

Nonetheless, there were a few discrepancies that were thought to be related to the T2 Shuffling technique, although not necessarily due to deficiencies in image quality. A case of lateral meniscus tear was seen in retrospect on T2 Shuffling but not as easily perceived compared to the conventional 2D images. This was felt to be due to external rotation of the patient's leg; as the T2 Shuffling sequence is prescribed in a true sagittal plane, evaluation of the meniscus was limited compared to proper oblique prescription for the 2D images. This was confirmed by viewing T2 Shuffling images in a sagittal plane relative to the knee, which clearly delineated the tear.

One case of patellar cartilage defect was less obviously seen on T2 Shuffling compared to conventional 2D, in part due to the presence of a hemarthrosis, where the signal of cartilage matched the complex signal of the hemorrhagic joint fluid on T2 Shuffling, decreasing the conspicuity of the cartilaginous defect. Lastly, a finding of discoid lateral meniscus was not prospectively noted on T2 Shuffling by one reader. The MRI diagnosis of discoid meniscus is typically based on counting the number of standard sagittal 2D slices on which the body of the meniscus is visualized. However, when interpreting a volumetric sequence, this feature becomes less apparent due to the much-increased number of slices. These examples illustrate some of the inherent differences of this volumetric sequence compared to conventional 2D and suggest a learning curve in interpreting these images.

The results show that the image quality of T2 Shuffling is not inferior to conventional 2D sequences, with high quality delineation of anatomic structures, even those which are primarily viewed on reformatted planes. The isolated case where delineation of the MCL, LCL, and retinaculum was limited on T2 Shuffling was thought to be due to patient body habitus, for which T2 Shuffling was not optimized, but compensated in conventional 2D FSE. However, other artifacts such as the commonly encountered popliteal artery pulsation artifact which manifests as ghosting on 2D FSE was reduced to mild localized blurring on T2 Shuffling. While T2 Shuffling yields an image quality that is comparable to 2D FSE, it does have a qualitatively different appearance due to its volumetric nature, and a period of acclimation can be expected. Although there is a dramatic increase in through-plane resolution, the in-plane resolution of T2 Shuffling is slightly decreased compared to 2D FSE. The isotropic nature of T2 shuffling has a potential to be more advantageous in the evaluation of complex pathology where it can be reformatted in arbitrary planes and with high spatial resolution.

A limitation of the study is that readers could not be blinded to the sequence given its manifest volumetric nature. A second limitation is the small sample size in this pilot study. Additionally, the gold standard was necessarily imperfect due to the paucity of arthroscopic data.

In conclusion, this study has shown that T2 shuffling in the clinical setting produces sharp, multi-contrast images in a single acquisition, with similar diagnostic performance as conventional 2D FSE and non-inferior image quality, and shows potential for realization of a single-sequence pediatric knee MRI.

5.3 Targeted Rapid Knee MRI Exam

Between 2011 and 2014, approximately 5.6 million sports and recreation related injuries occurred in child and adolescent populations, accounting for 28% of emergency visits [144]. Among these visits, 42% were due to lower extremity injuries [145]. Many types of knee pathology are visualized with MRI. In pediatric populations, common injuries are meniscal pathology, cruciate ligament tears, and cartilage abnormalities [123].

In a typical setting, patients start in the care of an orthopedist. Based on the clinical indications, the orthopedist will order an MRI. After authorization is granted by the patient's insurance plan, the patient is scheduled for an MRI in the following weeks. The scheduling process is long and usually disruptive to work and school for parents and children. Following insurance authorization and exam registration, the patient returns on the scheduled date for the MRI. Based on the order, a time slot of 45 minutes to one hour is allocated. The MRI protocol for routine knee imaging consists of a series of two-dimensional 2D FSE scans, each at a different orientation and a target image contrast weighting. The exam typically takes between 30 minutes and one hour.

This study describes initial experience implementing into clinical practice a targeted rapid MRI exam for pediatric knee imaging. The targeted knee exam is a 10-minute protocol consisting of a localizer and two FSE scans. The first FSE scan is a fat-suppressed sagittal T2 Shuffling, as it provides multiple PD and T2 weighted images and multi-planar reformattability. This was chosen based on the prior study presented in Section 5.2, which indicated that missing clinically relevant pathology is unlikely [118]. After the T2 Shuffling scan, a T1 coronal 2D FSE scan is run. The T1 weighted contrast scan provides complementary structural information to visualize bone and knee anatomy. The distributed reconstruction described in Chapter 4 was used to achieve processing times of about 90 seconds. Thus, the T2 Shuffling images are reconstructed and available on the scanner console before the 2D scan finishes. This enables the technologist to check image quality before the patient leaves the table. As the table time is significantly reduced compared to the conventional protocol, patients are scheduled for the targeted exam between existing slots. The targeted knee MRI is billed with a reduced charge modifier code, lowering the cost by a factor of three compared to the standard lower extremity joint MRI.

5.3.1 Materials and Methods

Patient Recruitment

With institutional review board approval and informed patient consent and/or assent, a targeted knee MRI exam for evaluating pediatric knee pain was implemented at Lucile Packard Children's Hospital on two 3T MRI scanners (MR750, GE Healthcare, Waukesha, WI USA) with 16 channel flex coil array. Insurance plans were reviewed to sub-select those plans with high likelihood of waived authorization, rapid authorization, or retroactive insurance authorization. Patients with indications of anterior knee pain, suspicion of

osteochondral lesions, anterior cruciate ligament tears, and meniscal pathology covered by sub-selected insurance plans were eligible for the rapid protocol knee MRI exam. A mechanism of communication between the orthopedic clinic and an adjacent MRI scanner was established. When a targeted knee MRI was ordered, a Current Protocol Terminology (CPT) modifier code of 52 was included to indicate a reduced charge relative to the routine lower extremity joint without contrast (CPT 73721). From October 2016 to May 2018, 47 subjects (mean age 14, 19 male, 28 female) satisfying the inclusion criteria were recruited. Two patients had requests for bilateral knee MRIs, resulting in 49 targeted knee MRI orders.

Data Acquisition

A 10-minute targeted knee MRI protocol was designed, consisting of a localizer (1 minute), fat-suppressed sagittal T2 Shuffling (7 minute), and a 2D coronal FSE T1 (3 minute). Based on technologist evaluation, scans were repeated or added if necessary, e.g. due to motion artifacts or incorrectly prescribed field of view. The typical prescribed scan parameters, shown in Table 5.6, were modified from those used in Section 5.2. The TR for the T2 Shuffling scan was lowered from 1400 ms to 1200 ms to increase scan efficiency, as the increased T1 weighting was found to be tolerable. In addition, a partial Fourier fraction of 65% was used as it was found to reduce residual incoherent artifacts [121].

Table 5.6: Targeted knee MRI median scan parameters across.

	3D Sagittal T2 Shuffling FS	2D Coronal T1
TR (ms)	1200	700
TE (ms)	—	18
Echo spacing (ms)	6	8
Echo train length	82	3
Bandwidth (kHz)	62.5	50
Fat Saturation Efficiency	0.85	-
Field of view (cm ²)	16 × 14.4	14 × 12.6
Matrix size	288 × 288	416 × 288
Resolution (mm ²)	0.6 × 0.6	0.3 × 0.4
Slices	240	28
Slice thickness (mm)	0.6	2.5
Slice gap (mm)	—	0.5
Partial Fourier fraction	0.65	—
Scan time (min)	7:04	2:50

Image Reconstruction

The T2 Shuffling reconstruction was initially implemented as described in Section 4.2, and the main reconstruction parameters are shown in Table 5.7. End-to-end processing was integrated into the hospital environment so that the DICOM images could be auto-transferred to the scanner console and to the PACS. Using the Orchestra software development kit (GE Healthcare, Waukesha, WI), the raw data are converted to a data format compatible with BART [88]. The reconstruction, including coil compression [78], coil sensitivity calibration [87], iterative compressed sensing [11], and Homodyne filtering [122], were carried out with BART. Following the reconstruction, gradient non-linearity correction was applied, and the result was written to DICOM images using Orchestra.

Table 5.7: T2 Shuffling reconstruction parameters.

Description	Value
Coil compression channels [78]	7
Soft-SENSE maps [87]	2
Basis coefficient images [117]	4
Locally low rank patch size [117]	12×12
Locally low rank regularization [117]	0.0125
Iterations [117]	250

Ongoing work during the study leveraged the benefits of the efficient, distributed reconstruction implementation described in Section 4.3. A four-node high-performance computing cluster was deployed at the hospital in January 2018, and connected to a local network shared by the scanner and the hospital PACS through a one-gigabit-per-second (Gbps) connection. The compute cluster had an internal 10 Gbps network. Each compute node had a dual-socket Intel Xeon Platinum 8180 processor with 28 cores and 192 GB of RAM (Intel, Santa Clara, CA, USA). Reconstructed sagittal source images corresponding to contrast-equivalent TEs [40] of 20 ms (PD), 50 ms (intermediate weighted), and 100 ms (T2 weighted) were sent to the scanner and to PACS. Following the reconstruction, images were reformatted by the technologist on the scanner console into axial, sagittal, and coronal planes with 2.5 mm slice thickness and 50% overlap between slices.

Evaluation

Date and time information was recorded for MRI referral and registration. Start and end times were logged for each scan in the exam period and used to assess total exam times. Descriptive statistics related to the patient cohort and the scans were analyzed, including instances where scans were repeated. Clinical indications leading to the MRI order were also noted. End-to-end reconstruction times were evaluated before and after the deployment of the distributed cluster.

5.3.2 Results

As a result of the CPT modifier code, the cost of the targeted knee MRI was one third lower than the cost of the routine knee MRI. Figure 5.8 shows a bar chart of the primary clinical indication for each exam order. The most common reason for exam referral was to evaluate meniscal tear, followed by knee pain. Figure 5.9 shows coronal T1 and T2 Shuffling reformatted coronally with T2 contrast of a patient who presented with knee pain. Clinical indication pointed to suspicion of meniscal tear, but the MRI finding was a bone bruise. The edema in the bone is well-visualized in the coronal T2 weighted reformat of the T2 Shuffling images. Figure 5.10 shows the MRI findings for a patient who presented with indications of meniscal tear. The T2 Shuffling images were reformatted into PD sagittal, T2 coronal, and intermediate weighting axial images. The reformats acutely visualized the complex tissue structure in the knee and confirmed the diagnosis of meniscal tear.




Figure 5.8: Bar chart of the primary clinical indication for each exam order.




Figure 5.9: Patient presented with knee pain and clinical indication of meniscal tear. (a) Coronal T1, and (b) T2 Shuffling reformatted to coronal T2 MRIs are shown. The MRI finding was bone bruise, as indicated by the yellow arrows.






Figure 5.10: A 13-year-old female patient evaluated for internal derangement of the left knee. T2 Shuffling images reformatted into (a) sagittal PD, (b) coronal T2, and (c) axial intermediate weighting. Clinical indication was meniscal tear, and was confirmed with MRI (yellow arrows). Additional findings were discoid meniscus (red arrow) and bone marrow edema (white arrow).

Figure 5.11a shows a histogram of the number of days between exam order to exam registration. Of the 49 targeted knee MRI orders, 20 exams (41%) were completed on the same day as requested. The median time from order to completion was 6 days (minimum of 0 days, maximum of 54 days). The cumulative distribution in Figure 5.11b shows that

the majority of cases were completed in less than one week, and only a small number of cases took more than two weeks to complete.



(a)



(b)

Figure 5.11: (a) Histogram and (b) cumulative distribution of number of days between exam order and exam completion. Mean and median days are marked with solid red and dashed black lines, respectively.

The histogram and cumulative distribution of exam times are shown in Figure 5.12. Median time from registration to exam completion was 18:44 minutes (minimum of 10:28 minutes, maximum of 37:14 minutes). Mean and median exam times are similar, indicating few outliers. There were a large number of exams that took under 15 minutes to complete. This corresponds to 20 cases (41%) that did not require additional or repeated image sequences. Nonetheless, a number of cases had repeated scans, most often due to

motion artifacts. In 13 cases the coronal T1 scan was repeated, and in 5 cases the T2 Shuffling scan was repeated. Additionally, 7 cases included a 2D sagittal T2 fat suppressed scan, likely to increase confidence in the MRI findings in the case of non-diagnostic T2 Shuffling images.

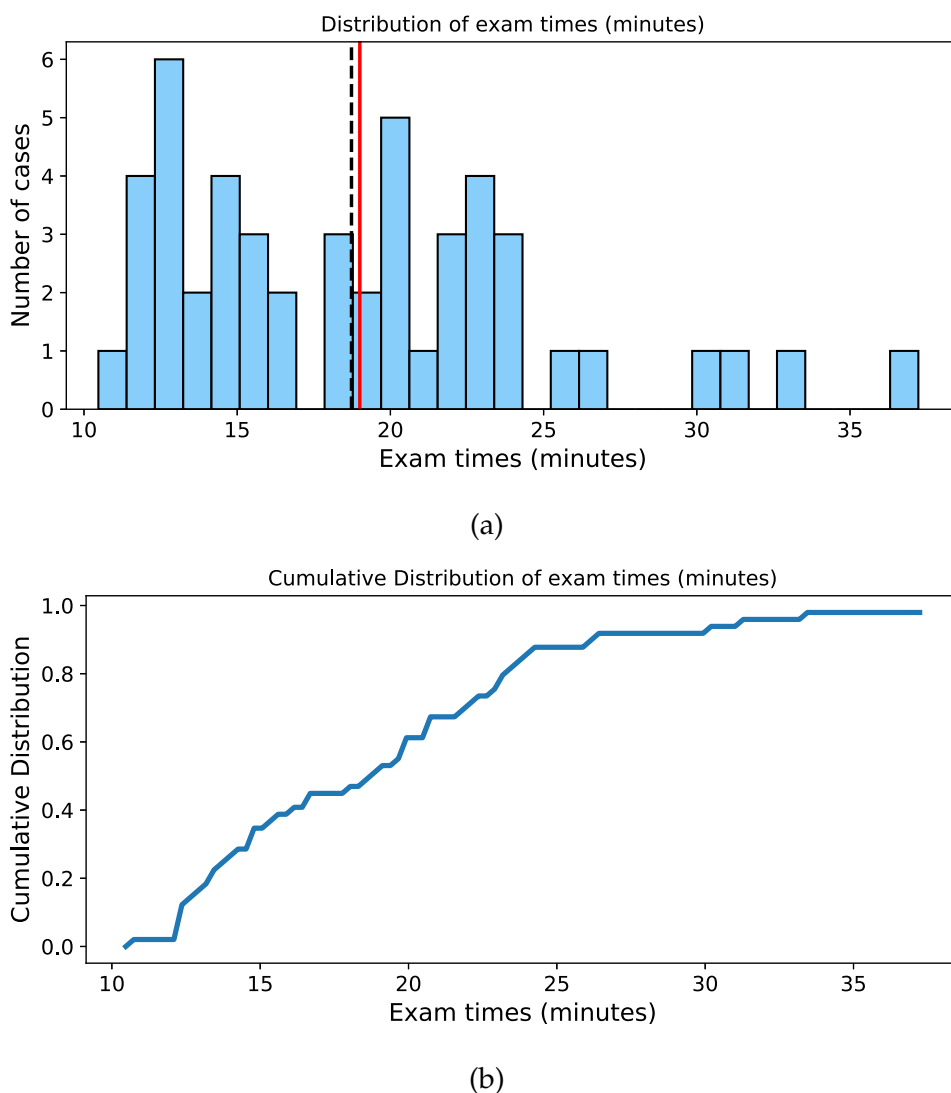


Figure 5.12: (a) Histogram and (b) cumulative distribution of time between exam registration and completion (minutes). Mean and median exam times are marked with solid red and dashed black lines, respectively.

Figure 5.13 shows the histogram and cumulative distribution of the T2 Shuffling processing times before the introduction of the fast distributed reconstruction. The median processing time was 18:48 minutes (minimum of 15:06 minutes, maximum of 83:35 minutes). The processing time was longer compared to the baseline implementation in de-

scribed Chapter 4 because it ran on older generation hardware and was not optimized for the specific architecture. There were two cases where the processing times exceeded 26 minutes. These were a result of using an older backup machine for the reconstruction during machine maintenance. After deploying the new reconstruction, the median processing time was reduced to 1:35 minutes (minimum of 1:09 minutes, maximum of 3:00 minutes), seen in Figure 5.14. The dramatic speedup in processing times was primarily due to (i) the use of multiple machines for each reconstruction, and (ii) optimizing core components of the iterative reconstruction algorithm.

