



Theses and Dissertations

2019-08-01

Robust Coil Combination for bSSFP MRI and the Ordering Problem for Compressed Sensing

Nicholas Brian McKibben
Brigham Young University

Follow this and additional works at: <https://scholarsarchive.byu.edu/etd>



Part of the [Engineering Commons](#)

BYU ScholarsArchive Citation

McKibben, Nicholas Brian, "Robust Coil Combination for bSSFP MRI and the Ordering Problem for Compressed Sensing" (2019). *Theses and Dissertations*. 7596.
<https://scholarsarchive.byu.edu/etd/7596>

This Thesis is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of BYU ScholarsArchive. For more information, please contact scholarsarchive@byu.edu, ellen_amatangelo@byu.edu.

Robust Coil Combination for bSSFP MRI and the Ordering Problem
for Compressed Sensing

Nicholas Brian Patrick McKibben

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of
Master of Science

Neal K. Bangerter, Chair
Brian D. Jeffs
Scott R. Burt

Department of Electrical and Computer Engineering
Brigham Young University

Copyright © 2019 Nicholas Brian Patrick McKibben
All Rights Reserved

ABSTRACT

Robust Coil Combination for bSSFP MRI and the Ordering Problem for Compressed Sensing

Nicholas Brian Patrick McKibben

Department of Electrical and Computer Engineering, BYU
Master of Science

Balanced steady-state free precession (bSSFP) is a fast, SNR-efficient magnetic resonance (MR) imaging sequence suffering from dark banding artifacts due to its off-resonance dependence. These banding artifacts are difficult to mitigate at high field strengths and in the presence of metallic implants. Recent developments in parametric modelling of bSSFP have led to advances in banding removal and parameter estimation using multiple phase-cycled bSSFP [1], [2]. With increasing number of coils in receivers, more storage and processing is required [3]. Coil combination is used to reduce dimensionality of these datasets which otherwise might be prohibitively large or computationally intractable for clinical applications [4]. However, our recent work demonstrates that some combination methods are problematic in conjunction with elliptical phase-cycled bSSFP [5].

This thesis will present a method for phase estimation of coil-combined multiple phase-cycled bSSFP to reduce storage and computational requirements for elliptical models. This method is general and works across many coil combination techniques popular in MR reconstruction including the geometric coil combine and adaptive coil combine algorithms [6], [7]. A viable phase estimate for the sum-of-squares is also demonstrated for computationally efficient dimension reduction. Simulations, phantom experiments, and in vivo MR imaging is performed to validate the proposed phase estimates.

Compressed sensing (CS) is an increasingly important acquisition and reconstruction framework. CS MR allows for reconstruction of datasets sampled well-under the Nyquist rate and its application is natural in MR where images are often sparse under common linear transforms [8]. An extension of this framework is the ordering problem for CS, first introduced in 2008 [9]. Although the assumption is made in CS that images are sparse in some specified transform domain, it might not be maximally sparse. For example, a signal ordered such that it is monotonic is maximally sparse in the finite differences domain. Knowledge of the correct ordering of an image's pixels can lead to much more sparse and powerful regularizers for the CS inverse problem. However, this problem has met with little interest due to the strong dependence on initial image estimates.

This thesis will also present an algorithm for estimating the optimal order of a signal such that it is maximally sparse under an arbitrary linear transformation without relying on any prior image estimate. The algorithm is combinatoric in nature and feasible for small signals of interest such as T_1 mapping time curves. Proof of concept simulations are performed that validate performance of the algorithm. Computationally feasible modifications for in vivo cardiac T_1 mapping are also demonstrated.

Keywords: MRI, bSSFP, compressed sensing, coil combination

ACKNOWLEDGMENTS

I am grateful to my committee and advisers who have provided both useful feedback and support. I have appreciated the freedom Neal has given me in exploring my oftentimes wacky ideas as well as gentle nudges in the right direction, ultimately resulting in this thesis. I will fondly remember my many hours in Dr. Jeffs' lectures as well as advice and encouragement during office hours. The insightful and interesting feedback as well as attention to detail from Dr. Burt has been invaluable to the quality of this thesis. Ed DiBella has been a constant source of direction and inspiration, always reminding me to focus on the important things.

Many thanks belong to the Bangerter lab: Grayson for many useful and interesting discussions, Michael for long hours together in the lab and on the road, Merry for always being available for my many questions, and Laurel for always letting me tag along to learn.

Finally, I am especially grateful for the support and encouragement of my family. Dad has been there every step of the way, proofreading my abstracts, papers, and thesis. My wife has been a constant source of love, encouragement, and understanding of long nights in the lab.

TABLE OF CONTENTS

Title Page	i
Abstract	ii
Acknowledgments	iii
List of Tables	vii
List of Figures	viii
Chapter 1 Introduction	1
1.1 Contributions	1
1.2 Outline	2
Chapter 2 Principles of Magnetic Resonance Imaging	6
2.1 Physics of MRI	6
2.1.1 Nuclear Spin and Magnetic Moments	6
2.1.2 Larmor Frequency	8
2.1.3 Net Magnetization	9
2.1.4 Excitation	11
2.1.5 Relaxation	11
2.1.6 Bloch Equations	12
2.2 Spatial Localization	14
2.2.1 Slice-Selective Excitation	15
2.2.2 Frequency and Phase Encoding	15
2.2.3 Signal Equation	17
2.2.4 k -space	18
2.3 Real World Imaging	18
2.3.1 Magnetic Field Inhomogeneity	19
2.3.2 Chemical Shift	19
2.3.3 Noise	20
Chapter 3 Balanced Steady-State Free Precession	21
3.1 Introduction	21
3.2 Pulse Sequence	21
3.3 Off-resonance and the bSSFP Spectrum	22
3.4 Signal Model	24
3.5 Elliptical Signal Model	25
Chapter 4 Robust Coil Combination for the bSSFP Elliptical Signal Model	27
4.1 Introduction	27
4.2 Theory	29

4.3	Methods	35
4.3.1	Simulation	36
4.3.2	Phantom	36
4.3.3	In vivo	37
4.4	Results	37
4.4.1	Simulation	37
4.4.2	Phantom	39
4.4.3	In vivo	41
4.5	Conclusion	41
Chapter 5 Compressed Sensing		46
5.1	Introduction	46
5.2	Problem Statement	47
5.3	MR Problem Statement	50
5.4	Algorithms	50
5.4.1	Gradient Descent	51
5.4.2	Iterative Soft Thresholding	52
Chapter 6 The Ordering Problem for Compressed Sensing		55
6.1	Introduction	55
6.2	Theory	58
6.3	Methods	63
6.3.1	Simulations	63
6.3.2	In vivo	65
6.4	Results	66
6.4.1	Simulations	66
6.4.2	In vivo	68
6.5	Discussion	69
References		73
Appendix A Coil Combination for ESM bSSFP		82
A.1	Introduction	82
A.2	Maximum SNR Coil Combination	82
A.3	Stochastic Matched Filter Coil Combination	83
Appendix B Deep Convolutional Neural Networks for Estimation of Pre-Reconstruction Data Reorderings		85
B.1	Purpose	85
B.2	Methods	85
B.3	Results	86
B.4	Discussion	87
B.5	Conclusion	88

Appendix C Synthetic Banding for bSSFP Data Augmentation	89
C.1 Introduction	89
C.2 Theory	90
C.3 Methods	91
C.4 Results	91
C.5 Conclusion and Discussion	92
C.6 Acknowledgements	93
Appendix D Deep Learning Super-FOV for Accelerated bSSFP Banding Reduction . .	94
D.1 Introduction	94
D.2 Theory	94
D.3 Methods	95
D.4 Results	96
D.5 Conclusion and Discussion	96
Appendix E Generation of Arbitrary Spectral Profiles Using Orthonormal Basis Com-	
 binations of bSSFP MRI	97
E.1 Introduction	97
E.2 Theory	97
E.3 Methods	98
E.4 Results	99
E.5 Discussion	100
E.6 Acknowledgements	101

LIST OF TABLES

2.1	Longitudinal and transverse relaxation constants for some common organic tissues at 1.5T.	13
6.1	Summary of results for exhaustive simulation of optimal ordering simulation. Bold entries indicate the best values. We find that the optimal ordering was found using the proposed algorithm, i.e., the minimum cost ordering corresponding to a comparable minimum NRMSE ordering.	68

LIST OF FIGURES

2.1	Amperian loop: a point charge with charge e travels in circular loop with radius r and velocity \vec{v}	7
2.2	Red arrows represent that magnetic moments of hydrogen nuclei. On the left, when no magnetic field is applied, the moments are randomly oriented and the net magnetization of large groups of moments is zero. However, as depicted on the right, in the presence of a magnetic field, the moments tend to align parallel or anti-parallel to the field. Slightly more will align with the field leading to a non-zero net magnetization in the direction of the field.	10
2.3	The longitudinal and transverse components of the magnetization vector \vec{M} are governed by equations (2.23) and (2.25), respectively. These exponential curves are illustrated here with the inclusion of a visualization of the longitudinal and transverse components of \vec{M} . At $t = 0$ the magnetization vector is tipped fully into the transverse plane followed by relaxation back to thermal equilibrium.	13
2.4	We often want to excite spins in just a slice or slab only, rather than exciting the entire volume. Application of the linear gradient field G_z varies the resonant frequencies as a function of position along the z -axis. When a band-limited RF pulse is applied in conjunction with this gradient, only spins resonating at frequencies within the bandwidth of the RF pulse will be excited, “selecting” the slice as shown.	16
3.1	A typical bSSFP pulse sequence diagram. Depicted are four waveforms: one for each the RF transmitter and linear gradient fields. The RF is played repeatedly with a period TR. Notice that the gradients’ waveforms integrate to zero, eliminating coherences TR to TR.	22
3.2	Magnitude and phase components of a simulated bSSFP spectral profile using TR = 12ms, $\alpha = 30^\circ$, $T_1 = 830$ ms, and $T_2 = 80$ ms. The horizontal axis measures off-resonance ranging between $\pm 1/\text{TR}$ Hz. Two phase components are shown. The phase predicted directly after excitation is different than when measured at TE = TR/2 as the magnetization vector at TE/2 has already gone through half of its total precession angle. Its not necessary to do so, but we can compensate for this by subtracting half of the phase-cycle value from the phase at $t = \text{TE}/2$ to recover the phase at $t = 0$ [10].	23
4.1	The geometric solution to the bSSFP elliptical signal model describes the cross point, I_d , of two pairs of 180° phase-cycled pixels plotted in the complex plane as the demodulated value.	29
4.2	Phased-array receivers produce an ellipse per channel in the complex plane. The GCC method is nonlinear and leads to unpredictable combined phase-cycle points. The ellipse is mapped chaotically through the coil combination and the phase-cycle points show that a solution to the IGS does not exist for this pixel, i.e., there is no cross-point. But the SMF is a linear mapping and the elliptical model is preserved after a phase-unwrapping step. The artifact at the bottom of the SMF ellipse is due to errors during phase-unwrapping.	32

4.3	Steps to form a composite ellipse. The underlying true ellipse is unknown. With a phased-array receiver, multiple coil ellipses are observed. The coil ellipses are related to the true ellipse by the complex coil sensitivities. Coil ellipse registration is performed which maps each coil ellipse to an arbitrarily scaled and rotated reference of the same shape. The inverse affine transform describes how to map from the registered reference to the coil ellipse. This transformation is interesting because the magnitude of the transformation gives us a relative confidence measure to weight the linear combination of registered coil ellipses. It serves the same function coil sensitivity maps do in the optimal coil combination reconstruction proposed by Roemer et al. Therefore, we use the same combination scheme to form the composite ellipse.	34
4.4	This figure shows the phase of the composite coil sensitivity map constructed by taking the phase of the maximum intensity coil sensitivity at each pixel for the simulated coil sensitivity map. Phase discontinuities can clearly be seen at the coil boundaries.	35
4.5	The left image shows the SOS coil combined 0° phase cycle (SNR=47). The right image shows the simple phase substitution for the SOS reconstruction with the phase discontinuities predicted in the Theory section. The discontinuities lie along the coil intersections as well as along the horizontal bSSFP “null.”	38
4.6	This figure shows the comparative results of the residual error images for a simulation run with SNR=47. At the bottom of each image the error is given along with a contrast factor that normalizes the residual image. The true reference image is the single coil IGS with no additive noise of a uniform circular phantom. The columns are labeled according to the reconstruction method: “coil-by-coil” IGS, “naive” (meaning that the coil combination and IGS were performed without any phase substitution), the proposed “full” phase substitution, and the proposed “simple” phase substitution. The proposed methods consistently outperform the naive approach. The red arrow points out a singularity in the naive SOS residual.	39
4.7	NRMSE plotted as a function of SNR using a logarithmic scale for prior coil combination with and without the proposed phase substitutions. Nonlinear coil combination techniques greatly benefit from the simple phase substitution and all coil combination methods benefit from the proposed methods. When the full composite ellipse is constructed, the quality of the reconstruction for the SOS coil combined images are superior to the SMF. SOS with the phase substitutions provide the best reconstruction in terms of raw NRMSE over the SNR range simulated.	40
4.8	SSIM plotted as a function of SNR using for prior coil combination with and without the proposed phase substitutions. The GCC solution with the proposed phase substitutions provide the best similarity score over the range simulated.	41
4.9	This figure shows the comparative results of the uniform phantom experiment. The naive and proposed methods are compared to the coil-by-coil IGS reconstruction. The columns are labeled according to the reconstruction method. The proposed methods consistently outperform the naive approach, showing virtually no artifacts in this uniform phantom. The red arrow points out an artifact due to phase discontinuities of the naive SMF coil combination. The SMF may have artifacts if the combined image’s phase is non-smooth while the proposed method exhibits no comparable artifacts.	42

4.10	This figure shows the comparative results of the uniform phantom experiment along with the NRMSE between the results using each coil combination method. Notice the scaling applied to the images for more obvious contrast. All naive methods produce varying levels of error: the SOS and GCC, as expected, produce images of unusable quality. Errors due to a phase singularity in the naive SMF can be seen. The full phase substitution consistently provides the best performance with the simple phase substitution performing not much worse.	43
4.11	This image shows the simple phase substitution of the 90° phase-cycled image for the uniform phantom experiment. The phase discontinuities clearly shown at the boundaries of the 4 coils matching the results of the simulation.	43
4.12	This figure shows the comparative results of the in vivo brain experiment. The naive and proposed methods are compared to the coil-by-coil IGS reconstruction. The columns are labeled according to the reconstruction method. The proposed methods clearly outperform the naive approaches for the SOS and GCC.	44
4.13	This figure shows the comparative residual error in the form of NRMSE for the in vivo brain experiment. As expected, massive error is found for the naive SOS reconstruction where all phase information has been discarded. Performance of the naive GCC also matches that of the simulation and phantom experiments.	44
4.14	This image shows the simple phase substitution of the 90° phase-cycled image for the in vivo brain experiment. Interestingly, the phase discontinuities cannot be obviously seen from the 4 coils in the brain matter, but are clear in the outer skull-shell, matching the results of the phantom and simulation experiments.	45
5.1	Iterative Method for CS Recovery	48
6.1	Mean Gini index for several transforms of 1000 optimally ordered T_1 time curves. FD, DCT, and wavelet (WVLT) transforms see improved sparsity over no transformation at all. The DCT empirically sparsifies the time curve best in this experiment. The error bars show the standard deviations.	61
6.2	Gini index of the DCT of T_1 relaxation curves across a range of 1000 T_1 values. Monotonic ordering, optimal ordering, and the three most common orderings are plotted. The most common orderings are near-optimal over certain ranges of T_1 values. Notice that while the first two most common optimal ordering are distinct, they are equivalent orderings: they produce the same sparsity measure. The first and second most common orderings are near-optimal for much of the entire simulated range. Some common orderings are only optimal over a small range of T_1 values and become worse than monotonic ordering for some ranges.	63
6.3	Frequency of the three most common optimal orderings. The ordering indices are shown along the x -axis. The two most common orderings account for almost 20% of all orderings.	64
6.4	Numerical phantom used to simulate inversion recovery pulse sequence to validate the proposed theory of prior-less ordering estimation for compressed sensing.	65
6.5	Cartesian k -space undersampling pattern for one time frame of the numerical phantom simulation.	66

6.6	Corrupted phantom image due to random Cartesian undersampling in the phase-encode direction.	67
6.7	Comparative reconstruction of the last time frame of the simulated T_1 mapping dataset. The reconstruction using monotonic ordering is shown on the left with residual underneath, while the reconstruction on the right uses the best cost orderings.	68
6.8	NRMSE plotted as a function of IST iteration for monotonically ordered and best cost ordered reconstructions.	69
6.9	Comparative reconstruction results for a representative pixel's time curve. The naive inverse fast Fourier transform (IFFT) corresponds to a zero-filled reconstruction. The curves corresponding to the monotonic ordering and best cost/NRMSE orderings appear close, with the best cost/NRMSE slightly better aligning with the true time curve. These small improvements for each pixel lead to improved reconstruction fidelity on the whole.	70
6.10	The normalized cost function over a fixed number of iterations for a single-slice cardiac reconstruction. The reconstruction using the average optimal ordering shows reduced cost function for all computed iterations.	71
6.11	The image on the left is a T_1 map generated from reconstructed retroactively undersampled cardiac data. No data reordering was applied during the reconstruction. The image on the right shows the residual image generated using the fully sampled cardiac data. The red bounding box shows an area of interest with mean error $1.03e-2 \pm 0.25s$. The scales are indicated in seconds.	71
6.12	The image on the left is a T_1 map generated from reconstructed retroactively undersampled cardiac data. The average optimal orderings for the expected range of T_1 values was applied during the reconstruction. The image on the right shows the residual image generated using the fully sampled cardiac data. The red bounding box shows an area of interest with mean error $3.20e-3 \pm 0.26s$. The scales are indicated in seconds.	72
B.1	A U-net architecture with 16 ray undersampled images as input and reordering matrices (real and imaginary) as output.	86
B.2	Results generated by the neural network compared to a magnitude image of the 72 ray NUFFT reconstruction.	87
C.1	An illustration of the model architecture.	90
C.2	An illustration of model verification using the geometric solution from elliptical signal model. (Left) Band reduced image is generated using four acquired images. (Right) Band reduced image is generated using two acquired images and two synthetically generated images.	91
C.3	bSSFP knee images. (a–d) are the acquired bSSFP phase-cycled images (0° , 180° , 90° and 270°) (e–f) are the synthetically generated bSSFP phase-cycled images (90° and 270°).	92
C.4	bSSFP band reduced images using the geometric solution to the elliptical signal model. (a) Inputs into the model are four acquired bSSFP phase-cycled images. (b) Inputs into the model are two acquired and two synthetically generated bSSFP phase-cycled images.	92

D.1	Undersampled k -space data from two phase-cycled acquisitions is fed to the model as input. The output is compared using mean squared error to the geometric solution to the elliptical signal model generated from four fully sampled phase-cycled acquisitions.	95
D.2	(a–b) is the combined k -space and image domain representation, respectively, from the two undersampled acquisitions (even lines from first phase-cycled image, odd lines from the second) as it was input to the neural network for training. (c–d) show the k -space and image space representations of the geometric solution to the elliptical signal model which was used as ground truth for training. (d–e) shows the model’s estimate of the banding-free image.	96
E.1	Magnitude plots of six basis functions generated from a high tip angle (90°) SSFP simulation. The shape of this spectral profile was observed in images from the fat phantom. The top three basis functions are similar to the sine terms of the Fourier series, while the bottom three basis functions are similar to the cosine terms of the Fourier series.	98
E.2	Magnitude plots of six basis functions generated from a moderate tip angle SSFP simulation. The shape of this spectral profile was observed in images from the water phantom. The top three basis functions are odd functions, while the bottom three basis functions are even functions. Despite the unusual shape of this set of basis functions, it can be used to generate a variety of arbitrary off-resonance spectra after it has been orthogonalized with Gram-Schmidt orthogonalization.	99
E.3	Generated spectra from a fat phantom using the high flip angle basis functions. A shim gradient was applied from the top to the bottom of each image to create a range of off-resonance frequencies. The top row shows the spectral profile generated using a notch spectrum while the bottom row shows the spectral profile generated from a band-pass spectrum.	100
E.4	Generated spectra from a water phantom using the moderate flip angle basis functions. A shim gradient was applied from the top to the bottom of each image to create a range of off-resonance frequencies. The top row shows the spectral profile generated using a notch spectrum while the bottom row shows the spectral profile generated using a band-pass spectrum.	101

CHAPTER 1. INTRODUCTION

Magnetic resonance (MR) imaging is a powerful imaging modality with virtually limitless possibilities for interesting and useful contrast generation. However, with great power comes great imaging times, perhaps the most troublesome issue plaguing the whole of the entire MR ecosystem. With each advancement in fast imaging new possibilities emerge and new issues arise requiring further study and innovation. With each publication in the field of MR signal processing and advanced pulse sequence design, one gets the feeling that these researchers will not be satisfied: they have a need for speed.

The present work is concerned with two specific areas: 1) the intersection of coil combination and multiple phase-cycled balanced steady-state free precession (bSSFP) and 2) the ordering problem for compressed sensing (CS). Both of these areas are deeply concerned with speeding up the MR experiment. The former considers speeding up data processing through data compression and the latter considers reconstruction techniques that assist in potential speed-ups during data acquisition.

This introduction describes the specific contributions made by the thesis in the aforementioned areas as well as informing the reader about where those contributions can be found.

1.1 Contributions

The main contributions of this thesis are as follows:

- We propose a novel phase-substitution method for coil combination of multiple phase-cycled bSSFP imaging. This method can be applied in conjunction with many coil combination strategies. The method is a direct answer to the question we first asked in [5] where we demonstrated that upstream coil combination is detrimental to downstream processing results relying on the bSSFP elliptical signal model. There are two phase estimation variants presented: a more computationally demanding but more accurate full phase estimate and a

faster, less accurate simple phase estimate. In simulation and under a variety of SNR conditions, both phase substitutions outperform standard coil combination techniques. Phantom and in vivo experiments validate the method under real imaging conditions and show that the simple, more computationally efficient method performs very well and is viable for compression of large datasets during preprocessing for phase sensitive bSSFP elliptical signal model algorithms.

- We present a prior-less version of the ordering problem for CS and an algorithm designed to find the best ordering for any arbitrary sparsifying linear transformation. The ordering problem is a niche area in signal processing for CS MR image reconstruction that has conventionally relied on prior image estimates in the form of low resolution priors, reference datasets, or first-pass CS solutions. This has made it unattractive as a mainstream solution for the CS MR reconstruction problem as these prior image estimates are sometimes not good indicators of order relations between pixels. Herein we recast the problem using an implicit model defined by the order relations to show that optimal ordering estimates can be found without an image prior. An exhaustive simulation is performed to validate the proposed theory using a computationally tractable T_1 mapping problem. Inclusion of additional priors for in vivo cardiac T_1 mapping is demonstrated as a computationally feasible extension of the proposed framework.

1.2 Outline

The thesis is organized as follows:

Principles of Magnetic Resonance Imaging

Chapter 2 introduces the basic concepts of magnetic resonance imaging (MRI) that will be assumed in later chapters.

Balanced Steady-State Free Precession

Chapter 3 presents the balanced steady-state free precession sequence and the basic terminology of MR pulse sequence design. The bSSFP signal equations and the elliptical signal model are derived. Understanding of this sequence is fundamental for Chapter 4 and many appendices.

Robust Coil Combination for the bSSFP Elliptical Signal Model

Chapter 4 presents the proposed method for phase estimation to be used in conjunction with existing coil combination methods. The problem of multiple phase-cycled bSSFP coil combination is motivated and demonstrated. The proposed method is derived theoretically in the context of existing coil combination methods and parallel imaging methods. The method is validated using simulations over a range of SNRs as well as phantom and in vivo experiments to demonstrate the validity of the phase estimate in real imaging scenarios.

Compressed Sensing

Chapter 5 presents the theory of compressed sensing. The general problem is described and contextualized for the CS MR inverse imaging problem. Two algorithms are detailed and their updates derived: gradient descent and iterative soft thresholding. These algorithms will reappear in Chapter 6.

The Ordering Problem for Compressed Sensing

Chapter 6 introduces the ordering problem for CS MRI. A brief literature review is followed by the proposal of a new class of ordering problem. Previously, orderings have been generated using image priors. However, these priors are often suspect. This chapter proposes a combinatorial algorithm to find an optimal ordering for the ordered CS problem without an image prior. The algorithm is validated through simulation of a T_1 mapping time curve reconstruction. Modifications are made to incorporate additional priors for computationally feasible in vivo cardiac T_1 mapping.

Appendices

The appendices contain derivations which are not strictly necessary to the development of the chapters as well as selected conference proceedings related to the topics in the present thesis. I am the first author or an equally contributing author on all included conference proceedings except when otherwise indicated.

Coil Combination for ESM bSSFP

Appendix A details the Roemer maximum SNR coil combination and Walsh’s adaptive coil combination methods [7], [11]. Derivations are presented for the nonlinear geometric relationships when applying the elliptical signal model for multiple phase-cycled bSSFP.

Deep Convolutional Neural Networks for Estimation of Pre-Reconstruction Data Orderings [5]

Appendix B contains an accepted proceeding to the International Society of Magnetic Resonance in Medicine’s (ISMRM) 2018 Workshop on Machine Learning which is related to the ordering problem for CS MRI. Dynamic contrast enhanced cardiac MR datasets using a golden angle radial sampling pattern were fed into a convolutional neural network which learned the orderings for the real and imaginary parts of the images. This architecture generates orderings directly from corrupted image estimates and compares favorably to those generated using orderings from simple low resolution priors.

Synthetic Banding for bSSFP Data Augmentation [12]

Appendix C contains an accepted proceeding to the 2019 ISMRM annual conference. Many algorithms exist that can effectively remove banding artifacts from bSSFP images, typically by requiring multiple phase-cycled acquisitions, which increases scan time. While some of the algorithms can suppress some banding with two sets of phase-cycled acquisitions, much more accurate band suppression is typically achieved with at least four phase-cycled acquisitions. In this work, we present a deep learning method for synthesizing additional phase-cycled images from a

set of at least two phase-cycled images that can then be used with existing band reduction techniques in order to reduce scan time. These synthesized phase-cycles are shown to be effective when used in conjunction with the bSSFP elliptical signal model.

