

NOVEL METHODS FOR STEADY-STATE
NEUROIMAGING

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Karla L. Miller
April 2004

© Copyright by Karla L. Miller 2004
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

John M. Pauly
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Brian A. Hargreaves

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Dwight G. Nishimura

Approved for the University Committee on Graduate Studies.

Abstract

Magnetic Resonance Imaging (MRI) has had a major impact on neuroscience due to the flexibility of MRI contrast. In particular, functional MRI (fMRI) and diffusion-weighted imaging (DWI) have provided neuroscientists with unique and powerful methods for studying brain function and structure. However, neither of these methods has the signal-to-noise ratio (SNR), spatial resolution or image quality achieved by many other MRI methods. This work presents new techniques for fMRI and DWI that are based on steady-state free precession (SSFP) imaging. SSFP achieves unusually high SNR efficiency by allowing the magnetization to persist from one RF excitation to the next. In addition, SSFP methods tend to have unusual signal properties that can be used to manipulate image contrast.

The standard method for fMRI, using the Blood Oxygenation Level Dependent (BOLD) effect, has significant limitations that result from the coupling of functional contrast to sources of image artifact. We have developed a method that uses the SSFP phase profile to invert the signal in deoxygenated blood relative to oxygenated blood. The resulting Blood Oxygenation Sensitive Steady-state (BOSS) signal decouples functional contrast from imaging, enabling significantly better image quality than BOLD fMRI. BOSS fMRI is very SNR-efficient, achieves strong functional contrast and is relatively immune to susceptibility gradients.

Standard DWI methods suffer from low SNR, poor image resolution and severe image artifacts. Although it has long been known that SSFP DWI is one of the most efficient DWI pulse sequences, this method has traditionally suffered from intense

sensitivity to brain motion. Motion in DWI causes phase offsets in the magnetization that create image artifacts in multi-shot acquisitions. This work has developed a method of removing phase artifacts from SSFP DWI using navigator echoes. Navigated SSFP DWI produces high-resolution diffusion images with high SNR and no visible artifact.

A major component of this work was the development of a novel navigator correction that removes phase artifacts caused by non-rigid motion based on an approximation to the least-squares optimal reconstruction. This method, referred to as a refocusing reconstruction, produces significantly better images than standard rigid-body corrections and should be applicable to other navigated DWI sequences.

Acknowledgements

A great many people deserve acknowledgement for their contributions to the work presented in this thesis.

The greatest thanks go to my adviser, John Pauly. John's experience and talent has contributed to every page of this thesis. He has been incredibly patient and generous with his time, and throughout my graduate studies has strived to prepare me for all aspects of a research career. I am continually buoyed by John's confidence in me, and I am honored to be his first doctoral student.

If I had a second adviser at Stanford, it was Brian Hargreaves. Brian is a talented researcher and a perceptive teacher, and he provided an excellent role model for a great many students in the lab, myself included. The effort that Brian puts into developing intuition for a problem makes him an effective and insightful adviser, and this thesis has benefitted greatly from the enormous amount of time that Brian has spent with me.

The many bright and engaging individuals at MRSRL have created a fantastically stimulating environment, and I cannot imagine a better place to begin a research career. While I extend my gratitude to all of my labmates at MRSRL, a few individuals deserve special thanks. Like so many other young engineers at Stanford, I have benefitted from Dwight Nishimura's skills as a teacher, both in formal coursework and informal advising. Dwight has imparted perspective on the breadth of the field, and I am continually impressed by his curiosity and astuteness. Dwight was kind enough to read this thesis, for which I am particularly grateful. Steve Connolly has been a

constant source of interesting ideas and helpful criticisms, in addition to thoughtful career advice. Bill Overall informed many of my thoughts on steady-state imaging and set high standards for those of us who followed him. I have also benefitted from the experience and wisdom of Craig Meyer, Greig Scott, Al Macovski, Bob Hu and Krishna Nayak. Finally, my fellow students and the other members of MRSRL deserve thanks for the many ways that they have made my time at Stanford more enjoyable and fulfilling.

I am also very grateful to members of the international MRI community who have provided outside perspective on this work. Most prominently, Rick Buxton has been a father figure throughout my research experiences. I deeply respect his integrity as a scientist and his clarity of thought, and I have often attempted to model my approach to research after his. Larry Frank, Eric Wong, Jim Pipe and Peter Jezzard have also provided stimulating viewpoints on my work, and their insights have often led me to explore new areas of a research problem.

Many thanks are to be extended to my collaborators. Christopher deCharms, Jongho Lee and David Ress were helpful in the initial BOSS fMRI experiments, and Garry Gold was a patient and positive collaborator on the knee imaging project.

Finally, I would like to dedicate this thesis to my mother. Mum has been my most ardent supporter, and her faith in me has been a continual source of strength.

Contents

Abstract	iv
Acknowledgements	vi
1 Introduction	1
1.1 A Brief History of MRI	2
1.2 Motivation	3
1.3 Outline	4
2 Magnetic Resonance Imaging	6
2.1 MR Physics	6
2.1.1 Polarization	7
2.1.2 Precession	8
2.1.3 Excitation	9
2.1.4 Relaxation	10
2.1.5 The Bloch Equation	12

2.2	Image Formation	12
2.2.1	Gradient Fields	12
2.2.2	Spatial Encoding	13
2.2.3	Encoding Trajectories	15
2.3	Steady-state MRI	16
2.3.1	Steady State Formation	16
2.3.2	Balanced SSFP	17
2.3.3	Matrix Representation of Balanced SSFP	18
2.3.4	Coherence Pathway Analysis	20
3	Neuroimaging With MRI	24
3.1	Neuroscience Applications of MRI	24
3.2	Functional MRI	25
3.2.1	Functional Neuroimaging Methods	25
3.2.2	BOLD Contrast	26
3.2.3	fMRI Experiments	26
3.2.4	Limitations of BOLD fMRI	28
3.3	Diffusion-Weighted Imaging	30
3.3.1	Diffusion Weighting	30
3.3.2	Pulse Sequences	32
3.3.3	Diffusion Tensor Imaging	34

3.3.4	Limitations of DWI	34
4	Balanced SSFP for Functional MRI	36
4.1	Overcoming the Limitations of BOLD	36
4.2	BOSS FMRI	37
4.2.1	FMRI with Frequency-sensitive Nulling	38
4.2.2	FMRI with Frequency-sensitive Inversion	39
4.2.3	Inversion vs. Nulling	41
4.3	Experimental Validation	41
4.3.1	Experimental Methods and Data Analysis	41
4.3.2	Single- and Multi-frequency Measurements	43
4.3.3	Negatively-correlated Signal Changes	45
4.3.4	High-resolution BOSS	45
4.4	Theoretical Comparison with BOLD	48
4.4.1	Functional Contrast	49
4.4.2	Image Quality	50
4.4.3	Signal-to-noise Ratio	52
4.5	BOSS Signal Characteristics	53
4.5.1	Contrast-to-noise Ratio	53
4.5.2	BOSS Contrast Optimization	56
4.5.3	Balanced-SSFP Phase Transition	57

4.5.4	BOSS Imaging Parameters	61
4.6	Summary	62
5	Navigated Diffusion Imaging with SSFP	63
5.1	SSFP-DWI	63
5.1.1	Diffusion Contrast	64
5.1.2	Relaxation Effects	66
5.1.3	Quantification	67
5.2	Navigated SSFP-DWI	70
5.2.1	Spiral-navigated Spinwarp Trajectories	71
5.2.2	Self-navigated Spiral Trajectories	73
5.2.3	Factors Affecting the Choice of Navigated Trajectory	73
5.3	SSFP-DWI of the Brain	76
5.3.1	Implementation	76
5.3.2	2DFT Images	76
5.3.3	Self-navigated Spiral Images	77
5.3.4	Cardiac-gated Images	78
5.4	SSFP-DWI of Cartilage	79
5.4.1	Navigated 3D SSFP-DWI for Cartilage	81
5.4.2	Cartilage Imaging Results	82
5.5	Summary	84

6	Non-linear Phase Correction of Diffusion MRI	86
6.1	Motion in DWI	86
6.2	Navigator Correction of Non-linear Phase Artifacts	89
6.2.1	Least-squares Reconstruction	89
6.2.2	Refocusing Reconstruction	93
6.3	Implementing the Refocusing Reconstruction	94
6.3.1	Efficient Reconstruction	94
6.3.2	Cardiac Synchronization	97
6.4	Experiments and Results	98
6.4.1	Experiments	98
6.4.2	Refocusing Reconstruction	101
6.4.3	Cardiac Triggering	101
6.4.4	Diffusion Tensor Imaging	103
6.5	Characterization of the Refocusing Reconstruction	103
6.5.1	Fidelity of the Refocusing Reconstruction	103
6.5.2	Cardiac Synchronization	106
6.5.3	Computational Efficiency	108
6.5.4	Weighted Refocusing Reconstruction	110
6.6	Extending the Refocusing Reconstruction	112
6.6.1	Generality of Reconstruction	112
6.6.2	Application to Other Sequences and Trajectories	113

6.7	Summary	114
7	Summary and Recommendations	115
7.1	Summary of Contributions	115
7.2	Recommendations for Future Work	116
7.2.1	BOSS FMRI	116
7.2.2	Navigated SSFP-DWI	118
7.2.3	Refocusing Reconstruction of Navigated DWI	119
A	Early diagnosis of osteoarthritis	122
B	Matrix representation of motion corruption	124
	Bibliography	128

List of Tables

4.1	Relative SNR efficiency of BOLD and BOSS images.	54
-----	--	----

List of Figures

2.1	Polarization of spins in an external field.	7
2.2	Precession of spins.	8
2.3	Excitation of magnetization.	10
2.4	Basic pulse sequence.	14
2.5	Common k-space trajectories.	15
2.6	Approach to steady-state signal.	17
2.7	Balanced-SSFP imaging sequence.	18
2.8	Simplified balanced-SSFP imaging sequence.	19
2.9	Balanced-SSFP signal profile.	21
2.10	Banding of balanced-SSFP images.	22
2.11	Coherence pathway analysis of steady-state magnetization.	23
3.1	BOLD contrast mechanism.	27
3.2	BOLD artifacts near susceptibility boundaries.	29
3.3	Bipolar diffusion gradients.	31
3.4	DWI pulse sequences.	33

4.1	Balanced-SSFP pulse sequence.	37
4.2	Balanced-SSFP signal profiles.	40
4.3	Single- and multi-frequency BOSS FMRI results.	44
4.4	Positive and negative correlations in BOSS FMRI.	46
4.5	High resolution BOSS FMRI.	47
4.6	Improving coverage with multi-frequency measurements.	48
4.7	Comparison of BOSS and BOLD image quality.	51
4.8	BOSS functional CNR profiles.	55
4.9	Functional CNR profile comparison.	57
4.10	Dynamics of SSFP phase transition.	59
4.11	Parameters affecting the SSFP phase profile.	60
5.1	Diffusion contrast in SSFP-DWI.	65
5.2	Relaxation in SSFP-DWI.	67
5.3	Sensitivity of diffusion coefficient fit to flip angle and T_1	69
5.4	Navigated SSFP-DWI pulse sequences.	72
5.5	Comparison of spiral and EPI navigators.	75
5.6	Ungated, navigated 2DFT SSFP-DWI images.	77
5.7	Ungated, self-navigated spiral SSFP-DWI images.	78
5.8	Gated, navigated 2DFT SSFP-DWI images.	80
5.9	3D SSFP-DWI images of knee cartilage.	83

5.10	3D SSFP-DWI images before and after navigator correction.	84
5.11	Diffusion-coefficient maps of knee cartilage.	85
5.12	Cross-section of the diffusion coefficient across knee cartilage.	85
6.1	Effects of motion on diffusion-weighted magnetization.	88
6.2	Least-squares reconstruction of navigated diffusion data.	91
6.3	The refocusing reconstruction in image-space and k -space.	95
6.4	Effect of 2DFT phase encode ordering on k -space.	99
6.5	Navigated steady-state DWI pulse sequence.	100
6.6	Navigated steady-state DWI images (ungated).	102
6.7	Gated, navigated steady-state DWI images (tensor axes).	104
6.8	Motion-corrected trace and FA images.	105
6.9	Effect of averaging on mixing impulse responses.	107
6.10	Effect of cardiac synchronization on the mixing impulse response.	109

Chapter 1

Introduction

Magnetic resonance imaging (MRI) is a method for imaging living anatomy and function that has rapidly become one of the most powerful diagnostic tools in modern medicine. The chief advantage of MRI over other medical imaging technologies (such as CT, ultrasound or PET) lies in its flexibility. Relative to other imaging methods, MRI is unconstrained physically: MRI can obtain images at arbitrary angles, selectively image subvolumes within a larger object and does not suffer from “shadowing” due to opaque structures. Perhaps more impressive is the contrast in MRI images, which is remarkably flexible. While MRI was initially used to visualize basic anatomical structures, many of the techniques developed for nuclear magnetic resonance (NMR) spectroscopy were adapted and enhanced to enable the visualization of physiological processes and function. Finally, MRI is generally non-invasive, but can also be used in conjunction with contrast agents or other imaging modalities to further increase its flexibility.

1.1 A Brief History of MRI

Magnetic resonance imaging (MRI) is based on nuclear magnetic resonance (NMR) spectroscopy, which is a method for inferring molecular structure based on the intrinsic angular momentum (or, “spin”) possessed by subatomic particles. NMR grew out of optical methods, with the major innovation of radiofrequency detection based on technology developed during World War II. This spectroscopy method was developed independently by Edward Purcell [1] and Felix Bloch [2] (who referred to the method as “nuclear induction”). NMR quickly became a popular method for elucidating molecular structures due to its flexibility and precision. The importance of NMR spectroscopy to chemistry is reflected in the number of Nobel Prizes awarded for NMR methodology: Bloch and Purcell’s joint Nobel Prize in Physics (1952), Ernst’s Nobel Prize in Chemistry (1991), and Wüthrich’s Nobel Prize in Chemistry (2002).

Imaging with NMR was first conceived of in 1973 by Paul Lauterbur, who proposed adding spatially-varying magnetic fields (called gradients) to an NMR spectrometer to induce spatially-varying resonance frequency [3]. Using detection coils tuned to the resonance frequency of hydrogen, the density of water in a given location could be determined by simply “listening” to the correct frequency. By making a number of measurements while varying the strength of the gradient fields, Lauterbur was able to make an image. Sir Peter Mansfield subsequently developed methods for analysis and rapid acquisition of images with this technique, which eventually became known as Magnetic Resonance Imaging (MRI). Later work extended this basic method to other nuclei, to create a broader field which is generally known as spectroscopic imaging.

MRI has since become one of the most powerful tools for radiological diagnosis. Currently, more than 60 million MRI scans are performed worldwide every year. A few months prior to the writing of this thesis, Lauterbur and Mansfield were awarded the 2003 Nobel Prize in Physiology or Medicine for “pioneering contributions, which later lead to the applications of magnetic resonance in medical imaging.”

1.2 Motivation

MRI has had a major impact on clinical neurology by providing a variety of tools for studying the structure and function of the human brain. The majority of clinical MRI scans currently performed are of the nervous system; in addition, MRI is increasingly used within cognitive science to non-invasively study the functioning human brain. Although the latter type of imaging has had a major impact on neuroscience, for various reasons the methods developed have not been able to achieve the image quality or statistical robustness that MRI is capable of in general. Two of the major advances in MR neuroimaging have been the development of functional MRI (fMRI) and diffusion-weighted imaging (DWI), both suffer from low signal-to-noise ratio (SNR), poor spatial resolution and image warping or blurring. While solutions to one or another of these problems have been suggested, the proposed approaches often incur large increases in scan time or exacerbate other problems.

Recent advances in gradient technology have advanced fast imaging methods and enabled the realization of a particular type of sequence known as “steady-state” imaging. Steady-state sequences are characterized by RF and gradient pulses which are repeated sufficiently quickly to strongly influence the effect of subsequent pulses. Because the magnetization has not fully recovered from the effect of one RF pulse before the next occurs, fairly complicated signal dynamics develop that do not occur in slower sequences. The resulting equilibrium (or “steady-state”) magnetization has unusual properties. Often, one simply wishes to use these sequences to image quickly and efficiently, but the unusual properties of these sequences can also be used to manipulate image contrast. In either case, the strong signal and speed of these sequences makes them an attractive option for imaging applications that suffer from low signal levels or inefficient readouts.

1.3 Outline

This thesis describes new methods for neuroimaging based on SSFP imaging. The methods developed have specific application to cognitive neuroscience and provide tools for studying the relation of brain structure to cognitive function. In particular, new methods for fMRI and DWI of the brain based on SSFP imaging are introduced. Both fMRI and DWI suffer from low SNR, poor spatial resolution and image artifacts. The high SNR efficiency of SSFP and the ability to acquire images quickly without requiring long readout windows suggest that it might be able to address the limitations of standard fMRI and DWI methods. However, traditional fMRI and DWI contrast both require long echo times, which is incompatible with SSFP. Instead, the methods presented in this work achieve similar contrast based on the unusual spin-manipulation properties of SSFP imaging. fMRI contrast is based on the sensitivity of SSFP to resonance frequency. DWI contrast is based on a plurality of diffusion-weighted echoes. The latter sequence is strongly sensitive to motion, and a new method for removing motion artifacts from DWI is presented.

This thesis is organized as follows:

Chapter 2: Magnetic Resonance Imaging The basic concepts necessary to understand MRI are presented briefly. These ideas are explored in more detail with specific application to imaging sequences based on steady-state free precession. Methods for calculating the steady-state magnetization are presented, and the chapter concludes with a discussion of the advantages and issues of SSFP imaging.

Chapter 3: Neuroimaging with MRI This chapter presents background on the application of MRI to specific neuroimaging problems. Overviews of the contrast mechanisms of conventional fMRI and DWI are given, and the relevant clinical and neuroscience applications are outlined. Some limitations of the methods are discussed and the potential advantages of SSFP-based methods are presented.

Chapter 4: Balanced SSFP for Functional MRI We introduce a novel method for fMRI that obtains functional contrast based on the sensitivity of the SSFP signal phase to resonance frequency. This method, called BOSS fMRI, inverts the signal from deoxygenated blood relative to oxygenated blood, creating a signal level that depends on local oxygen metabolism. BOSS has several advantages over the more conventional BOLD contrast mechanism, including higher SNR, fewer image artifacts, and reduced sensitivity to susceptibility boundaries. BOSS is more sensitive to spatial and temporal drift in the resonance frequency, and methods for overcoming this sensitivity are discussed.

Chapter 5: Navigated Diffusion Imaging with SSFP In SSFP-DWI, the signal is made strongly sensitive to water diffusion based on a steady-state consisting of a plurality of echoes. This chapter presents navigated SSFP-DWI, which has reduced sensitivity to motion and enables high-resolution diffusion images with strong diffusion contrast. In addition to high-quality diffusion images of the brain, this method has been used to obtain the first clinical-quality *in vivo* diffusion images of knee cartilage. A method for quantifying diffusion with this sequence is also presented.

Chapter 6: Nonlinear Phase Correction of Diffusion MRI Sensitivity to subject motion is a major issue for all diffusion imaging sequences since sensitivity to small, diffusive motion necessarily implies sensitivity to larger bulk motions. Motion artifacts can be removed from DWI data using a phase reference known as a navigator. This chapter presents an improved method of navigator correction, called a refocusing reconstruction, which is able to remove more complicated motion artifacts than previous reconstructions. Methods for combining the refocusing reconstruction with cardiac gating are presented, and the correction is found to remove artifacts from SSFP-DWI more effectively than standard corrections.

Chapter 7: Summary and Future Work The major contributions described in this thesis are reiterated, and recommendations for further work are presented.

Chapter 2

Magnetic Resonance Imaging

2.1 MR Physics

Nuclear magnetic resonance arises from the intrinsic angular momentum (spin) possessed by all subatomic particles. Protons, neutrons and electrons all possess a quantized amount of angular momentum which can be either positive or negative (corresponding to the two possible directions of rotation) and is usually referred to as a spin of $\pm 1/2$. Spin within a nucleus is additive, so that nuclei with multiple particles can cancel to create no net angular momentum. Only nuclei with an odd number of protons or neutrons will possess a net angular momentum. While the quantization of angular momentum cannot be derived from Newtonian physics, most phenomena related to magnetic resonance can be explained using classical physics with the postulate of quantized angular momentum. This section presents a classical description of MR physics.




Figure 2.1: Polarization of spins in an external magnetic field. (a) Spins are normally randomly oriented, so that an ensemble of spins have no net magnetic moment. (b) In an external field, spins will align parallel or anti-parallel to the field, with a net magnetic moment parallel to the field.

2.1.1 Polarization

Visualizing a proton as a rotating unit of charge, any nucleus with net angular momentum will also have an associated magnetic moment, causing it to behave like a tiny magnet. NMR deals with the behavior of this magnetic moment in external magnetic fields. Normally, the magnetic moment of a nucleus has no preferred direction, so that an ensemble of spins will be randomly oriented relative to each other, producing no net magnetic moment (as in Fig. 2.1a). However, an external magnetic field (usually denoted B_0) will create a torque on spins and cause them to align parallel or anti-parallel to the field. Because the parallel alignment has lower energy, there is a slight preference for spins to align parallel to the field creating a net magnetic moment (as in Fig. 2.1b), which we will refer to as the *magnetization of the ensemble*, or simply the *magnetization*. At equilibrium, the magnetization is aligned parallel to the B_0 field, which is also referred to as the longitudinal axis. The plane perpendicular to the longitudinal axis is referred to as the transverse plane.




Figure 2.2: Precession of spins in an external magnetic field. Out of alignment with the main magnetic field, spins follow a helical path in returning to energetic equilibrium along the longitudinal (z) axis.

2.1.2 Precession

If the magnetization is out of alignment with the main magnetic field, torque forces will tend to push it back parallel with B_0 . However, because the magnetization possesses angular momentum, it does not directly rotate into alignment, but instead follows a helical pathway (see Fig. 2.2). This behavior, in which the magnetization is said to *precess* about the axis of B_0 , arises from the same type of forces that slow the fall of a spinning top (when torque is applied to a spinning object, the change in direction is about an axis that is orthogonal to both the axis of rotation and the axis of torque). The speed of precession depends on the strength of the torque force, which is proportional to the magnetic field strength:

$$\omega_0 = \gamma B_0$$

The proportionality constant γ is known as the gyromagnetic ratio of the given nucleus, and the frequency of precession is known as the Larmor frequency, or simply the resonance frequency.

This precessing magnetization is the source of the NMR signal, which is measured by placing a coil near the sample. The coil experiences an oscillating magnetic field due to the precessing magnetization (provided the coil does not lie in the transverse plane), and a current is induced. The current is proportional to the size of the magnetization and therefore provides a measure of the number of spins in the sample.

The prevalence of water in the body makes hydrogen the nucleus most commonly studied with MRI. For the magnetic field strengths typically used in NMR (0.1–10 T), the resonance frequency of most nuclei is in the radiofrequency (RF) range. For this reason, the detection coils used in NMR are typically referred to as RF coils.

2.1.3 Excitation

We now consider how spins may be *excited* out of alignment with the main field so that signal can be detected. Excitation is achieved with a second magnetic field, denoted B_1 , which oscillates in the transverse plane. Unlike the main field, this oscillating field can be turned on and off. The B_1 field adds to the main field B_0 , with the result that the net magnetic field now rotates about the longitudinal axis.

If the frequency of oscillation of this field is different from the Larmor frequency, it has little effect on the magnetization. However, if it is matched to the Larmor frequency, a resonance condition is established which results in excitation. As the net field rotates, the magnetization begins to precess into alignment with the new net field. On-resonance, the rate at which the magnetization moves to align to the net field is exactly matched by the rate that the field is changing, so that the magnetization tracks the net field as it changes. The result is that the magnetization slowly spirals out of alignment with B_0 , as shown in Fig. 2.3.

Since the resonance frequency of nuclei is in the radiofrequency range, the time-varying B_1 field is usually referred to as an RF pulse, and the effect of the RF pulse is typically described by the angle between the final magnetization and the longitudinal axis (α in Fig. 2.3). The B_1 field is created by a coil similar to that used for detection,

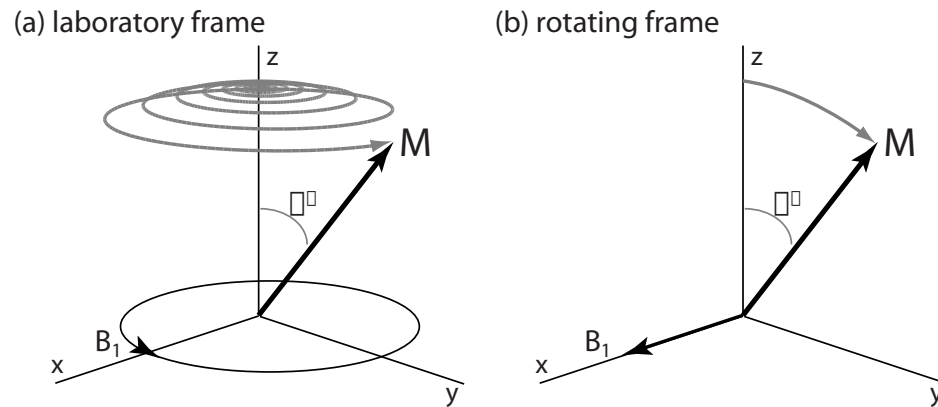


Figure 2.3: (a) Excitation of magnetization with an oscillating magnetic field. The magnetization precesses to track the RF field B_1 and gradually tips out of alignment with the main magnetic field. (b) In a frame of reference that rotates with the B_1 field, the magnetization simply appears to tip about the rotation axis defined by the RF field.

and in many experiments, a single coil is used for both excitation and detection.

2.1.4 Relaxation

The return of excited magnetization to energetic equilibrium is known as *relaxation* and is described by two separate time constants: T_1 is used to describe the rate at which the magnetization returns to alignment with the longitudinal axis, and T_2 describes the rate at which the magnetization is reduced in the transverse plane. Relaxation rates are related to molecular motion. As nuclei in the sample move due to thermal energy, a given nucleus experiences field fluctuations due to the magnetic moments of other nuclei. These fluctuations lead to relaxation. To understand the two different types of relaxation, we must consider the energy of field fluctuations at different frequencies, which is related to the energy of the motional spectrum at those frequencies.

In order to convert transverse magnetization into longitudinal magnetization, field

fluctuations must produce slight tips of the transverse magnetization (i.e., they must act as brief RF pulses). This is accomplished by fluctuations in the transverse plane at the Larmor frequency which slowly drive the magnetization back into alignment with the longitudinal axis. T_1 relaxation rates therefore depend on the energy of motion at the Larmor frequency. Since the Larmor frequency is proportional to B_0 and the motional spectrum is usually monotonically decreasing, T_1 will generally increase with field strength.

Fluctuations at the Larmor frequency will also reduce the transverse magnetization. However, this is not the only source of loss of magnetization in the transverse plane. Field offsets along the z direction create slight perturbations in the rate of precession of a nucleus, causing it to accrue phase relative to its buddies. Over time, these phase offsets cause spins in the transverse plane to lose phase coherence, and the net transverse magnetization is reduced. Since this dephasing is only significant for slow fluctuations, T_2 depends on the energy of motion at low frequencies. Since T_2 decay is caused by both low frequency motion and fluctuations at the Larmor frequency, T_2 is generally shorter than T_1 . Also, since T_2 tends to be dominated by low-frequency motion, it varies less with field strength than T_1 .

In addition to T_2 decay, there are sources of transverse magnetization loss which are reversible, generally referred as T_2^* effects. Stable offsets in the main magnetic field across the imaging sample cause different locations to slowly accrue phase relative to each other. Unlike T_2 decay, the resultant signal loss can be reversed by applying a 180° RF pulse, which inverts the phase of the magnetization. After this inversion pulse, spins will continue precessing until they are once again in phase. This phase inversion method is known as a “spin echo”, and its development by Hahn was a major advance in NMR [4].

Relaxation is an important aspect of the MR signal because it is highly tissue-dependent. For example, tissues containing large molecules or fixed structure will have little motional energy at the Larmor rate, but a concentration of energy at low frequencies. The result will be long T_1 and short T_2 (for example, in cartilage at

1.5 T, $T_1 \approx 800$ ms and $T_2 \approx 40$ ms). Tissues with small, unrestricted molecules will have more energy at the Larmor rate, and the T_1 and T_2 will be long (for example, in cerebrospinal fluid at 1.5 T, $T_1 \approx 1500$ ms and $T_2 \approx 250$ ms).

2.1.5 The Bloch Equation

The effects discussed above can be combined in a single differential equation that describes the behavior of magnetization in an NMR experiment. This formulation was first described by Felix Bloch in his seminal paper, and has become known as the Bloch equation [2]:

$$\frac{d\mathbf{M}}{dt} = (\mathbf{M} \times \gamma\mathbf{B}) - \frac{(M_x\hat{i} - M_y\hat{j})}{T_2} + \frac{(M_0 - M_z)\hat{k}}{T_1} \quad (2.1)$$

where $\mathbf{M} = [M_x \ M_y \ M_z]^T$ is the magnetization, \mathbf{B} is the total magnetic field, and T_1 and T_2 are the relaxation times.

2.2 Image Formation

The physics described above is a general description of magnetic resonance that applies to both NMR spectroscopy (a tool for chemical analysis) and MRI (a device for imaging objects). This section describes how MR physics can be employed to create images of the magnetization in an object.