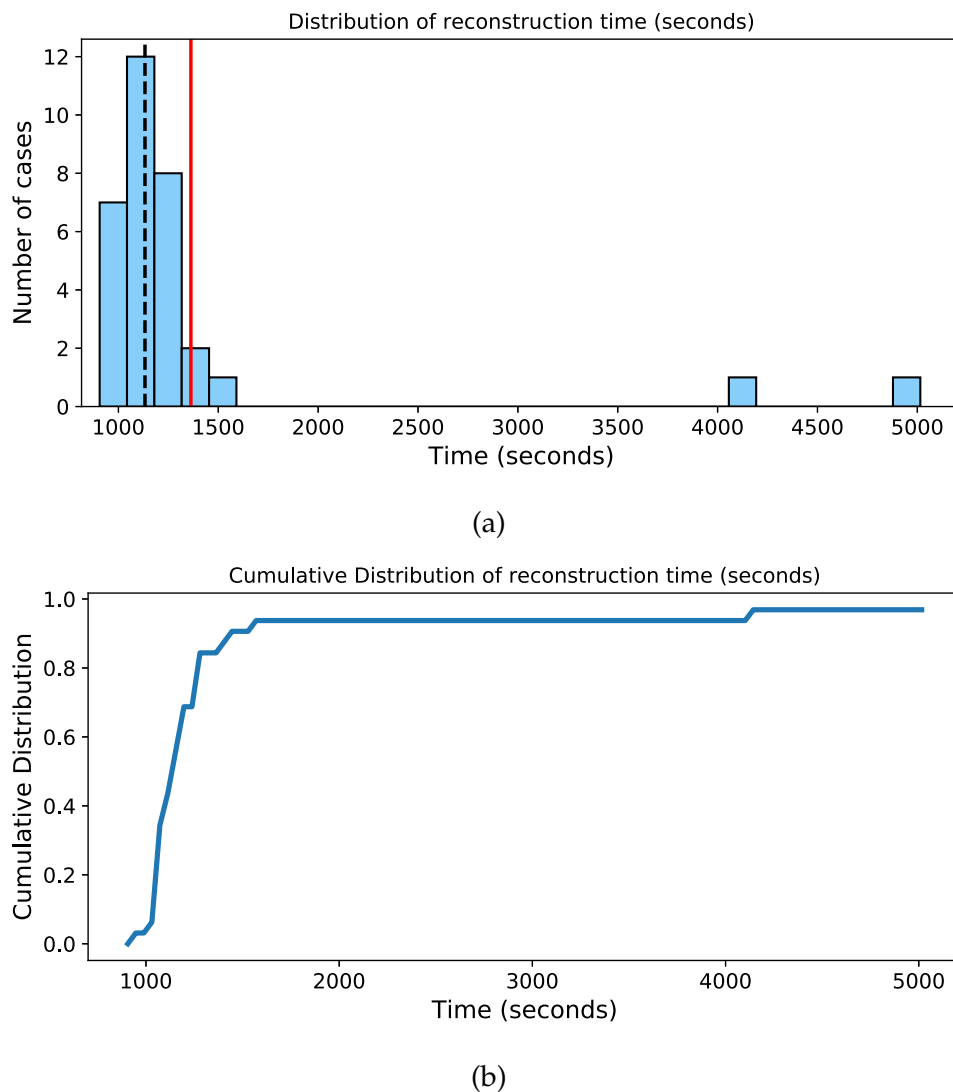


Figure 5.13: (a) Histogram and (b) cumulative distribution of T2 Shuffling reconstruction time (seconds) before deployment of the distributed processing. Mean and median reconstruction times are marked with solid red and dashed black lines, respectively.

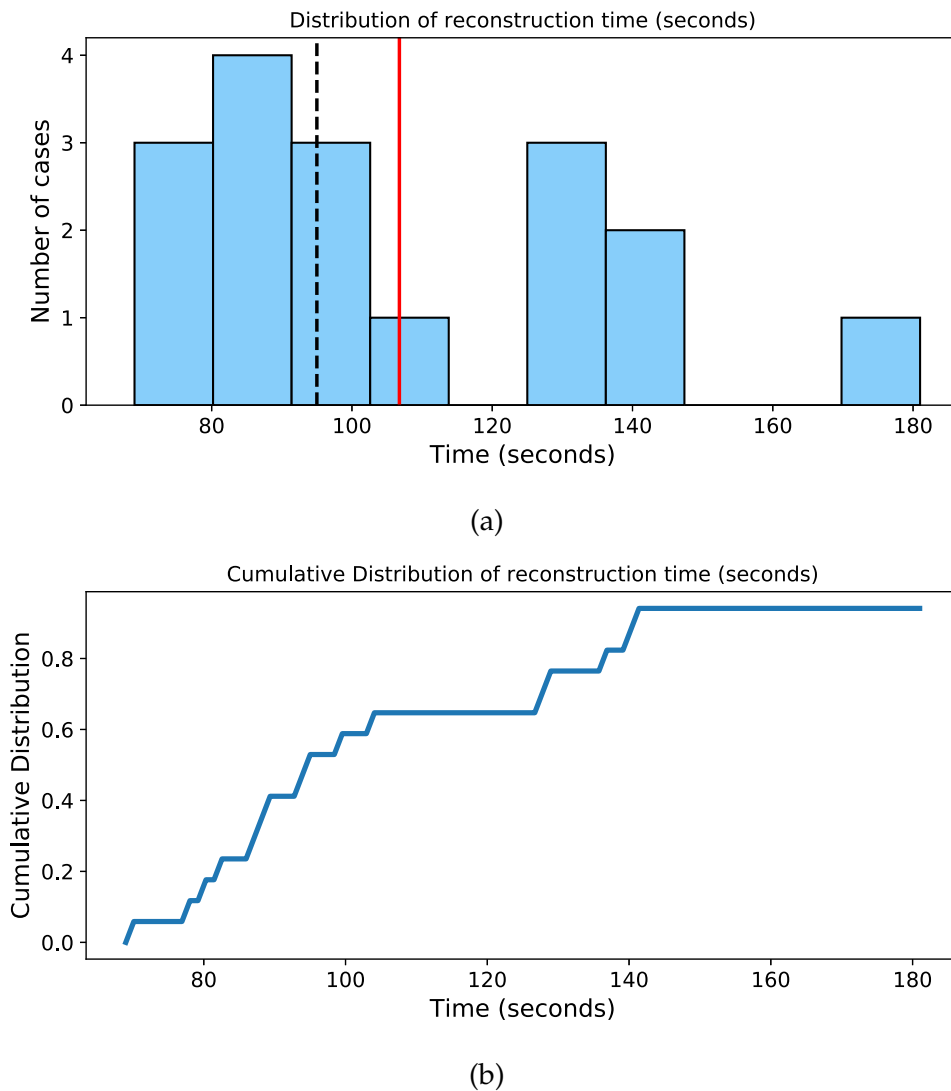


Figure 5.14: (a) Histogram and (b) cumulative distribution of T2 Shuffling reconstruction time (seconds) after deployment of the distributed processing. Mean and median reconstruction times are marked with solid red and dashed black lines, respectively.

5.3.3 Discussion

Many works aim to bring accelerated scanning methods to the clinic in an effort to reduce exam times and increase patient throughput. An important component to the MR value equation is the cost to the patient. This study evaluated a targeted knee MRI exam that was designed for specific clinical indications in a pediatric population with the goal of reducing time, cost, and burden to the patient and family. The exam was based on T2 Shuffling as a mechanism for obtaining fat-suppressed images in multiple orientations and

contrasts. The T2 Shuffling scan was paired with a 2D T1 coronal scan in order to obtain complementary information about bone and anatomy. As the imaging time was significantly reduced compared to the conventional knee MRI, a reduced charge CPT modifier code was used for insurance billing, lowering the cost by a factor of three compared to the routine knee MRI.

The targeted knee MRI was designed to fit in between already scheduled patients during the day. The purpose for this decision was so that patients could be seen the same day as their initial visit to the orthopedist. In this study, 20 of the 47 patients were scanned on the same day. Thus, a number of situations arise in which (i) the schedule could not accommodate the MRI, or (ii) the patients and parents elected to return on a different day as a matter of convenience. Although there was no obvious pattern based on age or time of year, the time of the scan was a positive predictor for likelihood of same-day scan: 15 of the 18 cases with exam registration between 2 pm and 6 pm were same-day orders. Conversely, only 5 out of 31 cases were same-day orders outside this window. This may indicate that patients who are seen by an orthopedist in the morning are more likely to be subsequently scheduled for their MRI later that day.

Since the deployment of the distributed reconstruction, the reconstruction times for T2 Shuffling no longer impacted the technologists' workflow, as the images were available on the scanner console before the coronal T1 scan finished. In contrast to the study presented in Section 5.2, the technologist could immediately review the images for diagnostic quality and initiate the multi-planar reformats. The distributed cluster used four high-performance compute nodes, each with a cost of about \$10,000. Although the high cost of the servers may pose a barrier to the use of the MRI exam, their amortized cost is expected to be low, as similar targeted MRI exams can be developed to leverage T2 Shuffling and other computational MRI methods. Such a scenario was simulated in Section 4.5.

A major challenge with pediatric imaging is patient cooperation during the exam. Although the targeted knee MRI reduces the exam time, a large number of cases still required additional scanning. Surprisingly, the 2D coronal scan was repeated more often than the T2 Shuffling scan. This may be a result of the low risk in re-acquiring a three-minute scan versus the longer seven-minute T2 shuffling scan. A number of cases used additional 2D imaging sequences following the targeted protocol, likely to increase confidence in the MRI diagnosis. As a result, the total scan time was greater than 20 minutes in about half the cases. The need to repeat scans poses a risk in the "on-demand" scheduling model, but it is not unique to the targeted knee MRI experience.

A limitation of this study is that arthroscopic data was not available to corroborate the MRI findings. A second limitation is that there is little control over the interplay between the orthopedists, radiologists, technologists, and the patients. Thus, although the study presents a glimpse into the true hospital workflow, it reduces the ability to accurately attribute cause and effect for scan times and patient throughput. A third limitation was the relatively small number of participants. The targeted knee exam was intentionally limited to a small subset of clinical indications and insurance pre-authorizations during

its initial evaluation.

In conclusion, this study has shown that a targeted knee MRI exam based on T2 Shuffling has the potential to reduce exam times and costs to the patients in a pediatric population.

5.4 Conclusion

The path to clinical adoption requires extensive engineering and validation effort to vet new technological approaches. This chapter presented a systematic approach to clinically evaluate T2 Shuffling in the pediatric knee through comparisons to existing established methods and developments of newly designed protocols. The results demonstrate the potential for the application of T2 Shuffling as a single-sequence comprehensive knee MRI protocol, with multi-planar reformattable images in PD, intermediate, and T2-weighted reconstructions. The prospect of a fast, single-sequence MRI protocol is especially well-suited for pediatric imaging, improving the likelihood for high quality scans and decreasing the need for sedation. Although the volumetric sequence is lengthier than any single 2D sequence, and hence may be susceptible to motion during the scan, motion artifacts did not appear to limit the images any more than in the 2D scans.

Although T2 Shuffling scored favorably as a single-sequence alternative to the conventional knee MRI protocol, there may be a strong preference among some practitioners to obtain at least one nonfat-suppressed sequence. As a result, the targeted knee MRI exam was developed, consisting of fat-suppressed T2 Shuffling and a 2D coronal T1 scan without fat suppression. Based on initial results, the targeted knee exam provides a faster and cheaper alternative for pediatric patients with specific clinical indications. The fast distributed reconstruction obviates the need to change the workflow, as images are available directly on the scanner console and in PACS in a timely fashion.

Chapter 6

T1-T2 Shuffling

6.1 Introduction

The focus of this dissertation has been on T2 Shuffling, an accelerated volumetric FSE acquisition that provides multiple image contrasts from a single scan. T2 Shuffling aims to reduce blurring commonly seen in 3D FSE by accounting for relaxation during the echo trains. Images are reconstructed at each TE in the echo train, resulting in a time series of images with contrast starting at PD and increasing to heavy T2 weighting. Since the acquisition is volumetric with isotropic resolution, the images can be retrospectively reformatted into arbitrary planes.

In addition to PD and T2 imaging, T1 weighted imaging plays an important role in diagnosis. Routine brain imaging protocols at 3T consist of T1, T2, and T2-FLAIR sequences [35]. T1 and T2 are also used for the detection of ischemia in fetal imaging [146]. Musculoskeletal imaging uses T1 sequences to visualize bone anatomy, infection, and tumors [123, 130]. A major challenge with FLAIR-based imaging is the timing of the inversion recovery time (TI). In brain imaging, the TI should be chosen to null the cerebrospinal fluid (CSF) and provide a high contrast-to-noise ratio between gray and white matter [35]. Exact timing can be difficult as it depends both on the scanning environment (e.g. field strength, sequence) and patient variability.

A number of MRI methods have been developed with the aim of simultaneously quantifying PD, T1, and T2 from a single scan. Simultaneous T1 and T2 mapping has been applied to 2D myocardial imaging based on T2 preparation and multiple inversion recoveries [147], as well as with MR fingerprinting [148]. Recently, the authors in [149] have introduced an approach to free-breathing, cardiac and respiratory resolved quantitative imaging that incorporates advances in reconstruction and high-dimensional signal modeling. Quantitative approaches based on FSE have been successfully applied to 2D brain imaging [96] and have been used to retrospectively generate synthetic contrast images [150]. However, due to the conventionally low scan efficiency of FSE, quantitative 3D approaches have typically relied on gradient echo-based sequences [151–153], but the

effective resolutions still remain low – on the order of 1.5-2 mm in-plane and 3-6 mm slice thickness. The approaches above can all be used to generate synthetic contrast images by passing the quantitative maps through the MRI signal equation. However, artifacts due to partial volume, flow, and other un-modeled effects raise questions about their diagnostic utility [154]. Non-parametric modeling based on deep learning has been used to bypass the quantitative step and directly map the under-sampled MRI to synthetic contrast images [155, 156]. These directions show promise at integrating advances in MR sequences with conventional MR image quality. Other methods aim to acquire multiple contrasts in a single scan without modeling [157, 158] by interleaving different types of magnetization preparation blocks throughout the sequence.

This chapter explores an extension to T2 Shuffling that incorporate T1 sensitivity into the acquisition. The volumetric reconstruction can be retrospectively synthesized into PD, T1 weighted, and T2 weighted image contrasts representing different combinations of scan parameters. The method, termed T1-T2 Shuffling [159], is also compatible with other FSE modules, including fast recovery through driven equilibrium [160, 161], multi-point Dixon-based water-fat separation [162], fat suppression through chemical saturation, and T1-prepared inversion recovery.

T1-T2 Shuffling adds the capability to change the TR for each echo train, effectively sampling the saturation recovery curve at multiple time points. The inclusion of multiple TRs adds a new temporal dimension to the data. Ideas from compressed sensing are used to reduce the sample complexity within this multi-dimensional dataset. The T2 relaxation and T1 recovery are explicitly coupled with a low-dimensional relaxation subspace constraint. Similarly to the T2 Shuffling reconstruction, the signal evolutions are further constrained through locally low rank regularization. The low-rank constraints are derived from EPG-based simulations. These simulations are also used to find linear combination weights of the relaxation coefficients to create target synthetic contrast images. The end result is a series of volumetric synthetic contrast images representing different combinations of retrospectively chosen scan parameters.

6.2 Theory

6.2.1 Signal Equation

In Chapter 2, the FSE signal equation was developed for the case of multiple TEs along an echo train and a single TR duration. In practice, the TR duration is not long enough to reach thermal equilibrium. For T2 imaging, the signal weighting due to the finite TR duration will affect each TE equally during the echo train. However, if the TR duration is changed during the acquisition, the signal at each voxel will be weighted by the amount of longitudinal recovery. As a result, the signal evolution of the transverse component at

the i^{th} TE and for a specific TR duration is given by

$$S_{iT_s}(\text{TR}) = M_0(\text{TR}) f_{iT_s}(T_1, T_2) \frac{1 - E_1(\text{TR})}{1 - E_1(\text{TR}) f_{T \cdot T_s}(T_1, T_2)}, \quad (6.1)$$

where $M_0(\text{TR})$ is the equilibrium longitudinal magnetization for a specific TR duration, $f_t(T_1, T_2)$ is the relaxation signal at time t , $E_1(\text{TR}) = \exp\left(-\frac{\text{TR} - T \cdot T_s}{T_1}\right)$, T_s is the echo spacing, and T is the ETL. It is assumed that the echo times are integer multiples of the echo spacing. The reason that a denominator appears in the longitudinal recovery term is because of a -90° fast recovery pulse that tips the remaining transverse magnetization (given by $f_{T \cdot T_s}$) back into the longitudinal axis [160]. The relaxation signal $f_t(T_1, T_2)$ depends on the variable refocusing flip angle scheme. Figure 6.1a shows the transverse and longitudinal magnetization for different combinations of T1 and T2 and the variable flip angles shown in Figure 6.1b. The figure shows the main components comprising the signal. The transverse magnetization is dominated by T2 relaxation; however, due to the variable flip angles, T1 recovery during the echo train will also influence the signal. Following the refocusing flip angles is a fast recovery tip-up pulse that acts to exchange the transverse and longitudinal magnetization components. Finally, the TR duration will limit the longitudinal recovery and dictate the equilibrium magnetization level for the next excitation.

In the simplest case where a single TR duration is repeated for multiple echo trains, the equilibrium magnetization will reach a steady state that can be calculated analytically [161],

$$M_0(\text{TR}) = \frac{1 - E_1(\text{TR})}{1 - f_{T \cdot T_s}(T_1, T_2) E_1(\text{TR})}. \quad (6.2)$$

In practice the equilibrium steady state is reached after a small number of repetitions [161]. Assuming the same TR duration is used for multiple consecutive echo trains, it is convenient to assume that (6.2) holds. Then, (6.1) can be separately analyzed for different TR segments, defined as blocks of consecutive echo trains with the same TR duration.

6.2.2 Low Rank Modeling

The signal evolution behavior given by (6.1) can be modeled for a particular (T1, T2) distribution through the EPG formalism. Given a selection of the TE, flip angles, ETL, and TR segments, an ensemble of signal evolutions will result in the third-order tensor $\mathcal{X} \in \mathbb{C}^{N \times T \times R}$, where N is the number of simulated signals, T is the ETL, and R is the number of TR segments. Figure 6.2 shows the construction of the tensor and its unfolding along the first mode [163]. In this formulation, the flip angles and ETL are identical across TR segments. The (T1, T2) distribution can be produced through conventional T1 and T2 mapping procedures or from values reported in the literature.




Figure 6.1: (a) Simulated transverse (solid) and longitudinal (dashed) magnetization signal evolution curves for FSE with fast recovery across a TR of 1500 ms using variable refocusing flip angles shown in (b). The T1 and T2 values (in ms) are shown to the right of their corresponding signal levels at the end of the TR.

