QUANTITATIVE MAGNETIC RESONANCE IMAGING OF CELLULAR DENSITY WITH TURBOSPI

by

James Rioux

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

at

Dalhousie University
Halifax, Nova Scotia
August 2012

DALHOUSIE UNIVERSITY

DEPARTMENT OF PHYSICS AND ATMOSPHERIC SCIENCE

The undersigned hereby certify that they have read and recommend to the Faculty of Graduate Studies for acceptance a thesis entitled "QUANTITATIVE MAGNETIC RESONANCE IMAGING OF CELLULAR DENSITY WITH TURBOSPI" by James Rioux in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Dated: August 1, 2012

External Examiner: _____
Dr. Blaine Chronik

Research Supervisors: _____
Dr. Chris Bowen

_____
Dr. Kevin Hewitt

Examining Committee: _____
Dr. Steven Beyea

_____
Dr. Laurent Kreplak

Departmental Representative: _____
Dr. Jordan Kyriakidis

DALHOUSIE UNIVERSITY

DATE: August 1, 2012

AUTHOR: James Rioux

TITLE: QUANTITATIVE MAGNETIC RESONANCE IMAGING OF CELLULAR DENSITY WITH TURBOSPI

DEPARTMENT OR SCHOOL: Department of Physics and Atmospheric Science

DEGREE: Ph.D. CONVOCATION: October YEAR: 2012

Signature of Author

To my wife Kim:

my truth,

my guide,

my life.

Table of Contents

List of Tables . . . . . **ix**

List of Figures . . . . . **x**

Abstract . . . . . **xiv**

List of Abbreviations and Symbols Used . . . . . **xv**

Acknowledgements . . . . . **xviii**

Chapter 1 Introduction **1**

1.1 Cellular Imaging with MRI 1
1.1.1 Contrast Agents for Cellular MRI 3
1.1.2 Techniques for Cellular MRI 6
1.1.3 Examples of Cellular MRI 8

1.2 Quantitative Cellular MRI 13
1.2.1 Signal-Based Techniques 13
1.2.2 Relaxometry Techniques 17

1.3 Project Overview and Hypotheses 20
1.3.1 Contributions to Research 21

Chapter 2 Background **23**

2.1 Magnetic Susceptibility in Materials 23

2.2 Nuclear Magnetic Resonance 25
2.2.1 Classical Description 25
2.2.2 Quantum Mechanical Description 28
2.2.3 Signal Detection and Relaxation 32
2.2.4 Spin and Stimulated Echoes 38

2.3 Magnetic Resonance Imaging 42
2.3.1 Spatial Encoding 42
2.3.2 Image Acquisition 43
2.3.3 Properties of k-space 49
2.3.4 Acquisition Hardware 55

Chapter 3 Implementation and Optimization of TurboSPI 57

3.1 Single Point Imaging Concepts 57
3.1.1 TurboSPI Sequence 59

3.2 Alignment of Stimulated Echoes 63
3.2.1 Calibration Procedure 66
3.2.2 Calibration Testing - Methods 68
3.2.3 Calibration Testing - Results 69

3.3 Discussion and Summary 69

Chapter 4 *In Vitro* Quantification of SPIO with TurboSPI 73

4.1 Relaxation Behaviour of SPIO 73
4.1.1 Relaxation in the Static Dephasing Regime 75
4.1.2 FID Experiment with Spherical Perturbers 77
4.1.3 Spin-Echo Experiment 79
4.1.4 Conditions for the Static Dephasing Regime 80

4.2 Methods 82
4.2.1 Sample Preparation 82
4.2.2 Imaging Parameters 84
4.2.3 Quantification Procedure 86

4.3 Results 87

4.4 Discussion and Summary 92

Chapter 5 Extended Analytical Description of SPIO Relaxation . 95

5.1 Relaxation Outside of the Static Dephasing Regime 95
5.1.1 Predicted Behavior 96
5.1.2 Overview of Alternate Models 98

5.2 Extension of the Slow Diffusion Model 100
5.2.1 Description of the Model 100
5.2.2 Approximate Forms 105

5.3 Validation of the Model 111
5.3.1 Monte Carlo Simulations 111
5.3.2 Experimental Data 114

5.4 Applications 115
5.4.1 Improved Quantification 115
5.4.2 Robustness to Field Inhomogeneity 118

5.5 Summary 120

Chapter 6 Acceleration of TurboSPI 121

6.1 Theory - MRI with Compressed Sensing 121

6.2 Compressed Sensing with TurboSPI - Methods 125
6.2.1 Reconstruction with Prior Information 125
6.2.2 Phase Encode Ordering and Undersampling 131
6.2.3 Assessment of Reconstruction Quality 134
6.2.4 Assessment of Quantification Accuracy 135

6.3 Compressed Sensing with TurboSPI - Results 135
6.3.1 Quality of Reconstructed Images 135
6.3.2 Impact on Quantification 140

6.4 Discussion and Summary 141

Chapter 7 *In Vivo* Imaging with TurboSPI 143

7.1 Challenges of Animal Imaging 143

7.2 Methods 144
7.2.1 Fat Saturation 144
7.2.2 *In Vivo* Imaging Methods 146

7.3 Results 147
7.3.1 Fat Saturation - Effect on Image Quality 147
7.3.2 Fat Saturation - Effect on Compressed Sensing 149
7.3.3 Fat Saturation - Effect on Quantification 151
7.3.4 Detection of Labeled Cells 151
7.3.5 Discussion and Summary 153

Chapter 8 Future Work and Conclusions 157

8.1 Future Work 157
8.1.1 Modeling Relaxation Behavior 157
8.1.2 Compressed Sensing 159
8.1.3 *In Vivo* Imaging 160

8.2 Conclusions 161

Appendix A 163

A.1 Field of a Magnetized Sphere 163

A.2 Static Dephasing Condition for Spin-Echo Experiment 165

A.3 Propagator Solution of the Bloch-Torrey Equation 167
A.3.1 Propagator Solution for FID Experiment 169
A.3.2 Propagator Solution for Spin-Echo Experiment 169

Bibliography **172**

List of Tables

Table 4.1 Relaxivities and LMD of all sample groups. 91

Table 4.2 Properties of each sample group relevant to the static dephasing conditions. 93

Table 5.1 Parameters used for generating f^{SE} for comparisons. 108

Table 5.2 Parameters used for comparison to experimental data. 115

Table 6.1 Undersampling parameters for image quality assessment. 135

List of Figures

Figure 1.1 Classification of SPIO by particle size. 4
Figure 1.2 Examples of cellular MRI of stem cells. 9
Figure 1.3 Examples of cellular MRI of immune cells. 12
Figure 1.4 Examples of signal-based quantification. 14
Figure 1.5 Limitations of signal-based quantification. 16
Figure 1.6 Examples of relaxometry-based quantification. 19
Figure 1.7 Overview of project components. 20

Figure 2.1 Larmor precession of a nucleus in a magnetic field. 26
Figure 2.2 The rotating frame of reference. 27
Figure 2.3 Manipulation of magnetization with an RF pulse. 28
Figure 2.4 Examples of FIDs and NMR spectra. 33
Figure 2.5 Spin-lattice relaxation. 35
Figure 2.6 Spin-spin relaxation. 35
Figure 2.7 Effects of molecular motion on relaxation times. 36
Figure 2.8 FID signal decay and line broadening. 38
Figure 2.9 Formation of a spin echo. 39
Figure 2.10 CPMG echo train for T_2 measurement. 40
Figure 2.11 Formation of a stimulated echo. 41
Figure 2.12 Relationship between imaged object and k-space. 44
Figure 2.13 Frequency and phase encoding. 45
Figure 2.14 Shaped RF pulse for slice-selective excitation. 47
Figure 2.15 Relationship between RF pulse and slice parameters. 48
Figure 2.16 Examples of basic MRI pulse sequences. 49
Figure 2.17 Image formation from spatial frequencies. 50
Figure 2.18 Image blurring due to signal decay. 52

Figure 2.19 Effects of finite sampling in k-space. 53
Figure 2.20 Aliasing due to insufficient sampling density. 54
Figure 2.21 Block diagram of a basic MRI system. 56

Figure 3.1 Single Point Imaging pulse sequence. 58
Figure 3.2 Spin-Echo SPI pulse sequence. 59
Figure 3.3 Basic TurboSPI pulse sequence. 59
Figure 3.4 Extensions of the basic TurboSPI sequence. 60
Figure 3.5 Linear phase encode ordering. 61
Figure 3.6 Example of a TurboSPI image series. 62
Figure 3.7 Denoising of TurboSPI time courses. 63
Figure 3.8 Crusher pattern for stimulated echo suppression. 64
Figure 3.9 Crusher pattern for stimulated echo retention. 65
Figure 3.10 Artifacts created by improper alignment of echoes. 65
Figure 3.11 Automatic calibration of TurboSPI gradient and phase. 67
Figure 3.12 Effect of automatic calibration on TurboSPI images. 70
Figure 3.13 TurboSPI image quality throughout the time course. 71

Figure 4.1 Relaxivity as a function of particle size. 74
Figure 4.2 Cylindrical phantom for SPIO quantification. 84
Figure 4.3 Procedure for MR susceptometry. 85
Figure 4.4 R_2^* quantification using TurboSPI. 87
Figure 4.5 TurboSPI time course from SPIO-loaded cells. 88
Figure 4.6 Example of a TurboSPI R_2^* map. 88
Figure 4.7 Relaxation rates for SPIO-loaded C3 cells. 89
Figure 4.8 Relaxation rates for MPIO particles. 90
Figure 4.9 Variability of fitted R_2^* values. 92

Figure 5.1 Comparison of full and approximate f^{SE}. 107

Figure 5.2 Predicted versus calculated peak shifts. 109
Figure 5.3 Predicted peak shifts from empirical expression. 109
Figure 5.4 Comparison of full and revised approximate f^{SE}. 110
Figure 5.5 Comparison of model with Monte Carlo simulations. 113
Figure 5.6 Comparison of model and Monte Carlo with TurboSPI data. . 116
Figure 5.7 Modulation introduced by macroscopic gradients. 119
Figure 5.8 Removal of modulation by macroscopic gradients. 120

Figure 6.1 Incoherent artifacts from random undersampling. 123
Figure 6.2 Conceptual illustration of compressed sensing. 124
Figure 6.3 TurboSPI with a readout gradient added. 126
Figure 6.4 TurboSPI image compared to matched FSE image. 127
Figure 6.5 CS reconstruction of a TurboSPI image series. 128
Figure 6.6 CS reconstruction of a single TurboSPI image. 130
Figure 6.7 Centric phase encode ordering. 131
Figure 6.8 Centric PE ordering in three dimensions. 132
Figure 6.9 Undersampled centric-order phase encoding. 133
Figure 6.10 Undersampling based on matched FSE image. 133
Figure 6.11 Comparison of reconstructed images at t = TE = 8 ms. 137
Figure 6.12 Error of reconstructed TurboSPI images. 138
Figure 6.13 Comparison of reconstructed images at t = 9.5 ms. 139
Figure 6.14 Effect of CS acceleration on TurboSPI time courses. 140
Figure 6.15 Effect of CS acceleration on measured R_2^*. 141

Figure 7.1 TurboSPI pulse sequence with fat saturation. 145
Figure 7.2 Effect of fat saturation on TurboSPI images. 147
Figure 7.3 Effect of fat saturation away from t=TE. 148
Figure 7.4 Time courses with and without fat saturation. 149

Figure 7.5 Effect of fat saturation on CS reconstruction. 150
Figure 7.6 Effect of fat saturation on image nRMSE. 150
Figure 7.7 Effects of fat saturation on quantification. 151
Figure 7.8 Images of SPIO-labeled cells in a mouse tumor. 152
Figure 7.9 Basic relaxometry of *in vivo* TurboSPI data. 153
Figure 7.10 Potential differences in image specificity. 154

Abstract

Magnetic Resonance Imaging (MRI) is capable of detecting and tracking cells which have been labeled with contrast agents such as superparamagnetic iron oxide (SPIO). For the evaluation of emerging cellular therapies and other applications, it is desirable for MRI to produce quantitative, longitudinal measurements of cellular density in living organisms. Such quantification is challenging with many traditional imaging methods. In this work, an MRI technique called TurboSPI is proposed for the quantification of cellular systems. TurboSPI acquires data that can be used to generate maps of the MRI relaxation rate R_2', which is directly related to the concentration of SPIO and therefore the density of cells.

The first stage of this project involved demonstrating the accuracy and range of this quantification using micron-sized iron oxide particles and cells labeled with SPIO. These experimental results generally confirmed the predictions of the established theory but also showed some behavior that was not predicted. Based on this, an extended analytical description of MRI signal relaxation in the presence of spherical magnetic perturbers was developed. This extended model compared favorably to Monte Carlo simulations and experimental data, and may serve as the basis for more sophisticated quantification strategies.

The main drawback of TurboSPI is its slow acquisition speed, which can be overcome with the use of compressed sensing, a signal processing technique for reconstructing data that has been significantly undersampled. TurboSPI is well suited to acceleration with compressed sensing, particularly because it permits the collection of prior information that can constrain the reconstruction. It is demonstrated that acceleration factors of 10 to 15 are possible without significant loss of image quality or quantification ability. TurboSPI is then demonstrated in animal models, with examples showing detection of iron and improved specificity with respect to traditional techniques. This represents the first application of TurboSPI to cellular imaging *in vivo*.

List of Abbreviations and Symbols Used

Abbreviations

b-SSFP - Balanced steady state free precession

CPMG - Carr, Purcell, Meiboom and Gill sequence

CS - Compressed sensing

CSI - Chemical shift imaging

DC - Dendritic cell

FID - Free induction decay

FOV - Field of view

FSE - Fast spin-echo

GESFIDE - Gradient echo sampling of FID and echo

GESSE - Gradient echo sampling of spin echo

GRASP - Gradient echo acquisition for superparamagnetic particles

GRE - Gradient echo

LMD - Local magnetic dose

MPIO - Micron-sized particles of iron oxide

MRI - Magnetic resonance imaging

NMR - Nuclear magnetic resonance

nRMSE - Normalized root-mean-square error

PET - Positron emission tomography

PDF - Probability density function

PSF - Point spread function

RF - Radio frequency

ROI - Region of interest

SE - Spin echo

SDR - Static dephasing regime

SNR - Signal-to-noise ratio

SPI - Single point imaging

SPIO - Super paramagnetic iron oxide

TE - Echo time

TR - Repetition time

TurboSPI - Turbo single point imaging

USPIO - Ultra-small particles of iron oxide

UTE - Ultrashort Echo Time

Symbols

χ - Magnetic susceptibility

$\Delta\chi$ - Magnetic susceptibility difference

$\delta\omega$ - Characteristic frequency

γ - Gyromagnetic ratio

λ - Dimensionless diffusion coefficient

ω_0 - Larmor frequency

τ - Dimensionless time variable

τ_E - Dimensionless echo time

ζ - Volume fraction (of magnetic perturbers)

B - Magnetic field

B_0 - Homogeneous external magnetic field

B_1 - Oscillating transmit magnetic field

D - Diffusion coefficient

f - Signal attenuation function

G - Magnetic field gradient

k - Reciprocal space position

M - Magnetization

R - Radius (of magnetic perturber)

R_1 - Longitudinal (spin-lattice) relaxation rate

R_2 - Irreversible transverse (spin-spin) relaxation rate

R_2' - Reversible transverse relaxation rate

R_2^* - Effective transverse relaxation rate

$s(t)$ - NMR signal as a function of time

T_1 - Longitudinal (spin-lattice) relaxation time

T_2 - Irreversible transverse (spin-spin) relaxation time

T_2^* - Effective transverse relaxation time

Acknowledgements

As with so many endeavors, this thesis was the result of several factors coming together at the right time, and I was simply lucky enough to be in the right place and surrounded by the right people to take advantage of that opportunity. For this I am continually grateful.

When I was finishing up my term as a technical officer with the NRC's Institute for Biodiagnostics (Atlantic) and preparing to return to grad school, I was fortunate to have a choice of interesting projects and enthusiastic supervisors. It was Chris Bowen's field of cellular and molecular imaging that caught my interest the most, and he gave me the chance to apply what I'd learned at NRC and during my undergraduate degrees to the problem of cellular MRI. Since that time, Chris has always been a vocal champion of my work, making efforts at various conferences to ensure it reached a wide audience, and he has been a tremendous help in deciding the future direction of my career. He always tried his best to reassure me in difficult times and keep me on track when I was getting bogged down. Hopefully someday I will be a good supervisor because of what I learned from working with him.

My other unofficial supervisor for this project was Steven Beyea, who had just left Bruce Balcom's MRI Research Centre as I was starting there as a co-op student. I didn't know it at the time, but TurboSPI, the sequence that would become the core of my Ph.D. research, had been recently developed at UNB by Bruce and Steven. When I came to work at the NRC, Steven was my guide in the unfamiliar world of human imaging, and throughout my time as both a technical officer and graduate student he has been an invaluable mentor and a good friend. I have a hard time picturing him without that characteristic smile on his face, and I am honored that I've been able to take the tool he invented and teach it a new trick or two.

Though Steven was my day-to-day supervisor while I worked for the NRC, the chance to work there in the first place came from Ryan D'Arcy, the group leader of IBD (Atlantic). Most of what I little I know about neuroimaging and the human brain is due to Ryan, who always found a way to get me thinking outside of my comfort

zone. We haven't always seen eye to eye on some of the realities of commercialization versus pure research, and perhaps we never will, but through our interactions I've certainly learned to think about these matters from a perspective that I wouldn't have considered otherwise, and for that I thank him.

Though my interactions with the Physics Department at Dalhousie were often limited, due to my research generally taking place off-campus, I am grateful for the contributions of several members of that department. For the first several years of my graduate studies, Gerhard Stroink was a valuable member of my supervisory committee, expressing an early interest in my project and staying involved even after his retirement. Kevin Hewitt was just as enthusiastic and helpful in the later stages of my work. Krista Cullymore and Tanya Timmins, administrators for the graduate program, were always helpful and kind, and quite generous in their constant reminders of very important things that I would have nonetheless found a way to forget.

The majority of my research was done at the NRC's Biomedical MRI Research Lab, which was just opening as I was beginning my program. The entire team at the BMRL, both past and present, contributed to this research in some way. Christa Davis, Iulia Dude and Kerry Lake were all involved in animal handling and preparation for my *in vivo* imaging experiments, and though I did eventually learn the mechanics of mouse MRI to an acceptable level, I was always glad to have their experience and assistance. At times, Drew DeBay basically kept the lab running through sheer force of will, and Kirk Feindel (who seems to know a little bit about everything) was happy to answer questions and offer advice.

Several other students have passed through IBD (Atlantic) during my time there, and as fellow travelers on the road, we learn to rely on each other as best we can. There are far too many to name here, but a few stand out in my mind.

Ahmed Elkady and Alison Hill laid a lot of the groundwork for the *in vitro* quantification of SPIO, figuring out the best way to make and image the MPIO and cell samples. Ahmed was also responsible for the first "shotgun" phantom used for imaging groups of tubes, and for providing the Matlab code for processing NMR susceptometry data. Ahmed was constantly cheerful and optimistic, and always a pleasure to work with.

Erin Mazerolle was one of the first students at the Neuroimaging Research Lab and later the BMRL, and as I write this she has just successfully defended her own Ph.D thesis. Her road to completion was a far longer and rockier one than mine, but she endured to the very end. Her continued enthusiasm and kind words for everyone around her were always appreciated.

Nicole Pelot was a recurring presence in the BMRL, switching projects with each co-op term but always making it seem like she'd been doing each one forever, such was her skill and knowledge. I know I had to remind myself more than once that she was still an undergrad; her competence, drive and energy were the equal of any grad student in the lab. Her example inspires all of us to try a little harder.

Steve Patterson originally wrote much of the Monte Carlo simulation code used for verification of the analytical model, for a project which involved modeling signal changes in functional brain imaging - as with the model itself, I simply changed the system geometry from blood-vessel-approximating cylinders to spheres. Steve is a great guy whose laid-back attitude conceals a sharp and discerning mind, who can talk about financial policy as readily as hockey playoffs.

Speaking of the analytical model, the description of relaxation near spherical perturbers that I developed as part of this thesis was a natural extension of the model for cylinders proposed in 1999 by Valerij Kiselev of the University Medical Center in Freiburg, Germany. I contacted Dr. Kiselev in 2010 when I first began developing the model and he was exceptionally helpful at explaining some of the derivations and concepts that were not included in his published papers. When I asked to meet with him at the ISMRM conference in Stockholm to ask about a particular formula, he ended up re-deriving it from memory on the flight there, and added a way to reproduce the model's major results using a simple physical argument (which is reproduced in Section 5.1.1). I look forward to working further with him as we ready the spherical model for publication.

The blessing of excellent colleagues is matched only by that of extraordinary teachers. Grant Williams first sparked my interest in physics, which only grew under Stephen Ross, Colan Linton, Zong-Chao Yan, Dennis Tokaryk - and yes, even Abdelhaq Hamza, who in the end was as fair as he was demanding. Bruce Balcom, Ben

Newling and Bryce MacMillan all shaped the way I think about MRI, even to this day. Thanks too to Stanimir Bonev and Andrew Rutenberg, for demonstrating that physics gets more interesting as you learn more about it and start to see how all the different pieces fit together in new and unexpected ways.

Of course, during my education I have learned more than just the theory and practice of my chosen field. Among many others, the discipline of meditation has been very helpful at keeping me grounded and focused when I would otherwise have been confused and distracted. I thank my teachers in the various aspects of meditation: Tim, who taught me of mindfulness; Daniel, who taught me of gratitude; Laura, who taught me of acceptance; Bob, who taught me of stillness; and Paul, who taught me of letting go. Through them I have learned that everything I was looking for was already within me.

My family, both immediate and extended, has been a wellspring of support and guidance throughout my life. In particular, my parents, Gary and Brenda Rioux, have provided the solid foundation on which all my accomplishments are constructed, and no matter where my path leads I know they will continue to encourage me. We have never been a family that talks much about our feelings and our desires, and ironically I find myself unable to fully express in words the extent of my gratitude to them for all they have taught and given me. But as I begin to contemplate a future in which I too may become a parent, I find myself aspiring to follow their example, and I believe that this is the most sincere thanks I can give them.

The bulk of the writing of this thesis was done at home where I was, shall we say, *assisted* by our cats, Valerie and Bellus. Valerie took it upon herself to act as my occupational health and safety officer, ensuring that I never worked uninterrupted for too long, and Bellus served as morale officer, insisting that belly rubs made her feel better, and if that didn't work for me I should keep trying until it did.

Finally, there is one person whose story has interwoven itself with mine, starting at UNB, then NRC, through Dalhousie and the BMRL, and whose end is far from written. Kim Brewer has been my love, my friend and my companion for nine years, five of which we've spent as husband and wife. She of course has also been a valuable colleague; in her role as a research scientist at Immunovaccine she was responsible

for preparing the SPIO-loaded cells and organizing many of the animal experiments described in this work. But her ongoing support, patience and kindness have been just as essential. In every dark night she has been a hand to hold and a light to follow, and in every bright day I have been thankful to have gotten the chance to know her and to love her. She never ceases to inspire me with her generosity, her creativity, her depth of emotion, and her strength of will and character. Kim, through your love and example, you have turned an arrogant, prideful boy into a confident and proud man - confident that I can do anything with you at my side, and proud to have you there with me. Wherever we go, as long as I am with you, I am home.

Chapter 1

Introduction

Since its development in the 1970s, Magnetic Resonance Imaging (MRI) has established itself as an invaluable tool for non-invasive imaging, with a tremendous range of applications from testing of materials to mapping function in the brain. A more recent development has been the use of MRI for cellular imaging, which has the goal of detecting, tracking and quantifying numbers of cells within living organisms. Such imaging will be instrumental in the development and assessment of cellular therapies, which have the potential to revolutionize the treatment of a variety of medical conditions. Evaluating these therapies will require monitoring and quantitative measurement of their effectiveness *in vivo*, which can be challenging even with established cellular MRI methods, many of which cannot quantify the high cell concentrations used in therapeutic contexts.

In this work we present an MRI technique called TurboSPI which will be applied to cellular imaging for the first time. After an overview of existing cellular MRI techniques, potential applications in cellular imaging, and some of the underlying physical theory, we will demonstrate that the TurboSPI technique is well suited to quantification of cellular density or concentration over a wide range. Newly developed theoretical models will allow fuller use to be made of the data that TurboSPI acquires, which could benefit a variety of applications. Data collection and reconstruction methods will be presented that allow TurboSPI to be used effectively in living organisms. We will conclude with a demonstration of cellular detection and potential quantification *in vivo*.