Deep Learning Super-FOV for Accelerated bSSFP Banding Reduction [13]

Appendix D contains an accepted proceeding to the 2019 ISMRM annual conference. We present a machine learning technique for bSSFP band removal using two undersampled phase-cycled bSSFP image acquisitions. A deep convolutional neural network was trained to solve a generalized SENSE reconstruction problem where bSSFP banding sensitivities are used as opposed to coil sensitivity maps. We demonstrate that a deep neural network can reduce banding artifacts in multiple acquisition bSSFP with comparable performance to the elliptical signal model and reduce overall scan time by requiring half as many phase-cycled images.

Generation of Arbitrary Spectral Profiles Using Orthonormal Basis Combinations of bSSFP MRI [14]

Appendix E contains an accepted proceeding to the 2018 ISMRM annual conference on which I was the second author. We present a technique for generating an arbitrary spectral profile by using multiple-acquisition bSSFP. Multiple phase-cycled bSSFP images with increasing TRs were acquired and Gram-Schmidt orthogonalization was applied to spectral basis functions to generate an orthonormal basis. We demonstrate simple proof-of-concept spectrum filtering.

CHAPTER 2. PRINCIPLES OF MAGNETIC RESONANCE IMAGING

In this chapter we will outline the basic principles of magnetic resonance (MR) imaging as they apply to the present thesis. We will not shy away from simple classical derivations, but necessarily this discussion will be abbreviated. The exceptionally curious reader is directed to the excellent references: [15], [16]. Our discussion of MR physics below will focus on a classical understanding of MR. For a more nuanced presentation including the underlying quantum mechanical effects governing the nuclear magnetic resonance (NMR) phenomenon, Levitt’s wonderful book is recommended: [17].

2.1 Physics of MRI

Unless otherwise stated, we will assume a constant magnetic field in the \hat{k} direction for the external main magnetic field, i.e., $\vec{B} = B_0\hat{k}$.

2.1.1 Nuclear Spin and Magnetic Moments

Because hydrogen is abundant in humans and we often want to look inside humans, we are interested in single protons. Protons (as well as electrons and neutrons) possess intrinsic angular momentum called “spin.” In fact, nuclei with an odd number of protons and neutrons have an unpaired component of spin angular momentum resulting in a net nonzero value. This property can be measured and quantified, its existence first inferred by the results of the Stern-Gerlach experiment [18]. Enigmatically, a stationary proton’s spin gives rise to a magnetic dipole moment, usually associated with a spinning or rotating charge. Thus it is useful (though not entirely accurate) to think of a proton as a spinning charge generating a current, which, in turn, produces a magnetic moment.

Sticking to our classical model, we here introduce the concept of an Amperian loop. After Hans Christian Ørsted discovered that electric currents induce magnetic fields and Andre-Marie




Figure 2.1: Amperian loop: a point charge with charge e travels in circular loop with radius r and velocity \vec{v} .

Ampere found that electric currents, like magnets, attract and repel each other, Ampere proposed that all magnetic fields are due to electric current loops [19]. In this model, the magnetic dipole resulting from nuclear spin is modeled as a sufficiently small Amperian loop with current I . Thus we consider a point charge with charge e and set it in motion with velocity v around a circle with radius r and area S as in Figure 2.1. This induces a magnetic dipole moment $\vec{\mu}$:

$$\vec{\mu} = I\vec{S}, \quad (2.1)$$

where the direction of \vec{S} is consistent with the right hand rule. Recall that current is simply charge per unit time:

$$I = \frac{\text{charge}}{\text{time}} = \frac{e}{t}. \quad (2.2)$$

We know the total time the point charge takes to make a complete orbit:

$$t = \frac{\text{distance}}{\text{velocity}} = \frac{2\pi r}{v}. \quad (2.3)$$

Putting these all together, we get an expression for the moment:

$$I = \frac{e}{\frac{2\pi r}{v}} = \frac{ev}{2\pi r} \quad (2.4)$$

$$\mu = IS = \frac{ev}{2\pi r} \pi r^2 = \frac{e}{2m} mvr. \quad (2.5)$$

We will call attention to the quantity $\gamma = \frac{e}{2m}$. This is the gyromagnetic ratio (also called the magnetogyric ratio), a fundamental property of the nucleus. We want the gyromagnetic ratio in terms of the so called nuclear magneton μ_N corresponding to a proton with g-factor $g \approx 5.6$ as

found in standard tables [20]:

$$\gamma = \frac{g\mu_N}{\hbar} \quad (2.6)$$

$$\mu_N = \frac{e\hbar}{2m_p}. \quad (2.7)$$

Importantly, we get for a proton:

$$\frac{\gamma}{2\pi} = 42.58 \frac{\text{MHz}}{\text{T}}. \quad (2.8)$$

Looking back at equation (2.5), we also find $L = mvr$ is the apparent angular momentum. Thus we have the simple relationship for the dipole:

$$\vec{\mu} = \gamma\vec{L}. \quad (2.9)$$

Surprisingly enough, this 18th century classical model is consistent with modern observations of dipole and angular momentum measurements.

2.1.2 Larmor Frequency

This section derives the famous Larmor frequency, but to get there, we must first briefly touch on the Bloch equations. The Bloch equations were derived by Bloch himself in his 1946 seminal paper [21]. They describe the dynamics of a magnetic moment in response to external magnetic fields and continue to be foundational in modern MR analyses and reconstruction techniques [19]. In a uniform magnetic field, \vec{B} , a magnetic dipole $\vec{\mu}$ will experience some torque $\vec{\tau}$ as represented by the cross product:

$$\vec{\tau} = \vec{\mu} \times \vec{B}. \quad (2.10)$$

Analogous to Newton's second law, we also have a simple relationship between angular momentum and torque:

$$\frac{d\vec{L}}{dt} = \vec{\tau}. \quad (2.11)$$

We can now derive an expression for $\frac{d\vec{\mu}}{dt}$:

$$\gamma \frac{d\vec{L}}{dt} = \frac{d\gamma\vec{L}}{dt} = \frac{d\vec{\mu}}{dt} = \gamma\vec{\mu} \times \vec{B}. \quad (2.12)$$

This last result informs us that $\vec{\mu}$ will precess, or nutate, when the cross product is non-zero. We are interested in how fast the moment is precessing or spinning due to the torque, i.e., we want to find the precession angular velocity. This precession frequency is ω_{Larmor} , also known as the Larmor frequency. By simply noting that ω_{Larmor} is the time derivative of ϕ (the angle traveled in the direction of \vec{v} along the Amperian loop) we get the following:

$$\tau = \frac{dL}{dt} \quad (2.13)$$

$$= \frac{L \sin \theta d\phi}{dt} \quad (2.14)$$

$$= |\vec{\mu} \times \vec{B}| \quad (2.15)$$

$$= |\mu B \sin \theta| \quad (2.16)$$

$$= \gamma L B \sin \theta. \quad (2.17)$$

Then, after a bit of algebra, we get the desired expression for the Larmor frequency:

$$\omega_{\text{Larmor}} = \frac{d\phi}{dt} = \gamma B. \quad (2.18)$$

We will be returning to the Bloch equations to revise them in section 2.1.6. For now, we have shown that $\vec{\mu}$ is acted on by torque $\vec{\tau}$ due to the external magnetic field \vec{B} resulting in precessing at the Larmor frequency.

2.1.3 Net Magnetization

In tissue, we are almost always dealing with a large number of nuclei that give rise to a large number of magnetic moments. If a strong static external field is applied to this large number of nuclei, they will align either parallel or anti-parallel to the field, corresponding to low and high energy states, respectively. The parallel orientation is slightly preferred over the anti-parallel one as it is the lower energy state, however thermal energy is significant enough to keep

the two populations nearly balanced. The small majority in the parallel orientation leads to a net magnetization in the direction of the applied static field. This polarization in presence of an external magnetic field is depicted in Figure 2.2. Since we can practically only deal with huge numbers of nuclei at a time, we consider only the net magnetization, \vec{M} , or the sum of all moments in a given volume (usually a voxel):

$$\vec{M} = \sum_{\text{volume}} \vec{\mu}. \quad (2.19)$$

Most literature will refer to these net magnetization vectors as “spins” as we can visualize it precessing in a stationary frame according to equation (2.20). We will also refer to these as “magnetization vectors.” By linearity, these spins behave for the most part as if they were a large magnetic dipole moment in the magnetic field \vec{B} :

$$\frac{d\sum \vec{\mu}}{dt} = \frac{d\vec{M}}{dt} = \gamma \vec{M} \times \vec{B}. \quad (2.20)$$

At this point we abandon the notion of dipole moments and substitute spins as the underlying MR signal of interest.




Figure 2.2: Red arrows represent that magnetic moments of hydrogen nuclei. On the left, when no magnetic field is applied, the moments are randomly oriented and the net magnetization of large groups of moments is zero. However, as depicted on the right, in the presence of a magnetic field, the moments tend to align parallel or anti-parallel to the field. Slightly more will align with the field leading to a non-zero net magnetization in the direction of the field.

2.1.4 Excitation

From equation (2.20), it is clear that when a spin has a nonzero angle with an external magnetic field \vec{B} , the spin will precess at the Larmor frequency given by equation (2.18). In fact, this is what we mean when we say \vec{M} exhibits nuclear magnetic resonance: when \vec{M} precesses, a circularly polarized radio-frequency (RF) wave is emitted at the Larmor frequency. Conversely, when \vec{M} is aligned with \vec{B} , there is no precession and \vec{M} is said to be at thermal equilibrium.

Dealing with RF, it is natural to expect reciprocity of the resonance: when an RF wave at the Larmor frequency with nonzero component orthogonal to \vec{B} is applied to \vec{M} , the spin is “knocked” out of alignment. This is called excitation. The RF wave itself is more commonly referred to as an RF pulse in MR pulse sequence design. Assume for convenience that the applied RF pulse $\vec{B}_1(t)$ is completely orthogonal to \vec{B} , that is, $\vec{B}_1(t) = B_{1,x}\hat{i} + B_{1,y}\hat{j}$. After the RF pulse has been applied to \vec{M} for a time T , we say that \vec{M} has been tipped into the transverse plane (which by convention is the xy -plane, as $\vec{B} = B_0\hat{k}$ is in the longitudinal direction). The angle that \vec{M} has been tipped is called the flip angle α and is given by the following relationship:

$$\alpha = \int_0^T \gamma \vec{B}_1(t) dt. \quad (2.21)$$

Now that \vec{M} has been tipped, it has a nonzero angle with \vec{B} and a receiver coil tuned to the Larmor frequency may be used to record the wave generated by the transverse magnetization, M_{xy} , of the resulting precession. This is the MR signal. Notice that the larger M_{xy} is, the stronger the resulting signal will be. As the strength of the signal that can be detected is proportional to the transverse component of the magnetization, maximum signal is obtained by using $\alpha = 90^\circ$ to place all signal in the transverse plane.

2.1.5 Relaxation

Once an RF pulse has been applied to \vec{M} , contrary to equation (2.20), it does not precess forever. In fact, in the absence of an RF pulse, \vec{M} will return to thermal equilibrium at a rate which is characteristic of the tissue that was excited. The process of returning to thermal equilibrium is called relaxation. There are two components of relaxation: longitudinal relaxation and transverse relaxation. Longitudinal relaxation is also known as spin-lattice relaxation and is described by the

equation:

$$\frac{\partial M_z}{\partial t} = \frac{-(M_z - M_0)\hat{k}}{T_1}, \quad (2.22)$$

where M_z is the longitudinal magnetization and M_0 is the magnitude of the magnetization at thermal equilibrium. Note that M_z cannot be directly measured using commercial MR scanners. We call T_1 the longitudinal relaxation constant. The differential equation (2.22) has the following solution:

$$M_z(t) = M_0 + (M_z(0) - M_0)e^{-t/T_1}, \quad (2.23)$$

where the time elapsed, t , is set to 0 when relaxation begins.

Transverse relaxation or spin-spin relaxation is governed by the following differential equation:

$$\frac{\partial M_{xy}}{\partial t} = \frac{-M_{xy}}{T_2}, \quad (2.24)$$

where M_{xy} is the transverse magnetization and T_2 is called the transverse relaxation constant. Again, there is a simple solution:

$$M_{xy}(t) = M_{xy}(0)e^{-t/T_2}. \quad (2.25)$$

Figure 2.3 shows representative relaxation curves for T_1 and T_2 along with an illustration of the magnetization vector \vec{M} . We should also note that T_2 is (almost always) less than T_1 . T_1 and T_2 are tissue specific and vary with field strength. For example, Table 2.1 gives values for common organic tissues at 1.5T (see Appendix II of [22]) We now have a mathematical model of how relaxation affects the dynamics of the magnetization vector \vec{M} and we are now ready to present the Bloch equations incorporating relaxation effects.

2.1.6 Bloch Equations

We introduce transverse (T_2) and longitudinal (T_1) exponential relaxation time constants to describe how the net magnetization reaches thermal equilibrium when perturbed. Assuming an externally applied static magnetic field $\vec{B} = B_0\hat{k}$, we have the modified equation (2.20), known as

Table 2.1: Longitudinal and transverse relaxation constants for some common organic tissues at 1.5T.

Tissue	T_1 (in ms)	T_2 (in ms)
Brain white matter	790	90
Brain grey matter	920	100
Cerebral-spinal fluid	≥ 4000	~ 2000
Myocardium	870	60
Skeletal muscle	870	50
Lipids	260	80

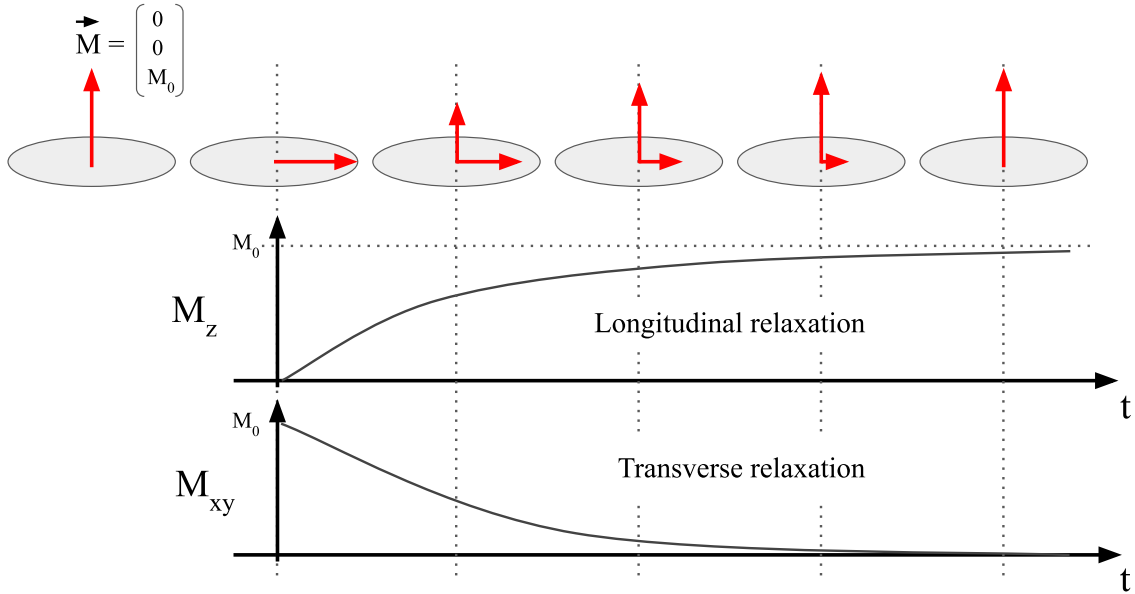


Figure 2.3: The longitudinal and transverse components of the magnetization vector \vec{M} are governed by equations (2.23) and (2.25), respectively. These exponential curves are illustrated here with the inclusion of a visualization of the longitudinal and transverse components of \vec{M} . At $t = 0$ the magnetization vector is tipped fully into the transverse plane followed by relaxation back to thermal equilibrium.

the Bloch equations, describing the dynamics of \vec{M} in an external magnetic field:

$$\frac{d\vec{M}}{dt} = \gamma \vec{M} \times B_0 \hat{k} - \frac{M_x \hat{i} + M_y \hat{j}}{T_2} - \frac{(M_z - M_0) \hat{k}}{T_1}. \quad (2.26)$$

Note that the unit vectors $(\hat{i}, \hat{j}, \hat{k})$ are used to represent the three coupled partial differential equations present in the above equation. The Bloch equations are often written in matrix form:

$$\frac{d}{dt} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} -\frac{1}{T_2} & \gamma B_z & -\gamma B_y \\ -\gamma B_z & -\frac{1}{T_2} & \gamma B_x \\ \gamma B_y & -\gamma B_x & -\frac{1}{T_1} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{M_0}{T_1} \end{bmatrix}, \quad (2.27)$$

with obvious simplifications when $\vec{B} = B_z \hat{k}$. The MR signal M_{xy} is the result of projecting \vec{M} onto the transverse plane.

2.2 Spatial Localization

An essential concept in MR imaging is that the Larmor frequency, ω_{Larmor} , can be made to be position-dependent. Imagine a scenario where we place a person in a magnetic field \vec{B} and an RF pulse is applied. Every spin in the body would be excited regardless of its spatial location. When we receive the resulting MR signal, the only information we have is a relative measure of how much hydrogen was in their body. This is not very interesting information. If the Larmor frequency was position-dependent, say, a unique frequency corresponded to the leg, then when we apply the RF pulse we would only excite and receive signals from spins in the leg. With more precise control of the Larmor frequency as a function of position, more interesting detail can be acquired from the MR signal. According to equation (2.18), the Larmor frequency depends on magnetic field strength, i.e., $\omega_{\text{Larmor}} \propto B_0$. That suggests that if we can make the external magnetic field position-dependent, then the Larmor frequency will be as well.

In an MR scanner, the main field is maintained by a superconducting magnet designed to produce a very uniform field (usually within a few parts per million) across an imaging region, the center of which is referred to as the isocenter. Position-dependence is accomplished with extra conventional electromagnets, often just called “gradients” and denoted G . MR scanners are usually equipped with three linear gradients: G_x , G_y , and G_z , one along each axis. “Linear” refers to the fact that the amplitude of the gradient varies linearly with position. The combined gradient fields

can be expressed as a time-varying function:

$$\vec{G}(t) = G_x(t)\hat{i} + G_y(t)\hat{j} + G_z(t)\hat{k}, \quad (2.28)$$

and the total field is given by:

$$B(\mathbf{r}, t) = B_0 + \vec{G}(t) \cdot \mathbf{r}, \quad (2.29)$$

where \mathbf{r} is the position vector or current (x, y, z) spatial coordinates. Then by manipulating linear gradients we can easily change the Larmor frequency as a function of location. Below we will discuss common strategies for spatial localization of the MR signal using the linear gradient magnets.

2.2.1 Slice-Selective Excitation

Slice-selective excitation uses a band-limited RF pulse and the G_z gradient to select and excite just a single slice of the imaging volume. Application of the G_z gradient leads to spatially varying Larmor frequency along the z -axis as described above. Applying an RF pulse carrying power only in those frequencies corresponding to a narrow band around $\omega_0 = \gamma B_0$ leads to only nuclei lying in the slice at those frequencies being selected and excited. This is illustrated in Figure 2.4. This is in contrast to non-selective excitation where everything in the imaging volume receives RF power.

2.2.2 Frequency and Phase Encoding

The signal from a single slice is still not very informative. Frequency and phase encoding allow us to resolve signals coming from individual voxels, or three-dimensional volumes from within the object being imaged. Phase encoding relies on another orthogonal gradient, for example the G_y gradient. This gradient is applied for a time T before readout, during which it is switched off. This means that nuclei along the y -axis have precessed either more quickly or slowly than $\omega_0 = \gamma B_0$ when the gradient was turned on. Once the G_y gradient is switched off, all nuclei resonate at ω_0 proper, but now the relative phases between them have changed as a function of




Figure 2.4: We often want to excite spins in just a slice or slab only, rather than exciting the entire volume. Application of the linear gradient field G_z varies the resonant frequencies as a function of position along the z -axis. When a band-limited RF pulse is applied in conjunction with this gradient, only spins resonating at frequencies within the bandwidth of the RF pulse will be excited, “selecting” the slice as shown.

position along the y -axis. Thus position in y has been encoded in the phase of MR signals emitted from each nuclei.

Frequency encoding uses another orthogonal gradient, for example, G_x , to encode spatial information during readout. The G_x gradient is played during signal acquisition, meaning that in the x direction, nuclei are precessing at different frequencies. In fact, after acquisition we can decompose the data into its frequency components and map the frequency bins directly to spatial locations. Thus using both phase and frequency encoding, the data are sums of sinusoids at different frequencies and phases corresponding to spatial location. Chapter 3 includes a pulse sequence diagram detailing how G_x , G_y , and G_z are applied to achieve slice-selection and frequency/phase encoding.

2.2.3 Signal Equation

The discussion of spatial localization leads us to the realization that the data being received are sums of sinusoids. We can easily represent these sinusoids as complex exponentials spinning in the complex plane:

$$M_{xy}(\mathbf{r})e^{-j\phi(\mathbf{r},t)}, \quad (2.30)$$

where $M_{xy}(\mathbf{r})$ is the transverse magnetization at position \mathbf{r} . Assuming a uniformly sensitive RF receive coil and negligible T_2 decay effects, the received signal, $s(t)$, is given by integrating over the volume, V :

$$s(t) = \int_{\mathbf{r} \in V} M_{x,y}(\mathbf{r})e^{-j\phi(\mathbf{r},t)} d\mathbf{r}. \quad (2.31)$$

Instantaneous frequency is $\omega = \frac{d\phi}{dt}$, so we can integrate $\omega = \gamma B$ to find ϕ :

$$\phi(\mathbf{r},t) = \int_0^t \omega(\mathbf{r},t') dt' = \gamma \int_0^t B(\mathbf{r},t') dt'. \quad (2.32)$$

In the presence of linear gradients we have the combined magnetic field $B(\mathbf{r},t)$, given by equation (2.29). Using these results, we find $\tilde{s}_r(t)$, the received modulated signal equation:

$$\tilde{s}_r(t) = e^{-j\gamma B_0 t} \int_{\mathbf{r} \in V} M_{xy} e^{-j\gamma \int_0^t \vec{G}(t') \cdot \mathbf{r} dt'} d\mathbf{r}. \quad (2.33)$$

We then demodulate this signal by multiplying by the conjugate complex exponential at the Larmor frequency to get $s_r(t)$, the demodulated received MR signal as a function of time:

$$s_r(t) = e^{j\gamma B_0 t} \tilde{s}_r(t) = \int_{\mathbf{r} \in V} M_{xy} e^{-j\gamma \int_0^t \vec{G}(t') \cdot \mathbf{r} dt'} d\mathbf{r}. \quad (2.34)$$

2.2.4 k -space

Consider the Fourier transform of M_{xy} :

$$F[M_{xy}] = \mathcal{M}(k_x, k_y, k_z), \quad (2.35)$$

$$\mathcal{M}(k_x, k_y, k_z) = \int_{x,y,z} M_{xy} e^{-j2\pi(k_x x + k_y y + k_z z)} dx dy dz, \quad (2.36)$$

and compare it to the expression obtained for $s_r(t)$ in equation (2.34). Comparing terms, we find:

$$s_r(t) = \mathcal{M}\left(\frac{\gamma}{2\pi} \int_0^t G_x(t') dt', \frac{\gamma}{2\pi} \int_0^t G_y(t') dt', \frac{\gamma}{2\pi} \int_0^t G_z(t') dt'\right). \quad (2.37)$$

Beautifully, each measurement acquired using our gradients corresponds to filling k -space at some position (k_x, k_y, k_z) . This not only suggests how we may use our linear gradients to “drive around” in k -space, but it tells us that we may simply perform an inverse Fourier transform of our received signal to recover M_{xy} if we have acquired the data on a Cartesian grid and have satisfied the Nyquist sampling criterion. In fact, turning on the G_x gradient corresponds to traveling in the k_x direction in k -space and similarly for G_y and G_z . Higher gradient amplitude corresponds to traveling more quickly through k -space. Then by manipulating our gradients, we move in arbitrary trajectories around k -space.

Without any real surprise, we note that the patterns that we sample k -space with are called “ k -space trajectories.” One of the most common trajectories is Cartesian, where we sample k -space line-by-line, after which the inverse Fourier transform may be directly applied as described above. If more exotic trajectories are used, such as radial, spiral, or Poisson, more sophisticated reconstruction techniques must be used. Often the strategy is to resample or regrid the data onto a Cartesian grid so we can apply the FFT algorithm. Chapters 5 and 6 will explore reconstruction methods for specialized trajectories meant to save time by sampling below the Nyquist rate.

2.3 Real World Imaging

The development of MR imaging thus far has expected a lot out of the world: perfectly homogeneous magnetic fields, flawless gradient performance, sharp RF power spectra, perfectly

homogeneous tissue samples, no motion, unitary receive coil sensitivities, etc. These are all, of course, not the case when acquiring data on the job. This section will briefly detail some confounders while admitting that this thesis does not have enough room to tackle every struggle with non-ideal imaging conditions.

2.3.1 Magnetic Field Inhomogeneity

The condition of a homogeneous main magnetic field is very important for spatial localization and signal generation as it can influence the Larmor frequency unpredictably. Unfortunately, inhomogeneities inevitably exist across the imaging volume. This is due to hardware imperfections, magnetic susceptibility interfaces that introduce complicated variations to the magnetic field, and even subtle changes like temperature drift or slight motion by the subject inside the magnet. Some large imperfections may be unavoidable, such as when imaging a subject with metallic implants. In order to achieve as homogeneous a magnetic field as possible, extra “shim” coils are built into MR scanners which can be tuned to counteract field inhomogeneities. This becomes even more difficult at higher field strengths as the magnitude of susceptibility mismatch artifacts grows with field strength. Correcting for these imperfections is a sub-field in its own right; [23] provides a good review of current correction strategies.