2.2.1 Gradient Fields

The essential modification to NMR spectrometers that enabled the formation of images was the addition of yet another set of magnetic fields. In addition to B_0 and B_1 ,

MRI scanners incorporate a set of magnetic fields known as gradients that are parallel to B_0 and vary linearly along some spatial direction. Typically, MRI scanners have 3 separate gradient coils, denoted G_x , G_y and G_z , which create linear variation in the magnetic field along x , y and z , respectively. These fields are used to create spatial variation in the resonance frequency which enables the signal to be localized. By manipulating the amount of current running through the gradient coils, the strength of the gradient field (i.e., the slope of the field) is changed. The 3 separate gradient coils combine linearly to enable a gradient field along any given direction which adds to the B_0 field:

$$B(\mathbf{r}) = B_0 + \mathbf{G} \cdot \mathbf{r}$$

where \mathbf{r} is a (3-dimensional) spatial location and \mathbf{G} is a vector of gradient field strengths (e.g., in mT/m). The manner in which the subsequent spatial variation in resonance frequency ($\omega(\mathbf{r}) = \gamma B(\mathbf{r})$) can be used to create an image is described below. Gradient fields are typically generated by an electromagnetic coil and are small relative to the main field (for example, current technology will typically impose a 40 mT/m gradient field on top of a 1.5 T main field).

2.2.2 Spatial Encoding

To see how gradients can be used to create an image of the magnetization, it is helpful to express the MR signal in the presence of gradients. If we ignore the static component of the field (e.g., by demodulating the signal at the Larmor frequency, $\omega_0 = \gamma B_0$), the MR signal $s(t)$ during a gradient waveform $\mathbf{G}(t)$ is given by:

$$s(t) = \int \mathbf{M}(\mathbf{r}) e^{i\gamma \int_0^t \mathbf{G}(\tau) \cdot \mathbf{r} d\tau} d\mathbf{r}$$

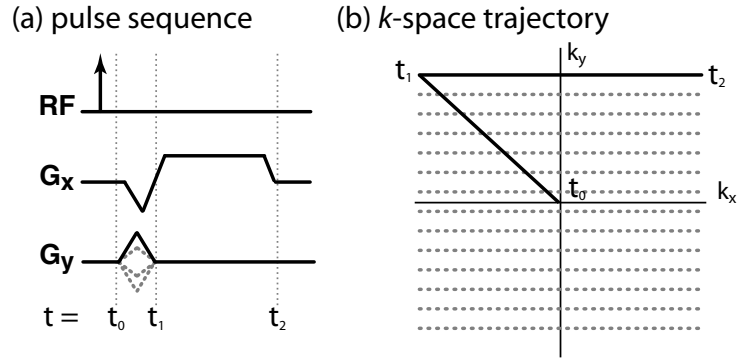


Figure 2.4: A basic MR imaging pulse sequence and the associated k -space trajectory. (a) The series of RF and gradient pulses for the sequence. These waveforms are repeated multiple times with the G_y waveform changing each repetition as indicated by the dashed lines. (b) The resulting k -space trajectory traces out a different line along k_x each repetition.

$$= \int \mathbf{M}(\mathbf{r}) e^{i2\pi \mathbf{k}(t) \cdot \mathbf{r}} d\mathbf{r} \quad (2.2)$$

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(\tau) d\tau \quad (2.3)$$

Equation 2.2 is suggestive of a Fourier transform, with the spatial frequency parameter $\mathbf{k}(t)$ given in Eq. 2.3. If $\mathbf{k}(t)$ covered the entire spatial frequency domain, an image of the magnetization $\mathbf{M}(\mathbf{r})$ could be recovered using an inverse Fourier transform.

Formation of an MR image consists of reading out the MR signal during a time-course of gradients for which Eq. 2.3 covers all of the spatial frequency domain, referred to as k -space. Note that the entirety of k -space does not have to be sampled following a single RF pulse, but can be acquired a piece at a time following multiple excitations. One possible method for sampling k -space is shown in Fig. 2.4, along with the pattern of gradients and RF pulses that would create this trajectory through k -space.




Figure 2.5: Three of the more common k -space trajectories include (a) 2DFT (also known as spinwarp), (b) echo-planar imaging (EPI) and (c) spiral. The portion of k -space collected during a single repetition period is shown by a black, solid line, and subsequent readouts are shown as gray, dashed lines.

2.2.3 Encoding Trajectories

The trajectory for covering k -space shown in Fig. 2.4 is only one possibility. Any traversal of k -space that covers the central region of k -space can be used to obtain an image provided that it satisfies Nyquist sampling requirements. Three trajectories with specific application to this work are shown in Fig. 2.5. The 2DFT trajectory (presented previously) is the simplest and the most common trajectory in clinical MRI. One of the benefits of this trajectory is the ease with which an image can be reconstructed: because data is collected along a Cartesian grid, a 2D fast Fourier transform (FFT) provides an excellent reconstruction of the data. The major disadvantage of 2DFT is the large number of excitations that are required to gather an entire image, which causes 2DFT scans to be fairly slow.

A number of alternative trajectories have been proposed which allow images to be gathered more quickly. The echo-planar (EPI) trajectory was developed early in the development of MRI by Mansfield [5] and has the advantage of collecting multiple lines in each readout. Initially, EPI images were reconstructed with a simple 2DFT, but this reconstruction tends to contain artifacts due to small errors in sequence timing and slight offsets in the magnetic field. More sophisticated methods have been

developed, and EPI reconstruction is still an area of active research. Another fast readout trajectory that was developed more recently is spiral imaging [6]. Reconstruction of spiral data requires resampling the data onto a Cartesian grid, typically referred to as *gridding*. Timing errors and field offsets tend to cause radial blurring in spiral images.

While all of the trajectories shown in Fig. 2.5 collect the image over multiple readouts (“multi-shot” imaging), EPI and spiral can potentially collect the entire image at once (“single-shot” imaging). Single-shot imaging is useful for imaging moving objects (such as the heart) or gathering time-series of images (typically in physiological imaging or functional studies).

2.3 Steady-state MRI

In this section we describe the physics underlying the steady-state pulse sequences developed in this thesis. These sequences have different signal dynamics that warrant further discussion.

2.3.1 Steady State Formation

An MR experiment consists of a series of RF pulses interspersed with data collection. In general, RF pulses are spaced regularly in time and repeat often enough that the magnetization will not fully recover from one RF pulse before the next occurs. As shown in Fig. 2.6, the result is that the magnetization evolves over several RF pulses to an equilibrium state (or *steady state*).

Most MRI pulse sequences use a large gradient, known as a *crusher* or *spoiler*, to dephase any remaining magnetization at the end of the T_R . This prevents the remaining magnetization from creating signal in subsequent repetition periods. Spoiling can lead to a significant increase in signal stability in short- T_R sequences, but also reduces

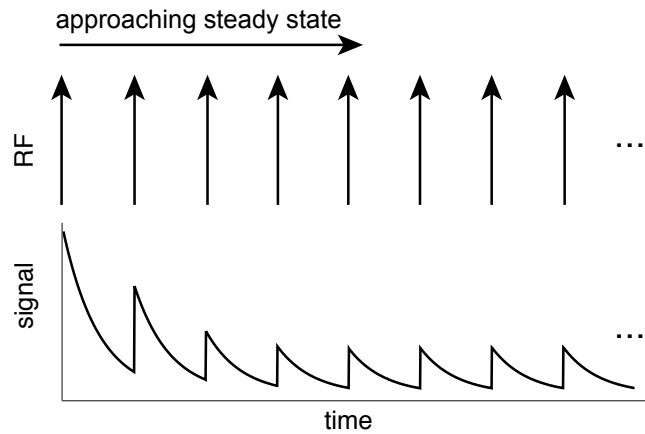


Figure 2.6: Approach to steady-state signal of a spoiled MR pulse sequence. Initially, the magnetization is aligned with the main field. Following the first RF pulse, the magnetization does not fully recover along the longitudinal axis before the next RF pulse, so that less signal is created. The signal created by the train of pulses decays exponentially to an asymptotic, or steady-state, value.

the signal level. In spoiled sequences, the evolution of the steady state is simply an asymptotic decay, as shown in Fig. 2.6.

2.3.2 Balanced SSFP

The key characteristic of the sequences developed in this work is that the magnetization is not spoiled at the end of the T_R . In fact, there is no net gradient along any axis between RF pulses, as in Fig. 2.7. In this type of sequence, the only source of magnetization phase is off-resonance precession during the repetition period. The resultant steady state has a strong dependence on off-resonance precession [7–10], which is referred to as a steady-state free precession (SSFP) condition [11–13]. Since sequences without perfectly rebalanced gradients can exhibit SSFP effects, we will refer to the sequence shown in Fig. 2.7 as a balanced-SSFP sequence.

One of the advantages of imaging under SSFP conditions is the ability to acquire




Figure 2.7: Balanced-SSFP imaging sequence, characterized by zero net gradient area along every gradient direction. Typically, these sequences use $T_R \ll T_2$.

images quickly without significant signal loss. Spoiled sequences discard the magnetization excited by a given RF pulse at the end of the repetition time, so that this type of sequence steadily loses signal as the repetition time is decreased. When the balanced-SSFP sequence shown in Fig. 2.7 is used with a short T_R (i.e., tens of milliseconds or less), the magnetization excited by a given RF pulse persists strongly over many repetition periods. This allows SSFP sequences to effectively re-use the magnetization excited in a given repetition period, resulting in a sequence that is able to efficiently acquire images with strong signal. However, the resultant images are sensitive to other effects such as off-resonance, as discussed below.

2.3.3 Matrix Representation of Balanced SSFP

While the differential Bloch equation given in Eq. 2.1 is the only expression that is able to fully explain the behavior of spins in a magnetic field, most periodic sequences can be more conveniently expressed using matrix formalism [14]. This section presents a method for calculating the steady-state magnetization of a balanced-SSFP pulse sequence. This representation has the advantage of providing a closed-form solution for the state of the magnetization at any point in the pulse sequence. In addition, the different forces affecting the magnetization (RF pulses, off-resonance and relaxation)




Figure 2.8: Pictorial representation of a balanced-SSFP pulse sequence used in deriving a matrix model for the steady-state magnetization. The pulse sequence consists of RF excitations (about the x axis with flip angle α) interspersed with periods of off-resonance precession (through an angle θ) and relaxation (according to the given T_1 and T_2).

are each represented by a separate matrix, resulting in an intuitive expression.

A simplified version of the balanced-SSFP pulse sequence is depicted in Fig. 2.8, in which RF pulses are interspersed with periods of off-resonance precession and relaxation. We represent the magnetization at any time t by a vector $\mathbf{M}(t) = [M_x(t) \ M_y(t) \ M_z(t)]^T$, which has a (normalized) equilibrium value $\mathbf{M}_0 = [0 \ 0 \ 1]^T$. Precession and excitation are simply rotations about z and x , respectively, and can be represented with rotation matrices $\mathbf{R}_x(\alpha)$ (excitation with flip angle α) and $\mathbf{R}_z(\theta)$ (precession through an angle θ over the course of the T_R). Signal decay due to relaxation takes the form of a diagonal matrix $\mathbf{E}(t) = \text{diag}(\exp(-t/T_2), \exp(-t/T_2), \exp(-t/T_1))$ and T_1 recovery toward equilibrium magnetization \mathbf{M}_0 is according to $1 - \exp(-T_R/T_1)$. We can then find a simple expression for the magnetization at the different time points shown in Fig. 2.8:

$$\begin{aligned}
 \mathbf{M}(t_2) &= \mathbf{R}_x(\alpha)\mathbf{M}(t_1) \\
 \mathbf{M}(t_3) &= \mathbf{E}(T_R)\mathbf{R}_z(\theta)\mathbf{M}(t_2) + (1 - e^{-T_R/T_1})\mathbf{M}_0 \\
 \mathbf{M}(t_4) &= \mathbf{R}_x(\alpha)\mathbf{M}(t_3)
 \end{aligned} \tag{2.4}$$

The steady-state magnetization at any time point can be found by noting that the magnetization will exactly repeat its motion over the T_R . We can find a closed-form

expression for $\mathbf{M}(t)$ from Eq. 2.4 by solving either for the magnetization immediately prior to the RF pulse $\mathbf{M}^- \triangleq \mathbf{M}(t_1) = \mathbf{M}(t_3)$, or immediately after $\mathbf{M}^+ \triangleq \mathbf{M}(t_2) = \mathbf{M}(t_4)$. The solution for \mathbf{M}^+ is:

$$\begin{aligned}\mathbf{M}^+ &= \mathbf{R}_x(\alpha)\mathbf{E}(T_R)\mathbf{R}_z(\theta)\mathbf{M}^+ + (1 - e^{-T_R/T_1})\mathbf{M}_0 \\ &= [\mathbf{I} - \mathbf{R}_x(\alpha)\mathbf{E}(T_R)\mathbf{R}_z(\theta)]^{-1}(1 - e^{-T_R/T_1})\mathbf{M}_0\end{aligned}\quad (2.5)$$

We can then express the signal at any time point t in the T_R as:

$$\mathbf{M}(t) = \mathbf{E}(t)\mathbf{R}_z\left(\frac{t}{T_R}\theta\right)\mathbf{M}^+ + (1 - e^{t/T_1})\mathbf{M}_0\quad (2.6)$$

Figure 2.9 shows the steady-state signal magnitude calculated from Eq. 2.6 for a range of off-resonance frequencies and flip angles. The signal magnitude is a strong function of off-resonance and flip angle, and the magnitude profile for a given flip angle repeats every T_R^{-1} Hz. The decreases in signal at some frequencies are typically referred to as nulls. These signal nulls cause bands of attenuation in images when there is frequency drift across the object, as can be seen in Fig. 2.10.

2.3.4 Coherence Pathway Analysis

While the matrix calculations presented above can be extended to other sequences, not all sequences are easily expressed using this formalism. In these cases, a qualitative understanding of steady-state signal dynamics can be achieved using coherence pathway analysis. The coherence pathway representation separates the magnetization into the longitudinal and transverse components that result from a given RF pulse. In other words, any RF pulse can be treated as a combination of a 90° pulse which excites a fraction of the magnetization, and a 0° pulse applied to the remaining


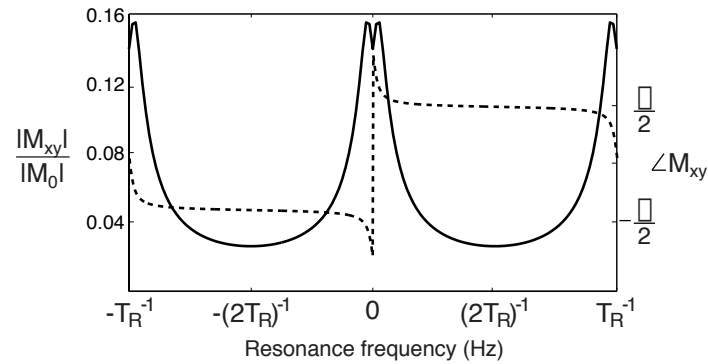


Figure 2.9: Balanced-SSFP signal profile, calculated from Eq. 2.6 for a range of flip angles and off-resonance frequencies. Other parameters included $T_1 = 1000$ ms, $T_2 = 100$ ms and $T_R = 10$ ms. Signal magnitude is a strong function of off-resonance and repeats every T_R^{-1} Hz.

fraction of the magnetization. (In actuality, there is also a 180° component of the RF pulse, but for the current discussion this can be neglected). The linearity of the Bloch equations means that these transverse and longitudinal components can be analyzed independently.

As shown in Fig. 2.11, application of an RF pulse splits the existing transverse and longitudinal components, with each pulse doubling the number of independent magnetization components. Hence, a train of RF pulses creates a plurality of concurrent “coherence pathways”, each of which is completely described by the state of the magnetization, transverse or longitudinal, during each T_R . The fraction of the magnetization in a given pathway depends on the flip angle α and relaxation times T_1 and T_2 . In every T_R , $\cos(\alpha)$ of the magnetization experiences a flip angle of 0° and the remaining $\sin(\alpha)$ of the magnetization experiences a flip angle of 90° . In addition, the magnetization in a given pathway experiences an attenuation $E_2 = \exp(-T_R/T_2)$ during each transverse period and $E_1 = \exp(-T_R/T_1)$ during each longitudinal period. These terms can be cascaded to calculate the fraction of the magnetization in

(a) balanced-SSFP signal profile



(b) magnitude image

(c) phase image




Figure 2.10: Balanced-SSFP images of the brain, exhibiting banding artifacts due to drift in the resonance frequency across the head. (a) Calculated signal profiles for the imaging parameters used in the scan ($\alpha = 5^\circ$ and $T_R = 40$ ms), $T_1 \approx 1000$ ms and $T_2 \approx 100$ ms. In addition to the magnitude effects shown in Fig. 2.8 (solid lines), the signal phase also exhibits strong sensitivity to off-resonance (dashed lines). These calculations are manifested in the images as (b) magnitude banding and (c) phase discontinuities.

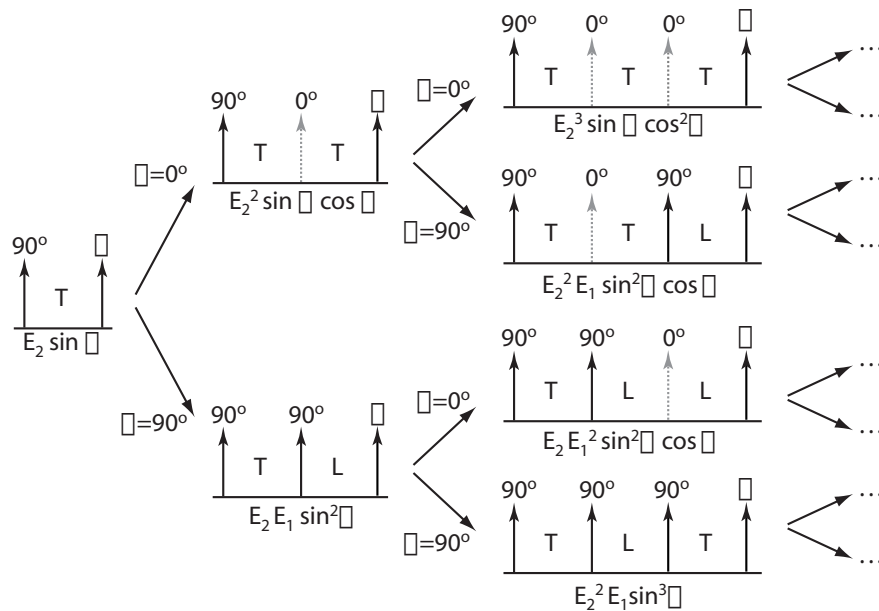


Figure 2.11: Coherence pathway analysis describing the magnetization resulting from a series of RF pulses. Each RF pulse splits the magnetization into transverse and longitudinal components, creating a spreading tree of state pathways experienced by different components of the magnetization. Starting with a transverse component created by one RF pulse, this figure depicts the 4 pathways created by 2 subsequent RF pulses. The effective RF pulse experienced by a given pathway in each T_R is shown above the RF pulse, and the orientation of the magnetization (longitudinal or transverse) during each of the 3 T_R s is denoted with either “L” or “T”. The fraction of the magnetization in each pathway is indicated under the pathway diagram.

a given pathway, as suggested in Fig. 2.11.

The steady-state is reached when the magnetization created by an RF pulse is exactly cancelled by the decay of existing pathways. The total signal in a given repetition period is the summation of the signal contribution from all coherence pathways. In balanced SSFP, the magnetization in a given pathway will create signal (form an echo) during every transverse period. Other steady states (such as SSFP-DWI, discussed in Chapter 5) may form slightly different pathways or only form echoes in a subset of transverse pathways.

Chapter 3

Neuroimaging With MRI

3.1 Neuroscience Applications of MRI

While MRI has found application throughout the body, nowhere has this imaging modality had greater impact than in the study of the brain. In clinical neurology, MRI has proven to be a powerful diagnostic tool for studying a wide range of neuropathologies including acute stroke [15] and cerebral tumors [16]. In addition, MRI has shown potential as a clinical marker for demyelinating disease [17] and many forms of dementia [18–20]. It has also has generated considerable interest as a method for assessing new neuropharmacological agents and other therapies [21].

While MRI technical development has in general been driven by clinical methods, its application to pure cognitive science has rapidly expanded since brain activity was first detected with functional MRI (fMRI) [22, 23]. Recently, the ability to infer directionality of white matter tracts from diffusion-weighted imaging (DWI) has created the potential for direct study of brain connectivity [24]. In turn, many of these methods are beginning to be evaluated for potential clinical applications [25].

fMRI and DWI provide complementary information about the brain: the functional activity of gray matter and the structural nature of white matter, respectively.

However, the state-of-the art in both methods suffers from low signal-to-noise ratio (SNR), poor resolution and severe image artifacts. These limitations produce image quality that is far worse than that attainable with other MRI methods. This thesis presents new methods for fMRI and DWI based on steady-state imaging which have the potential to overcome these limitations.

3.2 Functional MRI

In the 10 years since its inception, fMRI has revolutionized neuroscience by mapping activity throughout the brain without the use of radioactive tracers, electrical probes or other invasive procedures. The non-invasive nature of fMRI, along with its ability to visualize deep structures, makes for a powerful tool with which neuroscientists can study human brain function *in vivo*. In particular, Blood Oxygenation Level Dependent (BOLD) fMRI [22, 23, 26] has become a dominant tool for functional neuroimaging.

An overview of functional MRI and related neurophysiology can be found in Buxton's text [27]. More detailed descriptions of various technical issues and specific applications of fMRI can be found in the anthologies [28] and [29].

3.2.1 Functional Neuroimaging Methods

While a number of other mapping modalities exist, fMRI has the unique ability to map activity at its source without use of invasive procedures or tracers. PET is able to map a number of relevant physiological parameters [30, 31], but dosage concerns severely limit its use in human subjects. EEG [32, 33], MEG [34] and ERP [35, 36] are non-invasive methods with high temporal resolution, but have severely limited spatial resolution. Direct electrical recordings and optical flowmetry [37–39] are intractable for large-scale mappings and generally impractical for human studies.

In addition to being non-invasive, fMRI is capable of higher spatial resolution than other methods for global mapping such as PET. While BOLD is the most popular method for fMRI, several other methods of using MRI for functional imaging have been explored. The most prominent of these is Arterial Spin Labeling (ASL), which measures cerebral blood flow at the level of capillary perfusion [40–42]. While ASL may have some advantages over BOLD, it has lower SNR than BOLD, which is itself a low-SNR method. Measurements of cerebral blood volume have also been considered [43–45] but these measurements have proven difficult without the use of contrast agents.

3.2.2 BOLD Contrast

A number of localized physiological changes accompany neural activity, including increases in blood flow, blood volume and oxygen metabolism. BOLD imaging senses changes in blood oxygenation by exploiting the paramagnetism of deoxyhemoglobin. During activation the cerebral blood flow increases more than oxygen consumption [46], resulting in a dilution of the deoxyhemoglobin concentration. Because deoxyhemoglobin experiences a resonance frequency shift relative to water [47], the presence of deoxyhemoglobin in a blood vessel introduces a spread in resonance frequency around the vessel that causes the signal to dephase more quickly [48] (see Fig. 3.1). In BOLD fMRI, data acquisition is delayed following RF excitation (typically by 30-60 ms) to allow signal dephasing to occur, resulting in signal levels that depend on the local concentration of deoxyhemoglobin [22]. Since blood has a lower concentration of deoxyhemoglobin during activation, the signal during activation is larger than the signal at rest [49]. This signal change is the source of BOLD contrast.

3.2.3 fMRI Experiments

In a typical fMRI experiment, a time-series of images is gathered while the subject performs a task or experiences a stimulus. Brain activity is typically detected using




Figure 3.1: Illustration of the contrast mechanism in BOLD fMRI. The presence of deoxyhemoglobin within a blood vessel causes a chemical shift in the tissue surrounding the vessel that exhibits the characteristic dipole pattern depicted on the left. Within a voxel, this pattern represents a spread in resonance frequency (i.e., a broadening of the voxel spectrum). Over time, magnetization from different spatial locations lose phase coherence, resulting in phase dispersion and therefore signal attenuation that increases over time. The result is signal dependence on local blood oxygenation.

a measure of the difference in signal between control/rest and task/activation conditions. One simple method is to correlate the signal time-course for each voxel with a simple model for the expected signal response in regions of activation. Activated voxels will have a large correlation coefficient.

FMRI has been used to study a wide range of cortical and subcortical systems and mental processes. One of the most extensively studied parts of the brain is the visual system, which makes it a well-characterized testbed for new methods. The FMRI experiments described in this work are limited to the visual cortex, but the methods described are expected to be easily extensible to the same breadth of cortical systems as BOLD FMRI.

3.2.4 Limitations of BOLD FMRI

BOLD FMRI has a number of important limitations that result from the coupling of the BOLD contrast mechanism to image artifacts and signal loss. A key problem arises from other sources of off-resonance in the image which also cause signal dephasing. Susceptibility boundaries (such as those found near the sinus cavities) can cause severe image distortion [50,51] and signal dropout [52], as shown in Fig. 3.2. The severity of such artifacts is highly problematic in areas adjacent to the sinuses. A second issue with BOLD FMRI is the need for a long echo time. This causes significant signal loss due to transverse relaxation, decreasing the SNR relative to short-TE imaging [27]. In addition, in order to acquire images with reasonable temporal resolution (typically 2–3 seconds), long single-shot readouts are typically used. Off-resonance and T_2^* decay over the course of these long readouts causes warping and loss of resolution (blurring) [53], which is particularly pronounced at long echo times. This coupling of BOLD contrast to signal dropout, low SNR and image artifacts is a fundamental limitation of the BOLD method. Chapter 4 will present an alternative to BOLD FMRI based on SSFP that decouples functional contrast from the mechanism of image formation, creating the potential for independent optimization of image quality.




Figure 3.2: The presence of a susceptibility boundary (such as the sinus cavities) creates a sharp dropoff in the local magnetic field. Voxels near these boundaries contain a broad range of frequencies, and therefore tend to experience phase dispersion over time. Noting the similarity with Fig. 3.1, it can be seen that if a long echo time is used to enhance BOLD contrast, susceptibility-based image artifacts will subsequently increase. BOLD images typically experience intense signal loss and warping in tissue immediately above the sinuses, severely compromising the ability to detect signal in these regions.

3.3 Diffusion-Weighted Imaging

Diffusion-weighted imaging (DWI) achieves contrast based on the self-diffusion of water molecules. Water diffusion changes in the presence of tissue and reflects various aspects of tissue structure. In general, the presence of tissue structures reduces the diffusion coefficient (D) by introducing boundaries that restrict the distance that molecules tend to move in a given time. In DWI, the diffusion coefficient in a voxel can be calculated from the signal level. The primary application of DWI has been as a diagnostic of acute stroke, where changes in the diffusion coefficient can be used to differentiate acute infarcts from chronic lesions and healthy tissue [15, 54].

A basic but thorough introduction to DWI, with specific focus on neurological applications, can be found in Buxton's text [27]. A detailed discussion of the physics of diffusion and its effect on the NMR signal is given by Callaghan [55].

3.3.1 Diffusion Weighting

Diffusion-weighted contrast can be added to a pulse sequence with the inclusion of a bipolar set of gradients, i.e., two sequential gradients with the same area but opposite polarity (see Fig. 3.3). The first gradient finely encodes position in the magnetization phase and the second gradient applies the negative of this phase. If spins do not move between the encoding and decoding gradients, the magnetization experiences no net phase change. However, any displacement of a given spin between the two gradients results in a phase offset. Stochastic motion due to diffusion causes spins to accrue random phase, which results in phase dispersion and therefore signal attenuation within a voxel. Signal attenuation increases with greater diffusion, creating diffusion-weighted contrast. This contrast mechanism is reminiscent of the BOLD contrast mechanism discussed above in that it is based on phase dispersion within a voxel. However, the specific mechanisms that cause phase dispersion are quite different, giving rise to distinct issues and artifacts.

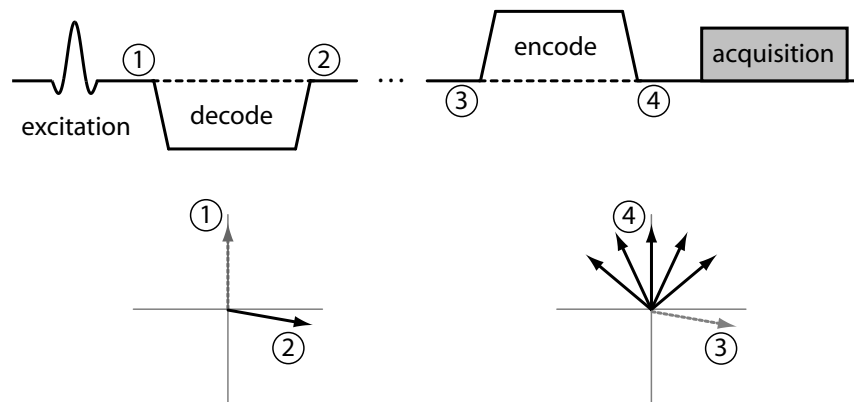


Figure 3.3: Diffusion contrast based on a bipolar set of gradients. The simplified pulse sequence shown above consists of RF excitation, bipolar diffusion gradients (denoted “encode” and “decode”) and acquisition. The ellipsis represents a temporal delay between the two diffusion gradient pulses, which is often long to allow spins to diffuse. The location of a representative subset of diffusing spins in the transverse plane is shown below at 4 different time points in the sequence, before and after each of the diffusion gradients. All of the spins within a voxel are coherent until the second gradient decodes spatial position, at which point any change in location between the diffusion gradients results in a phase offset. Diffusive motion causes dephasing within a voxel, and therefore signal attenuation. Note that this treatment assumes that no diffusion occurs during the diffusion gradient, a simplification which is commonly made but not necessarily correct.