1.1 Cellular Imaging with MRI

The development of cellular imaging has been spurred by the need to observe and understand complex cellular processes, which often involve different types of cells interacting over a long period of time, and which may be sensitive to the number of

cells present or to other biological conditions [1]. While histological studies of biopsies or fixed organs of animal models can be used to study these processes to some degree [2], ultimately they will have to be monitored *in vivo* if they are to be successfully translated into human populations in the form of cellular therapies.

To be suitable for *in vivo* cellular imaging, a particular modality should be non-invasive, non-toxic, reasonably rapid and quantitative, i.e. able to offer numerical data concerning the fate of therapeutic cells [3]. A number of modalities have been proposed for cellular imaging, with varying degrees of success due to their capabilities and limitations. Fluorescent markers can be added to biological molecules that will target particular types of cells, but fluorescence microscopy generally uses very small fields of view (making it difficult to scan entire organs) and has limited penetration depth [1]. Nuclear imaging methods such as Position Emission Tomography (PET) can be highly specific to particular tissues, but their spatial resolution is generally limited to 1 mm or greater, and longitudinal studies are made difficult by the short half-lives of most radionuclides and the need for repeated radiation doses.

Given these requirements, MRI has emerged as a promising candidate for a viable cellular imaging modality. MRI is widely available in most clinical settings and has been deployed at an increasing number of pre-clinical imaging laboratories. It is non-invasive, uses no ionizing radiation, can image deep into opaque tissues with arbitrary orientations, and offers reasonable acquisition times. MR images can be produced with a variety of contrasts based on relaxation times or other physical parameters.

Relaxation times will be discussed in more detail in Chapter 2, but briefly, they relate to the molecular processes that drive the evolution of the MRI signal. The relaxation times T_2 and T_2' describe loss of signal through irreversible (e.g. molecular motion, diffusion) and reversible (e.g. static inhomogeneities in the magnetic field) processes, respectively. These are often combined into an effective relaxation time, T_2^*. The return of magnetization to equilibrium after excitation is described by a different relaxation time, T_1. Depending on the particular MRI technique being used, the resulting image can have contrast weighting based on any of these parameters.

Though the highest attainable spatial resolution of most clinical MRI systems is on the order of 100 μm, which is an order of magnitude larger than many cells of interest, the use of MR contrast agents has allowed even single cells to be detected

in vivo using a clinical MRI system [4]. These agents reduce relaxation times in their vicinity, allowing labeled cells to be distinguished from surrounding tissue. Furthermore, this effect is often dependent on the concentration of contrast agent [5], providing a potential avenue for quantification.

In this section we will describe the contrast agents which make cellular MRI possible, and review some of the basic MRI methods used for cellular imaging. We will then summarize some of the advances that have been made towards MRI of cellular processes in living organisms, with a focus on those applications which relate to the imaging and monitoring of cellular therapies.

1.1.1 Contrast Agents for Cellular MRI

One of the main limitations of MRI is its low sensitivity compared to other techniques such as PET or fluorescence imaging. As will be described in the next chapter, only a small fraction of the water molecules in the body contribute to the formation of an MR image, as opposed to other modalities where every radioactive or fluorescent molecule present in the area of interest can be detected. To circumvent this problem, nearly all cellular imaging studies employ a contrast agent of some kind, which can accumulate within a cell and make the cell conspicuous on the final MRI image.

The most commonly used contrast agents for cellular imaging are iron oxide nanoparticles of varying sizes and compositions, which are collectively referred to as Superparamagnetic Iron Oxide or SPIO. The physical mechanism of superparamagnetism will be discussed in Chapter 2; for the purposes of this overview, when placed within an external magnetic field such as that used for MRI, SPIO creates distortions in that field at a far greater intensity and extent than most other substances at the same concentration. This distortion of the magnetic field creates areas of increased T_2 and T_2' relaxation around the SPIO, reducing the MRI signal within a volume far larger than that occupied by the SPIO itself [6] and creating regions of negative contrast. SPIO particles are biodegradable, can be detected by other modalities, and can be loaded into cells by a variety of mechanisms [2, 7].

SPIO is commercially available in a wide range of formulations, most of which contain a mixture of magnetite (Fe_3O_4) and maghemite (γ-Fe_2O_3) encased in a biocompatible coating composed of polymers such as dextran, styrene or polyethylene

glycol [8, 9]. These formulations can be classified by the size of the coated particle (see Figure 1.1). USPIO are ultra-small particles of iron oxide with a diameter in the 10-50 nm range, and USPIO with a single iron oxide crystal at their core are sometimes classified as monocrystalline iron oxide nanoparticles or MION. "Standard" SPIO particles have an inner core composed of multiple crystals and an outer diameter on the order of 50-150 nm. Particles above 300 nm are called micron-sized particles of iron oxide or MPIO, and can reach sizes of 5.8 μm or larger [10, 11].

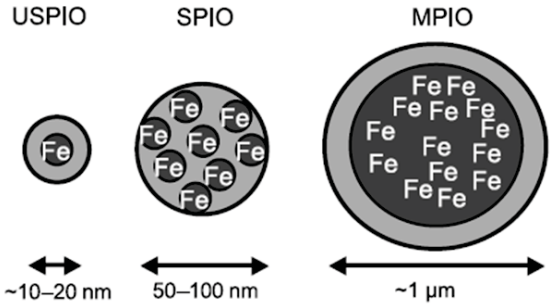

Figure 1.1: Superparamagnetic iron oxide particles can be classified by their effective size into USPIO (ultra-small), SPIO, and MPIO (micron-sized). Figure from [11].

Several formulations of SPIO such as Resovist (60 nm), Sinerem (30 nm) and Feridex (150 nm) have been approved for clinical use[1], though only for intravenous administration in applications such as liver and lymph node imaging [9]. In these cases, phagocytic cells which naturally take up SPIO-sized particles - Kupffer cells in the liver, and monocytes or macrophages in the bloodstream, spleen and lymph nodes - can be passively labeled *in vivo* [8]. Non-phagocytic cells can ingest SPIO and USPIO through pinocytosis, but the rate of uptake is limited for most cell types, particularly those of interest to cellular imaging such as stem cells [12].

Instead of *in vivo* labeling, most cellular imaging studies collect and label cells *in vitro* before injection or implantation into a living organism. The options for such labeling are more diverse, since a variety of agents can be used to increase the efficiency and reproducibility of the uptake, particularly in those cell types which do not take up SPIO naturally. The surface coating of the SPIO can be modified with peptides [13] or smaller particles [14] to greatly increase uptake, sometimes to levels of several million particles per cell or concentrations of 20 pg Fe/cell. Transfection agents with a net electrical charge such as poly-L-lysine can also bind to nanoparticles and

[1]The manufacture and distribution of many of these agents has been discontinued after their withdrawal from the clinical market due to low sales.

chaperone them into a cell [7], achieving similar loading levels. Micron-sized particles are sometimes naturally taken up into cells even without these modifications, and have the added advantage of rendering cells detectable with much smaller numbers of particles; in some cases cells containing a single MPIO particle are detectable, and remain visible within daughter cells after repeated divisions [10, 15].

Once within a cell, SPIO particles will either distribute themselves throughout the cytoplasm or become clustered into a small number of vesicles or endosomes [10, 16]. Regardless of how they are distributed, compartmentalization within the cell has a critical effect on the magnetic properties of the labeled cell, particularly the MRI relaxation times T_2 and T_2', which will be discussed in Chapter 4. In any event, multiple studies have demonstrated that neither the type nor the quantity of SPIO within a cell has any significant effect on the cell's ability to function properly or to reproduce, with only a few exceptions [7, 11].

Contrast agents other than SPIO have been used for cellular MR imaging. A standard contrast agent approved for clinical use is gadolinium (Gd^{3+}), which reduces the relaxation time T_1 and creates areas of increased signal or positive contrast. This leads to increased specificity as compared to the negative contrast often generated by SPIO. Though the gadolinium chelates used for clinical applications are not readily taken up into cells, gadolinium can be attached to other ligands that will target cells of interest [17], or loaded into cells in the form of nanoparticles [18]. One drawback of gadolinium is its low molar relaxivity compared to iron-oxide-based agents, though novel nanostructures containing gadolinium are being developed to address this limitation.

Reporter genes are also a potential source of contrast. If a cell can be engineered to produce an MR-visible substance, such as the protein ferritin which acts as a naturally formed iron nanoparticle, this ability will be passed on to the cell's progeny. This enables the longitudinal study of rapidly dividing cell populations and forgoes the need for *in vitro* labeling or *in vivo* injection of contrast agents [19].

Despite these alternatives and others, SPIO remains the most commonly used contrast agent for cellular MRI due to its high relaxivity, ease of use and commercial availability. This work will focus on imaging and quantifying systems containing cells that have been labeled with SPIO in some way.

1.1.2 Techniques for Cellular MRI

While the presence of SPIO affects all MR relaxation times, the largest effect is on T_2^* [5], and many of the methods used for imaging of SPIO-loaded cells are T_2^*-weighted to provide maximum sensitivity. Three-dimensional gradient-echo (GRE) techniques are often used for cellular imaging and, with sufficient resolution and at high field strengths, can even detect individual cells [20, 21]. However, these techniques are also prone to artefacts caused by other sources of magnetic field distortions, such as those found around air-tissue interfaces. T_2-weighted imaging methods such as Fast Spin-Echo (FSE) are more robust to these effects, but their reduced sensitivity means that they are most effective at very high magnetic field strengths and may not be suitable for widespread clinical implementation.

Balanced Steady-State Free Precession (b-SSFP, also known as FIESTA or True-FISP) has received attention for its ability to combine some of the benefits of gradient-echo and spin-echo sequences. Though not as sensitive as a GRE technique, it is far more sensitive than spin-echo techniques, [22], while providing resistance to artefacts and a high signal-to-noise ratio. The first successful MRI-based detection of single cells *in vivo* at clinical field strengths used b-SSFP [4].

In all of these techniques, the presence of SPIO reduces the MRI signal in its vicinity, producing negative contrast which cannot always be distinguished from other sources of reduced signal. The signal voids around SPIO may also obscure anatomical details in their vicinity. A number of methods have been proposed that instead generate positive contrast, i.e. signal around SPIO appears brighter than the background. This can be done by only exciting MRI signal in regions of inhomogeneous magnetic fields [23], suppressing all signal which is not near field inhomogeneities (a method referred to as IRON) [24], or deliberately applying magnetic field gradients to remove signal from areas not containing SPIO (which is called GRASP) [25]. All of these lead to reduced background signal and increased intensity around accumulations of SPIO, though signal at the center of such regions is still reduced. These techniques remain sensitive to other sources of field inhomogeneity and require some foreknowledge of the expected field distributions, though the use of positive contrast does slightly increase the specificity of these methods.

Another way to generate positive contrast is to extract information in post-processing about SPIO-generated magnetic susceptibility gradients in the image. One example is Susceptibility Gradient Mapping or SGM [26], which uses a local or "short-term" Fourier transform over groups of a few voxels to calculate a susceptibility gradient vector, and a map of this vector's magnitude generates a positive-contrast image. SGM can distinguish low-T_2 materials from SPIO, though they appear identical on normal images, because the former does not induce susceptibility distortions. Phase Gradient Mapping (PGM) [27] is a similar technique that instead calculates susceptibility gradients based on the image's complex phase. The main limitation of these methods is that they require some signal in an imaging voxel to calculate the gradient vector, making then unsuitable for imaging very high concentrations of iron with short T_2^* decay times.

Ultrashort Echo Time or UTE imaging has recently demonstrated its effectiveness in detecting SPIO using a different source of positive contrast. UTE techniques minimize the delay between excitation of MR signal and data acquisition, allowing them to image tissues such as cartilage that have very short T_2 relaxation times [28]. The resulting images have little T_2 contrast and are dominated by T_1 effects, which give positive contrast in the presence of SPIO. This contrast can be further increased by subtracting a background image acquired with a non-UTE sequence, leaving only those parts of the image with increased relaxation due to SPIO [29]. UTE shares the advantages of other positive-contrast methods in terms of specificity and potentially improves on them in terms of resistance to spurious signal from field inhomogeneities. It also directly images the signal from areas containing SPIO instead of signal on their periphery. The main drawback of UTE is its technical difficulty, often requiring specialized hardware and more sophisticated data reconstruction.

Finally, as an alternative to contrast agents which affect relaxation rates, the use of non-proton MRI to image cells is attracting attention. Since the body contains little signal from other MR-sensitive nuclei like ^{19}F, cells labeled with compounds containing ^{19}F will appear readily with little to no background signal, similar to a fluorescence image. If combined with traditional proton MRI, the sites of ^{19}F accumulation can be localized on anatomical images. This requires additional hardware, and due to its limited concentration the resulting ^{19}F MR images are often at lower

resolution, but this remains an emerging technique which will likely undergo significant development as it is more widely adopted. For example, an *in vivo* cell tracking method for tracking fluorine-labeled dendritic cells has been demonstrated in [30].

1.1.3 Examples of Cellular MRI

The techniques described in the previous section have been employed to study a variety of cellular systems, primarily in preclinical settings but occasionally in patient populations. The most widespread target of such study is the emerging field of cellular therapy, in which cells from the patient or a suitable donor are used to repair damaged tissue, incite the immune system to attack tumors, or provide other therapeutic effects without the use of drugs. These therapies may soon allow treatment of neurological disorders, cardiac disfunction, spinal cord injuries, and various forms of cancer. Before they can reach the clinic, these treatments have to be rigorously evaluated in preclinical settings and in clinical trials. The non-invasive and longitudinal monitoring of cellular therapies is a natural application of cellular MRI.

Stem Cell Therapies

Though not the main focus of this thesis, any discussion of cellular therapy would be incomplete without mention of treatments based on the use of stem cells. A stem cell is simply an undifferentiated cell that can give rise to one or more specialized cell types. Embryonic stem cells have the capacity to generate all the different types of cells in the body, but stem cells exist in the adult body as well - neuronal stem cells differentiate into neurons and glial cells, hematopoietic stem cells create all of the cells found in blood, and mesenchymal stem cells give rise to bone marrow, muscle, cartilage, and other tissues [31]. Because of their ability to create new cells, even of types that are difficult or impossible to repair by conventional methods, stem cells have the potential to be used therapeutically in the treatment of numerous pathologies.

With any such therapy, the clinical outcome will depend on the proper engraftment and long-term survival of the grafted cells, which in turn will depend on the delivery route (i.e. direct implantation versus injection into a vein or natural cavity), the initial distribution pattern or extent, as well as the effectiveness of migration to sites of

pathology [32]. Proper monitoring of these processes will be critical to the evaluation of emerging therapies [33], and cellular imaging could replace or supplement the use of biopsies and postmortem analysis to gain such information about the implanted cells.

For example, MRI has been used to monitor preclinical stem cell treatments of neurological disorders such as multiple sclerosis, Alzheimer's disease and Parkinson's disease, all of which could potentially be treated by grafts of stem cells [31]. Migration of SPIO-labeled neural progenitor cells has been imaged in rats with demyelinated neurons using T_2^*-weighted MRI (see Figure 1.2(a)) and correlated with histological evidence of cell differentiation and remyelination [33]. In a rodent model of stroke, Feridex-labeled mesenchymal stem cells were observed to migrate to the site of infarction even when implanted in the contralateral hemisphere, and were visualized for up to 10 weeks [34]. Stem cell therapies for treating cardiac disease and repairing damaged tissue after infarctions are also being explored. Imaging near the heart can be challenging, but clusters of labeled cells have been visualized using both negative-contrast [35, 36] and positive-contrast [37] techniques, with the latter being particularly useful for increasing specificity as shown in Figure 1.2(b).

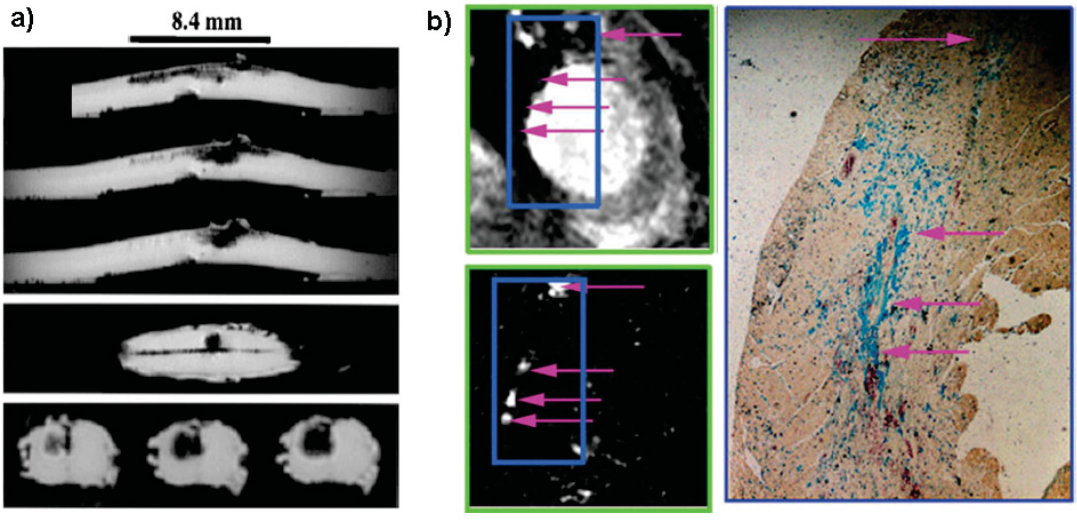

Figure 1.2: Examples of cellular MRI of stem cells. (a) Images showing grafts of neural progenitor cells in a mouse spinal cord, with significant migration of cells visible [33]. (b) Images of SPIO-labeled mesenchymal stem cells in mouse myocardium, acquired 1 week after implantation with a T_2^*-weighted image (top left) and a positive-contrast GRASP image (bottom left), as well as correspondence with histology (right) [37].

One of the major limitations of cellular MRI for long-term monitoring of stem cell therapies is dilution of the contrast agent after repeated cell divisions [32] or after migration of cells. Several studies have demonstrated a loss of visible contrast in the injection sites after two to three weeks [36, 38]. Effective long-term studies will require the use of larger SPIO particles [35, 15] or higher loading levels, as well as a technique capable of imaging both the high initial concentration of iron, and the increasingly reduced concentration as time passes.

Immunotherapies

Another growing subset of cell-based treatments are those which use the body's immune cells to achieve a therapeutic effect, such as the eradication of a tumor. Indeed, cancer-targeting immunotherapies are particularly promising, though several important obstacles must be overcome for such treatments to be effective clinically. Most cancers are not targeted by the immune system because they appear as "self" cells, and this natural tolerance is often supplemented by immunosuppressive elements such as regulatory T cells within the tumor [39]. In order to generate an effective anti-tumor response, antigen-presenting cells such as dendritic cells (DCs) must be activated by a tumor-associated antigen and then migrate to a lymph node, where cytotoxic T cells are primed by the DCs. These T cells migrate to the tumor and, in sufficient quantities, can overcome the local immunosuppressive cells and eliminate the tumor. However, a disruption of any stage in this process will halt the immune response [40], leaving the tumor intact.

Immunotherapies seek to manipulate this process in a number of ways to improve its chances of success. Adoptive cell transfer involves harvesting and *in vitro* activation of dendritic cells or T cells, followed by re-implantation into the patient [39]. Bypassing the natural activation pathway in this manner can considerably increase the strength of the immune response. Alternatively, a vaccine-based approach can be used to introduce large quantities of a tumor-specific antigen in a way that is more readily taken up by the immune system [41]. In either case, cellular MRI could be used to monitor each stage of the immune response [19]. This may allow better biomarkers to be developed for evaluating the success of these therapies, which would be especially valuable because traditional measures such as tumor volume may

not correlate well with eventual outcomes, and more reliable measures of immune response are still being sought [40, 42].

One challenge faced by immunotherapies is that even in the case of adoptive therapies where large numbers of dendritic cells are implanted, only a small fraction (1-10%) of them typically arrive at the draining lymph node [43]. Because dendritic cells can be labeled effectively with a variety of iron oxide formulations, including both SPIO and micron-sized particles [44], monitoring of these cells is relatively straightforward. This could allow the impact of different injection sites, timings, concentrations and formulations to be compared, increasing the final treatment efficacy [19].

For example, the migration of SPIO-labeled DCs from a mouse footpad to a lymph node has been visualized over the course of several days with a multi-echo spin-echo sequence at 4.7 Tesla, with histology used to confirm that approximately 2000 of the 10^6 injected cells arrived successfully [43]. Another study visualized DC migration with a positive-contrast field-mapping technique (shown in Figure 1.3(a)) and determined that SPIO labeling did not affect cell function [45]. Imaging of dendritic cells has also been performed in human subjects (see Figure 1.3(b)), in a study which also demonstrated that injections of cells into the lymph node were only effective in approximately half of the attempts, further emphasizing the need for effective monitoring of these therapies [46].

The behavior of activated T cells can also be monitored, though such studies are more challenging due to the reduced ability of T cells to take up SPIO; special peptides [3] or novel particle formulations [20, 47] are often necessary to achieve high enough loading levels. This difficulty can be bypassed by using a label such as ^{19}F, as done in [48] for longitudinal T cell tracking, for example. However, even with SPIO labeling, the homing of tumor-antigen-specific T cells has been imaged in a mouse model, in which only one of two tumors presented the desired antigen [3]. The tumors were imaged serially, as shown in Figure 1.3(c), and three times as many T cells arrived at the antigen-presenting tumor. As well, when multiple groups of cells were injected, each group homed to a slightly different location, showing that multiple injections may be helpful in clinical treatments.

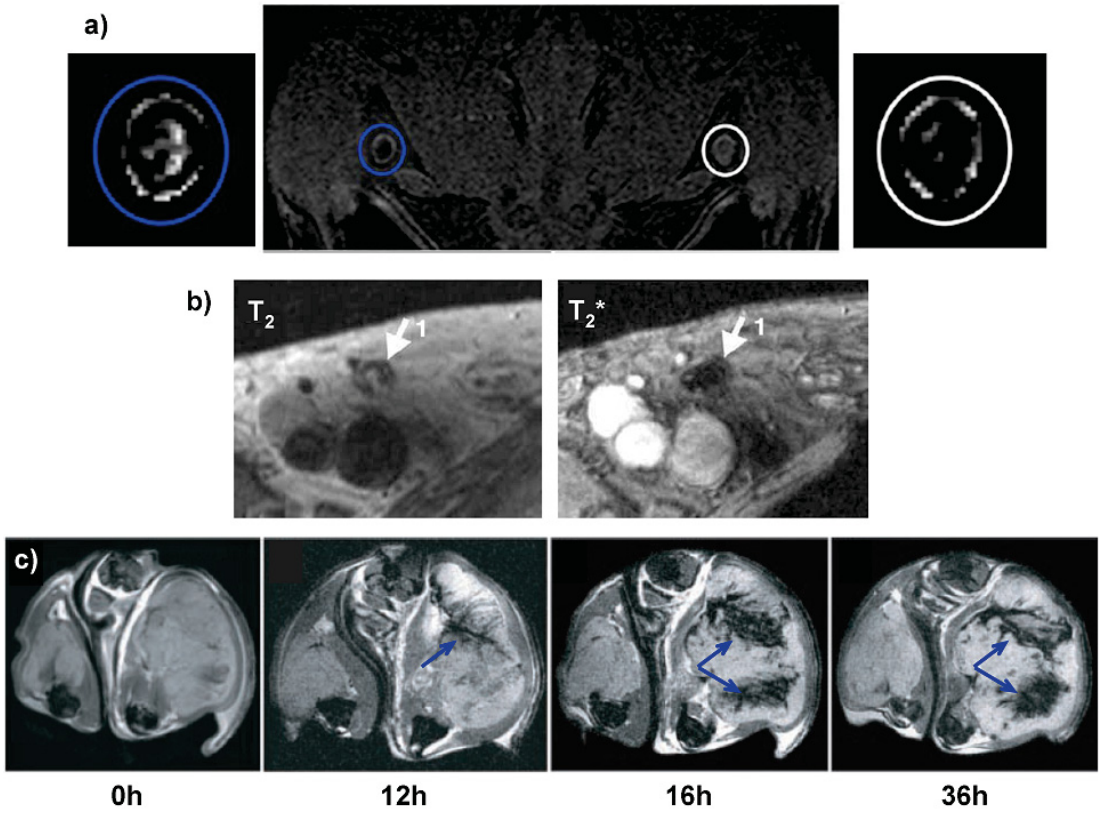

Figure 1.3: Examples of cellular MRI of immune cells. (a) Detection of dendritic cells in a mouse lymph node using a gradient-mapping technique. Susceptibility maps generated from gradient-echo data show increased signal in the lymph node containing DCs (blue) compared to the control (white) [45]. (b) Detection of labeled dendritic cells in a human lymph node using T_2-weighted spin-echo (left) and T_2^*-weighted gradient-echo (right) imaging; in this case, both sensitivity and specificity are higher in the T_2^*-weighted gradient-echo image [46]. (c) Time course of SPIO-labeled T cells (indicated by blue arrows) homing to a tumor in a mouse model, with images before injection, after 12h, 16h, and 36h [3].