Like in the matrix case, the tensor \mathcal{X} will exhibit low rank structure due to the correlation in the signal evolutions across the TEs and TRs. However, there is not necessarily a single “best-fit” low-rank tensor decomposition, and it may not be computationally tractable [163]. The authors in [164] define the nuclear norm of an n^{th} order tensor as a convex combination of the nuclear norms of the unfolded tensor along each mode,

$$\|\mathcal{Z}\|_* = \sum_{i=1}^n \beta_i \|\mathcal{Z}_{(i)}\|_*, \quad (6.3)$$

where $\beta_i \geq 0$, $\sum_i \beta_i = 1$, and $\mathcal{Z}_{(i)}$ is the unfolded tensor along the i^{th} mode. This definition is attractive because it reduces to the matrix nuclear norm in the two-dimensional

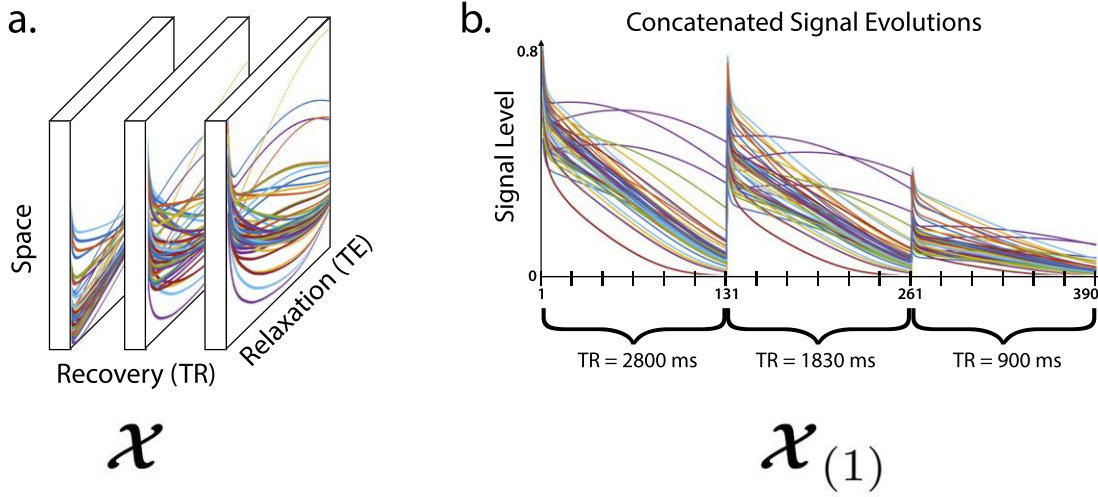


Figure 6.2: Tensor construction for T1-T2 Shuffling. (a) The tensor $\mathcal{X} \in \mathbb{C}^{N \times T \times R}$ consists of signals representing a distribution of T1 and T2 values that follow the signal equation given by (6.1). In this example, $N = 32$, $T = 130$, and $R = 3$. (b) The unfolded tensor along the first mode, $\mathcal{X}_{(1)}$, is a matrix with the relaxation and recovery dimensions concatenated.

case [163, 164]. Under this framework, low rank structure can be implicitly enforced by regularizing the unfolded tensors.

It is also possible to explicitly enforce low rank structure by constraining the unfolded tensors to a low-dimensional subspace [149]. Although in general there exist different subspaces for each mode, this work only considers the unfold along the first mode, as it captures the spatio-temporal relationship in the data. In other words, $\beta_1 = 1$ and $\beta_2 = \beta_3 = 0$. The unfolded tensor along the first mode is constructed by concatenating the signal relaxation curves of all the TR segments into a single temporal curve. Denote this ensemble of signal evolutions as

$$\mathcal{X}_{(1)} \equiv \mathbf{X} = [\mathbf{x}_1 \ \cdots \ \mathbf{x}_N], \quad (6.4)$$

where $\mathbf{x}_n \in \mathbb{R}^{RT}$ is the n^{th} concatenated, simulated signal evolution. A relaxation subspace $\Phi_K \in \mathbb{C}^{RT \times K}$ can be computed through SVD by taking the first K left principal vectors. Neglecting model error, the signal evolutions belong to the subspace, i.e.

$$\mathbf{X} \approx \Phi_K \Phi_K^* \mathbf{X}. \quad (6.5)$$

The relaxation coefficients of a signal evolution \mathbf{x} are given by

$$\boldsymbol{\alpha} = \Phi_K^* \mathbf{x}. \quad (6.6)$$

Figure 6.3 shows the concatenated signal evolutions from the unfolded tensor and the relaxation basis curves generated through SVD.




Figure 6.3: The data tensor is unfolded along the first mode by concatenating the signal evolutions from each TR segment. A subspace is then computed through SVD. In this example, the ETL is $T = 130$, the subspace size is $K = 4$, and there are $R = 3$ TR segments.

6.2.3 Synthetic Contrast

The T1-T2 Shuffling sequence encodes T1 and T2 sensitivity by sampling k-space at multiple TEs and TRs. Due to the subspace relationship, the relaxation coefficients can be used to synthesize images along the relaxation and recovery curves. One approach is to directly solve for the intrinsic relaxation parameters at each voxel given the non-linear signal model [55, 67, 96, 151]. From there, the signal equation is used to retrospectively synthesize images representing different choices of scan parameters. This work considers an alternative approach in which the relaxation coefficients are (linearly) combined to form an approximate target image contrast. The use of linear combinations eliminates the need to solve for the physical parameters, which would involve either non-linear least-squares optimization or grid-based search. In addition, SNR tradeoffs are better understood for linear filtering [94].

Using different linear combinations, weights are pre-calibrated by fitting the relaxation coefficients of the ensemble of simulated signal evolutions to a simulated target image contrast,

$$\mathbf{c}_l = \underset{\mathbf{c}}{\operatorname{argmin}} \frac{1}{2N} \sum_{n=1}^N \left| g_l \left(T_1^{(n)}, T_2^{(n)} \right) - \mathbf{c}^* \boldsymbol{\alpha}^{(n)} \right|^2, \quad (6.7)$$

where $\mathbf{c}_l \in \mathbb{C}^K$ are the weights for the l^{th} target image contrast, $\boldsymbol{\alpha}^{(n)} \in \mathbb{C}^K$ are the relaxation coefficients of the n^{th} simulated signal evolution, and g_l is the l^{th} target contrast function. For example, a spin-echo contrast with prescribed TE and TR durations is simulated through

$$g_{\text{SE}}(T_1, T_2) = e^{\left(-\frac{\text{TE}}{T_2}\right)} \left(1 - e^{\left(-\frac{(\text{TR}-\text{TE})}{T_1}\right)} \right). \quad (6.8)$$

Similarly, an inversion-recovery spin-echo contrast is given by

$$g_{\text{IRSE}}(T_1, T_2) = \left(1 - 2e^{-\frac{\text{TI}}{T_1}} + e^{-\frac{\text{TR}}{T_1}}\right) e^{-\frac{\text{TE}}{T_2}}, \quad (6.9)$$

where TI is the inversion time. To obtain the l^{th} synthetic MR image, the coefficients are linearly combined with the weighting vector at each voxel,

$$x_l(\mathbf{r}) = \mathbf{c}_l^* \boldsymbol{\alpha}(\mathbf{r}), \quad (6.10)$$

where $x_l(\mathbf{r})$ and $\boldsymbol{\alpha}(\mathbf{r})$ are the target contrast and relaxation coefficient images at voxel \mathbf{r} , respectively.

For each desired target contrast, a different linear combination must be computed. However, the weights can be pre-computed offline based on the ensemble of signal evolutions. The goodness of fit will depend on the subspace size, sequence parameters, and SNR. The fit will exhibit the same model error vs. noise error tradeoff discussed in Appendix B. In principle, the target contrast need not be physically realizable. For example, a linear combination could be created that aims to null a specific band in T1-T2 space such as white matter [94]. The fit in (6.7) is an ℓ_2 norm. Other objective functions could also be used, for example to penalize outliers or weight specific signal bands [165].

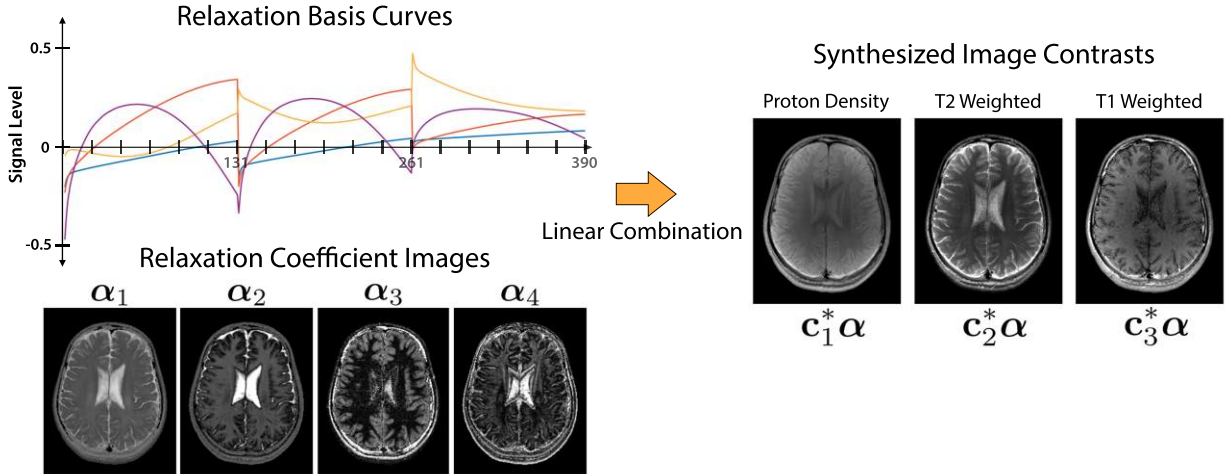


Figure 6.4: Using the subspace and the known T1 and T2 values of the simulated signal evolutions, a linear combination is computed and applied to the basis coefficient images to synthesize arbitrary image contrasts.

6.2.4 Optimized Selection of Repetition Time Durations

The inclusion of multiple TR segments is intended to encode sensitivity to T1 during the acquisition. In effect, this amounts to sampling the saturation recovery (SR) curve at

multiple time points. In the most general case, each echo train could be assigned a unique TR duration. The scan time is then equal to the sum of all the TRs. For example, the echo trains could be chosen to have linearly or logarithmically spaced TR durations starting at TR_{\min} and ending at TR_{\max} . However, other schemes could also be used with potential benefits. For example, scan efficiency could be increased by spending more scan time on short TR durations.

To motivate the TR selection, consider quantitative T1 mapping based on the signal equation (6.1). For N_{trains} echo trains, the goal is to select the TR duration for each echo train that results in the best T1 estimate. Because variable refocusing flip angles are used, longitudinal magnetization will recover throughout the echo train. In addition, the use of fast recovery pulses makes the longitudinal recovery dependent on T2. In this analysis, these effects are ignored as their influence on T1 estimation is small. Neglecting the influence of T1 recovery during the echo train and considering the signal value at a particular TE, the signal equation can be re-cast as a three-parameter model:

$$y_n(a, b, T_1) = a \frac{1 - E_1(\text{TR}_n)}{1 - bE_1(\text{TR}_n)} + w_n, \quad n = 1, \dots, N_{\text{trains}}, \quad (6.11)$$

where $a = M_0 f_{\text{TE}}(T_1, T_2)$, $b = f_{T, T_s}(T_1, T_2)$, and w_n is additive complex-valued white Gaussian noise. The Cramer Rao lower bound, defined as the inverse of the Fisher Information, provides a lower bound on the variance of any unbiased estimator [166, 167]. Thus, a reasonable strategy is to choose the number of echo trains and the corresponding TR durations such that the bound is minimized with respect to T_1 .

There are a number of challenges to the optimization. Under the steady-state equilibrium assumption (6.2), the ordering of the TRs will not influence the Fisher Information. In other words, any permutation of the echo trains will result in the same approximation bound. Denoting the Fisher Information with respect to T1 as $I(T_1)$, this means that

$$I(T_1, (\text{TR}_1, \dots, \text{TR}_{N_{\text{trains}}})) = I\left(T_1, \left(\text{TR}_{\pi_1}, \dots, \text{TR}_{\pi_{N_{\text{trains}}}}\right)\right), \quad (6.12)$$

where $(\pi_1, \dots, \pi_{N_{\text{trains}}})$ is a permutation. Therefore, each local and global minimum will appear $N_{\text{trains}}!$ times. Secondly, the number of echo trains in a given scan is typically in the hundreds, which is challenging to optimize over due the many local and global minima.

For a small number of echo trains, the optimal TR sequence could be found through exhaustive search. For example, Table 6.1 shows the top three solutions (excluding permutations) when selecting TR durations for five echo trains from 29 candidate TR values between 900 ms and 2800 ms. Also shown is the solution corresponding to linearly spaced TR values. In the optimal solution, three out of the 29 TR durations were selected, and the short and medium TRs were each repeated two times. The other top solutions had more than three unique TR durations, but this is likely due to discretization error. Note that this already required searching through $29^5 \approx 20$ million solutions, and the evaluation took 2.5 minutes on a 20-core machine. A similar optimization over 50 possible durations took

Table 6.1: Optimal choices for TR durations found through exhaustive search. The TR durations for five echo trains were selected from 29 candidate values linearly spaced between 900 ms and 2800 ms. Other parameters: $T_1 = 1000$ ms, $T_2 = 100$ ms, $T = 132$, $T_s = 5$ ms, $a = 1$, $b = \exp(-T \cdot T_s/T_2) = 0.0014$, $\sigma^2 = .001$.

	TR Values (ms)	Loss (dB)
1	900, 900, 1850, 1850, 2800	64.4429
2	900, 900, 1782, 1850, 2800	64.4714
3	900, 900, 1850, 1918, 2800	64.4739
\vdots	\vdots	\vdots
616,111*	900, 1375, 1850, 2325, 2800	66.3682

*Corresponds to linearly spaced TR values.

37 minutes, highlighting the computational difficulty in the exhaustive search. Based on this observation, consider the following conjecture:

Conjecture 6.1 *Suppose the measurement model is*

$$y_n(\boldsymbol{\theta}) = f(\boldsymbol{\theta}, u_n) + w_n, \quad n = 1, \dots, N,$$

where $\boldsymbol{\theta} \in \mathbb{R}^R$ is vector of unknown parameters, $\mathbf{u} \in \mathbb{R}^N$ is a vector of control inputs, f is the model, and w_n is additive white Gaussian noise. Then the optimal controls that minimize the inverse of the Fisher Information, $1/I(\boldsymbol{\theta})$, are given by

$$\mathbf{u}^* = [\mathbf{u}_1^\top \quad \dots \quad \mathbf{u}_R^\top]^\top,$$

where $\mathbf{u}_i = \{u_i\}^{N_i}$ for $i = 1, \dots, R$, and $\sum_{i=1}^R N_i = N$. That is, the Cramer Rao lower bound is minimized by repeating R unique control inputs.

Although not proved here, others have reported similar findings [168].

Based on Conjecture 6.1, the TR selection is reduced to choosing $R = 3$ unique TR segments and the number of echo trains in each segment. To this end, the optimization is cast as a mixed integer program:

$$\begin{aligned}
 & \underset{\{\text{TR}_i \in \mathbb{R}, N_i \in \mathbb{Z}\}_{i=1}^R}{\text{minimize}} && \sum_{l=1}^L \frac{1}{I(T_1^{(l)})} \\
 & \text{subject to} && \text{TR}_i \leq \text{TR}_{\max}, \quad i = 1, \dots, R \\
 & && \text{TR}_i \geq \text{TR}_{\min}, \quad i = 1, \dots, R \\
 & && \sum_{i=1}^R \text{TR}_i N_i \leq \text{Scan Time}
 \end{aligned} \tag{6.13}$$

The sum in the objective function is across L target T1 values (e.g. drawn from a distribution of likely values in the brain), and the TR durations are upper and lower bounded. The number of echo trains in the i^{th} TR segment is given by N_i . Different TR schedules can be compared by evaluating (6.13) across the full set of signal relaxation curves as a relative loss, i.e.

$$\mathcal{L}(\{\text{TR}_i\}_{i=1}^R) = 10 \log_{10} \sum_{n=1}^N \frac{1}{I(T_1^{(n)}, \{\text{TR}_i\})} \quad [\text{dB}], \quad (6.14)$$

where N is the number of T1 values considered and $T_1^{(n)}$ is the n^{th} T1 value.

6.3 Methods

6.3.1 Signal Modeling

A joint distribution of (T1, T2) was produced through conventional T1 and T2 mapping procedures, as described in Subsection 6.3.4. An ensemble of $N = 10240$ samples were chosen from the distribution and simulated using EPG [4, 29]. A relaxation basis of size $K = 4$ was constructed from the ensemble of signal evolutions using SVD. The subspace size was chosen as it resulted in less than 1% approximation error across the ensemble. The linear combinations for each desired synthetic image contrast were computed from the ensemble according to (6.7) for different values of TE and TR in (6.8) and TI in (6.9).

6.3.2 Acquisition

The CUBE 3D FSE pulse sequence (GE Healthcare, Waukesha, WI, USA) was previously modified for T2 Shuffling to support randomly shuffled and re-sampled phase encodes during the echo train [117]. The RF pulse sequence timing follows the diagram shown in Figure 2.9, and the sampling patterns at each TE are shown in Figure 3.9. The sequence used refocusing pulses with varying flip angles [3, 40]. This work added the capability to specify the TR duration independently for each echo train. The main scan prescription matched that of conventional 3D FSE, e.g. resolution, matrix size, and field of view. To support multiple TR durations, the maximum scan time was a controllable parameter, and the acceleration factor was automatically determined based on the optimized TR schedule. For a specific anatomy, protocol, and scan time, the TR segments and their corresponding fraction of scan time (based on number of echo trains) were found by solving (6.13) using a genetic algorithm (MATLAB, MathWorks, Santa Clara, California, USA). The optimization was solved with 100 T1 values drawn from a T1 distribution corresponding to a healthy subject's brain imaged at 3T, as well as for $T_1 = 876$ ms, which corresponded to the dominant mode of the distribution. The maximum and minimum allowable TR durations were set to 900 ms and 2800 ms, respectively, and the scan time was

set to 7 minutes. The Fisher Information of the optimized TR schedule was also compared to linearly and logarithmically spaced TR schedules over the same range of TR durations and the same total scan time.

Variable-density Poisson disc sampling patterns were prescribed for each TE and for each TR segment. This enabled good coverage of k-space for each TR segment in a manner that is compatible with both parallel imaging and compressed sensing [11]. After prescribing the sampling patterns, echo train views were ordered for each TR segment. Although there is tremendous flexibility in deciding the echo train ordering, in this work the TRs were played out sequentially in order of largest to smallest TR duration. This scheme was chosen so that the pre-scan could assign the receiver dynamic range based on maximal signal amplitude, and so that the longitudinal magnetization would quickly reach a steady-state equilibrium. The end result was a set of under-sampled k-space data at each TE and TR segment.