The diffusion gradient can be applied using any combination of the 3 gradient axes (although this is not shown explicitly in Fig. 3.3). This means that the bipolar set sensitizes the signal to diffusion along a particular direction, given by the vector sum of the diffusion gradients. For example, if diffusion gradients were played simultaneously on the x and y axes with the amplitude of the gradient on the x axis twice as large as the gradient on the y axis, the sequence would be sensitive to diffusive motion along the direction $[2\ 1\ 0]^T$. The ability to measure the component of water diffusion along different directions is an important property of DWI. In the presence of isotropic diffusion, making measurements along multiple dimensions can reduce bias in the measured diffusion coefficient caused by imperfections in the gradient hardware. More importantly, measurements along multiple dimensions can be used to detect diffusion anisotropy, which sometimes reflects tissue structure, as discussed below.

3.3.2 Pulse Sequences

A number of different sequences have been proposed for diffusion imaging. The most prevalent method for DWI uses a simple spin echo (SE) sequence with identical diffusion gradients on opposite sides of the 180° refocusing pulse [56] (see Fig. 3.4a). A similar method of diffusion encoding can be used in conjunction with a fast spin echo (FSE or CPMG) readout, with accompanying sensitivity to CPMG effects [57].

Stimulated echoes (STE) have also been proposed for DWI, typically with identical diffusion gradients inserted after the first and third 90° pulses [58] (see Fig. 3.4b). Although the signal in STE is lower than SE, STE-DWI has the potential advantage over SE-DWI that the magnetization relaxes according to T_1 during the “mixing” time between diffusion gradients. Since the amount of diffusion weighting depends both on the area under the diffusion gradients and the mixing time, the use of STE-DWI allows an increase in diffusion weighting without incurring significant signal loss due to T_2 decay, which is beneficial for short T_2 species.

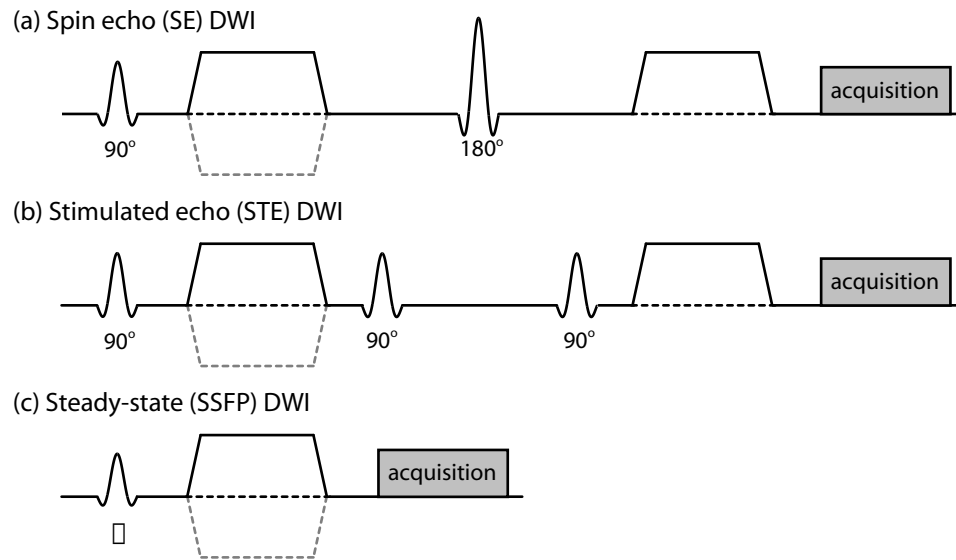


Figure 3.4: Three common DWI pulse sequences (a single T_R is depicted for each). Each sequence achieves diffusion weighting by effectively creating a pair of gradients functionally equivalent to the bipolar pair shown in Fig. 3.3 (where the effective negative gradient is indicated by the gray dashed lines). (a) Spin echo (SE) DWI places the two diffusion gradients on opposite sides of the 180° refocusing pulse. An effective bipolar set is formed because the refocusing pulse implicitly negates the first diffusion gradient. (b) Stimulated echo (STE) DWI has a similar mechanism to SE-DWI, where the second and third 90° pulses refocus the echo in 2 steps, negating the first diffusion gradient. (c) Steady-state (SSFP) DWI uses a single, unbalanced diffusion gradient in a given TR. As described in Chapter 5, this diffusion gradient forms an implicit bipolar pair with itself over multiple repetition periods.

The third major method for diffusion imaging uses SSFP to create a diffusion-weighted steady state [59, 60] (see Fig. 3.4c). Most commonly, a single diffusion gradient is inserted between excitation and readout. As described in Chapter 5, the signal in this sequence is the summation of a plurality of diffusion-weighted echoes. For a given size of diffusion gradient, SSFP-DWI can achieve greater diffusion weighting than SE-DWI and is considerably more efficient than STE-DWI. However, the strong sensitivity of this method to motion has largely precluded its use. SSFP-DWI and the removal of motion artifacts is the topic of the Chapter 5.

3.3.3 Diffusion Tensor Imaging

Restricted diffusion due to the presence of tissue boundaries usually does not exhibit any directional dependence, so that diffusion in most tissues is isotropic [61]. However, some tissues contain higher-level structure that restricts diffusion more along certain directions, causing anisotropic water diffusion. This effect has been observed in white matter, where the presence of myelin restricts diffusion across, but not along, white matter fibers [62, 63]. By measuring the diffusion coefficient along a number of different directions, the directionality of white matter fibers can be inferred [24]. This type of imaging is commonly termed diffusion tensor imaging (DTI), from early models in which the directional dependence of the diffusion coefficient was modeled as a tensor. DTI has garnered considerable attention in neuroscience as a potential method for tracing white matter tracts running from one region of the brain to another, which cannot be accomplished in living humans with any other method.

3.3.4 Limitations of DWI

The single greatest limitation in state-of-the-art DWI is the sensitivity to object motion, including both subject movement (rigid motion) and tissue deformation (non-rigid motion). Sensitivity to small diffusive motion is necessarily accompanied by sensitivity to other sources of motion. Bulk motion is often considerably larger than diffusive motion: for example, in brain tissue diffusive motion during a 50 ms window will on average displace water molecules by about $10 \mu m$, whereas the brainstem will displace by as much as 75–100 μm during the same period in peak systole [64, 65].

Bulk motion introduces phase shifts to the magnetization which can interfere to cause ghosting and signal loss in multi-shot acquisitions [66, 67]. Most DWI sequences use a single-shot acquisition (typically EPI) to reduce sensitivity to subject motion. However, single-shot imaging has severely limited resolution and image artifacts due to eddy currents and off-resonance. Reduction and/or removal of motion artifacts from DWI data is currently an active area of research. Proposed solutions include

trajectories with reduced motion sensitivity [68], prospective rephasing of the magnetization [69] and retrospective correction of phase artifacts using navigation [66,70–73] or multi-shot, full-FOV acquisitions [57]. Retrospective correction provides the most flexibility and will be discussed in Chapter 6.

Chapter 4

Balanced SSFP for Functional MRI

4.1 Overcoming the Limitations of BOLD

As discussed in Chapter 3, BOLD fMRI suffers from poor image quality due to the coupling of BOLD contrast to sources of image artifact and signal loss. This coupling causes signal dropout near susceptibility boundaries and image distortion in the presence of off-resonance. In addition, single-shot readouts are typically used to allow data to be gathered at a reasonable temporal rate. This forces a tradeoff between BOLD functional contrast (which increases with longer T_E) and image quality (which rapidly degrades as the echo time increases). The resultant images have very poor quality relative to that which is achievable with other MRI techniques.

This chapter introduces a new method for functional MRI using balanced-SSFP imaging that may overcome these issues. The sensitivity of the balanced-SSFP signal to resonance frequency can be used to directly detect the deoxyhemoglobin frequency shift associated with brain activity [?]. The sequence presented here uses the phase profile of SSFP [74] to invert signal from the deoxyhemoglobin frequency shift relative to water, establishing a Blood Oxygenation Sensitive Steady-state (BOSS) signal.




Figure 4.1: Balanced-SSFP pulse sequence. All gradients axes have zero total area over a repetition period, which causes a steady-state to form that has a strong dependence on resonance frequency.

This method is intrinsically sensitive to the deoxyhemoglobin frequency shift, decoupling the functional contrast mechanism from sources of artifact. The major issue with the BOSS approach is that the resonance frequency of oxygenated blood must be accurately aligned with the balanced-SSFP frequency response. To cover large regions of interest, such as the visual cortex, we can combine careful linear shimming with acquisition at a small number of frequencies. In addition to reduction of image artifacts and signal dropout, this method offers higher SNR and greater functional contrast than traditional BOLD FMRI.

4.2 BOSS FMRI

As discussed in Chapter 2, balanced SSFP consists of a rapidly-repeating series of RF pulses and fully-refocused imaging gradients (see Fig. 4.1). For short repetition times, the magnetization persists from one RF pulse to the next, creating a steady state that is strongly dependent on the amount of phase the magnetization accrues during the T_R . Since all gradients axes have zero total area, the only source of phase accrual between RF pulses is off-resonance precession, creating signal that is strongly dependent on local resonance frequency.

The frequency response of the balanced-SSFP signal has a unique and fairly complicated dependence on tissue and sequence parameters. As discussed in Chapter 2, certain frequencies experience drastically reduced signal, creating frequency bands of signal attenuation commonly referred to as “nulls”. For example, balanced-SSFP imaging usually uses a large flip angle ($\alpha \geq 30^\circ$) with alternating sign ($\pm \alpha$), resulting in a sharp null for species that precess 180° over each T_R . If, instead, the sign of the RF pulse does not alternate, this null occurs on-resonance (solid line in Fig. 4.2a).

4.2.1 FMRI with Frequency-sensitive Nulling

Scheffler *et al.* previously proposed using the magnitude null at large flip angles to detect the dynamic frequency shifts associated with brain activation [?]. This method carefully places the center frequency relative to the unshifted oxyhemoglobin frequency and the shifted deoxyhemoglobin frequency. For example, the signal at the deoxyhemoglobin frequency (which we will refer to as “deoxygenated signal”) is placed in or near the null, while the signal at the oxyhemoglobin frequency (or, “oxygenated signal”) is shifted slightly out of the null, as shown in Fig. 4.2a. As blood oxygenation levels rise during activation, the fraction of spins dwelling in the null decreases and the signal increases. Alternatively, the oxygenated signal could be placed in or near the null, causing a signal decrease during activation. This contrast mechanism has the advantage of directly detecting the activation-induced frequency shift, whereas BOLD detects the shift indirectly as signal dephasing. We will refer to this functional contrast mechanism as “frequency-sensitive nulling”, or simply “nulling”.

The functional contrast of this sequence can be thought of as the vector difference between the oxygenated and deoxygenated signal, which gives the signal change ΔM shown in Fig. 4.2c. This technique favors short, multi-shot readouts to maintain the steady-state, which results in improved image quality over BOLD imaging. This contrast mechanism has three main drawbacks. First, signals are deliberately acquired in the low-signal part of the SSFP signal profile. Second, it requires a very homogeneous magnetic field. Finally, because the slope of the null is inversely proportional to the

T_R , functional contrast is coupled to the T_R (favoring $T_R \approx 20-40$ ms) [?].

4.2.2 FMRI with Frequency-sensitive Inversion

We propose a different but related approach that creates functional contrast based on the phase of the balanced-SSFP magnetization profile. In addition to the magnitude effects described above, the magnetization phase undergoes an abrupt change of 180° over a narrow band of frequencies near resonance. At $T_E = T_R/2$, the phase profile outside this transition band is flat with positive and negative frequency offsets separated by 180° (dashed lines in Fig. 4.2). These flat regions of the phase profile reflect spin echoes on the positive and negative real axes [?]. Consequently, the signals for positive and negative frequency offsets have opposing signs. In the region of the transition, the signal phase is extremely sensitive to resonance frequency, while outside the transition the phase is largely independent of off-resonance. At small flip angles, the signal magnitude peaks in the phase transition [13] (solid line in Fig. 4.2b).

We can use the signal profile in Fig. 4.2b for functional contrast by placing the oxygenated and deoxygenated signals on opposite sides of the phase change (see Fig. 4.2b). In this arrangement, oxygenated signal has the opposite sign of deoxygenated signal (see Fig. 4.2d). Within a voxel, deoxygenated signal will subtract from the larger oxygenated signal. Activation causes a drop in the fraction of the signal that is deoxygenated (negative) and a rise in the fraction that is oxygenated (positive). Assuming that the majority of the magnetization is oxygenated, the signal increases during activation. If the majority of spins were frequency-shifted to the deoxyhemoglobin frequency, reducing the concentration of deoxyhemoglobin would cause signal loss. We will refer to this functional contrast mechanism as Blood Oxygenation Sensitive Steady-state (BOSS) FMRI. (Note that, although the method proposed by Scheffler [?] could also be considered to be a BOSS method, in this work the term BOSS will refer specifically to the phase inversion method).

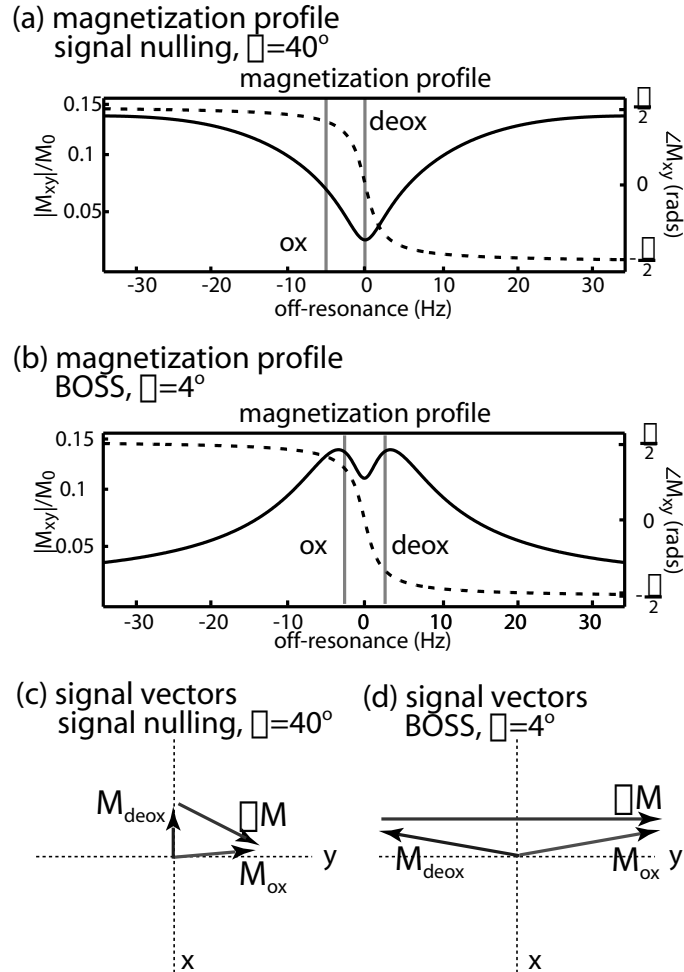


Figure 4.2: Balanced-SSFP signal magnitude (solid) and phase (dashed) as a function of resonance frequency. (a) At high flip angles, the signal magnitude has a null on-resonance. (b) At low flip angles the signal magnitude peaks on-resonance. A large phase transition (or sign change) occurs near resonance (0 Hz) for all flip angles. Placement of the center frequency relative to oxyhemoglobin (ox) and deoxyhemoglobin (deox) is shown for (a) the signal nulling technique, and (b) the BOSS technique introduced here. The functional contrast ΔM is shown in the vector diagrams for the two methods (c and d). BOSS has considerably higher functional contrast than the nulling method.

4.2.3 Inversion vs. Nulling

Like the nulling method, BOSS FMRI is intrinsically a multi-shot acquisition with short readouts. The use of multi-shot readouts drastically reduces the image artifacts and signal dropout compared to BOLD FMRI. In addition, both balanced-SSFP methods have the high SNR efficiency (i.e., SNR per unit of imaging time) characteristic of balanced SSFP. Both methods also require careful shimming and placement of the center frequency.

There are two key differences between these methods. First, the BOSS inversion exhibits a larger signal change than the nulling technique. This can be seen by comparing the vector diagrams in Fig. 4.2c and d, which show deoxygenated signal (M_{deoxy}), the oxygenated signal (M_{oxy}) and the resulting difference (ΔM) for the two methods. Second, whereas the nulling method couples functional contrast to lower signal levels, the BOSS method allows functional signal to be acquired in the highest-signal portion of the off-resonance profile, resulting in higher SNR. If the physiological noise is the same for the two methods (which may not be true), these factors suggest that BOSS should have greater functional contrast-to-noise (CNR) and therefore provide a more robust functional signal.

4.3 Experimental Validation

4.3.1 Experimental Methods and Data Analysis

The BOSS FMRI method was implemented on a 1.5 T GE Signa scanner with gradients capable of 40 mT/m and 150 mT/m/ms maximum slew rate. A quadrature T/R birdcage head coil was used in all experiments. The subject was presented with a visual stimulus consisting of a 10-Hz contrast-reversing annulus grating. The stimulus was presented in eight 30-s blocks consisting of 15 s of stimulus followed by 15 s of rest. The subject was instructed to fixate at all times on cross-hairs located in the center

of the visual FOV. A sagittal slice through the occipital pole (24 cm FOV, 128×64 , $\alpha=5^\circ$) was gathered during stimulation. A 2DFT trajectory was used with the center of k -space acquired halfway through the T_R (which varied across the experiments, and was generally chosen to yield a given spatial resolution at an acceptable frame rate). Linear shimming was targeted to the caudal occipital pole. Experiments were usually repeated at a range of center frequencies, and often were repeated multiple times at each of these frequencies.

Analysis was performed using either custom MATLAB software or the FSL software package (<http://www.fmrib.ox.ac.uk/fsl>). Analysis began with motion correction, high-pass filtering and averaging of the repetitions at a given frequency. Individual frequency measurements were separately processed using standard fMRI analysis methods based on a general linear model (GLM) for activation patterns [29]. In this analysis, the time-course for each voxel is fit to a simple model for the hemodynamic response to the stimulus. Activated pixels are identified by thresholding the correlation coefficient of the voxel time-course to the response model (or, equivalently, by thresholding the p - or Z -statistic calculated within the GLM). Results from fMRI experiments are presented using an “activation map”, where activated voxels are displayed as a color overlay on top of an anatomical reference. The reference image is usually obtained with a different pulse sequence, but in this work is typically the average image from the functional experiment to explicitly demonstrate the image quality of the underlying data.

In several of the experiments, the data at multiple frequencies was also processed as a group using a custom analysis. For each voxel, the time-course from the frequency with maximum stimulus correlation (or Z -statistic) was retained (i.e., a maximum contrast projection (MCP) of the multi-frequency measurements). This analysis results in a single data set that is pieced together from the data sets at multiple frequencies. This data set was thresholded and displayed as described above.

4.3.2 Single- and Multi-frequency Measurements

A pilot BOSS study was run in which a single sagittal slice (24 cm FOV, 128×64 , $\alpha=5^\circ$) was gathered every 0.5 seconds during stimulation ($T_R/T_E=7.8/3.9$ ms). The experiment was repeated 3 times at each of 4 frequency offsets spaced 4 Hz apart (where 4 Hz roughly corresponds to the width of the phase transition band for gray matter as shown in Figure 4.2). These data sets were analyzed separately and as a group using the analysis methods described above.

The results for the single-frequency experiment exhibiting the greatest area of activation are shown in Fig. 4.3a and c. Stimulus-correlated signal changes of 4–5% were found exclusively in the occipital lobe. These changes are 2–3 times larger than typical BOLD signals (1–2% at 1.5 T) for similar full-contrast visual stimuli [27]. The multi-frequency analysis identified a similar ROI (Fig. 4.3b) with even larger signal changes of 7–8% (Fig. 4.3d). Given that the activation ROIs are largely overlapping, the larger signal changes found by the multi-frequency analysis indicate that some voxels were more strongly activated at a frequency other than that shown in Fig. 4.3a and c. Because no temporal or spatial filtering was performed on the data, physiological fluctuations are more prevalent in the multi-frequency time-course.

The activation masks in Fig. 4.3 are overlaid on the balanced-SSFP images gathered during the experiment. These images contain no warping or blurring. The signal variation across these images is due to drift of the resonance frequency across the brain. The center frequency was matched to the resonance frequency of the occipital pole so that the visual cortex was near the phase transition. Consequently, anterior regions of the brain lay in the low-signal portion of the off-resonance profile (see Fig. 4.2b). This signal variation is not the same signal dropout as is found in BOLD images, which is unrecoverable. In BOSS, high signal can be obtained in any given region by matching the center frequency to the resonance frequency in that region.

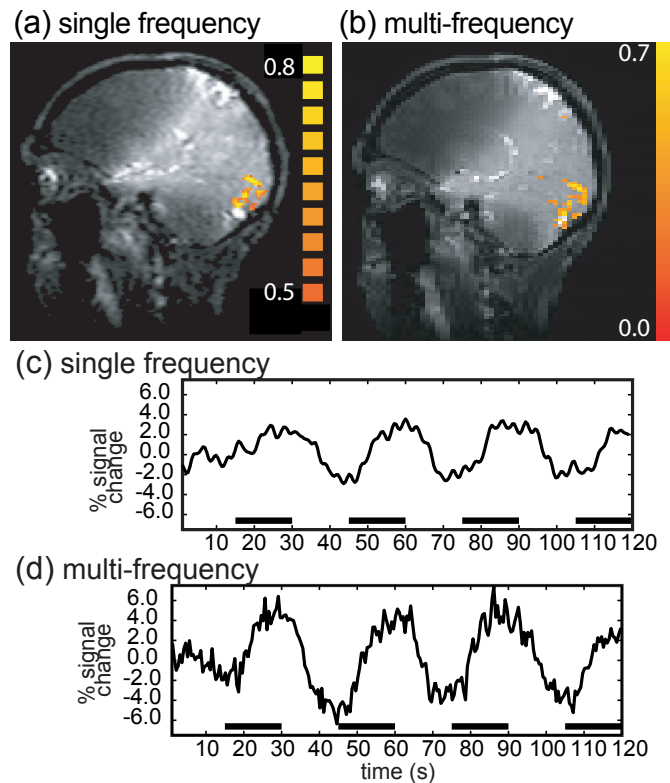


Figure 4.3: Results from the single- and multi-frequency analyses of the BOSS activation experiments. The activation masks for the (a) single- and (b) multi-frequency analyses are strikingly similar. These ROIs have stimulus-correlated signal changes of (c) 4–5% at a single-frequency and (d) 7–8% for the multi-frequency experiment. These changes are significantly larger than the 1–2% BOLD signal changes at 1.5 T. Temporal filtering was performed in (c) but not in (d), which exhibits larger physiological noise fluctuations. Masks are shown on the balanced-SSFP images from the experiment (the signal variation is due to the magnitude profile in Fig. 4.2).

4.3.3 Negatively-correlated Signal Changes

In addition to the positively-correlated voxels discussed above, we also observed negatively-correlated voxels, i.e., voxels that lost signal during stimulus periods. Like the positively correlated activations, these voxels had a strong, stimulus-correlated signal change of 6–7% in the multi-frequency MCP analysis. These activations were mostly remote from the visual cortex (see Fig. 4.4). Many voxels, both positively and negatively correlated, had a strong stimulus correlation at multiple frequencies. Most of these voxels (46 of 48 activated voxels) tended to exhibit a characteristic pattern of strictly negative or strictly positive correlations.

These negatively-correlated voxels may represent activation-induced frequency shifts outside of the phase transition (i.e., off resonance) due to the magnitude profile (akin to the functional contrast mechanism employed in [?]). Although it is possible to detect frequency shifts off-resonance due to the magnitude profile, both the SNR and functional contrast off-resonance are considerably worse than on-resonance. The large signal changes in these negatively-correlated voxels is therefore surprising, and suggests that this may not be simply a magnitude effect. A second possible cause of these negative correlations is an unusually long delay between stimulus and hemodynamic response, causing apparent negative correlations. A third possibility is that these voxels may represent large veins that exhibit inverted BOSS dynamics due to a larger fraction of frequency-shifted spins. Further experimentation is necessary to interpret these negative correlations.

4.3.4 High-resolution BOSS

One potential application of BOSS FMRI is in the acquisition of high resolution FMRI, an area of current research [75–79]. BOSS FMRI experiments were performed at 2 high-resolution voxel sizes, $1.25 \times 1.25 \times 3 \text{ mm}^3$ and $1 \times 1 \times 2 \text{ mm}^3$. A different subject was used for the 2 experiments. The sequence was timed to give an image every 2 seconds ($1.25 \times 1.25 \times 3 \text{ mm}^3$) or 3 seconds ($1 \times 1 \times 2 \text{ mm}^3$). The same visual

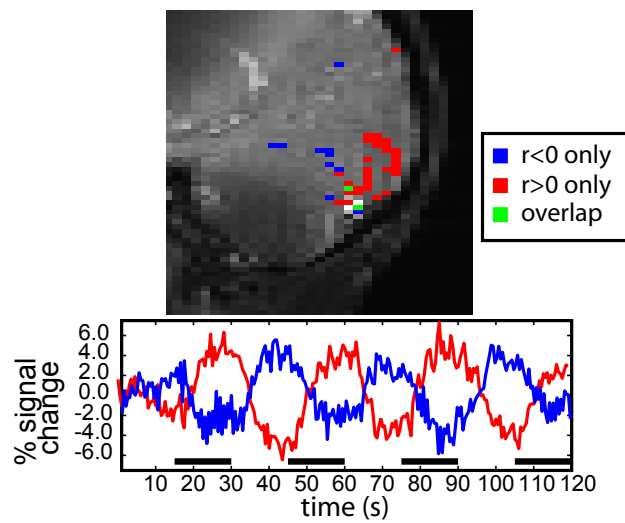


Figure 4.4: Positive and negative correlations in BOSS FMRI of visual cortex (zoomed). Comparing voxel correlations at different center frequencies, 33 voxels were strictly positively correlated (red), 13 voxels were strictly negatively correlated (blue) and 2 were positively correlated at one frequency and negatively correlated at a different frequency (green). Positive and negative correlations also tended to occur in distinct locations. Further experimentation will be required to determine the source of the negative correlations.




Figure 4.5: High-resolution fMRI with BOSS at voxel sizes of (a) $1.25 \times 1.25 \times 3 \text{ mm}^3$ and (b) $1 \times 1 \times 2 \text{ mm}^3$. Activation maps are overlaid on anatomical reference data. Activated voxels are confined to the gray matter off the calcarine fissure in primary visual cortex.

stimulus was presented as in the previous experiments. The 2-minute experimental runs were repeated 2 ($1.25 \times 1.25 \times 3 \text{ mm}^3$) or 6 ($1 \times 1 \times 2 \text{ mm}^3$) times. Analysis was performed using FSL/FEAT (<http://www.fmrib.ox.ac.uk/fsl>). Repeated runs were registered, averaged and processed using FEAT with cluster thresholding (cluster threshold $Z \geq .3$, cluster significance threshold $P \leq 0.01$).

The activation maps for the experiments are shown in Fig. 4.5, overlaid on high-resolution T1-weighted structural images. In both experiments, activation is confined to the gray matter and is focused on the calcarine fissure. In each data set, the percent signal change due to activation was calculated for the voxel with the greatest Z-statistic. The lower-resolution scan ($1.25 \times 1.25 \times 3 \text{ mm}^3$) had a signal change of 9.1% and the higher-resolution scan ($1 \times 1 \times 2 \text{ mm}^3$ data) had increased signal change of 15.7%. Previous work has reported increased signal change at smaller voxel size due to reduction of partial volume effects [76–79].

Although primary visual cortex (located along the calcarine fissure) is expected to activate strongly in response to the stimulus, higher visual areas in the occipital pole also tend to activate in response to this type of stimulus. The small region of activation in Fig. 4.5 likely results from frequency drift across the occipital lobe,




Figure 4.6: Demonstration of improved spatial coverage in high-resolution BOSS FMRI with multi-frequency measurements. Activation maps for an experiment with a voxel size of $1.25 \times 1.25 \times 3 \text{ mm}^3$ is shown with (a) a single frequency measurement and (b) the combined data from measurements at 2 center frequencies.

such that when the center frequency is matched to the calcarine fissure, other visual areas are not matched. As a result, the BOSS signal is only sensitive to areas near the calcarine fissure. If this is the case, spatial coverage could be improved with measurements at multiple center frequencies as described above. Figure 4.6 shows a multi-frequency MCP of 2 frequency measurements at $1.25 \times 1.25 \times 3 \text{ mm}^3$ resolution. For the same level of thresholding, the multi-frequency activation map exhibits a greater region of activation.

These high-resolution results were acquired at 1.5 T (relatively low-field) without special hardware. It is expected that these results may be improved upon at higher field and with the use of surface coils. The high SNR of these results suggests that even smaller voxel sizes should be obtainable with the current experimental setup.