1.2 Quantitative Cellular MRI

Most of the techniques and results highlighted in the previous section involve the detection of labeled cells, but to properly monitor, compare and evaluate various cellular therapies, some kind of quantitative imaging is needed. Though many detection techniques can be adapted for quantification to some degree, these approaches often carry drawbacks or require assumptions that may not always be valid. The use of specialized quantitative techniques is highly preferred due to their increased range and specificity. This section will provide an overview of existing quantitative imaging techniques and their use in cellular imaging studies, and will discuss some applications where they could be particularly helpful.

1.2.1 Signal-Based Techniques

The most straightforward way to obtain quantitative information from an MR image is to assume that either the total volume of the area showing contrast (whether negative or positive) or the change in signal intensity with respect to the background tissue is directly proportional to some quantity of interest, such as the concentration of SPIO. Distinguishing the volume affected by SPIO from its surroundings can be done through manual or automatic segmentation of the image; the former is most effective where an entire organ contains SPIO throughout its volume, while the latter is useful in tissues containing many small clusters of SPIO [49, 50]. Once a volume is distinguished, its extent or intensity can be compared to similar tissue in an unaffected area or from a reference scan obtained before the introduction of contrast.

Under proper conditions the accuracy of such techniques can be fairly good; in one study, the relative intensity of a region containing labeled cells was used to estimate cell densities to within 15% [51], though this was not done *in vivo*. Examples of *in vivo* quantification using void-based measurements include monitoring of vaccine clearance in mice vaccinated with SPIO-labeled antigens [52], tracking of dendritic cell migration [44], monitoring of disease progression in mice with multiple sclerosis [49] and correlation with antibody expression in a model of post-surgical inflammation [50] (an example of which is shown in Figure 1.4(a)).

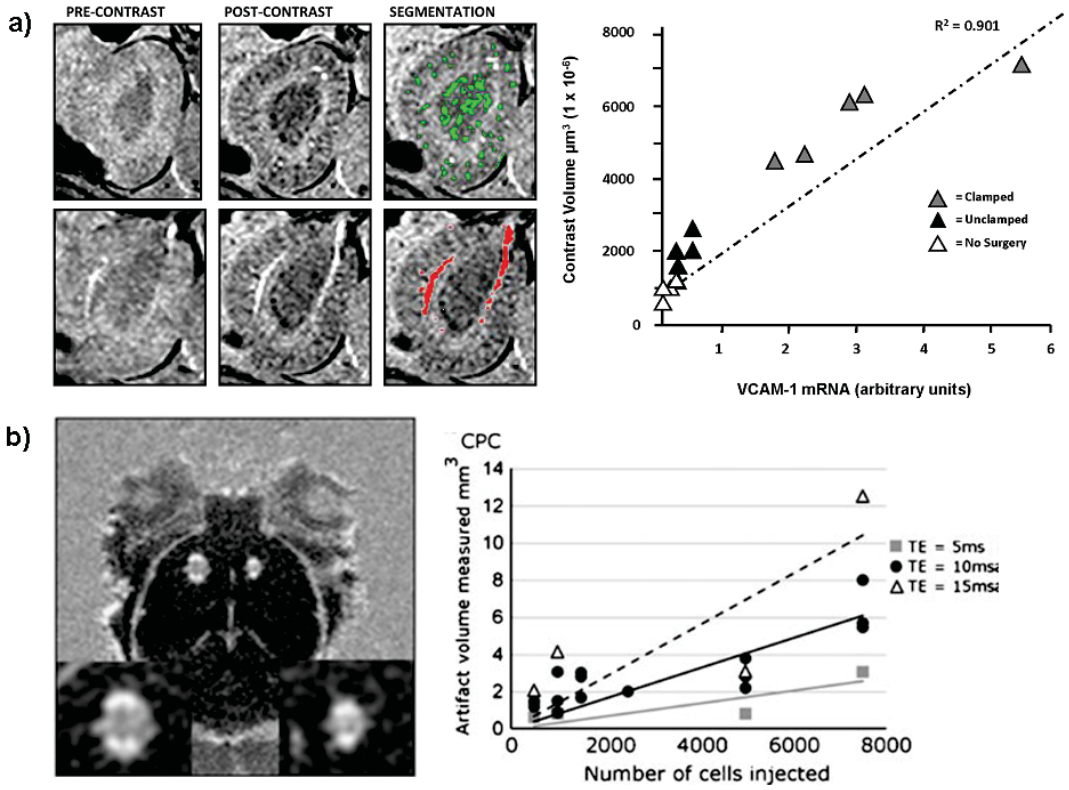

Figure 1.4: Examples of quantification using methods based on analysis of signal changes due to contrast agent. (a) Quantification based on automatic segmentation of signal voids in mouse kidney; areas showing significant negative contrast are highlighted in green. The total volume showing contrast is correlated with antibody expression measured *ex vivo* [50]. (b) Quantification based on a Combined Positive Contrast subtraction image, in which contrast volume correlates to the number of labeled cells injected into a mouse brain [53].

Signal-based analysis can be applied to positive-contrast images as well; both selective excitation of off-resonance MRI signal and suppression of on-resonance signal yield areas of positive contrast in rough proportion to the number of labeled cells present [24, 23]. However, very high concentrations of cells (10^7 - 10^8 cells/mL) were required to achieve these results at typical SPIO loading levels. The GRASP technique has also been used to quantify highly compartmentalized iron in macrophages, again at high concentrations [25]. A newer technique called Combined Positive Contrast subtracts a normal background image from the sum of two GRASP acquisitions, as shown in Figure 1.4(b), yielding images that can be used to successfully quantify smaller numbers of labeled cells [53]. The positive contrast produced by UTE techniques can also be used for quantification, though residual T_2 contrast may limit accuracy at high SPIO concentrations [54].

Limitations

Though effective in certain applications, signal-based analysis methods all share a number of drawbacks. The range of quantification is limited by the image contrast available; there must be enough SPIO to distinguish a region of interest from the background, and in the case of negative contrast methods, there must not be so much SPIO that no signal remains within the void.

A practical example of this limitation is illustrated in Figure 1.5. In this study, mice were vaccinated with a DepoVax [41] formulation in which the tumor antigen had been conjugated to 20 nm SPIO. In principle, this allows the clearance of the vaccine from the injection site to be calculated based on void volume and intensity [52], while also permitting dendritic cells which have taken up the vaccine to be visualized when they arrive at a draining lymph node. However, if the concentration of SPIO is low enough to permit quantification of the clearance rate through volume/intensity measurements of the depot, no significant changes in lymph node intensity are observed. High SPIO concentrations allow reliable visualization of contrast in the lymph node, but at such concentrations, the signal in the void is effectively zero, and remains so throughout the clearance process, such that no change can be quantified. This underscores the need for an MRI technique which provides more robust quantitative monitoring of SPIO concentrations over a large dynamic range.

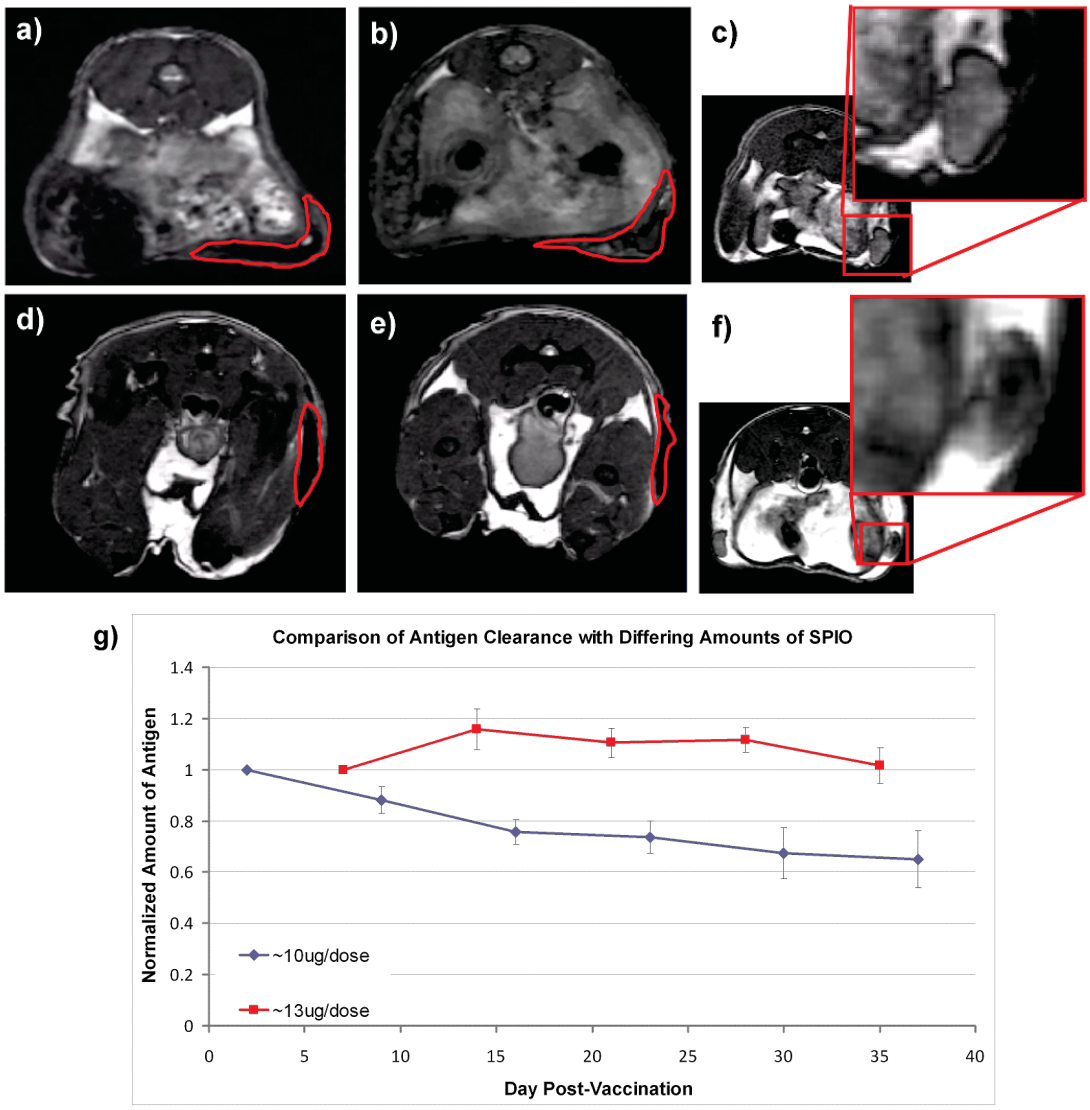

Figure 1.5: Quantification of SPIO based on signal void volume and intensity is often limited by the available contrast. (a) The red outline shows a depot of low-concentration SPIO-labeled vaccine in a mouse's flank, 1 week after injection. (b) Depot in the same mouse 4 weeks after injection, showing a significant difference in contrast. (c) No significant contrast in the draining lymph node is visible at this iron concentration. (d) Depot of a high-concentration SPIO-labeled vaccine, 5 days after injection. (e) Depot in the same mouse 33 days after injection, with no visible change in contrast. (f) The increased iron concentration allows negative contrast to be seen in the draining lymph node. (g) Clearance of SPIO-labeled vaccine based on signal intensity within the depot, relative to quantity at week 1. Due to insufficient signal, no change is seen at higher iron concentration (red).

This is not the only limitation of signal-based quantification. Signal-based analysis can also lack specificity, since contrast which is unrelated to SPIO (such as that caused by magnetic field inhomogeneities or natural variations in tissue composition) can be misidentified. Similarly, it is often assumed that any changes in contrast volume or intensity are solely due to changes in SPIO concentration, and that the background signal to which the contrast is compared will remain constant. However, other factors may influence both the region volume and the overall signal intensity, impacting the accuracy of quantification to an unpredictable degree. Finally, in many cases the concentration of SPIO is not measured directly, but correlated with another quantity of interest. Since many positive-contrast techniques measure distortions around SPIO-containing regions rather than the regions themselves, positive-contrast quantification is often an even more indirect measure of actual iron content.

Despite these limitations, signal analysis of either negative- or positive-contrast images remains a commonly used technique due to the simplicity of analysis and technical implementation.

1.2.2 Relaxometry Techniques

The other major category of techniques for SPIO quantification are those which allow relaxometry, or spatially resolved measurements of MRI relaxation times such as T_2 or T_2^*. Because the relaxation rates $R_2 = 1/T_2$ and $R_2^* = 1/T_2^*$ vary in direct proportion to the concentration of contrast agent in a particular area, relaxometry techniques which can accurately generate maps of these rates are being increasingly applied to the problem of quantitative cellular imaging.

Measuring R_2 or R_2^* generally requires a number of acquisitions with different contrast, which is altered by varying the echo time of the sequence (to be described further in Chapter 2). The relaxation rate is calculated based on an exponential fit to the signal decay in a particular region. Since R_2^* is more sensitive to changes in SPIO than R_2, often by 1-2 orders of magnitude [5], the majority of quantification studies use gradient-echo sequences that produce R_2^* contrast, such as the example shown in Figure 1.6(a). However, there is some benefit to R_2 quantification as well. For example, as demonstrated in Figure 1.6(b), R_2 and R_2^* maps of the same anatomy can distinguish free SPIO from SPIO within cells *in vivo*, since the free SPIO has a

similar R_2 but a much smaller R_2^* [55, 56]. This could be helpful in monitoring cell death, upon which a cell's SPIO is released into the surrounding tissue.

Though a traditional single-echo scan can be performed repeatedly with different echo times, it is more common to use a multiple-echo technique which offers increased imaging speed by collecting data at many echo times during a single scan. Examples of such techniques include Gradient Echo Sampling of FID and Echo (GESFIDE [58]) and Gradient Echo Sampling of Spin Echo (GESSE [59]), which repeatedly acquire data following signal excitation or during a spin echo. As an added advantage, both of these techniques permit simultaneous R_2 and R_2^* quantification, though this is not always fully exploited. At low concentrations of SPIO it is also possible to perform simultaneous quantification of R_2 and $R_1 = 1/T_1$ by using a multiple-echo version of b-SSFP [60], which is more robust against large-scale field inhomogeneities.

The range of SPIO concentrations that can be successfully quantified with these techniques depends on the number of images collected (which impacts the accuracy of the calculated relaxation rate) as well as the minimum time between signal excitation and collection of the first image, which must be kept short to retain signal near high SPIO concentrations. The number and spacing of points that can be achieved with a multiple-gradient-echo sequence depends on available hardware, as does the minimum echo time. Further decreasing the echo spacing entails repeating the scan with slightly different echo times, a strategy that has been employed to quantify R_2^* up to 1000 s^{-1} at the cost of increased scan duration [61]. Another refinement was demonstrated by Seevinck et.al. [62], who quantified R_2^* values up to 1500 s^{-1} by combining shifted echoes with an estimate of the initial signal based on background intensity. Though this resulted in a high upper limit of quantification, it also introduces bias into the fits and requires that the background signal be relatively uniform. Simultaneous measurements of R_2^* and R_2 have not been reported with this type of sequence.

Because relaxometry computes a physical quantity without regard to region volume, overall signal intensity or a background reference, these techniques generally have improved accuracy compared to signal-based methods and would be more desirable for longitudinal or multi-site studies. The measured relaxation rates are proportional to the total iron load, making quantification straightforward. However, few examples exist of relaxometry directly applied to the quantification of labeled cells *in*

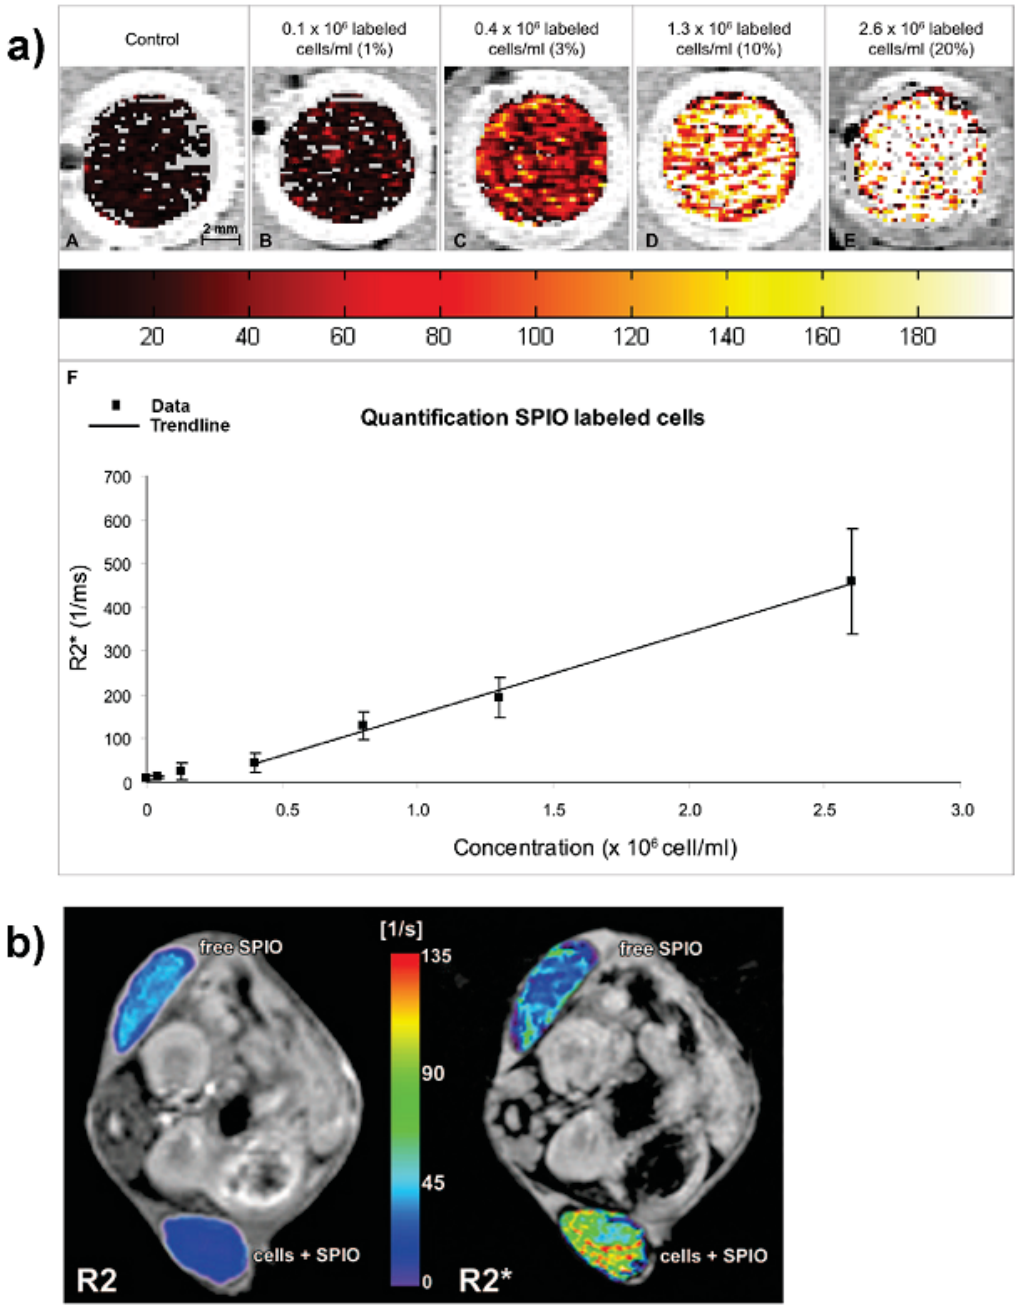

Figure 1.6: Examples of quantification using methods based on relaxometry. (a) Quantification of SPIO-labeled mesenchymal stem cells within cartilage defects, imaged *in vitro* [57]. (b) Simultaneous R_2 and R_2^* maps of a mouse containing pellets of free SPIO (upper) and SPIO loaded within carcinoma cells (lower); significant differences in relaxation behavior are seen despite identical SPIO concentrations [55].

vivo, potentially due to the increased technical difficulty of relaxometry techniques or a perception that signal-based analysis is adequate for existing applications.

1.3 Project Overview and Hypotheses

While a number of MRI techniques have been proposed and applied to the clinically relevant problem of imaging cells labeled with SPIO, techniques capable of quantitative cellular imaging are rare. Relaxometry is a promising avenue, but quantification of high iron concentrations remains challenging. In this work, we propose the use of a technique called TurboSPI which, though initially developed for imaging non-biological materials, is also well suited to imaging systems of labeled cells.

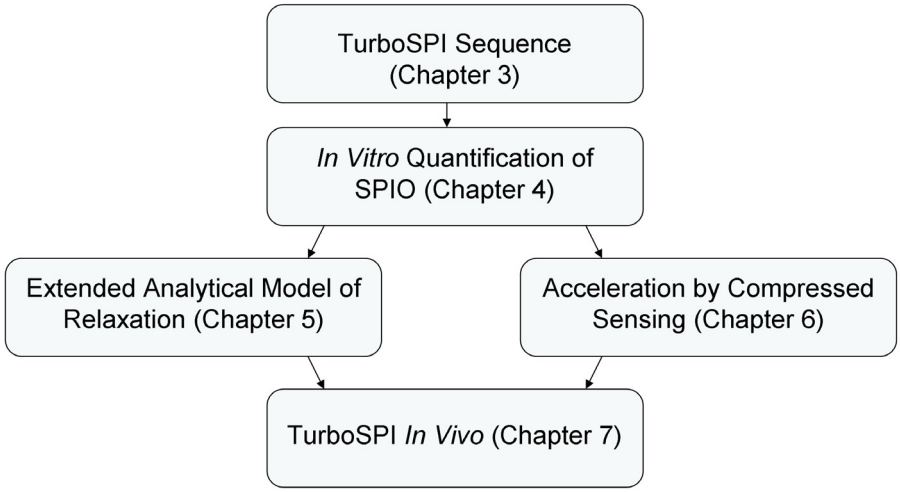

Figure 1.7: Overview of significant components in the project described in this work.

The process of developing and applying TurboSPI to cellular imaging is illustrated in Figure 1.7. After implementation and optimization of the TurboSPI sequence, as described in Chapter 3, its ability to quantify SPIO *in vitro* was tested, as shall be demonstrated in Chapter 4. TurboSPI acquires relaxometry data at far higher sampling rates than most gradient-echo methods, which we believe will permit accurate R_2^* quantification of micron-sized iron oxide particles and labeled cells over a wide dynamic range.

TurboSPI also allows exploration of MRI signal behavior near SPIO at a temporal resolution that has not been previously demonstrated. Physical theories describing

the evolution of the MRI signal may therefore be tested and verified. In the case of relaxation near SPIO, experimental data obtained using TurboSPI shows signal behavior that visibly differs from theoretical predictions, even in the case of SPIO-loaded cells. Since no existing model was suitable for fully describing this behavior, a new analytical model was developed to predict the complete signal time course near systems of spherical magnetic particles. In Chapter 5, this model will be compared with simulated and experimental data, and approximate forms will be computed which we hypothesize to be suitable for curve fitting and extracting estimates of relevant physical parameters. For example, it may be possible to separate the contributions of cell density and cell loading to the overall iron concentration.

Our implementation of TurboSPI retains signal from stimulated echoes, using a novel automatic calibration routine to avoid artifact formation and improve imaging speed. However, even with this modification, acquisition times are still too slow for most *in vivo* applications involving animals under anesthesia. Chapter 6 of this thesis demonstrates the use of a signal processing technique called Compressed Sensing which allows significant acceleration of TurboSPI data collection. As a purely phase encoded sequence it was anticipated that TurboSPI would be particularly amenable to compressed sensing, and that prior information could be used to help constrain the image reconstruction algorithm and further improve performance. High resolution 3D TurboSPI images should be capable of acceleration by factors of 15 or more without significant loss of image quality; this is a remarkably high acceleration factor for an image acquired with a single-channel RF coil.

With these refinements to imaging speed and quantification ability, TurboSPI should be capable of quantitative imaging of cellular density in living organisms, and we will conclude this work in Chapter 7 with a demonstration of iron detection and R_2^* relaxometry in a mouse tumor model. This represents the first application of the TurboSPI sequence to cellular imaging *in vivo*.