6.3.3 Reconstruction

The reconstruction is similar to that of T2 Shuffling and builds off the L1-ESPIRiT algorithm, which combines parallel imaging and compressed sensing in a single reconstruction framework [87]. The data were represented at each acquired TE and TR with the forward model

$$\mathbf{y} = \mathbf{PFSx}, \quad (6.15)$$

where \mathbf{x} are the complex-valued and concatenated signal evolutions at each image voxel, \mathbf{S} are sensitivity maps estimated using ESPIRiT, \mathbf{F} consists of the Fourier transform operators applied to each TE and TR image, \mathbf{P} are the sampling operators that select the acquired k-space samples at each TE and TR, and \mathbf{y} are the acquired complex-valued data. As in T2 Shuffling, the forward model is relaxed with a relaxation subspace constraint (6.5). The reconstruction is then formulated as

$$\underset{\alpha}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{y} - \mathbf{PFS}\Phi\alpha\|_2^2 + \lambda \sum_{\mathbf{r}} \|R_{\mathbf{r}}(\alpha)\|_*, \quad (6.16)$$

where α are the relaxation coefficient images. Locally low rank regularization [69, 117] is applied to the relaxation coefficient images to reduce the number of degrees of freedom. After reconstruction, the relaxation coefficients are linearly combined according to (6.10) to form the target synthetic image contrasts.

The iterative reconstruction is solved with the FISTA algorithm [21] and implemented in BART [19, 88]. The operator commutativity described by (3.19) was used to reduce computation when solving the normal equations. The synthetic contrast linear combination fitting was implemented in Python.

6.3.4 In Vivo Experiments

All experiments were performed under institutional review board approval and informed consent, and scan parameters are shown in Table 6.2. A volunteer's brain was scanned at 3T (Siemens Healthineers, Erlangen, Germany) using conventional 2D multi spin-echo and inversion recovery FSE, and the data were used to obtain reference T1 and T2 maps and a joint (T1, T2) distribution through fitting. Another volunteer's brain was scanned at 3T (GE Healthcare, Waukesha, WI, USA) with a 32-channel head coil using CUBE T1 and CUBE T2 to obtain reference T2 weighted and T1 weighted images, respectively, followed by T1-T2 Shuffling. Three TR segments and their respective fractions of scan time were chosen based on the Cramer-Rao lower bound (6.13) for a fixed scan time of 7 minutes. The data were reconstructed according to (6.16) and linearly combined to generate synthetic contrast images according to (6.7).

Table 6.2: Scan and imaging parameters used for the in vivo experiments.

	Multi Spin-Echo	Inversion Recovery FSE	CUBE-T2	CUBE-T1	T1-T2 Shuffling
Acquisition Type	2D interleaved	2D interleaved	3D non-selective	3D non-selective	3D non-selective
Readout Direction	Axial	Axial	Sagittal	Sagittal	Sagittal
TR (ms)	3000	3000	3000	600	900, 1830, 2800
TR Scan Time Fraction	—	—	—	—	0.34, 0.4, 0.26
TE (ms)	13.2 - 422.4	9	97	13	—
TI (ms)	—	200, 400, 800, 1200	—	—	—
Echo Spacing (ms)	13.2	8.9	4.7	4.4	4.7
Echo Train Length	32	3	130	28	132
Receive Bandwidth (kHz)	33.3	64.5	62.5	62.5	62.5
Field of View	$22 \times 17.5 \text{ cm}^2$	$22 \times 17.5 \text{ cm}^2$	$26 \times 23.4 \times 17.8 \text{ cm}^3$	$26 \times 23.4 \times 17.8 \text{ cm}^3$	$26 \times 23.4 \times 17.8 \text{ cm}^3$
Acquisition Voxel	$0.86 \times 0.86 \text{ mm}^2$	$0.86 \times 0.86 \text{ mm}^2$	$0.8 \times 0.8 \times 1.2 \text{ mm}^3$	$0.8 \times 0.8 \times 1.2 \text{ mm}^3$	$0.8 \times 0.8 \times 1.2 \text{ mm}^3$
Acquisition Matrix	256×204	256×204	$320 \times 288 \times 148$	$320 \times 288 \times 148$	$320 \times 288 \times 148$
Slice Thickness (mm)	5	5	1.2	1.2	1.2
Scan Time (min)	10	14	4	3	7

6.4 Results

Figure 6.5 shows T1 and T2 maps derived from the multi-SE and IR-FSE scans of the first volunteer's brain and the corresponding joint (T1, T2) distribution. The distribution is far from uniform in both directions, emphasizing the importance of tailoring the signal modeling to the 2D distribution of the specific scan anatomy. The Fisher Information optimization (6.13) was solved using 100 values drawn from the T1 marginal probability distribution. Additionally, the optimization was also solved with a single T1 of 876 ms, which was the dominant mode in the T1 marginal distribution. Table 6.3 lists the optimized, linear, and logarithmic TR schedules, the number of echo trains in each schedule, the fraction of scan time for each TR, and the relative loss. The loss was computed over 128 T1 values linearly spaced between 469 ms and 2732 ms. These values were chosen because they captured 90% of the T1 values in the distribution (5% tails). The optimized schedules allocated more echo trains to the short TR segments. This allocation matches intuition: short TR durations afford a higher scan efficiency since more echo trains can

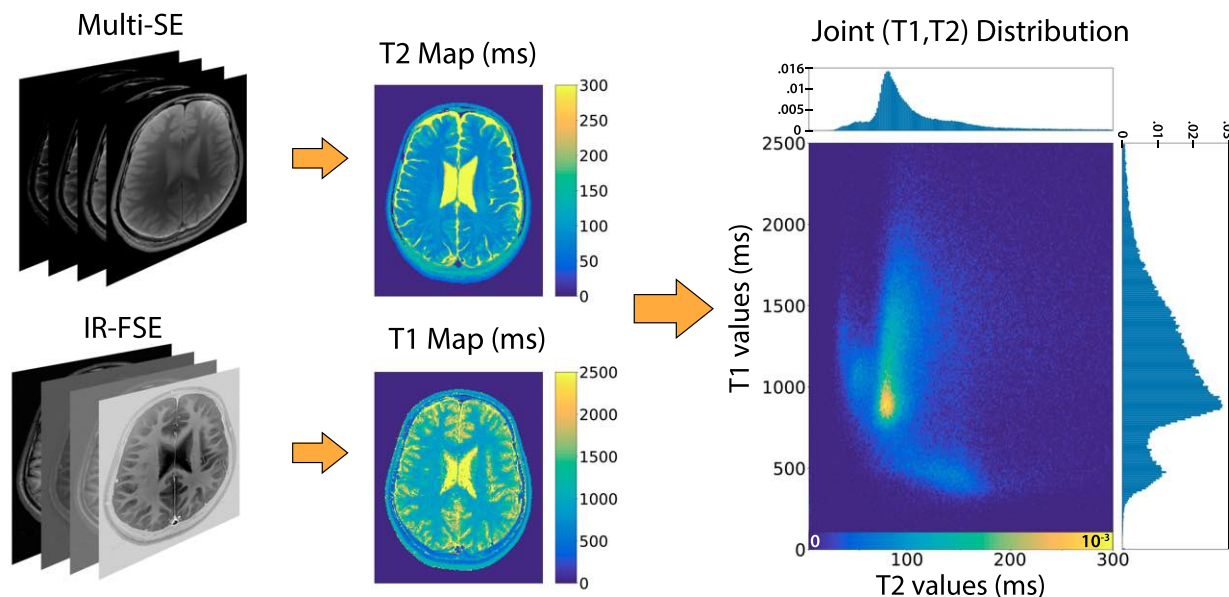


Figure 6.5: Process for creating anatomy-specific joint (T1, T2) distribution from conventional quantitative imaging. Fully sampled brain scans of 2D multi spin-echo and inversion recovery fast spin-echo at six slice locations were used to estimate T1 and T2 maps of the brain at 3T. These maps were used to derive a joint (T1, T2) distribution specific to brain imaging.

Table 6.3: Optimized, linearly spaced, and logarithmically spaced TR schedules and their respective relative losses for a scan time of 7 minutes.

Schedule Type	TR Values (ms)	Echo Trains	Scan Time Fraction	Loss (dB)
Optimized over 100 T1 values	2800, 1830, 900	292	0.26, 0.40, 0.34	88.94
Optimized for $T_1 = 876$ ms	2800, 1791, 962	278	0.27, 0.41, 0.33	89.18
Linearly spaced	900 to 2800	228	—	92.60
Logarithmically spaced	900 to 2800	251	—	92.26

Maximum scan time was set to 7 minutes.

Relative loss was computed over 128 T1 values between 469 ms and 2732 ms.

be acquired in a fixed scan time, but also exhibit lower SNR due to partial longitudinal magnetization recovery. Figure 6.6 shows the relative loss of the optimized TR schedules over the T1 range. Both optimized schedules outperformed the linear and logarithmic schedules, reaching a maximum of 4 dB gain across the range of T1 values. The two optimized schedules had very similar losses, indicating that the main T1 mode dominates the optimization.

The sampling patterns used in the acquisition are shown in Figure 6.7. The ETL was set to 132 to match the default 3D FSE T2 brain protocol (ETL = 130) with the addition




Figure 6.6: Relative loss of the TR schedules over T1 values ranging from 469 ms to 2732 ms (lower is better). The optimized schedules (solid blue line and dashed red line with circles) have uniformly lower loss compared to the linearly and logarithmically spaced TR schedules (dotted yellow line and dashed purple line with plus signs, respectively).

of two initial skipped echoes for ESPIRiT calibration (note the conventional acquisition skips the first four to eight initial echoes depending on the target contrast). The number of samples in each TR segment was set to match the optimized TR schedule from the first row of Table 6.3.

Figure 6.8 shows SE, SR-SE, and IR-SE simulations of 16 out of the 10240 samples drawn from the $(T1, T2)$ distribution, as well as their synthetic contrast fits based on the pre-determined linear combinations. The first two simulations followed (6.8) and the third followed (6.9). Below the simulations are the normalized root mean-squared errors (NRMSE) across all 10240 simulations for each time point. The SE synthetic contrast showed close agreement, with error initially decreasing until $TE = 28$ ms and increasing thereafter, reaching a maximal error of 0.035 at $TE = 100$ ms. The SR-SE synthetic contrast showed the reverse trend: the highest error was at the lowest TR value of 400 ms; the error decreased until reaching a minimum at $TR = 1221$ and increased thereafter. The IR-SE synthetic contrast had significantly higher error, reaching a maximal error of 0.176 at $TI = 631$ ms. Despite the high error, the fits were qualitatively similar to the inversion recovery behavior at short echo times.

Figure 6.9 shows T1-T2 Shuffling images from the second volunteer after reconstructing and linearly combining the relaxation coefficient images. The left column shows sagittal reformats at a synthetic TR of 7000 ms and a synthetic TE increasing from 5 ms to 200 ms. The middle column shows coronal reformats of saturation recovery at a synthetic TE of 10 ms and synthetic TR increasing from 200 ms to 700 ms. The right column shows axial reformats of inversion recovery with a synthetic TE of 10 ms, synthetic TR of 7000




Figure 6.7: Variable-density Poisson disc sampling patterns are generated for each unique TE and TR, resulting in a grid of under-sampled k-space. The number of samples in the i^{th} TR segment is equal to N_i calculated from (6.13). In general, shorter TRs are allocated more samples to increase scan efficiency and SNR. The first two TEs (not shown) follow a center-out ordering for ESPIRiT calibration.

ms, and synthetic TI increasing from 50 ms to 7000 ms. Qualitatively, the images follow the expected relaxation behavior.

Figure 6.10 compares conventional T1 weighted and T2 weighted 3D FSE images obtained from two separate CUBE acquisitions and the equivalent synthetic contrasts images obtained with T1-T2 Shuffling. Both the T1-T2 Shuffling scan time and the combined CUBE-T1 and CUBE-T2 scans were 7 minutes. However, the T1-T2 Shuffling reconstruction can produce multiple synthetic image contrasts to mimic different combinations of scan parameters. There is also a clear increase in image sharpness as a result of modeling the long echo train. Nonetheless, there is less contrast between gray and white matter in the T1-T2 Shuffling reconstruction, likely due to the limited dynamic range in a saturation-recovery based technique.




Figure 6.8: Simulations (top) of SE (left), SR-SE (middle), and IR-SE (right) for 16 combinations of T1 and T2 (solid lines) plotted against the linear combination contrast synthesis (dashed circled lines). Normalized root mean squared error plots across the 10240 simulated signal evolutions for each time point are shown in the second row.

6.5 Discussion

Standard clinical practice demands images at target contrasts and orientations. A single-sequence, rapid scan that satisfies this clinical need could potentially accelerate a number of MRI protocols. Multi-parametric mapping [149, 162, 169–171] combined with Synthetic MR [96, 150, 172, 173] provides one avenue toward this goal. This indirect approach – first solving for the intrinsic physical parameters, and then creating a synthetic image contrast – has been demonstrated in a number of settings and already seen clinical use [150]. However, parametric mapping methods are prone to error due to partial voluming or other unaccounted factors. This is often the case with synthetic FLAIR contrast, in which the CSF is surrounded by a hyperintense signal artifact [154, 155]. Additionally, the indirect methods are typically implemented in 2D slice-interleaved imaging to reach a high scan efficiency, which necessitates a separate scan for each desired imaging plane.




Figure 6.9: T1-T2 Shuffling synthetic contrast images at several contrasts and orientations. Left, top to bottom: sagittal reformats at a synthetic TR of 7000 ms and synthetic TE increasing from 5 ms to 200 ms show the T2 relaxation curve. Middle, top to bottom: coronal reformats at a synthetic TE of 10 ms and synthetic TR increasing from 200 ms to 7000 ms show the T1 saturation recovery curve. Right, top to bottom: axial reformats at a synthetic TE of 10 ms, synthetic TR of 7000 ms, and synthetic TI increasing from 50 ms to 7000 ms show the T1 inversion recovery curve.

In contrast to explicit parametric methods, T1-T2 Shuffling computes the synthetic contrast images directly from the reconstructed data. This is accomplished by linearly combining the relaxation coefficient images. Implicit in this processing is the assumption that the synthetic contrast images live in the same subspace as the data. In effect, the synthetic contrast images represent interpolated points within the linear subspace. In practice, this is only an approximation for a number of reasons. Firstly, the nonlinear dynamics that govern the signal equation are difficult to fully express with linear filtering only. Secondly, the linear combination is often operating in an extrapolation regime, e.g. when generating synthetic PD and T2-FLAIR contrasts from measurements that did not include inversions or short echo times. Third, the linear combination is sensitive to un-modeled system imperfections such as noise, eddy current effects, and B1 inhomogeneity. In principle, some of these effects could be accounted for when pre-computing the linear combination weights. A nonlinear combination that directly maps the reconstructed images to the desired target synthetic contrast may provide the fitting accuracy needed to adequately address these shortcomings. This could be accomplished through kernel regression [169], classification [174], or deep learning [155] methods.

T1-T2 Shuffling inherits many of the benefits seen with T2 Shuffling [117]. This includes reduced image blurring, as the echo train time progression is accounted for in the reconstruction. Although 3D FSE blurring is not a significant issue in brain imaging, it has been a major limitation for musculoskeletal imaging [40]. However, many of these protocols require the inclusion of both fat-suppressed and nonfat-suppressed imaging [118]. To accommodate this setting, it may be necessary to explore Dixon-based water-fat separation methods within the T1-T2 Shuffling acquisition and reconstruction [162].

There is a large degree of flexibility in choosing the T1-T2 Shuffling scan parameters. This work used saturation recovery T1 mapping to motivate the TR selection and played each TR sequentially, from largest to smallest. More complex variations could be explored, e.g. interleaved TRs, as they may impact the signal evolution and exhibit different properties with respect to motion. Additionally, every echo train used the same ETL and refocusing flip angles, including the fast recovery tip-up pulse. To increase T1 sensitivity, it may be beneficial to use shorter ETLs for the shorter TRs with tip-down pulses. In effect, this would provide an inversion-like prep pulse that is T2-dependent. Scan efficiency could also be increased by moving to non-Cartesian trajectories such as radial stack-of-stars, which have been shown to be effective in multi-contrast imaging settings [175].

The T1-T2 Shuffling reconstruction collapsed the relaxation and recovery dimensions through tensor unfolding. This provided a convenient reuse of the T2 Shuffling reconstruction, which also uses a single explicit subspace constraint with locally lower rank regularization. However, since the two relaxation dimensions are governed by different intrinsic physical parameters, it should be possible to further constrain the relaxation signals, e.g. with full tensor decompositions [164]. This has been shown to reduce the sample complexity in the context of free-breathing multi-parametric cardiac imaging [149].

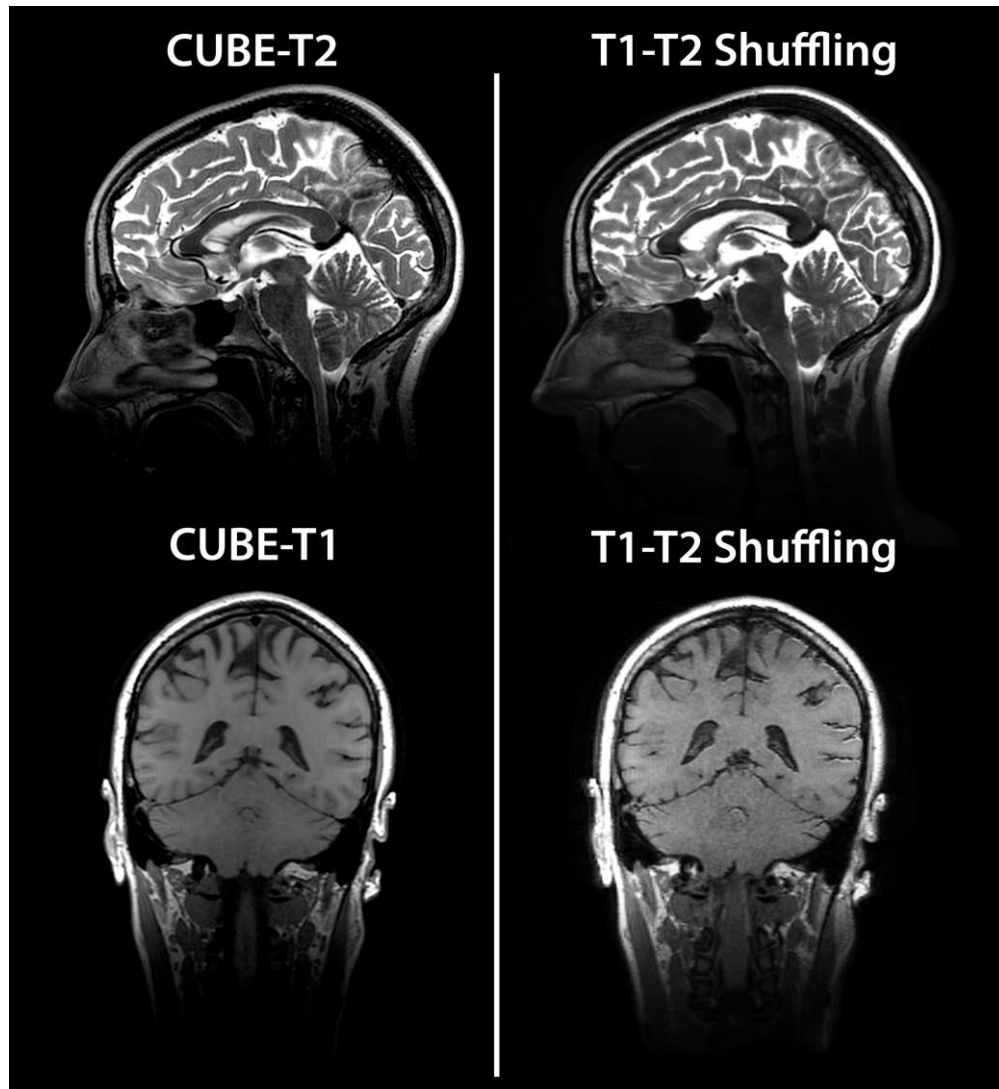


Figure 6.10: Comparison of conventional T1 weighted and T2 weighted images obtained from two separate CUBE acquisitions with the synthetic contrast image using T1-T2 Shuffling (single acquisition).