4.4 Theoretical Comparison with BOLD

Although BOLD and BOSS FMRI both measure signal changes related to the deoxy-hemoglobin frequency shift, BOSS FMRI measures this shift directly, whereas BOLD

detects the shift indirectly as signal dephasing. This gives BOSS several advantages over BOLD including improved image quality and higher SNR. In addition, BOSS detects activations in the visual cortex with higher functional contrast than is typically found in BOLD because the BOSS signal change is based on signal inversion rather than attenuation. The following sections analyze various aspects of the BOSS signal, and further compare BOSS with BOLD FMRI.

4.4.1 Functional Contrast

BOSS FMRI is expected to have improved functional contrast over BOLD imaging. Since the BOSS signal change is due to a pool of spins that switch from deoxygenated (negative signal) to oxygenated (positive signal), the signal change is roughly twice the size of the exchanging pool (see the vector diagram in Fig. 4.2d). In comparison, the signal change in BOLD imaging is an attenuation of the exchanging pool, resulting in a smaller percentage signal change. We therefore expect improved functional contrast in BOSS FMRI over BOLD imaging. This prediction is supported by our initial results that found 2–3 times the functional contrast of BOLD FMRI in a single frequency measurement and more than 4 times the functional contrast when multiple measurements are combined.

Thus far we have considered the magnetization to consist of two discrete pools of spins with shifted resonance frequencies. A more realistic model of BOSS contrast would consider a spectrum of resonance frequencies as described by Scheffler [?]. This spectrum is usually modeled as a summation of Lorentzian distributions centered at the deoxyhemoglobin and water frequencies [80]. Increasing the concentration of deoxyhemoglobin within a voxel causes this spectrum to broaden [81,82]. This concept of spectral broadening was introduced in Chapter 3 as the source of BOLD contrast. The signal measured in SSFP imaging is the integral of the spectrum multiplied by the SSFP off-resonance profile [?]. Since BOSS FMRI inverts a fraction of the spectrum, small changes in the size of this fraction can be detected as a change in the signal magnitude. While such a model of BOSS signal dynamics is more complete, the

simpler description above is sufficient for the present discussion.

4.4.2 Image Quality

BOSS and BOLD imaging differ significantly in terms of image quality. With conventional BOLD imaging, a long echo time is required for the T_2^* contrast to evolve, resulting in lost signal due to dephasing and relaxation [53,83]. In addition, the long readouts used in BOLD produce image distortion (in EPI) or blurring (in spiral imaging) due to off-resonance and susceptibility boundaries [50–52, 83, 84]. These effects couple BOLD contrast to image degradation. In comparison, BOSS functional contrast is inherent to the steady-state signal, allowing independent optimization of the image acquisition to achieve high SNR and minimize image artifacts. At $T_E=T_R/2$, the BOSS signal is similar to a spin echo in that it has no T_2^* effects and therefore no signal dropout [?, 85]. Provided the signal is collected in the temporal vicinity of this spin echo, the image will be essentially free of T_2^* artifacts. Furthermore, with long readouts, BOSS would be effectively immune to off-resonance since the functional contrast occurs only near resonance. Overall, BOSS FMRI is expected to exhibit significantly improved image quality over BOLD imaging.

Figure 4.7 compares BOLD and BOSS image quality for the same coverage (16 cm FOV, 128×128) and frame rate (1 s). These phantom images were acquired at 1.5 T with typical FMRI protocols for EPI-BOLD ($\alpha=70^\circ$, $T_R = 1000$ ms, $T_E = 50$ ms) and 2DFT-BOSS ($\alpha=5^\circ$, $T_R = 7.8$ ms, single frequency). The EPI-BOLD images exhibit significant warping and ghosting, while the BOSS images have no such artifacts. In addition, the EPI-BOLD images appear to have lower resolution than the BOSS images (insets of Fig. 4.7) even though the k -space coverage is the same. The blurring of the EPI image is due to T_2^* decay over the single-shot readout [53].

The balanced-SSFP images shown in Fig. 4.3 exhibit signal intensity variations due to spatial drift in the magnetic field (this effect can also be seen at the bottom of the SSFP phantom image in Fig. 4.7). This signal attenuation is quite different




Figure 4.7: Comparison of typical SSFP images for BOSS FMRI ($T_R/T_E=7.8/3.9$ ms, 2DFT, 128×128 , $\alpha = 5^\circ$) with typical GRE images used in BOLD studies ($T_R/T_E=1$ s/60 ms, EPI, 128×128 , $\alpha = 65^\circ$) in a phantom. The BOLD images are more distorted and have lower apparent resolution than the BOSS images due to T_2^* decay over the readout.

from BOLD signal dropout near regions of susceptibility changes (such as the sinus cavities). While BOLD signal dropouts are unrecoverable, the specific spatial pattern of BOSS magnitude variations is dictated by the choice of center frequency. In the BOSS FMRI experiments presented above, the center frequency was chosen to achieve high signal (and therefore functional contrast) in the occipital lobe. High signal can easily be achieved in other brain areas by shifting the center frequency to match the resonance frequency of the region of interest.

Complete coverage should be achievable in BOSS imaging with acquisition at multiple center frequencies. Linear combination of multi-frequency acquisitions has been used previously to achieve high-quality anatomical images with balanced SSFP [86]. Whereas susceptibility boundaries introduce unrecoverable signal loss in BOLD imaging, they simply represent steep off-resonance gradients for BOSS imaging. As with other sources of off-resonance, regions of high susceptibility can be addressed in BOSS using multi-frequency measurements.

4.4.3 Signal-to-noise Ratio

Another advantage of BOSS FMRI is the high SNR of balanced SSFP. The SNR efficiency (η) of a pulse sequence depends on the (normalized) signal magnitude $|M_{xy}|/M_0$ and the duty cycle (the readout time T_{read} divided by the imaging time T_{im}) [87]:

$$\eta \propto \sqrt{\frac{T_{read}}{T_{im}}} \cdot \frac{|M_{xy}|}{M_0} \quad (4.1)$$

Balanced-SSFP imaging generally has high SNR efficiency because a large fraction of the total imaging time can be spent collecting data. The duty cycle in anatomical SSFP imaging tends to be limited by the need for a short T_R , which dictates short readouts. BOSS FMRI can obtain a particularly high readout duty cycle because the T_R can be longer than for most balanced-SSFP sequences (10–20 ms, as discussed

below), allowing the majority of each T_R to be dedicated to acquisition. In contrast, the long- T_E GRE sequences used in BOLD FMRI have fairly low SNR efficiency. There are two reasons for this. First, a significant portion of the signal must decay in order for BOLD contrast to develop. Second, BOLD FMRI requires long T_{RS} (0.5–3 s) to allow T_1 recovery, and only a small fraction of this time can be dedicated to data acquisition. These factors lead to higher SNR efficiency (η) for BOSS FMRI.

Table 4.1 compares the SNR of BOSS and BOLD for a standard FMRI experiment at 1.5 T with the same volume coverage (12 slices with a 64×64 matrix) and frame rate (1 s). In BOSS, multiple slices would be acquired using a 3D acquisition (which does not disrupt the steady-state, unlike 2D multi-slice), allowing a high duty cycle in which the entire magnetization contributes signal every T_R . GRE-BOLD sequences do not gain much from 3D acquisitions due to incomplete T_1 recovery between RF pulses [88, 89], so BOLD data is usually collected using multi-slice acquisitions. However, for completeness, a 3D BOLD analysis is also included in Table 4.1. In the 2D BOLD acquisition, each slice is excited and read out separately, so the readout time used in the SNR analysis is the time to read out a single slice. For the 3D BOLD acquisition, a 3D stack-of-EPI trajectory is assumed with the same EPI trajectory for each z phase encode as was used in the 2D BOLD acquisition. In the comparison of Table 4.1, 3D BOSS FMRI has roughly 3 times the SNR efficiency of 2D BOLD FMRI and about 4.5 times the SNR efficiency of 3D BOLD. This gain for 3D BOSS over 2D BOLD depends on the duty cycle of the BOSS readout, and ranges from a factor of 2 for a 40% duty cycle to 3 for a 90% duty cycle.

4.5 BOSS Signal Characteristics

4.5.1 Contrast-to-noise Ratio

This section considers the functional contrast-to-noise (CNR) of BOSS and possible optimizations of sequence parameters based on CNR. The CNR of BOSS has a

	$ M_{xy} $	T_{read} (ms)	duty cycle	η
3D BOLD	0.05	384	38.4%	1
2D BOLD	0.26	32	3.2%	1.5
3D BOSS	0.15	820	82%	4.4

Table 4.1: Relative SNR efficiency (η) for BOLD and BOSS images at 1.5 T with volume acquisition every 1 s. The signal $|M_{xy}|$ is expressed as a fraction of M_0 ($T_1=880$ ms, $T_2=80$ ms, $T_2^*=60$ ms). The BOLD magnetization is calculated at the Ernst angle, and the BOSS magnetization is taken as the peak magnitude in Fig. 4.2. In 1 second, each method could acquire 12 slices with a 64×64 matrix. In BOSS: 3D interleaved EPI, $ETL=8$, $\alpha=5^\circ$, $T_R/T_E = 10$ ms/5 ms, 16 μ s samples, $T_{read}=12 \times 64^2 \times 16 \mu$ s = 820 ms. In 2D BOLD: multi-slice, single-shot EPI, $T_{read}=32$ ms, $\alpha=65^\circ$, $T_R/T_E = 1000$ ms/60 ms. In 3D BOLD: stack-of-EPI with the same EPI readout and T_E as 2D BOLD and the 3D slab excitations separated by $T_R = 1000/12 = 84$ ms. 3D BOSS is 3 times more SNR efficient than 2D BOLD and 4.4 times more efficient than 3D BOLD.

strong dependence on resonance frequency not found in most sequences. Assuming an activation-induced frequency shift Δf , the functional contrast at a given off-resonance frequency f is:

$$\text{Contrast}(f) = \frac{|M_{xy}(f - \frac{\Delta f}{2}) - M_{xy}(f + \frac{\Delta f}{2})|}{M_0} \quad (4.2)$$

where normalization by M_0 expresses the contrast as a fraction of the total possible signal. For a given imaging setup, the thermal noise (σ) is independent of resonance frequency, allowing direct comparison of the CNR at different frequencies based on Eq. 4.2:

$$\text{CNR}(f) = \frac{\text{Contrast}(f)}{\sigma} \propto \text{Contrast}(f) \quad (4.3)$$

In other words, the proportionality in Eq. 4.3 results from the fact that σ is a constant

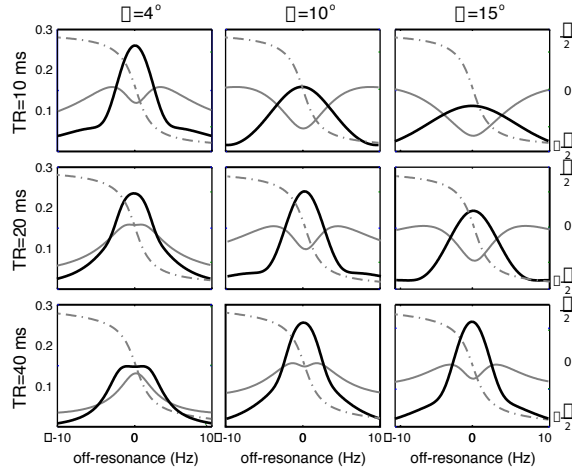


Figure 4.8: BOSS functional CNR for several flip angles ($\alpha=4^\circ$, 15° , 25°) and T_{RS} ($T_R=10$, 20 , 40 ms). The gray lines show the balanced-SSFP signal profile (solid=magnitude, dashed=phase) and the black lines represent BOSS functional CNR (Eq. 4.3). CNR calculations assume an activation-induced frequency shift of 5 Hz. The plots along the diagonal have the highest peak CNR for the given T_R . As the T_R increases, the flip angle with highest peak CNR also increases.

for a given experimental setup. The calculations presented here express CNR as a fraction of the total signal, as suggested by Eq. 4.2. Note that the CNR calculations presented here consider only thermal noise, which is not necessarily the dominant noise in fMRI experiments, which tend to be sensitive to “physiological noise” (e.g., cardiac and respiratory effects).

Figure 4.8 plots BOSS CNR over a band of off-resonance frequencies for several choices of T_R and flip angle, assuming an activation-induced frequency shift of 5 Hz [?, 27]. For most choices of flip angle and T_R , functional contrast peaks on-resonance where the phase transition has greatest slope. For the approximate parameters used in this study (upper left-hand panel in Fig. 4.8), BOSS exhibits strong functional CNR over a band of frequencies near-resonance, outside of which the functional CNR is considerably reduced and largely independent of resonance frequency.

4.5.2 BOSS Contrast Optimization

The plots in Fig. 4.8 suggest that the flip angle can be optimized for a given T_R based on CNR. One obvious criterion for optimization would be peak CNR. The profiles depicted along the diagonal of Fig. 4.8 represent the flip angle with greatest peak CNR for the given T_R . As the T_R is lengthened, the flip angle with greatest peak CNR increases. As indicated by the solid gray lines, this increase essentially maintains the magnitude profile from Fig. 4.2b.

This type of optimization can provide significant gains in peak CNR. For example, Fig. 4.9 demonstrates the potential gain in peak CNR that can be obtained for the T_R used in [?]. Figure 4.9b represents the functional CNR used in their study, where the dashed line indicates their placement of the center frequency (chosen as a region of high slope in the balanced-SSFP magnitude profile). Functional CNR of this imaging setup can be increased in 2 ways. First, functional CNR for the same flip angle ($\alpha = 50^\circ$) can be increased by setting the center frequency to gather data at the CNR peak (i.e., by sliding the dashed line in Fig. 4.9b to 0 Hz off-resonance). Second, decreasing the flip angle to $\alpha = 15^\circ$ increases the signal magnitude near-resonance, further increasing the peak functional CNR (as shown by the dashed line in Fig. 4.9a). As can be seen by comparing the dashed lines in Fig. 4.9a and b, proper choice of flip angle and center frequency could roughly triple the functional CNR.

A second possible criterion for optimizing the flip angle is signal stability. Most of the profiles in Fig. 4.8 have a fairly sharp peak in CNR, which could be problematic in the presence of fluctuations in voxel resonance frequency. A flatter CNR profile (at the expense of peak CNR) may be preferable to maintain signal stability. For example, the profile obtained with $T_R=40$ ms and $\alpha = 4^\circ$ (lower left-hand panel in Fig. 4.8) would be more immune to small drifts in the resonance frequency since the CNR plateaus near resonance. This trade-off of peak CNR for signal stability offers an additional design parameter in BOSS experiments that is not present in BOLD FMRI.

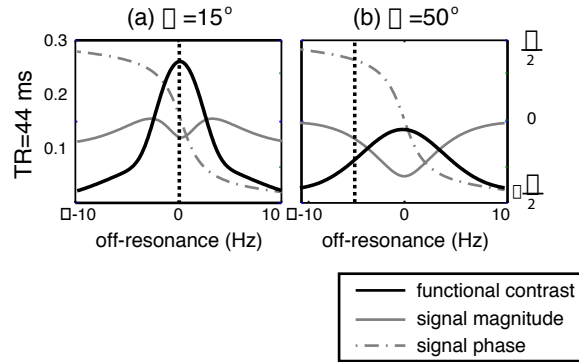


Figure 4.9: Functional CNR (dark solid lines) and balanced-SSFP signal profiles (gray solid=magnitude, gray dashed=phase) for the T_R used in [?]. Curves are shown for the flip angle used in that study (50°) and the flip angle with highest peak CNR (15°). The method of [?] sets the center frequency off-resonance (e.g., the vertical dashed line in the right-hand panel) to obtain functional contrast based on the magnitude profile. However, the CNR peaks on-resonance due to the phase transition. For this T_R , the peak CNR is even higher for a flip angle of 15° (vertical dashed line in left-hand panel).

In general, the ability to optimize BOSS parameters based on CNR increases flexibility in the design of fMRI experiments. For example, measurements of physiological fluctuations may favor a more peaked CNR, whereas activation experiments would benefit from a more stable signal to reduce sensitivity to physiological noise. The latter optimization may improve the functional CNR over that found in our data.

4.5.3 Balanced-SSFP Phase Transition

BOSS functional contrast is based on the abrupt phase transition balanced-SSFP sequences exhibit near-resonance. This phase change is an unusual effect that warrants discussion.

First, we consider the source of the SSFP phase profile. Species in the flat portions of the phase profile at $T_E = T_R/2$ form a steady-state in which off-resonance precession during the T_R is canceled by the RF tip between T_R s. Assuming the RF tip is about

the imaginary axis (as is used in our simulations), spins that precess counter-clockwise form a spin echo on the positive imaginary axis halfway through the T_R (Fig. 4.10a) [?]. Similarly, spins that precess clockwise form a similar echo on the negative imaginary axis (Fig. 4.10c). However, spins in the phase transition form a very different steady-state. These spins are near-resonance, and therefore precess very little over a T_R . For these species, relaxation is the dominant effect during the T_R , so the RF tip primarily serves to cancel relaxation (Fig. 4.10b). Because the RF pulse tips along the imaginary axis, this steady-state is formed along the real axis. A given off-resonance frequency forms a steady-state that reflects the relative magnitude of relaxation and precession effects during a T_R . These steady-states lead to the phase profile shown in Fig. 4.2.

In order for BOSS FMRI to be a reliable method for detecting activation, it is critical that the SSFP phase profile be robust across a range of likely imaging conditions. In particular, we desire a stable transition region, since the transition is the source of functional contrast. The effects of various acquisition and physiological parameters on the phase profile are shown in Fig. 4.11. The most important characteristic of the phase profile is that the transition region depends only on T_2 , and is unaffected by flip angle, T_R , T_E and T_1 . The width of the phase transition increases for decreasing T_2 . However, this effect is minor for the relaxation times likely to be encountered in brain imaging. Lengthening the T_R pushes the phase transitions closer together while preserving the absolute bandwidth over which the transition occurs. Use of $T_E \neq T_R/2$ changes the phase profile outside of the transition region (reflecting the formation of spin echoes halfway through the T_R [?]). Spins in the phase transition, which are near-resonance, precess very little during each T_R and their phase is essentially unaffected by the choice of echo time. The phase profile is completely independent of flip angle and T_1 . The fairly simple dependence of the phase profile on tissue and pulse sequence parameters indicates that BOSS contrast should be robust under a wide range of imaging conditions.

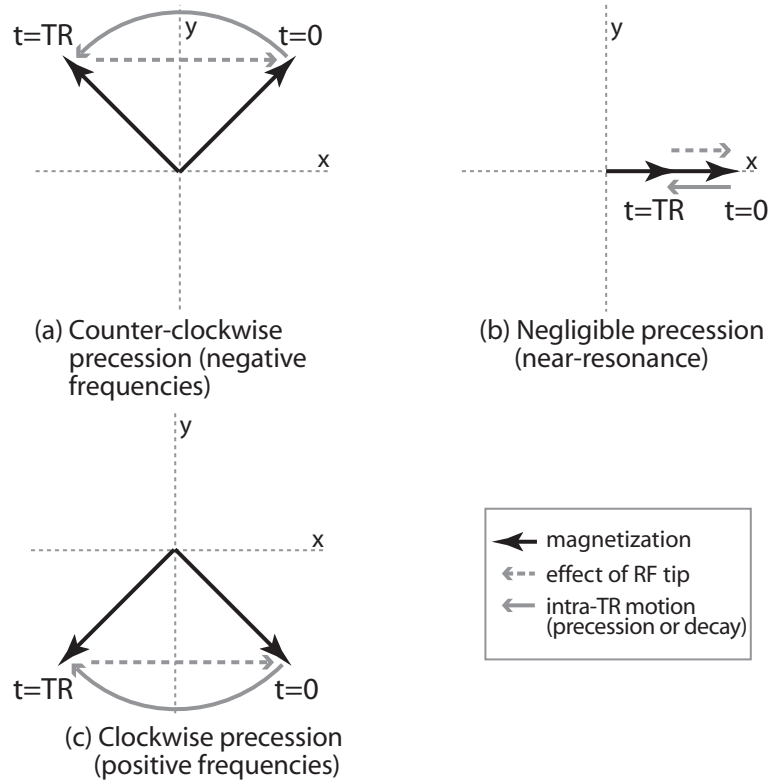


Figure 4.10: Steady-state dynamics giving rise to the SSFP phase transition, assuming the RF pulse is applied along the imaginary (y) axis. The transverse magnetization is shown immediately after ($t=0$) and before ($t=T_R$) the RF pulse. In all cases, a steady-state is formed such that the RF pulse (dashed arrow) exactly cancels the motion of the magnetization during the T_R (solid arrow). Species for which off-resonance precession is the dominant intra- T_R motion form the steady-states shown in (a) and (c). The RF pulse serves to exactly cancel precession. These species form spin echoes on the positive or negative imaginary axis, depending on the direction of precession. Species for which relaxation is the dominant effect, usually because precession is negligible, form the steady-state shown in (b). Each diagram shows only the dominant intra- T_R motion, precession or relaxation.

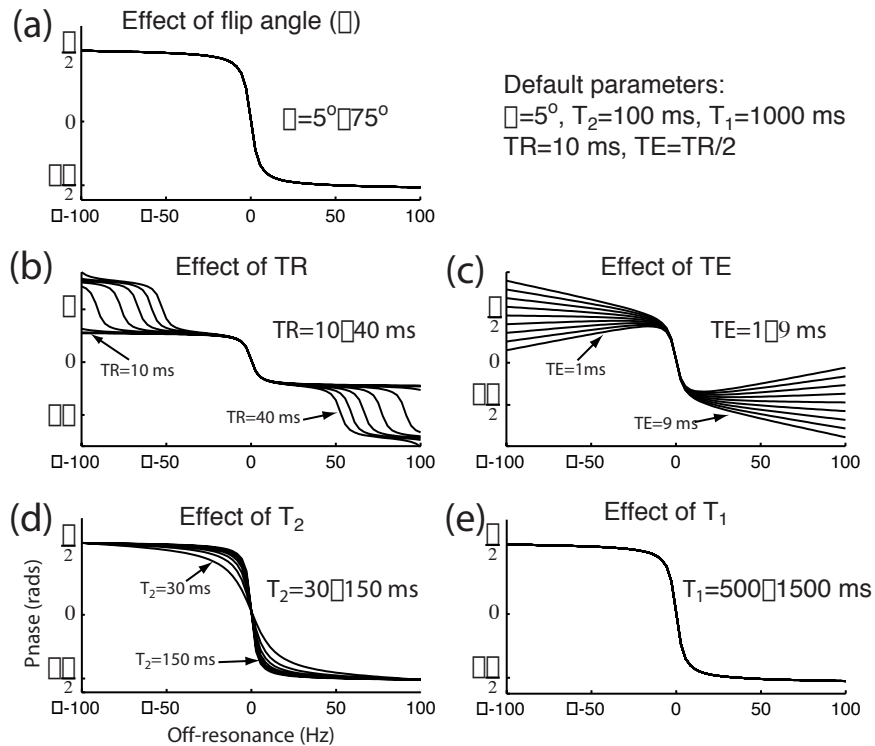


Figure 4.11: The effect of acquisition and physiological parameters on the SSFP phase profile. Each plot shown varies one of the parameters (flip angle, T_R , T_E , T_1 and T_2) while holding the others constant at the default parameters. Flip angle and T_1 have no effect on the phase profile. Increasing the T_R decreases the distance between phase transitions without changing the transition itself. The echo time changes the slope of the area outside the transition band. The only parameter that affects the phase transition itself is T_2 , which has the effect of sharpening the transition for longer T_2 .

4.5.4 BOSS Imaging Parameters

BOSS imaging diverges from both BOLD fMRI and balanced SSFP in the choice of optimal imaging parameters. Unlike BOLD fMRI, the functional contrast is independent of T_E since the phase transition is preserved throughout the T_R . The formation of spin echoes at $T_R/2$ makes this a good choice of echo time since image artifacts will be minimal. However, shifted echo times simply introduce a slight gradient in the off-resonance signal phase as shown in Fig. 4.11c while the on-resonance region where functional contrast is maximized is essentially independent of echo time. This lack of dependence on echo time is different from BOLD fMRI, where the echo time must be long to allow functional contrast to evolve. BOSS contrast allows flexibility in the choice of echo time, and therefore readout trajectory.

BOSS fMRI also has several important differences from anatomical SSFP imaging. As discussed above, the functional CNR of BOSS can be used to optimize the flip angle for a given T_R . Additionally, the T_R has a very different impact in BOSS imaging as compared to anatomical SSFP imaging. In balanced-SSFP imaging (using high flip angles), the T_R is usually minimized to increase the distance between the signal magnitude nulls (shown in Fig. 4.2a). With sufficient shimming, this allows the object to be placed within the high-signal portion of the off-resonance profile, avoiding the characteristic “banding artifacts” caused by the magnitude nulls. In BOSS fMRI, activation can only be detected near the phase transition, meaning that only this portion of the off-resonance profile is useful. Increasing the T_R pushes the phase transitions closer without affecting the transition width (see Fig. 4.11b), making BOSS contrast compatible (and perhaps superior) with long T_R s. The ability to image with long T_R s and small flip angles should enable BOSS fMRI at high field where conventional balanced-SSFP methods are limited due to banding artifacts and SAR considerations.

4.6 Summary

We have presented a new method for functional MRI that directly detects dynamic changes in the deoxyhemoglobin concentration using a Blood Oxygenation Sensitive Steady-state (BOSS) signal. This steady state inverts the frequency-shifted deoxyhemoglobin signal relative to the water signal so that changes in the deoxyhemoglobin concentration during activation cause a signal modulation. BOSS imaging requires careful placement of the center frequency relative to the water and deoxyhemoglobin resonance frequencies, which can be achieved with careful shimming and acquisition at a small number of center frequencies. Additionally, BOSS may be more sensitive to physiological fluctuations than BOLD imaging. BOSS FMRI has several advantages over traditional BOLD FMRI. First, BOSS has fewer image artifacts because signal is acquired at a spin echo with short, multi-shot readouts. Second, BOSS data has higher SNR than BOLD due to the large fraction of imaging time that is dedicated to data acquisition. Third, BOSS functional contrast is greater than is usually seen in BOLD imaging because the BOSS signal change is produced by a signal inversion rather than an attenuation. These advantages make BOSS FMRI a powerful alternative for functional neuroimaging.

Chapter 5

Navigated Diffusion Imaging with SSFP

SSFP-DWI has long been known to achieve strong diffusion contrast with relatively small diffusion gradients and without extensive signal decay. In addition to being SNR efficient, SSFP is a natural framework for 3D acquisitions, which implies potential for high-resolution DWI. Because SSFP-DWI is fundamentally a multi-shot method, it is extremely sensitive to motion, which has precluded it as a practical method. In recent years, various groups have developed methods for removing motion artifacts from DWI data using phase references known as navigators. This work presents navigated SSFP-DWI, with the ultimate goal of high-resolution DWI.

5.1 SSFP-DWI

As discussed in Chapter 3, diffusion weighting is usually added to a pulse sequence with a pair of large gradients with the same area but opposite polarity. Diffusive motion results in phase dispersion, and therefore signal attenuation. SSFP-DWI is able to obtain similar diffusion weighting with a single, unbalanced gradient (Fig. 3.4c)

[59]. The source of this diffusion weighting is based on the formation of a steady-state in which an effective bipolar set is formed from this single diffusion gradient over multiple T_{RS} .

5.1.1 Diffusion Contrast

Coherence pathway analysis (introduced in Chapter 2) makes it easier to see that the diffusion gradient in SSFP-DWI can form an implicit bipolar set over multiple T_{RS} . Instead of forming an echo during every transverse period, the presence of the unbalanced diffusion gradient allows echoes to form only after an even number of transverse periods [90]. Figure 5.1a shows an example echo pathway that forms an echo in the final repetition period. In this echo, the magnetization excited during the first transverse period is dephased by the diffusion gradient, tipped longitudinal by the next RF pulse, and re-enters the transverse plane $4 T_{RS}$ later. The magnetization only forms an echo in this final repetition period when the diffusion gradient rephases the signal (see Fig. 5.1b). The total signal is a summation of echoes similar to that depicted in Fig. 5.1a, each with a different number of longitudinal periods.

To understand why this sequence results in strong diffusion weighting, consider the signal dynamics diagrammed in Fig. 5.1. The diffusion gradient finely encodes position during the first transverse period and decodes position during the second transverse period. If spins have diffused during the intervening longitudinal periods, the magnetization is incompletely refocused, resulting in diffusion contrast. The amount of attenuation a given echo experiences depends on the length of time spins are allowed to diffuse, which is proportional to the number of longitudinal periods. Echoes with a large number of longitudinal periods experience heavy diffusion weighting. These echoes are the primary source of diffusion weighting in SSFP-DWI.

Since the SSFP-DWI signal is a weighted summation of many diffusion-weighted echoes, the signal attenuation due to diffusion is more complicated than in standard DWI methods (which sample a single echo). In standard SE-DWI or STE-DWI, the




Figure 5.1: Signal dynamics for an example echo pathway in SSFP-DWI. The pulse sequence is approximated as a series of RF pulses with an intervening diffusion gradient (so that imaging gradients are ignored). (a) The example echo pathway consists of 2 transverse periods separated by 4 longitudinal periods. (b) The diffusion gradient dephases the signal in the first transverse period and rephases the signal is refocused to form an echo in the second transverse period. Diffusion motion during the longitudinal periods causes signal attenuation.

effect of diffusion gradients is a simple attenuation of the unweighted signal by a factor $\exp(-bD)$ where b depends only on the diffusion gradient amplitude and timing. In SSFP-DWI, each coherence pathway has a different amount of diffusion weighting, creating a more complicated signal attenuation which depends on T_1 , T_2 , D , α and the diffusion gradients.