1.3.1 Contributions to Research

The TurboSPI sequence was initially developed by Steven Beyea [63], and its potential application to cellular imaging was proposed by Steven Beyea and Chris Bowen. Implementation and testing of the TurboSPI sequence on the Varian console

was performed by James Rioux, as was the development of a calibration routine for improving image quality, based on earlier work by Bowen and colleagues [64].

For testing *in vitro* quantification, samples were prepared by James Rioux and, in the case of the SPIO-loaded cell samples, Kim Brewer, who was responsible for growing and labeling the cells. All data used for quantification was collected by James Rioux, who also performed all of the data analysis including the writing of Matlab programs for R_2^* mapping.

The analytical model for spin-echo relaxation near spherical perturbers, including the full integral forms and polynomial approximations, was developed by James Rioux based on the existing cylindrical model of Kiselev and Posse [65]. Valerij Kiselev provided clarification on several formulae and methods used to devise the earlier model, and verified many of the calculations. Relaxation behavior was simulated via Monte Carlo methods using Matlab code provided by Steve Patterson, which was modified for systems of spherical objects by James Rioux.

The modifications to the TurboSPI pulse sequence required for the implementation of Compressed Sensing were carried out by James Rioux, as was all data acquisition using the accelerated TurboSPI sequence. The reconstruction algorithm was based on that of Vaswani and Lu [66], and modified for iterative time series reconstruction by James Rioux. The *in vivo* imaging experiment was performed by James Rioux with animal handling assistance from Drew DeBay, Iulia Dude and Kim Brewer, and using cells prepared by Kim Brewer.

Both the calibration procedure and the demonstration of *in vitro* quantification with TurboSPI have recently been published in the Journal of Magnetic Resonance [67], and these results are reproduced in Chapters 3 and 4 in accordance with Elsevier's policy on intellectual property rights[2]. Other publications concerning the analytical model and compressed sensing acceleration are currently in preparation.

[2]http://www.elsevier.com/wps/find/authorsview.authors/rights

Chapter 2

Background

2.1 Magnetic Susceptibility in Materials

Before proceeding to an overview of the physics of magnetic resonance imaging, it will be instructive to briefly consider how different materials respond to the presence of an external magnetic field. In particular, this response is the source of the contrast produced by SPIO which enables cellular imaging with MRI.

Fundamentally, magnetism in materials arises from the electrical charges within atoms and molecules. Any charged particle with angular momentum, either intrinsic (i.e. spin) or due to motion (such as an orbiting electron), will have a magnetic dipole moment. In most materials, in the absence of an external magnetic field, these dipoles are randomly oriented and there is no net magnetization. However, if an external field is applied, the dipoles will interact with the field and the material will become magnetized to some degree.

The resulting magnetization is typically proportional to the applied field,

$$\vec{M} = \chi \vec{B} \tag{2.1}$$

with the proportionality constant χ the *magnetic susceptibility* of the material. The magnitude and sign of χ depend on the type of material and, in particular, the nature of its magnetic dipoles.

Consider first the case of a dipole produced by orbital angular momentum, which can be modeled as a charged particle moving in a circle with the centripetal acceleration provided by the electrostatic interaction [68, p.261]. If an external magnetic field is applied, the moving electron will experience an additional force that alters its dipole moment in a way that opposes the applied field, regardless of its orientation. Summed over the entire material, the result is a bulk magnetization in the direction opposite the applied field. This is referred to as *diamagnetism*, and typical diamagnetic materials have negative susceptibilities on the order of 10^{-9} to 10^{-5}.

Diamagnetism is a property of all atoms, but is only the dominant response when there are no unpaired electrons and the overall spin of the molecule is zero. In atoms with unpaired electrons, the dipole moment μ experiences a torque

$$\vec{\tau} = \vec{\mu} \times \vec{B} \tag{2.2}$$

which tends to align the dipole with the external field B. Again, over the entire ensemble of atoms the result is a net magnetization parallel to the external field. This *paramagnetic* behavior is characterized by a positive susceptibility, with a magnitude typically on the order of 10^{-5} to 10^{-2}, such that the diamagnetic response is no longer evident.

In both cases, the removal of the external magnetic field will remove the material's magnetization, as the dipoles return to their random alignment. However, if the interactions between neighboring nuclei are sufficiently strong, dipoles will tend to retain their alignments even after the external field is removed, forming extended regions (or magnetic domains) where all of the dipoles are aligned in parallel. Applying an external magnetic field to such a material will further align the domains, to the extent that a net magnetization will remain even when the external field is absent. Such a material is a *ferromagnet*, with a susceptibility that generally varies with the external field but can be much larger than 1.

The iron oxide contrast agents discussed in Chapter 1 are examples of *superparamagnetic* materials. Superparamagnetism is a property of nanometer-size crystals of ferromagnetic material, each of which forms a single magnetic domain and can be treated as an individual magnetic dipole. Such crystals are possible when the formation of the walls which separate adjacent domains (called Bloch walls) is thermodynamically unfavorable; in iron oxide, for example, this occurs on a length scale of 15-50 nanometers, depending on the precise composition [9].

Because the crystals are subject to thermal randomization, including the process of Neel relaxation which can randomly reverse the orientation of an entire magnetic domain [69], superparamagnetic materials have no net magnetization in the absence of an applied field. However, when the individual domains are influenced by an external field, they align to an extent comparable to that observed in ferromagnets, with similar susceptibility.

2.2 Nuclear Magnetic Resonance

Magnetic resonance imaging is founded on the principles of Nuclear Magnetic Resonance (NMR), which was first described by Bloch and Purcell in 1946 [70, 71]. NMR lends itself well to classical descriptions under many circumstances, and for this overview of NMR physics we will use these more intuitive representations when possible. However, since NMR is fundamentally a quantum mechanical phenomenon it will sometimes have to be considered in that framework.

2.2.1 Classical Description

While the previous section dealt with magnetic effects on electrons, atomic nuclei can also act as magnetic dipoles if they have a net angular momentum $\vec{L}$. The source of this angular momentum is quantum-mechanical and will be considered shortly, but regardless of its origin, this angular momentum gives the nucleus a magnetic dipole moment $\vec{\mu}$,

$$\vec{\mu} = \gamma \vec{L} \tag{2.3}$$

The gyromagnetic ratio γ is characteristic of a given nucleus; for example, the hydrogen nucleus ^1H has $\gamma = 2.68$ x 10^8 rad/s/T (or $\frac{\gamma}{2\pi} = 42.58$ MHz/T).

In a system at thermal equilibrium and in the absence of any external magnetic fields, the ensemble of nuclei are randomly oriented and there is no net magnetization. If an external field $\vec{B_0}$ is applied, just as in the case of electrons, each nucleus will experience a torque

$$\vec{\tau} = \vec{\mu} \times \vec{B_0} \tag{2.4}$$

However, because the torque applied to a system equals the rate of change of its angular momentum, this implies

$$\frac{d\vec{L}}{dt} = \gamma \vec{L} \times \vec{B_0} \tag{2.5}$$

This equation of motion describes a nucleus whose magnetic moment precesses around the applied field $\vec{B_0}$ with an angular frequency

$$\omega_0 = \gamma B_0 \tag{2.6}$$

This frequency is called the *Larmor frequency* of the particular nucleus at a given field strength; for example, for a proton in a 3 Tesla magnetic field the Larmor frequency is approximately 128 MHz, which is in the radio-frequency (RF) region of the electromagnetic spectrum.

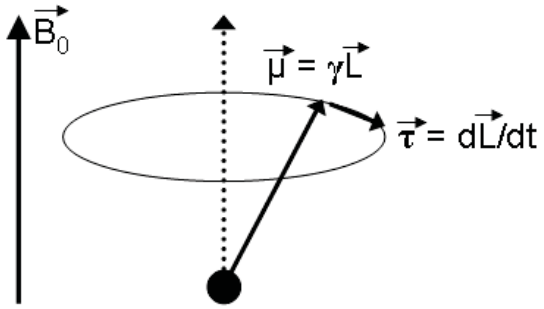

Figure 2.1: Larmor precession of a nucleus in a magnetic field. The torque applied by the external field $\vec{B_0}$ rotates the magnetic moment $\vec{\mu}$ around the $\hat{B_0}$ axis at a frequency $\omega_0 = \gamma B_0$.

At typical field strengths, the tendency of nuclei to precess about and align with the external magnetic field is dominated by thermal effects which randomize molecular motions. Nonetheless, the presence of the external field does produce a small net magnetization in the direction of the field, though without any transverse component due to cancelation of phase throughout the ensemble. Rather than considering individual nuclei we can now consider manipulation of this net magnetization $\vec{M}$.

Our eventual goal is to measure $\vec{M}$ experimentally, but it must first be moved away from the longitudinal axis to distinguish it from $\vec{B_0}$. This can be done by the application of a magnetic field perpendicular to $\vec{B_0}$, to change the effective axis around which the magnetization precesses. The nature of this perturbing field is best considered in the so-called rotating reference frame, whose z-axis coincides with the stationary (laboratory) frame, but whose x' and y' axes rotate at a frequency ω in the direction of precession, as illustrated in Figure 2.2.

In this frame, if the rotation of the axes is described by a vector $\vec{\omega} = \omega\hat{z}$, the new equation of motion in the rotating frame is

$$\left(\frac{d\vec{L}}{dt}\right)' = \gamma\vec{L} \times \vec{B}_{eff} \tag{2.7}$$

with the effective magnetic field

$$\vec{B}_{eff} = \vec{B}_0 - \frac{1}{\gamma}\vec{\omega} \tag{2.8}$$

Evidently, if the rotating frame is chosen to rotate at the Larmor frequency, with $\omega = \omega_0$, the effective magnetic field is zero and the magnetization is stationary, even if it has a component in the transverse plane.

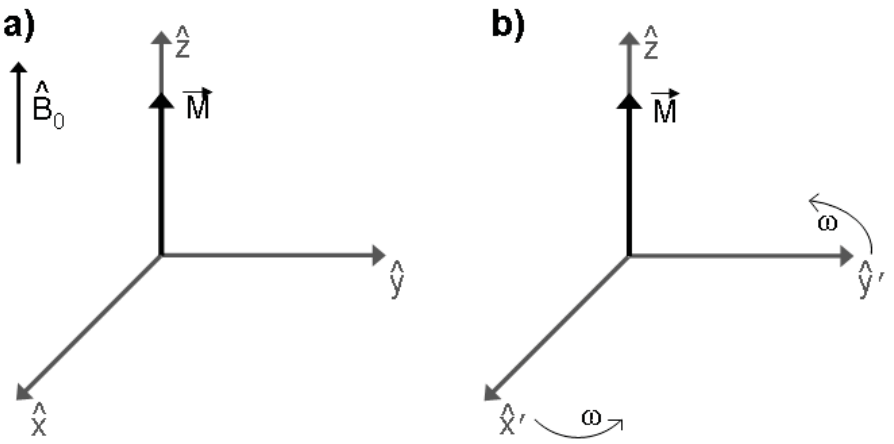

Figure 2.2: The rotating frame of reference for an NMR experiment. (a) In the laboratory frame, magnetization created by the external field B_0 aligns with the $\hat{z}$ axis. Any component in the transverse plane would rotate at a frequency ω_0. (b) In a frame where the $\hat{x}'$ and $\hat{y}'$ axes rotate at ω_0, magnetization is stationary, and the effective field is zero.

We can now consider the application of a magnetic field $\vec{B}_1$ along one of the transverse axes in the rotating frame, say $\hat{x}'$. The magnetization will precess about that axis at a frequency γB_1, such that in a time t the magnetization will have moved away from the $\hat{z}$ axis and toward $\hat{y}'$ by a *flip angle* of

$$\theta = \gamma B_1 t \tag{2.9}$$

In the laboratory frame, the perpendicular field $\vec{B}_1$ corresponds to a magnetic field that oscillates at the Larmor frequency. Such a field can be applied by a suitably oriented transmitter coil, which is generally referred to as an RF coil due to the frequencies involved. The magnitude of $\vec{B}_1$ is typically much less than the main field $\vec{B}_0$; in many practical applications, B_1 is on the order of 10^{-6} T, which requires an *RF pulse* lasting only a few milliseconds to rotate the magnetization into the transverse

plane. For illustration we will assume a flip angle of 90 degrees, though different flip angles are often used.

The effect of the RF pulse in the laboratory frame can be regarded as a gradual widening of the angle at which the magnetization precesses, such that it traces a spiral path toward the x-y plane as shown in figure 2.3. This occurs only if the pulse frequency is at the Larmor frequency, or 'on-resonance'. This is the resonance alluded to in the term Nuclear Magnetic Resonance.

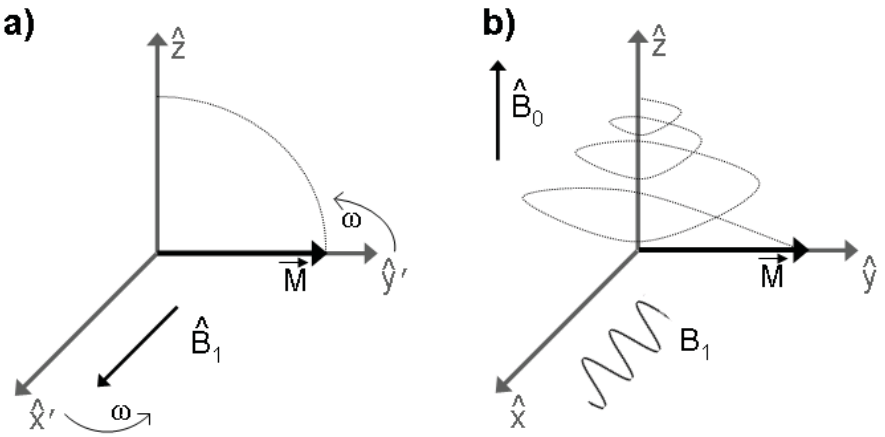

Figure 2.3: Manipulation of magnetization with an RF pulse. (a) In the rotating frame, applying a constant magnetic field B_1 along $\hat{x}'$ rotates the magnetization towards the $\hat{y}'$ axis. (b) In the laboratory frame, B_1 oscillates at the Larmor frequency, and the magnetization is rotated into the transverse plane in a spiraling path.

The time evolution of magnetization can be described by the equation

$$\left(\frac{d\vec{M}}{dt}\right) = \gamma\vec{M} \times \vec{B}_{ext} \tag{2.10}$$

where B_{ext} is the total external magnetic field, combining the static (B_0) and RF (B_1) contributions. This forms the core of the Bloch equations, which also include modifications to describe the return of the system to equilibrium, and which will be considered later.

2.2.2 Quantum Mechanical Description

The preceding classical description of NMR is intuitively appealing, and its macroscopic predictions are correct, but a clearer understanding of the precise physical

mechanisms can only be achieved through the use of quantum mechanics. For example, the magnetic moment of the nucleus is not due to physical rotation but to its intrinsic angular momentum, or spin. Any particle with a net non-zero spin $\vec{S}$, such as the spin-1/2 hydrogen nucleus ^1H, will have a magnetic moment

$$\vec{\mu} = \gamma \vec{S} \tag{2.11}$$

When placed inside a magnetic field $\vec{B_0}$, the nucleus acquires potential energy described by the Hamiltonian

$$H = -\vec{\mu} \cdot \vec{B_0} \tag{2.12}$$

If we constrain the external field to lie along a particular direction, such as $\hat{z}$, then

$$H = -\gamma B_0 S_z \tag{2.13}$$

with S_z an appropriate spin matrix. For example, in the case of a spin-1/2 nucleus, S_z is proportional to the Pauli spin matrix σ_z. This matrix has two eigenstates with eigenvalues $\pm\frac{\hbar}{2}$, where $\hbar$ is Planck's constant, $\hbar = \frac{h}{2\pi} = 1.067 \times 10^{-34} J \cdot s$. These eigenstates correspond to "spin-up" and "spin-down" orientations, in which the magnetic moment is aligned parallel or anti-parallel to the external field B_0. By inspection of eqn. 2.13 the two eigenstates of H are the same as those of S_z, each with a different energy:

$$\begin{aligned} E_+ &= -\hbar\gamma B_0/2 \quad \text{for the spin-up state } |+\rangle \\ E_- &= +\hbar\gamma B_0/2 \quad \text{for the spin-down state } |-\rangle \end{aligned}$$

This distinction of otherwise degenerate states due to the presence of a magnetic field is referred to as Zeeman splitting [72, p.308]. The energy difference between the Zeeman states is

$$\Delta E = \gamma \hbar B_0 = \hbar \omega_0 \tag{2.14}$$

The Larmor frequency ω_0 derived in the previous section appears once again, but with a different significance; now it relates to the energy required to make a nucleus transition from one state to the other.

Because the spin-up state has lower energy, more nuclei will tend to populate that state, but thermal effects will ensure that the higher-energy state is populated

as well. The populations of these states can be determined with Maxwell-Boltzmann statistics, which describe ensembles of nuclei in thermal equilibrium at relatively high temperatures. The expected number of nuclei N_s at a given energy E_s is

$$N_s = \frac{N}{Z} \exp^{-E_s/kT} \tag{2.15}$$

where N is the total number of nuclei, Z is the partition function, $Z = \sum_s \exp^{-E_s/kT}$, T is the temperature of the system and $k = 1.38 \times 10^{-23} J/K$ is Boltzmann's constant. The number of excess spins in the spin-up alignment is

$$\begin{aligned} N_+ - N_- &= N \frac{e^{\gamma\hbar B_0/2kT} - e^{-\gamma\hbar B_0/2kT}}{e^{\gamma\hbar B_0/2kT} + e^{-\gamma\hbar B_0/2kT}} \\ &= N\, tanh(\gamma\hbar B_0/2kT) \\ &\approx \frac{N\Delta E}{2kT} \end{aligned} \tag{2.16}$$

For typical experimental parameters of $B_0 = 3T$ and $T = 300K$, this gives an excess of 1 nucleus in the spin-up state for every 10^6 nuclei in the ensemble. This small population difference leads to an overall net magnetization along the direction of the applied field, as in the classical case.

Next, we consider the time evolution of this magnetization. The evolution of a particular quantum state $|\Psi\rangle$ is given by the time-dependent Schrödinger equation,

$$i\hbar \frac{\partial}{\partial t} |\Psi\rangle = H|\Psi\rangle \tag{2.17}$$

which has the solution

$$|\Psi\rangle(t) = |\psi\rangle e^{-\frac{iEt}{\hbar}} \tag{2.18}$$

provided that the stationary portion $|\psi\rangle$ of the state also satisfies the time-independent Schrödinger equation, $H|\psi\rangle = E|\psi\rangle$.

In a system of nuclei within an external field, each nucleus in the ensemble exists in a superposition of the $|+\rangle$ and $|-\rangle$ states, which evolves in time according to

$$|\Psi\rangle(t) = c_1|+\rangle e^{-\frac{iE_+t}{\hbar}} + c_2|-\rangle e^{-\frac{iE_-t}{\hbar}} \tag{2.19}$$

The complex coefficients c_1 and c_2 are normalized; $|c_1|^2 + |c_2|^2 = 1$. It should be noted that when the energies $E_\pm$ are substituted into the above expression, the factors of $\hbar$

will cancel, leading to classically relevant behavior. Indeed, it can be calculated [72, p.161] that the expectation value of the spin in the $\hat{z}$ direction is

$$\langle S_z \rangle = \langle \Psi | S_z | \Psi \rangle = \frac{\hbar}{2}(c_1^2 - c_2^2) \tag{2.20}$$

while that of the transverse components of the spin is

$$\begin{aligned} \langle S_x \rangle &= \hbar c_1 c_2 \cos(\omega_0 t) \\ \langle S_y \rangle &= \hbar c_1 c_2 \sin(\omega_0 t) \end{aligned} \tag{2.21}$$

Examination of these equations allows us to identify $\langle S_z \rangle$ with the longitudinal magnetization, while $\langle S_x \rangle$ and $\langle S_y \rangle$ describe transverse magnetization which rotates around the $\hat{z}$ axis at a frequency ω_0. This replicates the classical result but offers additional insight into the true physical origin of the Larmor precession; it reflects the time evolution of the wave function in the presence of the magnetic field Hamiltonian.

An RF pulse of intensity B_1, frequency ω and duration t can also be modeled as a perturbation of the system, which takes the form of a time-dependent potential [72, p.320]

$$V(t) = \frac{\gamma \hbar B_1}{2} \left(e^{i\omega t} |+\rangle\langle -| + e^{-i\omega t} |-\rangle\langle +| \right) \tag{2.22}$$

Each term of this potential contains an operator that changes a spin from one state to the other, hinting at the mechanism by which the pulse operates. Under the influence of such a potential, the coefficients c_1 and c_2 that describe the superposition of states evolve according to

$$\frac{d}{dt} c_n(t) = \sum_m V_{nm} e^{i(E_n - E_m)t/\hbar} c_n(t) \tag{2.23}$$

If we supply the initial condition $c_1 = 1$ at $t = 0$, i.e. the system consists of the excess nuclei in the spin-up state, the solution to the coupled system of equations is given by Rabi's formula:

$$\begin{aligned} |c_2(t)|^2 &= \frac{\gamma^2 B_1^2}{\gamma^2 B_1^2 + (\omega - \omega_0)} \sin^2(\sqrt{\gamma^2 B_1^2/4 - (\omega - \omega_0)^2/4})t \\ |c_1(t)|^2 &= 1 - |c_2(t)|^2 \end{aligned} \tag{2.24}$$

In the case of an on-resonance RF pulse we can make the assumption $\omega = \omega_0$, and these expressions reduce even further to

$$\begin{aligned} |c_1(t)|^2 &= cos^2(\gamma B_1 t/2) \\ |c_2(t)|^2 &= sin^2(\gamma B_1 t/2) \end{aligned} \tag{2.25}$$

The population of spins oscillates sinusoidally from spin-up to spin-down and back again. Substituting these expressions into our earlier results for the expectation values of S along the three axes, we find

$$\begin{aligned} \langle S_z \rangle &= \frac{\hbar}{2} cos(\gamma B_1 t) \\ \langle S_x \rangle &= \frac{\hbar}{2} sin(\gamma B_1 t) \cos(\omega_0 t) \\ \langle S_y \rangle &= \frac{\hbar}{2} sin(\gamma B_1 t) \sin(\omega_0 t) \end{aligned} \tag{2.26}$$

from which we recover the earlier expression for the flip angle,

$$\theta = \gamma B_1 t \tag{2.27}$$

It should be noted that this explanation does not directly address the origin of transverse magnetization, which to fully treat in a quantum mechanical framework is outside the scope of this work. Briefly, the RF pulse affects the system's density matrix, creating quantum coherences among the various states that make up the ensemble of particles. These coherences are what allow transitions between states, and also give rise to a non-zero expectation value in the transverse plane. While it is intuitively appealing to think of the RF pulse aligning the phases of randomly oriented individual nuclei, it is actually the global state of the system that changes in response to an RF pulse, and the behavior of individual nuclei cannot be inferred.

That said, for most practical applications this full quantum description is not necessary. Indeed, now that the central results of the classical picture of NMR have been confirmed with quantum mechanics, we are more justified in using a classical framework to explore NMR physics away from the atomic scale. A quasi-classical approach is generally more instructive when modeling microscopic phenomena such as relaxation of the excited signal, and in other cases we only need to consider the net macroscopic magnetization of the ensemble.

2.2.3 Signal Detection and Relaxation

We wish to measure the magnetization of a system of nuclei in an applied magnetic field, and thereby obtain information on the properties of that system. After the application of an RF pulse, the resulting magnetization in the transverse plane will

precess at the Larmor frequency. Maxwell's equations state that any time-varying magnetic field will produce an electric field, and thereby induce current flow into a suitably oriented receiver. The resulting voltage becomes the measured NMR signal $s(t)$ and, if measured immediately after the excitation pulse, is often referred to as a Free Induction Decay, or FID.