2.3.2 Chemical Shift

Chemical shift may be understood as interactions between electron clouds and B_0 . The errors introduced by chemical shift are due to a shift in resonant frequency induced by bonding electrons. The bonding electrons are diamagnetic and are magnetized in the presence of a magnetic field. This means that using the same magnetic field, some nuclei in molecules of type A experience potentially different magnetic fields than nuclei in a molecule of type B. For example, water molecules are made up of two hydrogen and one oxygen and experience little electronic shielding. In contrast, fat molecules’ electronic interactions cause the protons to feel a weaker magnetic field, causing the resonant frequency to drop by about 3.5 parts per million [16]. Because of this, care must be taken when imaging regions containing both water and fat.

2.3.3 Noise

The MR signal we measure is a voltage induced in one or more coils. In fact, quadrature detection of the MR signal yields both magnitude and phase information in the form of real and imaginary channels. As such, we must expect some amount of electrical noise. However, usually the dominant source of noise in an MR experiment is that which originates from the subject, unless poor coils are used or imaging is done at low field strengths. Brownian motion of electrons and random motion of ions in the fluids of the subject being imaged contribute to body noise. The aggregate noise is well-modeled by a bivariate (complex) zero-mean Gaussian process with a uniform power spectral density (i.e., white noise) [16].

CHAPTER 3. BALANCED STEADY-STATE FREE PRECESSION

3.1 Introduction

Due to its high speed, efficient SNR, and unique contrast, balanced steady-state free precession (bSSFP) has become a popular MR imaging sequence with desirable contrast for applications such as cardiac imaging, abdominal imaging, angiography, functional MRI, and musculoskeletal imaging [24]–[28]. As Chapter 4 relies heavily on bSSFP imaging, we will briefly review the pulse sequence design and signal models that characterize this important steady-state sequence.

3.2 Pulse Sequence

The bSSFP sequence was first described by Carr in his 1958 NMR paper [29]. The sequence itself is fundamentally very simple: a rapid train of identical excitation pulses in the presence of “balanced” gradients that induce no net phase to the magnetization by the end of each repetition. The time from one RF pulse to the next is called the repetition time, or TR. Figure 3.1 shows a typical pulse sequence diagram for a 2D bSSFP sequence. The diagram describes the timing of the RF and gradient pulses used. Notice that the “echo time” is the time from the RF pulse to the center of the gradient being played during data acquisition, or the “readout gradient.” The echo time is notated “TE” and, for a bSSFP experiment, is always set to $TE = TR/2$.

BSSFP is also sometimes called “fully-refocused” SSFP, calling attention to the fact that the gradient waveforms are balanced, or refocused, such that the area of each gradient waveform integrates to zero across any given TR. We also say that the zeroth moments of the gradient fields are nulled. This implies that the gradients have no net effect on spin precession from one excitation to the next: any phase accumulation due to the gradients across a TR is undone prior to subsequent excitation.




Figure 3.1: A typical bSSFP pulse sequence diagram. Depicted are four waveforms: one for each the RF transmitter and linear gradient fields. The RF is played repeatedly with a period TR. Notice that the gradients’ waveforms integrate to zero, eliminating coherences TR to TR.

Repetition time, TR, is usually chosen to be short; that is, the repeated excitation pulses tip the spins at regular, short intervals. Between excitations, the spins freely precess – hence, “free precession.” But as the spins decay, they are re-excited quickly enough that there is very little relaxation, and after a short time the spins reach steady-state. The time required to reach steady-state is usually several times the longest T_1 present in the imaging sample, but there are many methods designed to reduce the transient response and “catalyze” the steady-state, for example, see [30]–[33].

3.3 Off-resonance and the bSSFP Spectrum

The bSSFP signal has well-known off-resonance dependence (see Figure 3.2) which is perhaps one of the most interesting and problematic properties of the sequence. A plot of signal magnitude vs off-resonance is called the bSSFP spectrum or profile and depicts how off-resonant precession affects the measured transverse magnetization. The shape of the spectral profile is dominated by off-resonance frequency and flip angle rather than T_1 , T_2 , or other properties that

commonly dictate MR contrast. It must be noted that bSSFP also yields T_2/T_1 contrast in common conditions [27].




Figure 3.2: Magnitude and phase components of a simulated bSSFP spectral profile using $TR = 12\text{ms}$, $\alpha = 30^\circ$, $T_1 = 830\text{ms}$, and $T_2 = 80\text{ms}$. The horizontal axis measures off-resonance ranging between $\pm 1/TR$ Hz. Two phase components are shown. The phase predicted directly after excitation is different than when measured at $TE = TR/2$ as the magnetization vector at $TE/2$ has already gone through half of its total precession angle. Its not necessary to do so, but we can compensate for this by subtracting half of the phase-cycle value from the phase at $t = TE/2$ to recover the phase at $t = 0$ [10].

An important property of RF pulse design for bSSFP imaging is phase-cycling [34]. BSSFP spectra are periodic in off-resonance (with a period of $2/TR$) and typically have a pass-band region where the profile is relatively constant and a transition band region where signal nulls are introduced. These nulls are often referred to “banding artifacts.” Interestingly, by linearly varying the phase of the RF pulse each TR (i.e., adding a constant phase increment each excitation) the bSSFP spectra are shifted along the off-resonance axis by a corresponding amount. There are many interesting applications of this so-called phase-cycled bSSFP, one of which will be explored in Chapter 4.

3.4 Signal Model

Here we derive the bSSFP steady-state signal model. This derivation closely follows that of Freeman and Hill and other similar discrete-time system derivations due to Jaynes [32], [35], [36].

We start by noting that each pulse tips the spin vectors through a tip angle α . Assume for this discussion that the flip angle rotates the spins about the x -axis. Thus we can find that the magnetization directly after the pulse, \vec{M}^+ , is related to the magnetization directly before the pulse, \vec{M}^- , by the rotation operator:

$$\vec{M}^+ = R_x(\alpha)\vec{M}^-. \quad (3.1)$$

Between excitations and during free precession, the spins precess through an angle θ given by their off-resonance, Δf_0 (in Hz), and phase-cycle increment, $\Delta\theta$ (in radians):

$$\theta = 2\pi\text{TR} + \Delta\theta. \quad (3.2)$$

θ can also take into account non-idealities such as chemical shift, but for simplicity we will consider these negligible. During this same period, relaxation also occurs according to the dynamics of the Bloch equations. Longitudinal and transverse relaxation are applied using the following relaxation operator:

$$S(\text{TR}, T_1, T_2) = \begin{bmatrix} E_2 & 0 & 0 \\ 0 & E_2 & 0 \\ 0 & 0 & E_1 \end{bmatrix}, \quad (3.3)$$

where $E_i = e^{-\text{TR}/T_i}$. Realizing that precession is a rotation in z , we have the following magnetization, \vec{M}^\dagger , at the end of the free precession period:

$$\vec{M}^\dagger = R_z(\theta)S(\text{TR}, T_1, T_2)\vec{M}^+ + (1 - E_1)M_0, \quad (3.4)$$

where M_0 is the magnetization at thermal equilibrium. For the steady-state expression, we know $\vec{M}^\dagger = \vec{M}^-$:

$$\vec{M}^\dagger = R_z(\theta)S(\text{TR}, T_1, T_2)R_x(\alpha)\vec{M}^\dagger + (1 - E_1)M_0 \quad (3.5)$$

$$= \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} E_2 & 0 & 0 \\ 0 & E_2 & 0 \\ 0 & 0 & E_1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & \sin \alpha \\ 0 & -\sin \alpha & \cos \alpha \end{bmatrix} \vec{M}^\dagger + (1 - E_1)M_0 \quad (3.6)$$

$$= \begin{bmatrix} E_2 \cos \theta & E_2 \cos \alpha \sin \alpha & E_2 \sin \alpha \sin \theta \\ -E_2 \cos \theta & E_2 \cos \alpha \cos \theta & E_2 \cos \theta \sin \alpha \\ 0 & -E_1 \sin \alpha & E_1 \cos \alpha \end{bmatrix} \vec{M}^\dagger + (1 - E_1)M_0. \quad (3.7)$$

Then, solving for the transverse components directly after exciting, we arrive at the familiar bSSFP signal equations:

$$M_x^+ = \frac{M_0(1 - E_1)E_2 \sin \alpha \cos \theta}{(1 - E_1 \cos \alpha)(1 - E_2 \cos \theta) - E_2(E_1 - \cos \alpha)(E_2 - \cos \theta)}, \quad (3.8)$$

$$M_y^+ = \frac{M_0(1 - E_1) \sin \alpha (1 - E_2 \cos \theta)}{(1 - E_1 \cos \alpha)(1 - E_2 \cos \theta) - E_2(E_1 - \cos \alpha)(E_2 - \cos \theta)}. \quad (3.9)$$

3.5 Elliptical Signal Model

Equations (3.8) and (3.9) appear to be a complicated function of T_1 , T_2 , α , and θ . However, when we plot the complex valued $I = M_x + jM_y$ in the complex plane, we find that locus of the magnetization vectors simply traces out an ellipse. In fact it has a well known parameterization given in [37] and summarized in [1]:

$$M = \frac{M_0(1 - E_1) \sin \alpha}{1 - E_1 \cos \alpha - E_2^2(E_1 - \cos \alpha)}, \quad (3.10)$$

$$a = E_2, \quad (3.11)$$

$$b = \frac{E_2(1 - E_1)(1 + \cos \alpha)}{1 - E_1 \cos \alpha - E_2^2(E_1 - \cos \alpha)}, \quad (3.12)$$

$$I = M \frac{1 - ae^{j\theta}}{1 - b \cos \theta}. \quad (3.13)$$

Thus we have an ellipse I parameterized by the θ -independent quantities M , a , and b . This formulation is known as the elliptical signal model for bSSFP and gives useful insight on multiple phase-cycled bSSFP image combination and parameter estimation [1], [2].

CHAPTER 4. ROBUST COIL COMBINATION FOR THE BSSFP ELLIPTICAL SIGNAL MODEL

4.1 Introduction

Since the inception of magnetic resonance imaging (MRI), improvements to both scanner hardware and software have significantly improved imaging time and quality. Among these improvements is the receive array coil, used to acquire signals simultaneously from several locations in order to provide greater coverage and higher signal to noise ratio (SNR) than an equivalently sized single surface coil [11]. Phased-array coils have also allowed the implementation of parallel imaging (PI) techniques to reduce the required scan time [38], [39]. Combined, these advances have revolutionized image acquisition techniques while also posing new challenges to clinicians and researchers.

Today, with the large number of coil channels in use, large amounts of raw data need to be acquired and processed when reconstructing images. In order to reduce the complexity and time required for computation, the amount of data is usually compressed by combining coil channel information using one of a variety of linear or nonlinear combination methods, generally with negligible effects on image quality [6], [7], [11], [40].

Balanced steady-state free precession (bSSFP) acquisition is a fast, SNR-efficient sequence that suffers from dark banding artifacts due to its off-resonance dependence [29]. These banding artifacts can be mitigated in a number of ways, usually by acquiring multiple linear phase-cycled images and constructively combining them [1], [34], [41]. The elliptical signal model (ESM) is a parametric representation of the bSSFP signal equation where phase-cycles correspond to points along an ellipse in the complex plane defined by tissue parameters (i.e., T_1 , T_2 , proton density, etc.) [1], [42]. Chemical shift, off-resonance, and RF phase correspond to rotations of the ellipse in the complex plane about the origin. The linearized geometric solution (IGS) was presented by Xiang et al. as an analytic solution for the single de-banded image given four appropriately phase-

cycled images [1]. Although first proposed for fast imaging near metallic implants, the IGS is useful for imaging at high fields where low TRs are not feasible, leading to more frequent banding, and where off-resonance is more difficult to correct [43].

The IGS is only defined for single coil data sets, requiring large data sets with many coil channels to be solved independently before coil combination at the end of the reconstruction pipeline [44]. This may be prohibitively long for clinical applications, such as fast imaging near metallic implants or at high field strengths. In order to reduce reconstruction times, coil combination prior to the IGS reconstruction would be desirable, but unfortunately due to the nonlinear nature of the IGS and nonlinearities introduced by some coil combination methods, the results produce extremely poor quality images with large artifacts [5]. These artifacts arise largely due to distortion of the required elliptical relationships between pixels of phase-cycled images. Coil compression, in general, succeeds on each phase-cycled image individually, creating artifact-free bSSFP-contrast phase-cycled images. However, the elliptical relationships between the pixels of each combined phase-cycled image may not be preserved, violating the IGS’s assumptions. This is especially true when coil combination methods apply phase alignment or other techniques to produce smoothly varying coil sensitivity maps [6], [45], [46]. Note that coil compression in conjunction with phase-sensitive auto-calibrating PI techniques will succeed on non-ESM datasets, as only one set of virtual coil images are produced: there is in general no relationship between multiple sets of virtual coils to maintain [6], [39], [47]. Application of auto-calibrating PI methods to phase-cycled bSSFP datasets precludes upstream coil combination as elliptical relationships between the interpolated datasets are crucial for downstream processing.

Preserving accurate phase information during coil combination has been much studied in fields including susceptibility weighted imaging, quantitative susceptibility mapping, MR thermometry, and phase contrast imaging methods [48]–[54]. Some of these techniques, while effective, are unsuited for use with multiple acquisition phase-cycled bSSFP. Multi-echo techniques including COMPOSER and ASPIRE require additional images acquired at short TEs, multiplying the number of images required for use with the IGS [54], [55]. Virtual body coil and virtual reference coil techniques break down in areas of low SNR, such as around banding artifacts in bSSFP images, and for larger fields of view at high magnetic field strengths, e.g., at 7T [4], [50], [53]. They also do not guarantee compatible reference phases that preserve the elliptical relation-

ship between the combined phase-cycled images. We do not consider methods with explicit prior knowledge of the complex coil sensitives in this paper and restrict our attention to maximum coil compression, that is, a single image from multiple coil images.

In this work, we will describe the performance of several well-known coil combination methods to describe their performance when applied prior to the IGS. We will also present a phase substitution scheme that successfully mitigates these artifacts.

4.2 Theory

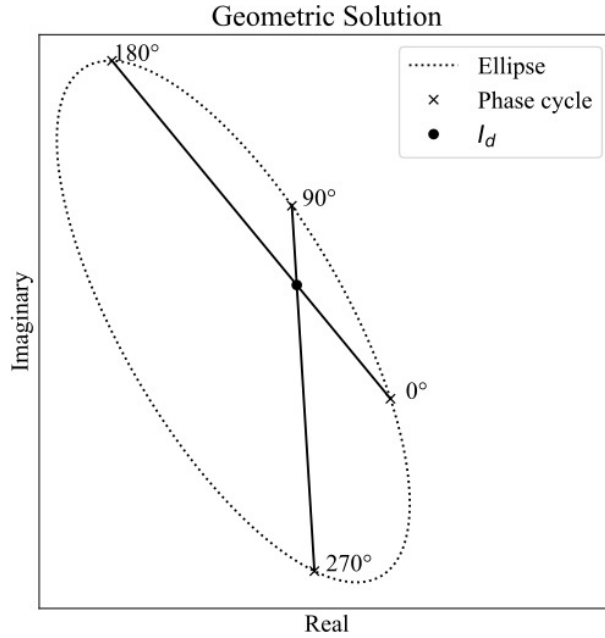


Figure 4.1: The geometric solution to the bSSFP elliptical signal model describes the cross point, I_d , of two pairs of 180° phase-cycled pixels plotted in the complex plane as the demodulated value.

Given phase-cycled images I_i , $0 \leq i < 4$, the direct geometric solution to the ESM is the result of a nonlinear function given pixel-wise by:

$$I_d(x,y) = \frac{\text{Im} [M_0^* M_2] (M_1 - M_3) - \text{Im} [M_1^* M_3] (M_0 - M_2)}{\text{Re} [M_0 - M_2] \text{Im} [M_1 - M_3] + \text{Re} [M_1 - M_3] \text{Im} [M_2 - M_0]}, \quad (4.1)$$

where $M_i = I_i(x,y)$ is the complex-valued pixel of the i th phase-cycled image [1]. This can be graphically understood as the cross point between two pairs of 180° phase-cycles as shown in

Figure 4.1. As there are often singularities due to noise in the direct solution, we assume that I_d has been regularized by replacing errant pixels with the complex sum solution. Xiang et al. also propose a second pass solution, recognizing that the single coil demodulated bSSFP signal can be represented as a weighted sum with weights $0 < w_{0,1}(x, y) < 1$ applied to each pixel. The linearized geometric solution is given by:

$$I_{\text{IGS}} = \frac{1}{2} [w_0 I_0 + (1 - w_0) I_2 + w_1 I_1 + (1 - w_1) I_3], \quad (4.2)$$

where $w_{0,1}$ are computed using the nonlinear direct solution, I_d , such that the regional differential energy is minimized:

$$w_0(x, y) = \frac{\sum_{\text{region}} (M_2 - M_d)^* (M_2 - M_0) + (M_2 - M_0)^* (M_2 - M_d)}{2 \sum_{\text{region}} (M_0 - M_2)^* (M_0 - M_2)}, \quad (4.3)$$

and similarly for w_1 within small regions around (x, y) , e.g., a 5×5 pixel patch.

Consider I_{ij} to be the i th phase-cycle acquired by the j th coil using a phased-array receiver with N elements. Then we rewrite the IGS reconstruction with multiple coils:

$$\begin{bmatrix} I_{\text{IGS},0} \\ I_{\text{IGS},1} \\ \vdots \\ I_{\text{IGS},N-1} \end{bmatrix} = \frac{1}{2} \left(w_0 \begin{bmatrix} I_{0,0} \\ I_{0,1} \\ \vdots \\ I_{0,N-1} \end{bmatrix} + (1 - w_0) \begin{bmatrix} I_{2,0} \\ I_{2,1} \\ \vdots \\ I_{2,N-1} \end{bmatrix} + w_1 \begin{bmatrix} I_{1,0} \\ I_{1,1} \\ \vdots \\ I_{1,N-1} \end{bmatrix} + (1 - w_1) \begin{bmatrix} I_{3,0} \\ I_{3,1} \\ \vdots \\ I_{3,N-1} \end{bmatrix} \right), \quad (4.4)$$

which we can write more succinctly, adopting the following matrix notation:

$$I_{\text{IGS}} = \frac{1}{2} [w_0 I_0 + (1 - w_0) I_2 + w_1 I_1 + (1 - w_1) I_3]. \quad (4.5)$$

Notice that each coil image corresponding to the same phase-cycle shares the same weights, as the coil ellipses (the ellipse generated by each coil image) are similar, i.e., scaled and rotated versions of each other (see Figure 4.2). There may also be a translation if any coil element adds a DC offset. The absolute rotation and scaling of the coil combined ellipse is determined by both off-resonance and coil sensitivity. As we only have N measurements, off-resonance and coil-dependent phase

($N + 1$ unknowns) cannot be determined uniquely unless we have accurate coil sensitivity maps or additional phase-cycles (> 4) [2]¹. We also note that the accuracy of the weights are determined by the quality of the corresponding direct solution, I_d , and the properties of the steady-state at which the phase-cycles were acquired. Since we know weights are shared between coils, the optimal strategy for a coil-by-coil IGS is to solve for the weights $w_{0,1}^j$ for the j th coil and then take the variance weighted mean before using them to solve equation (4.5):

$$w_{0,1}(x, y) = \frac{\sum_j w_{0,1}^j(x, y) \sigma_j^{-2}}{\sum_j \sigma_j^{-2}}, \quad (4.6)$$

where σ_j^2 is the j th coil's noise variance estimate. (The coil noise covariance matrix may be substituted if known [11]).

Equation (4.5) describes the IGS being solved N times: once for each coil. Coil combination must then be accomplished after evaluating equations (4.6) and (4.5). We desire a solution that requires the weights to be solved for only once, that is, solve the coil combination problem first as opposed to last. This is desirable for large datasets, as the regional differential energy minimization step includes pixel-wise computations that cannot be vectorized. The challenge is to choose an appropriate coil combination method and choice of weights, $w_{0,1}$.

It is clear that any linear coil combination method² can be used with equation (4.2) using the substitution $\hat{I}_i = I_{i, \text{coil combined}}$. Phase-unwrapping of the phase-cycled images may also be necessary, adding another complication into the image reconstruction pipeline as 4 points are usually not sufficient to reliably unwrap the ellipse [2]. Appropriate weights can be obtained by evaluating equation (4.6) across a reduced number of averages, i.e., evaluate a reduced number of coils. If a nonlinear coil combination method is used, equation (4.2) cannot be evaluated using the substitution $\hat{I}_i = I_{i, \text{coil combined}}$, as the elliptical relationship between the phase-cycled images in the complex plane may no longer hold. Figure 4.2 shows the effect of the nonlinear GSS and linear stochastic matched filter (SMF) coil combinations on the resulting combined coil ellipses [6],

¹The IGS requires only 4 phase-cycled images while an ellipse is fully defined by 5 points. In this paper we refer to the 4 phase-cycle points plotted in the complex plane as ‘‘coil ellipses.’’ Perhaps they are better described as ‘‘coil quadrilaterals.’’ However, because they are points generated from an underlying elliptical model, we will persist in calling them ellipses.

²This is an abuse of terminology. Virtually all coil combination methods are linear combinations of coil images. However, when considering coil combination as an operator, the nonlinearities introduced during coefficient estimation become more clear. A more full discussion is found in Appendix A.

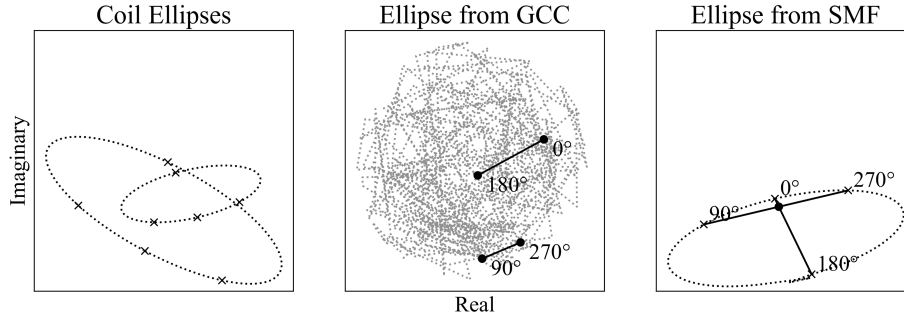


Figure 4.2: Phased-array receivers produce an ellipse per channel in the complex plane. The GCC method is nonlinear and leads to unpredictable combined phase-cycle points. The ellipse is mapped chaotically through the coil combination and the phase-cycle points show that a solution to the IGS does not exist for this pixel, i.e., there is no cross-point. But the SMF is a linear mapping and the elliptical model is preserved after a phase-unwrapping step. The artifact at the bottom of the SMF ellipse is due to errors during phase-unwrapping.

[7]. The nonlinear geometric coil combine (GCC) loses the elliptical relationship and does not reliably have a cross-point (i.e., no solution exists), whereas the linear SMF retains the elliptical relationship from which a cross-point can be found.

For any coil combination method (linear or nonlinear), the magnitude of the coil combined phase-cycles are generally good estimates of the magnitude of the equivalent single coil phase-cycle:

$$|I_{i,\text{coil combined}}| \approx |I_{i,\text{single coil}}|, \quad (4.7)$$

for example: sum-of-squares (SOS), maximum SNR coil combination, adaptive coil combination, GCC, etc. [6], [7], [11], [46]. Then, only the phase of the combined images is problematic. If we can estimate the phase of the combined phase-cycled images, then nonlinear coil combination methods may become viable as a preprocessing step to the IGS. To estimate the phase for each phase-cycle value, a phase substitution may be made by taking the phase of each phase-cycle point on a composite ellipse at each pixel. This composite ellipse is constructed by finding the inverse transformation $T_j(x,y)$ required to co-register the j th coil ellipse at pixel (x,y) and treating these

as the ‘‘sensitivities’’ relating intra-phase-cycle coil images:

$$\angle \begin{bmatrix} I_{0,\text{composite}}(x,y) \\ I_{1,\text{composite}}(x,y) \\ I_{2,\text{composite}}(x,y) \\ I_{3,\text{composite}}(x,y) \end{bmatrix} = \angle \left(\frac{T^H(x,y)\boldsymbol{\sigma}^{-2}}{T^H(x,y)\boldsymbol{\sigma}^{-2}T(x,y)} \begin{bmatrix} I_0(x,y) \\ I_1(x,y) \\ I_2(x,y) \\ I_3(x,y) \end{bmatrix} \right). \quad (4.8)$$

This is in fact the SENSE reconstruction with an acceleration factor of 1 using the inverse affine transformations as the sensitivity weightings [11], [38]. Coil ellipses in the complex plane are formed through the linear transformation, $T_j(x,y)$. $T_j(x,y)$ is a complex number giving the proper rotation and scaling of the phase-cycle points on the coil ellipse to match an arbitrary reference ellipse. Assuming additive white Gaussian noise, we use the least squares solution to the following for each pixel:

$$T_j(x,y) = \begin{bmatrix} I_{0,\text{ref}}(x,y) \\ I_{1,\text{ref}}(x,y) \\ I_{2,\text{ref}}(x,y) \\ I_{3,\text{ref}}(x,y) \end{bmatrix}^T \begin{bmatrix} I_{0,j}(x,y) \\ I_{1,j}(x,y) \\ I_{2,j}(x,y) \\ I_{3,j}(x,y) \end{bmatrix}^* \left(\begin{bmatrix} I_{0,j}(x,y) \\ I_{1,j}(x,y) \\ I_{2,j}(x,y) \\ I_{3,j}(x,y) \end{bmatrix}^T \begin{bmatrix} I_{0,j}(x,y) \\ I_{1,j}(x,y) \\ I_{2,j}(x,y) \\ I_{3,j}(x,y) \end{bmatrix}^* \right)^{-1}. \quad (4.9)$$

The reference should be the same shape and location in the complex plane as the unknown true ellipse (up to a complex scaling factor). This means that for each pixel the reference ellipse is, in general, a differently shaped ellipse. As in practice we do not know the true ellipse, a simple choice for the reference coil ellipse is the largest ellipse at each pixel, i.e., choose the coil ellipse with greatest average phase-cycle magnitude. However, equation (4.9) requires a least squares solution for each pixel, negating some of the savings of coil combination prior to the IGS. Figure 4.3 shows the steps to form the composite ellipse. For increased computational savings, a simple approximation to the full phase substitution scheme is to choose the dominant term of equation (4.8), i.e., choose the phase for each phase-cycle to be the phase of the coil image that has the highest maximum intensity at that pixel. Only a max operation is required across the coil dimension which leads to significant computational savings compared to both the full composite ellipse described in (4.9) and the coil-by-coil IGS. We will dub the phase substitution technique described

in Figure 4.3 as the “full phase substitution” and the simple approximation to this as the “simple phase substitution.”