A simple but useful approximation of the signal attenuation was suggested by Buxton [90]. While the full signal expression includes all echo pathways with an even number of transverse periods, the signal is dominated by echoes that experience only 2 transverse periods [90] (i.e., spin and stimulated echoes). By explicitly considering only these echoes, the signal attenuation A_M can be approximated as:

$$A_M = \frac{M(G\tau)}{M(0)} \approx \frac{A_D(1 + A_D E_1)(1 - E_1 \cos(\alpha))}{(1 + E_1)(1 - A_D E_1 \cos(\alpha))} \quad (5.1)$$

where $A_D = \exp(-(\gamma G\tau)^2 T_R D)$ is the diffusion attenuation factor for a diffusion gradient of strength G and duration τ , and $E_1 = \exp(-T_R/T_1)$. Equation 5.1 is independent of T_2 since all echo pathways have 2 transverse periods and therefore experience identical T_2 decay. This approximation is accurate provided $T_R \geq T_2$ and the flip angle is small. As was illustrated in Fig. 2 of [90], the need for small flip angle is lessened as the T_R becomes long, albeit with an accompanying loss of SNR.

Equation 5.1 has a different form than standard DWI methods. For convenience of description, we will refer to the amount of diffusion weighting in SSFP-DWI using the b -value which would give equivalent attenuation in a tissue of interest (i.e., $b_{eff} : A_M = \exp(-b_{eff} D)$). This allows direct comparison of the diffusion weighting in SSFP-DWI with standard DWI methods.

5.1.2 Relaxation Effects

In addition to strong diffusion weighting, SSFP-DWI has attractive relaxation properties that result from the signal dynamics discussed above. If we consider only echo pathways like those shown in Fig. 5.2, the signal in a given echo undergoes T_2 decay during 2 repetition periods, and relaxes according to T_1 during the remainder of the time the echo is evolving. This property is important since T_2 signal decay tends to be the factor that ultimately limits the amount of diffusion attenuation a given pulse sequence can achieve.

In SSFP-DWI, diffusion weighting can be increased by biasing the signal more to echo components with heavy diffusion weighting (those with a large number of longitudinal periods). One way to achieve this is to reduce the flip angle [90], which increases the fraction of magnetization that remains longitudinal following a given RF pulse and therefore increases the fraction of the magnetization in echo pathways with more longitudinal periods. Since all pathways experience the same amount of T_2 decay, diffusion weighting is increased without incurring additional T_2 decay. By comparison, increasing the diffusion weighting in SE-DWI usually necessitates




Figure 5.2: Relaxation effects for the same SSFP-DWI echo pathway shown in Fig. 5.1. (a) The example echo pathway consists of 2 transverse periods separated by 4 longitudinal periods. (b) The echo decays according to T_2 during the 2 transverse periods and T_1 during the 4 longitudinal periods. The signal decay due to T_2 is minimized for the given diffusion gradients. Since the majority of water diffusion occurs during the longitudinal periods, this allows long diffusion times without significant signal loss due to T_2 decay.

lengthening the echo time to accommodate longer diffusion gradients. Hence, in SE-DWI, increasing diffusion contrast results in more signal loss due to T_2 decay.

These relaxation properties makes SSFP-DWI a promising sequence for imaging short- T_2 species (e.g., cartilage), which undergo severe signal decay with standard DWI methods such as SE-DWI. These properties are also useful with longer- T_2 species (e.g., white matter) when unusually heavy diffusion weighting is desired. SSFP-DWI may have significant advantages for applications such as high angular resolution DTI which require strong diffusion weighting.

5.1.3 Quantification

There have been several schemes for quantification of D using SSFP-DWI, which vary in complexity and accuracy. In an early paper, Buxton attempted to fit the full expression for the SSFP-DWI signal using assumed relaxation values [90], but

found consistent overestimations of the diffusion coefficient. This overestimation was hypothesized to result either from underestimation of gradient moments or (more likely) variation in the flip angle over the slice. A second group proposed calibrating the SSFP-DWI signal using a phantom with similar relaxation values to a tissue of interest and approximating signal attenuation as exponential [91]. Unfortunately, this method also assumes specific relaxation values, would require calibration for all possible imaging parameters and would also be sensitive to flip angle. Recent work has proposed inversion of the full signal expression with measured relaxation and proton density values [92]. The signal expression is transformed into a 5th order polynomial which can be solved for the diffusion coefficient. While this method is complete and was shown in phantoms to provide valid quantification, the polynomial expression is quite complicated and is likely to be sensitive to noise.

Given T_1 , Eq. 5.1 involves a single unknown, D , which can be fit using standard non-linear optimization methods. This procedure should be simpler and less error-prone than previous methods. The requirement that $T_R \geq T_2$ makes this method useful primarily for short- T_2 species. The SNR efficiency of SSFP-DWI would be compromised if this constraint were placed on brain imaging (where $T_2 \approx 100$ ms). However, for short- T_2 species such as cartilage ($T_2 \approx 20\text{--}40$ ms), this quantification method is much more appropriate.

Equation 5.1 can be fit to data acquired with several different levels of diffusion weighting using nonlinear optimization. We use a nonlinear, bounded optimization based on the Golden Section search algorithm [93] with a RMS-error (l_2 -norm) cost function and bounds $0.0 \leq D \leq 0.01$ mm²/s (which are loose bounds relative to the range of diffusion coefficients found in tissue). To evaluate the sensitivity of this fitting procedure to errors in the flip angle and T_1 , fits were made to simulated data generated from the full expression for the SSFP-DWI signal [90]. Simulations used imaging parameters appropriate for cartilage ($T_R = 30$ ms, $G = 7.8, 15, 27.5, 40$ mT/m, $\tau = 5.5$ ms, $\alpha = 25^\circ$), and the approximate tissue parameters for cartilage ($T_2 = 30$ ms, $T_1 = 800$ ms, $D = 0.0015$ mm²/s). In each fit, either the flip angle or the T_1 used in the fit was varied to simulate the effect of an incorrect estimate of the given parameter.




Figure 5.3: Sensitivity of the diffusion coefficient fit (D_{fit}) to the flip angle and T_1 used in the fit. Data was simulated as described in the text. The dashed lines in the plots above represent the “actual” values used to generate the simulated data. The diffusion coefficient was fit to the simulated data for a range of (incorrect) values for α and T_1 to evaluate the sensitivity of the fit to errors in these parameters.

Each fit was repeated 200 times with varying Gaussian noise at $\text{SNR} \approx 110$ (the SNR measured from the cartilage imaging experiments described below).

The results of these fitting experiments are shown in Fig. 5.3. The fitting procedure is sensitive to errors in the flip angle (Fig. 5.3a) due to the strong dependence of the diffusion weighting on flip angle. Underestimates of α lead to underestimates of D since at smaller flip angles the expected attenuation due to diffusion is greater (i.e., if $\alpha_1 < \alpha_2$, $A_M(\alpha_1) > A_M(\alpha_2)$). Accurate quantification will require careful calibration of the flip angle. In addition, slice selective acquisitions would need to account for variations in the effective flip angle due to slice profile imperfections. The latter complication is largely avoided in 3D acquisitions where the slab can be assumed to have constant flip angle except at outer slices, which are usually discarded.

The fitting procedure is largely unaffected by the value of T_1 used in the fit, as shown in Fig. 5.3b. Underestimating the T_1 does introduce a slight overestimate of D since increased longitudinal decay is interpreted as diffusion-related attenuation.

However, for tissues with relatively long T_1 the amount of increased decay is small, leading to fairly insignificant fitting errors. These simulations indicate that the fit should be very robust against errors in T_1 since the range of T_1 values considered in these simulations is larger than that arising due to physiological variation, even in disease states. The greatest fit error in our simulation (by halving the T_1) only introduces a 17% error in the fit, but it is unlikely that a reasonable estimate of T_1 will be this inaccurate.

5.2 Navigated SSFP-DWI

As discussed in Chapter 3, motion is a significant issue in all DWI sequences because sensitivity to small, diffusive motion of water molecules necessarily implies sensitivity to larger movements of the underlying object. Bulk motion in the presence of diffusion gradients results in phase offsets in the magnetization, creating image artifacts in multi-shot acquisitions when this phase varies from one readout to another [66]. The most common way to avoid motion artifacts in DWI data is to acquire the entire image in a single acquisition (for example, with single-shot EPI). However, single-shot readouts have severely limited image resolution and are highly sensitive to off-resonance. SSFP-DWI is only SNR efficient with the use of multi-shot readouts (and, in particular, 3D acquisitions), making motion particularly problematic for this sequence.

One method for reducing motion artifacts in multi-shot DWI is to measure and remove phase errors created by motion. This requires the acquisition of a low-resolution, full-FOV phase reference known as a navigator [66, 70, 72]. Each excitation, the navigator is acquired along with the current high-resolution data frame. Provided the navigator data (a) samples the same phase error as the high-resolution data, and (b) covers k -space sufficiently to resolve this phase error, the navigator can be used to correct phase errors in the high-resolution data before the individual readouts are combined to make an image. Although early work used 1D navigators [66, 70], 2D

navigators have since been shown to be much more useful for removal of phase artifacts. In the brain, motion is in general spatially non-linear due to deformation of the brain during arterial pulsatility, necessitating a full-FOV phase reference. Navigator correction of navigated DWI data is discussed in Chapter 6. We now present several navigated trajectories with specific application to SSFP-DWI.

5.2.1 Spiral-navigated Spinwarp Trajectories

The basic sequence for navigated SSFP-DWI is shown in Fig. 5.4a. This sequence consists of a slice-selective RF excitation followed by a large, unbalanced diffusion gradient (which can be applied along any direction). This gradient is followed by two readouts, a low-resolution 2D spiral navigator and a high-resolution 2DFT readout. Each high-resolution phase encode is corrected based on the image-space phase measured by the corresponding navigator. The use of a spinwarp trajectory is well-suited to SSFP-DWI, but may be less useful with other sequences. Although spinwarp trajectories are not generally efficient and therefore tend not to be favored in most magnetization-prepared sequences, the short T_{RS} employed in SSFP sequences increase efficiency.

As discussed above, volumetric SSFP-DWI is most efficient when used with a 3D acquisition. One straightforward 3D extension of the 2DFT sequence is shown in Fig. 5.4b. This sequence is identical to the first except for the addition of phase encoding along the z direction. As can be seen on the right of Fig. 5.4b, this method acquires 3D high-resolution data along with a 2D navigator (that is, a navigator that is spatially resolved in x and y and is a projection along z). While this navigator is not sufficient to correct for all possible motion artifacts, it has been found to be highly effective when motion is likely to be restricted along one dimension, as is discussed below. The acquisition of a full 3D navigator is not a straightforward extension of 2D navigation due to the considerable increase in the necessary scan time. Implementation of 3D navigators will be considered in future work.

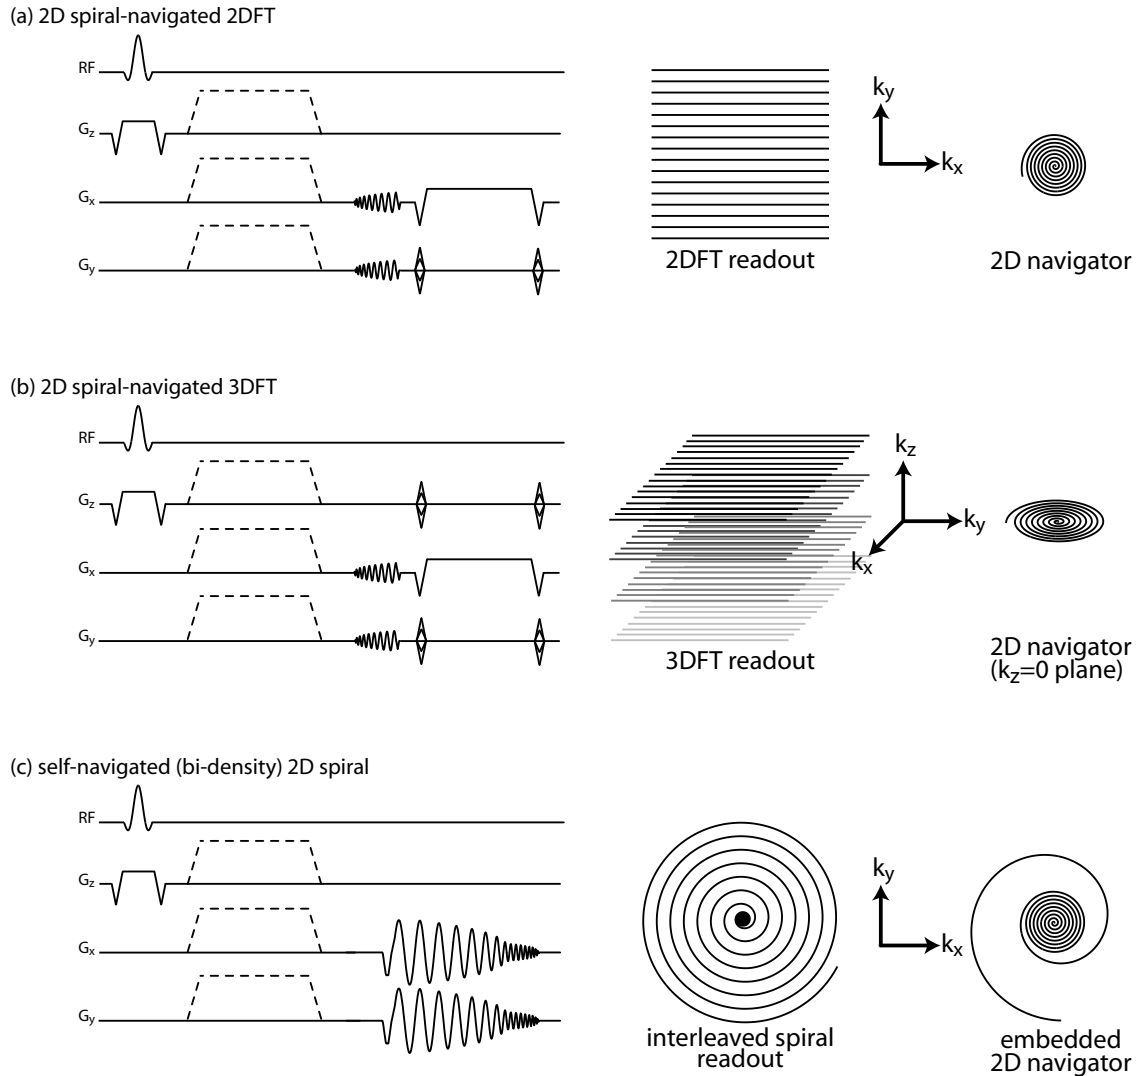


Figure 5.4: Navigated SSFP-DWI pulse sequences implemented in this work and their corresponding trajectories. All sequences employ the same 2D spiral navigator in conjunction with different high-resolution trajectories. The top 2 sequences use a separate navigator readout to correct the data from (a) 2DFT and (b) 3DFT high-resolution trajectories. The bottom sequence (c) uses a self-navigated spiral in which the inner, high-density portion of the trajectory serves as a navigator for the outer, multi-shot portion of the spiral.

5.2.2 Self-navigated Spiral Trajectories

Methods that acquire navigators separate from high resolution data are somewhat inefficient, as the navigator constitutes a sizable fraction of the readout time but does not contribute to the final image SNR. An appealing alternative is the use of a self-navigated trajectory in which the navigator is embedded in the high-resolution interleave. Such a trajectory is SNR-efficient while allowing full correction of motion artifacts. In addition, because self-navigated corrections rewind the phase to the real axis, the real part of the image can be displayed to further increase SNR [94].

A self-navigating trajectory should both (a) provide a low resolution image at the full FOV in each acquisition, and (b) provide a high resolution image over multiple acquisitions (interleaves). Because spiral trajectories are characterized by a sampling density that varies with k-space radius, they are well-suited to self-navigation. The bidensity spiral interleave shown in Fig. 5.4c has a high inner density and a low outer density with a smoothly varying transition region. This trajectory is self-navigating, with the inner portion serving as the navigator while the entire readout contributes to the high resolution image. The trajectory spirals inward to acquire the center of k-space close to the echo at $t=T_R$ [90] (note that this echo time is quite different from standard refocused SSFP, which forms an echo at $T_R/2$).

5.2.3 Factors Affecting the Choice of Navigated Trajectory

A well-designed sequence will dedicate hardware resources and imaging time primarily to the high-resolution trajectory rather than the navigator trajectory or the diffusion gradient. This section considers the realization of this goal within SSFP-DWI and its impact on navigation.

It is ideal for the majority of the repetition period to be dedicated to acquisition of the high-resolution data frame rather than the navigator acquisition or the diffusion gradient. Most importantly, the navigator readout should not be longer than the high-resolution readout since it does not contribute to the SNR of the final image.

This is one reason why self-navigated trajectories are particularly compelling: since the high-resolution readout contains the navigator, self-navigating trajectories by definition meet this constraint. Similarly, it is preferable not to dedicate a large portion of the repetition period to diffusion gradients. The ability to obtain strong diffusion weighting without requiring long-duration diffusion gradients is a compelling strength of SSFP-DWI. This enables a greater readout duty cycle as compared to other DWI sequences, which tend to spend the majority of the repetition period achieving diffusion weighting.

A second consideration that places practical constraints on navigated DWI trajectories is gradient hardware. DWI in general tends to be limited by the capabilities of gradient hardware, and the gradients required for heavy diffusion weighting often lead to heating of both the gradient amplifiers, as well as the gradient coils. Navigated trajectories increase heating by attempting to acquire a reasonably large portion of k -space as quickly as possible (e.g., a 16×16 region at the center of k -space). Ideally, gradient heating would be dominated by either the diffusion gradient or the high-resolution readout since these factors directly relate to image quality. For the specific hardware used in this work, amplifier heating tends to severely limit the efficiency of the SSFP-DWI pulse sequence.

These two considerations (allocation of imaging time and gradient heating constraints) can be used to qualitatively analyze the navigated trajectories presented above. The self-navigated trajectory proposed above is theoretically more efficient for 2D imaging than 2DFT since it constitutes a more efficient use of scan time. However, in practice, this trajectory requires considerable “dead time” in which the gradient amplifier is allowed to cool. In comparison, a spiral-navigated 2DFT trajectory introduces similar amplifier heating, but can be acquired over an extended duration (i.e., with reduced gradient amplitude and readout bandwidth), resulting in increased SNR.

Similarly, a spiral navigator is not obviously more efficient than other low-resolution single-shot trajectories. Spirals tend to be inefficient at the center of k -space where




Figure 5.5: Example (a) spiral and (b) EPI navigators with the same coverage (FOV=24 cm, 16×16 matrix). These navigators were designed for the gradient system used in this study ($G_{max}=40$ mT/m, maximum slew rate 150 mT/m/ms, $4 \mu\text{s}$ sampling). The navigators have similar duration (3.0 ms spiral vs 3.6 ms EPI), but the spiral distributes gradient heating over 2 axes and oversamples at the center of k -space.

the velocity of the k -space traversal is limited by the slew rate of the gradient coils. The navigators used here are slew-limited throughout the trajectory and therefore require longer readouts than would seem to be necessary based on the size of the reconstructed image. For example, a reconstructed 16×16 spiral navigator over a 24 cm FOV has only 256 voxels but requires 758 time samples. EPI might seem to be achievable fewer samples and therefore require less time than a spiral, but in fact suffers from similar problems. A ramp-sampled EPI navigator with the same resolution and FOV as the spiral above actually requires 898 samples, and the resultant waveform is also slew-limited (i.e., most of the sampling occurs on the ramps of the waveform). In addition, EPI has the disadvantage that the majority of the workload is placed on a single gradient axis, unlike spiral trajectories. In spiral trajectories, heating is distributed between two gradient axes, which is a significant advantage in DWI. Finally, slew-limited trajectories cluster samples in regions of k -space where the trajectory velocity is the slowest. In spirals, this occurs at the center of k -space where the most signal is, whereas slew-limited EPI trajectories cluster samples at the outer portions of k -space. As such, spiral navigators should also have an SNR advantage

over EPI navigators.

5.3 SSFP-DWI of the Brain

5.3.1 Implementation

Brain images were acquired on healthy volunteers with the navigated steady-state diffusion-weighted imaging (SSFP-DWI) pulse sequences shown in Fig. 5.4a and 5.4c. Experiments were performed on a 1.5 T GE Signa CV/i research scanner (General Electric Co., Milwaukee, WI) with 40 mT/m gradients switch-able at 150 mT/m/ms and a transmit/receive quadrature birdcage head coil. Diffusion gradients could be placed along any direction, and in the data presented here were placed along the gradient axes (corresponding in the brain to right/left, anterior/posterior and superior/inferior). To minimize large motions, padding was placed at the subject's temples. The refocusing reconstruction described in Chapter 6 was used for navigator correction.

5.3.2 2DFT Images

For the spiral-navigated 2DFT experiments, readout consisted of a spiral navigator (FOV=24 cm, 16x16 matrix, $15 \times 15 \text{ mm}^2$ in-plane, receive BW 62.5 kHz) followed by a 2DFT high-resolution acquisition (FOV=24 cm, 192x192 matrix, $1.25 \times 1.25 \text{ mm}^2$ in-plane, receive BW 7.8125 kHz). Diffusion gradients had amplitude $G=40 \text{ mT/m}$ and duration $\tau=6.5 \text{ ms}$ ($b_{eff}=980 \text{ s/mm}^2$). Other imaging parameters were: $TR=40 \text{ ms}$, $\alpha=30^\circ$, thick=6 mm, $n_{ex}=48$, for a total scan time of 6:15 per image.

Typical images with this sequence are shown in Fig. 5.6. The R/L- and A/P-weighted images are relatively free of artifacts and exhibit strong diffusion contrast. In particular, a number of fairly small white matter structures are visible, including the external capsule and cortical white matter tracts. The S/I-weighted image has




Figure 5.6: Ungated, navigated 2DFT SSFP-DWI images. Labels indicate the direction of diffusion weighting in each image.

lower SNR than the other 2 images, which can be seen by noting that the signal from the scalp is roughly the same in all 3 images, but that the S/I-weighted signal is considerably lower in the brain. In addition, there is a clear loss of signal in medial areas containing the thalamus. This signal loss is due to residual motion artifacts. These artifacts are related to arterial pulsatility during the systolic portion of the cardiac cycle. An approach to cardiac gating with SSFP-DWI is briefly presented below and discussed further in Chapter 6.

5.3.3 Self-navigated Spiral Images

For the self-navigated spiral experiments, brain images were acquired using a bidensity spiral with FOV=24 cm, covered in a 16x16 matrix at inner k-space in each interleave and in a 256x256 matrix at outer k-space over 16 interleaves (in-plane resolution $0.9375 \times 0.9375 \text{ mm}^2$). Diffusion gradients had amplitude $G=40 \text{ mT/m}$ and duration $\tau=5.5 \text{ ms}$ ($b_{eff}=835 \text{ s/mm}^2$). Other scan parameters were: $TR=40 \text{ ms}$, $\alpha=25^\circ$, thick=6mm, $n_{ex}=240$ and total scan time of 2:40.

Figure 5.7 shows a set of SSFP-DWI images acquired with this self-navigating




Figure 5.7: Ungated, self-navigated spiral SSFP-DWI images.

trajectory. These images show good diffusion contrast but have considerably lower SNR than the 2DFT images shown in Fig. 5.6. The SNR difference is likely a combined effect of the smaller voxel volume and the decrease in total acquisition time, which together would be expected to reduce the SNR by nearly a factor of 3.

While this trajectory is potentially more efficient than the 2DFT trajectory, it has not been explored further for two main reasons. First, there is not a straightforward extension to 3D imaging, which limits its use with SSFP-DWI. Second, gradient heating severely reduces the SNR benefits by requiring significant increase in the sequence T_R .

5.3.4 Cardiac-gated Images

The images shown in Fig. 5.6 can be improved with cardiac gating. During systole, the brain deforms about the ventricles and displaces inferiorly through the foramen magnum [64,65,95], such that medial regions tend to experience large displacements. In the presence of diffusion-weighting gradients, these large motions can introduce intra-voxel dephasing, which cannot be removed with post-processing.

To reduce artifacts in DWI it is often necessary to use cardiac gating information

to synchronize acquisition to the cardiac cycle. Since most navigator corrections can only remove the effects of rigid-body motion, data is only acquired in the small portion of the cardiac cycle during which motions are reasonably approximated as rigid-body. This introduces significant increases in scan time. The refocusing reconstruction developed in Chapter 6 removes non-rigid motions as well, enabling acquisition during a larger portion of the cardiac cycle. With this correction, it is only necessary to avoid peak systole when intra-voxel dephasing occurs. The relatively short T_{RS} used in SSFP-DWI along with the ability to correct for most non-rigid motions allow this sequence to use cardiac gating without major increases in scan time.

To synchronize SSFP-DWI to the cardiac cycle, a plethysmograph was used to detect the onset of systole. Unlike most DWI methods, SSFP-DWI sequences cannot simply stop imaging following a systolic trigger because this would disrupt the steady state of the magnetization. Instead, the steady state is maintained by acquiring dummy cycles in which no data is acquired but the RF pulse and gradients continued to run as normal. Data acquisition continues following a pre-set delay of 80–120 ms (the portion of systole in which intra-voxel dephasing was found to occur). This scheme allows us to avoid intra-voxel dephasing with a minimal increase in scan time. As discussed in Chapter 6, careful ordering of phase encode acquisitions relative to the cardiac cycle was also found to be important.

The images shown in Fig. 5.8 were acquired with the same 2DFT sequence discussed above with the addition of cardiac gating. These images contain little visible artifact, even in the S/I-weighted image, and have increased conspicuity of small white matter structures compared with Fig. 5.6. Additional cardiac-gated images are shown in Figs. 6.7 and 6.8.

5.4 SSFP-DWI of Cartilage

In addition to the neurological applications discussed above, DWI is a potentially useful diagnostic for musculoskeletal pathology, and has shown particular promise for




Figure 5.8: Cardiac-synchronized, spiral-navigated 2DFT images acquired with the phase encode interleaving scheme shown in Fig. 6.4 and discussed in Chapter 6. Diffusion weighting is along the (a) A/P, (b) R/L and (c) S/I directions.

early detection of osteoarthritis (see the literature overview in Appendix A). DWI studies of excised cartilage samples suggest that the diffusion coefficient of cartilage may be an early indicator of breakdown of the cartilage matrix [96,97]. However, standard DWI methods (e.g., SE-DWI) have not been applicable *in vivo* due to the short T_2 of cartilage (20–40 ms [98]). As discussed above, SSFP-DWI is able to achieve diffusion weighting without incurring extensive signal loss due to T_2 relaxation, making it a strong candidate for DWI of cartilage and other short- T_2 species.

Stimulated echo (STE-DWI) sequences have similar potential to image short- T_2 species (which can be easily seen by noting that the SSFP-DWI signal is simply a weighted summation of stimulated echoes). However, STE-DWI is less SNR efficient and does not naturally extend to 3D acquisitions, limiting the achievable resolution and hence the applicability of STE-DWI to cartilage.

5.4.1 Navigated 3D SSFP-DWI for Cartilage

This work uses the 2D-navigated 3DFT SSFP-DWI sequence diagrammed in Fig. 5.4b. In addition to being favorable for short- T_2 species like cartilage, SSFP-DWI has the potential to acquire high-resolution 3D data. Articular knee cartilage is only a few millimeters thick, with considerable structural changes occurring over this short distance. For this reason, sub-millimeter voxel sizes are crucial for meaningful diagnosis. To achieve high-resolution 3D SSFP-DWI, it is necessary to perform navigator correction. While motion in the knee is likely to be less pervasive than in the brain, motion artifacts can nevertheless compromise image quality, as can be seen in the uncorrected image in Fig. 5.10.

As discussed above, removal of all possible phase errors in a 3D data set would require acquisition of a 3D navigator every T_R . However, considering the specific problem of knee imaging, motion may be restricted if the subject is largely immobilized during scanning. The primary motion during scanning is likely to be in the axial plane, where the knee might rotate or shift. These motions will introduce low-order phase in the axial plane [66]. The knee might also shift in the superior-inferior (S/I) direction, but rotations or other complex motions in this direction are less likely. Depending on the direction of diffusion encoding, a shift in the S/I direction may cause a constant phase offset, but no higher-order terms [70]. We use a 2D navigator in the axial plane (i.e., at $k_z=0$), which measures varying phase along the R/L and A/P directions, and the mean phase along S/I. We found a 2-turn axial spiral navigator to be sufficient to correct motion-induced phase errors.

For comparison with the existing *in vitro* literature, we refer to the amount of diffusion weighting in SSFP-DWI with the b -value which would give equivalent attenuation in cartilage ($D \approx 0.00145 \text{ mm}^2/\text{s}$ [61, 98]). The short T_2 of cartilage allows the application of the quantification method proposed in Eq. 5.1. In this study, we assume $T_1=800 \text{ ms}$ [61, 99] and fit D to Eq. 5.1 from several measurements with varying diffusion gradient area.