6.6 Conclusion

In conclusion, this chapter presented and developed a novel extension to T2 Shuffling that includes variable TRs during the acquisition. The method, called T1-T2 Shuffling, can flexibly reconstruct multiple PD, T1 weighted, and T2 weighted synthetic contrast images. Each reconstructed image is a 3D volume that can be reformatted into different orientation. As the method is based on 3D FSE, it can be extended to include various magnetization preparation sequences and modules, including inversion recoveries, water-fat separation, reduced field-of-view imaging, fat suppression, and partial Fourier acquisi-

tion. This work explored the qualitative properties of T1-T2 Shuffling as it compared to 3D FSE. More rigorous analysis including diagnostic comparison through clinical validation is needed.

Chapter 7

Conclusions and Future Work

7.1 Summary

This dissertation presented new accelerated MRI methods targeted toward real clinical adoption. Specifically, T2 Shuffling was presented as an extension to volumetric fast spin-echo (3D FSE) that models the temporal aspect of the MRI acquisition. The effects of accounting for the dynamics were two-fold. First, blurring commonly seen in FSE due to T2 decay was significantly reduced. Second, multiple 3D images were reconstructed from a single scan, corresponding to different contrasts along the signal relaxation curve. The approach leveraged advances in sparse and low rank modeling to reduce the scan time necessary for a well-posed reconstruction. Using incoherent sampling schemes, a low-dimensional subspace constraint, and locally low rank regularization, T2 Shuffling is able to reconstruct a time series of virtual echo time images from a single scan that is only slightly longer than a conventional 3D FSE scan.

To put the pieces together in a clinical environment, a fast and distributed reconstruction was designed and fully integrated into Lucile Packard Children’s Hospital (LPCH), achieving end-to-end reconstruction times of about 90 seconds. The optimized reconstruction took advantage of massive parallelism and multiple high-performance computers to achieve the target latency goal. The reconstruction tools and framework were implemented in the Berkeley Advanced Reconstruction Toolbox (BART) [19, 88], and were designed in such a way to enable re-use in other computationally expensive parallel imaging and compressed sensing-based methods.

The T2 Shuffling sequence was evaluated clinically for knee imaging through rigorous comparison to conventional protocols. The results of the comparisons indicated that missing clinically relevant pathology using only T2 Shuffling was unlikely. Altogether, the confidence in the method’s diagnostic benefits has led to its integration into routine clinical practice at LPCH. T2 Shuffling has replaced several of the 2D fat-suppressed sequences in the lower extremity knee protocol, and continues to be evaluated daily in the ankle, shoulder, wrist, elbow, and prostate [176]. Figure 7.1 highlights examples of these

other anatomy scanned with T2 Shuffling. Based on the initial clinical experience using T2 Shuffling, a targeted, rapid knee MRI exam was developed and offered to pediatric patients at one third the cost of the conventional exam. Patients with indications of knee pain can be referred to an MRI and scanned on the same day as their initial hospital visit, eliminating the multi-week turnaround typically seen with MRI scheduling. At the time of writing, over 2,400 T2 Shuffling patient scans have been performed.



Figure 7.1: T2 Shuffling scans of different anatomy on pediatric and adult patients. The top row shows a proton-density reconstruction and the bottom row shows a T2-weighted reconstruction. From left to right: brain, shoulder, wrist, ankle, and prostate scans.

The approach to modeling signal dynamics flexibly extends to other dimensions. Building off of T2 Shuffling, a method called T1-T2 Shuffling was presented that additionally captures T1 dynamics by sampling the saturation recovery curve at multiple repetition times (TRs). The relaxation subspace is extended to model a 2D distribution that includes T1 recovery. The selection of the TR durations was approached through the lens of decision theory applied to T1 mapping, and scan efficiency was increased by allocating more scan time to the short TR durations. By incorporating T1 sensitivity into the acquisition, the reconstruction affords synthetic contrast images corresponding to different target T1 weighted and T2 weighted scans. The method was demonstrated in vivo for brain imaging and shows promise toward creating a fast, single-sequence MRI exam.

7.2 Suggestions for Future Work

A number of innovations were presented in this dissertation; nevertheless, *work remains*. The topics in this dissertation provide exciting avenues for future directions. This section summarizes some of these ideas and speculates on their use.

7.2.1 Improved Low-Dimensional Modeling

The techniques used in T2 Shuffling and its extensions heavily use low-rank representations of spatio-temporal data. The primary approach taken was to represent the data as a low-dimensional subspace with the singular value decomposition (SVD). Although the SVD provides a powerful low-dimensional representation, it is ultimately limited in expressibility. Since the dynamics are generated from physical parameters through a non-linear process, a linear subspace representation often necessitates a large number of coefficients to achieve an adequate approximation. A promising alternative to the linear subspace model that may better represent the signal relaxation dynamics is a low-dimensional manifold. Several works explore this notion and learn the manifold representation from simulation or from the data itself [177–179], showing improved image quality over SVD-based subspace methods. The inclusion of a low-dimensional manifold in the reconstruction, e.g. through local linear embeddings [180], would maintain many of the features of the T2 Shuffling reconstruction. Thus, it could still leverage the methods presented for fast and distributed reconstructions.

When multiple parameter directions are imaged, the low-dimensional structure is well-described by a tensor [149]. In T1-T2 Shuffling, the tensor data were unfolded along the first mode by concatenating the signal relaxation curves and treated thereafter as a low rank matrix. Although this provides a convenient parallel to T2 Shuffling, it masks additional low-dimensional structure. The inclusion of a full tensor model in the reconstruction, e.g. through the Tucker decomposition [163], could provide additional avenues for acceleration and image quality.

7.2.2 Multi-Parameter Shuffling

In T2 Shuffling, the sequence is identical for each echo train. T1-T2 Shuffling added the capability to vary the TR duration for each echo train by adding dead time after the last echo. Nonetheless, the remaining sequence parameters were left unchanged. Motivated by the success of MR fingerprinting [63], additional sensitivity to relaxation parameters could be accomplished by varying the flip angles and echo train lengths during the scan. In particular, one direction worth exploring is the use of different fast recovery flip angles [160]. For example, a $+90^\circ$ fast recovery pulse could be used to tip the remaining transverse magnetization into the negative longitudinal direction, effectively adding T2-dependent inversion pulses to the spin dynamics. This concept is already used in 3D T1 weighted imaging based on FSE. In theory, each echo train could contain a unique TR, flip angle train, fast recovery flip angle, echo spacing, and magnetization preparation module [181]. The reconstruction could then enforce the low-dimensional manifold structure implicit in the data, as it comes from a small number of intrinsic physical parameters.

The inclusion of multiple parameters could extend beyond T1, and T2. Bipolar gradient readouts could be included to support water-fat separation [162]. Longer echo spacings could be used to sensitize the sequence to T2* relaxation and diffusion. Phase cycling

techniques could be incorporated to correct for phase errors [182, 183].

7.2.3 Parameter Mapping

This dissertation focused on generating clinical contrast in short exam times. This was accomplished by designing sequences that capture information about the relaxation dynamics. One direction worth exploring is extending the processing methods to quantitative parameter mapping. This could be accomplished through fitting each voxel in the reconstructed images to a physical model, either through non-linear least squares [51, 54] dictionary-based methods [63, 64], or machine learning [169]. The fits could even be performed directly in the subspace domain. Currently, quantitative imaging is not as common in a clinical setting, though the landscape is rapidly changing.

7.2.4 Optimal Spatio-Temporal Sampling

The sampling schemes used in T2 Shuffling and its extensions relied on ideas from parallel imaging and compressed sensing. In the end, each time point was sampled with a variable-density Poisson disc sampling pattern. Although this is an established approach for incoherent sampling, it may not be the best choice given the large degree of redundancy in the spatio-temporal data. Alternative sampling schemes that incorporate complementary information in each temporal frame [184] and the subspace basis functions themselves [185] could provide better incoherence properties. It may also be possible to learn an optimal sampling trajectory, either through decision theory [186] or through deep learning [187]. Additionally, non-Cartesian [175] and other encoding approaches may provide better incoherence properties [188].

7.2.5 Learned Reconstructions

Recently, reconstruction methods have become increasingly data-driven. Using a large corpus of available training data, end-to-end reconstructions are learned with and without physical modeling [189, 190]. The approaches learn transforms and non-linear activation functions that are tailored to the input data. As a result, new low-dimensional representations are automatically learned and used in the reconstruction. Data-driven approaches to MR Shuffling have been explored, with promising initial results [174]. An unrolled optimization method could be used to further learn the reconstruction steps. One difficulty to applying these methods is the availability of ground-truth data. Although there are thousands of clinical T2 Shuffling data sets, none are fully sampled. Future directions should investigate reference-less and semi-supervised approaches to directly learn from the under-sampled data [191].

Appendix A

Low Rank Regularization

Compressed sensing enables sparse recovery of vector-valued objects under incoherent linear measurements through ℓ_1 -norm minimization. An analogue in the matrix case is recovery of low-rank matrices through nuclear norm minimization [12]. Matrix-based objects appear in applications such as collaborative filtering and system identification [8]. In MRI, it is natural to represent dynamic imaging as a spatio-temporal matrix, where the rows represent voxels and the columns represent temporal dynamics [90]. The temporal dynamics can be a result of signal relaxation [57], contrast enhancement [18], cardiac motion [91], respiratory motion [16], or a combination [149].

Global Low Rank

Suppose low rank regularization is imposed on the measurement matrix $\mathbf{Y} = \mathbf{X} + \mathbf{W}$, where \mathbf{X} is a low-rank matrix and \mathbf{W} is a measurement noise matrix. Then a low-rank matrix \mathbf{Z} that best approximates \mathbf{X} can be found solving the unconstrained convex optimization problem

$$\operatorname{argmin}_{\mathbf{Z}} \frac{1}{2} \|\mathbf{Y} - \mathbf{Z}\|_F^2 + \lambda \|\mathbf{Z}\|_*, \quad (\text{A.1})$$

where $\|\mathbf{Z}\|_*$ is the nuclear norm, defined as the sum of the singular values of \mathbf{Z} . The solution to (A.1) is given by the singular value soft threshold (SVT) function [192, 193]. In words, the SVT function applies soft thresholding to the singular values of \mathbf{Y} . The soft threshold function is defined as

$$\text{SoftThresh}(\mathbf{y}, \lambda) := (|\mathbf{y}| - \lambda)_+ \frac{\mathbf{y}}{|\mathbf{y}|}, \quad (\text{A.2})$$

where $(\mathbf{z})_+$ sets the negative elements of \mathbf{z} to zero and leaves the positive elements unchanged.

Denote the singular value decomposition (SVD) of \mathbf{Y} as

$$\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*, \quad (\text{A.3})$$

where $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_K)$ are the ordered singular values. Define $\hat{\Sigma}$ as the diagonal matrix of the singular values after soft-thresholding, i.e.

$$\hat{\Sigma} = \text{diag} \left(\text{SoftThresh} \left([\sigma_1 \ \dots \ \sigma_K]^\top, \lambda \right) \right). \quad (\text{A.4})$$

Then the solution to (A.1) is given by

$$\mathbf{Z} = \text{SVT}(\mathbf{Y}, \lambda) = \mathbf{U} \hat{\Sigma} \mathbf{V}^*. \quad (\text{A.5})$$

Local Low Rank

Instead of global low rank, it is possible to impose local low rank (LLR) constraints on a matrix. This implies that sub-matrices of \mathbf{Y} are separately low rank. This is suitable in dynamic MRI because local regions of an image often exhibit similar dynamics [69, 71]. The LLR optimization problem is given by

$$\underset{\mathbf{Z}}{\text{argmin}} \quad \frac{1}{2} \|\mathbf{Y} - \mathbf{Z}\|_F^2 + \lambda \sum_{\mathbf{r}} \|R_{\mathbf{r}}(\mathbf{Z})\|_*. \quad (\text{A.6})$$

The operation $R_{\mathbf{r}}(\mathbf{Z})$ extracts a sub-matrix of \mathbf{Z} at position \mathbf{r} . When the sub-matrices are non-overlapping, the solution to this problem is given by Algorithm A.1, called Local SVT (LSVT). Because the operation is not translation invariant, blocking artifacts will occur at the boundaries of sub-matrices. One way to remove these artifacts is to use overlapping blocks of sub-matrices, equivalent to applying LSVT to shifts of the matrix [82]. In this case Step (3) will include an average over all shifts of position \mathbf{r} . LLR regularization can be

Algorithm A.1 Procedure for applying LSVT.

Inputs: $\mathbf{Y} \in \mathbb{C}^{N \times T}$ —data, λ —regularization parameter, b —block size

Outputs: $\mathbf{Z} \in \mathbb{C}^{N \times T}$ —LLR output

- 1: Decimate the rows of \mathbf{Y} to form sub-matrices of size $b \times T$ centered at each row \mathbf{r} . Call each sub-matrix $\mathbf{Y}_{\mathbf{r}} = R_{\mathbf{r}}(\mathbf{Y}) \in \mathbb{C}^{b \times T}$. By construction, $\mathbf{Y} = \sum_{\mathbf{r}} \mathbf{Y}_{\mathbf{r}}$.
 - 2: Apply SVT to each sub-matrix: $\mathbf{Z}_{\mathbf{r}} = \text{SVT}(\mathbf{Y}_{\mathbf{r}}, \lambda)$.
 - 3: Sum the result: $\mathbf{Z} = \sum_{\mathbf{r}} \mathbf{Z}_{\mathbf{r}}$.
-

flexibly applied to 2D and 3D imaging applications by extracting $b \times b$ spatial blocks from each temporal frame and concatenating them as a single column. as illustrated in Figure A.1. As previously mentioned, LSVT applied to non-overlapping blocks will lead to blocking artifacts. When using LLR in an iterative inverse problem, it can be prohibitively expensive to apply LSVT to all image translations. An alternative to applying translation-invariant shifting is to apply a random shift to the image in each iteration [83]. This has the effect of averaging out the blocking artifacts over the iterations while only applying LSVT once per iteration.




Figure A.1: Visualization of the block-wise matrix operator for LLR with $T = 5$ temporal frames. Blocks from each temporal frame are reshaped into columns to form small matrices. The transform components are the singular vectors of the matrices.

Appendix B

Subspace Noise Propagation

In this appendix it is shown that the noise variance of a fully sampled, subspace-constrained linear reconstruction is equal to $\sigma^2 K$, where K is the dimension of the subspace [84]. Consider a noisy time series signal evolution at voxel \mathbf{r} , as described by (3.4) and (3.5):

$$x_i(\mathbf{r}) = \rho(\mathbf{r}) f_{iT_s}(\mathbf{r}), \quad (\text{B.1})$$

$$y_i(\mathbf{r}) = x_i(\mathbf{r}) + w_i(\mathbf{r}), \quad i = 1, \dots, T, \quad (\text{B.2})$$

where

$$\mathbf{w}(\mathbf{r}) = [w_1(\mathbf{r}) \ \cdots \ w_T(\mathbf{r})]^\top \sim \mathcal{N}_c(0, \sigma^2 \mathbf{I}) \in \mathbb{C}^T. \quad (\text{B.3})$$

Each noise observation $w_i(\mathbf{r})$ is an independent, zero-mean complex Gaussian random variable with variance σ^2 . Thus,

$$\mathbf{y}(\mathbf{r}) \sim \mathcal{N}_c(\mathbf{x}(\mathbf{r}), \sigma^2 \mathbf{I}). \quad (\text{B.4})$$

The measurements $\mathbf{y}(\mathbf{r})$ can be approximated by projecting them onto the temporal subspace spanned by the semi-unitary matrix $\Phi_K \in \mathbb{C}^{T \times K}$. Let $\alpha(\mathbf{r}) = \Phi_K^* \mathbf{y}(\mathbf{r})$. The random vector $\alpha(\mathbf{r})$ has covariance matrix given by

$$\Sigma_\alpha = \Phi_K^* \Sigma_y \Phi_K \quad (\text{B.5})$$

$$= \sigma^2 \Phi_K^* \Phi_K \quad (\text{B.6})$$

$$= \sigma^2 \mathbf{I} \in \mathbb{C}^{K \times K}, \quad (\text{B.7})$$

i.e. the temporal coefficients remain independent and Gaussian, and the total variance is reduced to $\|\Sigma_\alpha\|_F^2 = \sigma^2 K$. After projecting back to the time series, the estimate is $\hat{\mathbf{y}}(\mathbf{r}) = \Phi_K \alpha(\mathbf{r})$. Since the matrix Φ_K has orthonormal columns, the total variance in $\hat{\mathbf{y}}(\mathbf{r})$ is unchanged; however, the noise at each time point is no longer independent nor identically distributed. The covariance matrix of $\hat{\mathbf{y}}(\mathbf{r})$ is $\sigma^2 \Phi_K \Phi_K^*$. Figure B.1 shows the noise

variance at each echo time after back-projecting, i.e. $\text{diag}(\Phi_K \Phi_K^*)$. In general, earlier time points will have higher noise variance after the projection since the vectors given by Φ_K represent signal decay. The expected error due to noise is lowest at the central echo times corresponding to T2 weighted images.