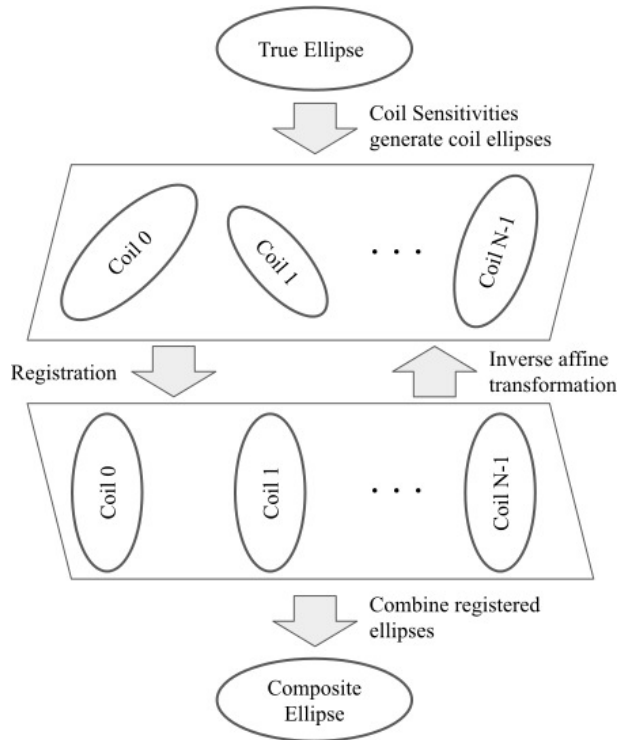


Figure 4.3: Steps to form a composite ellipse. The underlying true ellipse is unknown. With a phased-array receiver, multiple coil ellipses are observed. The coil ellipses are related to the true ellipse by the complex coil sensitivities. Coil ellipse registration is performed which maps each coil ellipse to an arbitrarily scaled and rotated reference of the same shape. The inverse affine transform describes how to map from the registered reference to the coil ellipse. This transformation is interesting because the magnitude of the transformation gives us a relative confidence measure to weight the linear combination of registered coil ellipses. It serves the same function coil sensitivity maps do in the optimal coil combination reconstruction proposed by Roemer et al. Therefore, we use the same combination scheme to form the composite ellipse.

A potential complication of the simple phase substitution method is that phase discontinuities will exist at coil boundaries where adjacent pixels of the composite ellipse take the phase from different coils. We expect the discontinuities to lie smoothly along the intersections of the coil sensitivity maps as shown in Figure 4.4. There will also be discontinuities due to the bSSFP acquisition, as the phase polarization flips at the nulls [34]. As these discontinuities are not rep-

representative of the underlying true image, we enforce a simple rule that, for any given pixel, the substitute phase must come from a common coil ellipse.

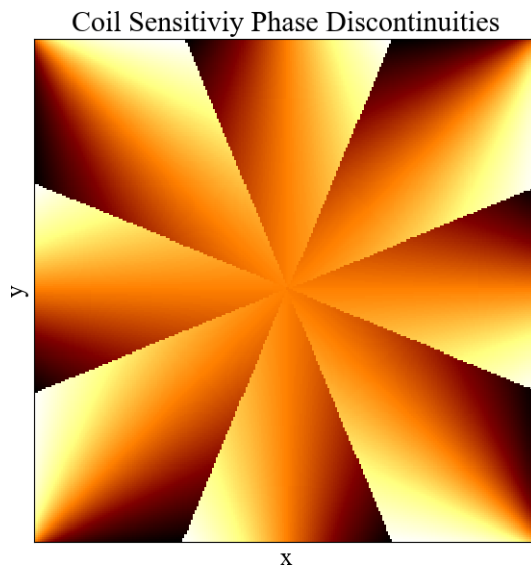


Figure 4.4: This figure shows the phase of the composite coil sensitivity map constructed by taking the phase of the maximum intensity coil sensitivity at each pixel for the simulated coil sensitivity map. Phase discontinuities can clearly be seen at the coil boundaries.

In the remainder of this paper, we will show that nonlinear coil combination techniques can use both the proposed full and simple phase substitution techniques to effectively mitigate artifacts and banding with acceptable image quality. The simple substitution is intended for use where quick post-processing is desired. The full substitution will be provided for comparison.

4.3 Methods

As we only want to validate the proposed methods for a prior coil combination step, exhaustive validation across coil combination techniques is outside the scope of this paper. We have selected three representative coil combination algorithms to use for the experiments we describe in this section: SOS, GCC, and SMF. The SOS and GCC methods are nonlinear whereas the SMF is linear. The SOS method is especially interesting as it clearly does not preserve phase information and cannot be applied alongside the IGS without some additional method for phase estimation.

The GCC is representative of nonlinear SVD-based methods. The SMF approach is representative of linear coil combination methods. None of the methods require prior coil sensitivity estimates.

4.3.1 Simulation

We have performed numerical simulations to validate the proposed phase substitution methods. The purpose of the simulations are to demonstrate the effectiveness of the methods across multiple coil combination techniques and to quantify error as compared to a coil-by-coil IGS reconstruction. Simulations used a 256×256 2D slice through a uniform cylindrical phantom with $T_1 = 1.5\text{s}$, $T_2 = 0.8\text{s}$, and $M_0 = 1$ (a.u.). Simple coil sensitivity maps for a 5 element array were simulated. The bSSFP acquisition simulation for 0° , 90° , 180° , and 270° phase-cycles used parameters: $\text{TR} = 3\text{ms}$, $\text{FA} = 30^\circ$. Off-resonance was simulated as a linear gradient along the x direction from $-1/\text{TR}$ Hz to $1/\text{TR}$ Hz to clearly show banding artifacts.

Coil combination was performed along the coil dimension of the acquisition and reconstructed both with and without the proposed phase substitutions. The simulation was also run under various levels of additive white Gaussian noise on both real and imaginary channels. The normalized root mean squared error (NRMSE) and structural similarity index (SSIM) are plotted as a function of SNR to compare reconstruction quality. The gold standard is taken to be the single coil IGS with no added noise.

4.3.2 Phantom

Axial slices of a cylindrical doped water phantom were acquired on a Siemens TIM Trio 3T scanner. A modified TrueFISP sequence with the following parameters was used to acquire 0° , 90° , 180° , and 270° phase-cycled bSSFP images: TE/TR : $3\text{ms}/6\text{ms}$, BW : 558 Hz/Px, matrix size: 256×256 , FOV : $150\text{mm} \times 150\text{mm}$, averages: 8, and FA : 70° . 200 dummy pulses were used to ensure steady-state was reached. A 4-channel head coil was used. Again, coil combination was performed and the images reconstructed with and without the proposed phase substitution methods. The gold standard is the coil-by-coil IGS, the solution we wish to achieve. Unlike in the simulations, we do not know the underlying true image, so each coil combination method is compared separately.

4.3.3 In vivo

Axial slices of a 24 year old, healthy female participant’s brain were also acquired on a Siemens TIM Trio 3T scanner. This participant had a permanent retainer on her bottom teeth. The participant gave written informed consent prior to participation according to institutional review board guidelines. The same modified TrueFISP sequence was employed to acquire 0° , 90° , 180° , and 270° phase-cycled bSSFP images using the following parameters: TR/TE: 10ms/5ms, BW: 543 Hz/Px, matrix size: 512×512 , FOV: $250\text{mm} \times 250\text{mm}$, averages: 16, and FA: 15° . Again, 200 dummy pulses were used to ensure steady-state was reached. A 4-channel head coil was used to acquire all data. The same post-processing steps as were used for the phantom experiment were applied. The gold standard was again taken to be the coil-by-coil IGS as no suitable true reference is available.

4.4 Results

4.4.1 Simulation

Figure 4.5 shows the phase substitution generated by the proposed simple method for the SOS coil combined 0° phase-cycled image. The magnitude image shows a well-combined composite image without any phase information available. The proposed simple phase estimate has discontinuities along the coil boundaries and bSSFP signal nulls as predicted. The results of the simulation are summarized in Figure 4.6. The error is given at the bottom of the images along with a normalization factor to create sensible contrast for the residual images, as the performance between methods varied wildly. The residual images are normalized to have a maximum of 1. This also means that the residual error for the naive SOS method appears lower than it is due to a singularity, despite being the worst performing solution. The singularity is pointed out with a red arrow in the figure.

The images generated without substitution and nonlinear coil combination have unusable image quality. Interestingly, the SOS coil combination with both full and simple phase substitutions provided the lowest NRMSE among all coil combinations simulated. This suggests that the proposed phase substitutions are reliable estimates, as the SOS provides the optimal magnitude estimate with respect to maximizing SNR [11]. This is desirable, as the SOS reconstruction is easily

computed compared to the SMF, the next best coil combination method with and without the proposed phase substitutions. GCC is greatly improved using both the proposed simple and full phase substitution methods but appears to be limited compared to the SOS and linear coil combination.

Figure 4.6 also depicts some biases in the reconstructions when using different upstream coil combination methods. The SOS method is the maximum SNR solution (while discarding phase information, of course) and produces a noise-like residual when coil-by-coil IGS is performed. This is expected, as the IGS has noise characteristics similar to the complex-sum solution [1]. We find that the desirable noise-like residual is found in both of the proposed methods' residuals with similar error to that of the coil-by-coil residual. The GCC solution has spatial bias with the edges consistently containing more error (except in the naive case where the image quality is unusable). The SMF consistently finds more error at the boundaries of the coils. This is not surprising as the method includes an implicit coil sensitivity map estimation in its algorithm.



Figure 4.5: The left image shows the SOS coil combined 0° phase cycle (SNR=47). The right image shows the simple phase substitution for the SOS reconstruction with the phase discontinuities predicted in the Theory section. The discontinuities lie along the coil intersections as well as along the horizontal bSSFP “null.”

Figures 4.7 and 4.8 show NRMSE and SSIM, respectively, plotted against various simulated SNRs. For nonlinear coil combination techniques, the proposed full and simple phase substitutions perform very well. The simple phase substitution reduces the NRMSE considerably for nonlinear coil combination methods while slightly improving the performance of the linear




Figure 4.6: This figure shows the comparative results of the residual error images for a simulation run with SNR=47. At the bottom of each image the error is given along with a contrast factor that normalizes the residual image. The true reference image is the single coil IGS with no additive noise of a uniform circular phantom. The columns are labeled according to the reconstruction method: “coil-by-coil” IGS, “naive” (meaning that the coil combination and IGS were performed without any phase substitution), the proposed “full” phase substitution, and the proposed “simple” phase substitution. The proposed methods consistently outperform the naive approach. The red arrow points out a singularity in the naive SOS residual.

method. All methods perform best over the range of SNRs simulated when the full phase substitution is used, indicating good estimation of phase. The simple phase substitution does not perform significantly worse than the full substitution, suggesting that the computationally more efficient simple substitution may be used without too much penalty.

4.4.2 Phantom

Figure 4.9 shows the results of the phantom experiments. The naive coil combination followed by the IGS demonstrates low reconstruction quality for nonlinear coil combination methods and artifacts due to phase singularities for the SMF (see red arrow in figure) [46]. These non-smooth phase issues did not arise during simulation where the off-resonance was simulated as a

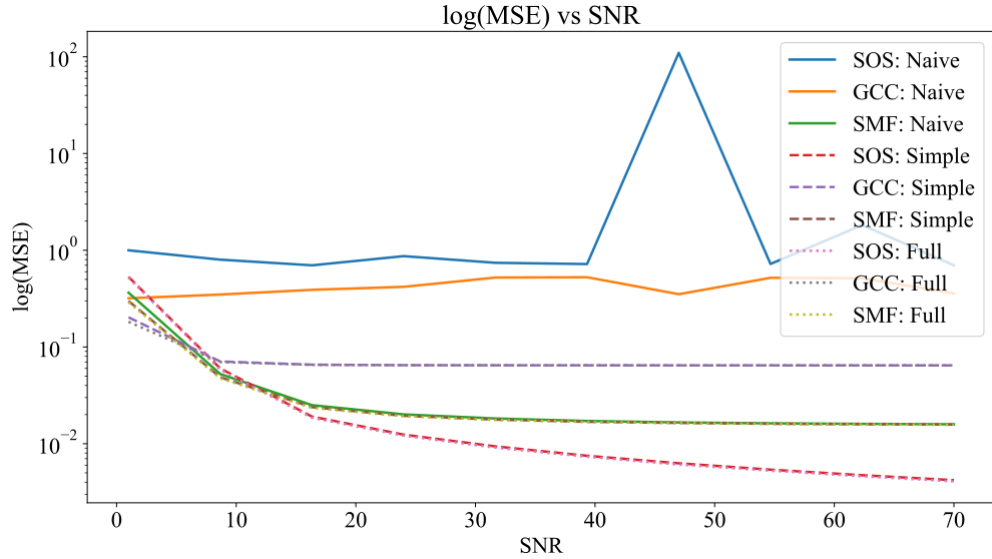


Figure 4.7: NRMSE plotted as a function of SNR using a logarithmic scale for prior coil combination with and without the proposed phase substitutions. Nonlinear coil combination techniques greatly benefit from the simple phase substitution and all coil combination methods benefit from the proposed methods. When the full composite ellipse is constructed, the quality of the reconstruction for the SOS coil combined images are superior to the SMF. SOS with the phase substitutions provide the best reconstruction in terms of raw NRMSE over the SNR range simulated.

smooth gradient. When using adaptive coil combination, singularities in phase can occur in areas of low SNR where differences in individual coil measurements lead to signal cancellation or where a better estimate of absolute phase is required for phase unwrapping [51], [53]. This in turn leads to distortion of the magnitude image (as seen in the figure) and difficulties in phase unwrapping. As seen in simulations, the naive SOS and GCC show serious artifacts making the image quality unusable.

Figure 4.10 shows the residual error due to coil combination. The NRMSE is shown at the bottom of each image. The gold standard is the first column which naturally has no error. Again, to provide reasonable contrast, the residual images have been scaled by the factor shown in the right hand corner of each image. The full phase substitution provides the closest reconstruction to the coil-by-coil target image with the proposed simple phase substitution only yielding slightly higher NRMSE. This is in close agreement with both the theory and simulated results. Figure 4.11 shows the proposed simple phase substitution for the phantom. The coil boundaries are clearly shown by the phase discontinuities, again agreeing with theoretical and simulated results.

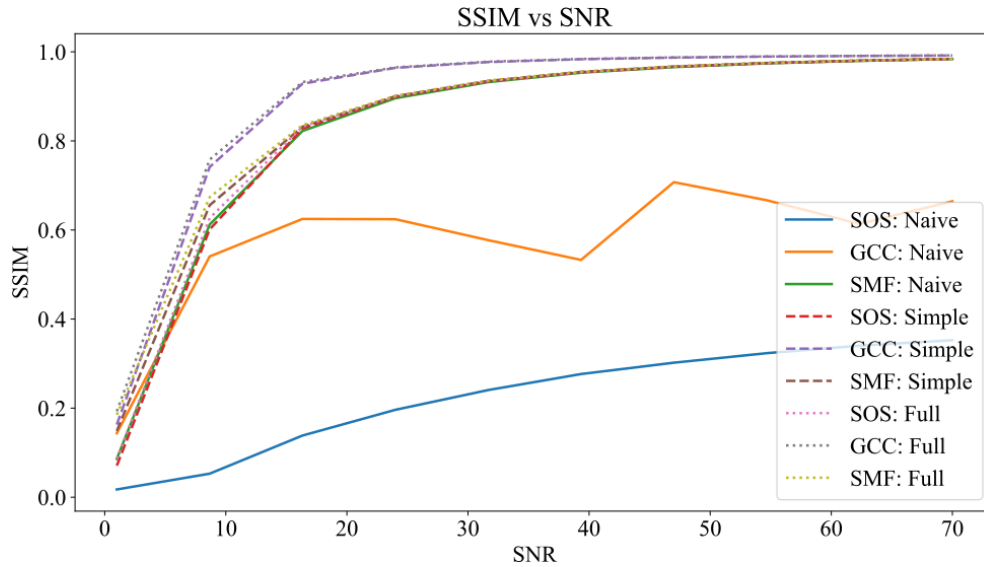


Figure 4.8: SSIM plotted as a function of SNR using for prior coil combination with and without the proposed phase substitutions. The GCC solution with the proposed phase substitutions provide the best similarity score over the range simulated.

4.4.3 In vivo

Figure 4.12 shows the comparative results for the in vivo axial brain reconstructions. The corresponding residual error results are depicted in Figure 4.13 with the same scaling convention as used in previous residual error figures. Again, the full substitution always outperforms the simple phase substitution for each coil combination method compared, although the NRMSE for the full and simple substitutions are similar. The SOS alongside the full phase substitution provided the closest reconstruction to the coil-by-coil IGS, consistent with previous results demonstrating the performance of the proposed substitutions.

Figure 4.14 shows the simple phase substitution. There are no obvious phase discontinuities between the 4 coils in the brain matter, but there appear to be discontinuities in the skull. Again, the proposed methods produce the most consistent and lowest error results.

4.5 Conclusion

The proposed phase substitution methods consistently produce faithful reconstructions using the IGS algorithm. While the full phase substitution was explored as a solution to the ellipse combination problem, the simple phase substitution was presented as a computationally efficient




Figure 4.9: This figure shows the comparative results of the uniform phantom experiment. The naive and proposed methods are compared to the coil-by-coil IGS reconstruction. The columns are labeled according to the reconstruction method. The proposed methods consistently outperform the naive approach, showing virtually no artifacts in this uniform phantom. The red arrow points out an artifact due to phase discontinuities of the naive SMF coil combination. The SMF may have artifacts if the combined image’s phase is non-smooth while the proposed method exhibits no comparable artifacts.

alternative that does not degrade performance significantly. Simulations were performed that revealed biases in the solutions for the IGS when employing the GCC and SMF coil combination methods. Not only was the SOS shown to be feasible using the proposed phase-substitution, but the reconstruction with phase substitution was found to avoid these systematic biases and produce solutions with the lowest error. This gives confidence that the proposed phase substitutions are accurate, as the SOS estimate is the optimal magnitude estimate. SOS and the proposed phase substitutions are also significantly easier to compute compared to the SMF and GCC methods. Interestingly, as shown in the simulation, similarity index was maximized using the GCC, which can be understood as exploitation of the spatial correlations between coil channels [6]. Both the proposed phase substitutions are promising avenues for answering the coil combination problem for multiple phase-cycled bSSFP.




Figure 4.10: This figure shows the comparative results of the uniform phantom experiment along with the NRMSE between the results using each coil combination method. Notice the scaling applied to the images for more obvious contrast. All naive methods produce varying levels of error: the SOS and GCC, as expected, produce images of unusable quality. Errors due to a phase singularity in the naive SMF can be seen. The full phase substitution consistently provides the best performance with the simple phase substitution performing not much worse.




Figure 4.11: This image shows the simple phase substitution of the 90° phase-cycled image for the uniform phantom experiment. The phase discontinuities clearly shown at the boundaries of the 4 coils matching the results of the simulation.




Figure 4.12: This figure shows the comparative results of the in vivo brain experiment. The naive and proposed methods are compared to the coil-by-coil IGS reconstruction. The columns are labeled according to the reconstruction method. The proposed methods clearly outperform the naive approaches for the SOS and GCC.




Figure 4.13: This figure shows the comparative residual error in the form of NRMSE for the in vivo brain experiment. As expected, massive error is found for the naive SOS reconstruction where all phase information has been discarded. Performance of the naive GCC also matches that of the simulation and phantom experiments.

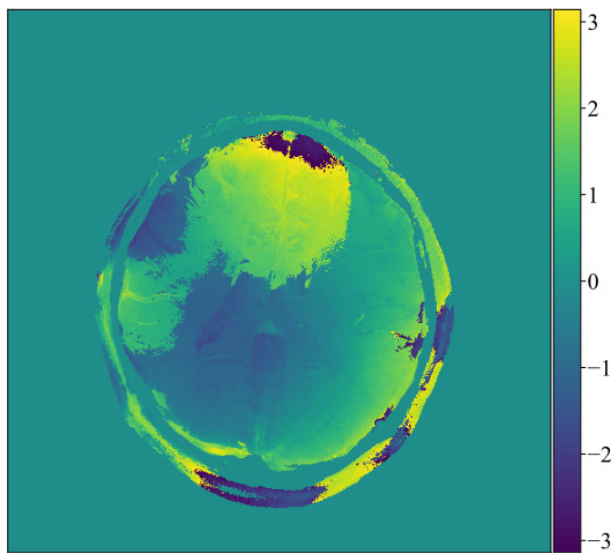


Figure 4.14: This image shows the simple phase substitution of the 90° phase-cycled image for the in vivo brain experiment. Interestingly, the phase discontinuities cannot be obviously seen from the 4 coils in the brain matter, but are clear in the outer skull-shell, matching the results of the phantom and simulation experiments.

CHAPTER 5. COMPRESSED SENSING

5.1 Introduction

Compressed Sensing (CS, also referred to in the literature as “compressive sensing”) is a growing paradigm in signal processing used to subvert the classical Shannon-Nyquist sampling theorem¹ [56]. Signals of interest can be reconstructed accurately with far fewer samples than the desired resolution. Since its creation, magnetic resonance imaging (MRI) has suffered from restrictive and costly acquisition times, relying mostly on coil improvements and clever pulse sequence programming to mitigate scan time and improve resolution (e.g., GRAPPA, SENSE, partial Fourier sampling, etc.) [57], [58]. The application of CS to MRI (CS MRI) has proven to be very effective, with k -space lending itself naturally as a sparsifying transform domain for the CS formulation. In practice, discrete wavelet transforms (DWT), the discrete cosine transform (DCT), and finite differences transformations, among others, have been found to reliably sparsify natural images, including those arising from MR [8], [9], [59].

It should be noted that better sparse approximations can be made by adapting the reconstruction using *a priori* learned information about the signal of interest, referred to as data adaptive CS [60]. The adaptive CS literature is divided into two groups: dictionary methods and non-dictionary methods. Dictionary methods, while they have been well studied, are largely being replaced by non-dictionary methods, especially those involving deep convolutional neural networks, and are fast gaining ground across many fields within MR due to increasing availability of computational power, comprehensive open-source software packages, and the growing number of large, anonymized, public datasets [61]–[66]. The obvious drawbacks to adaptive approaches are the need for reliable training data and inability to generalize reconstruction pipelines. A complete

¹Perhaps “subvert” is not quite the right word for it, as all the traditional notions of information theoretic criteria can be met for exact reconstruction while still not sampling at the Nyquist rate. Any deficit required by sub-sampling is rectified by imposing additional constraints and regularization terms during the reconstruction process. It is really a trade-off between taking the time to measure and make sure, or assuming that something is the case and cutting corners.

survey of adaptive CS is not in the scope of this thesis. Fortunately, much of the existing CS theory and methods are based on non-adaptive techniques, relying on transform bases arising naturally in specific application domains [67]. It is on this non-adaptive theory that the work in this Chapter will be focused.

CS was introduced into MR to combat the aforementioned trade off between acquisition time and resolution. It grew out of the pioneering work of Candès, Romberg, Tao, and Donoho, who demonstrated that n -dimensional signals with sparse representations can be reconstructed from a set of linear, non-adaptive measurements [67]–[73]. Further, they showed that in a wide variety of cases these reconstructions can exactly recover the desired signal. Thus, it is possible to sense sparse signals by taking far fewer measurements.

CS differs from classical sampling theory in a few important ways. First, classical sampling theory considers infinite dimensional, continuous-time signals. CS restricts itself to finite, n -dimensional signals (usually in \mathbb{R}^n , but also in \mathbb{C}^n). Secondly, CS formulations usually do not sample the signal at specific points in time as is done classically. Rather, measurements are acquired in the form of inner products between the signal and more general test functions (usually including some random component, although some regular sampling schemes have been found to satisfy the CS requirements) [67], [74]. Lastly, Nyquist-Shannon formulations rely on sinc interpolation, whereas CS signal reconstruction typically uses highly non-linear reconstruction and interpolation methods [67], [75].

The notion of sparse representation and reconstruction is actually a well-studied problem within mathematics and statistics as first described by Santosa and Symes in 1986 [76]. But the ramifications of the robustness of the reconstructions provided by this theory have only started to be realized in the past decade, with an emphasis on sparse sensing instead of retroactive application to already complete datasets, hence the name “compressed sensing” [73].

5.2 Problem Statement

Intuitively, we can understand CS as inducing a pseudo-noisefloor by randomly undersampling the signal of interest. While uniform undersampling leads to coherent artifacts, random undersampling leads to incoherent aliasing in the measurement domain. In the sparse domain (the domain where the signal is sparsely represented), the incoherent aliasing leads to a noise floor

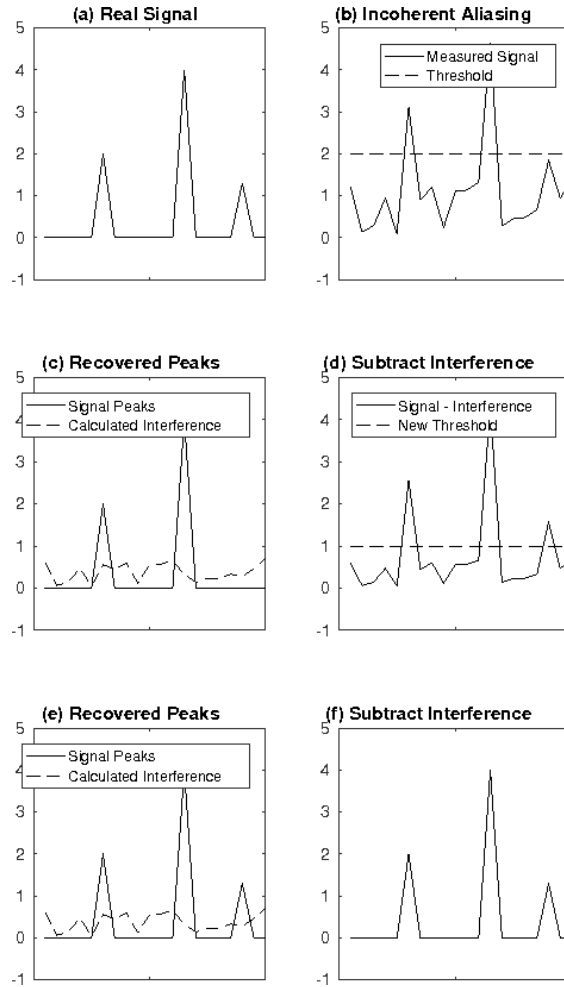


Figure 5.1: Iterative method for CS recovery. (a) shows the desired signal. (b) shows the measured signal with incoherent aliasing due to random undersampling. The threshold we will use to detect peaks is shown. (c) With the pseudo-noise thresholded away, we are left with two out of the three desired peaks. We use the two detected peaks to calculate their expected interference. (d) shows the calculated interference subtracted from the measured signal revealing another, third peak just above the noise threshold. The peaks' magnitudes are adjusted back to original measurements, ensuring data consistency iteration to iteration. (e) Now the interference caused by all three peaks is calculated. (f) shows this last calculated interference subtracted from the measured signal, effectively recovering the desired signal.

with identifiable high sparse coefficients representing true signal components. An iterative method based on thresholding, recovering the strong components, calculating the interference caused by them, and subtracting the calculated interference from the noise floor while retroactively enforcing data consistency can be used to understand the reconstruction process [8]. This intuitive, iterative method is presented in Figure 5.1.