In practice it is often more instructive to examine the frequency spectrum $S(\omega)$ of the FID, which can be obtained through a Fourier transform:

$$S(\omega) = \int_{-\infty}^{\infty} s(t) e^{2\pi i \omega t} dt \tag{2.28}$$

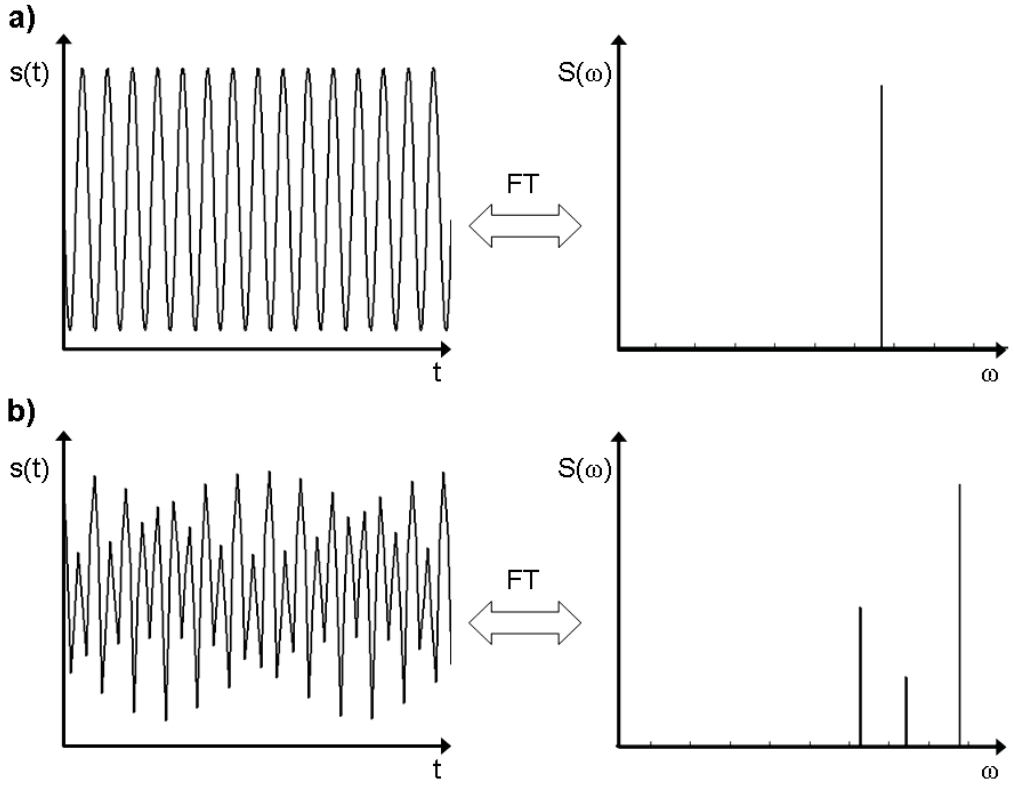

Figure 2.4: Examples of FIDs and NMR spectra, assuming no signal decay. (a) Sample of pure water. (b) Sample of pure ethanol (CH_3CH_2OH); each peak corresponds to a different proton chemical environment.

If all of the nuclei in the ensemble shared the same Larmor frequency and if the transverse magnetization did not decay with time, the resulting NMR signal would be a pure sinusoid of frequency ω_0, the Fourier transform of which would be a single sharp peak at ω_0 (Figure 2.4a). In the similar case of a sample containing protons

in different chemical environments, each species will have slightly different Larmor frequencies, because the electron cloud surrounding each nucleus partially shields it from the external field. The NMR signal then contains components at various frequencies, and the Fourier transform will yield a spectrum characteristic of the material (Figure 2.4b). This is the principle underlying *NMR spectroscopy*.

In practice, of course, even in the case of a sample of pure water, the magnetic environment of the nuclei is not homogeneous, and the NMR signal will not persist indefinitely. As discussed above, the RF pulse moves the spin system away from equilibrium in two ways: by inducing transitions of spins to the higher-energy spin-down state, decreasing the longitudinal magnetization of the system, and by introducing coherence that gives rise to transverse magnetization. After the RF pulse terminates, thermal effects begin to return the system to its equilibrium state, in which there is net magnetization only along the z axis, and the NMR signal created by the RF pulse vanishes. These effects are collectively referred to as *relaxation*.

On the microscopic scale, relaxation is driven primarily by interactions between neighboring nuclei as random motion brings their dipole fields into close proximity. These motions, chiefly translation and rotation of molecules with respect to one another, occur over a range of frequencies depending on the nature of the material, which can be characterized by a spectral density function, $J(\omega)$. The magnitude of $J(\omega)$ at particular frequencies affects the efficiency of the relaxation processes.

Restoration of the longitudinal magnetization requires that spins which were excited to the spin-down state emit a quantum of energy, $\hbar\omega_0$, and return to the spin-up state. Such emission can be stimulated by local field fluctuations due to molecular motions at frequency ω_0, and the emitted energy is transferred to neighboring molecules, which are referred to as the lattice. This *spin-lattice relaxation* causes recovery of longitudinal magnetization towards its equilibrium value, M_0, at a rate which depends on the existing magnetization M_z as

$$\frac{d}{dt}M_z = \frac{1}{T_1}(M_0 - M_z) \tag{2.29}$$

with T_1 the longitudinal relaxation time. The solution to this differential equation is

$$M_z(t) = M_0 + (M_z(0) - M_0)e^{-t/T_1} \tag{2.30}$$

and is depicted in Figure 2.5.

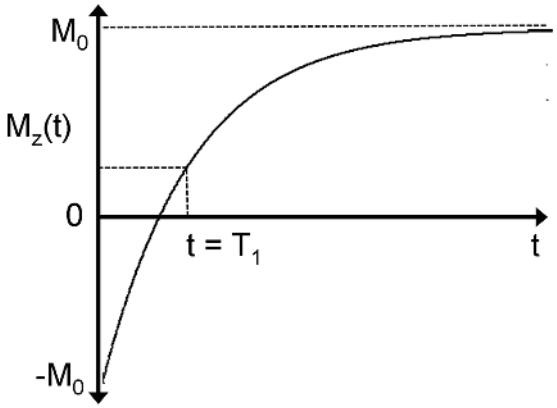

Figure 2.5: Spin-lattice relaxation. Longitudinal magnetization recovers exponentially with time constant T_1 toward the equilibrium value of $M_z = M_0$ from an initial value of $M_z(0) = -M_0$.

Spin-lattice relaxation will be most efficient (i.e. T_1 is short) in materials which have significant molecular motion at the Larmor frequency, which typically requires molecules of intermediate size or with somewhat restricted motion. Systems with small molecules that are free to move (such as pure water) or large molecules in rigid arrangements (such as polymers) will have inefficient relaxation and long T_1.

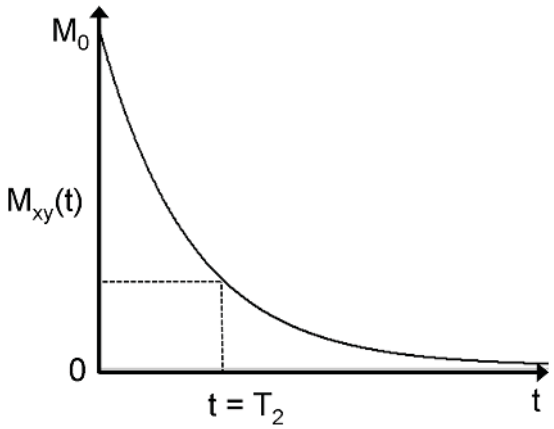

Figure 2.6: Spin-spin relaxation. Starting from a value of $M_{xy} = M_0$ following an RF excitation, transverse magnetization decays exponentially with time constant T_2.

A different mechanism leads to loss of transverse magnetization, which can be regarded as a measure of the phase coherence between the precessing particles. As molecular motions cause neighboring dipoles to interact, the magnetic field experienced by each spin will fluctuate, causing its precession frequency to increase or

decrease. As these interactions continue, the ensemble of nuclei will eventually lose phase coherence entirely, leading to a loss of transverse magnetization according to

$$\frac{d}{dt}M_{xy} = -\frac{1}{T_2}M_{xy} \tag{2.31}$$

which describes an exponential decay (see Figure 2.6) of the form

$$M_{xy}(t) = M_{xy}(0)e^{-t/T_2} \tag{2.32}$$

In contrast to T_1 relaxation, T_2 or *spin-spin relaxation* is facilitated by slow molecular motions, since fast motions will lead to a rapidly fluctuating field that, averaged over time, has no net effect on the precession of nearby spins. T_2 relaxation is slowest in fluids like pure water where molecular motion is unrestricted, and is fastest in systems with restricted motion, such as solids, that have a significant fraction of their spectral density near zero. Because spin-lattice relaxation processes also lead to loss of phase coherence, T_2 is never longer than T_1, and depending on the physical system, T_2 can potentially be much shorter than T_1.

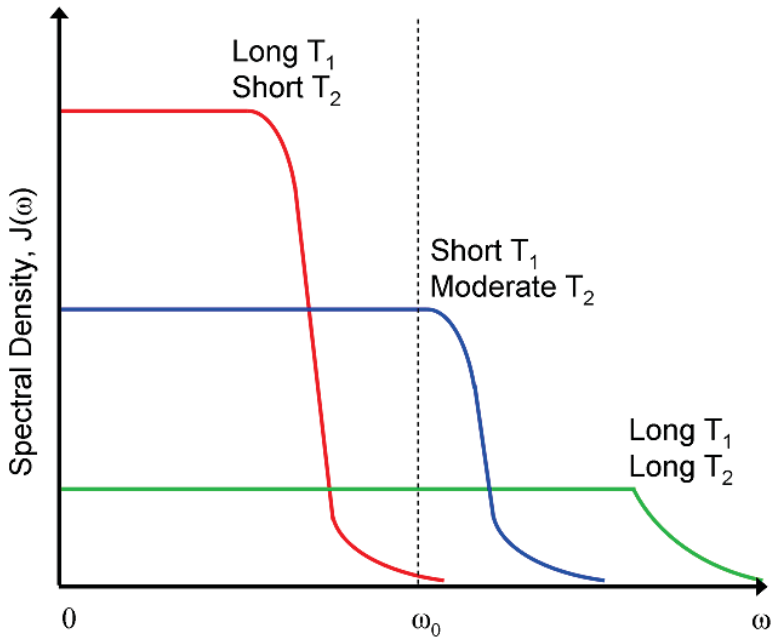

Figure 2.7: Effects of molecular motion on the T_1 and T_2 relaxation times. Systems with highly restricted motion (red line) will have long T_1 and short T_2. Systems with unrestricted motion (green line) will have long T_1 and long T_2. Intermediate motion regimes (blue line) lead to short T_1 and moderate T_2, though $T_2 \leq T_1$ in all cases.

Reference has already been made in the introduction to contrast agents, which are substances that have a significant effect on T_1 or T_2 in their vicinity. Adding even a small quantity of such an agent can reduce relaxation times dramatically. In such contexts, it is more common to refer to the relaxation rates $R_1 = 1/T_1$ and $R_2 = 1/T_2$, which typically vary linearly with the concentration of the contrast agent. The change in a particular relaxation rate produced by a given concentration of contrast agent is called the *relaxivity*, with typical units of $\mathrm{s}^{-1}/\mathrm{mM}$. T_1 contrast agents contain paramagnetic ions such as gadolinium or manganese, and protons in the vicinity of these ions experience a much stronger spin-lattice relaxation. T_2 contrast agents are often superparamagnetic substances like iron oxide, which significantly dephase protons in their vicinity. By definition a T_1 agent must shorten T_2 as well, but the effect on T_1 is relatively larger than in a T_2 agent.

T_1 and T_2 relaxation can be added to equation (2.10), to form the Bloch equations that describe the evolution and relaxation of magnetization, as

$$\left(\frac{d\vec{M}}{dt}\right) = \gamma\vec{M} \times \vec{B}_{ext} + \frac{M_0 - M_z}{T_1} - \frac{M_{xy}}{T_2} \tag{2.33}$$

As a result of T_2 relaxation, the ideal signal depicted in Figure 2.4 will decay with time, as shown in Figure 2.8. In the spectral representation this corresponds to a broadening of the line, i.e. the peak at ω_0 now has a non-zero width inversely proportional to the T_2 of the sample. In many cases this peak is further broadened by macroscopic field inhomogeneities, caused by perturbations of the main magnetic field due to the presence of the sample being imaged. It is customary to define an effective spin-spin relaxation time T_2^*, as well as the corresponding relaxation rate $R_2^* = 1/T_2^*$, to account for these larger inhomogeneities.

Though the precise definition of T_2^* depends on the physical system, and will be explored in more detail later for some cases of interest, for now we can give a general form that separates the total effective relaxation T_2^* into contributions from microscopic (T_2) and macroscopic (T_2') field inhomogeneities, which combine as

$$\begin{gathered} \frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \\ \text{or} \\ R_2^* = R_2 + R_2' \end{gathered} \tag{2.34}$$

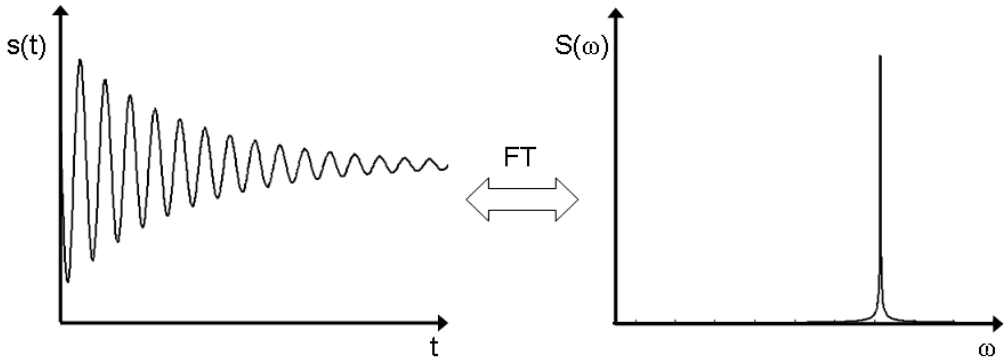

Figure 2.8: FID signal decay and line broadening due to relaxation. The induced NMR signal decays due to T_2^* effects, and the corresponding spectral lines are broadened by an amount proportional to the relaxation rate R_2^*.

Here it is assumed that T_2' relaxation also produces an exponential decay of signal, which is not always the case, but it is instructive to assume so in most cases.

2.2.4 Spin and Stimulated Echoes

The decay of the NMR signal due to T_2 relaxation is unrecoverable, since it results from interactions between nuclei that are driven by random molecular motions, but T_2' relaxation due to macroscopic magnetic fields can be reversed. In a classical picture, inhomogeneities in the magnetic field will cause neighboring nuclei to precess at different rates and lose phase coherence. However, since these fields are time-invariant, if the phases of the nuclei are reversed, phase coherence can be regained and signal recovered. This reversal of phase can be created by applying appropriate RF pulses to a system which already has transverse magnetization.

Most commonly, a pulse is applied which rotates the magnetization 180 degrees about either the $\hat{x}'$ or $\hat{y}'$ axis, as depicted in Figure 2.9. (The rotating frame introduced in section 2.2.1 is used for convenience.) Spins which had accumulated a positive phase with respect to those at ω_0 will now have a negative phase, and vice-versa, but they will continue to accumulate phase as they did before the 180 degree pulse, or *refocusing pulse*. If the refocusing pulse is applied at a time TE/2 after the excitation pulse, then at time TE the spins should have regained their original phase coherence (except for the irreversible losses from T_2 decay). This is called a *spin echo*, and the time TE is the echo time.

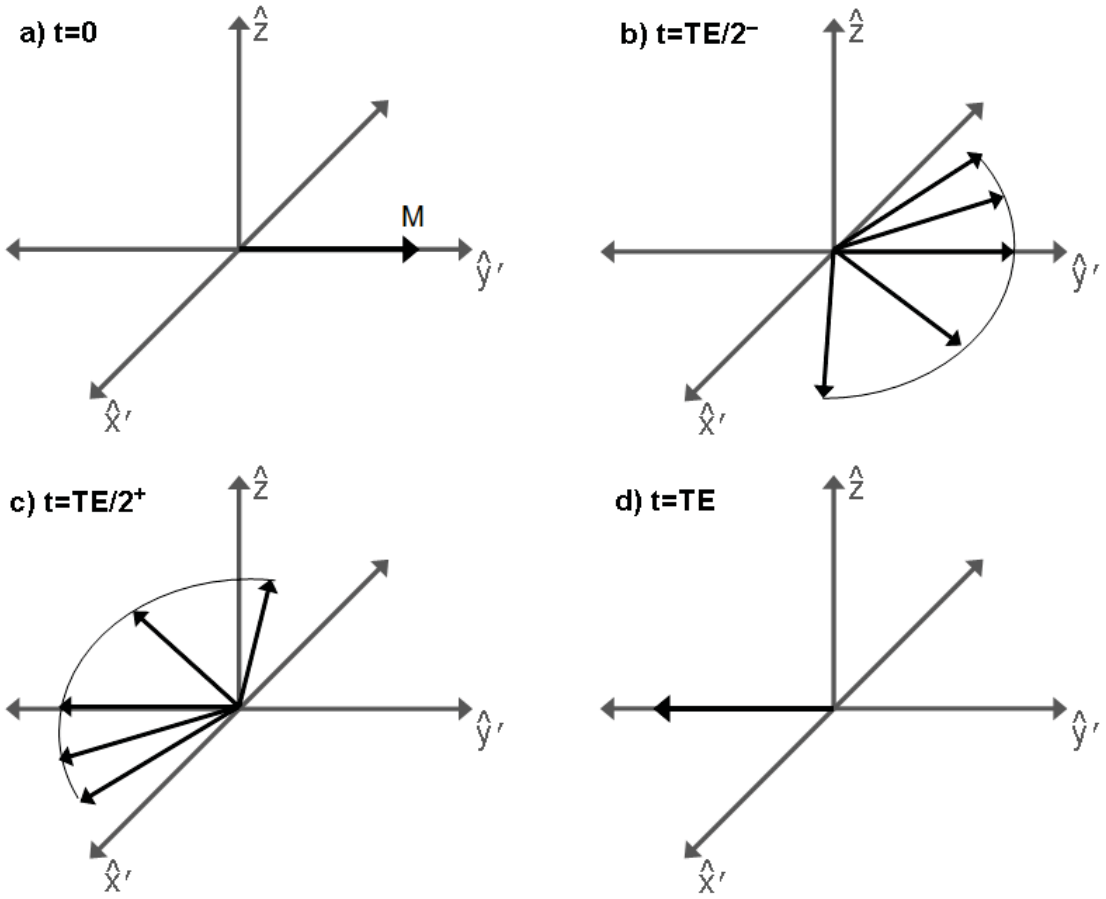

Figure 2.9: Formation of a spin echo after a 180 degree refocusing pulse. (a) Magnetization is initially coherent ($M = M_0\hat{y}'$). (b) After evolving for a time TE/2, magnetization dephases in the transverse plane. (c) An RF pulse rotates the magnetization 180 degrees, about $\hat{x}$ in this example. (d) The magnetization rephases along $-\hat{y}$ at time TE, forming a spin echo.

Spin echoes are useful in the study of systems with significant field inhomogeneities (and therefore short T_2'), since the NMR signal will decay too quickly after the excitation pulse for it to be collected, but the signal can be mostly recovered with a spin echo. Repeated applications of a 180 degree pulse can be used to continually refocus the magnetization, creating a train of echoes whose envelope decays with T_2, making this technique effective for measuring the T_2 of a material. The most common method for performing this measurement is to apply a 90 degree pulse along one axis (say $\hat{x'}$) and refocusing pulses on the perpendicular axis ($\hat{y'}$). This set of pulses self-compensates for small errors in the flip angles of the refocusing pulses, and is known as a CPMG sequence after its inventors: Carr, Purcell, Meiboom and Gill [73, 74].

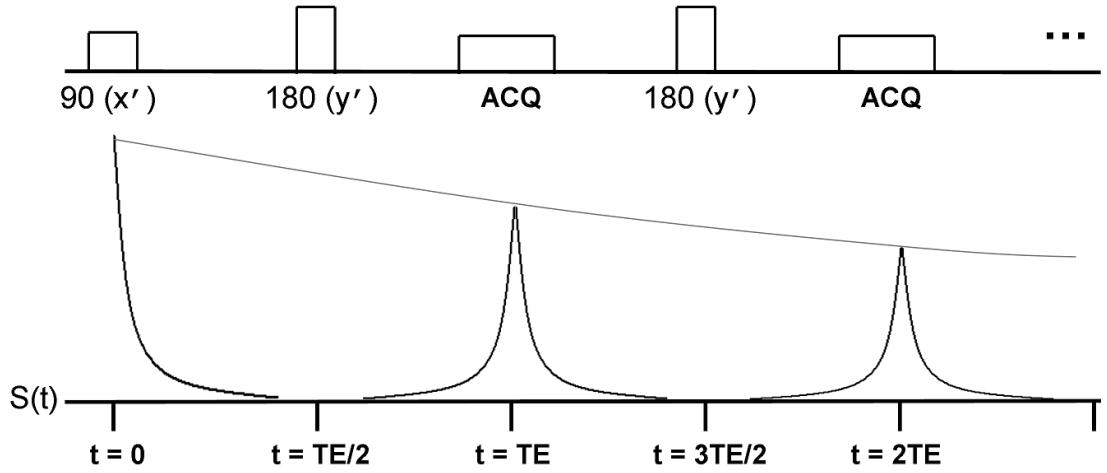

Figure 2.10: CPMG echo train for T_2 measurement. The set of 180° RF pulses, applied at times TE/2 + $n\cdot$ TE, leads to the formation of spin echoes (black lines) at times t = $n\cdot$ TE, with amplitudes that decay as a function of T_2 (gray line).

A slightly different form of echo can be created by applying a second 90 degree pulse at time TE/2 after the initial excitation instead of a 180 degree pulse. As shown in Figure 2.11, this rotates the transverse magnetization back towards the longitudinal axis, though components of the magnetization remain in the transverse plane. These transverse components will actually form a partial echo (called a Hahn echo [75, p.486]) at time TE as in the case of a spin echo, though its amplitude is reduced with respect to the spin echo. The magnetization that was returned to the longitudinal axis will experience T_1 relaxation, and can be returned to the transverse plane with another 90 degree excitation, where it will form a *stimulated echo* at a

time TE/2 after the third pulse. This echo also has reduced amplitude with respect to a spin echo, but between the second and third pulses the magnetization decays with T_1, not T_2, which can be advantageous for systems with long T_1 but short T_2.

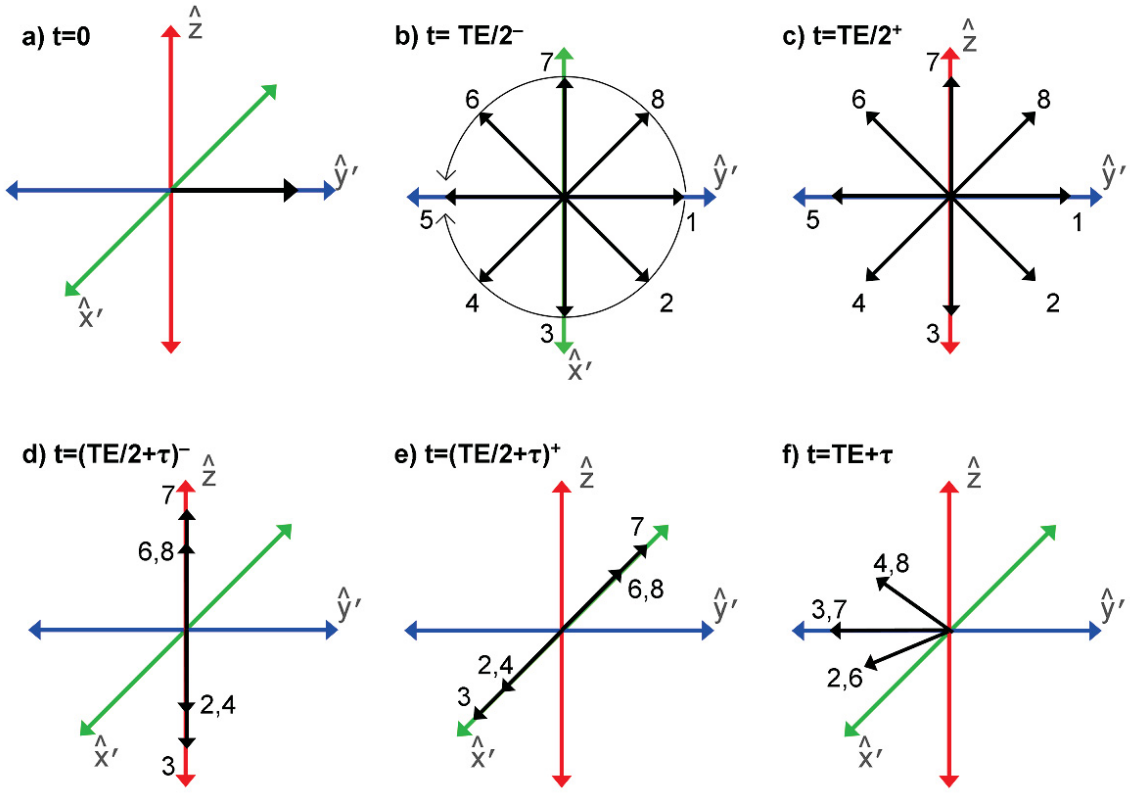

Figure 2.11: Formation of a stimulated echo after three 90 degree pulses. (a) Magnetization is excited onto the $\hat{y}'$ axis with a 90 degree pulse. (b) Magnetization dephases in the $\hat{x}'$-$\hat{y}'$ plane for a time TE/2. (c) A second 90 degree pulse rotates magnetization back into the $\hat{y}'$-$\hat{z}$ plane. (d) After a time τ, transverse components have decayed with T_2. (e) Remaining magnetization is rotated back to the transverse plane with a third 90 degree pulse. (f) At a time TE/2 after the third pulse, at time TE + τ, all magnetization is on one side of the $\hat{x}'$ axis, and a stimulated echo forms.