Figure B.1: Noise variance in the time series after back-projecting through the basis. The noise is Gaussian but no longer white. The initial time points have a higher noise variance consistent with the behavior of the signal decay.

Bias vs. Noise:

The distribution of the error $\mathbf{z}(\mathbf{r}) := \hat{\mathbf{y}}(\mathbf{r}) - \mathbf{x}(\mathbf{r})$ is given by

$$\mathbf{z}(\mathbf{r}) \sim \mathcal{N}_c((\Phi_K \Phi_K^* - \mathbf{I}) \mathbf{x}(\mathbf{r}), \sigma^2 \Phi_K \Phi_K^*), \quad (\text{B.8})$$

From this, the mean-squared error is

$$\mathbb{E} \|\mathbf{z}(\mathbf{r})\|_2^2 = \mathbf{x}(\mathbf{r})^* (\mathbf{I} - \Phi_K \Phi_K^*) \mathbf{x}(\mathbf{r}) + \sigma^2 K \quad (\text{B.9})$$

$$= \|\mathbf{x}(\mathbf{r})\|_2^2 - \|\Phi_K^* \mathbf{x}(\mathbf{r})\|_2^2 + \sigma^2 K. \quad (\text{B.10})$$

The first two terms represent the error due to subspace modeling and the third term represents error due to noise. There is an inherent tradeoff between model error and noise with subspace size. As K increases, the model error decreases and the noise variance increases. The SNR of the measured signal is T_2 -dependent, and is given by

$$\text{SNR}(\mathbf{r}) = 10 \log_{10} \frac{\|\mathbf{x}(\mathbf{r})\|_2^2}{\sigma^2 K} \quad [\text{dB}]. \quad (\text{B.11})$$

The root mean-squared error (RMSE) is plotted in Figure 3.8 for $\sigma = 0.07$, which corresponds to an SNR of about 15 dB for $T_2 = 14$ ms.

Appendix C

T2 Shuffling Transform Point Spread Function

The transform point spread function (TPSF) can be used to empirically assess the interference caused by a single transform coefficient [11]. For T2 Shuffling, locally low rank (LLR) matrices are used as a low-dimensional transform [69]. The LLR transform operates on temporal coefficient image blocks, as described in Appendix A and illustrated in Figure A.1. Since there are K temporal coefficient images, the maximum number of coefficients (rank) for each matrix is equal to K .

One can analyze the effect of a single coefficient in the LLR domain by constructing a system input with a uniform patch in the center of a single temporal coefficient image, and zero elsewhere. The TPSF is then computed by passing the system input through the forward operator and its adjoint:

$$\alpha_{\text{TPSF}} = \Phi_K^* \mathbf{S}^* \mathbf{F}^* \mathbf{P} \mathbf{F} \mathbf{S} \Phi_K \alpha_{\text{input}}. \quad (\text{C.1})$$

The procedure is repeated K times; each repetition places the uniform patch in the subsequent temporal coefficient image. The final TPSF is a K -by- K grid of temporal coefficient images. The TPSF depends on the spatio-temporal sampling scheme used. In Figure 3.5, both the center-out and the randomly shuffled sampling patterns were constructed such that every phase encode was sampled exactly once (i.e. no parallel imaging acceleration or repeated phase encode re-sampling). The center-out TPSF shows coherent interference and blurring, and the coefficients are coupled. This demonstrates the difficulty of deconvolution. The randomly shuffled TPSF shows incoherent interference and maintains a sharp central patch. The coefficients are decoupled, and the interference manifests as benign noise-like artifacts.

Appendix D

Sampling Pattern and Echo Train Formation

In T2 Shuffling, phase encodes are randomly shuffled and re-acquired throughout the echo trains. To leverage compressed sensing-based approaches, each echo time is sampled with a variable-density Poisson disc sampling pattern, as shown in Figure 3.9. The first two echoes follow a conventional center-out phase encode ordering so that they can be used for coil sensitivity estimation. The number of samples in each sampling pattern is equal to the number of echo trains, N_{trains} , and is

$$R = \sqrt{\frac{\pi}{4} \frac{N_y \times N_z}{N_{\text{trains}} \times \tau}}, \quad (\text{D.1})$$

where (N_y, N_z) are the k-space phase encode dimensions and $\tau \approx 1.1$ is a “fudge factor” so that the number of actual points is greater than or equal to N_{trains} . The $\pi/4$ factor is included because the corners of k-space are not acquired [40]. After generating the masks each with acceleration factor R , points are randomly pruned until each mask contains exactly N_{trains} samples. The total number of acquired samples is equal to $T \cdot N_{\text{trains}}$, where T is the echo train length (ETL). Each mask can incorporate additional sampling strategies, such as partial Fourier acquisition [121], fully sampled center, etc.

Echo trains are formed by choosing one sample from each sampling pattern. Though it is possible to assign phase encodes to echo trains arbitrarily, this could result in eddy current effects due to the large, random jumps in k-space. To reduce these effects, phase encodes are grouped locally to form an echo train. The general problem of assigning phase encodes based on local distances takes the form of the linear bottleneck assignment problem and can be solved with the Hungarian Algorithm [194]. Instead, a simpler greedy approach based on k-nearest neighbors is used, shown in Algorithm D.1. Starting with one point from the mask of the first TE that is fully sampled at the center, the closest neighboring point is chosen from the mask of the next TE. This process is repeated for each sampling pattern and these T points comprise the first echo train. The process repeats until all echo trains are formed.

Algorithm D.1 Randomly shuffled sampling pattern and echo train formation.

Inputs: $N_y, N_z, N_{\text{trains}}, T, \tau$

Outputs: $\text{trains} [N_{\text{trains}}, T]$ – array of phase encodes to acquire

```

1:  $\text{trains} \leftarrow \text{zeros}(N_{\text{trains}}, T)$ 
2:  $\text{masks} \leftarrow \text{zeros}(N_y, N_z, T)$ 
3:  $R \leftarrow \text{Eq. (D.1)}$ 
4: for  $i = 1 : T$  do
5:    $\text{mask} \leftarrow \text{genVDPoissonMask}(N_y, N_z, R, \tau)$ 
6:    $\text{masks}[:, :, i] \leftarrow \text{mask.randprune}(N_{\text{trains}})$ 
7: end for
8: for  $n = 1 : N_{\text{trains}}$  do
9:    $\text{trains}[n, 1] \leftarrow \text{RandomPoint}(\text{masks}[:, :, 1])$ 
10:   $\text{RemovePointFromMask}(\text{trains}[n, 1], \text{masks}[:, :, 1])$ 
11:  for  $j = 2 : T$  do
12:     $\text{trains}[n, j] \leftarrow \text{FindClosestPoint}(\text{trains}[n, j-1], \text{masks}[:, :, j])$ 
13:     $\text{RemovePointFromMask}(\text{trains}[n, j], \text{masks}[:, :, j])$ 
14:  end for
15: end for

```

Bibliography

- [1] D. G. Nishimura. **Principles of magnetic resonance imaging**. Stanford University, 1996.
- [2] J. Hennig, A. Nauerth, and H. Friedburg, **RARE imaging: a fast imaging method for clinical MR**. *Magn Reson Med* 1986, 3(6):823–833. DOI: 10.1002/mrm.1910030602.
- [3] J. P. Mugler 3rd, **Optimized three-dimensional fast-spin-echo MRI**. *J Magn Reson Imaging* 2014, 39(4):745–767. DOI: 10.1002/jmri.24542.
- [4] R. F. Busse, H. Hariharan, A. Vu, et al., **Fast spin echo sequences with very long echo trains: design of variable refocusing flip angle schedules and generation of clinical T2 contrast**. *Magn Reson Med* 2006, 55(5):1030–1037. DOI: 10.1002/mrm.20863.
- [5] D. K. Sodickson and W. J. Manning, **Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays**. *Magn Reson Med* 1997, 38(4):591–603.
- [6] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, et al., **SENSE: sensitivity encoding for fast MRI**. *Magn Reson Med* 1999, 42(5):952–962.
- [7] M. A. Griswold, P. M. Jakob, R. M. Heidemann, et al., **Generalized autocalibrating partially parallel acquisitions (GRAPPA)**. *Magn Reson Med* 2002, 47(6):1202–1210. DOI: 10.1002/mrm.10171.
- [8] V. Chandrasekaran, B. Recht, P. A. Parrilo, et al., **The Convex Geometry of Linear Inverse Problems**. *Found Comput Math* 2012, 12(6):805–849. DOI: 10.1007/s10208-012-9135-7.
- [9] D. L. Donoho, **Compressed sensing**. *IEEE Trans Inf Theory* 2006, 52(4):1289–1306. DOI: 10.1109/TIT.2006.871582.

- [10] E. J. Candes, J. Romberg, and T. Tao, **Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information.** *IEEE Trans Inf Theory* 2006, 52(2):489–509. DOI: 10.1109/TIT.2005.862083.
- [11] M. Lustig, D. Donoho, and J. M. Pauly, **Sparse MRI: The application of compressed sensing for rapid MR imaging.** *Magn Reson Med* 2007, 58(6):1182–1195. DOI: 10.1002/mrm.21391.
- [12] B. Recht, M. Fazel, and P. Parrilo, **Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization.** *SIAM Rev* 2010, 52(3):471–501. DOI: 10.1137/070697835.
- [13] Z. Liang, **Spatiotemporal imaging with partially separable functions.** *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Apr. 2007. p. 988–991. DOI: 10.1109/ISBI.2007.357020.
- [14] M. S. Hansen and P. Kellman, **Image reconstruction: an overview for clinicians.** *J Magn Reson Imaging* 2015, 41(3):573–585. DOI: 10.1002/jmri.24687.
- [15] J. Gouveia, P. A. Parrilo, and R. R. Thomas, **Lifts of Convex Sets and Cone Factorizations.** *Mathematics of OR* 2013, 38(2):248–264. DOI: 10.1287/moor.1120.0575.
- [16] W. Jiang, F. Ong, K. M. Johnson, et al., **Motion robust high resolution 3D free-breathing pulmonary MRI using dynamic 3D image self-navigator.** *Magn Reson Med* 2018, 79(6):2954–2967. DOI: 10.1002/mrm.26958.
- [17] L. Feng, R. Grimm, K. T. Block, et al., **Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI.** *Magn Reson Med* 2014, 72(3):707–717. DOI: 10.1002/mrm.24980.
- [18] T. Zhang, J. Y. Cheng, A. G. Potnick, et al., **Fast pediatric 3D free-breathing abdominal dynamic contrast enhanced MRI with high spatiotemporal resolution.** *J Magn Reson Imaging* 2015, 41(2):460–473. DOI: 10.1002/jmri.24551.
- [19] J. I. Tamir, F. Ong, J. Y. Cheng, et al. **Generalized Magnetic Resonance Imaging using The Berkeley Advanced Reconstruction Toolbox.** *ISMRM Scientific Workshop - Data Sampling and Image Reconstruction*, Sedona, 2016.
- [20] M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, **An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems.** *IEEE Trans Image Process* 2011, 20(3):681–695. DOI: 10.1109/TIP.2010.2076294.

- [21] A. Beck and M. Teboulle, **A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems**. *SIAM J Imaging Sci* 2009, 2(1):183–202. DOI: 10.1137/080716542.
- [22] S. Boyd, N. Parikh, E. Chu, et al., **Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers**. *Foundations and Trends® in Machine Learning* 2011, 3(1):1–122. DOI: 10.1561/22000000016.
- [23] A. Chambolle and T. Pock, **A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging**. *J Math Imaging Vis* 2011, 40(1):120–145. DOI: 10.1007/s10851-010-0251-1.
- [24] K. King, **HyperSense Enables Shorter Scan Times Without Compromising Image Quality**. *GE SIGNA Pulse of MR* 2016, 22(Autumn 2016):48–51.
- [25] C. Forman, J. Wetzl, C. Hayes, et al., **Compressed Sensing: a Paradigm Shift in MRI**. *MAGNETOM Flash* 2016, 6619.
- [26] L. Geerts-Ossevoort, E. de Weerd, A. Duijndam, et al., **Compressed SENSE: Speed done right. Every time**. *Philips FieldStrength Magazine* 2018, (January 2018):1–16.
- [27] M. Lustig, D. L. Donoho, J. M. Santos, et al., **Compressed Sensing MRI**. *IEEE Signal Process Mag* 2008, 25(2):72–82. DOI: 10.1109/MSP.2007.914728.
- [28] R. Tibshirani, **Regression Shrinkage and Selection Via the Lasso**. *J R Stat Soc Series B Stat Methodol* 1994, 58:267–288.
- [29] J. Hennig, M. Weigel, and K. Scheffler, **Calculation of flip angles for echo trains with predefined amplitudes with the extended phase graph (EPG)-algorithm: principles and applications to hyperecho and TRAPS sequences**. *Magn Reson Med* 2004, 51(1):68–80. DOI: 10.1002/mrm.10658.
- [30] M. Weigel, **Extended phase graphs: dephasing, RF pulses, and echoes - pure and simple**. *J Magn Reson Imaging* 2015, 41(2):266–295. DOI: 10.1002/jmri.24619.
- [31] J. V. Hajnal, B. De Coene, P. D. Lewis, et al., **High signal regions in normal white matter shown by heavily T2-weighted CSF nulled IR sequences**. *J Comput Assist Tomogr* 1992, 16(4):506–513.
- [32] E. L. Hahn, **Spin Echoes**. *Phys Rev* 1950, 80(4):580–594. DOI: 10.1103/PhysRev.80.580.

- [33] H. Y. Carr and E. M. Purcell, **Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments**. *Phys Rev* 1954, 94(3):630–638. DOI: 10.1103/PhysRev.94.630.
- [34] S. Meiboom and D. Gill, **Modified Spin-Echo Method for Measuring Nuclear Relaxation Times**. *Rev Sci Instrum* 1958, 29(8):688–691. DOI: 10.1063/1.1716296.
- [35] H. Lu, L. M. Nagae-Poetscher, X. Golay, et al., **Routine clinical brain MRI sequences for use at 3.0 Tesla**. *J Magn Reson Imaging* 2005, 22(1):13–22. DOI: 10.1002/jmri.20356.
- [36] G. M. Bydder and I. R. Young, **MR imaging: clinical use of the inversion recovery sequence**. *J Comput Assist Tomogr* 1985, 9(4):659–675.
- [37] G. E. Gold, R. F. Busse, C. Beehler, et al., **Isotropic MRI of the knee with 3D fast spin-echo extended echo-train acquisition (XETA): initial experience**. *AJR Am J Roentgenol* 2007, 188(5):1287–1293. DOI: 10.2214/AJR.06.1208.
- [38] R. Kijowski and G. E. Gold, **Routine 3D magnetic resonance imaging of joints**. *J Magn Reson Imaging* 2011, 33(4):758–771. DOI: 10.1002/jmri.22342.
- [39] U. Tariq, P. Lai, M. Lustig, et al. **Impact of Sampling Pattern on Combined Parallel Imaging Compressed Sensing Volumetric Knee MRI**. *Proc Intl Soc Mag Reson Med 21*, Salt Lake City, 2013. p. 1664.
- [40] R. F. Busse, A. C. S. Brau, A. Vu, et al., **Effects of refocusing flip angle modulation and view ordering in 3D fast spin echo**. *Magn Reson Med* 2008, 60(3):640–649. DOI: 10.1002/mrm.21680.
- [41] M. P. Recht, R. Otazo, L. Rybak, et al. **3D TSE Imaging using Sparse-SENSE Acceleration: Comparison with Conventional 2D TSE Imaging for Detection of Internal Derangement of the Knee**. *Proc Intl Soc Mag Reson Med 23*, Toronto, 2015. p. 1198.
- [42] S. Reid, K. F. King, D. J. Lomas, et al. **Imaging of the knee using 3D fast spin echo with compressed sensing**. *Proc Intl Soc Mag Reson Med 23*, Toronto, 2015. p. 1197.
- [43] X. Zhou, Z. P. Liang, G. P. Cofer, et al., **Reduction of ringing and blurring artifacts in fast spin-echo imaging**. *J Magn Reson Imaging* 1993, 3(5):803–807.
- [44] K. Oshio and M. Singh, **Correction of T2 distortion in multi-excitation RARE sequence**. *IEEE Trans Med Imaging* 1992, 11(1):123–128. DOI: 10.1109/42.126919.

- [45] T. Çukur, M. Lustig, E. U. Saritas, et al., **Signal compensation and compressed sensing for magnetization-prepared MR angiography**. *IEEE Trans Med Imaging* 2011, 30(5):1017–1027. DOI: 10.1109/TMI.2011.2116123.
- [46] H. Takahashi, K. Ogawa, K. Oshio, et al., **A proposal of removal method for T2 decay effects in RARE sequence**. *IEEE Trans Nucl Sci* 1995, 42(4):1343–1347. DOI: 10.1109/23.467860.
- [47] C. Santa-Marta, J. Lafuente, J. J. Vaquero, et al., **Resolution recovery in Turbo Spin Echo using segmented Half Fourier acquisition**. *Magn Reson Imaging* 2004, 22(3):369–378. DOI: 10.1016/j.jmri.2004.01.009.
- [48] K. T. Block, M. Uecker, and J. Frahm, **Model-based iterative reconstruction for radial fast spin-echo MRI**. *IEEE Trans Med Imaging* 2009, 28(11):1759–1769. DOI: 10.1109/TMI.2009.2023119.
- [49] T. J. Sumpf, M. Uecker, S. Boretius, et al., **Model-based nonlinear inverse reconstruction for T2 mapping using highly undersampled spin-echo MRI**. *J Magn Reson Imaging* 2011, 34(2):420–428. DOI: 10.1002/jmri.22634.
- [50] T. J. Sumpf, A. Petrovic, M. Uecker, et al., **Fast T2 mapping with improved accuracy using undersampled spin-echo MRI and model-based reconstructions with a generating function**. *IEEE Trans Med Imaging* 2014, 33(12):2213–2222. DOI: 10.1109/TMI.2014.2333370.
- [51] N. Ben-Eliezer, D. K. Sodickson, T. Shepherd, et al., **Accelerated and motion-robust in vivo T2 mapping from radially undersampled data using bloch-simulation-based iterative reconstruction**. *Magn Reson Med* 2016, 75(3):1346–1354. DOI: 10.1002/mrm.25558.
- [52] F. H. Petzschner, I. P. Ponce, M. Blaimer, et al., **Fast MR parameter mapping using k-t principal component analysis**. *Magn Reson Med* 2011, 66(3):706–716. DOI: 10.1002/mrm.22826.
- [53] C. Huang, C. G. Graff, E. W. Clarkson, et al., **T2 mapping from highly undersampled data by reconstruction of principal component coefficient maps using compressed sensing**. *Magn Reson Med* 2012, 67(5):1355–1366. DOI: 10.1002/mrm.23128.
- [54] C. Huang, A. Bilgin, T. Barr, et al., **T2 relaxometry with indirect echo compensation from highly undersampled data**. *Magn Reson Med* 2013, 70(4):1026–1037. DOI: 10.1002/mrm.24540.