While this intuitive method is appealing, it is not sufficiently generalizable nor efficient to be effective and we require a more rigorous definition. Formally, CS can be understood as a method of finding an exact (or close) solution to the highly under-determined system $Ax = b$, with A being an $m \times n$ matrix such that $m \ll n$. To find a unique solution, the system is regularized by imposing sparsity using the ℓ_0 pseudo-norm in the constrained, highly non-convex optimization problem:

$$\min_{\forall x \in \mathbb{R}^n} \|x\|_0 \text{ subject to } Ax = b \quad (5.1)$$

as in [77], [78], with Lagrangian form given in [67] as:

$$\min_{\forall x \in \mathbb{R}^n} \{ \|b - Ax\|_2^2 + \lambda \|x\|_0 \}. \quad (5.2)$$

Solutions to equation (5.2) incorporating the ℓ_0 quasi-norm are known to be NP-hard² [79]. In nearly all practical CS algorithms and implementations, the ℓ_0 quasi-norm is substituted with its more tractable convex relaxation, the ℓ_1 norm [77]. This leads to favorable forms of solutions and algorithms for the class of L1-regularized problems (see Section 5.4 for a larger discussion of algorithms for CS signal recovery).

We note briefly here that while we assume the matrix A (sometimes called the “sensing matrix”) samples the signal randomly or pseudo-randomly, it must in fact satisfy certain criteria for the inverse problem defined by equation (5.2) to be solvable via linear programming. Although there are a number of ways to say that A has the right properties, the most common is the restricted isometry property (RIP) [67], [80]. If A satisfies the RIP, then the CS problem is well defined. Unfortunately, verifying that A does indeed satisfy the restricted isometry condition is NP-hard. Thus, we usually choose A to be something we know does satisfy or with high probability satisfies the RIP. In this way, the randomness of A is too important to be left up to chance [81].

²In fact, for any fixed L_p norm ($0 < p < 1$), finding the global minimum of the problem is strongly NP-hard [79].

5.3 MR Problem Statement

We can adapt equation (5.2) into a form more natural for the MR experiment:

$$\hat{x} = \operatorname{argmin}_x \frac{1}{2} \|F_u^{-1}x - y\|_2^2 + \lambda \|Tx\|_1, \quad (5.3)$$

where \hat{x} is the n -dimensional reconstructed image, F_u^{-1} is the undersampled inverse Fourier transform operator, y is the acquired k -space data, T is the sparsifying transformation operator, e.g., the DWT operator, and λ is the tunable regularization parameter giving relative weights between the fidelity and sparsity terms. This form takes into account k -space formalism by recognizing F_u^{-1} as the sensing matrix A . Here we have somewhat relaxed our mathematical notation without becoming ambiguous. We, and much of the literature, will prefer this more relaxed style.

5.4 Algorithms

Algorithms designed to solve equations (5.2), (5.3), or variants of these can be broadly classified under the following headings as in [82]:

- Convex
- Greedy
- Thresholding
- Combinatorial/Graph-based
- Non-Convex
- Bayesian

Convex approaches, such as linear programming (including the well-known basis pursuit, basis pursuit denoising, and LASSO algorithms), have been most prevalent in MR due their simplicity of implementation, robustness to noise, and existing software solutions [83]–[86]. However, these algorithms tend to scale poorly leading to increasing computational requirements as datasets get larger. Greedy and thresholding algorithms are faster and scale well but have stricter sparsity constraints; that is, the assumption of sparsity is stronger, with some algorithms requiring signals

to be k -sparse (contain at most k non-zero coefficients in the sparsifying transform domain), and as a result have less prevalence in CS MR reconstruction.

A complete review of algorithms for CS signal recovery is outside the scope of this thesis, however, the interested reader is directed to: [87]. We will now briefly describe gradient descent (a convex method) for its intuition and simplicity as well as iterative soft thresholding (a thresholding method) for its desirable computational efficiency.

5.4.1 Gradient Descent

One of the simplest first-order algorithms for unconstrained optimization is gradient descent (also called “steepest descent”) [88]. The gradient descent algorithm has the form of an update:

$$x_{t+1} = x_t - \lambda_t \nabla f(x_t), \quad (5.4)$$

where $f(x)$ is the objective or cost function that we wish to minimize, x_t is the value of its argument at iteration t , and λ_t is the step size for the current iteration. This update is evaluated either for a fixed number of iterations or until some end condition is satisfied, e.g., $\|\nabla f(x_t)\|_2$ is sufficiently small. From an initial starting point $x = x_0$, the minimum of $f(x)$ is found by continually taking steps in the negative direction of the gradient at each point until a stationary point is reached – hopefully, a local minimum. The negative gradient, $-\nabla f(x)$, gives the direction that $f(x)$ decreases at the fastest rate.

In the case of CS MR reconstruction, the cost function $f(x)$ is equation (5.3) and step size is usually constant, chosen by trial and error or using a technique similar to the L-curve method [89]. We recognize that the gradient operator is linear and may be applied to each term of equation (5.3) individually. The first term is derived using simple matrix calculus:

$$\frac{\partial}{\partial x} \frac{1}{2} \|F_u^{-1}x - y\|_2^2 = (F_u^{-1})^H (F_u^{-1}x - y) \quad (5.5)$$

$$= F_u (F_u^{-1}x - y). \quad (5.6)$$

The second regularizing term is a little more troublesome as it includes an arbitrary operator T inside the ℓ_1 norm. The ℓ_1 norm is defined as the sum of absolute values. Unfortunately, the

absolute value function is non-smooth and not well-defined for all argument values. So, as any good engineer would, we fudge the relation to create a smooth function [8]:

$$|x| \approx \sqrt{x^H x + \varepsilon}, \quad (5.7)$$

where ε is some small smoothing factor usually in the interval $[10^{-15}, 10^{-6}]$. Then we have the approximation:

$$\frac{\partial |x|}{\partial x} = \frac{x}{\sqrt{x^* x + \varepsilon}}. \quad (5.8)$$

We can then define the diagonal matrix M with elements $M_i = \sqrt{(Tx)_i^H (Tx)_i + \varepsilon}$ and get the following result:

$$\nabla \|Tx\|_1 = T^H M^{-1} Tx. \quad (5.9)$$

It has been suggested that applying a finite differences approximation to the gradient can also be effective [90], [91]. When the finite differences transform is chosen for T , we recognize the second term as simply the total variation (TV) [8]. Total variation is a popular regularizing term for both reconstruction and image enhancement, famously introduced by Rudin, Osher, and Fatemi (ROF) in 1992 for image denoising [90]. While the ROF formulation derives the gradient using variational calculus, TV regularization terms have well-known, though cumbersome, finite differences gradient expressions, for example, calculated explicitly in [91]. When other sparsifying transforms are used, other algorithms tend to be favored, i.e., the famous Split-Bregman iterations or more general linear programming solvers [86]

5.4.2 Iterative Soft Thresholding

Iterative soft thresholding (also known as “iterative shrinkage thresholding”) is a member of the proximal gradient descent family and was introduced first in 2004 by Daubechies [92]. It may be thought of as an extension of the gradient descent algorithm. IST is strongly related to iterative hard thresholding (IHT) algorithms, both of which have the form:

$$x_{t+1} = S_\lambda(x_t + F_u(y - F_u^{-1} x_t)), \quad (5.10)$$

where $S_\lambda(\cdot)$ is the soft thresholding operator for IST and hard thresholding operator for IHT. The soft thresholding operator was originally derived using variational calculus, but a more straightforward derivation is found using proximal mapping since equation (5.3) is in the form of the sum of two convex functions [92], [93]. A proximal mapping is a generalization of the concept of a projection operator and was defined by Moreau in the 1960s as [93], [94]:

$$\text{prox}_g(x) = \underset{u}{\text{argmin}} g(u) + \frac{1}{2} \|u - x\|_2^2. \quad (5.11)$$

Then for $g(x) = \lambda|x|, \lambda > 0$, we know that prox_g minimizes the function:

$$h(u) = \begin{cases} h_1(u) = \lambda u + \frac{1}{2} \|u - x\|_2^2, & u > 0, \\ h_2(u) = -\lambda u + \frac{1}{2} \|u - x\|_2^2, & u \leq 0. \end{cases} \quad (5.12)$$

Thus for $u > 0$ we minimize $h_1(u)$ by taking the derivative and setting it equal to zero:

$$0 = h'_1(u) = \lambda + u - x \quad (5.13)$$

$$\implies u = x - \lambda. \quad (5.14)$$

Therefore, for $x > \lambda$, $\text{prox}_g(x) = x - \lambda$. A similar argument finds that for $x < -\lambda$, $\text{prox}_g(x) = x + \lambda$. Now for $|x| \leq \lambda$, $\text{prox}_g(x)$ must be the only point of nondifferentiability of h which is at 0. We now define the soft thresholding function as prox_g :

$$S_\lambda(x) = \text{prox}_g(x) = \begin{cases} x - \lambda, & x \geq \lambda, \\ 0, & |x| < \lambda, \\ x + \lambda, & x \leq -\lambda. \end{cases} \quad (5.15)$$

We may write $S_\lambda(x)$ more succinctly as:

$$S_\lambda(x) = \max(|x| - \lambda, 0) \text{sign}(x). \quad (5.16)$$

We also have a suitable counterpart for complex-valued functions presented here without derivation (see Remark 2.5 in [92]): $S_\lambda(re^{j\theta}) = S_\lambda(r)e^{j\theta}$.

For the (not strictly) convex function $\|Tx\|_1$ in equation (5.3), we notice that the ℓ_1 norm is a summation of functions of the form $g(x)$. Thus by Theorem 6.6 (proximal mapping of separable functions) in [94], we can write the proximal mapping of $\|Tx\|_1$ element-wise as:

$$\text{prox}_{\lambda\|\cdot\|_1}(Tx) = S_\lambda(Tx). \quad (5.17)$$

Now we view IST in the proximal forward-backward framework, i.e., minimization of the residual followed by a “projection” back into a sparse subspace via the proximal mapping $S_\lambda(x)$ [93]. Notice that we now consider the step-size from gradient descent as the threshold for the soft thresholding function. Convergence of this class of algorithms is well-studied and guaranteed for an appropriately chosen threshold [92], [93], [95]. An appropriate threshold is usually chosen empirically (as is custom with gradient descent) with an additional gradual contraction factor, i.e., $\lambda_{t+1} = \mu\lambda_t, 0 < \mu < 1$ [82].

CHAPTER 6. THE ORDERING PROBLEM FOR COMPRESSED SENSING

6.1 Introduction

Compressed sensing is an image reconstruction framework that allows for reconstruction of datasets sampled well-under the Nyquist rate as described in Chapter 5. Its application is natural in MR where images are often sparse under common transforms such as variants of the discrete wavelet transform (DWT) and discrete cosine transform (DCT). While prior information about the signal of interest is often incorporated in the inverse problem through the use of a sparsifying linear transformation, other forms of prior information exist; for example: we might expect that the underlying image has small total variation (TV) measure or we might like to enforce a low-rank solution by minimizing the nuclear norm.

An incredibly interesting and non-standard prior was proposed very early on in CS MR by Adluru and DiBella in [9]. They proposed that knowledge of pixel order relations gives strong prior information from which to regularize the inverse problem. These order relations contain information about relative pixel intensities of the true image. By “true image” we mean the image that would exist if it were fully sampled, or equivalently, if we had the perfect reconstruction. The orderings are applied to the current image estimate for sparsity constraints, i.e., the image estimate is sorted before a sparsifying transform is applied and unsorted after inverse transformation. Ordering of both pixel magnitudes and real/imaginary pixel intensities have been considered as powerful regularizers of the ill-posed inverse problem. Monotonic ordering has been the most common in the literature for its simplicity and effectiveness under a variety of transforms.

It has been found that enforcing monotonic pixel ordering is incredibly effective when the orderings of the true image are known a priori, especially in conjunction with TV [9], [59]. Unfortunately, the true image’s monotonic orderings are difficult to estimate from the corrupted initial image estimates. In fact, reconstruction quality can degrade when ordering estimates are not correct [96]. Thus some researchers have proposed using an ordering-based reconstruction as a

second-pass solution: obtain a first-pass solution using conventional CS reconstruction techniques and feed this result into a second reconstruction using the orderings estimated from the first-pass solution [97]. This can be effective but it is still sensitive to errors in the first-pass result, i.e., the first-pass result might not contain order relations completely consistent with the true image.

Other researchers have asked the question: which ordering is the “best” ordering? In the literature it is understood that the best ordering is that which most enhances the sparsity of the sparsifying linear transform, T . This takes the ordering information and presents it as a new sparsity assumption: the true image is sparse under T after proper ordering. Initial results showed that prior-based monotonic sorting was simple and led to impressive results for such a small investment. It was also recognized that searching for the most effective ordering is equivalent to the infamous traveling salesman problem (TSP), which is known to be NP-hard [98], [99]. The cities in this case are pixels or patches from a prior image estimate and the distances between the patches are some measure of smoothness between them. The TSP solution defines a permutation of the signal. It is natural that simple monotonic orderings of prior image estimates have been used: sorting such that the signal is monotonic is a useful heuristic, as it removes high frequency content from the signal. It can also be used as a simple way to find a solution to the TSP when represented as a complete graph. Sorting the weights and choosing the next vertex that minimizes the cost at each step is very similar to Kruskal’s algorithm for finding the minimum spanning tree [100].

Further, Wu et al. applied the ordering problem to parallel imaging CS also using a TV constraint during reconstruction [59]. This work also demonstrated feasibility of the ordered DCT for sparsifying 1D signals. The theory of using relative orderings for reconstruction priors was formalized by Ramirez et al. as a Technique for Reconstruction Using Intensity Order (TRIO) [97]. This framework, again, only considers monotonically ordered data but incorporates smoothness conditions on phase and allows for block orderings, i.e., smaller groupings of pixels that maintain monotonic relationships rather than enforcing global monotonicity. Ram et al. introduced a more general signal processing framework based on smooth ordering of patches (not necessarily monotonic) and first recognized the ordering step as a solution to the TSP [98]. Using patches is desirable if the statistics of the underlying image cannot be estimated well using single pixels, a condition we find ourselves with in CS MRI. Adluru et al. also considered orderings of the Casorati matrix for multi-image reconstructions using low-rank regularizers via the nuclear norm [96].

Various wavelet transformations have been considered in both [99] and [101]. The former takes orderings similar to Ram et al., i.e., via an approximation to the solution of the TSP on a graph of image patches, while the latter uses an appropriate image prior to train permutations of patches for a locally adaptive DWT. Similar work including orderings has been presented in [102] and [103] but seem to be more related to the convenience of sorting by pixel intensity when computing the Gini index, an alternate measure of sparsity [104], [105]. The easy-path wavelet transform can also be viewed in the context of locally optimized orderings at each decomposition level [106].

As mentioned, previous works have chosen monotonic orderings or trained orderings based on low-resolution or CS reconstruction priors. Most previous work has restricted attention to 2D or 3D image sets which present computational challenges for solutions to the TSP or any brute force training approaches. Multi-image datasets have received attention mainly focusing on large datasets such as dynamic cardiac perfusion and diffusion imaging [96]. To the author's knowledge, no previous work has considered cases where orderings are known a priori in one or more dimensions or on datasets small enough to allow computationally challenging combinatorial searches.

Here we depart from the existing CS ordering literature by presenting a prior-less formulation for order estimation based on implicit models defined by ordering constraints. The goal is to decouple the orderings from potentially untrustworthy image priors. It also includes imaging situations where orderings are known by virtue of physical principles but for which simple ordering schemes do not produce appreciable improvement to the reconstruction quality. Examples include T_1 and T_2 mapping where the acquisition is comprised of k -space data collected at various time points. T_1 mapping datasets have time curves which are strictly increasing whereas T_2 mapping datasets have strictly decreasing. Both of these properties are due to the underlying physical relaxation properties described in Chapter 2. Importantly for our development, the number of time points is small (~ 7). As we know that the ordering problem is NP-hard, then even for small images the computational burden of using a brute force method for finding the optimal order is not feasible. However, for time curves with a small number of points, combinatoric searches are possible.

T_1 mapping datasets have naturally monotonically sorted time curves. However, the time curves are exponential, leading to poor sparsity under the gradient operator which also lead to

largely unimproved TV regularization. Sparsity of ordered exponential functions under other sparsifying transforms have been largely unstudied. Thus we require a method for selecting optimal orderings for datasets with few time frames. We can do this without a prior image estimate by basing the ordering on minimization of the final reconstruction objective function. The assumption that the objective function is a good surrogate for reconstruction quality is at the heart of any CS reconstruction algorithm. By finding the most desirable objective function value, we can assume that we have found the most desirable reconstruction and optimal ordering.

In this chapter we will consider the CS ordering problem where one dimension is known to be monotonically varying. In the Theory section we will formally state the problem and present a simple algorithm to find the optimal ordering for an arbitrary sparsifying transformation. We will also present a computationally tractable method to find the best ordering on average for T_1 mapping experiments with more than one T_1 species. To validate the theory we will present the results of simulated T_1 mapping experiments on simple phantoms and an in vivo reconstruction experiment using cardiac MR datasets.

6.2 Theory

CS reconstruction algorithms typically consider two types of terms in their objective functions: fidelity terms and sparsity regularization terms. In this paper we consider CS MR reconstruction algorithms of the form:

$$x = \operatorname{argmin}_x \frac{1}{2} \|F_u^{-1}x - y\|_2^2 + \lambda \|Tx\|_1, \quad (6.1)$$

where x is the current image estimate, y is the complex k -space measurement, F_u^{-1} is the under-sampled inverse Fourier transformation, T is the sparsifying linear transform, and λ is the regularization constant that determines the trade-off between fidelity and sparsity. The sparsifying transform may be any linear transform operator. The ℓ_1 regularization term is often replaced by the TV measure, as it has the form of an ℓ_1 norm. In fact, in one dimension, it is the ℓ_1 norm of the finite difference:

$$\operatorname{TV}(x) = \sum_i |x_{i+1} - x_i| = \|\nabla x\|_1, \quad (6.2)$$

over valid indices of the signal x .

The ordering problem restates the basic CS formulation by simply inserting the permutation operator Π which maps pixels to their ordered locations:

$$x = \operatorname{argmin}_x \frac{1}{2} \|F_u^{-1}x - y\|_2^2 + \lambda \|T\Pi x\|_1. \quad (6.3)$$

The goal is then to construct an operator Π_j that produces, in general:

$$\|T\Pi_i x\|_1 > \|T\Pi_j x\|_1, \forall i \neq j, \quad (6.4)$$

that is, the permutation operator that yields the lowest ℓ_1 norm. For TV, it has long been known that monotonic orderings produce good results but it has never been proved optimal [9]. Here we present a simple argument that monotonic ordering is the best such ordering for TV minimization.

To maximize the number of zeros under the finite differences operator, simply group the elements with the same value. Transition coefficients between groups are minimized when the distance between groups is minimized. We accomplish this by sorting the groups. But this is just monotonic ordering. Notice that the set of transition coefficients obtained in this way has cardinality less than or equal to that of any other ordering, since the number of zeros was maximized at the start. The transition coefficients are a subset of the set of coefficients for all other orderings. In fact, it is the subset of smallest coefficients. Thus the ℓ_1 norm is less than or equal to that of any other ordering. Therefore monotonic ordering is an optimal solution for TV minimization.

However, other transforms do not have obvious solutions for optimal sortings, leading researchers to consider patch-based TSP solutions, advanced data-learning techniques, or heuristics [59], [99], [101].

T_1 time curves have exponential shapes. Under the finite differences transform, an exponential retains the same shape with a change of multiplicative factor, which ideally should not change the sparsity measure [105]. This is actually not true for the ℓ_1 norm, as a scaling factor < 1 will reduce the norm. For this reason, we choose to utilize the Gini index as proposed by both [105] and [104]. Thus we have that T_1 curves are inherently monotonically sorted and that TV regularization is potentially ill-suited for the reconstruction of exponential curves. For these

reasons we must consider other transforms and, in general, non-monotonic orderings generated without prior image estimates.

Here we present a simple brute force algorithm that finds the orderings which minimize the objective function described in equation (6.3): For all possible orderings Π , reconstruct x according to equation (6.3) and choose the reconstruction which minimizes equation (6.3). In fact, this is equivalent to a model based reconstruction under the same assumptions. The sparsity assumption defines an implicit model that selects signals which can be represented sparsely after permutation. Adluru and DiBella proposed monotonic ordering for TV which defined an implicit model enforcing contrast between disparate regions of the image. However, other transformations, after ordering, do not have as easily divivable implicit models, such as those presented by Lai et al. and Lui et al. [99], [101]. However, if the implicit model is reasonable (whether intuitively or empirically), then the order corresponding to the minimum value of the objective function should be chosen, as it has simultaneously minimized the data fidelity term and maximized sparsity under equivalent reconstruction conditions, satisfying the implied model. If the ordering that produces the minimum mean squared error (MSE) differs from that which minimizes the cost function, then the original assumption of sparsity was flawed and the implicit model does not model the signal well. All existing CS ordering literature has considered empirical evidence of enhanced sparsity sufficient justification for their proposed models.

In the case of T_1 mapping, the sparsity constraint implies that exponential time curves arising in MR have sparse permutations under the transformation T . While it is not clear that this must be the case, using numerical simulations, we have confirmed that this implicit model is reasonable across a variety of tissues and leads to improved reconstruction quality compared to similar reconstructions with and without ordering. Figure 6.1 shows the results of the Gini sparsity measure for optimal orderings found according to our simple algorithm over 1000 simulations for 7-point T_1 time curves of tissue parameters with T_1 simulated from 100ms to 2000ms. A T_1 exponentially decaying model was chosen:

$$s(t) = A - Be^{-t/T_1}, \quad (6.5)$$

where $s(t)$ is the simulated time curve at inversion times t and A and B are model parameters chosen to be 1 and 2, respectively.

The following transforms were compared: the identity transformation, one-dimensional finite differences (FD), DCT, and the one-dimensional Cohen-Daubechies-Feauveau bi-orthogonal wavelet (a common wavelet transform in MR [107]). As expected, all optimal orderings for the FD transform were monotonically ordered. The orderings for the other transforms were not predictable, but the DCT did lead to greater sparsity as measured by the Gini index. This transform is interesting for the present discussion because the best orderings are not obvious and it is efficient to compute. For these reasons, we choose T to be the Type II DCT for the remainder of our discussion. The fact that orderings other than the identity permutation were chosen indicates that ordering improved the sparsity measure. We consider this significant, as sparse representation is usually more difficult for small datasets (in this case a single 7-point time curve).

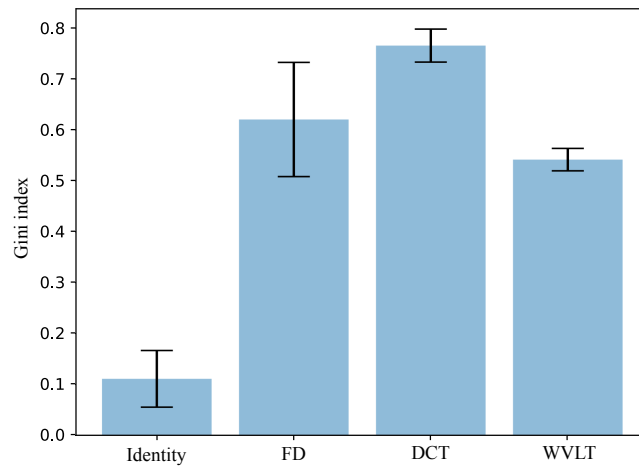


Figure 6.1: Mean Gini index for several transforms of 1000 optimally ordered T_1 time curves. FD, DCT, and wavelet (WVLT) transforms see improved sparsity over no transformation at all. The DCT empirically sparsifies the time curve best in this experiment. The error bars show the standard deviations.

As the DCT does not lend itself to an obvious solution for optimal ordering, in the Methods section we will apply the aforementioned brute force algorithm to present proof-of-concept simulation results for the prior-less CS ordering problem. We predict that we will find more accurate

reconstructions by using our proposed method than from the standard CS prior. Notice that the standard CS prior corresponds to the identity permutation which is the best theoretical ordering for finite differences.