5.4.2 Cartilage Imaging Results

In vivo images of the patellar-femoral joint were acquired in 3 healthy subjects on a 1.5 T Signa LX scanner (40 mT/m maximum gradient amplitude, 150 T/m/s maximum gradient slew rate) using a 7.5 cm surface coil centered on the patella. Axial images were gathered through the articular cartilage covering a $14 \times 14 \times 4.8$ cm³ FOV in a $256 \times 192 \times 16$ matrix for $0.5 \times 0.7 \times 3.0$ mm³ voxel resolution. Spiral navigators were acquired in the $k_z=0$ plane over a 14×14 cm² FOV in a 4×4 matrix. The SSFP-DWI sequence utilized $TR=30$ ms and $\alpha=25^\circ$ to satisfy the constraints on Eq. 5.1. Diffusion gradients of strength $G = 7.8, 15, 27.5$ and 40 mT/m were applied for $\tau = 5.5$ ms, causing signal attenuation in cartilage equivalent to $b = 30, 115, 350, 630$ s/mm². To avoid large bulk motions (which can necessitate image registration), the subject was stabilized using a knee-high plastic brace loosely mounted on the patient bed. Three averages were acquired for a total scan time for each b -value of 4:40. The data was fit to Eq. 5.1 using the same nonlinear search algorithm described in the simulations above. The excitation pulse was carefully calibrated to improve the quality of the fit.

Figure 5.9 shows 3D SSFP-DWI knee images with varying levels of diffusion weighting obtained on a healthy volunteer. These images have been navigator-corrected to remove motion artifacts. The cartilage has high signal at low levels of diffusion weighting (SNR \approx 110 in the $b=30$ s/mm² image), and is attenuated at increasing b -values due to the relatively high diffusion coefficient of cartilage. Navigator correction was crucial in generating usable images, particularly at higher levels of diffusion weighting. Figure 5.10 shows a typical $b=630$ s/mm² image before and after navigator correction. The navigator correction removes the ghosting and blurring artifacts found in the uncorrected image.

Figure 5.11 shows diffusion coefficient (D) fits for 3 contiguous slices from a 3D SSFP-DWI data set. The D -maps have high SNR in cartilage, indicating a well-conditioned fit to the data. The average diffusion coefficient over an ROI of pure cartilage (433 voxels identified from the $b=115$ s/mm² anatomical image) is 0.0016 mm²/s, slightly higher than *in vitro* literature values of about 0.00145 mm²/s [61, 98]. This




Figure 5.9: 3D diffusion-weighted knee images acquired using SSFP-DWI with 3 levels of diffusion weighting. The cartilage signal is attenuated with increasing b .

slight overestimate may be due to a bias in the fitting procedure itself, or may simply reflect the existence of synovial fluid in the ROI.

Previous work has reported characteristic spatial heterogeneity in cartilage, with low D and T_2 at the bone-cartilage interface which increase across the cartilage to peak at the articular surface [98,100]. Figure 5.12 shows the fitted diffusion coefficient (solid lines) and the T_2 -weighted signal from the $b=115$ s/mm² images (dashed lines) across the cartilage (both femoral and patellar). These results are in good qualitative agreement with previous findings [98,100].

The current implementation of 3D SSFP-DWI requires approximately 5 minutes of scan time per level of diffusion weighting. To be clinically useful, this time should be reduced so that multiple b -values can be acquired in under 10 minutes. For ease of implementation, our current version of this pulse sequence uses a 3DFT imaging trajectory. More efficient 3D trajectories should significantly reduce scan time.




Figure 5.10: 3D diffusion-weighted knee images before and after navigator correction. Arrows indicate motion artifacts removed by the correction. The navigator correction removes ghosting and blurring artifacts from the cartilage. These images are windowed harshly to accentuate artifacts.

5.5 Summary

Diffusion imaging with SSFP has long been known to be capable of strong diffusion weighting without requiring the large gradients or long echo times of more conventional DWI sequences (such as SE-DWI). However, its sensitivity to motion has traditionally precluded the use of SSFP-DWI. This work has presented a method for navigated SSFP-DWI that is able to overcome these issues by enabling the correction of motion-induced phase artifacts. This method can be combined with cardiac gating to get high-resolution images with strong diffusion contrast and essentially no artifact due to motion. Because diffusion contrast does not require long echo times, SSFP-DWI is also useful for imaging short- T_2 species or long- T_2 species with very strong diffusion contrast. In addition, the use of $T_R \geq T_2$ allows quantification of the diffusion coefficient via a fairly simple approximation of the signal attenuation.




Figure 5.11: Diffusion maps for 3 representative slices in the 3D SSFP-DWI acquisition (bottom row) along with the $b=115$ s/mm² images for anatomical references (top row). The values fit for the diffusion coefficient (D) are in the range of literature values (0.00145 mm²/s).




Figure 5.12: Average cross-sections for 5 adjacent voxel columns through the patellar and femoral cartilage from the images shown in Fig. 5.11. Shown are the fitted diffusion coefficient D (solid lines, in mm²/s) and the T_2 -weighted signal from the $b=115$ s/mm² images (dashed lines, in arbitrary units). The outer pixels in each plot correspond to the bone-cartilage interface, and the dips in the center occur at the cartilage surface. Both T_2 and D are low at the bone interface and high at the cartilage surface, as has been previously reported *in vitro*.

Chapter 6

Non-linear Phase Correction of Diffusion MRI

6.1 Motion in DWI

A key problem in diffusion-weighted imaging is non-diffusive bulk motion, which causes phase shifts in the voxel signal [66]. When data is gathered in a multi-shot acquisition, phase shifts during different excitations cause destructive interference to create serious image artifacts. A number of methods have been proposed to deal with motion artifacts in DWI. One simple way to avoid motion-induced artifacts is to gather the image data in a single acquisition, in which case the image magnitude is uncorrupted. However, these single-shot acquisitions have limited resolution and are sensitive to susceptibility and eddy-current effects, as described in Chapter 3. Another approach is to correct motion-induced phase artifacts before combining the data from multiple acquisitions. Phase correction imposes no intrinsic limits on resolution but does require the acquisition of additional navigator data that resolves the phase errors in image space. Early work acquired 1D projection navigators [66, 70, 71] or pairs of orthogonal projections [101], and subsequent studies have extended this approach to a variety of pulse sequences and trajectories [69, 102–107]. An important disadvantage

of 1D navigation is the inability to correct for phase shifts orthogonal to the direction of navigation. Recent studies have addressed this problem with 2D navigation [57, 72, 73].

Navigated DWI methods typically assume rigid-body motion, the effects of which have been well characterized [66, 67]. However, a significant component of brain motion is non-rigid [64, 65, 95]. During systole, the brain deforms about the ventricles and displaces inferiorly through the foramen magnum as though being pulled by the spinal cord. The more superior and lateral regions of the brain are nearly motionless, whereas the more inferior and medial regions tend to experience large displacements. Even during the relative quiescence of diastole, motion induced by the cardiac cycle follows this nonlinear spatial pattern as the brain slowly returns to its resting position.

In the presence of diffusion-weighting gradients, these deformations cause phase corruptions that vary nonlinearly across the object magnetization. The effect of brain motion on the DWI signal is shown in Fig. 6.1. The signal phase and amplitude are correlated with the cardiac cycle and exhibit nonlinear spatial variation. The loss in signal amplitude during peak systole is due to intravoxel dephasing caused by intense motion. The highly focal nature of this dephasing highlights the spatial nonlinearity of phase corruptions. Note that the dephasing occurs primarily in the medial regions, which experience the most displacement, as discussed above.

These phase errors are problematic for multi-shot acquisitions because each readout experiences a different image-space phase corruption ($\phi_j(\mathbf{r})$ for the j^{th} readout). This causes phase interference when readouts are combined to form a high-resolution image. Navigated techniques correct phase errors with low-resolution navigator data that is acquired along with each high-resolution data frame [70]. Based on the assumption of rigid-body motion, standard navigator methods correct for 0^{th} and 1^{st} order spatial-phase terms [66]. In the presence of the non-rigid motions described above, a rigid-body correction is a first-order approximation to the desired high-order correction. Rigid-body corrections leave residual errors that are strongly correlated

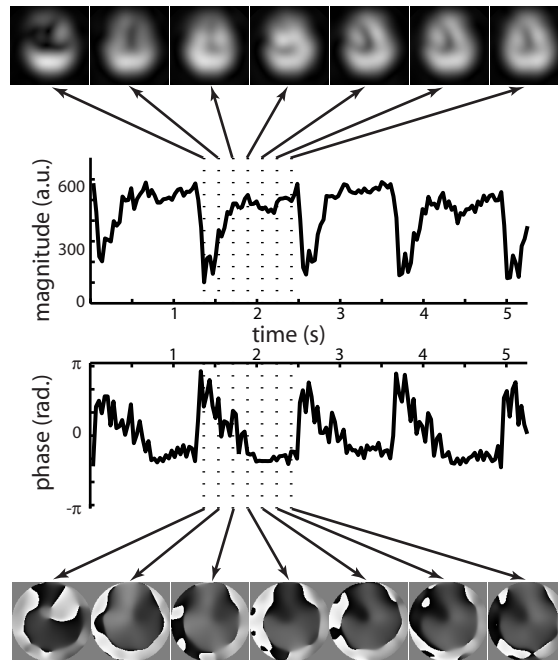


Figure 6.1: Illustration of the effects of motion on diffusion-weighted magnetization. The magnitude (top) and phase (bottom) of a series of low-resolution axial single-shot DWI images (2D navigators) are shown over a single cardiac cycle. The magnitude and phase of a representative voxel are plotted in the middle over four cardiac cycles. The spatial variation in the image magnitude and phase is nonlinear and strongly correlated to the cardiac cycle.

to the cardiac cycle [107]. Given that a 2D navigator contains high-order phase information, image reconstruction should be improved by correcting for high-order phase terms.

This work presents a generalized reconstruction method that corrects for non-rigid motion based on least-squares estimation. The full least-squares estimate requires prohibitive computation, but suggests an efficient approximation called the “refocusing reconstruction”. We consider the conditions for this approximation, and propose efficient algorithms for the refocusing reconstruction. We also present a time-efficient approach to cardiac gating that improves the performance of the refocusing reconstruction. These methods are applied to SSFP-DWI [90], a potentially powerful DWI method that is particularly sensitive to motion-induced phase corruptions.

6.2 Navigator Correction of Non-linear Phase Artifacts

In this section, we examine the effect of nonlinear phase corruption on diffusion-weighted data and derive a nonlinear phase correction based on least-squares estimation. Computational issues require that the least-squares reconstruction be approximated, and an approximation called the “refocusing reconstruction” is presented. This reconstruction can be performed in either image-space or k -space.

6.2.1 Least-squares Reconstruction

Phase corruption can be expressed at each point in image space as a multiplication of the magnetization by a phase function and can be described conveniently using matrix formalism. For an arbitrary trajectory that covers k -space in C interleaves of length R , corresponding to an $N \times N$ image, the acquired data is:

$$\mathbf{d} = \mathbf{G}\mathbf{F}\mathbf{m} = \mathbf{R}\mathbf{P}\mathbf{m} \quad (6.1)$$

where \mathbf{d} is a $CR \times 1$ vector containing the k -space data, \mathbf{m} is a $N^2 \times 1$ vector containing the object magnetization in Cartesian coordinates, and \mathbf{P} ($CN^2 \times N^2$), \mathbf{F} ($CN^2 \times CN^2$) and \mathbf{G} ($CR \times CN^2$) are matrices respectively representing image-space phase corruption, discrete Fourier transform, and resampling from a Cartesian grid onto the k -space trajectory. The phase corruption matrix \mathbf{P} is a stack of diagonal matrices containing the phase corruption for each readout. This phase corruption is measured by the navigator phase, which must be interpolated to match the high-resolution readout (since the $n \times n$ navigator is of lower resolution than the high-resolution readout). The matrix $\mathbf{R} = \mathbf{G}\mathbf{F}$ describes the basis set used in the k -space trajectory. Expressions for these vectors and matrices are given in Appendix B.

The magnetization can be estimated from the data in Eq. 6.1 using least-squares. Assuming an accurate measurement of the phase errors \mathbf{P} , the least-squares reconstruction is:

$$\begin{aligned} \hat{\mathbf{m}} &= (\mathbf{P}^*\mathbf{R}^*\mathbf{R}\mathbf{P})^{-1}\mathbf{P}^*\mathbf{R}^*\mathbf{d} \\ &= \mathbf{M}^{-1}\mathbf{P}^*\mathbf{R}^*\mathbf{d} \end{aligned} \quad (6.2)$$

where $*$ denotes conjugate transpose. It is helpful to consider this reconstruction as occurring in three steps: (1) reconstruction \mathbf{R}^* of the individual data frames in image-space, which we call the “interleave reconstruction”; (2) phase correction \mathbf{P}^* , which we call the “refocusing operator”; and (3) removal of remaining artifacts by $\mathbf{M}^{-1} = (\mathbf{P}^*\mathbf{R}^*\mathbf{R}\mathbf{P})^{-1}$, which we call the “unmixing operator”. These steps are diagrammed in Fig. 6.2 using the notation described in this section and in Appendix B.

The data is first partially reconstructed using the interleave reconstruction $\mathbf{R}^* =$

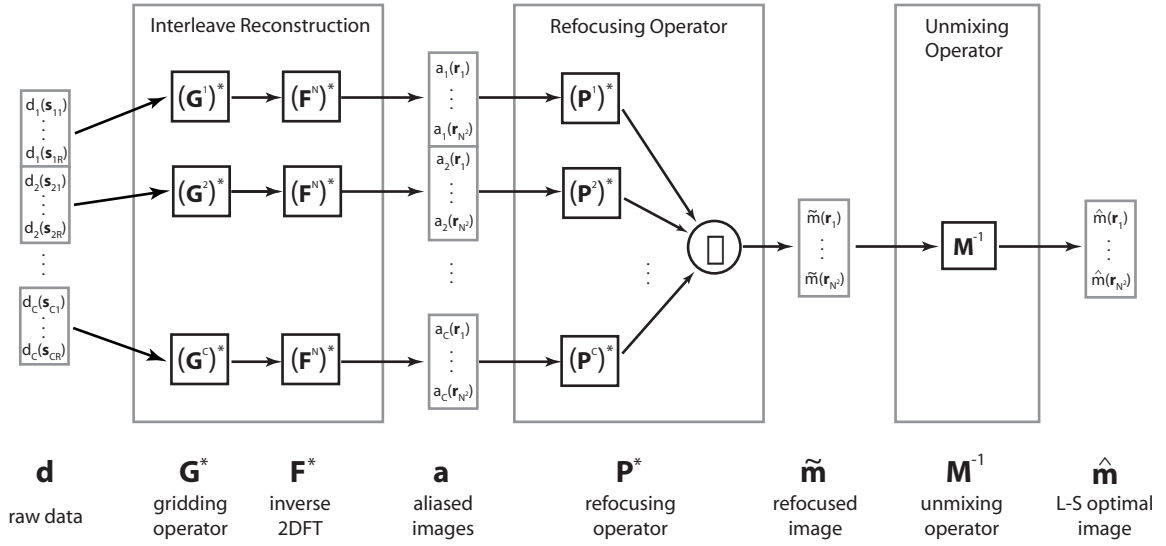


Figure 6.2: Diagram of the least-squares matrix reconstruction given in Eq. 6.2. Each raw data interleaves \mathbf{d}_j is gridded and inverse Fourier-transformed by the interleaves reconstruction to form an aliased image \mathbf{a}_j . These images are phase-corrected and summed by the refocusing operator to form the refocused image $\tilde{\mathbf{m}}$. Remaining aliased energy can be removed by the unmixing operator \mathbf{M}^{-1} to give the least-squares optimal image $\hat{\mathbf{m}}$. If the aliased energy in the refocused image is sufficiently low, the reconstruction can stop at $\tilde{\mathbf{m}}$, avoiding the computationally expensive inverse calculation \mathbf{M}^{-1} . In this case, we call the algorithm a “refocusing reconstruction”, which is the reconstruction used on all images in this work.

$\mathbf{F}^*\mathbf{G}^*$, which separately re-samples each readout onto a Cartesian grid at the high resolution readout size and inverse Fourier transforms the re-sampled data without combining the interleaves. The resulting vector $\mathbf{a} = \mathbf{R}^*\mathbf{d}$ contains the image-space data corresponding to each readout of k -space data. Unless each readout covers a full field-of-view (FOV), the images in the vector \mathbf{a} are aliased.

In the next step, the refocusing operator \mathbf{P}^* is applied to the interleave images to yield the refocused image $\tilde{\mathbf{m}} = \mathbf{P}^*\mathbf{a}$. The refocusing operator performs the essential task of rephasing the interleaves based on the phase corruption measured by the navigator. Each interleave image is multiplied by the phase conjugate of the navigator image to rephase the unaliased component. The sum of these rephased interleaves is the refocused image. Provided the navigator associated with each interleave is a stable measure of the object phase, the unaliased components will add coherently in this summation. However, the aliased components will not necessarily cancel since the orthogonality of the Fourier encoding has been disturbed. The refocused image is thus the desired image plus some aliased artifact. To the extent that the phase errors are random, this aliased energy will be incoherent and will tend to cancel. The characterization and minimization of the aliased signal is discussed later.

The final step of the least-squares reconstruction is the application of the unmixing operator \mathbf{M}^{-1} to the refocused image. The unmixing operator removes two artifacts from the refocused image: non-uniformities in the sampling trajectory and motion-induced aliasing (see Appendix B). The compensation for trajectory non-uniformities, a standard correction in the least-squares reconstruction of non-Cartesian data, will be neglected here. More important to our discussion is the removal of the aliased energy discussed above. The unmixing operator uses a linear combination of voxel signals to remove this coupling between voxels. If the refocused image contains little aliasing, the mixing matrix is close to the identity matrix.

6.2.2 Refocusing Reconstruction

Calculation of the unmixing operator \mathbf{M}^{-1} is problematic due to its computational complexity. Not only is the $N^2 \times N^2$ inversion computationally expensive, it also cannot be pre-computed because it depends on the phase corruptions. Further, the calculation of the mixing matrix \mathbf{M} itself requires $CRN^2 \approx N^4$ operations for each element of the $N^2 \times N^2$ matrix (see Eq. B.7). These computational requirements make direct calculation of the unmixing matrix (and hence the full least-squares reconstruction) intractable.

If the sources of coupling are sufficiently low, the mixing matrix has little off-diagonal energy and is close to the identity matrix. Approximating the mixing matrix as the identity (i.e., leaving out the unmixing step) yields the “refocusing reconstruction”:

$$\tilde{\mathbf{m}} = \mathbf{P}^* \mathbf{R}^* \mathbf{d} \quad (6.3)$$

Provided the unmixing matrix is nearly diagonal, the refocusing reconstruction is a good approximation to the least-squares estimate which is much faster to compute. The refocusing reconstruction is also conceptually elegant: before the readouts are combined, each readout is refocused in image space by multiplication with the navigator phase conjugate. This interpretation is not only computationally useful, but also suggests a more efficient algorithm than the matrix computation in Eq. 6.3 for refocusing the data. This algorithm will be discussed later.

Some important properties of the refocusing reconstruction can be understood by considering the problem in k -space. Corruption by a spatial phase function ($e^{i\phi(\mathbf{r})}$) is equivalent to convolution by the transform of that phase function ($\mathcal{F}\{e^{i\phi(\mathbf{r})}\}$) in the spectral domain. If all of k -space were convolved with the same phase corruption kernel, we could restore the original spectrum by deconvolving with this kernel.

However, in motion-corrupted DWI data, each readout samples a spectrum that has been convolved with a different kernel ($\mathcal{F}\{e^{i\phi_j(\mathbf{r})}\}$ for the j^{th} readout).

Interpreted in k -space, the refocusing operator deconvolves each interleave separately using the phase-corruption kernel from that excitation. Refocusing uses the best-estimate deconvolution kernel for points from the same excitation. For data from different excitations, the deconvolution will work well if the kernels are similar for adjacent positions in k -space, but will be unable to remove artifacts if the kernels are dissimilar. Hence, it is desirable for the kernel to change slowly and smoothly over k -space so that the deconvolution is locally consistent. An additional concern is periodicity in the phase corruptions that cause a periodic weighting of the data in k -space. Any residual periodicity left by the refocusing operator will cause coherent ghosts, which are a more disruptive artifact than the diffuse aliasing associated with discontinuities. The refocusing reconstruction will produce high fidelity images when phase corruptions are smooth and non-periodic in k -space.

6.3 Implementing the Refocusing Reconstruction

6.3.1 Efficient Reconstruction

Although the matrix formulation given in the previous section is a useful way to derive the refocusing algorithm, it is an inefficient method for calculating the refocused image. Directly implementing the matrix formulation has prohibitive memory requirements and is unlikely to take full advantage of the highly structured nature of these matrices. It is much more efficient to use operations equivalent to those given in the matrix formulation, since these operations often have highly-optimized algorithms (such as the Fast Fourier Transform). The implementation of image-space refocusing shown in Fig. 6.3a can be calculated more efficiently than the matrix calculation of $\tilde{\mathbf{m}}$ shown in Fig. 6.2. Refocusing can also be implemented as a k -space deconvolution, as shown in Fig. 6.3b.

(a) image-space refocusing reconstruction




 (b) k -space refocusing reconstruction


Figure 6.3: Block diagram of two versions of the refocusing reconstruction in which refocusing is performed in (a) image space and (b) k -space. The boxed expressions indicate the operations of gridding (\mathcal{G}_j grids the j^{th} readout trajectory), inverse Fourier transform (\mathcal{F}^{-1}) and refocusing (multiplication by $\mathcal{P}_j^*(\mathbf{r}) = e^{-i\phi_j(\mathbf{r})}$ in image-space or convolution by $\mathcal{F}\{\mathcal{P}_j^*\}$ in k -space).

The image-space refocusing algorithm in Fig. 6.3a involves the following steps:

1. Initialize a zero-filled $N \times N$ matrix (where we will accumulate the refocused image).
2. For each interleave:
 - (a) Reconstruct the navigator data in k -space ($n \times n$).
 - (b) Zero-pad (a) to the final image resolution ($N \times N$).
 - (c) Inverse Fourier transform (b) (to get the navigator image).
 - (d) Reconstruct the high-resolution data frame in k -space at the final image resolution ($N \times N$).
 - (e) Inverse Fourier transform (d) (to get the interleave image).
 - (f) Multiply (e) by the phase conjugate of (c) (refocusing step).
 - (g) Add (f) to the sum accumulating in (1).
3. The refocused image is the cumulative sum from (g).

For every interleave, this reconstruction requires two Fourier transforms and a matrix multiplication at the high-resolution matrix size ($N \times N$). Each Fourier transform requires time $N^2 \log N^2$ and the matrix multiplication requires time N^3 for a total complexity of $O(N^2 \log N^2 + N^3) \approx O(N^3)$. If the reconstruction of the navigator in k -space requires a gridding operation, this will incur an additional cost of $O(r^2 n^2)$ for a kernel of width r . This additional cost will in general be much smaller than the terms discussed above since $n < N$ and r is also small.

The k -space refocusing algorithm (Fig. 6.3b) is:

1. Initialize a zero-filled $N \times N$ matrix (where we will accumulate the refocused spectrum).
2. For each interleave:

- (a) Reconstruct the navigator data in k -space ($n \times n$).
 - (b) Zero-pad (a) by a factor of 2 ($2n \times 2n$).
 - (c) Inverse Fourier transform (b).
 - (d) Calculate the phase conjugate of (c).
 - (e) Fourier transform (d) (to get the refocusing kernel).
 - (f) Reconstruct the high-resolution data frame in k -space at the final image resolution ($N \times N$).
 - (g) Convolve (f) with (e) (refocusing step).
 - (h) Add (g) to the sum accumulating in (1).
3. Inverse Fourier transform the cumulative sum from (h) to get the refocused image.

The refocusing kernel must be calculated at an extended bandwidth (step 2b) because normalizing in image-space (step 2d) extends the kernel in k -space. Calculating the kernel at twice the navigator bandwidth avoids Gibbs ringing in the refocused image due to this extension. For every interleave, the k -space algorithm calculates two $2n \times 2n$ Fourier transforms and one convolution of an $N \times N$ matrix with a $2n \times 2n$ kernel. This algorithm runs in approximately $O(N^2n^2 + n^2 \log n^2) \approx O(N^2n^2)$ time (where constant multipliers have been dropped, as is standard in O-notation).

In most imaging situations, the k -space algorithm will perform better than the image-space algorithm since we will usually have $n^2 < N$. However, this is not always the case, as will be discussed later.

6.3.2 Cardiac Synchronization

If minor subject restraints are used to reduce bulk motion, the primary source of motion artifacts in the brain is deformation related to the cardiac cycle (see Fig. 6.1). Cardiac triggering information can be used to synchronize interleave ordering

to the cardiac cycle to reduce discontinuities and periodicities in the k -space data [108,109]. For example, in a spinwarp trajectory phase encodes are normally gathered incrementally, resulting in a periodic weighting over k -space with large discontinuities (see Fig. 6.4a). These effects can be lessened by choosing a phase encode increment that traverses the extent of k -space during a single cardiac cycle. By interleaving the phase encodes over multiple cardiac cycles in this manner, adjacent lines in k -space are collected at the same portion of the cardiac cycle and the resultant k -space weighting is smooth and non-periodic (see Fig. 6.4b). This method for cardiac synchronization of phase encoding should not require large increases in scan time since data can be acquired throughout a large portion of the cardiac cycle. Cardiac synchronization in this manner reduces the types of k -space modulations that are most problematic for the refocusing reconstruction, discontinuity and periodicity, without incurring the large increases in scan time that usually accompany cardiac gating.

6.4 Experiments and Results

6.4.1 Experiments

Brain images were acquired on healthy volunteers with the navigated SSFP-DWI pulse sequence shown in Fig. 6.5. To minimize large motions, padding was placed at the subject's temples. Readout consisted of a spiral navigator ($FOV=24$ cm, 16×16 matrix, 15×15 mm² in-plane, receive BW 62.5 kHz) followed by a spinwarp high-resolution acquisition ($FOV=24$ cm, 192×192 matrix, 1.25×1.25 mm² in-plane, receive BW 7.8125 kHz). Experiments were performed on a 1.5 T GE Signa CV/i research scanner (General Electric Co., Milwaukee, WI) with 40 mT/m gradients switch-able at 150 mT/m/ms and a transmit/receive quadrature birdcage head coil. Diffusion-weighted images were acquired using diffusion gradients with $G=40$ mT/m and $\tau=6.5$ ms ($b_{eff}=980$ s/mm² [90]) which could be applied along any direction. Some studies also acquired an image with minimal diffusion weighting ($G=0.3$ mT/m and $\tau=6.5$ ms). Other imaging parameters were: $T_R=40$ ms, $\alpha=30^\circ$ and 6 mm slice

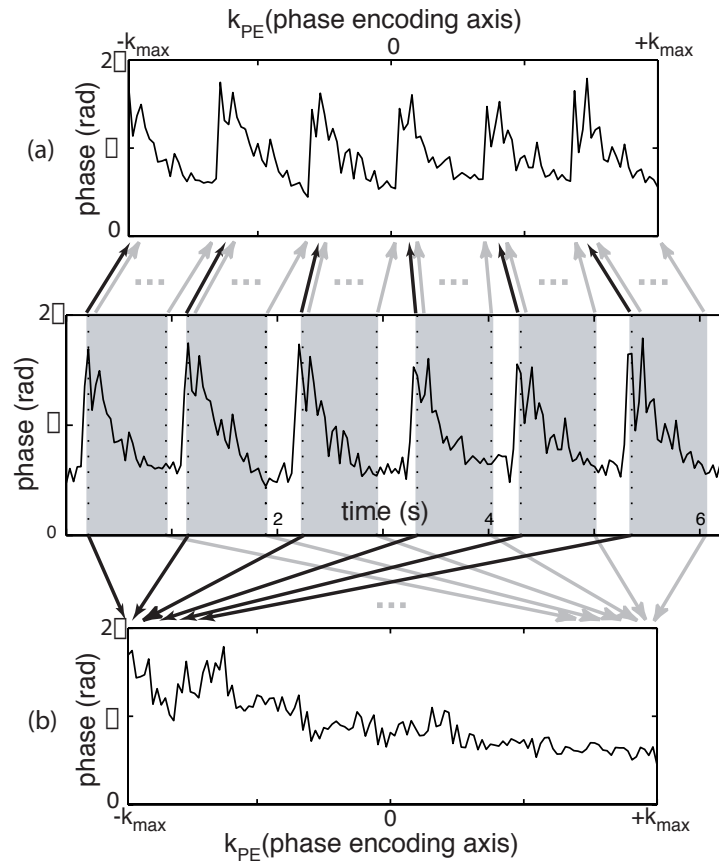


Figure 6.4: The effect of 2DFT phase encode ordering on k -space. The phase history of a representative voxel over time is shown in the middle plot. The gray swathes represent cardiac-gated imaging windows. The k -space modulation imparted by this phase history is shown for (a) standard incremental phase encode ordering and (b) an alternate phase encode order. The latter ordering interleaves the data from different cardiac cycles to achieve a more smooth modulation.




Figure 6.5: Navigated SSFP-DWI pulse sequence used in this study. The sequence is fully refocused except for the diffusion gradient (with amplitude G and duration τ), which rephases diffusion-weighted echoes over multiple T_{RS} .

thickness. The sequence was repeated 48 times for a total scan time of 6:15 per image. To date 10 subjects have been scanned with variants of this sequence, 5 of which were scanned with the prescription specified above.