Since any RF pulse after the first can create an echo, the number of potential echoes increases geometrically with the number of RF pulses used, and often multiple echoes can form simultaneously [76]. As will be seen, in most experimental applications these echoes will have to be dealt with properly or they will contaminate the desired signal.

2.3 Magnetic Resonance Imaging

Having established the fundamentals of NMR we can now describe how they are used to obtain images of an object, and explore some of the practical limitations on the speed and resolution of MR imaging.

2.3.1 Spatial Encoding

It has already been mentioned that natural variations in the magnetic field experienced by the nuclei in a material can be used to examine, for instance, its chemical composition. The cornerstone of MRI is the use of deliberately applied variations of B_0 to encode other information about a sample, such as the distribution of nuclei in space. The simplest variation that can be applied is a linear gradient of the magnetic field's $\hat{z}$ component along an arbitrary direction,

$$\vec{G} = \frac{\partial B_z}{\partial x}\hat{x} + \frac{\partial B_z}{\partial y}\hat{y} + \frac{\partial B_z}{\partial z}\hat{z} \tag{2.35}$$

which introduces a spatial dependence into the Larmor frequency,

$$\omega(\vec{r}) = \gamma B(\vec{r}) = \gamma(B_0 + \vec{G}\cdot\vec{r}) \tag{2.36}$$

Such gradients are straightforward to generate with appropriate arrangements of current-carrying wires, called a *gradient coil*. Multiple independent sets of wires are generally integrated into a single unit, allowing the simultaneous application of gradients along all three spatial axes, as described above. For simplicity we will now consider a gradient applied only along $\hat{z}$. If the MRI signal is measured in the presence of such a gradient, it can be expressed as an integral over the density of spins in the $\hat{z}$ direction, weighted by their phases:

$$s(t) = \int_{-\infty}^{\infty} \rho(z) e^{i\omega(z)t} dz \tag{2.37}$$

The function $\rho(z)$ actually represents an effective density of spins, incorporating effects due to the sensitivity of the RF receiver coil and other system electronics, the longitudinal magnetization at t=0, etc. Nonetheless, for this treatment we assume $\rho(z)$ is proportional to the actual number of nuclei in a particular location. The effect of relaxation on the measured image will be added later.

Ignoring the constant modulation at ω_0, which in practice can be easily mixed out, the remaining signal is

$$s(t) = \int_{-\infty}^{\infty} \rho(z) e^{i\gamma G_z t z} dz \tag{2.38}$$

This underscores the Fourier relationship between $\rho(z)$ and $s(t)$, mediated by the Larmor equation that equates changes in position with changes in frequency. It also allows a convenient change of variables, first proposed by Mansfield and Grannell [77], which combines t and G into a single variable, k:

$$k = \frac{1}{2\pi} \gamma G t \tag{2.39}$$

In the more general case when G is not a constant it can instead be shown that

$$k = \frac{1}{2\pi} \gamma \int_0^t G(t') dt' \tag{2.40}$$

In either event, k is effectively a modified time coordinate, which can also be thought of as a *spatial frequency* due to its Fourier relationship with the spatial coordinate z:

$$s(k) = \int_{-\infty}^{\infty} \rho(z) e^{2\pi i k z} dz = \mathcal{F}[\rho(z)] \tag{2.41}$$

The generation of an MR image consists of acquiring data in the spatial frequency domain (or *k-space*) followed by a Fourier transformation back into image space, yielding an estimate of the spatial distribution of nuclei,

$$\rho(z) = \mathcal{F}^{-1}[s(k)] = \int_{-\infty}^{\infty} s(k) e^{-2\pi i k z} dk \tag{2.42}$$

This expression can be easily generalized to higher dimensions, describing the collection of 2D or 3D images through the acquisition of 2D or 3D k-space.

$$\rho(\vec{r}) = \int_{-\infty}^{\infty} s(\vec{k}) e^{-2\pi i \vec{k} \cdot \vec{r}} d\vec{k} \tag{2.43}$$

The relationship between the various domains involved in MR imaging is summarized in Figure 2.12.

2.3.2 Image Acquisition

To acquire an MR image of an object involves building up a representation of that object in k-space, which is done by collecting the MRI signal after appropriate

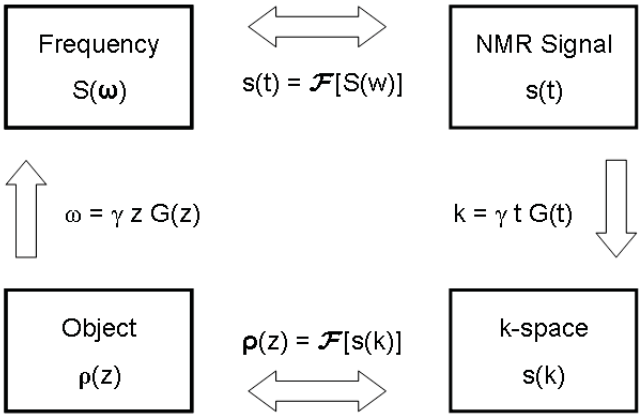

Figure 2.12: Summary of the relationships between the object being imaged and its k-space representation in the context of MRI. A gradient G applied to an object gives it a spatially-dependent Larmor frequency, which can be measured as the MRI signal and expressed in terms of k-space. The Fourier transform of k-space is an image of the object.

application of gradients. It is helpful to think of using a combination of gradients applied over time to traverse k-space according to equation (2.40). For example, if a constant gradient $G_x = G\hat{x}$ is applied following the excitation of signal, and multiple data points are sampled at intervals Δt, the acquired data will fall along a line in k-space starting at the origin,

$$k_{xn} = G_x \cdot n\Delta t, n = (0, 1, ...N_x/2) \tag{2.44}$$

which is illustrated in Figure 2.13a. This is called *frequency encoding*, with G_x the *readout gradient*. In most implementations a negative gradient lobe is used to move to the edge of k-space before acquisition, as shown in Figure 2.13b, such that an entire line is acquired at once,

$$k_{xn} = -G_x N_x/2\Delta t + G_x \cdot n\Delta t, n = (0, 1, ...N_x) \tag{2.45}$$

This forms a *gradient echo* at the center of k-space, as the signal which had been dephased by the negative rewinding lobe is rephased and then dephased by the readout gradient. The amplitude of the echo peak is determined by T_2^* decay, as opposed to T_2 decay as in a spin echo.

While it is possible to use frequency encoding in arbitrary directions and sample

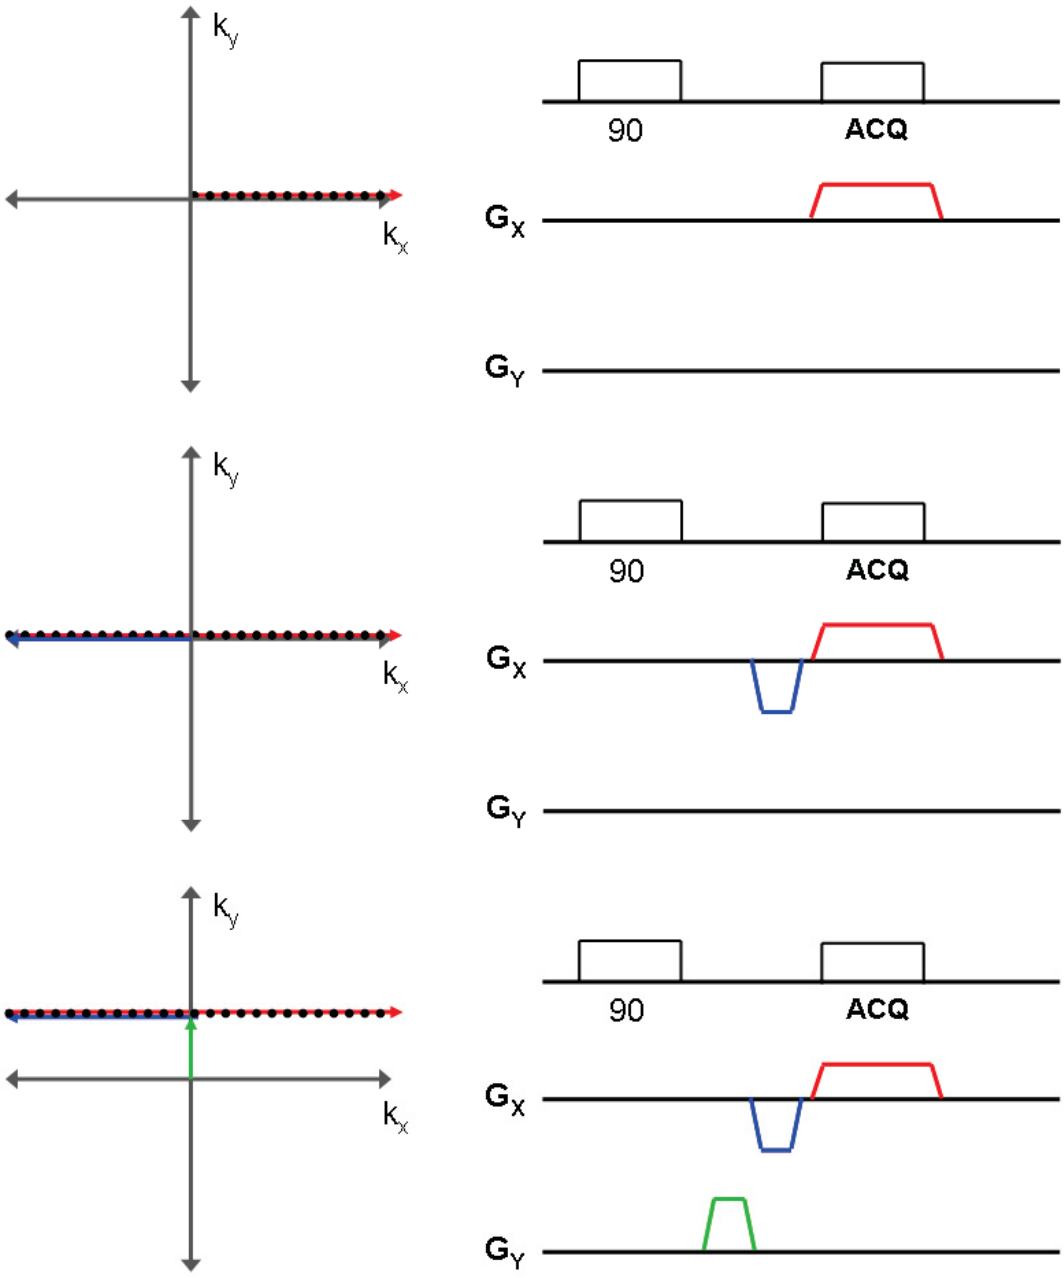

Figure 2.13: Examples of basic frequency and phase encoding. (a) A set of k-space points along $k_y = 0$ acquired using a readout gradient (red) in the $\hat{x}$ direction. (b) A rewinding lobe (blue) enables acquisition of an entire line of k-space during a gradient echo. (c) A phase encode gradient lobe (green) changes the k_y coordinate to acquire a different line.

k-space radially, with subsequent reconstruction of these projections, it is more conventional to sample k-space in a rectilinear fashion, collecting Cartesian data that can be processed with a fast Fourier transform. To change the k_y coordinate, a gradient lobe $G_y = G\hat{y}$ can be applied for a fixed duration prior to data acquisition. After the collection of a line with frequency encoding, the experiment can be repeated with different value of G_y to collect lines at different k_y positions:

$$k_{yn} = t \cdot n\Delta G_y, n = (-N_y/2, ... - 1, 0, 1, ...N_y/2) \tag{2.46}$$

G_y is referred to as a *phase encoding* gradient, and is illustrated in Figure 2.13c. If desired a second phase encoding gradient G_z could be added to allow 3D imaging.

Depending on the type of gradient used, the interval between points in k-space can be either $\Delta k = t\Delta G$ or $\Delta k = G\Delta t$. In either case, this interval is related to the field of view (FOV) in image space as

$$FOV = \frac{1}{\Delta k} \tag{2.47}$$

This is a consequence of the Fourier relationship between the two domains. Because the nominal resolution of the image Δx is simply the FOV divided by the number of acquired points, this is equivalent to

$$\Delta x = \frac{1}{N\Delta k} \tag{2.48}$$

which states that the resolution is inversely proportional to the extent of sampling in k-space. For a given FOV (and therefore Δk), obtaining a higher-resolution image (reducing Δx) therefore requires that more points be sampled in k-space, increasing the number of repetitions and the acquisition time. In a 3D acquisition with two phase encode gradients, the total number of repetitions is $N_y * N_z$, assuming a single line is acquired during each repetition.

As an alternative to true 3D acquisition, data at different positions along $\hat{z}$ can be acquired by using slice-selective RF pulses. To this point, we have assumed that the RF pulses used to excite magnetization are constant in amplitude and affect all nuclei in the ensemble, regardless of their Larmor frequency. This is in fact only true for broadband or *hard* RF pulses, and in most practical applications, a shaped or *soft* RF pulse is used. Such pulses are longer in duration, have an amplitude

that varies in time, and excite nuclei only within a limited range (or bandwidth) of Larmor frequencies. When a shaped pulse is applied with a gradient (called a *slice-select gradient*), the result is an excitation of nuclei within a specific slice instead of the entire sample.

At small excitation flip angles α, it can be shown [75, p.397] that the frequency profile excited by a shaped RF pulse is approximately equal to the Fourier transform of its amplitude. A common shaped RF pulse is the sinc pulse, which has an amplitude that varies as $\sin(t)/t$ and is usually truncated after a small number of zero crossings. Its excitation profile is approximately rectangular, with some deviations introduced by the truncation; these can be lessened by apodizing or smoothing the pulse. More sophisticated RF pulses can be designed through the use of the Shinnar-LeRoux (SLR) algorithm [78], which allows direct specification of parameters such as the transition bandwidth and passband ripple.

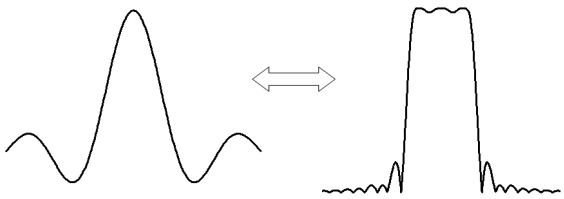

Figure 2.14: Example of a shaped RF pulse for slice-selective RF excitation. The sinc-shaped amplitude profile of the RF pulse (left) leads to an excited region which is approximately rectangular (right); magnetization outside this region is not significantly affected.

The strength of the applied gradient G_z determines the extent δz of the excited region, according to

$$\delta z = \frac{BW}{\gamma G_z} \tag{2.49}$$

where BW is the bandwidth of the pulse, which for most pulses is inversely proportional to its duration. Exciting thin slices then requires the use of a low-bandwidth pulse or, more commonly, the use of a strong slice-select gradient G_z.

The slice position is determined by changing the frequency offset of the RF pulse, so that instead of a carrier frequency at ω_0, it has a carrier frequency of

$$\omega_c = \gamma(B_0 + G_z \Delta z) \tag{2.50}$$

where Δz is the desired position of the slice center. Typically a number of slices are collected at a spacing that prevents overlap of the excited areas, allowing data to be collected from one slice while the nuclei in another slice are relaxing back to equilibrium.

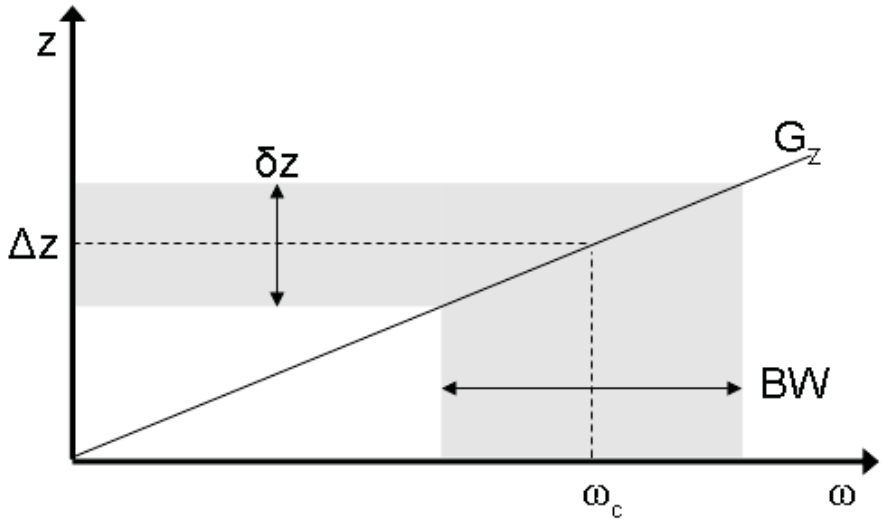

Figure 2.15: Relationship between RF pulse and slice parameters. The bandwidth of the RF pulse and the strength of the slice-selection gradient determine the width of the excited region, and the carrier frequency of the RF pulse affects the position of the excited region.

In general, MRI data is obtained by exciting the sample with an RF pulse (which can affect all of the nuclei within the RF coil or only those within a desired slice) and then, before the signal decays, acquiring some portion of k-space using a combination of frequency encoding and phase encoding. Depending on the available hardware and the relaxation rates of the material being imaged, any portion of k-space from a single point to an entire 3D volume can be collected following each excitation, though in many cases a single line is acquired per repetition. After allowing a suitable delay for recovery of longitudinal magnetization, this process can be repeated as many times as necessary to obtain the desired coverage of k-space.

The particular pattern of RF pulses and gradients used at each repetition is called a *pulse sequence*, and dozens of different sequences have been proposed, each of which offers a particular contrast weighting, k-space sampling strategy, or mechanisms for encoding other information about the sample such as diffusion weighting. The choice

of pulse sequence greatly influences the appearance and quality of the acquired image. Figure 2.16 shows some simple pulse sequences which are used to acquire gradient-echo images, with contrast determined by T_2^*, and spin-echo images, with contrast based on T_2.

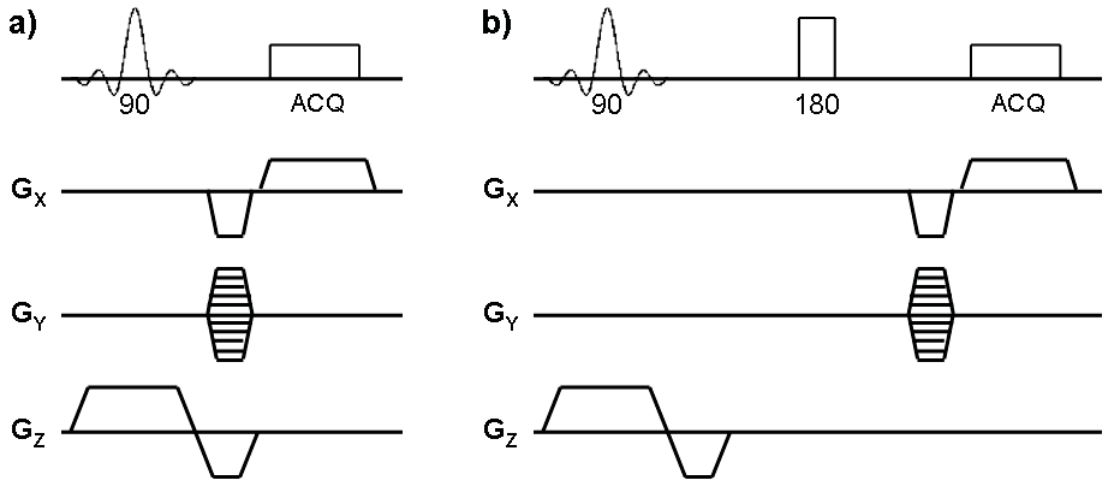

Figure 2.16: Examples of basic MRI pulse sequences. (a) Gradient-echo sequence. (b) Spin-echo sequence. Both use a slice-selective RF pulse and acquire a single k-space line per excitation. The negative lobe after the slice-select gradient reverses dephasing caused by the slice-select gradient.

2.3.3 Properties of k-space

Because MRI does not directly obtain images of a physical object, but acquires the k-space corresponding to that object and produces an image via Fourier transform, it is important to understand how the acquisition of k-space and the process of Fourier transformation affect the desired final image.

As stated earlier, k represents a spatial frequency, and the intensity of a particular location in k-space relates to how much the image varies at that frequency. The inner regions of k-space contain the low spatial frequencies, meaning that they define those parts of the image which do not change rapidly, such as large regions of generally uniform signal intensity. The outer regions of k-space, similarly, represent the higher spatial frequencies that describe edges and rapid transitions in image intensity. This is represented in Figure 2.17. The centre of k-space is proportional to the total signal

intensity of the image, as seen by setting

$$s(\vec{k}=0) = \int_{-\infty}^{\infty} \rho(\vec{x}) e^{2\pi i x \cdot 0} d\vec{x} = \int_{-\infty}^{\infty} \rho(\vec{x}) d\vec{x} \tag{2.51}$$

Therefore, acquisition of the centre of k-space is critical to properly defining the overall image contrast, but acquiring the extremities is necessary to obtain high resolution, as discussed in the previous section.

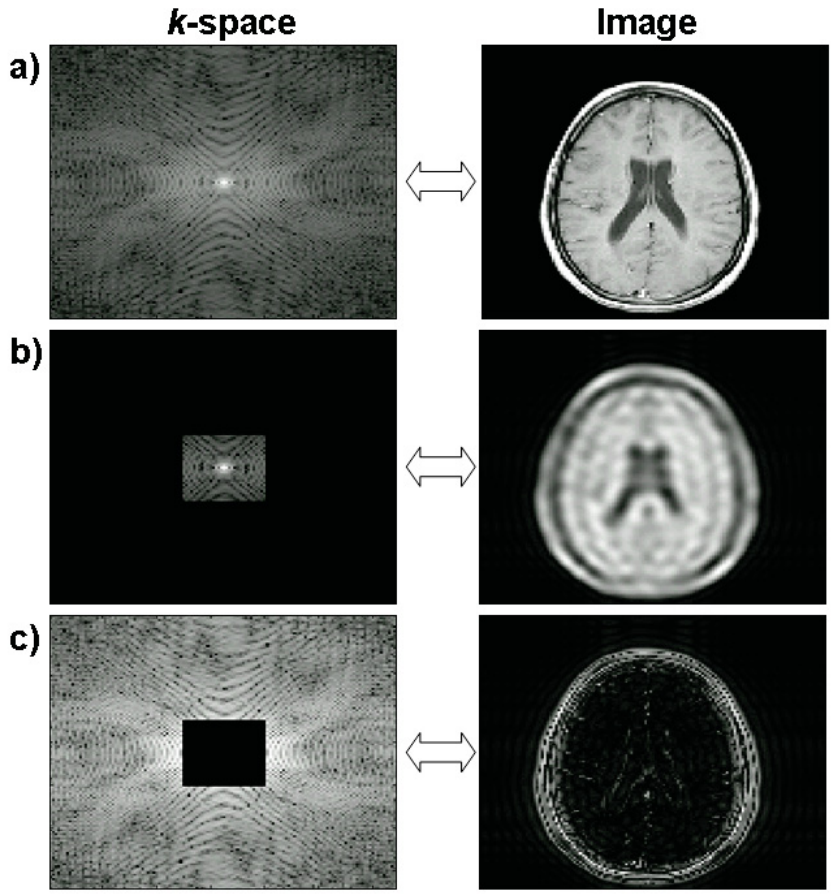

Figure 2.17: Regions of k-space contribute to different aspects of the corresponding image. (a) Complete k-space representation of a human brain, with the corresponding image after Fourier transform. (b) The central regions of k-space define the overall signal and contrast of the image. (c) The outer regions of k-space define edges and high-resolution features.

In the case that $\rho(\vec{r})$ is real, the properties of the Fourier transform also indicate that $s(\vec{k})$ should be conjugate-symmetric, i.e.

$$s(\vec{k}) = s^*(-\vec{k}) \tag{2.52}$$

This symmetry can be exploited to reduce the amount of k-space coverage needed to produce an image, since only half of the full k-space needs to be collected. However, in practice the assumption of conjugate symmetry is violated by phase variations in the image due to inhomogeneous magnetic fields and other effects. This can be compensated for to some degree by phase-correction methods such as homodyne reconstruction [79]. Nonetheless, the typical MR imaging experiment involves collection of a complete or nearly complete k-space.