- [55] B. Zhao, W. Lu, T. K. Hitchens, et al., **Accelerated MR parameter mapping with low-rank and sparsity constraints**. *Magn Reson Med* 2015, 74(2):489–498. DOI: 10.1002/mrm.25421.
- [56] M. B. Keerthivasan, A. Bilgin, D. R. Martin, et al. **Isotropic T2 Mapping using a 3D Radial FSE (or TSE) pulse sequence**. *Proc Intl Soc Mag Reson Med* 23, Toronto, 2015. p. 3237.
- [57] M. Doneva, P. Börnert, H. Eggers, et al., **Compressed sensing reconstruction for magnetic resonance parameter mapping**. *Magn Reson Med* 2010, 64(4):1114–1120. DOI: 10.1002/mrm.22483.
- [58] M. Hong, M. Lin, F. Liu, et al. **MR Image Reconstruction from under-sampled measurements using local and global sparse representations**. *Proc Intl Soc Mag Reson Med* 23, Toronto, 2015. p. 3422.
- [59] B. Zhao, J. P. Haldar, A. G. Christodoulou, et al., **Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints**. *IEEE Trans Med Imaging* 2012, 31(9):1809–1820. DOI: 10.1109/TMI.2012.2203921.
- [60] J. V. Velikina and A. A. Samsonov, **Reconstruction of dynamic image series from undersampled MRI data using data-driven model consistency condition (MOCCO)**. *Magn Reson Med* 2015, 74(5):1279–1290. DOI: 10.1002/mrm.25513.
- [61] S. Bhave, S. G. Lingala, C. P. Johnson, et al., **Accelerated whole-brain multi-parameter mapping using blind compressed sensing**. *Magn Reson Med* 2016, 75(3):1175–1186. DOI: 10.1002/mrm.25722.
- [62] S. J. Matzat, E. J. McWalter, F. Kogan, et al., **T2 Relaxation time quantitation differs between pulse sequences in articular cartilage**. *J Magn Reson Imaging* 2015, 42(1):105–113. DOI: 10.1002/jmri.24757.
- [63] D. Ma, V. Gulani, N. Seiberlich, et al., **Magnetic resonance fingerprinting**. *Nature* 2013, 495(7440):187–192. DOI: 10.1038/nature11971.
- [64] N. Ben-Eliezer, D. K. Sodickson, and K. T. Block, **Rapid and accurate T2 mapping from multi-spin-echo data using Bloch-simulation-based reconstruction**. *Magn Reson Med* 2015, 73(2):809–817. DOI: 10.1002/mrm.25156.
- [65] E. Y. Pierre, D. Ma, Y. Chen, et al., **Multiscale reconstruction for MR fingerprinting**. *Magn Reson Med* 2016, 75(6):2481–2492. DOI: 10.1002/mrm.25776.
- [66] B. Zhao, F. Lam, B. Bilgic, et al. **Maximum likelihood reconstruction for magnetic resonance fingerprinting**. *2015 IEEE 12th International Symposium on Biomed-*

- ical Imaging (ISBI)*, IEEE, Apr. 2015. p. 905–909. DOI: 10.1109/ISBI.2015.7164017.
- [67] J. Assländer, M. A. Cloos, F. Knoll, et al., **Low rank alternating direction method of multipliers reconstruction for MR fingerprinting**. *Magn Reson Med* 2018, 79(1):83–96. DOI: 10.1002/mrm.26639.
- [68] B. Zhao, K. Setsompop, E. Adalsteinsson, et al., **Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling**. *Magn Reson Med* 2018, 79(2):933–942. DOI: 10.1002/mrm.26701.
- [69] J. Trzasko, A. Manduca, and E. Borisch. **Local versus global low-rank promotion in dynamic MRI series reconstruction**. *Proc Int Symp Magn Reson Med*, Montreal, 2011. p. 4371.
- [70] J. D. Trzasko and A. Manduca. **Calibrationless parallel MRI using CLEAR**. *2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, Nov. 2011. p. 75–79. DOI: 10.1109/ACSSC.2011.6189958.
- [71] T. Zhang, J. M. Pauly, and I. R. Levesque, **Accelerating parameter mapping with a locally low rank constraint**. *Magn Reson Med* 2015, 73(2):655–661. DOI: 10.1002/mrm.25161.
- [72] G. Li, M. Zaitsec, M. Weigel, et al. **On the feasibility of hybrid acquisition in SPACE**. *Proc Intl Soc Mag Reson Med* 22, Milan, 2014. p. 1467.
- [73] S. Kecskemeti, A. Samsonov, S. A. Hurley, et al., **MPnRAGE: A technique to simultaneously acquire hundreds of differently contrasted MPRAGE images with applications to quantitative T1 mapping**. *Magn Reson Med* 2016, 75(3):1040–1053. DOI: 10.1002/mrm.25674.
- [74] T. Benkert, A. J. Bartsch, M. Blaimer, et al., **Generating multiple contrasts using single-shot radial T1 sensitive and insensitive steady-state imaging**. *Magn Reson Med* 2015, 73(6):2129–2141. DOI: 10.1002/mrm.25337.
- [75] J. P. Haldar, D. Hernando, and Z. Liang. **Super-resolution reconstruction of MR image sequences with contrast modeling**. *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, June 2009. p. 266–269. DOI: 10.1109/ISBI.2009.5193035.
- [76] N. N. Lukzen and A. A. Savelov, **Analytical derivation of multiple spin echo amplitudes with arbitrary refocusing angle**. *J Magn Reson* 2007, 185(1):71–76. DOI: 10.1016/j.jmr.2006.11.010.

- [77] J. I. Tamir, P. Lai, M. Uecker, et al. **Reduced blurring in 3D fast spin echo through joint temporal esprit reconstruction.** *Proc Intl Soc Mag Reson Med* 22, 22, Milan, 2014. p. 0616.
- [78] T. Zhang, J. M. Pauly, S. S. Vasanawala, et al., **Coil compression for accelerated imaging with Cartesian sampling.** *Magn Reson Med* 2013, 69(2):571–582. DOI: 10 . 1002/mrm.24267.
- [79] K. Kim, J. C. Ye, W. Worstell, et al., **Sparse-view spectral CT reconstruction using spectral patch-based low-rank penalty.** *IEEE Trans Med Imaging* 2015, 34(3):748–760. DOI: 10 . 1109/TMI . 2014 . 2380993.
- [80] J. V. Velikina, A. L. Alexander, and A. Samsonov, **Accelerating MR parameter mapping using sparsity-promoting regularization in parametric dimension.** *Magn Reson Med* 2013, 70(5):1263–1273. DOI: 10 . 1002/mrm.24577.
- [81] J. I. Tamir, W. Chen, P. Lai, et al. **T2 Shuffling: Multicontrast 3D Fast Spin Echo Imaging.** *Proc Intl Soc Mag Reson Med* 23, Toronto, 2015. p. 3399.
- [82] R. R. Coifman and D. L. Donoho. **Translation-Invariant De-Noising.** *Wavelets and Statistics*. Ed. by A. Antoniadis and G. Oppenheim. New York, NY, 1995. pp. 125–150. DOI: 10 . 1007/978-1-4612-2544-7_9.
- [83] M. A. T. Figueiredo and R. D. Nowak, **An EM algorithm for wavelet-based image restoration.** *IEEE Trans Image Process* 2003, 12(8):906–916. DOI: 10 . 1109 / TIP . 2003 . 814255.
- [84] J. I. Tamir, W. Chen, P. Lai, et al. **Improved Scan Efficiency of 3D Fast Spin Echo with Subspace-Constrained Reconstruction.** *Proc Intl Soc Mag Reson Med* 23, 23, Toronto, 2015. p. 2463.
- [85] M. E. Boesen, L. A. S. Maior Neto, A. Pulwinski, et al., **Fast spin echo imaging of carotid artery dynamics.** *Magn Reson Med* 2015, 74(4):1103–1109. DOI: 10 . 1002 / mrm.25494.
- [86] F. Huang, S. Vijayakumar, Y. Li, et al., **A software channel compression technique for faster reconstruction with many channels.** *Magn Reson Imaging* 2008, 26(1):133–141. DOI: 10.1016/j.mri.2007.04.010.
- [87] M. Uecker, P. Lai, M. J. Murphy, et al., **ESPIRiT—an eigenvalue approach to auto-calibrating parallel MRI: where SENSE meets GRAPPA.** *Magn Reson Med* 2014, 71(3):990–1001. DOI: 10 . 1002/mrm.24751.
- [88] M. Uecker and J. I. Tamir. *BART: version 0.4.02*. Nov. 2017. DOI: 10 . 5281 / zenodo.1066014.

- [89] M. Mani, M. Jacob, V. Magnotta, et al., **Fast iterative algorithm for the reconstruction of multishot non-cartesian diffusion data.** *Magn Reson Med* 2015, 74(4):1086–1094. DOI: 10.1002/mrm.25486.
- [90] Z.-P. Liang, H. Jiang, C. P. Hess, et al., **Dynamic imaging by model estimation.** *Int J Imaging Syst Technol* 1997, 8(6):551–557.
- [91] A. S. Gupta and Z.-P. Liang. **Dynamic imaging by temporal modeling with principal component analysis.** *Proc Intl Soc Mag Reson Med* 9, Glasgow, Scotland. p. 10.
- [92] J. Tsao, P. Boesiger, and K. P. Pruessmann, **k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations.** *Magn Reson Med* 2003, 50(5):1031–1042. DOI: 10.1002/mrm.10611.
- [93] R. Vidal, **Subspace Clustering.** *IEEE Signal Process Mag* 2011, 28(2):52–68. DOI: 10.1109/MSP.2010.939739.
- [94] L. Vidarsson, S. M. Conolly, K. O. Lim, et al., **Echo time optimization for linear combination myelin imaging.** *Magn Reson Med* 2005, 53(2):398–407. DOI: 10.1002/mrm.20360.
- [95] D. A. Reiter, P.-C. Lin, K. W. Fishbein, et al., **Multicomponent T2 relaxation analysis in cartilage.** *Magn Reson Med* 2009, 61(4):803–809. DOI: 10.1002/mrm.21926.
- [96] J. B. M. Warntjes, O. D. Leinhard, J. West, et al., **Rapid magnetic resonance quantification on the brain: Optimization for clinical usage.** *Magn Reson Med* 2008, 60(2):320–329. DOI: 10.1002/mrm.21635.
- [97] A. M. Loening, M. Saranathan, N. Ruangwattanapaisarn, et al., **Increased speed and image quality in single-shot fast spin echo imaging via variable refocusing flip angles.** *J Magn Reson Imaging* 2015, 42(6):1747–1758. DOI: 10.1002/jmri.24941.
- [98] J. Y. Cheng, K. Hanneman, T. Zhang, et al., **Comprehensive motion-compensated highly accelerated 4D flow MRI with ferumoxytol enhancement for pediatric congenital heart disease.** *J Magn Reson Imaging* 2016, 43(6):1355–1368. DOI: 10.1002/jmri.25106.
- [99] T. Zhang, S. Chowdhury, M. Lustig, et al., **Clinical performance of contrast enhanced abdominal pediatric MRI with fast combined parallel imaging compressed sensing reconstruction.** *J Magn Reson Imaging* 2014, 40(1):13–25. DOI: 10.1002/jmri.24333.

- [100] L. Feng, L. Axel, H. Chandarana, et al., **XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing.** *Magn Reson Med* 2016, 75(2):775–788. DOI: 10.1002/mrm.25665.
- [101] Y. Wang and L. Ying, **Compressed sensing dynamic cardiac cine MRI using learned spatiotemporal dictionary.** *IEEE Trans Biomed Eng* 2014, 61(4):1109–1120. DOI: 10.1109/TBME.2013.2294939.
- [102] R. M. Lebel, J. Jones, J.-C. Ferre, et al., **Highly accelerated dynamic contrast enhanced imaging.** *Magn Reson Med* 2014, 71(2):635–644. DOI: 10.1002/mrm.24710.
- [103] M. Murphy, M. Alley, J. Demmel, et al., **Fast l1-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime.** *IEEE Trans Med Imaging* 2012, 31(6):1250–1262. DOI: 10.1109/TMI.2012.2188039.
- [104] M. Uecker, P. Virtue, F. Ong, et al. **Software Toolbox and Programming Library for Compressed Sensing and Parallel Imaging.** *ISMRM Scientific Workshop - Data Sampling and Image Reconstruction*, Sedona, 2013.
- [105] J. Gai, N. Obeid, J. L. Holtrop, et al., **More IMPATIENT: A Gridding-Accelerated Toeplitz-based Strategy for Non-Cartesian High-Resolution 3D MRI on GPUs.** *J Parallel Distrib Comput* 2013, 73(5):686–697. DOI: 10.1016/j.jpdc.2013.01.001.
- [106] M. S. Hansen and T. S. Sørensen, **Gadgetron: an open source framework for medical image reconstruction.** *Magn Reson Med* 2013, 69(6):1768–1776. DOI: 10.1002/mrm.24389.
- [107] F. Knoll, A. Schwarzl, C. Diwoky, et al. **gpuNUFFT-an open source GPU library for 3D regridding with direct Matlab interface.** *International Society for Magnetic Resonance in Medicine: Scientific Meeting & Exhibition*, Milan, 2014. p. 4297.
- [108] M. Uecker, F. Ong, J. I. Tamir, et al. **Berkeley advanced reconstruction toolbox.** *Proc Intl Soc Mag Reson Med* 23. 23, Toronto, 2015. p. 2486.
- [109] T. Ou, F. Ong, M. Uecker, et al. **Nufft: fast auto-tuned gpu-based library.** *Proc Intl Soc Mag Reson Med* 25, Honolulu, Hawaii, 2017. p. 3807.
- [110] J.-M. Lin, **Python Non-Uniform Fast Fourier Transform (PyNUFFT): An Accelerated Non-Cartesian MRI Package on a Heterogeneous Platform (CPU/GPU).** *Journal of Imaging* 2018, 4(3):51. DOI: 10.3390/jimaging4030051.

- [111] S. S. Stone, J. P. Haldar, S. C. Tsao, et al., **Accelerating Advanced MRI Reconstructions on GPUs**. *J Parallel Distrib Comput* 2008, 68(10):1307–1318. DOI: 10.1016/j.jpdc.2008.05.013.
- [112] D. Kim, J. Trzasko, M. Smelyanskiy, et al., **High-Performance 3D Compressive Sensing MRI Reconstruction Using Many-Core Architectures**. *Int J Biomed Imaging* 2011, 2011473128. DOI: 10.1155/2011/473128.
- [113] P. Després and X. Jia, **A review of GPU-based medical image reconstruction**. *Phys Med* 2017, 4276–92. DOI: 10.1016/j.ejmp.2017.07.024.
- [114] H. Xue, S. Inati, T. S. Sørensen, et al., **Distributed MRI reconstruction using Gadgetron-based cloud computing**. *Magn Reson Med* 2015, 73(3):1015–1025. DOI: 10.1002/mrm.25213.
- [115] G. T. Kowalik, J. A. Steeden, and V. Muthurangu, **Implementation of a generalized heterogeneous image reconstruction system for clinical magnetic resonance: HETEROGENEOUS IMAGE RECONSTRUCTOR FOR MRI**. *Concurr Comput* 2015, 27(6):1603–1611. DOI: 10.1002/cpe.3349.
- [116] S. T. Ting, R. Ahmad, N. Jin, et al., **Fast implementation for compressive recovery of highly accelerated cardiac cine MRI using the balanced sparse model**. *Magn Reson Med* 2017, 77(4):1505–1515. DOI: 10.1002/mrm.26224.
- [117] J. I. Tamir, M. Uecker, W. Chen, et al., **T2 shuffling: Sharp, multicontrast, volumetric fast spin-echo imaging**. *Magn Reson Med* 2017, 77(1):180–195. DOI: 10.1002/mrm.26102.
- [118] S. Bao, J. I. Tamir, J. L. Young, et al., **Fast comprehensive single-sequence four-dimensional pediatric knee MRI with T2 shuffling**. *J Magn Reson Imaging* 2017, 45(6):1700–1711. DOI: 10.1002/jmri.25508.
- [119] P. B. Roemer, W. A. Edelstein, C. E. Hayes, et al., **The NMR phased array**. *Magn Reson Med* 1990, 16(2):192–225. DOI: 10.1002/mrm.1910160203.
- [120] S. S. Iyer, F. Ong, and M. Lustig, **Towards a Parameter Free ESPIRiT: Soft Weighting for Robust Coil Sensitivity Estimation**. *Proc Intl Soc Mag Reson Med* 24, Singapore, 2016. p. 1093.
- [121] J. I. Tamir, M. Uecker, S. S. Vasanawala, et al. **T2 Shuffling with Partial Fourier Acquisition and Reconstruction**. *ISMRM Scientific Workshop - Data Sampling and Image Reconstruction*, Sedona, 2016.