In some experiments, acquisition was synchronized to the cardiac cycle using plethysmographic triggering. To maintain the steady-state, dummy cycles without acquisition were gathered while waiting for cardiac triggers. Following a systolic trigger, data was collected in an interleaved order such that the phase encodes acquired during a single cardiac cycle were evenly distributed across k_y -space. This process was repeated for successive cardiac cycles until all phase encodes had been acquired. Additionally, 2-3 dummy cycles (80-120 ms) were inserted following detection of a trigger to avoid the largest motions associated with systole. Cardiac synchronization usually increased scan times by 1-2 minutes.

Images were reconstructed using 3 methods: no correction, the standard linear correction, and the k -space refocusing correction proposed in this work. Reconstruction of the data was implemented off-line in C code using the FFTW library for Fourier transforms [110]. Both the image-space and k -space versions were implemented, although all refocused images shown here were calculated with the k -space implementation.

6.4.2 Refocusing Reconstruction

The different navigator corrections are compared on untriggered data in Fig. 6.6. Motion-induced phase artifacts interfere to cause signal attenuation in areas with large motions, and little or no artifact outside the object. Motion artifacts are manifested as signal loss rather than the more familiar ghosting artifacts due to the large number of excitations in the acquisition. All of the uncorrected images (top row) are corrupted, with the most serious artifacts in the S/I-weighted image. Arrows in the top row indicate artifacted areas in the R/L- and A/P-weighted images. The linear correction (middle row) restores signal in some corrupted regions (e.g., posterior areas in the S/I-weighted image), but the images remain attenuated in the medial areas where motion is more severe. Additionally, the linear correction introduces attenuation to the frontal lobes in the R/L- and A/P-weighted images due to a poor linear fit to the nonlinear phase artifacts (arrows in middle row of Fig. 6.6). The refocused images (bottom row) have restored signal and minimal artifacts. In particular, the medial regions show structures that were partially or completely suppressed in the uncorrected images (e.g., the external and internal capsule, indicated by arrows). The images weighted along the primary axis of motion (S/I) have lower SNR due to the large motions in the S/I direction. In peak systole, these motions create sufficient phase gradients to cause intravoxel dephasing which cannot be corrected using post-processing.

6.4.3 Cardiac Triggering

A set of refocused images acquired with cardiac synchronization are shown in Fig. 5.8. These images show strong diffusion contrast with clear delineation of small white matter structures. The S/I-weighted image does not suffer from the degraded SNR seen in untriggered images. The A/P- and R/L-weighted images are also improved with triggering, showing fine detail such as cortical white matter that is less apparent in the untriggered images.




Figure 6.6: Untriggered spiral-navigated 2DFT images with varying levels of correction. Rows (from top to bottom) have no correction, the standard linear correction, and the refocusing correction. Columns (from left to right) show images with diffusion weighting along the A/P, R/L and S/I directions. Arrows in the A/P- and R/L-weighted uncorrected images (top row) indicate examples of dephasing due to motion. Arrows in the middle row indicate artifacts introduced by the linear correction. Arrows in the bottom row indicate structures with increased conspicuity after the refocusing correction.

6.4.4 Diffusion Tensor Imaging

A full set of diffusion tensor images acquired with the standard directions for diffusion tensor imaging [111] is shown in Fig. 6.7. These images were acquired with cardiac-synchronized phase encode ordering and reconstructed using the refocusing reconstruction. Intravoxel dephasing is minor in these images due to the systolic delay (120 ms) and the fact that diffusion weighting was never applied along the principal axis of brain motion (S/I). These images show strong diffusion weighting with little visible artifact. Figure 6.8 shows the trace image and fractional anisotropy (FA) map for this data before and after correction. The correction improves both images, although the most dramatic improvement is found in the FA map.

6.5 Characterization of the Refocusing Reconstruction

We have presented a method for correction of nonlinear phase corruption. This method does not require additional navigator data beyond that used for 2D linear corrections. Rather, the refocusing correction represents a more complete use of the information in the navigator. The refocusing reconstruction algorithm is straightforward to implement: refocusing is simply a k -space deconvolution that is performed before individual readouts are combined. Alternatively, it can be thought of as image-space multiplication of individual data frames by the navigator phase conjugate.

6.5.1 Fidelity of the Refocusing Reconstruction

The quality of reconstruction by the refocusing method depends on the assumption that the mixing operator does not contribute significantly to the least-squares optimal reconstruction. Hence, the impulse response of the mixing operator is a useful way to characterize the fidelity of the refocused image. The off-peak power in the impulse




Figure 6.7: Navigated SSFP-DWI images acquired with the spiral-navigated 2D spin-warp trajectory. Images were diffusion weighted along the standard directions for diffusion tensor. Images are cardiac-synchronized and were reconstructed using the refocusing reconstruction described in Chapter 6.




Figure 6.8: Trace and fractional anisotropy maps for the data in Fig. 6.7 before and after the refocusing correction. Images before and after correction are identically windowed. The trace image shows a moderate improvement in contrast. The FA map shows a more marked improvement. The uncorrected FA map shows apparent anisotropy in gray matter regions due to motion-induced signal loss in a subset of the images. The refocusing correction restores this signal to give a more accurate map of anisotropy.

response indicates the amount of aliased energy in the refocused image and therefore the relative benefit of applying the unmixing operator. Of particular interest are coherences in the voxel impulse responses, which indicate coherent ghosts in the refocused image. The impulse response of the mixing operator is shown in Figure 6.9a for several representative voxels in an ungated data set. Although there is significant energy off-center, most of the aliased signal has random phase and will add incoherently. The impulse responses also show widening of the central peak and coherent off-center peaks, which will lead to blurring and ghosting, respectively. These effects in the impulse response need to be reduced for the refocusing operator to provide accurate reconstructions.

The quality of the refocusing reconstruction is dependent on the number of excitations that make up an image. Aliased energy in the refocused image decreases as the number of excitations increases since the incoherent sources of aliasing are sampled more often. As shown in Fig. 6.9b, averaging suppresses the incoherent off-peak energy found in unaveraged data. In this case averaging also suppresses the coherent peaks since data was acquired asynchronous to the cardiac cycle. Some averaging is always necessary to correct strong phase corruptions, which warp the k -space trajectory. Acquiring multiple averages ensures that the perturbed encoding functions will cover k -space more fully. One side effect of this need to average is the improvement of the refocusing reconstruction.

6.5.2 Cardiac Synchronization

The need for cardiac gating has been a common criticism of navigator methods. Since linear corrections assume rigid-body motion, cardiac gating is required to limit imaging to the more stable portions of the cardiac cycle in which motions are better approximated as linear. With the nonlinear corrections introduced here, post-systolic delays only need to be long enough to avoid intravoxel dephasing during peak systole which cannot be corrected in post-processing. For the 2DFT SSFP-DWI sequence

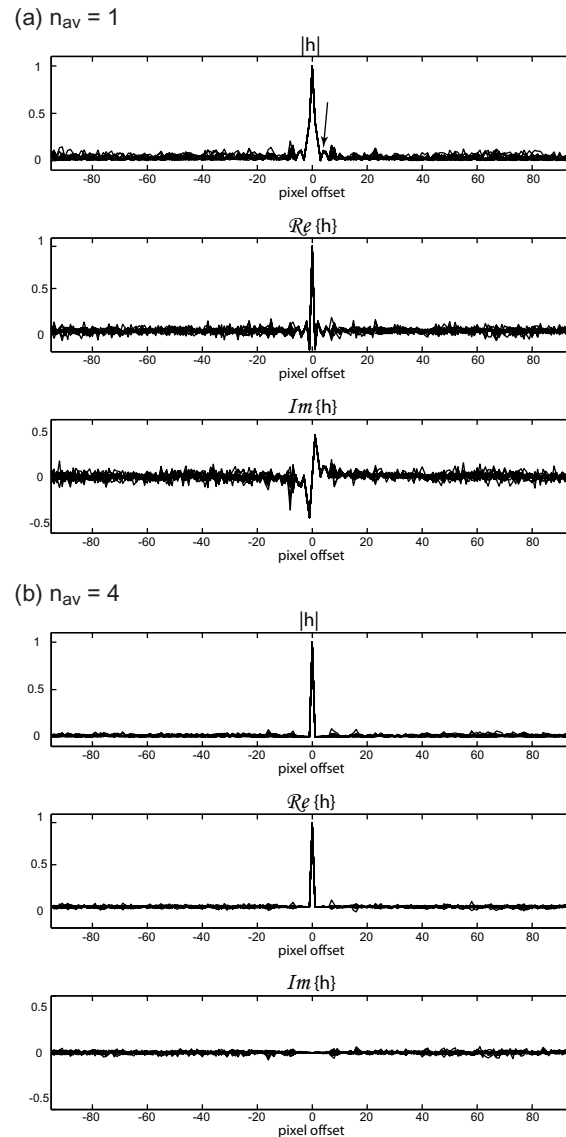


Figure 6.9: Impulse responses of the mixing operator for a 2DFT acquisition with (a) no averaging and (b) 4 averages. Each subfigure overlays the impulse responses for the same 10 voxels taken from a representative navigator data set. The arrow in (a) indicates coherences in the mixing impulse responses which lead to coherent ghosts. Averaging in (b) suppresses much of the off-peak energy. There is also a slight broadening of the main peak in the unaveraged impulse response (a) which would create a slight blurring of the image, and which is absent in the averaged response (b). Since both impulse responses are fairly uncorrupted, the unmixing step is unlikely to substantially improve image quality.

used here, cardiac synchronization requires only moderate increases in scan time (25-33%). This imaging efficiency is possible because cardiac information is being used to synchronize the acquisition to the cardiac cycle, rather than simply limiting the portion of the cardiac cycle in which data is gathered. Our current implementation approximates the heart rate for the entire scan. Monitoring changes in heart rate during scanning would further reduce the overhead time for cardiac synchronization.

Improvements to the refocused image due to cardiac synchronization can be seen by examining the impulse response of the mixing operator. The impulse response for the standard phase encode ordering scheme (Fig. 6.10a) has significant off-peak energy as well as a series of coherent peaks (arrows) due to the periodic weighting of k -space introduced by the cardiac cycle (several harmonics are clearly visible). Cardiac-synchronized ordering (Fig. 6.10b) eliminates off-center coherent peaks and produces a more well-behaved impulse response. The attenuation band around the impulse is due to the suppression of low frequencies in the k -space modulation due to the phase encode ordering.

6.5.3 Computational Efficiency

Although mathematically equivalent, the image-space and k -space refocusing reconstructions differ in terms of computational efficiency. The choice of algorithm depends on the relative sizes of the navigator and high-resolution acquisitions, as well as the details of the high-resolution trajectory. However, in almost all practical imaging scenarios, the k -space implementation is expected to outperform the image-space algorithm. As presented above, the k -space algorithm will perform better if $n^2 < N$ because the k -space convolution will be faster than the large matrix multiplication.

An additional consideration in the comparison of the two methods is the specific k -space trajectory used for the high-resolution acquisition. Our previous analysis of the k -space algorithm assumes that the convolution must be performed over the entire $N \times N$ matrix, which is not true for all trajectories. For some trajectories, the

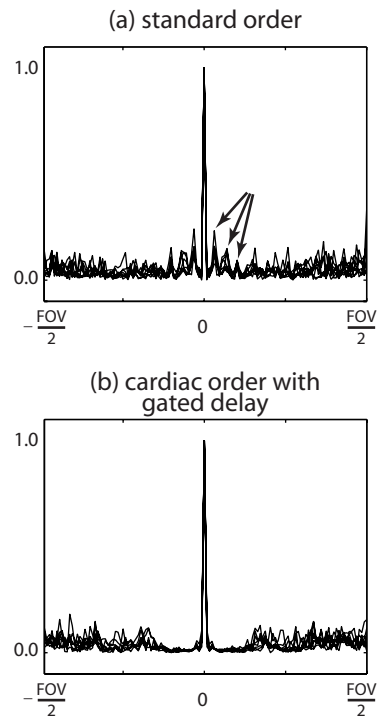


Figure 6.10: Simulated impulse response magnitude of the mixing operator for a 2DFT acquisition for with and without cardiac synchronization ($n_{av} = 1$). (a) The impulse responses for standard incremental order have strong coherences (arrows) that will lead to ghosting. (b) These coherences are eliminated when cardiac-synchronized phase-encode ordering is used.

convolution time can be significantly reduced by performing a local convolution. For example, with a 2DFT trajectory, the convolution need only include the area around each phase encode line, and the k -space reconstruction can therefore be performed in $O(Nn^2)$ time. Recalling that the image-space algorithm takes time $O(N^3)$, the k -space algorithm is expected perform better on 2DFT acquisitions when $n < N$, which will always be the case.

Thus far we have only considered 2D acquisitions. For 3D acquisitions, all 2D FFTs, matrix multiplications and convolutions are replaced with 3D operations. The k -space algorithm is expected to scale more favorably than the image-space algorithm for 3D acquisitions since because performing a local 3D convolution will be significantly faster than calculating large 3D Fourier transforms and matrix multiplications. Early work with 3D SSFP-DWI knee imaging has used a k -space algorithm with good results [?]. For this application, the image-space algorithm was found to be prohibitively slow.

For the experiments performed here (a spinwarp trajectory with $n = 16$, $N = 192$), the k -space algorithm would be expected to drastically outperform the image-space algorithm (roughly by a factor of $(N/n)^2=144$). Our implementation computed the k -space reconstruction in 2.6 s per average and the image-space reconstruction in 6.9 s per average. The modest improvement of the k -space algorithm is likely due to the use of highly-optimized FFT software and unoptimized convolution calculations. This mismatch between theory and implementation highlights the importance of code optimization and hidden time constants (e.g., constant multipliers which are usually ignored in complexity analysis).

6.5.4 Weighted Refocusing Reconstruction

Previous work on navigated DWI has usually included a weighting or thresholding of data frames that are suspected to contain uncorrectable phase artifacts. The method proposed here does not include such a correction, primarily because it has not proven

necessary under the current imaging circumstances. However, imaging methods that acquire fewer data frames are more impacted by artifacted data frames, and might benefit from a weighting scheme.

Examining Fig. 6.1, we see that during periods of the most intense motion (such as systole) navigator images exhibit signal loss due to intravoxel dephasing. Navigator magnitude can thus be used as an indicator of uncorrectable phase corruption and can be incorporated into the reconstruction. If we weight each image-space sample by its navigator magnitude prior to correction, the least-squares estimate will be biased to more reliable data. This can be accomplished by solving the preconditioned problem:

$$\mathbf{Qd} = \mathbf{QRPM} \quad (6.4)$$

where \mathbf{Q} is a preconditioning matrix that accomplishes the desired image-space weighting. Note that preconditioning is applied to the problem in k -space.

The least-squares solution to Eq. 6.4 is:

$$\hat{\mathbf{m}}_{\mathbf{Q}} = (\mathbf{P}^* \mathbf{R}^* \mathbf{Q}^* \mathbf{Q} \mathbf{R} \mathbf{P})^{-1} \mathbf{P}^* \mathbf{R}^* \mathbf{Q}^* \mathbf{Q} \mathbf{d} \quad (6.5)$$

$$(6.6)$$

The estimate in Eq. 6.6 filters the k -space data by $\mathbf{Q}^* \mathbf{Q}$ prior to refocusing. The magnitude effects of the filter are then removed by the unmixing matrix $(\mathbf{P}^* \mathbf{R}^* \mathbf{Q}^* \mathbf{Q} \mathbf{R} \mathbf{P})^{-1}$. To weight each image-space voxel by the navigator magnitude, $\mathbf{Q}^* \mathbf{Q}$ must convolve each high-resolution interleave with the Fourier transform of the navigator magnitude. In other words, $\mathbf{R}^* \mathbf{Q}^* \mathbf{Q} \mathbf{d}$ should contain the interleave images weighted by the navigator magnitude.

A weighted refocusing reconstruction can then be defined by approximating the

mixing matrix $\mathbf{P}^*\mathbf{R}^*\mathbf{Q}^*\mathbf{Q}\mathbf{R}\mathbf{P}$ as diagonal by zeroing the off-diagonal elements. The approximated unmixing matrix removes the weighting caused by the preconditioning but does not attempt to remove any aliasing. Regularizing can be used to improve the stability of the inversion in noise regions with low signal.

6.6 Extending the Refocusing Reconstruction

The derivation of the refocusing reconstruction presented here does not implicitly depend on the details of the DWI method used in a given experiment. However, the underlying assumptions made in approximating the least-squares estimate are likely to depend on the specifics of the imaging method. This section considers these dependences and discusses possible extensions of the refocusing reconstruction to address these issues.

6.6.1 Generality of Reconstruction

The least-squares reconstruction is general in the sense that its derivation does not make assumptions about the trajectory or pulse sequence used to gather the data. Although the refocusing reconstruction will not always be a good approximation of the least-squares solution, it is expected to perform well provided the k -space modulation induced by phase corruptions has minimal discontinuity and periodicity. In addition to the 2DFT data shown here, we have applied the refocusing reconstruction to the self-navigated spiral trajectory presented in Chapter 5 [112] and recent work by another group has found the refocusing reconstruction to out-perform standard linear corrections with a self-navigated EPI trajectory [113]. The latter work acquired relatively few excitations (4 interleaves with 8 repetitions), indicating that the requirement for a large number of excitations may not be a major limitation of the refocusing reconstruction. Pipe and colleagues have also had good success using a refocusing reconstruction with a FSE PROPELLER sequence [57] (although in their

case, the refocusing reconstruction is the same as the least-squares reconstruction since each readout is unaliased).

While full least-squares reconstruction is generally intractable due to the inversion of the mixing matrix, it may be manageable in special cases or with efficient approximations. For example, the mixing matrix \mathbf{M} will be block diagonal and sparse for trajectories in which the signal from a voxel can only alias onto a subset of the other voxels (see Appendix B). In this case, the matrix inversion of \mathbf{M} can be calculated by individually inverting the smaller block submatrices. In our experience, this calculation is still very time-consuming. Alternatively, the mixing matrix could be approximated with an iterative scheme similar to that proposed for multi-coil reconstructions [114].

6.6.2 Application to Other Sequences and Trajectories

The refocusing reconstruction will provide a good reconstruction of the data provided (a) a large number of excitations are acquired and (b) k -space modulations induced by motion are smooth and non-periodic.

The requirement for many excitations minimizes aliasing energy by allowing the incoherent phase errors to cancel. We have already discussed the improvement in the refocused image from averaging; other methods of increasing the number of excitations will have the same effect. For example, a 3D acquisition should be better-conditioned to a refocusing reconstruction than a 2D multi-slice acquisition. Simply increasing the number of interleaves in an acquisition will also improve the refocused image. The need for a large number of excitations is a disadvantage for some sequences such as spin echo which require long echo times. These methods may benefit from a weighted refocusing reconstruction which will tend to suppress these errors, particularly if regularization is used.

The requirement for smooth and non-periodic k -space modulations ensures that the refocusing operator will be able to remove phase corruptions. This requirement

can be addressed in most trajectories with cardiac synchronization. Some trajectories cannot simultaneously achieve smoothness and non-periodicity, in which case we prefer to meet the smoothness constraint since the refocusing reconstruction does a poor job of removing discontinuities. Spiral and radial scans can achieve a smooth k -space weighting by distributing the rotated interleaves from 0° to 180° over each cardiac cycle (this scheme would result in azimuthal periodicity in k -space). Segmented EPI scans can also be made smooth if we acquire each segment of a single image at roughly the same portion of the cardiac cycle. The cardiac cycle could then be filled out by acquiring more slices (in multi-slice), more z-phase encodes (for 3D) or more averages. All of these methods use cardiac information to achieve well-behaved modulations of the k -space data without requiring large increases in scan time.

6.7 Summary

This work presents a method for correcting nonlinear phase errors in diffusion-weighted data using 2D navigation. The refocusing reconstruction, which is an approximation to the computationally-expensive least-squares estimate, is a straightforward algorithm that makes complete use of the information in a 2D navigator. The heart of the reconstruction is the refocusing operator, which rephases the unaliased component of each acquisition via multiplication with the navigator phase conjugate. This reconstruction is a good approximation to the full least-squares estimate provided the phase errors cause smooth, non-periodic modulations of k -space and a sufficient number of excitations are acquired. Cardiac triggering information can be used to achieve desirable modulations through ordering of data interleaves. The refocusing reconstruction is expected to be applicable in a wide range of imaging conditions.

Chapter 7

Summary and Recommendations

7.1 Summary of Contributions

This thesis has presented new techniques for the important neuroimaging methods of functional MRI and diffusion-weighted imaging. The pulse sequences developed are based on SSFP imaging, which has high SNR efficiency and unusual signal properties. In addition, a new method for reconstruction of navigated DWI data was presented which provides significant improvement over existing methods. The major contributions of this work are summarized below.

- **BOSS FMRI:** A new method for functional MRI, BOSS FMRI, was developed. BOSS obtains functional contrast based on the sensitivity of the balanced-SSFP signal phase to off-resonance. The advantages of BOSS FMRI over the standard BOLD method include higher SNR efficiency, reduction of image artifacts, reduced sensitivity to susceptibility and stronger functional contrast. A method for increasing spatial coverage with measurements at multiple frequencies was also proposed.
- **Balanced-SSFP phase description:** The dependence of the balanced-SSFP signal phase on resonance frequency was characterized. An intuitive description

of the source of signal phase was developed, and the phase profile was calculated for a range of relaxation times, sequence timing parameters and flip angles.

- **Navigated SSFP-DWI:** A method for incorporating navigation into the SSFP-DWI pulse sequence was developed and implemented. Navigation significantly reduces the sensitivity of SSFP-DWI to motion, which had previously prevented the pulse sequence from being practically useful. Navigated SSFP-DWI was demonstrated at high resolution for 2D imaging of the brain and 3D imaging of articular knee cartilage.
- **Quantitative DWI of short- T_2 species:** A simple approximation for the SSFP-DWI signal was presented, which holds under certain imaging conditions. This approximation enables quantification of the diffusion coefficient in short- T_2 species. SSFP-DWI was demonstrated to have high SNR in these species, where conventional DWI methods fail. This work presented the first clinical-quality diffusion-weighted images of *in vivo* cartilage.
- **Refocusing reconstruction of navigated DWI:** A method for reconstructing navigated DWI based on the least-squares optimal reconstruction was formulated. An efficient approximation to this reconstruction, the refocusing reconstruction, was proposed and characterized. This method was shown to significantly reduce motion artifacts in SSFP-DWI data. Cardiac synchronization and frame reordering were shown to improve the conditioning of the refocusing reconstruction.

7.2 Recommendations for Future Work

7.2.1 BOSS FMRI

- **Comparison with BOLD FMRI:** The theoretical comparisons of BOSS and BOLD functional signals presented in Chapter 4 do not take physiological fluctuations into account. Physiological sources of signal variability tend to dominate

over thermal noise processes in BOLD imaging, particularly at high magnetic field strengths. The functional CNR of BOSS and BOLD FMRI should be compared over a number of subjects and cortical regions.

- **Extension to 3D:** The results presented in Chapter 4 are from single-slice acquisitions. In balanced-SSFP imaging, the most practical multi-slice acquisition schemes are true-3D trajectories (i.e., 3D Fourier-encoded acquisitions). To maintain frame rates suitable for studying brain function, this will mean fast-3D acquisitions such as stack-of-EPI or stack-of-spirals.
- **Reduced sensitivity to temporal frequency drift:** The frequency sensitivity of BOSS functional contrast means that any temporal drift in the center frequency of a voxel will change the sensitivity of the signal to activity. Temporal frequency drifts can result from scanner instability, subject motion or physiological effects such as respiration. One potential method for reducing the sensitivity to such drifts is optimization the balanced-SSFP sequence parameters, as discussed in Chapter 4. Another potential method is to measure the bulk center frequency during the experiment and feed this information back into the pulse sequence so that the transmit frequency can be adjusted to track frequency drifts.
- **Analysis of multi-frequency experiments:** The multi-frequency results presented in this work combine the data from multiple frequencies using a simple maximum contrast projection (MCP) measure. While this method is simple, it is not a maximally sensitive or statistically valid inference method. Future work should consider more careful methods for analysis of multi-frequency BOSS data.

7.2.2 Navigated SSFP-DWI

- **3D navigated SSFP-DWI of the brain:** The SSFP-DWI brain data presented in this work used 2D acquisitions, and required a large amount of averaging. A 3D sequence with the same voxel size could acquire a slice for every average without incurring any loss of SNR. Unlike the 3D cartilage data presented here, the complexity of motion artifacts in the brain is likely to require 3D navigation; unfortunately, acquiring a 3D navigator volume every T_R is intractable. Given that most of the motion in the brain arises from cardiac pulsatility, it may be possible to acquire a multi-shot 3D navigator, provided the acquisition window is fast relative to the cardiac cycle. For example, a 3D stack-of-spirals navigator could acquire a single z phase encode each period. The full navigator for a given period could then be reconstructed using the data frames that are temporally the closest to the current period (essentially a sliding window reconstruction). The navigator would be a valid representation of the phase corruption as long as the motion is varying more slowly than the navigator acquisition window, which is only likely to be valid for rapid acquisitions like SSFP.
- **High-resolution DTI of the brain:** One of the most compelling applications of SSFP-DWI is diffusion tensor imaging (DTI) or high angular resolution diffusion (HARD) imaging, in which directionality of white matter tracts are inferred from diffusion anisotropy. One advantage of SSFP-DWI with true 3D acquisitions would be the ability to achieve isotropic voxels. Anisotropic voxels can introduce bias to estimates of directionality due to partial volume effects and differential sensitivity to motion. Current DWI data also tends to have poor spatial resolution, enabling only the largest white matter fibre bundles to be visualized. SSFP-DWI has strong potential for high spatial resolution due to high SNR efficiency and the low levels of image distortion. High resolution, isotropic voxels are critically important for diffusion tractography analysis.
- **High q -value imaging:** There has been increasing interest in imaging with

strong diffusion weighting (also referred to as high q -value imaging, where $q = \gamma G \tau$). The apparent diffusion process measured in DWI is highly sensitive to q , with low q -values reflecting average diffusion and high q -values differentiating between separate components of the underlying diffusion process. Large q -values reveal information about the differential structure of tissue, and are in many ways more interesting than low q -values. However, most DWI methods do not have sufficient SNR to image at high q values. The SNR efficiency of SSFP-DWI make it well-suited to high q values. An interesting aspect of SSFP-DWI is the capability to increase the amount of diffusion weighting by decreasing the flip angle. This method of increasing the diffusion weighting maintains important aspects of the sequence (q value, echo time, eddy currents, etc), which may make it a robust method for studying the diffusion process.

7.2.3 Refocusing Reconstruction of Navigated DWI

- **Application to other trajectories:** As discussed in Chapter 6, the refocusing reconstruction is expected to be applicable to most navigated trajectories, provided the navigator covers a full FOV. Application of this reconstruction to other trajectories should provide improvements over standard linear corrections. In some cases (e.g., acquisitions with relatively few excitations) may reveal the limits of the refocusing approximation.
- **Algorithm variants:** Variations on the refocusing reconstruction may improve image quality in some circumstances. The preconditioning step proposed in Chapter 6 would suppress data with intra-voxel dephasing (similar to the thresholding used in most other navigator reconstructions). Similarly, in cases where the refocusing approximation is poor, methods for approximating the unmixing matrix (e.g., using iterative techniques or selectively reconstructing “important” components of the matrix) might be useful.
- **Multi-coil imaging:** Recent research in multi-coil imaging has resulted in the development of SENSE/SMASH reconstruction, which use the (complex) coil

sensitivity patterns as a second set of basis functions in addition to Fourier encoding. By reducing the amount of data that needs to be acquired, SENSE enables a reduction in the image acquisition time. Scan time reductions would be helpful in DWI, for example in HARD imaging where a large number of diffusion directions are acquired. However, the presence of motion-induced artifacts complicates SENSE reconstruction of multi-shot DWI data. SENSE reconstruction, which combines coil data which is weighted by different phase and magnitude spatial profiles, is strikingly similar to the navigated-DWI reconstruction problem: both motion-induced artifacts and coil sensitivity take the form of complex weighting in the object domain. A navigator measured by a given coil in a multi-coil array would reflect the composite effect of motion artifact and coil sensitivity, which should be removable with a single reconstruction similar to that developed in Chapter 6. In this case, the weighting to be removed would contain both phase and magnitude corruptions, suggesting that the weighted refocusing reconstruction would be appropriate. SENSE reconstructions are sensitive to the orthogonality of the coil sensitivity patterns, which could be reduced by motion-induced phase. Careful regularization of the reconstruction should be able to address this issue.

- **Estimation of directional bias:** The refocusing reconstruction has been shown to remove most of the motion artifacts in diffusion-weighted images. However, the resulting images contain residual effects which differ based on the direction of diffusion weighting. One error, which can be immediately noted from the images shown in Chapter 5, is that images weighted along directions with intense motion contain some residual attenuation (for example, the images with weighting along S/I are slightly attenuated relative to the images weighed only along A/P and R/L). This error creates a bias in measurements of diffusion attenuation in which the diffusion coefficient is overestimated along directions that exhibit the most motion. It may be possible to estimate and remove these errors from the navigators (similar to the weighted reconstruction discussed in Chapter 6).

- **Estimation of directional SNR:** A second residual effect in refocused images is variable SNR. In addition to the attenuation described above, the SNR of the refocused images will vary both across images (due to greater motion artifacts along certain directions) and spatially within an image (due to greater motion in certain locations). In addition to the diffusion-weighted images shown in Chapter 5, low-resolution SNR maps could be estimated from the navigator data. This additional information could be gracefully incorporated into tractography to appropriately weight a given measurement by an estimate of its reliability (e.g., Bayesian tractography).