A variety of irregularities or artifacts can appear in MR images, most of which manifest as modulations of the k-space. The effect of these modulations on the resulting image can be understood through the convolution theorem,

$$\mathcal{F}[fg] = \mathcal{F}[f] \otimes \mathcal{F}[g] \tag{2.53}$$

which states that multiplications in one Fourier domain correspond to convolutions in the other. The convolution operation $\otimes$ on functions f and g is defined by

$$f \otimes g = \int_{-\infty}^{-\infty} f(\tau)g(t-\tau)d\tau \tag{2.54}$$

Practically, any effect which modulates the "ideal" signal intensity in k-space - such as relaxation, finite sampling effects, inter-scan inconsistencies, or motion of the object during the scan - will have the effect of convolving the "ideal" image with the Fourier transform of the modulation. This transform is often referred to as a point spread function (PSF) since it describes how signal from an ideal point impulse is distributed to other areas of the image.

For example, since the signal decays according to T_2^* during a frequency-encoded readout, the k-space intensity is modulated by e^{-t/T_2^*} and the image is convolved with $\mathcal{F}[e^{-t/T_2^*}]$, which is a Lorentzian function with a width proportional to $R_2^* = 1/T_2^*$. This has the effect of blurring the image along the frequency-encode direction, an effect which increases as T_2^* shortens, as shown in Figure 2.18.

Another important effect is that of sampling in k-space. The actual object being imaged in any MRI experiment is continuous, as is its ideal k-space representation, but when MRI data is acquired, only discrete data points are sampled. This can be expressed as a multiplication of the ideal k-space by a sampling function,

$$U(k) = \sum_{n=-\infty}^{\infty} \delta(k - n\Delta k) \tag{2.55}$$

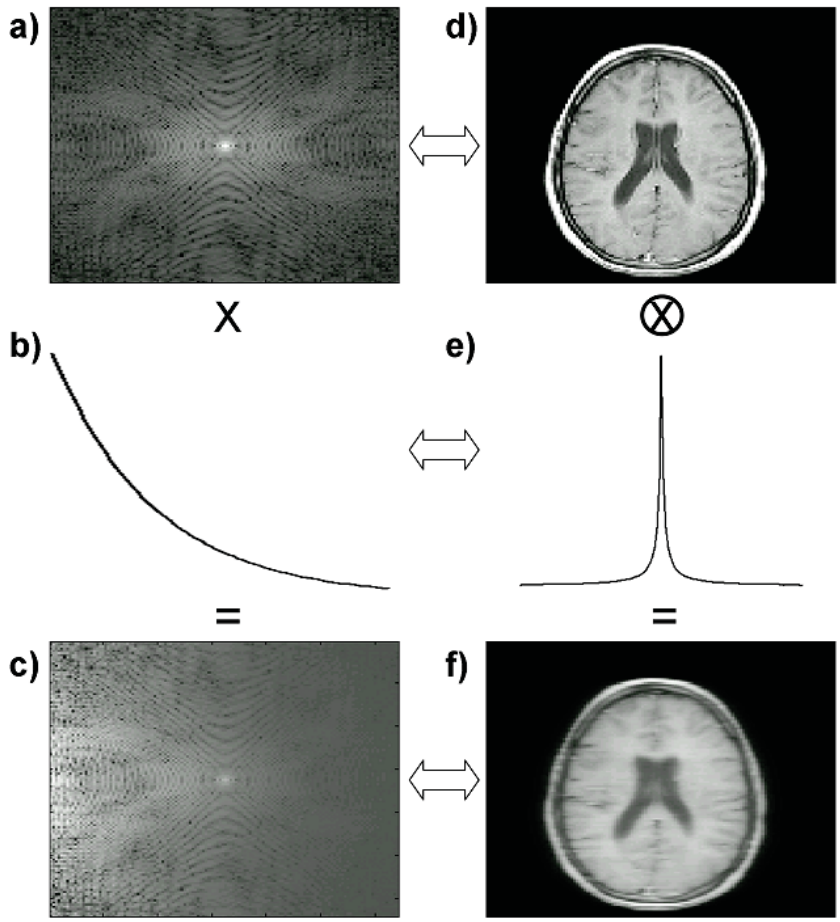

Figure 2.18: If the measured MRI signal decays significantly due to T_2^* effects during readout, the ideal k-space (a) is multiplied by an exponential function (b). The observed k-space (c) will be modulated in the frequency-encode direction (left-right). The effect on the corresponding image (d) is a convolution with the Fourier transform of the exponential decay, a Lorentzian (e). The result is an image blurred in the left-right direction (f). Note that there are no effects in the phase encode (up-down) direction.

with Δk the sampling interval size and $\delta(k)$ the Dirac delta function. The Fourier transform of such a sampling function is simply another sampling function with a spacing $1/\Delta k$. The result of convolving the image with this function is a replication of the image, creating a series of identical copies a distance $1/\Delta k$ apart. The unwanted copies or aliases can be filtered out as long as they do not overlap the primary image (Figure 2.19).

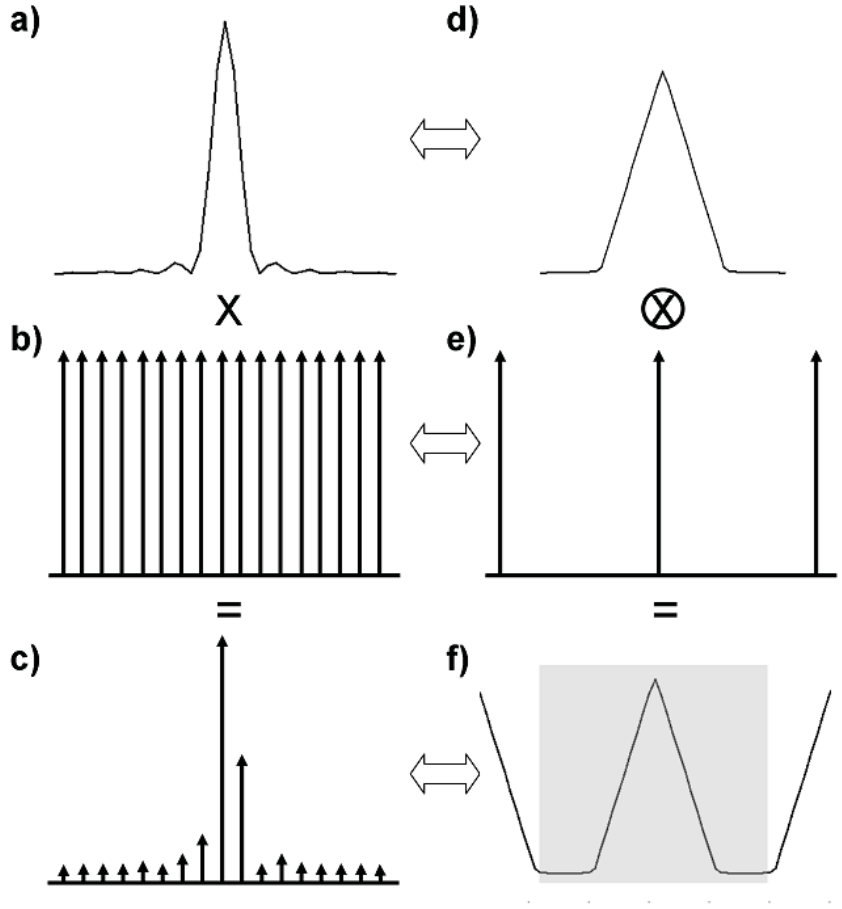

Figure 2.19: Effects of finite sampling in k-space. Sampling of an ideal, continuous k-space (a) is equivalent to multiplying by a sampling function (b), yielding a k-space with values only at a set of discrete locations (c). In image space, this is equivalent to convolving the ideal image (d) with a similar sampling function (e), creating multiple copies of the image (f). In this case the copies do not overlap, and one can be selected (gray overlay) to form the desired image.

However, if the spacing Δk is too large, the resulting convolution will create copies which are too close together. Effectively, those parts of the image which are outside

the desired field-of-view will be folded inward, creating aliasing artifacts as shown in Figure 2.20.

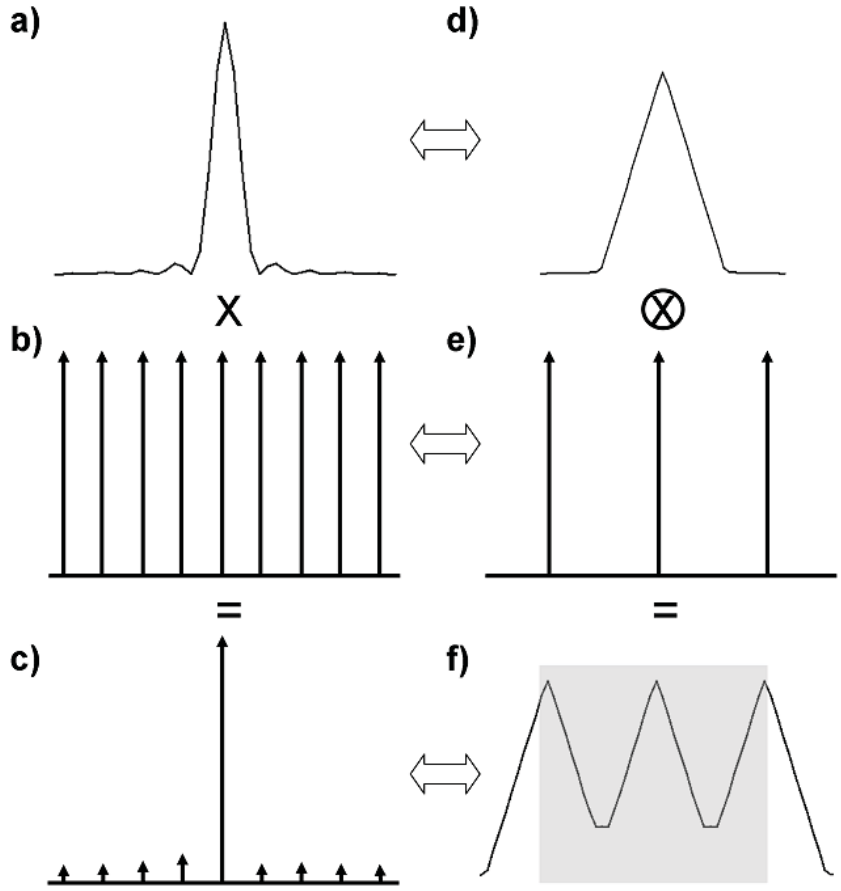

Figure 2.20: Aliasing due to insufficient sampling density. The ideal k-space (a) is sampled with an increasing spacing between points (b,c). The corresponding image (d) is therefore convolved with a sampling function of reduced spacing (e). The resulting copies overlap with each other (f), making it impossible to recover an unaliased image.

This imposes restrictions on the sampling rate of frequency-encoded images; for a fixed readout gradient, Δt must be small enough that the resulting field of view fully contains the object being imaged. This restriction is related to the well-known Nyquist criterion [75, p.236] which defines the minimum sampling rate f_s needed to unambiguously encode a signal with bandwidth BW:

$$f_s \geq 2 \cdot BW \tag{2.56}$$

Sampling below this rate means that high frequencies cannot be properly distinguished and they "wrap around" into the rest of the image[1]. In this case the bandwidth is the spread of Larmor frequencies introduced by the applied gradient, and the sampling rate is $f_s = 1/\Delta t$. Therefore,

$$\begin{aligned} \frac{1}{\Delta t} &\geq \gamma G \cdot FOV \\ \frac{1}{\gamma G \Delta t} &\geq FOV \\ \frac{1}{\Delta k} &\geq FOV \end{aligned} \tag{2.57}$$

Because Δk can also be thought of as the lowest spatial frequency encoded by the acquisition, it is necessary to keep this quantity small, such that large-scale variations in the object can be properly represented.

2.3.4 Acquisition Hardware

The typical MRI system consists of three primary components - a magnet used to generate the static field B_0, a gradient coil used to produce controlled magnetic field variations in three dimensions, and a radio-frequency transmitter/receiver coil to excite and collect signal. Modern MRI systems also include a variety of hardware to support these key systems. Figure 2.21 shows a basic block diagram of a typical MRI system.

The main magnetic field B_0 of clinical and preclinical imaging systems is generally in the range of 1.5 to 7 Tesla, and is most often provided by a superconducting magnet. The field of such magnets is most homogeneous in the center of their cylindrical bore, into which the gradient and RF coils are inserted. All imaging experiments described in this thesis were performed on a 3.0 T horizontal bore magnet.

Coaxially placed within the magnet bore is the gradient coil, which contains arrangements of current-carrying wires necessary to generate linear variations in B_0 along all three cartesian axes, as well as other non-linear variations used for "shimming" or improving overall field homogeneity. The gradient coil used for all experiments in this thesis was a 305/210 mm HD coil (Magnex Scientific, Oxford, UK) with

[1]Most MRI systems detect both the real and imaginary components of the complex MRI signal, such that the sign of the frequency can be determined. This removes the factor of 2 in the Nyquist criterion since the signal bandwidth is effectively halved.

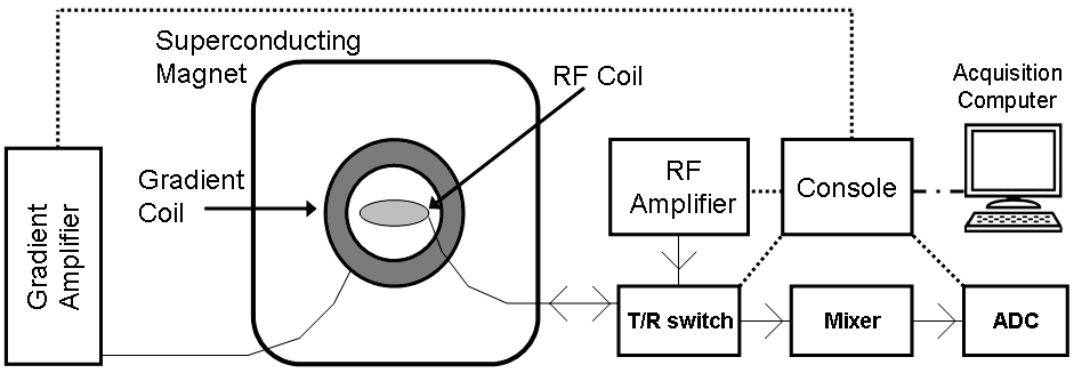

Figure 2.21: Block diagram of the key hardware components of a typical clinical or preclinical MRI system.

a maximum strength of 0.2 T/m and a slew rate of 475 T/m/s. The high currents needed to produce the desired gradients are supplied by an amplifier (Performance Controls, Inc.) that provides a maximum of 80 A per axis. An active water cooling system is used to maintain the temperature of the coil during operation.

A variety of RF coils can be used depending on the size of the object being imaged and other experimental needs. Coils can contain a single channel or two independent channels in quadrature, which effectively measure the real and imaginary component of the complex transverse magnetization. Though it is possible to use planar coils that are placed on the surface of the object being imaged, all of the experiments in this work use volume coils, within which the sample is positioned. The coils used for the experiments performed in this work included a homebuilt quadrature birdcage coil (50 mm inner diameter), a homebuilt single-channel solenoid coil (50 mm i.d.) and a Doty Scientific (Columbia, SC) quadrature Litz cage coil (25 mm i.d.).

When in transmit mode, the RF power supplied to the coil is generated by an amplifier whose output is in the kW range; in the setup used here, this was a 1 kW Varian amplifier. In receive mode, the acquired signal is pre-amplified, mixed to 20 MHz and then digitized. The collection of data, as well as timing of all inputs to the gradient and RF amplifiers, is performed by the system console and controlled by the acquisition computer. The 3.0 T MRI system used for these experiments is controlled by a Varian/Agilent DirectDrive console running VnmrJ 2.3 imaging software.

Chapter 3

Implementation and Optimization of TurboSPI

A number of MRI techniques have been used to image SPIO-loaded cells, as described in Chapter 1. While detection of SPIO is straightforward with traditional frequency-encoded techniques (see Section 2.3.2), these methods are prone to artifacts when R_2^* is large, which is the case near high concentrations of SPIO. A purely phase encoded technique would be more robust to these effects, and could greatly improve the range of effective quantification.

In this chapter we describe TurboSPI, a multiple-echo spin-echo technique originally developed for imaging non-biological materials, but which is equally applicable to biological systems, particularly those with large R_2^* and small to moderate R_2 relaxation rates such as tissues containing SPIO-loaded cells. We will also demonstrate a method to optimize the quality of TurboSPI images which increases the available signal and is robust enough to operate under *in vivo* imaging conditions.

3.1 Single Point Imaging Concepts

TurboSPI is an extension of Single Point Imaging (SPI), also known as Constant Time Imaging, which was first proposed by Emid and Creyghton [80] as a method for imaging materials with very large R_2 relaxation rates. Unlike the majority of clinical MRI techniques, SPI and other related sequences do not use frequency encoding to collect multiple k-space points along a line or other trajectory, but sample a single phase-encoded location in k-space following each RF pulse. This obviously creates a significant penalty in terms of imaging speed, since many more RF excitations are required to sample all of k-space. However, purely phase-encoded techniques have a number of advantages centered around their robustness to certain kinds of artifacts and their ability to measure samples with very large R_2 and R_2^* (short T_2 and T_2^*).

As outlined in the previous chapter, any modulation of the acquired data in k-space corresponds to a convolution of the desired image by an appropriate point

spread function. For instance, a sample in which the signal decays significantly during data collection will be blurred. Large changes in the sample phase due to field inhomogeneity or chemical shift can lead to misregistration of signal or other distortions in a frequency-encoded image. SPI and its derivatives are immune to these effects, and in areas of high signal decay or field perturbation the result is only a loss of signal intensity [81].

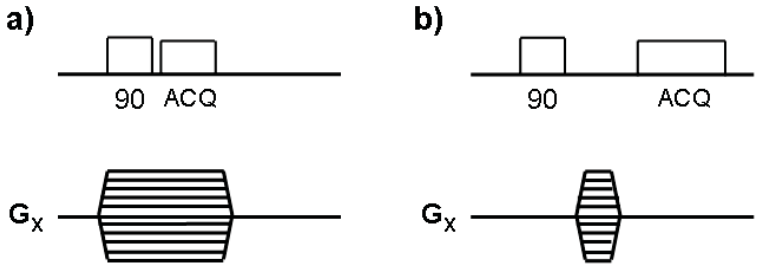

Figure 3.1: Pulse sequences for one-dimensional Single Point Imaging. (a) The signal can be excited and collected during the phase encoding gradient to minimize T_2^* decay. (b) Alternatively, data can be acquired after the phase encoding gradient is turned off, as in the case of Chemical Shift Imaging.

In the original description of SPI the data is collected while the phase encoding gradient is still turned on, to minimize the nominal echo time and reduce signal decay, as shown in Figure 3.1(a). It is also possible to collect data once the phase encode gradient has been switched off, as in Figure 3.1(b). This latter implementation allows the acquisition of multiple data points at each k-space location, and can be used to obtain a frequency spectrum at each pixel in the resulting image; this is known as Chemical Shift Imaging (CSI) [75]. In contrast, many SPI implementations of the type shown in Figure 3.1(a) only collect a single data sample at each k-space location and therefore produce a single image.

SPI and related techniques are most widely employed in the study of materials having short T_2^* relaxation times such as concrete, polymers and gases [82, 83]. If T_2 is somewhat longer than T_2^*, a refocusing pulse can be used to introduce T_2 contrast into the images without much loss of signal. Spin-Echo SPI or SE-SPI (see Figure 3.2) is useful in applications involving thin films or fluids within porous media [84, 85].

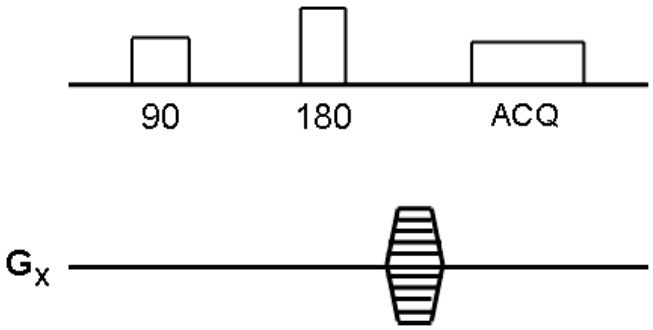

Figure 3.2: Pulse sequence for one-dimensional Spin-Echo SPI, which adds T_2 contrast with a refocusing pulse.

3.1.1 TurboSPI Sequence

A further modification for systems with much longer T_2 than T_2^* is the use of a multi-echo train of RF pulses similar to that used in RARE [86] or Fast Spin-Echo sequences. The result is Turbo Single Point Imaging or TurboSPI [63], with a very basic sequence diagram shown in Figure 3.3. TurboSPI shares the main advantages of SPI in terms of robustness to artifacts and the ability to image materials with very short T_2^*. The train of pulses allows the collection of multiple k-space locations following an RF excitation, which increases imaging speed.

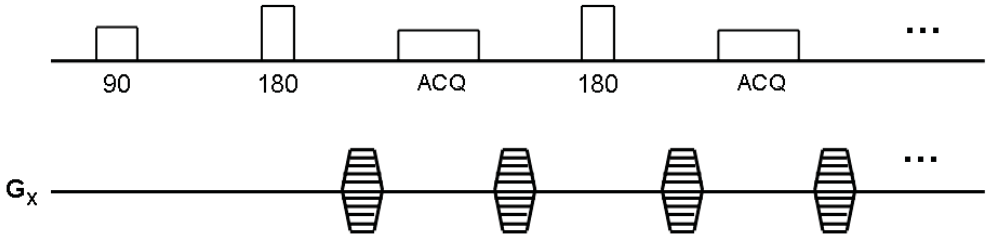

Figure 3.3: Pulse sequence for a one-dimensional non-selective TurboSPI acquisition. Two echoes are shown; typically a train of 8 or 16 echoes would be used depending on the echo time and sample T_2. The phase encoding gradient is rewound after each acquisition, so that no accumulated phase is carried over to the next echo in the train.

TurboSPI was originally proposed for systems such as porous aggregates which have very short T_2^* (< 1 ms) but comparatively long T_2 (> 50 ms), such that the resulting image is T_2-weighted but remains unaffected by short-T_2^* effects [63]. However, TurboSPI is well suited to imaging any system with $T_2 \gg T_2^*$, including cells

that have been labeled with SPIO. In particular, TurboSPI provides the opportunity to perform relaxometry on labeled cells at very high temporal resolution.

The version of TurboSPI illustrated in Figure 3.3 acquires samples along one spatial dimension, e.g. k_x, after a non-selective excitation. It is straightforward to modify the sequence to selectively excite signal from a particular slab with a shaped RF pulse (Figure 3.4(a)) and to add phase encoding gradients in all three dimensions to produce 3D images (Figure 3.4(b)). Other modifications to the sequence to optimize image quality will be discussed in the next section.