In the context of T_1 mapping, up to this point we have assumed a single-tissue imaging sample. This single species required $N!$ reconstructions to be performed to find the optimal ordering. In general, it is unreasonable to assume only a single species is present, but rather that each voxel may contain a unique T_1 relaxation value. Of course this translates to quickly escalating computational burden for the brute force method we have presented, as the optimization must consider all combinations of orderings for all pixels simultaneously. In lieu of this, we propose that simple priors derived from the imaging context can be used to find acceptable, if sub-optimal, orderings that still provide reduced objective function value. For this chapter, we propose that we know a range of expected T_1 values present in the sample, e.g., we expect each voxel will have $0.2s < T_1 < 2.0s$. The intuitive argument is that similar exponential curves have similar orderings, therefore if we optimize with respect to a range of values, then we can find an ordering that works well on average. Figure 6.2 shows the results of a simulation of 8-point time curves over a range of 1000 T_1 values using the optimal orderings (the orderings that produce the most sparse representation of the T_1 signal model), the monotonic ordering, and the three most common orderings. The frequency of the three most common orderings is shown in Figure 6.3. Note that the frequencies include occurrences of multiple optimal orderings, i.e., more than one ordering produces optimal results. We find that these average optimal orderings are effective at maximizing the sparsity over particular ranges of T_1 . The two most common orderings account for about 20% of all optimal orderings found. In this case, the first and second most common orderings were effective over the widest range of values and were at least as effective as monotonic ordering for almost the entire range. As the two orderings overlap, we consider them to be equivalent orderings, i.e., they lead to equivalent sparsity measures. We then take one of these to be the average optimal ordering to apply during subsequent CS reconstruction.

Using the average optimal ordering scheme, all simulations to determine the correct ordering are performed prior to application during the actual reconstruction. This reduces the computational burden significantly once an average optimal ordering is found for a particular imaging context: only one optimization must be performed. We also decouple assumptions about orderings

from the naive image priors based on acquired data, a property we found desirable when performing brute force cost function evaluations. We rely on the assumption that maximizing sparsity minimizes the cost function and that we have chosen the correct range of T_1 values over which to optimize orderings.

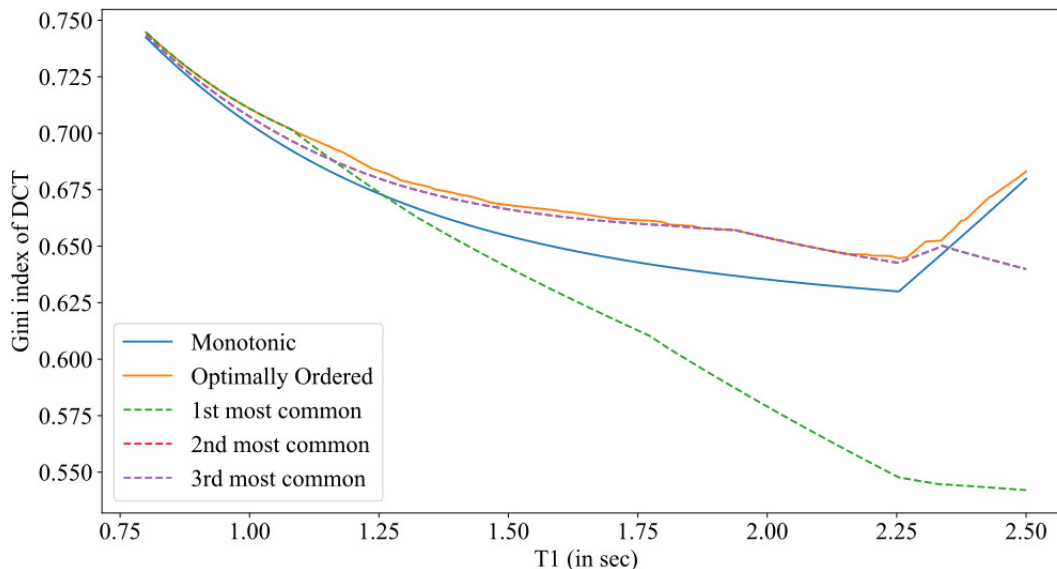


Figure 6.2: Gini index of the DCT of T_1 relaxation curves across a range of 1000 T_1 values. Monotonic ordering, optimal ordering, and the three most common orderings are plotted. The most common orderings are near-optimal over certain ranges of T_1 values. Notice that while the first two most common optimal ordering are distinct, they are equivalent orderings: they produce the same sparsity measure. The first and second most common orderings are near-optimal for much of the entire simulated range. Some common orderings are only optimal over a small range of T_1 values and become worse than monotonic ordering for some ranges.

The remainder of this chapter will describe numerical experiments (both simulated and in vivo) to validate the ordering methods proposed in the Theory section.

6.3 Methods

6.3.1 Simulations

We wish to demonstrate that orderings for the DCT can be found that outperform the monotonic ordering strategy. To this end, we choose a simple 128×128 numerical phantom consisting of seven circles with identical tissue parameters: $T_1 = 1200\text{ms}$, $T_2 = 100\text{ms}$, and $M_0 = 1$ (a.u.).




Figure 6.3: Frequency of the three most common optimal orderings. The ordering indices are shown along the x -axis. The two most common orderings account for almost 20% of all orderings.

These are approximate MR parameters for gray matter at 1.5T. This phantom is depicted in Figure 6.4. This choice of phantom is similar to conventional T_1 vial experiments, but in this case, all vials contain similar tissue. A spin-echo inversion recovery sequence was simulated with $TR = 2500\text{ms}$ and $TE = 10\text{ms}$ for seven time points corresponding to inversion times $TI = 1.14\text{s}, 1.55\text{s}, 1.93\text{s}, 2.26\text{s}, 2.54\text{s}, 2.76\text{s},$ and 2.93s . The resulting k -spaces were retroactively undersampled, keeping only 30% of the k -space phase-encoding lines while guaranteeing that the center 10 lines were always acquired. A representative sampling mask is shown in Figure 6.5 and the corresponding zero-filled naive inverse Fourier transform of the undersampled phantom is shown in Figure 6.6.

For the proposed method we performed reconstructions using an iterative soft thresholding (IST) algorithm as described in Chapter 5 with starting threshold $\lambda_0 = 0.09$ and threshold contraction parameter $\mu = 0.98$. T was chosen to be the one-dimensional Type II DCT across the time dimension. For $N = 7$ time points, we ran the reconstruction for each of the $N! = 5,040$ possible orderings. The IST algorithm was run for a fixed 50 iterations for each reconstruction. The ending value of the objective function was evaluated and the ordering corresponding to the minimum was chosen as the best ordering. For comparison, the cost function and normalized root mean square error (NRMSE) was evaluated for the chosen ordering, i.e., the ordering which produced the lowest NRMSE, and for the ordering corresponding to the standard CS recon. We also performed a nonlinear least squares estimate for T_1 using the Levenberg–Marquardt algorithm, as this is the actual value desired when reconstructing datasets for T_1 mapping.




Figure 6.4: Numerical phantom used to simulate inversion recovery pulse sequence to validate the proposed theory of prior-less ordering estimation for compressed sensing.

All iterative reconstructions for the simulations were run on Brigham Young University’s Mary Lou Fulton supercomputer.

6.3.2 In vivo

We wish to demonstrate that average optimal orderings for the DCT can outperform the monotonic ordering heuristic. Single-slice cardiac data was acquired on a 3T Siemens Magnetom Prisma Fit. The patient was imaged head first-supine using a Cartesian MOLLI 5(3)3 sequence [108] with the following parameters: TR/TE: 460.9ms/1.12ms, FA: 35°, averages: 1, FOV: 315mm × 360mm, slice thickness: 7mm, matrix size: 224 × 256, BW: 1080 Hz/Px, and inversion times at 117ms, 257ms, 1172ms, 1282ms, 2172ms, 2325ms, 3174ms, and 4189ms. The Levenberg–Marquardt algorithm was used to create a reference T_1 map pixel-by-pixel using the model adjustments for MOLLI:

$$s(t) = A - Be^{-t/T_1^*}, \quad (6.6)$$

$$T_1 \approx T_1^*(B/A - 1), \quad (6.7)$$




Figure 6.5: Cartesian k -space undersampling pattern for one time frame of the numerical phantom simulation.

where T_1^* is the apparent T_1 value from the modified Look-Locker experiment, t are the inversion times, and A and B are model parameters [109].

We retroactively under-sampled the fully acquired Cartesian 8-point cardiac dataset using the same Cartesian undersampling strategy as in the simulation but this time sampling 30% of k -space and ensuring the center 15 lines were always acquired. The average optimal ordering was chosen suitably from a simulation over the range: $0.25\text{s} < T_1 < 2.5\text{s}$. An IST reconstruction using regularization parameter $\lambda_0 = 1.5$ and contraction parameter $\mu = 0.98$ was run for a fixed 50 iterations for both monotonic orderings and average optimal orderings. After reconstruction, T_1 maps were generated and compared.

6.4 Results

6.4.1 Simulations

Table 6.1 summarizes the results of the simulation. The bold quantities indicate the best values in the given row. NRMSE decreased and T_1 estimates were closer to the true value (1.2s) when using the ordering generated from the proposed algorithm in a fixed number of iterations.




Figure 6.6: Corrupted phantom image due to random Cartesian undersampling in the phase-encode direction.

The best NRMSE and best cost orderings are different. However this is not surprising, as there are in general multiple optimal or near-optimal orderings. Both orderings lead to virtually the same NRMSE and cost scores as well as equal improvement over the monotonic ordering heuristic. The final average nonlinear least squares estimate for T_1 is also closer using the proposed method with lower mean error and standard deviation. Figure 6.7 shows the comparative reconstructions using both monotonic orderings and best cost ordering estimates for a representative time frame. The residual is consistently lower using the best cost estimate. Notice there are still residual artifacts seen inside the circles. These artifacts are due to the fact that the IST algorithm has not run enough iterations. A cap on maximum iterations was necessary for reasonable run-time. Figure 6.8 shows the NRMSE for both ordering strategies as a function of iteration. It is clear that the best cost ordering NRMSE trends below that of the monotonic ordering as iteration number goes past that of diminishing returns.

Figure 6.9 shows a representative pixel time curve from the variously-ordered reconstructions and the naive zero-filled inverse Fourier transform of the undersampled k -space. The differences in individual ordered time curves appear to be small, but on the whole, given every pixel's contribution, the result corresponding to the best cost function value yields improvement over the

Table 6.1: Summary of results for exhaustive simulation of optimal ordering simulation. Bold entries indicate the best values. We find that the optimal ordering was found using the proposed algorithm, i.e., the minimum cost ordering corresponding to a comparable minimum NRMSE ordering.

	Monotonic	Best NRMSE	Best Cost
Order	[0 1 2 3 4 5 6]	[2 4 6 5 3 1 0]	[3 2 1 5 4 0 6]
NRMSE	4.29e-2	3.97e-2	3.97e-2
Cost	620.38	568.47	568.47
T_1 estimate	1.23s \pm 0.083 s	1.21s \pm 0.062 s	1.21s \pm 0.062 s

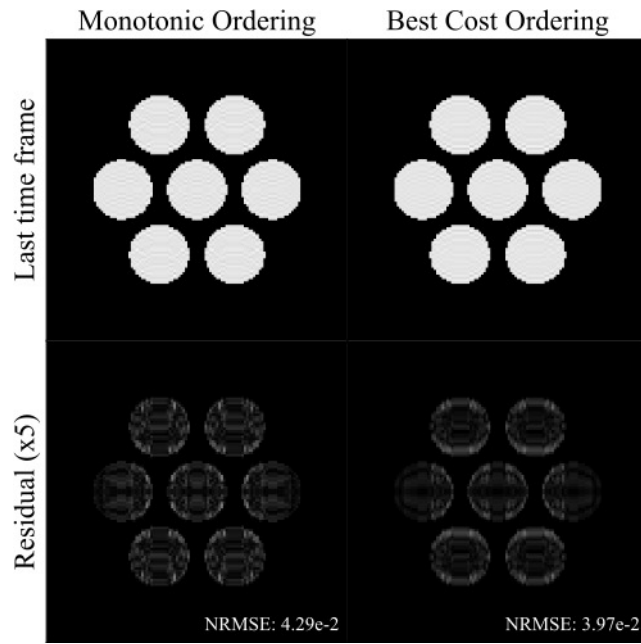


Figure 6.7: Comparative reconstruction of the last time frame of the simulated T_1 mapping dataset. The reconstruction using monotonic ordering is shown on the left with residual underneath, while the reconstruction on the right uses the best cost orderings.

monotonic heuristic, verifying that orderings other than monotonic can provide better reconstruction fidelity and they can be found without a prior image estimate.

6.4.2 In vivo

Figure 6.10 shows the normalized cost function value versus iteration for a fixed number of iterations with and without the proposed average optimal ordering. It is clear that the chosen




Figure 6.8: NRMSE plotted as a function of IST iteration for monotonically ordered and best cost ordered reconstructions.

ordering induces a lower cost function over the same iterations. This is consistent with the theory and performed simulations for finding the average optimal orderings. Figure 6.11 shows the reconstructed T_1 map of the cardiac data using monotonic ordering and a fixed number of iterations and its corresponding residual image. The residual image was generated by subtracting the reconstruction from the T_1 map generated from the fully sampled cardiac dataset. Correspondingly, Figure ?? shows the results of the reconstructed T_1 map using the average optimal ordering. Compared to the monotonic ordering, the average optimally ordered image is smoother and has smaller residual error, verifying that average optimal orderings can produce better quality reconstructions in a fixed number of iterations.

6.5 Discussion

At the cost of running the reconstruction $N!$ times, we achieve the goal of obtaining an ordering that is not tied to a prior estimate or reliant on the heuristic of sorting monotonically. We consider this an encouraging result for a first-of-its-kind study. We have also chosen a target application that has a small number of time points, ensuring computational feasibility of the brute force search method. The computational burden is further lightened by easy parallelization of the small reconstructions. Using a simple modification of the proposed framework, we can find




Figure 6.9: Comparative reconstruction results for a representative pixel’s time curve. The naive inverse fast Fourier transform (IFFT) corresponds to a zero-filled reconstruction. The curves corresponding to the monotonic ordering and best cost/NRMSE orderings appear close, with the best cost/NRMSE slightly better aligning with the true time curve. These small improvements for each pixel lead to improved reconstruction fidelity on the whole.

average optimal orderings across a range of T_1 values for T_1 mapping. These average optimal orderings were effective at reducing residual error and producing more accurate T_1 maps. It is left to future work to study the optimization of the proposed methods and extension to datasets with more time points and the complications of multiple tissue characteristics.




Figure 6.10: The normalized cost function over a fixed number of iterations for a single-slice cardiac reconstruction. The reconstruction using the average optimal ordering shows reduced cost function for all computed iterations.

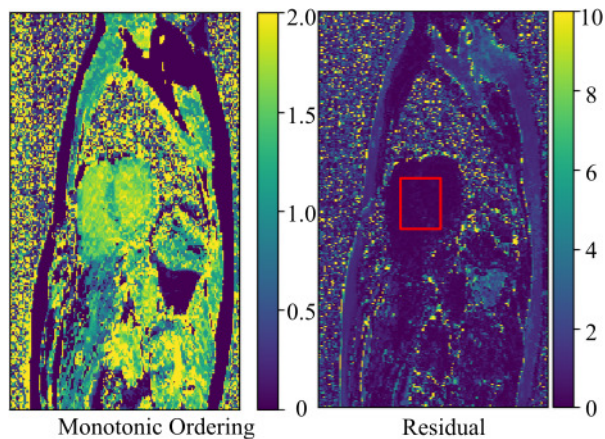


Figure 6.11: The image on the left is a T_1 map generated from reconstructed retroactively under-sampled cardiac data. No data reordering was applied during the reconstruction. The image on the right shows the residual image generated using the fully sampled cardiac data. The red bounding box shows an area of interest with mean error $1.03e-2 \pm 0.25s$. The scales are indicated in seconds.

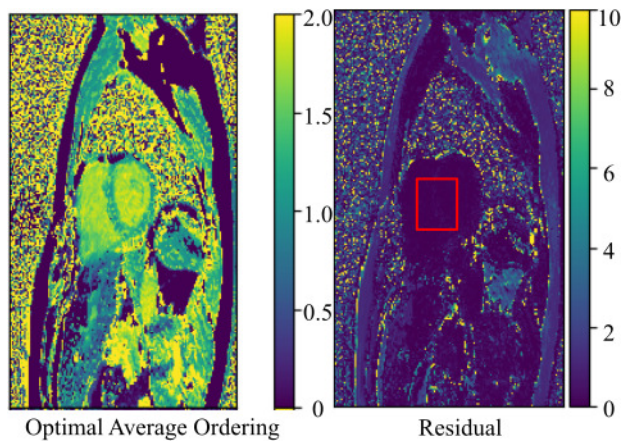


Figure 6.12: The image on the left is a T_1 map generated from reconstructed retroactively undersampled cardiac data. The average optimal orderings for the expected range of T_1 values was applied during the reconstruction. The image on the right shows the residual image generated using the fully sampled cardiac data. The red bounding box shows an area of interest with mean error $3.20e-3 \pm 0.26s$. The scales are indicated in seconds.

REFERENCES

- [1] Q. Xiang and M. N. Hoff, "Banding artifact removal for bSSFP imaging with an elliptical signal model," *Magnetic resonance in medicine*, vol. 71, no. 3, pp. 927–933, 2014. ii, 25, 26, 27, 28, 29, 38, 89, 92, 94, 97
- [2] Y. Shcherbakova, C. A. T. van den Berg, C. T. W. Moonen, and L. W. Bartels, "Planet: an ellipse fitting approach for simultaneous T1 and T2 mapping using phase-cycled balanced steady-state free precession," *Magnetic resonance in medicine*, vol. 79, no. 2, pp. 711–722, 2018. ii, 26, 31, 92
- [3] B. Keil, J. N. Blau, S. Biber, P. Hoecht, V. Tountcheva, K. Setsompop, C. Triantafyllou, and L. L. Wald, "A 64-channel 3T array coil for accelerated brain MRI," *Magnetic resonance in medicine*, vol. 70, no. 1, pp. 248–258, 2013. ii
- [4] M. Buehrer, P. Boesiger, and S. Kozerke, "Virtual body coil calibration for phased-array imaging," in *Proceedings of the 17th Annual Meeting of the ISMRM; Hawaii, USA, 2009*, p. 760. ii, 28
- [5] N. McKibben, G. Tarbox, M. A. Mendoza, and N. K. Bangerter, "Analysis of coil combination for bSSFP elliptical signal model," *Proceedings of the 27th Annual Meeting of the ISMRM; Montreal, Canada*, vol. 7213, 2019. ii, 1, 4, 28
- [6] T. Zhang, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "Coil compression for accelerated imaging with cartesian sampling," *Magnetic resonance in medicine*, vol. 69, no. 2, pp. 571–582, 2013. ii, 27, 28, 31, 32, 42
- [7] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magnetic resonance in medicine*, vol. 43, no. 5, pp. 682–690, 2000. ii, 4, 27, 31, 32, 83
- [8] M. Lustig, D. L. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic resonance in medicine*, vol. 58, no. 6, pp. 1182–1195, 2007. [Online]. Available: <http://dx.doi.org/10.1002/mrm.21391> ii, 46, 48, 52
- [9] G. Adluru, S. P. Awate, T. Tasdizen, R. T. Whitaker, and E. V. DiBella, "Temporally constrained reconstruction of dynamic cardiac perfusion MRI," *Magnetic resonance in medicine*, vol. 57, no. 6, pp. 1027–1036, 2007. ii, 46, 55, 59, 85, 86, 88
- [10] K. L. Miller, "Asymmetries of the balanced SSFP profile. Part I: theory and observation," *Magnetic resonance in medicine*, vol. 63, no. 2, pp. 385–395, 2010. viii, 23

- [11] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, “The NMR phased array,” *Magnetic resonance in medicine*, vol. 16, no. 2, pp. 192–225, 1990. 4, 27, 31, 32, 33, 37, 82
- [12] M. A. Mendoza, N. McKibben, G. Tarbox, and N. K. Bangerter, “Synthetic banding for bSSFP data augmentation,” *Proceedings of the 27th Annual Meeting of the ISMRM; Montreal, Canada*, vol. 4781, 2019. 4
- [13] N. McKibben, M. A. Mendoza, E. V. DiBella, and N. K. Bangerter, “Deep learning super-FOV for accelerated bSSFP banding reduction,” *Proceedings of the 27th Annual Meeting of the ISMRM; Montreal, Canada*, vol. 4791, 2019. 5
- [14] M. A. Mendoza, N. McKibben, and N. K. Bangerter, “Generation of arbitrary spectral profiles using orthonormal basis combinations of bSSFP MRI,” *Proceedings of the 26th Annual Meeting of the ISMRM; Paris, France*, vol. 6332, 2018. 5, 89, 93
- [15] E. M. Haacke, R. W. Brown, M. R. Thompson, R. Venkatesan *et al.*, *Magnetic resonance imaging: physical principles and sequence design*. Wiley-Liss New York:, 1999, vol. 82. 6
- [16] D. G. Nishimura, *Principles of magnetic resonance imaging*. Stanford University, 1996. 6, 19, 20
- [17] M. H. Levitt, *Spin dynamics: basics of nuclear magnetic resonance*. John Wiley & Sons, 2001. 6
- [18] S. Bandyopadhyay and M. Cahay, *Introduction to spintronics*. CRC press, 2015. 6
- [19] M. R. Mourino, “From Thales to Lauterbur, or from the lodestone to MR imaging: magnetism and medicine.” *Radiology*, vol. 180, no. 3, pp. 593–612, 1991. 7, 8
- [20] P. J. Mohr and B. N. Taylor, “CODATA recommended values of the fundamental physical constants: 2002,” *Reviews of modern physics*, vol. 77, no. 1, p. 1, 2005. 8
- [21] F. Bloch, “Nuclear induction,” *Physical review*, vol. 70, no. 7-8, pp. 460–474, Oct 1946. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.70.460> 8
- [22] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI pulse sequences*. Elsevier, 2004. 12
- [23] U. Vovk, F. Pernus, and B. Likar, “A review of methods for correction of intensity inhomogeneity in MRI,” *IEEE transactions on medical imaging*, vol. 26, no. 3, pp. 405–421, 2007. 19
- [24] D. C. Peters, D. B. Ennis, and E. R. McVeigh, “High-resolution MRI of cardiac function with projection reconstruction and steady-state free precession,” *Magnetic resonance in medicine*, vol. 48, no. 1, pp. 82–88, 2002. 21
- [25] Q. Peng, R. W. McColl, J. Wang, J. M. Chia, and P. T. Weatherall, “Water-saturated three-dimensional balanced steady-state free precession for fast abdominal fat quantification,” *Journal of magnetic resonance imaging*, vol. 21, no. 3, pp. 263–271, 2005. 21

- [26] N. K. Bangerter, T. Cukur, B. A. Hargreaves, B. S. Hu, J. H. Brittain, D. J. Park, G. E. Gold, and D. G. Nishimura, "Three-dimensional fluid-suppressed T2-prep flow-independent peripheral angiography using balanced SSFP," *Journal of magnetic resonance imaging*, vol. 29, no. 8, pp. 1119–1124, 2011. 21
- [27] K. L. Miller, "fMRI using balanced steady-state free precession (SSFP)," *NeuroImage*, vol. 62, no. 2, pp. 713–719, 2012. 21, 23, 96
- [28] N. K. Bangerter, M. Taylor, G. Tarbox, A. J. Palmer, and D. J. Park, "Quantitative techniques for musculoskeletal MRI at 7 Tesla," *Quantitative Imaging in Medicine and Surgery*, vol. 6, pp. 715–730, 2016. [Online]. Available: <http://dx.doi.org/10.21037/qims.2016.12.12> 21
- [29] H. Y. Carr, "Steady-state free precession in nuclear magnetic resonance," *Physical Review*, vol. 112, no. 5, p. 1693, 1958. 21, 27
- [30] K. Scheffler, "On the transient phase of balanced SSFP sequences," *Magnetic resonance in medicine*, vol. 49, no. 4, pp. 781–783, 2003. 22
- [31] B. A. Hargreaves, S. S. Vasanawala, J. M. Pauly, and D. G. Nishimura, "Characterization and reduction of the transient response in steady-state MR imaging," *Magnetic resonance in medicine*, vol. 46, no. 1, pp. 149–158, 2001. 22
- [32] P. W. Worters and B. A. Hargreaves, "Balanced SSFP transient imaging using variable flip angles for a predefined signal profile," *Magnetic resonance in medicine*, vol. 64, no. 5, pp. 1404–1412, 2010. 22, 24
- [33] V. S. Deshpande, Y. Chung, Q. Zhang, S. M. Shea, and D. Li, "Reduction of transient signal oscillations in true-FISP using a linear flip angle series magnetization preparation," *Magnetic resonance in medicine*, vol. 49, no. 1, pp. 151–157, 2003. 22
- [34] N. K. Bangerter, B. A. Hargreaves, S. S. Vasanawala, J. M. Pauly, G. E. Gold, and D. G. Nishimura, "Analysis of multiple-acquisition SSFP," *Magnetic resonance in medicine*, vol. 51, no. 5, pp. 1038–1047, 2004. 23, 27, 34, 89, 94, 96, 97
- [35] R. Freeman and H. D. W. Hill, "Phase and intensity anomalies in Fourier transform NMR," *Journal of magnetic resonance (1969)*, vol. 4, no. 3, pp. 366–383, 1971. 24
- [36] E. T. Jaynes, "Matrix treatment of nuclear induction," *Physical Review*, vol. 98, no. 4, p. 1099, 1955. 24
- [37] M. L. Lauzon and R. Frayne, "Analytical characterization of RF phase-cycled balanced steady-state free precession," *Concepts in Magnetic Resonance Part A: An Educational Journal*, vol. 34, no. 3, pp. 133–143, 2009. 25
- [38] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Magnetic resonance in medicine*, vol. 42, no. 5, pp. 952–962, 1999. 27, 33, 83
- [39] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic resonance in medicine*, vol. 47, no. 6, pp. 1202–1210, 2002. 27, 28