- [122] D. C. Noll, D. G. Nishimura, and A. Macovski, **Homodyne detection in magnetic resonance imaging**. *IEEE Trans Med Imaging* 1991, 10(2):154–163. DOI: 10.1109/42.79473.
- [123] D. R. Pai and P. J. Strouse, **MRI of the pediatric knee**. *AJR Am J Roentgenol* 2011, 196(5):1019–1027. DOI: 10.2214/AJR.10.6117.
- [124] K. G. Gill, B. A. Nemeth, and K. W. Davis, **Magnetic resonance imaging of the pediatric knee**. *Magn Reson Imaging Clin N Am* 2014, 22(4):743–763. DOI: 10.1016/j.mric.2014.07.010.
- [125] D. Jaramillo and T. Laor, **Pediatric musculoskeletal MRI: basic principles to optimize success**. *Pediatr Radiol* 2008, 38(4):379–391. DOI: 10.1007/s00247-007-0645-4.
- [126] R. Bhargava, G. Hahn, W. Hirsch, et al., **Contrast-enhanced magnetic resonance imaging in pediatric patients: review and recommendations for current practice**. *Magn Reson Insights* 2013, 695–111. DOI: 10.4137/MRI.S12561.
- [127] L. Sun, **Early childhood general anaesthesia exposure and neurocognitive development**. *Br J Anaesth* 2010, 105 Suppl 1i61–8. DOI: 10.1093/bja/aeq302.
- [128] A. J. Mueller, J. K. Kwon, J. W. Steiner, et al., **Improved Magnetic Resonance Imaging Utilization for Children with Musculoskeletal Infection**. *J Bone Joint Surg Am* 2015, 97(22):1869–1876. DOI: 10.2106/JBJS.O.00403.
- [129] M. Recht, M. Macari, K. Lawson, et al., **Impacting key performance indicators in an academic MR imaging department through process improvement**. *J Am Coll Radiol* 2013, 10(3):202–206. DOI: 10.1016/j.jacr.2012.08.008.
- [130] P. J. Strouse and K. Koujok, **Magnetic resonance imaging of the pediatric knee**. *Top Magn Reson Imaging* 2002, 13(4):277–294.
- [131] R. Kijowski, K. W. Davis, M. A. Woods, et al., **Knee joint: comprehensive assessment with 3D isotropic resolution fast spin-echo MR imaging—diagnostic performance compared with that of conventional MR imaging at 3.0 T**. *Radiology* 2009, 252(2):486–495. DOI: 10.1148/radiol.2523090028.
- [132] C. Q. Li, W. Chen, J. K. Rosenberg, et al., **Optimizing isotropic three-dimensional fast spin-echo methods for imaging the knee**. *J Magn Reson Imaging* 2014, 39(6):1417–1425.
- [133] R. Kijowski, K. W. Davis, D. G. Blankenbaker, et al., **Evaluation of the menisci of the knee joint using three-dimensional isotropic resolution fast spin-echo**

- imaging: diagnostic performance in 250 patients with surgical correlation.** *Skeletal Radiol* 2012, 41(2):169–178. DOI: 10.1007/s00256-011-1140-4.
- [134] V. Athalye, M. Lustig, and Martin Uecker, **Parallel magnetic resonance imaging as approximation in a reproducing kernel Hilbert space.** *Inverse Probl* 2015, 31(4):045008. DOI: 10.1088/0266-5611/31/4/045008.
- [135] S. Kim, J. Bosque, J. P. Meehan, et al., **Increase in outpatient knee arthroscopy in the United States: a comparison of National Surveys of Ambulatory Surgery, 1996 and 2006.** *J Bone Joint Surg Am* 2011, 93(11):994–1000. DOI: 10.2106/JBJS.I.01618.
- [136] N. Lefevre, J. F. Naouri, Y. Bohu, et al., **Partial tears of the anterior cruciate ligament: diagnostic performance of isotropic three-dimensional fast spin echo (3D-FSE-Cube) MRI.** *Eur J Orthop Surg Traumatol* 2014, 24(1):85–91. DOI: 10.1007/s00590-012-1135-4.
- [137] T. Ai, W. Zhang, N. K. Priddy, et al., **Diagnostic performance of CUBE MRI sequences of the knee compared with conventional MRI.** *Clin Radiol* 2012, 67(12):e58–63. DOI: 10.1016/j.crad.2012.07.020.
- [138] S. J. Ahn, Y. M. Jeong, B. G. Lee, et al., **Using three-dimensional isotropic SPACE MRI to detect posterolateral corner injury of the knee.** *Acta radiol* 2016, 57(10):1251–1260. DOI: 10.1177/0284185115626470.
- [139] J. Fritz, E. Raithel, G. K. Thawait, et al., **Six-Fold Acceleration of High-Spatial Resolution 3D SPACE MRI of the Knee Through Incoherent k-Space Undersampling and Iterative Reconstruction-First Experience.** *Invest Radiol* 2016, 51(6):400–409. DOI: 10.1097/RLI.0000000000000240.
- [140] S. H. Lee, Y. H. Lee, and J.-S. Suh, **Accelerating knee MR imaging: Compressed sensing in isotropic three-dimensional fast spin-echo sequence.** *Magn Reson Imaging* 2018, 46(0):90–97. DOI: 10.1016/j.mri.2017.10.018.
- [141] R. Kijowski, H. Rosas, A. Samsonov, et al., **Knee imaging: Rapid three-dimensional fast spin-echo using compressed sensing.** *J Magn Reson Imaging* 2017, 45(6):1712–1722. DOI: 10.1002/jmri.25507.
- [142] I. Gans, M. A. Bedoya, V. Ho-Fung, et al., **Diagnostic performance of magnetic resonance imaging and pre-surgical evaluation in the assessment of traumatic intra-articular knee disorders in children and adolescents: what conditions still pose diagnostic challenges?** *Pediatr Radiol* 2015, 45(2):194–202. DOI: 10.1007/s00247-014-3127-5.

- [143] D. L. Schub, F. Altahawi, A. F Meisel, et al., **Accuracy of 3-Tesla magnetic resonance imaging for the diagnosis of intra-articular knee injuries in children and teenagers.** *J Pediatr Orthop* 2012, 32(8):765–769. DOI: 10.1097/BPO.0b013e3182619181.
- [144] Centers for Disease Control and Prevention (CDC), **Nonfatal sports- and recreation-related injuries treated in emergency departments—United States, July 2000–June 2001.** *MMWR Morb Mortal Wkly Rep* 2002, 51(33):736–740.
- [145] Y. Sheu, L.-H. Chen, and H. Hedegaard, **Sports- and Recreation-related Injury Episodes in the United States, 2011–2014.** *Natl Health Stat Report* 2016, (99):1–12.
- [146] J. de Laveaucoupet, F. Audibert, F. Guis, et al., **Fetal magnetic resonance imaging (MRI) of ischemic brain injury.** *Prenat Diagn* 2001, 21(9):729–736. DOI: 10.1002/pd.135.
- [147] M. Akçakaya, S. Weingärtner, T. A. Basha, et al., **Joint myocardial T1 and T2 mapping using a combination of saturation recovery and T2-preparation.** *Magn Reson Med* 2016, 76(3):888–896. DOI: 10.1002/mrm.25975.
- [148] J. I. Hamilton, Y. Jiang, Y. Chen, et al., **MR fingerprinting for rapid quantification of myocardial T1, T2, and proton spin density.** *Magn Reson Med* 2017, 77(4):1446–1458. DOI: 10.1002/mrm.26216.
- [149] A. G. Christodoulou, J. L. Shaw, C. Nguyen, et al., **Magnetic resonance multitasking for motion-resolved quantitative cardiovascular imaging.** *Nature Biomedical Engineering* 2018, 2(4):215–226. DOI: 10.1038/s41551-018-0217-y.
- [150] L. N. Tanenbaum, A. J. Tsiouris, A. N. Johnson, et al., **Synthetic MRI for Clinical Neuroimaging: Results of the Magnetic Resonance Image Compilation (MAGiC) Prospective, Multicenter, Multireader Trial.** *AJNR Am J Neuroradiol* 2017, 38(6):1103–1110. DOI: 10.3174/ajnr.A5227.
- [151] S. Kvernby, M. J. B. Warntjes, H. Haraldsson, et al., **Simultaneous three-dimensional myocardial T1 and T2 mapping in one breath hold with 3D-QALAS.** *J Cardiovasc Magn Reson* 2014, 16(1):466. DOI: 10.1186/s12968-014-0102-0.
- [152] S. Kvernby, M. Warntjes, J. Engvall, et al., **Clinical feasibility of 3D-QALAS - Single breath-hold 3D myocardial T1- and T2-mapping.** *Magn Reson Imaging* 2017, 38:13–20. DOI: 10.1016/j.mri.2016.12.014.
- [153] D. Ma, Y. Jiang, Y. Chen, et al., **Fast 3D magnetic resonance fingerprinting for a whole-brain coverage.** *Magn Reson Med* 2018, 79(4):2190–2197. DOI: 10.1002/mrm.26886.

- [154] M. I. Vargas, J. Boto, and B. M. Delatre, **Synthetic MR Imaging Sequence in Daily Clinical Practice**. *AJNR Am J Neuroradiol* 2016, 37(10):E68–E69. DOI: 10 . 3174 / ajnr . A4895.
- [155] P. Virtue, J. I. Tamir, M. Doneva, et al. **Learning Contrast Synthesis from MR Fingerprinting**. *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 0676.
- [156] E. Gong, S. Banerjee, J. Pauly, et al. **Improved Synthetic MRI from Multi-echo MRI Using Deep Learning**. *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 2795.
- [157] J. Jeing, Y. Nam, and J. Lee. **Penta-contrast imaging: a Novel Pulse Sequence for Simultaneous Acquisition of Proton Density, T1, T2, T2* and FLAIR images**. *Proc Intl Soc Mag Reson Med* 25, Honolulu, Hawaii, 2017. p. 3881.
- [158] S. Skare, T. Sprenger, O. Norbeck, et al., **A 1-minute full brain MR exam using a multicontrast EPI sequence**. *Magn Reson Med* 2018, 79(6):3045–3054. DOI: 10 . 1002 / mrm . 26974.
- [159] J. I. Tamir, V. Taviani, S. S. Vasanawala, et al. **T1-T2 Shuffling: Multi-Contrast 3D Fast Spin-Echo with T1 and T2 Sensitivity**. *Proc Intl Soc Mag Reson Med* 25, Honolulu, Hawaii, 2017. p. 0451.
- [160] C. M. van Uijen and J. H. den Boef, **Driven-equilibrium radiofrequency pulses in NMR imaging**. *Magn Reson Med* 1984, 1(4):502–507. DOI: 10 . 1002 / mrm . 1910010409.
- [161] S. J. Malik, A. Beqiri, F. Padormo, et al., **Direct signal control of the steady-state response of 3D-FSE sequences**. *Magn Reson Med* 2015, 73(3):951–963. DOI: 10 . 1002 / mrm . 25192.
- [162] J. Ma, J. B. Son, Y. Zhou, et al., **Fast spin-echo triple-echo dixon (fTED) technique for efficient T2-weighted water and fat imaging**. *Magn Reson Med* 2007, 58(1):103–109. DOI: 10 . 1002 / mrm . 21268.
- [163] T. Kolda and B. Bader, **Tensor Decompositions and Applications**. *SIAM Rev* 2009, 51(3):455–500. DOI: 10 . 1137 / 07070111X.
- [164] J. Liu, P. Musialski, P. Wonka, et al., **Tensor completion for estimating missing values in visual data**. *IEEE Trans Pattern Anal Mach Intell* 2013, 35(1):208–220. DOI: 10 . 1109 / TPAMI . 2012 . 39.
- [165] C. K. Jones, Q.-S. Xiang, K. P. Whittall, et al., **Linear combination of multiecho data: short T2 component selection**. *Magn Reson Med* 2004, 51(3):495–502. DOI: 10 . 1002 / mrm . 10713.

- [166] R. W. Keener. **Theoretical Statistics**. Springer, Sept. 2010.
- [167] J. K. Barral, E. Gudmundson, N. Stikov, et al., **A robust methodology for in vivo T1 mapping**. *Magn Reson Med* 2010, 64(4):1057–1067. DOI: 10.1002/mrm.22497.
- [168] K. Li, Z. Zu, J. Xu, et al., **Optimized inversion recovery sequences for quantitative T1 and magnetization transfer imaging**. *Magn Reson Med* 2010, 64(2):491–500. DOI: 10.1002/mrm.22440.
- [169] G. Nataraj, J.-F. Nielsen, C. Scott, et al., **Dictionary-Free MRI PERK: Parameter Estimation via Regression with Kernels**. *IEEE Trans Med Imaging* 2018, 1–1. DOI: 10.1109/TMI.2018.2817547.
- [170] H. Qi, J. Sun, H. Qiao, et al., **Simultaneous T1 and T2 mapping of the carotid plaque (SIMPLE) with T2 and inversion recovery prepared 3D radial imaging**. *Magn Reson Med* 2018. DOI: 10.1002/mrm.27361.
- [171] R. Heule, Z. Celicanin, S. Kozerke, et al., **Simultaneous multislice triple-echo steady-state (SMS-TESS) T1 , T2 , PD, and off-resonance mapping in the human brain**. *Magn Reson Med* 2018, 80(3):1088–1100. DOI: 10.1002/mrm.27126.
- [172] S. A. Bobman, S. J. Riederer, J. N. Lee, et al., **Cerebral magnetic resonance image synthesis**. *AJNR Am J Neuroradiol* 1985, 6(2):265–269.
- [173] Y. Jiang, D. Ma, N. Seiberlich, et al., **MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout**. *Magn Reson Med* 2015, 74(6):1621–1631. DOI: 10.1002/mrm.25559.
- [174] J. I. Tamir, V. Taviani, S. Vasanawala, et al. **Data-Driven Image Contrast Synthesis from Efficient Mixed-Contrast Sequences**. *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 2793.
- [175] T. Benkert, J. P. Mugler 3rd, D. S. Rigie, et al., **Hybrid T2 - and T1 -weighted radial acquisition for free-breathing abdominal examination**. *Magn Reson Med* 2018. DOI: 10.1002/mrm.27200.
- [176] A. T. Roh, J. I. Tamir, J. Shaikh, et al. **Four-dimensional T2-weighted imaging in prostate MRI with T2 shuffling**. *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 0713.
- [177] S. Poddar and M. Jacob, **Dynamic MRI Using Smoothness Regularization on Manifolds (SToRM)**. *IEEE Trans Med Imaging* 2016, 35(4):1106–1115. DOI: 10.1109/TMI.2015.2509245.

- [178] U. Nakarmi, Y. Wang, J. Lyu, et al., **A Kernel-Based Low-Rank (KLR) Model for Low-Dimensional Manifold Recovery in Highly Accelerated Dynamic MRI.** *IEEE Trans Med Imaging* 2017, 36(11):2297–2307. DOI: 10.1109/TMI.2017.2723871.
- [179] U. Nakarmi, K. Slavakis, H. Li, et al. **MLS: Self-learned joint manifold geometry and sparsity aware framework for highly accelerated cardiac cine imaging.** *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 0441.
- [180] S. T. Roweis and L. K. Saul, **Nonlinear dimensionality reduction by locally linear embedding.** *Science* 2000, 290(5500):2323–2326. DOI: 10.1126/science.290.5500.2323.
- [181] P. Cao, X. Zhu, S. Tang, et al., **Shuffled magnetization-prepared multicontrast rapid gradient-echo imaging.** *Magn Reson Med* 2018, 79(1):62–70. DOI: 10.1002/mrm.26986.
- [182] Y. Zur and S. Stokar, **A phase-cycling technique for canceling spurious echoes in NMR imaging.** *J Magn Reson* 1987, 71(2):212–228. DOI: 10.1016/0022-2364(87)90051-5.
- [183] Y. Zur, **Analysis of the multi-echo spin-echo pulse sequence.** *Concepts Magn Reson Part A* 2017, 46A(1):e21402. DOI: 10.1002/cmr.a.21402.
- [184] E. Levine, B. Daniel, S. Vasanawala, et al., **3D Cartesian MRI with compressed sensing and variable view sharing using complementary poisson-disc sampling.** *Magn Reson Med* 2017, 77(5):1774–1785. DOI: 10.1002/mrm.26254.
- [185] E. Levine and B. Hargreaves, **On-the-Fly Adaptive k-Space Sampling for Linear MRI Reconstruction Using Moment-Based Spectral Analysis.** *IEEE Trans Med Imaging* 2018, 37(2):557–567. DOI: 10.1109/TMI.2017.2766131.
- [186] J. P. Haldar and D. Kim, **OEDIPUS: An Experiment Design Framework for Sparsity-Constrained MRI.** 2018.
- [187] B. Gozcu, R. K. Mahabadi, Y.-H. Li, et al., **Learning-Based Compressive MRI.** *IEEE Trans Med Imaging* 2018, 37(6):1394–1406. DOI: 10.1109/TMI.2018.2832540.
- [188] S. S. Iyer, B. Bilgic, and K. Setsompop. **Faster T2 Shuffling with Wave-encoding.** *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 0382.
- [189] K. Hammernik, T. Klatzer, E. Kobler, et al., **Learning a variational network for reconstruction of accelerated MRI data.** *Magn Reson Med* 2018, 79(6):3055–3071. DOI: 10.1002/mrm.26977.

- [190] B. Zhu, J. Z. Liu, S. F. Cauley, et al., **Image reconstruction by domain-transform manifold learning**. *Nature* 2018, 555(7697):487–492. DOI: 10 . 1038 / nature25988.
- [191] F. Ong and M. Lustig. **k-Space Aware Convolutional Sparse Coding: Learning from Undersampled k-space Datasets for Reconstruction**. *Proc Intl Soc Mag Reson Med* 26, Paris, 2018. p. 3378.
- [192] J. Cai, E. Candès, and Z. Shen, **A Singular Value Thresholding Algorithm for Matrix Completion**. *SIAM J Optim* 2010, 20(4):1956–1982. DOI: 10 . 1137 / 080738970.
- [193] E. J. Candès and B. Recht, **Exact Matrix Completion via Convex Optimization**. *Found Comput Math* 2009, 9(6):717. DOI: 10 . 1007 / s10208-009-9045-5.
- [194] R. Burkard, M. Dell’Amico, and S. Martello. **Assignment Problems**. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, Mar. 2009.