Appendix A

Early diagnosis of osteoarthritis

The development of tools for diagnosing and monitoring the progression of osteoarthritis (OA) is of considerable clinical interest [115]. OA is a leading cause of chronic disability in the United States, affecting 10% of the population over 30 years old [115]. In OA, the cartilage matrix slowly degrades, resulting in pain in the underlying bone during movement or load-bearing exercise. Collagen matrix breakdown may be one of the earliest markers of OA [116]. A number of MRI techniques have been proposed as potential OA diagnostics in recent years. T_2 -mapping can detect collagen loss, but requires long scan times [100]. Delayed gadolinium-enhanced MRI of cartilage (dGEMRIC) is sensitive to early changes in proteoglycan content, but is time-consuming and requires a contrast injection [117]. Sodium imaging is also useful for imaging proteoglycan depletion but has low SNR and requires special hardware due to the shifted resonance frequency of sodium [118].

DWI is a promising alternative to these methods [98, 119] since its sensitive to structural properties of tissue and may be able to detect early cartilage matrix damage [96, 97]. While a number of *in vitro* studies have been performed [61, 96–98, 119, 120], *in vivo* DWI of cartilage has proven difficult due to the need for high resolution and the short T_2 of cartilage (20–40 ms [98]). Standard DWI techniques such as spin echo (SE-DWI) require long echo times, causing these methods to suffer from poor

SNR due to heavy T_2 decay prior to readout. SNR has proven to be a significant challenge for *in vivo* DWI of cartilage [121]. Motion is also a significant obstacle for *in vivo* DWI. Motion-artifacts can be corrected in multi-shot data using a low-resolution phase reference known as a navigator [66, 70]. *In vivo* DWI of cartilage will require a method with high SNR for short- T_2 species and navigator correction of motion artifacts.

Appendix B

Matrix representation of motion corruption

The motion-corrupted magnetization during the c^{th} readout is denoted $m_c(\mathbf{r}) = |m_c(\mathbf{r})|e^{i\phi_c(\mathbf{r})}$.

The phase corruption during the c^{th} excitation is expressed by the $N^2 \times N^2$ matrix \mathbf{P}^c :

$$\mathbf{P}^c = \begin{pmatrix} e^{i\phi_c(\mathbf{r}_1)} & 0 & \dots & 0 \\ 0 & e^{i\phi_c(\mathbf{r}_2)} & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & e^{i\phi_c(\mathbf{r}_{N^2})} \end{pmatrix} \quad (\text{B.1})$$

$\mathbf{P}^c \mathbf{m}$ is the phase-corrupted magnetization that is the source of signal during the c^{th} readout. Each of the C interleaves has a different phase corruption, represented by stacking the \mathbf{P}^c 's to form \mathbf{P} :

$$\mathbf{P} = \begin{pmatrix} \mathbf{P}^1 \\ \mathbf{P}^2 \\ \vdots \\ \mathbf{P}^C \end{pmatrix} \quad (\text{B.2})$$

\mathbf{P} splits the magnetization vector \mathbf{m} into C separate phase-corrupted magnetization vectors. These vectors are stacked to form the $CN^2 \times 1$ vector \mathbf{Pm} representing the phase-corrupted magnetization that is sampled over the course of the experiment.

Each of the phase-corrupted magnetization vectors is separately Fourier transformed by the matrix \mathbf{F} ($CN^2 \times CN^2$):

$$\mathbf{F} = \begin{pmatrix} \mathbf{F}^N & 0 & \dots & 0 \\ 0 & \mathbf{F}^N & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \mathbf{F}^N \end{pmatrix} \quad (\text{B.3})$$

\mathbf{F} is composed of C 2DFT matrices \mathbf{F}^N ($N^2 \times N^2$) consisting of the Fourier kernels:

$$\mathbf{F}_{lm}^N = e^{i2\pi\mathbf{k}_l \cdot \mathbf{r}_m} \quad (\text{B.4})$$

The product \mathbf{FPm} is the $CN^2 \times 1$ vector representing the C phase-corrupted spectra, sampled in Cartesian coordinates.

Finally, the spectra given by \mathbf{FPm} are resampled onto the k -space sampling trajectories by the resampling matrix \mathbf{G} ($CR \times CN^2$):

$$\mathbf{G} = \begin{pmatrix} \mathbf{G}^1 & 0 & \dots & 0 \\ 0 & \mathbf{G}^2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \mathbf{G}^C \end{pmatrix} \quad (\text{B.5})$$

Each submatrix \mathbf{G}^c ($R \times N^2$) interpolates the Cartesian k -space points \mathbf{k}_l onto the sampling points for the c^{th} readout \mathbf{s}_m^c :

$$\mathbf{G}_{lm}^c = \text{sinc}(\mathbf{k}_l - \mathbf{s}_m^c) \quad (\text{B.6})$$

Note that for trajectories that collect data on Cartesian coordinates (such as spinwarp or interleaved EPI), the elements of \mathbf{G} are discrete delta functions.

The interpolation from the Cartesian coordinates onto the k -space trajectory is the inverse of the reconstruction problem, which seeks to resample the k -space data onto an evenly-spaced Cartesian grid. The latter operation is more mathematically intricate, however, since the data may be more densely sampled in some portions of k -space than in others. Least-squares reconstruction uses the pseudo-inverse of \mathbf{G} to reconstruct data. Previous work has shown that gridding reconstruction can also be expressed using matrix formalism and is an approximation to the least-squares reconstruction [122].

Using the notation developed above, the elements of the mixing matrix \mathbf{M} for an arbitrary trajectory are:

$$\mathbf{M}_{lm} = \sum_{c=1}^C e^{i(\phi_c(\mathbf{r}_m) - \phi_c(\mathbf{r}_l))} \sum_{r=1}^R \left[\sum_{n=1}^{N^2} \text{sinc}(\mathbf{k}_n - \mathbf{s}_r^c) e^{+i2\pi \mathbf{k}_n \cdot \mathbf{r}_m} \right] \left[\sum_{n=1}^{N^2} \text{sinc}(\mathbf{k}_n - \mathbf{s}_r^c) e^{-i2\pi \mathbf{k}_n \cdot \mathbf{r}_l} \right] \quad (\text{B.7})$$

The summation over each readout r describes non-idealities in the k -space trajectory that causes coupling between voxels within a readout. The phase difference describes the aliasing caused by motion-induced phase errors. These phase offsets are manifested as perturbations of the encoding functions used in the trajectory.

Bibliography

- [1] Purcell EM, Torrey HC, Pound RV. Resonance absorption by nuclear magnetic moments in a solid. *Phys Rev* 1946; 69:37.
- [2] Bloch F, Hansen WW, Packard ME. Nuclear induction. *Phys Rev* 1946; 69:127.
- [3] Lauterbur PC. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature* 1973; 242:190.
- [4] Hahn EL. Spin echoes. *Phys Rev* 1950; 80:580–594.
- [5] Mansfield P. Multi-planar image-formation using NMR spin echoes. *Journal of Physics C-Solid State Physics* 1977; 10:L55–L58.
- [6] Meyer CH, Hu BS, Nishimura DG, Macovski A. Fast spiral coronary artery imaging. *Magn Reson Med* 1992; 28:202–213.
- [7] Buxton RB, Edelman RR, Rosen BR, Wismer GL, Brady TJ. Contrast in rapid MR imaging– T1-weighted and T2-weighted imaging. *J Comput Assist Tomogr* 1987; 11:7–16.
- [8] Hawkes RC, Patz S. Rapid Fourier imaging using steady-state free precession. *Magn Reson Med* 1987; 4:9–23.
- [9] Zur Y, Stokar S, Bendel P. An analysis of fast imaging sequences with steady-state transverse magnetization refocusing. *Magn Reson Med* 1988; 6:175–193.

- [10] Patz S. Some factors that influence the steady-state in steady-state free precession. *Magn Reson Imaging* 1988; 6:405–413.
- [11] Carr HY. Steady-state free precession in nuclear magnetic resonance. *Phys Rev Lett* 1958; 112:1693–1701.
- [12] Oppelt A, Graumann R, Barfuss H, Fischer H, Hartl W, Shajor W. FISP— a new fast MRI sequence. *Electromedica* 1986; 54:15–18.
- [13] Buxton RB, Fisel CR, Chien D, Brady TJ. Signal intensity in fast NMR imaging with short repetition times. *J Magn Reson* 1989; 83:576–585.
- [14] Jaynes E. Matrix treatment of nuclear induction. *Phys Rev* 1955; 98:1099–1105.
- [15] Baird AE, Warach S. Magnetic resonance imaging of acute stroke. *J Cereb Blood Flow Metabol* 1998; 18:583–609.
- [16] Aronen HJ, Perkio J. Dynamic susceptibility contrast MRI of gliomas. *Neuroimaging Clin N Am* 2002; 12:501–523.
- [17] Arnold DL, Matthews PM. MRI in the diagnosis and management of multiple sclerosis. *Neurology* 2002; 48:S23–S31.
- [18] Scheltens P. Early diagnosis of dementia: neuroimaging. *J Neurol* 1999; 246:16–20.
- [19] Shenton ME, Dickey CC, Frumin M, McCarley RW. A review of MRI findings in schizophrenia. *Skeletal Rad* 2001; 49:1–52.
- [20] Lee BC, Mintun M, Buckner RL, Morris JC. Imaging of Alzheimer’s disease. *J Neuroimaging* 2003; 13:199–214.
- [21] Rudin M, Beckmann N, Porszasz R, Reese T, Bochen D, Sauter A. In vivo magnetic resonance methods in pharmaceutical research: current status and perspectives. *NMR Biomed* 1999; 12:69–97.

- [22] Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, Cheng HM, Brady TJ, Rosen BR. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci* 1992; 89:5675–5679.
- [23] Ogawa S, Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, Ugurbil K. Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci* 1992; 89:5951–5955.
- [24] Conturo TE, Lori NF, Cull TS, Akbudak E, Snyder AZ, Shimony JS, McKinstry RC, Burton H, Raichle ME. Tracking neuronal fiber pathways in the living human brain. *Proc Natl Acad Sci* 1999; 96:10422–10427.
- [25] Hammeke TA. Functional MRI in neurology. *in* “Functional MRI” (Moonen CTW, Bandettini PA, Eds.), chapter 39. Springer, 1 ed., 1999.
- [26] Frahm J, Bruhn H, Merboldt KD, Hanicke W, Math D. Dynamic MR imaging of human brain oxygenation during rest and photic stimulation. *J Magn Reson Imaging* 1992; 2:501–505.
- [27] Buxton RB, “Introduction to Functional Magnetic Resonance Imaging”. Cambridge University Press, Cambridge, 1st ed., 2002.
- [28] Moonen CTW, Bandettini PA, “Functional MRI”. Springer, Berlin, 1st ed., 1999.
- [29] Jezzard P, Matthews PM, Smith SM, “Functional MRI: An introduction to methods”. Oxford University Press, Oxford, 1st ed., 2001.
- [30] Raichle ME. Circulatory and metabolic correlates of brain function in normal humans. *in* “Handbook of physiology— the nervous system” (Mountcastle VB, Plum F, Geiger SR, Eds.), pp. 643–674. American Physiological Society, Bethesda, 1987.

- [31] Ollinger JM, Fessler JA. Positron-emission tomography. *IEEE Signal Processing Magazine* 1997; 14:43–55.
- [32] Gevins AS, Zeitlin GM, Doyle JC, Yingling CD, Schaffer RE, Callaway E, Yeager CL. Electroencephalogram correlates of higher cortical functions. *Science* 1979; 203:665–668.
- [33] Bronzino JD. Principles of electroencephalography. *in* “The Biomedical Engineering Handbook” (Bronzino JD, Ed.), Vol. 1, chapter 15. CRC Press, 2 ed., 2000.
- [34] Malmivuo J. Biomagnetism. *in* “The Biomedical Engineering Handbook” (Bronzino JD, Ed.), Vol. 1, chapter 16. CRC Press, 2 ed., 2000.
- [35] Brandeis D, Lehmann D. Event-related potentials of the brain and cognitive-processes— Approaches and applications. *Neuropsychologia* 1986; 24:151–168.
- [36] Kutas M, Hillyard SA. Event-related potentials in cognitive science. *in* “Handbook of Cognitive Neuroscience” (Gazzaniga MS, Ed.). Plenum Press, New York, 1984.
- [37] Stern MD. In vivo evaluation of microcirculation by coherent light-scattering. *Nature* 1975; 254:56–58.
- [38] Dirnagl U, Kaplan B, Jacewicz M, Pulsinelli W. Continuous measurement of cerebral cortical blood-flow by Laser-Doppler Flowmetry in a rat stroke model. *J Cereb Blood Flow Metabol* 1989; 9:589–596.
- [39] Villringer A, Haberl RL, Dirnagl U, Anneser F, Verst M, Einhaupl KM. Confocal laser microscopy to study microcirculation on the rat-brain surface in vivo. *Brain Research* 1989; 504:159–160.
- [40] Detre JA, Leigh JS, Williams DS, Koretsky AP. Perfusion imaging. *Magn Reson Med* 1992; 23:37–45.

- [41] Williams DS, Detre JA, Leigh JS, Koretsky AP. Magnetic-Resonance-Imaging of perfusion using spin inversion of arterial water. *Proc Natl Acad Sci* 1992; 89:212–216.
- [42] Edelman RR, Siewert B, Darby DG, Thangaraj V, Nobre AC, Mesulam MM, Warach S. Qualitative mapping of cerebral blood-flow and functional localization with echo-planar MR-imaging and signal targeting with alternating radio-frequency. *Radiology* 1994; 192:513–520.
- [43] Belliveau JW, Kennedy DN, Mckinstry RC, Buchbinder BR, Weisskoff RM, Cohen MS, Vevea JM, Brady TJ, Rosen BR. Functional mapping of the human visual-cortex by magnetic-resonance-imaging. *Science* 1991; 254:716–719.
- [44] Mandeville JB, Marota JJA, Kosofsky BE, Keltner JR, Weissleder R, Rosen BR, Weisskoff RM. Dynamic functional imaging of relative cerebral blood volume during rat forepaw stimulation. *Magn Reson Med* 1998; 39:615–624.
- [45] Liu TT, Luh WM, Wong EC, Frank LR, Buxton RB. A method for dynamic measurement of blood volume with compensation for T2 changes. In: *Proc 8th ISMRM, Denver, 2000*. p. 52.
- [46] Fox PT, Raichle ME. Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. *Proc Natl Acad Sci* 1986; 83:1140–1144.
- [47] Pauling L, Coryell C. The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbon monoxyhemoglobin. *Proc Natl Acad Sci* 1936; 22:210–216.
- [48] Thulborn KR, Waterton JC, Matthews PM, Radda GK. Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochim Biophys Acta* 1982; 714:265–270.
- [49] Ogawa S, Lee TM, Nayak AS, Glynn P. Oxygenation-sensitive contrast in magnetic resonance images of rodent brain at high magnetic fields. *Magn Reson Med* 1990; 14:68–78.

- [50] Gati JS, Menon RS, Ugurbil K, Rutt BK. Experimental determination of the BOLD field strength dependence in vessels and tissue. *Magn Reson Med* 1997; 38:296–302.
- [51] Jezzard P, Clare S. Sources of distortion in functional MRI data. *Human Brain Mapping* 1999; 8:80–85.
- [52] Glover GH. 3D z-shim method for reduction of susceptibility effects in BOLD fMRI. *Magn Reson Med* 1999; 42:290–299.
- [53] Farzaneh F, Riederer SJ, Pelc NJ. Analysis of T2 limitations and off-resonance effects on spatial-resolution and artifacts in echo-planar imaging. *Magn Reson Med* 1990; 14:123–139.
- [54] Moseley ME, Cohen Y, Mintorovich J, Chileuitt L, Shimizu H, Kucharczyk J, Wendland MF, Weinstein PR. Early detection of regional cerebral ischemia in cats: comparison of diffusion and T2-weighted MRI and spectroscopy. *Magn Reson Med* 1990; 14:330–346.
- [55] Callaghan PT, “Principles of nuclear magnetic resonance microscopy”. Oxford Science Publications, New York, 1st ed., 1991.
- [56] Stejskal EO, Tanner JE. Spin-diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J Chem Phys* 1965; 19:306–317.
- [57] Pipe JG, Farthing VG, Forbes KP. Multishot diffusion-weighted FSE using PROPELLER MRI. *Magn Reson Med* 2002; 47:42–52.
- [58] Merboldt KD, Hanicke W, Frahm J. Self-diffusion NMR imaging using stimulated echoes. *J Magn Reson* 1985; 64:479–486.
- [59] LeBihan D. Intravoxel incoherent motion imaging using steady-state free precession. *Magn Reson Med* 1988; 7:346–351.
- [60] LeBihan D, Turner R, Macfall JR. Effects of intravoxel incoherent motions (IVIM) in steady-state free precession (SSFP) imaging: application to molecular diffusion imaging. *Magn Reson Med* 1989; 10:324–337.

- [61] Henkelman RM, Stanisz GJ, Kim JK, Bronskill MJ. Anisotropy of NMR properties of tissues. *Magn Reson Med* 1994; 32:592–601.
- [62] Moseley ME, Cohen Y, Mintorovich J, Chileuitt L, Shimizu H, Kucharczyk J, Tsuruda J, Weinstein PR, Norman D. Evidence of anisotropic self-diffusion. *Radiology* 1990; 176:439–445.
- [63] Pierpaoli C, Basser PJ. Toward a quantitative assessment of diffusion anisotropy. *Magn Reson Med* 1996; 36:893–906.
- [64] Greitz D, Wirestam R, Franck A, Nordell B, Thomsen C, Stahlberg F. Pulsatile brain movement and associated hydrodynamics studied by magnetic resonance phase imaging. *Neuroradiology* 1992; 34:370–380.
- [65] Poncelet BP, Wedeen VJ, Weisskoff RM, Cohen MS. Brain parenchyma motion: Measurement with cine echo-planar MR imaging. *Radiology* 1992; pp. 645–651.
- [66] Anderson AW, Gore JC. Analysis and correction of motion artifacts in diffusion weighted imaging. *Magn Reson Med* 1994; 32:379–387.
- [67] Norris DG. Implications of bulk motion for diffusion-weighted imaging experiments: effects, mechanisms and solutions. *J Magn Reson Imaging* 2001; 13:486–495.
- [68] Trouard TP, Theilmann RJ, Altbach MI, Gmitro AF. High-resolution diffusion imaging with DIFRAD-FSE (diffusion-weighted radial acquisition of data with fast spin-echo) MRI. *Magn Reson Med* 1999; 42:11–18.
- [69] Norris DG, Driesel W. Online motion correction for diffusion-weighted imaging using navigator echoes: Application to RARE imaging without sensitivity loss. *Magn Reson Med* 2001; 45:729–733.
- [70] Ordidge RJ, Helpert JA, Qing ZX, Knight RA, Nagesh V. Correction of motion artifacts in diffusion-weighted MR images using navigator echoes. *Magn Reson Imaging* 1994; 12:455–460.

- [71] deCrespigny AJ, Marks MP, Enzmann DR, Moseley ME. Navigated diffusion imaging of normal and ischemic human brain. *Magn Reson Med* 1995; 33:720–728.
- [72] Butts K, Pauly J, Crespigny A, Moseley M. Isotropic diffusion-weighted and spiral-navigated interleaved EPI for routine imaging of acute stroke. *Magn Reson Med* 1997; 38:741–749.
- [73] Atkinson D, Porter DA, Hill DLG, Calamante F, Connelly A. Sampling and reconstruction effects due to motion in diffusion-weighted interleaved echo planar imaging. *Magn Reson Med* 2000; 44:101–109.
- [74] Freeman R, Hill HDW. Phase and intensity anomalies in Fourier transform NMR. *J Magn Reson* 1971; 4:366–383.
- [75] Clare S, Hykin J, Bowtell R, Coxon R, Morris PG. Sub-millimeter resolution functional brain imaging with multislice interleaved EPI at 3.0T. In: *Proc 4th ISMRM*, New York, 1996. p. 1821.
- [76] Thulborn KR, Chang SY, Shen GX, Voyvoddic JT. High-resolution echo-planar fMRI of human visual cortex at 3.0 Tesla. *NMR Biomed* 1997; 10:183–190.
- [77] Hoogenraad FG, Hofman MB, Pouwels PJ, Reichenbach JR, Rombouts SA, Haacke EM. Sub-millimeter fMRI at 1.5 Tesla: correlation of high resolution with low resolution measurements. *J Magn Reson Imaging* 1999; 9:475–482.
- [78] Hyde JS, Biswal BB, Jesmanowicz A. High-resolution FMRI using multi-slice partial k-space GR-EPI with cubic voxels. *Magn Reson Med* 2001; 46:114–125.
- [79] Kim SG, Ugurbil K. High-resolution functional magnetic resonance imaging of the animal brain. *Methods* 2003; 30:28–41.
- [80] Haacke EM, Brown RW, Thompson MR, Venkatesan R, “Magnetic Resonance Imaging: Physical Principles and Sequence Design”. Wiley-Liss, New York, 1st ed., 1999.

- [81] Matwiyoff NA, Gasparovic C, Mazurchuk R, Matwiyoff G. On the origin of paramagnetic inhomogeneity effects in whole-blood. *Magn Reson Med* 1991; 20:144–150.
- [82] Gasparovic C, Matwiyoff NA. The magnetic properties and water dynamics of the red blood cell: A study by proton-NMR lineshape analysis. *Magn Reson Med* 1992; 26:274–299.
- [83] Noll DC, Pauly JM, Meyer CH, Nishimura DG, Macovski A. Deblurring for non-2D Fourier-transform magnetic-resonance-imaging. *Magn Reson Med* 1992; 25:319–333.
- [84] Irarrazabal P, Meyer CH, Nishimura DG, Macovski A. Inhomogeneity correction using an estimated linear field map. *Magn Reson Med* 1996; 35:278–282.
- [85] Scheffler K. A pictorial description of steady-states in rapid magnetic resonance imaging. *Concepts Magn Reson* 1999; 11:291–304.
- [86] Vasanawala SS, Pauly JM, Nishimura DG. Linear combination steady-state free precession MRI. *Magn Reson Med* 2000; 43:82–90.
- [87] Macovski A. Noise in MRI. *Magn Reson Med* 1996; 36:494–497.
- [88] Yang YH, Glover GH, vanGelderens P, Patel AC, Mattay VS, Frank JA, Duyn JH. A comparison of fast MR scan techniques for cerebral activation studies at 1.5 Tesla. *Magn Reson Med* 1998; 39:61–67.
- [89] Lai S, Glover GH. Three-dimensional spiral fMRI technique: A comparison with 2D spiral acquisition. *Magn Reson Med* 1998; 39:68–78.
- [90] Buxton RB. The diffusion sensitivity of fast steady-state free precession imaging. *Magn Reson Med* 1993; 29:235–243.
- [91] Quest RA, McRobbie DW. Investigation of the visual pathways using diffusion-weighted PSIF. In: *Proc 9th ISMRM, Glasgow, 2001*. p. 1548.

- [92] Deoni SCL, Peters TM, Rutt BK. Quantitative diffusion imaging with steady-state free precession. *Magn Reson Med* 2004; 51:428–433.
- [93] Press WH, Teukolsky SA, Vetterling WT, Flannery BP, “Numerical Recipes in C”. Cambridge University Press, New York, NY, 2nd ed., 1994.
- [94] Noll DC, Nishimura DG. Homodyne detection in magnetic resonance imaging. *IEEE Trans Med Imaging* 1991; 10:154–163.
- [95] Enzmann DR, Pelc NJ. Brain motion: Measurement with phase-contrast MR imaging. *Radiology* 1992; 185:653–660.
- [96] Frank LR, Wong EC, Luh WM, Ahn JM, Resnick D. Articular cartilage in the knee: Mapping of the physiologic parameters at MR imaging with a local gradient coil— preliminary results. *Radiology* 1999; 210:241–246.
- [97] Mlynarik V, Sulzbacher I, Bittsansky M, Fuiko R, Trattnig S. Investigation of apparent diffusion constant as an indicator of early degenerative disease in articular cartilage. *J Magn Reson Imaging* 2003; 17:440–444.
- [98] Xia Y, Farquhar T, Burtonwurster N, Ray E, Jelinski LW. Diffusion and relaxation mapping of cartilage-bone plugs and excised disks using microscopic magnetic-resonance-imaging. *Magn Reson Med* 1994; 31:273–282.
- [99] Duewell SH, Ceckler TL, Ong K, Wen H, Jaffer FA, Chesnick SA, Balaban RS. Musculoskeletal MR imaging at 4 T and at 1.5 T: Comparison of relaxation times and image contrast. *Radiology* 1995; 196:551–555.
- [100] Mosher TJ, Dardzinski BJ, Smith MB. Human articular cartilage: influence of aging and early symptomatic degeneration on the spatial variation of T2— preliminary findings at 3 T. *Radiology* 2000; 214:259–266.
- [101] Butts K, deCrespigny A, Pauly JM, Moseley M. Diffusion-weighted interleaved echo-planar imaging with a pair of orthogonal navigator echoes. *Magn Reson Med* 1996; 35:763–770.

- [102] Mori S, van Zijl PCM. A motion correction scheme by twin-echo navigation for diffusion-weighted magnetic resonance imaging with multiple RF echo acquisition. *Magn Reson Med* 1998; 40:511–516.
- [103] Williams CFM, Redpath TW, Norris DG. A novel fast split-echo multi-shot diffusion-weighted MRI method using navigator echoes. *Magn Reson Med* 1999; 41:734–742.
- [104] Brockstedt S, Moore JR, Thomsen C, Holtas S, Stahlberg F. High-resolution diffusion imaging using phase-corrected segmented echo-planar imaging. *Magn Reson Imaging* 2000; 18:649–657.
- [105] Clark CA, Barker GJ, Tofts PS. Improved reduction of motion artifacts in diffusion imaging using navigator echoes and velocity compensation. *J Magn Reson* 2000; 142:358–363.
- [106] Bosak E, Harvey PR. Navigator motion correction of diffusion weighted 3D SSFP imaging. *MAGMA* 2001; 12:167–176.
- [107] Jiang H, Golay X, van Zijl PCM, Mori S. Origin and minimization of residual motion-related artifacts in navigator-corrected segmented diffusion-weighted EPI of the human brain. *Magn Reson Med* 2002; 47:818–822.
- [108] Bailes DR, Gilderdale DJ, Bydder GM, Collins AG, Firmin DN. Respiratory ordered phase encoding (ROPE): A method for reducing respiratory motion artifacts in MR imaging. *J Comput Assist Tomogr* 1985; 9:835.
- [109] Haacke EM, Patrick JL. Reducing motion artifacts in two-dimensional Fourier transform imaging. *Magn Reson Imaging* 1986; 4:359–376.
- [110] Frigo M, Johnson SG. FFTW: An adaptive software architecture for the FFT. *Proc. ICASSP* 1998; 3:1381–1384.
- [111] Basser PJ, Pierpaoli C. A simplified method to measure the diffusion tensor from seven MR images. *Magn Reson Med* 1998; 39:928–934.

- [112] Miller KL, Meyer CH, Pauly JM. Self-navigated spirals for high resolution steady-state diffusion imaging. In: Proc 10th ISMRM, Honolulu, 2002.
- [113] Nunes RG, Jezzard P, Clare S. Non-linear motion correction for diffusion imaging using a self-navigated cartesian-based sequence. In: Proc 12th ISMRM, Kyoto, 2004.
- [114] Preussmann KP, Weiger M, Bornert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med* 2001; 46:638–651.
- [115] Felson DT, Lawrence RC, Dieppe PA, Hirsch R, Helmick CG, Jordan JM, Kington RS, Lane NE, Nevitt MC, Zhang Y, Sowers M, McAlindon T, Spector TD, Poole AR, Yanovski SZ, Ateshian G, Sharma L, Buckwalter JA, Brandt KD, Fries JF. Osteoarthritis: new insights. Part 1: the disease and its risk factors. *Ann Intern Med* 2000; 133:635–46.
- [116] Maier CF, Tan SG, Hariharan H, Potter HG. T2 quantitation of articular cartilage at 1.5 T. *J Magn Reson Imaging* 2003; 17:358–364.
- [117] Bashir A, Gray ML, Burstein D. Gd-DTPA²⁻ as a measure of cartilage degradation. *Magn Reson Med* 1996; 36:665–673.
- [118] Reddy R, Insko EK, Noyszewski EA, Dandora R, Kneeland JB, Leigh JS. Sodium MRI of human articular cartilage *in vivo*. *Magn Reson Med* 1998; 39:697–701.
- [119] Burstein D, Gray ML, Hartman AL, Gipe G, Foy BD. Diffusion of small solutes in cartilage as measured by nuclear resonance (NMR) spectroscopy and imaging. *J Orthop Res* 1993; 11:465.
- [120] Knauss R, Schiller J, Fleischer G, Karger J, Arnold K. Self-diffusion of water in cartilage and cartilage components as studied by pulsed field gradient NMR. *Magn Reson Med* 1999; 41:285–292.

- [121] Gold GE, Butts K, Fechner KP, Bergman G, Beaulieu CF, Lang PK, Macovski A. *In vivo* diffusion-weighted imaging of cartilage. In: Proc 6th ISMRM, Sydney, 1998. p. 1066.
- [122] Sedarat H, Nishimura DG. On the optimality of the gridding reconstruction algorithm. IEEE Trans Med Imaging 2000; 19:306–317.