- [40] M. Buehrer, K. P. Pruessmann, P. Boesiger, and S. Kozerke, "Array compression for MRI with large coil arrays," *Magnetic resonance in medicine*, vol. 57, no. 6, pp. 1131–1139, 2007. 27
- [41] Y. Wang, X. Shao, T. Martin, S. Moeller, E. Yacoub, and D. J. J. Wang, "Phase-cycled simultaneous multislice balanced SSFP imaging with CAIPIRINHA for efficient banding reduction," *Magnetic resonance in medicine*, vol. 76, no. 6, pp. 1764–1774, 2016. 27, 89
- [42] R. Dharmakumar and G. A. Wright, "Understanding steady-state free precession: A geometric perspective," *Concepts in Magnetic Resonance Part A: An Educational Journal*, vol. 26, no. 1, pp. 1–10, 2005. 27
- [43] M. N. Hoff, "Approaches to inhomogeneity artifact correction in magnetic resonance imaging near metals," Ph.D. dissertation, University of British Columbia, 2012. 28
- [44] G. Tarbox, J. Valentine, M. Taylor, and N. K. Bangerter, "bSSFP elliptical signal model with GRAPPA parallel imaging," *Proceedings of the 25th Annual Meeting of the ISMRM; Honolulu, Hawaii, USA, 2017*. 28
- [45] M. Uecker and M. Lustig, "Estimating absolute-phase maps using ESPIRiT and virtual conjugate coils," *Magnetic resonance in medicine*, vol. 77, no. 3, pp. 1201–1207, 2017. 28
- [46] S. J. Inati, M. S. Hansen, and P. Kellman, "A solution to the phase problem in adaptive coil combination," *Proceedings of the 21st Annual Meeting of the ISMRM; Salt Lake City, Utah, USA, vol. 2672, 2013*. 28, 32, 39
- [47] F. Huang, S. Vijayakumar, Y. Li, S. Hertel, and G. R. Duensing, "A software channel compression technique for faster reconstruction with many channels," *Journal of magnetic resonance imaging*, vol. 26, no. 1, pp. 133–141, 2008. 28
- [48] E. M. Haacke, Y. Xu, Y. N. Cheng, and J. R. Reichenbach, "Susceptibility weighted imaging (SWI)," *Magnetic resonance in medicine*, vol. 52, no. 3, pp. 612–618, 2004. 28
- [49] Y. Ishihara, A. Calderon, H. Watanabe, K. Okamoto, Y. Suzuki, K. Kuroda, and Y. Suzuki, "A precise and fast temperature mapping using water proton chemical shift," *Magnetic resonance in medicine*, vol. 34, no. 6, pp. 814–823, 1995. 28
- [50] S. D. Robinson, K. Bredies, D. Khabipova, B. Dymerska, J. P. Marques, and F. Schweser, "An illustrated comparison of processing methods for MR phase imaging and QSM: combining array coil signals and phase unwrapping," *NMR in Biomedicine*, vol. 30, no. 4, p. e3601, 2017. 28
- [51] C. Ros, S. Witoszynskij, K. H. Herrmann, and J. R. Reichenbach, "Reconstruction of phase images for GRAPPA based susceptibility weighted imaging (SWI)," *Proceedings of the 16th Annual Meeting of the ISMRM; Toronto, Canada, p. 1265, 2008*. 28, 40
- [52] ———, "Reconstruction of phase images for GRAPPA accelerated magnetic resonance imaging," *4th European Conference of the International Federation for Medical and Biological Engineering*, pp. 803–806, 2009. 28

- [53] D. L. Parker, A. Payne, N. Todd, and J. R. Hadley, “Phase reconstruction from multiple coil data using a virtual reference coil,” *Magnetic resonance in medicine*, vol. 72, no. 2, pp. 563–569, 2014. 28, 40
- [54] K. Eckstein, B. Dymerska, B. Bachrata, W. Bogner, K. Poljanc, S. Trattnig, and S. D. Robinson, “Computationally efficient combination of multi-channel phase data from multi-echo acquisitions (ASPIRE),” *Magnetic resonance in medicine*, vol. 79, no. 6, pp. 2996–3006, 2018. 28
- [55] S. D. Robinson, B. Dymerska, W. Bogner, M. Barth, O. Zaric, S. Goluch, G. Grabner, X. Deligianni, O. Bieri, and S. Trattnig, “Combining phase images from array coils using a short echo time reference scan (COMPOSER),” *Magnetic resonance in medicine*, vol. 77, no. 1, pp. 318–327, 2017. 28
- [56] O. Holtz, “Compressive sensing: a paradigm shift in signal processing,” *arXiv preprint arXiv:0812.3137*, 2008. 46
- [57] D. J. Larkman and R. G. Nunes, “Parallel magnetic resonance imaging,” *Physics in medicine and biology*, vol. 52, no. 7, p. R15, 2007. 46
- [58] H. Jung, K. Sung, K. S. Nayak, E. Y. Kim, and J. C. Ye, “k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI,” *Magnetic resonance in medicine*, vol. 61, no. 1, pp. 103–116, 2009. 46
- [59] B. Wu, R. P. Millane, R. Watts, and P. J. Bones, “Prior estimate-based compressed sensing in parallel MRI,” *Magnetic resonance in medicine*, vol. 65, no. 1, pp. 83–95, 2011. 46, 55, 56, 59
- [60] R. Otazo and D. K. Sodickson, “Adaptive compressed sensing MRI,” *Proceedings of the 18th Annual Meeting of the ISMRM; Stockholm*, p. 4867, 2010. 46
- [61] J. Jiang, P. Trundle, and J. Ren, “Medical image analysis with artificial neural networks,” *Computerized Medical Imaging and Graphics*, vol. 34, no. 8, pp. 617 – 631, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0895611110000741> 46
- [62] J. Kleesiek, G. Urban, A. Hubert, D. Schwarz, K. Maier-Hein, M. Bendszus, and A. Biller, “Deep MRI brain extraction: A 3D convolutional neural network for skull stripping,” *NeuroImage*, vol. 129, no. Supplement C, pp. 460 – 469, 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1053811916000306> 46
- [63] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, “Tensorflow: A system for large-scale machine learning,” *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pp. 265–283, 2016. 46
- [64] D. Lee, J. Yoo, and J. C. Ye, “Deep residual learning for compressed sensing MRI,” *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pp. 15–18, 2017. 46

- [65] J. Sun, H. Li, Z. Xu *et al.*, “Deep ADMM-Net for compressive sensing MRI,” *Advances in neural information processing systems*, pp. 10–18, 2016. 46
- [66] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo *et al.*, “DAGAN: deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction,” *IEEE transactions on medical imaging*, vol. 37, no. 6, pp. 1310–1321, 2017. 46
- [67] Y. C. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012. 47, 49
- [68] R. G. Baraniuk, “Compressive sensing [lecture notes],” *IEEE Signal Processing Magazine*, vol. 24, no. 4, pp. 118–121, July 2007. 47
- [69] E. J. Candès *et al.*, “Compressive sampling,” *Proceedings of the international congress of mathematicians*, vol. 3, pp. 1433–1452, 2006. 47
- [70] E. J. Candès and J. Romberg, “Quantitative robust uncertainty principles and optimally sparse decompositions,” *Foundations of Computational Mathematics*, vol. 6, no. 2, pp. 227–254, 2006. 47
- [71] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on information theory*, vol. 52, no. 2, pp. 489–509, 2006. 47
- [72] ———, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on pure and applied mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006. 47
- [73] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006. 47
- [74] F. B. da Silva, R. von Borries, and C. J. Miosso, “Golden number sampling applied to compressive sensing,” *2018 IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI)*, pp. 1–4, 2018. 47
- [75] A. V. Oppenheim and R. W. Schaffer, *Digital signal processing*. Englewood Cliffs, New York, 1975. 47
- [76] F. Santosa and W. W. Symes, “Linear inversion of band-limited reflection seismograms,” *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 4, pp. 1307–1330, 1986. 47
- [77] Y. Tsaig and D. L. Donoho, “Breakdown of equivalence between the minimal l_1 -norm solution and the sparsest solution,” *Signal Processing*, vol. 86, no. 3, pp. 533–548, 2006, sparse Approximations in Signal and Image Processing. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0165168405002203> 49
- [78] B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM Journal on Computing*, vol. 24, no. 2, pp. 227–8, 04 1995. 49

- [79] D. Ge, X. Jiang, and Y. Ye, “A note on the complexity of L-p minimization,” *Mathematical programming*, vol. 129, no. 2, pp. 285–299, 2011. 49
- [80] E. J. Candès, “The restricted isometry property and its implications for compressed sensing,” *Comptes rendus mathématique*, vol. 346, no. 9-10, pp. 589–592, 2008. 49
- [81] R. R. Coveyou, “Random number generation is too important to be left to chance,” *Applied Probability and Monte Carlo Methods and modern aspects of dynamics. Studies in applied mathematics*, vol. 3, pp. 70–111, 1969. 49
- [82] M. Rani, S. B. Dhok, and R. B. Deshmukh, “A systematic review of compressive sensing: Concepts, implementations and applications,” *IEEE Access*, vol. 6, pp. 4875–4894, 2018. 50, 54
- [83] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM review*, vol. 43, no. 1, pp. 129–159, 2001. 50
- [84] W. Lu and N. Vaswani, “Modified basis pursuit denoising (modified-BPDN) for noisy compressive sensing with partially known support,” *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3926–3929, 2010. 50
- [85] R. Tibshirani, “Regression shrinkage and selection via the LASSO,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996. 50
- [86] T. Goldstein and S. Osher, “The split Bregman method for L1-regularized problems,” *SIAM journal on imaging sciences*, vol. 2, no. 2, pp. 323–343, 2009. 50, 52
- [87] A. Draganic, I. Orovic, and S. Stankovic, “On some common compressive sensing recovery algorithms and applications,” *arXiv preprint arXiv:1705.05216*, 2017. 51
- [88] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004. 51
- [89] P. C. Hansen, “The L-curve and its use in the numerical treatment of inverse problems,” 1999. 51
- [90] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: nonlinear phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992. 52
- [91] Y. Zhang, Y. Wang, and C. Zhang, “Total variation based gradient descent algorithm for sparse-view photoacoustic image reconstruction,” *Ultrasonics*, vol. 52, no. 8, pp. 1046–1055, 2012. 52
- [92] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004. 52, 53, 54
- [93] P. L. Combettes and V. R. Wajs, “Signal recovery by proximal forward-backward splitting,” *Multiscale Modeling & Simulation*, vol. 4, no. 4, pp. 1168–1200, 2005. 53, 54
- [94] A. Beck, *First-order methods in optimization*. SIAM, 2017, vol. 25. 53, 54

- [95] I. Bayram, “On the convergence of the iterative shrinkage/thresholding algorithm with a weakly convex penalty,” *IEEE Transactions on Signal Processing*, vol. 64, no. 6, pp. 1597–1608, 2015. 54
- [96] G. Adluru, Y. Gur, L. Chen, D. Feinberg, J. Anderson, and E. V. DiBella, “MRI reconstruction of multi-image acquisitions using a rank regularizer with data reordering,” *Medical physics*, vol. 42, no. 8, pp. 4734–4744, 2015. 55, 56, 57, 85, 88
- [97] L. Ramirez, C. Prieto, C. Sing-Long, S. Uribe, P. Batchelor, C. Tejos, and P. Irarrazaval, “TRIO a technique for reconstruction using intensity order: application to undersampled MRI,” *IEEE transactions on medical imaging*, vol. 30, no. 8, pp. 1566–1576, 2011. 56
- [98] I. Ram, M. Elad, and I. Cohen, “Image processing using smooth ordering of its patches,” *IEEE transactions on image processing*, vol. 22, no. 7, pp. 2764–2774, 2013. 56, 88
- [99] Z. Lai, X. Qu, Y. Liu, D. Guo, J. Ye, Z. Zhan, and Z. Chen, “Image reconstruction of compressed sensing MRI using graph-based redundant wavelet transform,” *Medical image analysis*, vol. 27, pp. 93–104, 2016. 56, 57, 59, 60, 85, 88
- [100] J. B. Kruskal, “On the shortest spanning subtree of a graph and the traveling salesman problem,” *Proceedings of the American Mathematical society*, vol. 7, no. 1, pp. 48–50, 1956. 56
- [101] S. Liu, J. Cao, H. Liu, X. Zhou, K. Zhang, and Z. Li, “MRI reconstruction via enhanced group sparsity and nonconvex regularization,” *Neurocomputing*, vol. 272, pp. 108–121, 2018. 57, 59, 60
- [102] F. Yanez and P. Irarrazaval, “Sorted compressed sensing in MRI,” *Proceedings of the 22nd Annual Meeting of the ISMRM; Milan, Italy*, vol. 22, p. 1592, 2014. 57
- [103] C. Prieto, G. Varela-Mattatall, C. Sing-Long, and P. Irarrazaval, “Gini reweighted l1 minimization for rapid MRI,” *Proceedings of the 26th Annual Meeting of the ISMRM; Paris, France*, vol. 3532, 2018. 57
- [104] D. Zonoobi, A. A. Kassim, and Y. V. Venkatesh, “Gini index as sparsity measure for signal reconstruction from compressive samples,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 5, pp. 927–932, 2011. 57, 59
- [105] N. Hurley and S. Rickard, “Comparing measures of sparsity,” *IEEE Transactions on Information Theory*, vol. 55, no. 10, pp. 4723–4741, 2009. 57, 59
- [106] G. Plonka, “The easy path wavelet transform: a new adaptive wavelet transform for sparse representation of two-dimensional data,” *Multiscale Modeling & Simulation*, vol. 7, no. 3, pp. 1474–1496, 2009. 57
- [107] Y. Zhang, J. Yang, J. Yang, A. Liu, and P. Sun, “A novel compressed sensing method for magnetic resonance imaging: exponential wavelet iterative shrinkage-thresholding algorithm with random shift,” *International journal of biomedical imaging*, vol. 2016, 2016. 61

- [108] P. Kellman, J. R. Wilson, H. Xue, M. Ugander, and A. E. Arai, “Extracellular volume fraction mapping in the myocardium, part 1: evaluation of an automated method,” *Journal of cardiovascular magnetic resonance*, vol. 14, no. 1, p. 63, 2012. 65
- [109] P. Kellman and M. S. Hansen, “T1-mapping in the heart: accuracy and precision,” *Journal of cardiovascular magnetic resonance*, vol. 16, no. 1, p. 2, 2014. 66
- [110] M. Taylor, S. Whitaker, M. N. Hoff, N. K. Bangerter *et al.*, “MRI field mapping using bSSFP elliptical signal model,” *Proceedings of the 25th Annual Meeting of the ISMRM; Honolulu, Hawaii, USA*, 2017. 89, 92
- [111] F. Liu, L. Feng, and R. Kijowski, “MANTIS: Model-Augmented Neural network with Incoherent k-space Sampling for efficient MR parameter mapping,” *Magnetic resonance in medicine*, 2019. 93
- [112] M. A. Mendoza, N. McKibben, L. Hales, M. Taylor, and N. K. Bangerter, “Banding artifact reduction in musculoskeletal bSSFP using deep learning,” *Proceedings of the 2018 ISMRM Workshop on Machine Learning II; Washington, DC*, vol. 61, 2018. 94
- [113] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, 2015. 94
- [114] M. Lustig, J. M. Santos, and J. M. Pauly, “A super-FOV method for rapid SSFP banding artifact reduction,” *Proceedings of the 13th Annual Meeting of the ISMRM; Miami Beach, Florida, USA*, p. 504, 2005. 94
- [115] M. Schär, S. Kozerke, S. E. Fischer, and P. Boesiger, “Cardiac SSFP imaging at 3 Tesla,” *Magnetic resonance in medicine*, vol. 51, no. 4, pp. 799–806, 2004. 96
- [116] S. S. Vasanawala, J. M. Pauly, and D. G. Nishimura, “Linear combination steady-state free precession MRI,” *Magnetic resonance in medicine*, vol. 43, no. 1, pp. 82–90, 2000. 97
- [117] —, “Fluctuating equilibrium MRI,” *Magnetic resonance in medicine*, vol. 42, no. 5, pp. 876–883, 1999. 97
- [118] B. A. Hargreaves, S. S. Vasanawala, K. S. Nayak, B. S. Hu, and D. G. Nishimura, “Fat-suppressed steady-state free precession imaging using phase detection,” *Magnetic resonance in medicine*, vol. 50, no. 1, pp. 210–213, 2003. 97
- [119] J. Leupold, J. Hennig, and K. Scheffler, “Alternating repetition time balanced steady state free precession,” *Magnetic resonance in medicine*, vol. 55, no. 3, pp. 557–565, 2006. 97
- [120] B. Quist, B. A. Hargreaves, T. Cukur, G. R. Morrell, G. E. Gold, and N. K. Bangerter, “Simultaneous fat suppression and band reduction with large-angle multiple-acquisition balanced steady-state free precession,” *Magnetic resonance in medicine*, vol. 67, no. 4, pp. 1004–1012, 2012. 97

APPENDIX A. COIL COMBINATION FOR ESM BSSFP

A.1 Introduction

General coil combination is simply a linear combination of coil images using some set of weights. The central conceit of Chapter 4 was that nonlinear coil combination methods (those that use nonlinear methods to derive weights) distort the elliptical relationship between multiple phase-cycled datasets that rely on the elliptical signal model (ESM). This appendix will show how these weights change for two of the most popular coil combination techniques: the Roemer method and the adaptive coil combine method (referred to in this thesis as the stochastic matched filter method). The Roemer method was not analyzed in Chapter 4 as it relies on prior coil sensitivity map estimates, but nonetheless it will be instructive to derive the geometric solution to see the complicated ways upstream coil combination can affect downstream computations.

A.2 Maximum SNR Coil Combination

Roemer et al. proposed the maximum SNR combination strategy when coil sensitivity maps, S_j , and noise correlation matrices, C_n , are known [11]. For a single phase-cycle, the combined image can be calculated pixel-wise as:

$$\hat{I}_i(x, y) = \frac{\sum_{j,k} S_j^*(x, y) C_{n,jk}^{-1} I_{i,k}(x, y)}{\sum_{j,k} S_j^*(x, y) C_{n,jk}^{-1} S_k(x, y)}, \quad (\text{A.1})$$

which can be written more succinctly when $S(x, y)$ and C_n are suitably defined matrices:

$$\hat{I}_i(x, y) = \frac{S^H(x, y) C_n^{-1}}{S^H(x, y) C_n^{-1} S(x, y)} I_i(x, y) \quad (\text{A.2})$$

$$= A I_i(x, y). \quad (\text{A.3})$$

Then A is simply the corresponding SENSE operator with an acceleration factor of 1 [38]. Notice that S and C_n do not vary with phase-cycle as they are only coil dependent.

Substituting equation (A.3) into the pixel-wise version of the linearized geometric solution (IGS) to the ESM (equation (4.2)), by linearity we have that the coil combined IGS solution is the same as coil combination prior to the IGS, as desired:

$$A I_{\text{IGS}}(x, y) = \frac{1}{2} \left[w_0(x, y) \hat{I}_0(x, y) + (1 - w_0(x, y)) \hat{I}_2(x, y) + w_1(x, y) \hat{I}_1(x, y) + (1 - w_1(x, y)) \hat{I}_3(x, y) \right]. \quad (\text{A.4})$$

What remains is to determine the weights from regional differential energy required for the IGS: $w_{0,1}$ (see Chapter 4). Substitution of equation (A.3) into the direct geometric solution (equation (4.1)) yields the following expression for $I_d(x, y)$:

$$\frac{\text{Im} [M_0^H A^H A M_2] A(M_1 - M_3) - \text{Im} [M_1^H A^H A M_3] A(M_0 - M_2)}{\text{Re} [A(M_0 - M_2, :)] \text{Im} [A(M_1 - M_3)] + \text{Re} [A(M_1 - M_3)] \text{Im} [A(M_2 - M_0)]}. \quad (\text{A.5})$$

It is clear that the resulting weights will differ from those found using a coil-by-coil strategy due to the nonlinear nature of this solution. Numerical evaluation is required to determine how the weights change and what effect this has on the final reconstruction.

A.3 Stochastic Matched Filter Coil Combination

When coil sensitivities and noise correlation matrices are unknown, the Roemer phased-array combination strategy cannot be used. To overcome this limitation, Walsh et al. proposed the maximum expected SNR solution to the coil combination problem [7]. For a single phase-cycle this solution has a simple linear form given as:

$$\hat{I}_i = m_i^H \begin{bmatrix} I_{i,0} \\ I_{i,1} \\ \vdots \\ I_{i,N-1} \end{bmatrix}. \quad (\text{A.6})$$

Each m_i is the eigenvector corresponding to the largest eigenvalue of the matrix $P = R_n^{-1} R_s$, where R_n and R_s are the autocorrelation matrices for the signal and noise processes, respectively. We

further recognize that the matched filter coefficients, m_i , are the same for each phase-cycle as the coil geometry has not been altered. Using the linearity of both the IGS reconstruction and the chosen matched filter coil combination method, it is simple to show that substituting in $I_i = \hat{I}_i$ into equation (4.2) yields the desired demodulated signal with only a single evaluation of the IGS.

Given K phase-cycled images and N coils we need to estimate m , the most significant eigenvector of $P = R_n^{-1}R_s$. We have K observations of autocorrelation matrices: R_n^l and R_s^l , $0 \leq l < K$. A first guess might be to perform coil combination separately on each phase-cycle set, i.e., find and apply each m_i . However, this is ignoring the fact that all m_i should be the same and, as a result, this method can perform poorly. Noting that the eigenvectors m_i are not additive, a simple and effective estimator of m is to use the average over the autocorrelation matrices. We recognize that the l th observation corresponds to a linear phase-cycle $\Delta\theta_l = \frac{l2\pi}{K}$, but the autocorrelation function is only sensitive to intra-phase-cycle variation; thus, we may neglect phase unwinding due to $\Delta\theta_l$:

$$R_{s,ij}^l = E \left[\left(I_i^l e^{-i\Delta\theta_l} \right) \left(I_j^l e^{-i\Delta\theta_l} \right)^* \right] \quad (\text{A.7})$$

$$= E \left[I_i^l I_j^{l*} \right], \quad (\text{A.8})$$

$$\hat{R}_s = \frac{1}{K} \sum_l R_s^l, \quad (\text{A.9})$$

and similarly for \hat{R}_n .

To find the weights, $w_{0,1}$, required for the IGS, we substitute equation (A.6) into equation (4.1) to find the direct solution, $I_d(x, y)$:

$$\frac{\text{Im} [M_0^H m m^H M_2] m^H (M_1 - M_3) - \text{Im} [M_1^H m m^H M_3] m^H (M_0 - M_2)}{\text{Re} [m^H (M_0 - M_2)] \text{Im} [m^H (M_1 - M_3)] + \text{Re} [m^H (M_1 - M_3)] \text{Im} [m^H (M_2 - M_0)]}. \quad (\text{A.10})$$

Again, due to this nonlinear solution, the weights will differ from those found using a coil-by-coil strategy and require numerical techniques for a more full analysis.

APPENDIX B. DEEP CONVOLUTIONAL NEURAL NETWORKS FOR ESTIMATION OF PRE-RECONSTRUCTION DATA REORDERINGS

B.1 Purpose

Successful reconstruction of highly undersampled datasets requires the successful navigation of an ill-posed inverse problem. Regularization is a process that introduces assumptions about the data in order to find an acceptable solution. As natural images are often piece-wise smooth, a reasonable assumption is that they are sparse in the finite differences domain, such that total variation (TV) is minimized. However, some MR datasets do not have this smoothness property. Several papers have proposed methods of predicting the pixel reordering that minimizes the reconstructed image's TV norm [9], [96], [99]. All previous methods use prior image estimates to infer the optimal reordering (e.g., reference images, low resolution estimates). It has been shown that using poor prior estimates can be detrimental to reconstructed image quality [9]. This work studies the use of a neural network to produce more optimal image reorderings for use in constrained reconstruction algorithms.

B.2 Methods

222 golden angle radially sampled 2D cardiac images from dynamic contrast enhanced cardiac perfusion MR datasets were retroactively undersampled from 72 rays to 16 rays. Non-uniform fast Fourier transform (NUFFT) reconstructions of the 72 ray datasets were used as the ground truth for training the network, which estimated the reordering for use as a prior in subsequent reconstruction. Only column reordering was considered. A U-net architecture was implemented with complex 16 ray k -space datasets as input and reordering matrices for both real and imaginary parts of the image as output (depicted in Figure B.1). The cost function for the network was the mean squared error between the true and prospective reorderings. 80% of the priors were fed as input to the neural network to train with the remaining reserved for evaluation of the network. The results

of the network were compared by reordering the columns of the 72 ray image reconstructions using both the true reorderings and reorderings based on the image priors, and then computing the anisotropic TV norm, similar to that used in [9]:

$$C(x) = \frac{1}{2} \|F_u^{-1}x - y\|_2^2 + \lambda \|\sqrt{\nabla_y x^2 + \varepsilon}\|_1, \quad (\text{B.1})$$

where ∇_y is the gradient operator across the columns, x is the image, y is the measured k -space, F_u^{-1} is the undersampled Fourier operator, λ is the regularization weight, and ε is some small positive constant to avoid singularities in the derivative of the functional.




Figure B.1: A U-net architecture with 16 ray undersampled images as input and reordering matrices (real and imaginary) as output.

B.3 Results

The network generated reorderings that reduced the anisotropic TV norm on average $13.4 \pm 1.2\%$ relative to a naive reordering taken from a low resolution prior. Figure B.2 shows (a) the magnitude image of the 72 ray NUFFT reconstruction, (b) the optimally reordered real part of the reconstruction ($\text{TV} = 1.42 \cdot 10^{-4}$), (c) the magnitude image of the prior image estimate, (d) the

reordered real part of the reconstruction based on the network’s output ($TV = 1.34 \cdot 10^{-3}$), and (e) the reordered real part of the reconstruction based on the prior image estimate ($TV = 1.68 \cdot 10^{-3}$). On each image a red line is drawn. Below the image, the real pixel intensity values along this line are plotted to demonstrate how the pixels changes across the column. The network results were consistently superior to the prior reordering estimate, although still an order of magnitude shy of the optimum.




Figure B.2: Results generated by the neural network compared to a magnitude image of the 72 ray NUFFT reconstruction.

B.4 Discussion

On inspection, it is apparent that the network-reordered image shows fewer misplaced high-intensity pixels and corresponds better to the optimal column reordering than the naive reordering. Both the network reordering and naive reordering fail to capture smooth transitions, perhaps stemming from the difficulty of sorting lower intensity pixels when highly corrupted image priors are used as input. However, it is clear from this study that more accurate reorderings may be inferred

by a neural network than from low quality image priors alone. It is expected that using larger training sets and more training iterations will yield more optimal reorderings.

B.5 Conclusion

This network demonstrates that accurate reorderings may be found for use in constrained image reconstructions that use total variation regularization [9], [96] and wavelet [99] regularization without requiring an initial high quality image estimate. This reduces preprocessing time and injects accurate prior information into the ill-posed inverse imaging problem. Preprocessing time is especially important in graph-based techniques that must solve NP-hard problems (e.g., traveling salesman) to find a reordering estimate [99]. This work only considers the simplest case of pixel reordering along one dimension. More research is needed to compare performance with other deep learning approaches, patch-based reordering approaches [98], and to study the network's performance along multiple dimensions, including graph-based reordering models [98].

APPENDIX C. SYNTHETIC BANDING FOR BSSFP DATA AUGMENTATION

Balanced Steady State Free Precession (bSSFP) MRI is a highly-efficient MRI pulse sequence but suffers from banding artifacts caused by its high sensitivity to magnetic field inhomogeneity. Many algorithms exist that can effectively remove these banding artifacts, typically by requiring multiple phase-cycled acquisitions, which increase scan time [34]. While some of the algorithms can suppress banding to some degree with two sets of phase-cycled acquisitions, much more accurate band suppression is typically achieved with at least four phase-cycled acquisitions. In this work, we present a deep learning method for synthesizing additional phase-cycled images from a set of at least two phase-cycled images that can then be used with existing band reduction techniques in order to reduce scan time.

C.1 Introduction

Banding reduction techniques for balanced steady state free precession (bSSFP) acquisition often require the collection of many phase-cycled images [34]. In general, a minimum of two phase-cycled images are required to reduce banding since with only one phase-cycled image, parts of the image may fall in a signal null. By having a combination of two different phase-cycled images it is possible to acquire signal for each voxel in a field-of-view. While a minimum of two phase-cycles is required, many banding reduction techniques and post-processing algorithms rely on the availability of at least four phase-cycled images [1], [14], [110]. However, in some applications such as dynamic contrast enhanced cardiac MR, it is not possible to collect four fully sampled phase-cycles for each slice or time point. Recently, methods have been proposed that collect multiple phase-cycled images simultaneously, greatly reducing the acquisition time required [41]. Inspired by these methods, we propose a deep learning model that generates phase-cycled images in order to reduce the number of acquired phase-cycled images. These generated images augment the existing phase-cycled datasets. The combination of acquired and augmented images can be

used in existing band reduction methods to reduce scan time. Additionally, by producing these supplemental images a variety of post processing techniques for parameter estimation such as T_1 , T_2 or field map estimates can be generated from augmented datasets.

C.2 Theory

A bSSFP data augmentation method was created using a deep neural network. Figure C.1 outlines the image reconstruction architecture for the method. A modified U-net architecture was used as the basis for the neural network. The mean squared error between the output and truth data was used as the cost function for training. As input to the network, two phase-cycled images with RF phase-cycles of 0° and 180° were used. The ground truth for the network was a pair of phase-cycled images with incremental RF phases of 90° and 270° . As verification of the experiment, the four phase-cycled images used in training were input into the geometric solution of the elliptical signal model (ESM) to generate a band free image. Then the combination of the acquired phase-cycled images (0° , 180°) and the generated phase-cycled images (90° , 270°) were input into the same ESM model. This verification procedure is outlined in Figure C.2. The resultant band reduced images were then compared.

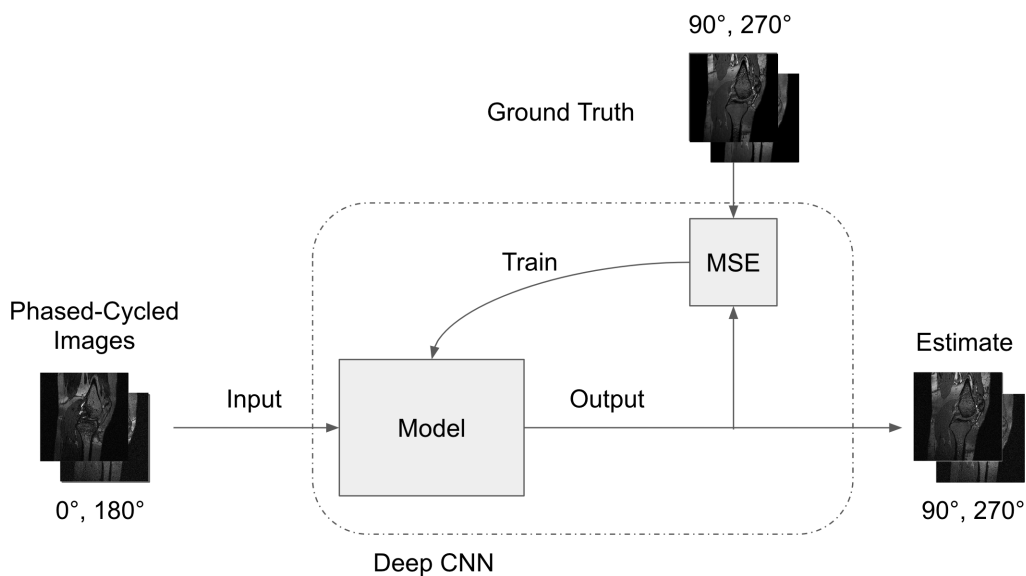


Figure C.1: An illustration of the model architecture.

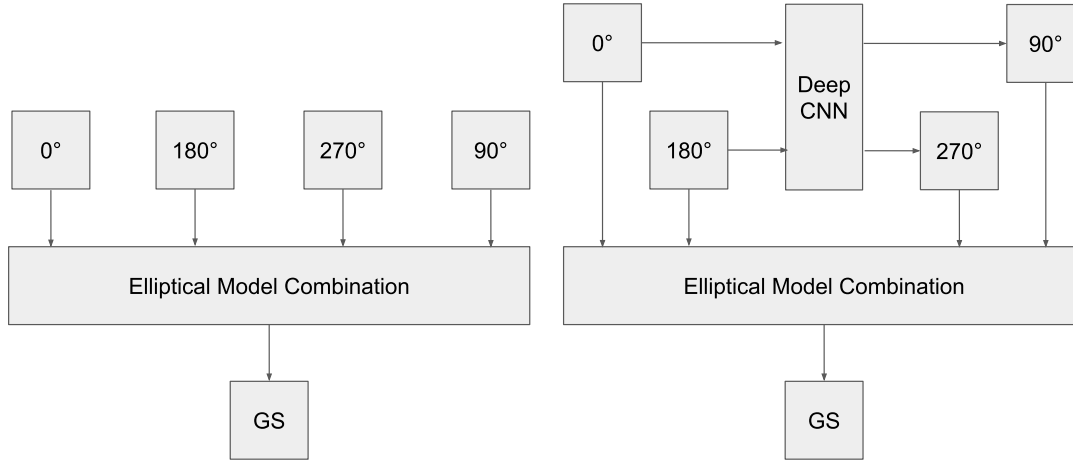


Figure C.2: An illustration of model verification using the geometric solution from elliptical signal model. (Left) Band reduced image is generated using four acquired images. (Right) Band reduced image is generated using two acquired images and two synthetically generated images.

C.3 Methods

To train the deep neural network, four phase-cycled 3D knee volumes were acquired on a Siemens 3T TIM Trio with increasing RF phase increments of 0° , 90° , 180° , and 270° . All images were acquired using a bSSFP pulse sequence with parameters $TR = 2.3\text{ms}$, $TE = 4.6\text{ms}$, a flip angle α of 50° , FOV 220mm, matrix size $256 \times 256 \times 128$, and 4 averages.

C.4 Results

The synthetically generated phase-cycled images are shown alongside the original images in Figure C.3. After training the model, the network successfully generated additional phase-cycles (90° , 270°) from the acquired phase-cycled images (0° , 180°). The generated images are comparable to the acquired images. The synthetically generated images compared to acquired images have average mean squared error of $(6.21 \pm 1.96) \cdot 10^{-4}$, mSSIM of 0.86987 ± 0.042 , and average PSNR of 38.35 ± 1.58 . As further verification of the method, Figure C.4 shows that the

reconstruction from the elliptical signal model of four acquired images exhibits band reduction similar to reconstruction from the elliptical signal model of two acquired images combined with the synthetically generated images.




Figure C.3: bSSFP knee images. (a–d) are the acquired bSSFP phase-cycled images (0° , 180° , 90° and 270°) (e–f) are the synthetically generated bSSFP phase-cycled images (90° and 270°).




Figure C.4: bSSFP band reduced images using the geometric solution to the elliptical signal model. (a) Inputs into the model are four acquired bSSFP phase-cycled images. (b) Inputs into the model are two acquired and two synthetically generated bSSFP phase-cycled images.

C.5 Conclusion and Discussion

We have demonstrated that given two phase-cycled bSSFP acquisitions, a deep convolutional neural network can be trained to augment the dataset to four phase-cycled images without the need for additional acquisitions. Using a set of four images consisting of acquired and augmented phase-cycled images, ESM-based algorithms can be used to eliminate banding [1]. Additionally, this data is suitable as input to algorithms estimating tissue parameters and field maps [2], [110].

We have also demonstrated that correct phase information is synthesized, as the solution to the ESM is very close to the true solution, allowing phase information to propagate through a post-processing pipeline. Future work could involve more complicated network architectures: recent work suggests that sharper images may be generated by generative adversarial networks [111]. Some algorithms also require or are improved with more than four input phase-cycled images, leading to the possibility of synthesizing more than two phase-cycled output images [14].

C.6 Acknowledgements

We would like to thank Bradley Bolster from Siemens for his support of the BYU MRI Research Facility.

APPENDIX D. DEEP LEARNING SUPER-FOV FOR ACCELERATED BSSFP BANDING REDUCTION

We present a technique for bSSFP band removal using two undersampled phase-cycled bSSFP image acquisitions.

D.1 Introduction

Balanced Steady State Free Precession (bSSFP) is an imaging technique capable of fast image acquisition with a high signal-to-noise ratio (SNR). However, bSSFP is highly sensitive to magnetic field inhomogeneities. This sensitivity is known to create characteristic banding artifacts which degrade image quality. Several methods for reducing banding artifacts have been developed using the combination of multiple phase-cycled images in order to effectively reduce the banding. The acquisition of multiple phased-cycled images increases scan time and is a limiting factor. We present a technique for state-of-the-art band reduction using undersampled bSSFP acquisitions combined with synthetic images generated from a deep neural network. This method uses fewer images compared to some techniques for band reduction, like the Elliptical Signal Model, and is capable of reducing scan time [34]. We validate our technique with in-vivo experimental results.

D.2 Theory

A banding artifact reduction method was created using a deep neural network. Previously it was shown that such a network could be designed to effectively reduce banding artifacts using a U-Net architecture. [112], [113]. The ground truth data was generated using four phased-cycled images combined with the geometric solution to the elliptical signal model in order to generate a band reduced image [1]. Using this work as a basis, this network was improved to utilize undersampled SSFP images as the inputs to the network to solve a generalized SENSE reconstruction. This improvement is analogous to previous work for SSFP banding artifact reduction Super-FOV [114].

In the Super-FOV method, the banding artifacts of SSFP are used as a sensitivity map in order to reconstruct a pair of undersampled phase-cycled bSSFP images. The sensitivity maps are used in a similar manner to other subsampled reconstruction techniques like GRAPPA and SENSE. Figure D.1 outlines the imaging and reconstruction process for this technique.

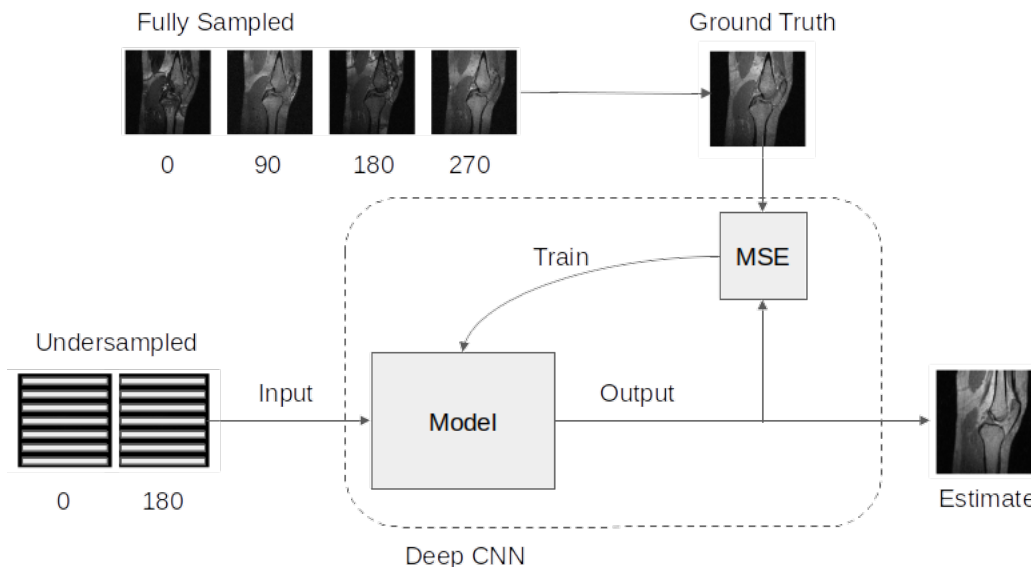


Figure D.1: Undersampled k -space data from two phase-cycled acquisitions is fed to the model as input. The output is compared using mean squared error to the geometric solution to the elliptical signal model generated from four fully sampled phase-cycled acquisitions.

D.3 Methods

We acquired sets of four phase-cycled images using a 3T Siemens TIM Trio scanner. This experiment produced the data necessary for both training and testing the model. The model was trained using two undersampled phase-cycled images and the truth data was generated from the results of the geometric solution to the elliptical signal model. As inputs to the elliptical model, four images corresponding to the RF phase increments of 0° , 90° , 180° , and 270° from the coronal slice of a knee were acquired. This data was acquired using a FOV of 250 mm, a $TR = 4.6\text{ms}$, a $TE = 2.3\text{ms}$, a tip angle of 55° , and an acquisition matrix size of $256 \times 256 \times 128$. Two of the bSSFP images with RF phase increments 0° and 180° were undersampled retroactively with an undersampling factor of $R = 2$.

D.4 Results

The results of the experiment are shown in Figure D.2. After training the model, the reconstruction exhibits band reduction similar to the geometric solution of the elliptical signal model. The network generated banding-free image compared to the true solution to the elliptical signal model has average mean squared error of $(5.38 \pm 2.06) \cdot 10^{-7}$, mSSIM of $0.9998 \pm 6.66 \cdot 10^{-5}$, and average PSNR of 69.02 ± 1.61 . We have demonstrated that a deep neural network can reduce banding artifacts in multiple acquisition bSSFP with comparable performance to the elliptical signal model and reduce overall scan time by requiring fewer phase-cycled images.




Figure D.2: (a–b) is the combined k -space and image domain representation, respectively, from the two undersampled acquisitions (even lines from first phase-cycled image, odd lines from the second) as it was input to the neural network for training. (c–d) show the k -space and image space representations of the geometric solution to the elliptical signal model which was used as ground truth for training. (d–e) shows the model’s estimate of the banding-free image.

D.5 Conclusion and Discussion

There are a variety of applications which could benefit from the excellent bSSFP contrast and SNR efficiency, but which cannot tolerate severe banding at low repetition times such as in musculoskeletal imaging [34], dynamic cardiac MR [115], and transition band functional MRI [27]. This technique might be improved by measuring an auto-calibration signal in the center of k -space in addition to the undersampled acquisitions taken. This would increase acquisition time, but allow for much better estimation of the phase sensitivity maps required for reconstruction. The model could also be improved with the addition of an additional decoder that reconstructs the fully sampled input images, allowing for the inclusion of a data fidelity term in the training cost function.

APPENDIX E. GENERATION OF ARBITRARY SPECTRAL PROFILES USING ORTHONORMAL BASIS COMBINATIONS OF BSSFP MRI

We present a technique for generating an arbitrary spectral profile by using multiple-acquisition bSSFP. Multiple phase-cycled bSSFP images with increasing TRs were acquired and Gram-Schmidt orthogonalization was applied to spectral basis functions to generate an orthonormal basis. This generated orthonormal basis was used to approximate an arbitrary spectral profile by using linear combinations of the calculated basis functions. A variety of spectral functions were simulated and used as a template to approximate spectral profiles in water and fat phantoms.

E.1 Introduction

Balanced steady-state free precession (bSSFP) is a high SNR, fast imaging technique suffering from banding artifacts which degrade image quality. These banding artifacts are a function of off-resonance frequency and magnetic field inhomogeneities. Over the past two decades, several methods have been developed using multiple-acquisition bSSFP to either reduce the appearance of banding artifacts [1], [34], [116] or to move the location of a banding artifacts to suppress undesired off-resonance frequencies [117]–[120].

This work focuses on developing a novel technique for generating an arbitrary off-resonance spectral profile using multiple-acquisition bSSFP.

E.2 Theory

Using multiple phase-cycled bSSFP images with increasing TR, it is possible to form an orthonormal basis from which to approximate any arbitrary spectral profile. An orthonormal basis is a set of linearly independent vectors that can be used to span the complex numbers. A linear combination of basis vectors, f_n , can be used to approximate any analytic function f according

to its basis expansion: $f \approx \hat{f} = \sum a_n f_n$. Increasing the number of terms in the basis expansion decreases the approximation error.

At large flip angles, bSSFP images have a sinusoidally-shaped spectral profiles for certain tissues with a pass band region and region of signal null. Since the bSSFP profile is similar to sines and cosines, we can use multiple phase-cycled acquisitions to generate a basis similar to a Fourier series. The spectra used to generate this basis are shown in Figure E.1. As this set of vectors is not exactly sinusoidal, it is not necessarily orthonormal. We therefore use Gram-Schmidt orthogonalization to generate an orthonormal basis. At moderate flip angles, the spectra generated appear less sinusoidal but can still form an acceptable approximating basis set (see Figure E.2).

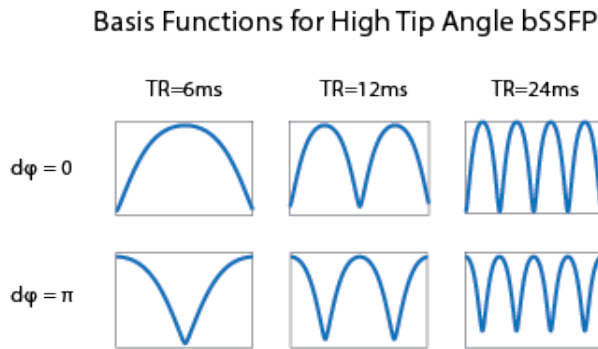


Figure E.1: Magnitude plots of six basis functions generated from a high tip angle (90°) SSFP simulation. The shape of this spectral profile was observed in images from the fat phantom. The top three basis functions are similar to the sine terms of the Fourier series, while the bottom three basis functions are similar to the cosine terms of the Fourier series.

Multiple-echo acquisitions can be used to manipulate the spectral profile of a bSSFP image in order to create different contrast features and to suppress unwanted off-resonant frequencies (e.g., fat suppression). Using this technique we can approximate an arbitrary band-pass and notch filter as shown in Figures E.3 and E.4.

E.3 Methods

Six phase-cycled bSSFP images were acquired of a fat and water phantom using a high tip angle (90°). These images were generated using two phase-cycled values ($d\phi = 0^\circ, 180^\circ$) at three

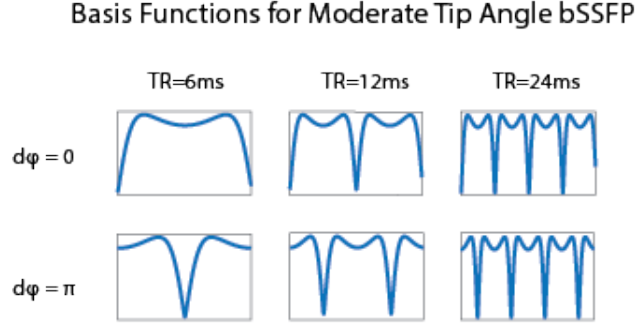


Figure E.2: Magnitude plots of six basis functions generated from a moderate tip angle SSFP simulation. The shape of this spectral profile was observed in images from the water phantom. The top three basis functions are odd functions, while the bottom three basis functions are even functions. Despite the unusual shape of this set of basis functions, it can be used to generate a variety of arbitrary off-resonance spectra after it has been orthogonalized with Gram-Schmidt orthogonalization.

different TR values, namely $TR = 6\text{ms}$, 12ms , and 24ms , with $TE = TR/2$. We then used Gram-Schmidt orthogonalization to generate an orthonormal basis. Using the desired spectral profile as a template function, the basis function coefficients were selected to approximate the desired bSSFP spectrum. These basis coefficients were computed by solving a linear system of equations, $Ac = f$, with the basis forming the column space of A , f the desired function, and c the coefficients.

E.4 Results

After generating spectra for both the notch and band-pass filters, the resulting linear combinations shown in Figures E.3 and E.4 show successful filtering of the phantoms. In both the fat and water phantoms, the spectral profile was observed to be similar to those shown in Figures E.1 and E.2, respectively.

In fat, the notch filter was well approximated by the simulated and measured data, with side ripple levels below 10%. The band-pass, with high frequency content, was approximated using a near sinusoidal basis, performing similarly to a Fourier series expansion. The resultant images show the spectrum profiles with strong induced signal peaks and nulls. The spectral approximations for the notch and band-pass filters in water were less accurate, with ripple levels in excess of 20%. Although the measured data deviates from the simulated spectra, the general profile is observed in both filters and the signal peaks and nulls are pronounced in the resultant image.

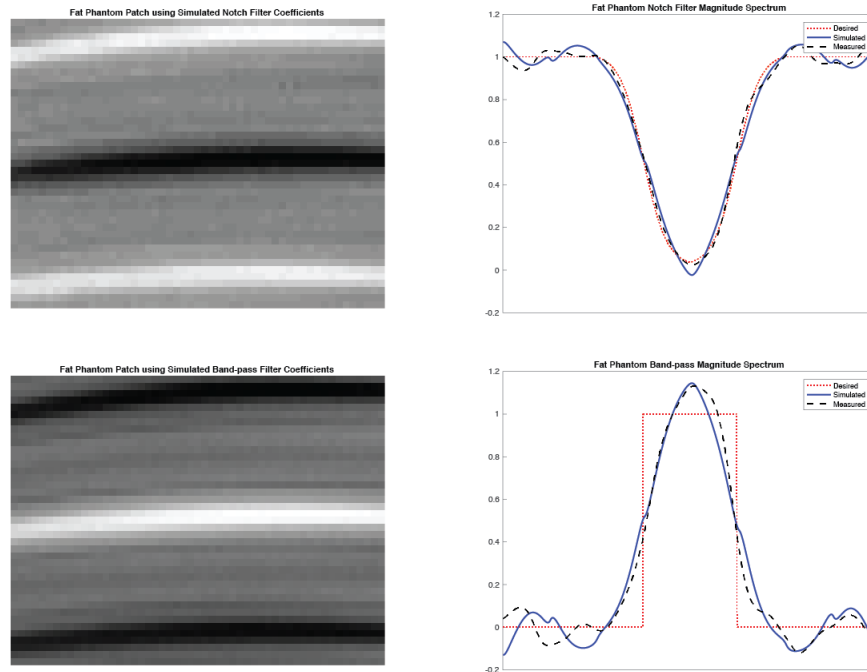


Figure E.3: Generated spectra from a fat phantom using the high flip angle basis functions. A shim gradient was applied from the top to the bottom of each image to create a range of off-resonance frequencies. The top row shows the spectral profile generated using a notch spectrum while the bottom row shows the spectral profile generated from a band-pass spectrum.

E.5 Discussion

The method successfully generates off-resonance spectral profiles of both notch and band-pass filters in fat and water using multiple-acquisition bSSFP. Coefficients may be generated for arbitrary functions and may be more accurately approximated using additional terms. However, each term in the basis expansion corresponds to a new bSSFP acquisition at an increased TR value. High resolution is required to resolve banding artifacts for high TR values, leading to long acquisition times for desired spectrum profiles with high frequency content.

Fat suppression could be achieved with this technique by creating a notch filter spectrum and shifting the notch to the off-resonance frequency needed to suppress fat. This process requires the recalculation of basis coefficients such that the location of the null is steered to the desired location, similar to beam forming in phased-array antennas. In addition, accurate pixel-by-pixel reconstructions require a field map due to field inhomogeneities, further increasing acquisition time.

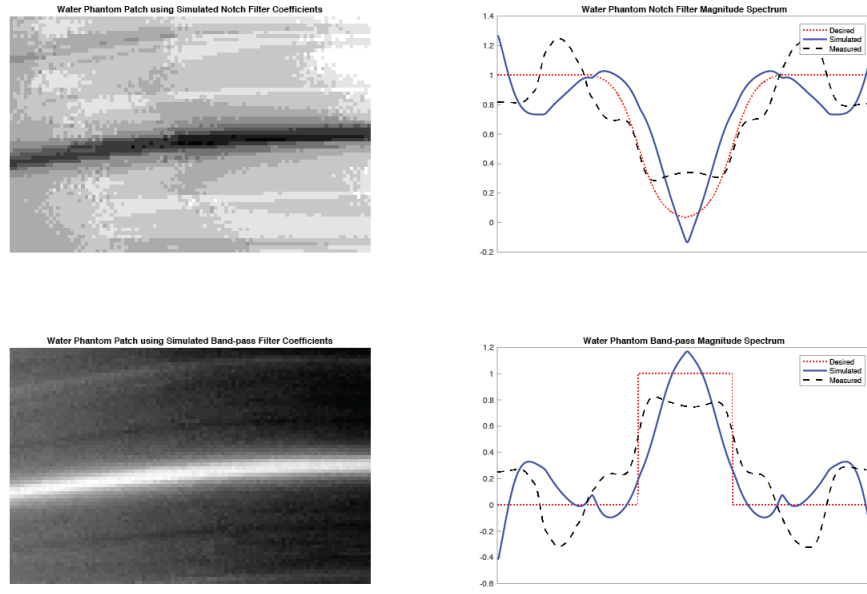


Figure E.4: Generated spectra from a water phantom using the moderate flip angle basis functions. A shim gradient was applied from the top to the bottom of each image to create a range of off-resonance frequencies. The top row shows the spectral profile generated using a notch spectrum while the bottom row shows the spectral profile generated using a band-pass spectrum.

E.6 Acknowledgements

We would like to thank Bradley Bolster from Siemens for his support of the BYU MRI Research Facility.