As a purely phase encoded sequence, TurboSPI offers considerable flexibility in choosing the order in which k-space locations are sampled. For the moment, we will

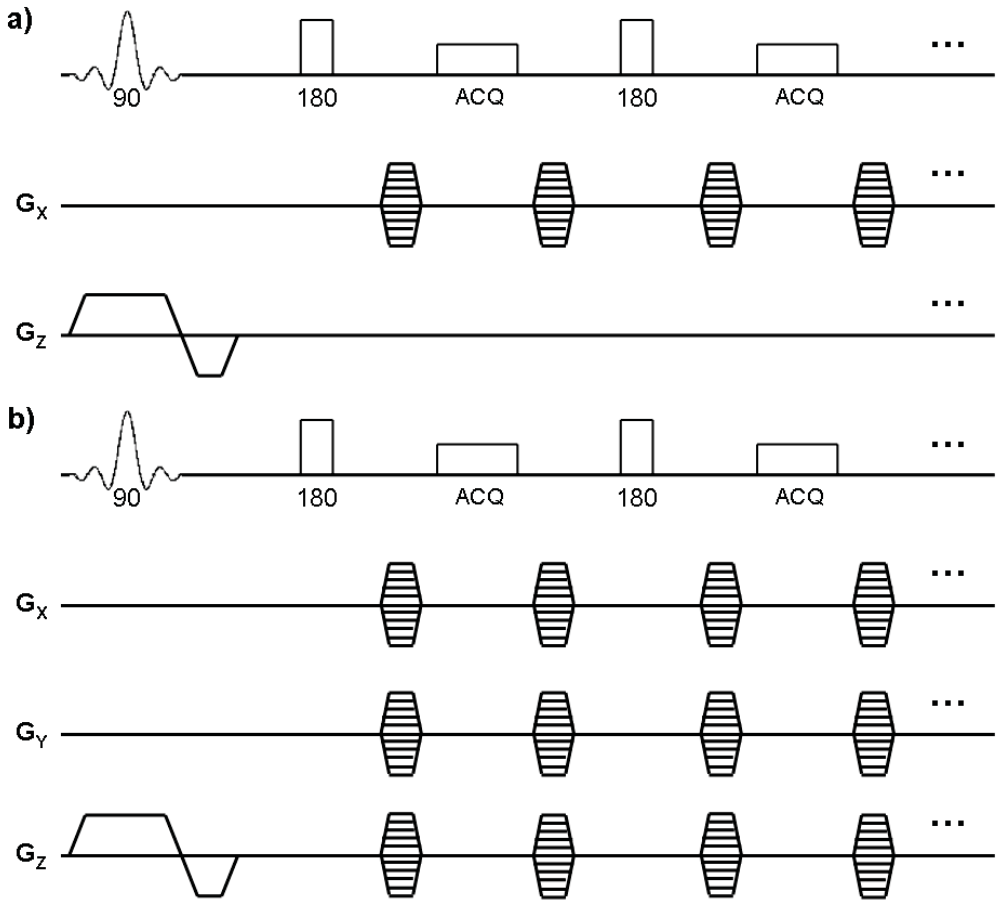

Figure 3.4: Extensions of the basic TurboSPI sequence. (a) A slice-selection gradient and a shaped excitation pulse are used to excite the magnetization, permitting slice-selective imaging. (b) Phase encode gradients can be added in up to three dimensions to allow 3D slab-selective imaging.

consider a linear ordering, which all of the phase encodes within a single line of k-space are collected before moving to the next line. Within a line, points closer to the center of k-space are collected during the first echoes of the multi-echo train, and points far from the center are collected later[1], as illustrated in Figure 3.5. This phase encoding order collects all of k-space and ensures that any artifacts due to signal modulations during the echo train are confined to one direction.

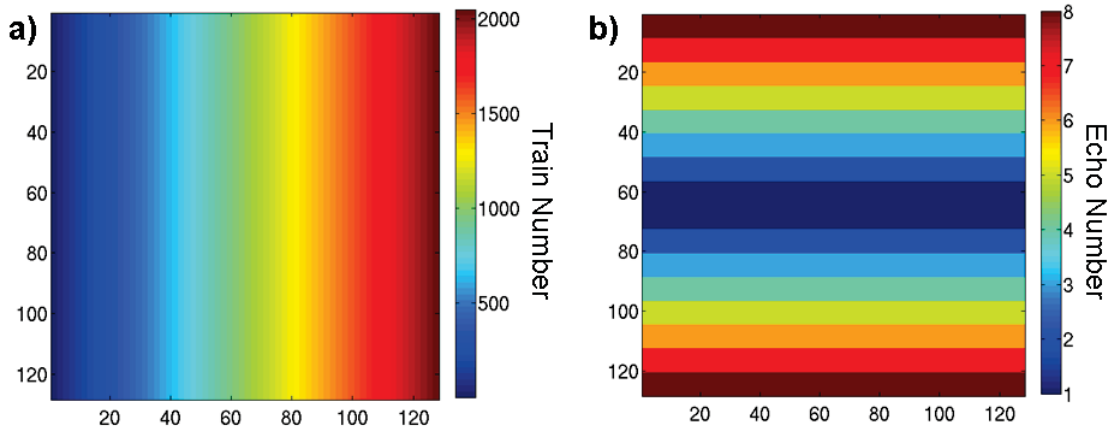

Figure 3.5: Linear phase encode ordering for a 2D TurboSPI acquisition, in which 2048 trains of eight echoes each are used to fully sample a 128x128 matrix (16384 total sampled k-space values). (a) Train number of each sampled point; trains are acquired beginning on the left side of k-space and proceeding to the right. (b) Echo number of each sampled point; the central region of k-space near $k_y = 0$ is sampled early in the train, with the outer regions sampled late in the train.

As in the case of Chemical Shift Imaging, data can sampled throughout the period between RF pulses, as the spin echo forms and then decays. The result will be a series of images which will describe the evolution of the system during the spin echo, as in Figure 3.6. Since there are no gradients applied during this period, the sampling has no restrictions in terms of bandwidth, and data can be sampled as densely as the system's digitizer will permit. For example, with a typical echo spacing of several milliseconds, hundreds of time points can be sampled at frequencies of 100 kHz or more, creating a series of images separated by tens of microseconds. For comparison, the minimum spacing between images acquired with a multiple-gradient-echo sequence like GESFIDE is several hundred microseconds, and to sample points more

[1]Phase encodes can also be ordered to enhance T_2 contrast by acquiring the centre of k-space later in the echo train [86, 63], but for our application this is not desirable.

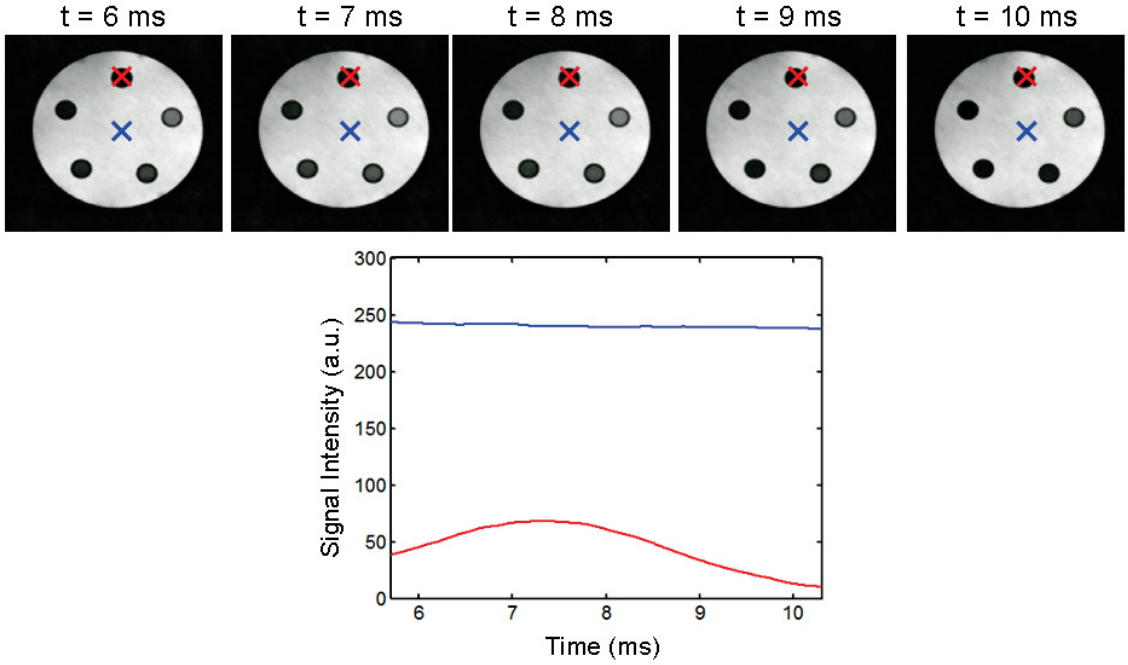

Figure 3.6: A TurboSPI image series of a cylindrical water-filled phantom containing 5 NMR tubes (described in more detail in Section 4.2.2). The water signal (blue crosses and line) has a long T_2^* and does not change significantly throughout the time course, but the contrast of the NMR tubes changes to varying degrees due to rephasing and dephasing of signal during the spin echo (red crosses and line). Only five images are shown; the complete time series typically contains 256 or 512 images.

densely spaced than this requires repeated acquisitions [62]. This gives TurboSPI a considerable advantage in terms of accurate characterization of the acquired time course data, which will be applied to R_2^* relaxometry in the next chapter.

Another advantage of SPI-type techniques like TurboSPI is the ability to filter the acquired data at far lower frequencies than is possible with a frequency-encoded image [81]. In sequences having a readout gradient, the minimum filter bandwidth is imposed by that gradient; since a sample of width δx in a gradient G_x will have a frequency bandwidth of $\gamma G_x \delta x$, the filter bandwidth must be larger than this to avoid aliasing. For practical applications this bandwidth is on the order of tens of kHz, and the inability to filter below this frequency often allows high-frequency noise into the resulting image. For a purely phase-encoded sequence like TurboSPI, no such restriction exists, and the only limit on the filter bandwidth is the frequency content of the time course. This allows significant denoising of the time course, as

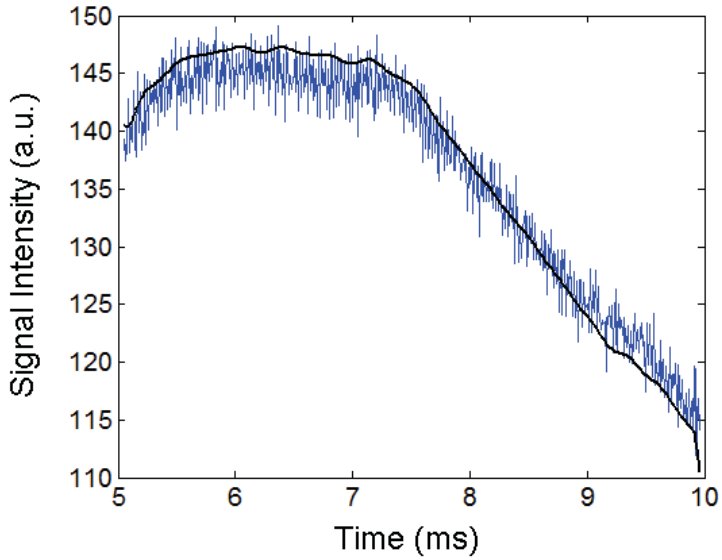

Figure 3.7: TurboSPI allows for significant denoising of time courses by lifting the restrictions on filter bandwidth which are found in frequency-encoded sequences. Blue line: Time course of a representative pixel, from an image acquired at a filter bandwidth of 100 kHz. The dominant frequency is approximately 50kHz and corresponds to noise produced by the gradient amplifier. Black line: Time course from the same pixel but with filter bandwidth of 10 kHz.

demonstrated in Figure 3.7, where a significant source of external noise is removed from the time course data without any effect on the underlying signal. In a frequency encoded image, this noise cannot be removed as easily and degrades image SNR.

3.2 Alignment of Stimulated Echoes

As discussed in Chapter 2, every RF pulse after the initial excitation will create an echo, either in the form of a spin echo or a stimulated echo. Depending on the sequence timing, these echoes can potentially overlap and interfere with each other, and in multiple-spin-echo sequences like TurboSPI this can lead to significant image artifacts.

The most common way to remove these undesired echoes is with the use of crusher or spoiler gradients surrounding each refocusing pulse in the train, as illustrated in Figure 3.8. The gradient lobe immediately after the refocusing pulse dephases any signal excited by the pulse if its flip angle is not exactly 180 degrees. The matched-amplitude lobe before the refocusing pulse ensures that signal which already exists

prior to the pulse is retained; it will be dephased before the pulse but rephased by an equal amount immediately afterward. Finally, by changing the amplitude of the matched lobes throughout the entire train, signal which is stored on the $\hat{z}$ axis by one pulse will not be fully rephased if a subsequent pulse rotates it back into the transverse plane, and therefore stimulated echoes will not form [87].

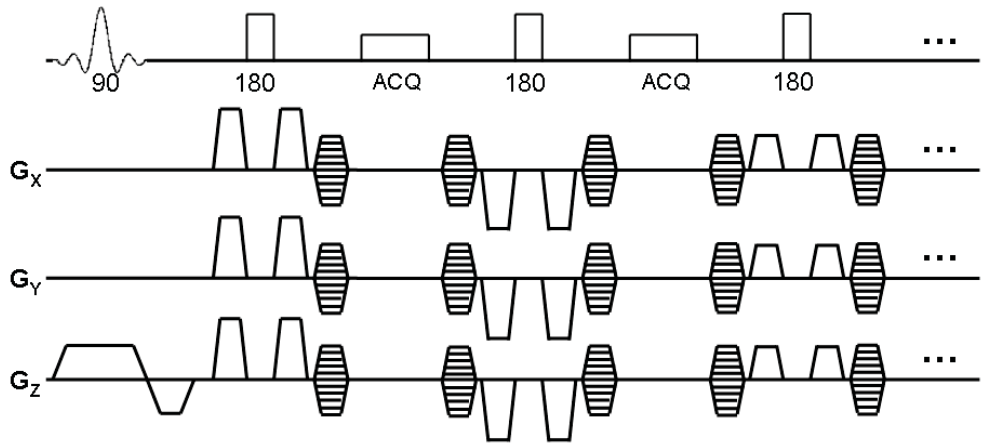

Figure 3.8: Stimulated echoes generated by the TurboSPI echo train can be suppressed with an appropriate crusher pattern, such as this one based on the method described in [87]. The crusher lobes suppress spurious signal excited by the 180 degree pulses and cause stimulated echoes to be improperly refocused, leaving only the desired spin echo intact.

Though effective, and though this method was used in previous implementations of TurboSPI [63], removing additional echoes in this way has disadvantages. If a stimulated echo has the same phase encoding as a spin echo, they can be combined constructively, creating an image with increased signal intensity compared to an image in which stimulated echoes are suppressed with crushers; this additional signal could permit higher resolution imaging and improve quantification. This is particularly true when the RF pulses are not spatially homogeneous, and more of the signal intensity forms stimulated echoes after experiencing non-ideal refocusing pulses.

In principle, alignment of echoes can be achieved with matched-amplitude crusher pairs which have the same amplitude from one echo to the next, combined with rewinding of phase encode gradients, as shown in Figure 3.9. This will still dephase signal excited by the refocusing pulses but rephase any stimulated echoes as well as the spin echo. However, this coherent summation of echoes will only be achieved if all

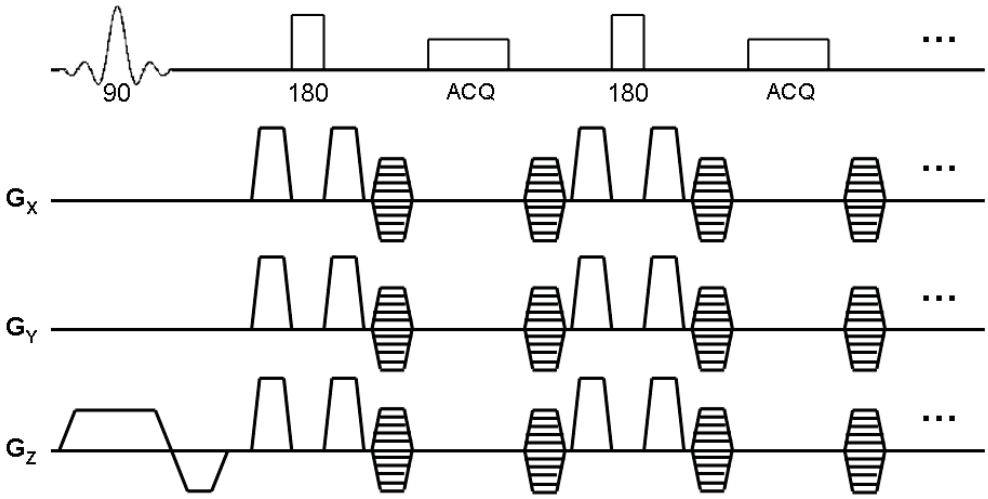

Figure 3.9: A crusher pattern which maintains a constant amplitude throughout the echo train will yield stimulated echoes that align with the spin echo, assuming that the effects of all gradients and refocusing pulses are identical.

echoes experience the same gradient area and then align in the transverse plane [64], which can be difficult in practice due to hardware limitations on gradient fidelity and RF pulse phase. Any incoherent summation of echoes will lead to severe ghosting artifacts in the phase-encode direction along which the echo train is acquired, and potentially other image distortions or artifacts, as illustrated in Figure 3.10.

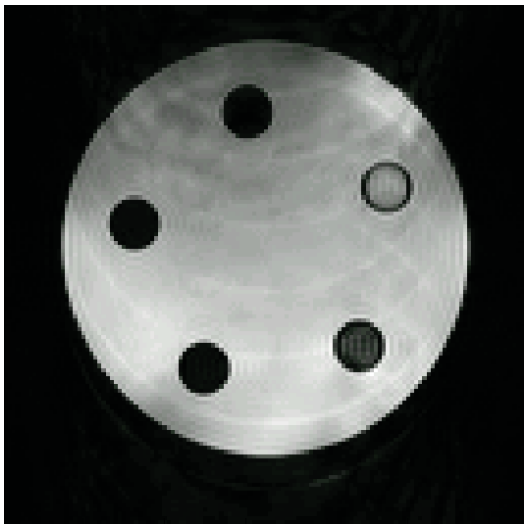

Figure 3.10: An improper alignment of spin and stimulated echoes leads to significant image artifacts, primarily ghosting in the phase encode direction (top to bottom in this image).

To compensate for these problems and optimize the quality of TurboSPI images, an automatic calibration routine was developed, similar to an existing procedure for Fast Spin-Echo sequences [64]. This calibration ensures that the necessary conditions for constructive signal summation (matched gradient-time integral and zero phase difference between refocusing pulses) are met as closely as possible.

3.2.1 Calibration Procedure

TurboSPI uses no readout gradient, rewinds the phase encoding gradients after each acquisition, and the amplitudes and durations of all crusher gradient lobes are matched throughout the sequence to prevent spurious excitation of signal. The only gradient which may be unbalanced (and which may lead to phase differences between echoes) is the slice-select gradient and its refocusing lobe, whose combined area should be zero but which may deviate from this ideal for a number of reasons. For example, while Bloch simulations can determine the theoretically required refocusing area for any RF excitation pulse [75], in practice, B_1 inhomogeneity and other factors influence the refocusing area in an unpredictable way. The correct amplitude for the slice-select refocusing gradient must therefore be determined empirically.

The first stage of the calibration routine accomplishes this by incrementally adjusting the amplitude of the refocusing gradient, and at each increment, acquiring an echo train with no phase encoding gradients. Each train is characterized by the peak intensity of the first echo, which can be plotted as a function of the refocusing gradient amplitude. The resulting curve is fitted to a polynomial (usually of order 4), with the peak of the fitted curve corresponding to the amplitude which produces the maximum signal across the train, and therefore to the optimal alignment of echoes (see Figure 3.11(a)). This amplitude produces a smooth, monotonic decay of the signal, as shown in Figure 3.11(b), and is passed back to the pulse sequence for use in all subsequent acquisitions.

The second stage of the calibration adjusts the phase of the first refocusing pulse, to account for any discrepancies between the phase accumulated prior to that pulse, and the phase accumulated between subsequent pairs of pulses. This can occur, for example, if the synthesizer phase is not properly referenced by the console during an off-resonance excitation. The refocusing pulse phase is adjusted incrementally from

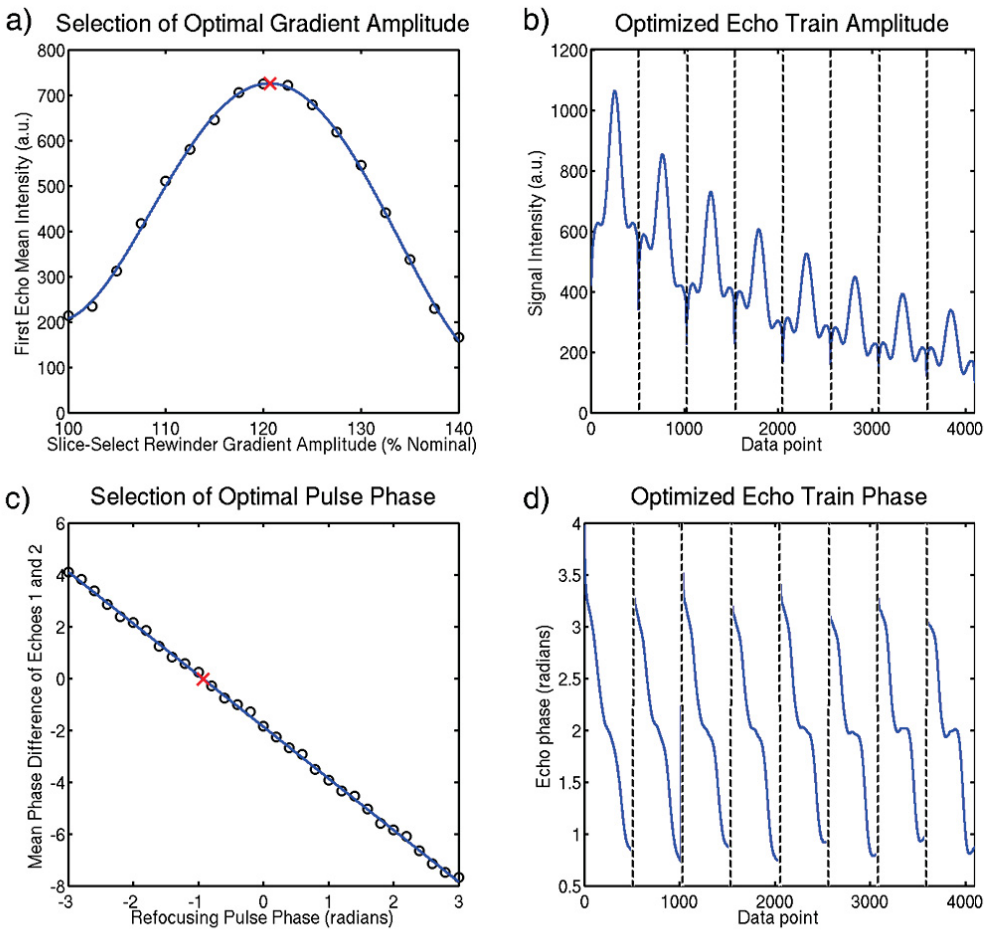

Figure 3.11: Example of automatic calibration of TurboSPI slice-select rewinder gradient amplitude and RF refocusing pulse phase. (a) Fit to the first echo peak intensity as a function of gradient amplitude, with calculated optimal value shown by the cross. (b) Magnitude of eight-echo train with no phase encoding and gradient amplitude set to optimal value. (c) Fit to the average phase difference between the first two echoes as a function of refocusing pulse phase, with x-intercept shown by the cross. (d) Phase of eight-echo train with no phase encoding and pulse phase set to optimal value.

$-\pi$ to $+\pi$, and echo trains are acquired with no phase encoding gradients. The mean phase difference between the first two readouts is computed as a function of the pulse phase, and the resulting curve is fitted to a straight line (Figure 3.11(c)). The x-intercept of this curve yields the phase which produces zero mean difference between the phase of successive readouts, corresponding to a constructive addition of spin and stimulated echoes with stable phase throughout the train, as shown in Figure 3.11(d).

Two sets of 20 to 30 trains, which can be acquired in several seconds with a repetition time of 250 ms, are typically sufficient. Processing of this calibration data is done using a compiled program written in Matlab 2008b (The Mathworks, Natick MA), which runs on the acquisition computer. The plots shown in Figure 3.11 were generated by this program.

3.2.2 Calibration Testing - Methods

To assess the performance of the automatic calibration procedure and the resulting image quality, tests were performed using a water-filled phantom containing tubes of iron oxide particles (to be further described in Section 4.2.2) and a rat imaged *in vivo*. Because of the high signal intensity and defined edges in the phantom, ghosting and other artifacts are easily observed when present, while performing the calibration on the rat will demonstrate its robustness in realistic and challenging experimental conditions.

TurboSPI images were obtained of a 5 mm thick 2D slice with a matrix size of 128x128, covering a 50x50 mm field of view. Images were acquired before and after calibration using an echo train length of 8 and a repetition time of 250 ms, with each acquisition requiring 8.5 minutes. For comparison, images were also obtained using a single-echo acquisition, which will not generate any stimulated echoes; these images were acquired in 68 minutes. Points near the center of k-space are acquired on the first echo, meaning that the effective echo time TE was equal to the inter-echo spacing of 8 ms. During each echo, 512 time points were collected at a rate of 100 kHz for a readout duration of 5.12 ms. RF pulse durations were 250 μs for the shaped excitation pulse and 50 μs for the hard refocusing pulses.

For the demonstration of *in vivo* imaging, a single male Long-Evans rat was imaged after anesthesia with an intraperitoneal injection of 1.6 g/kg urethane. The rat's

head was positioned within a stereotaxic restraint and secured with ear bars. The respiration rate was monitored, and temperature was maintained at 37°C with a feedback-controlled warm air heater (Small Animal Instruments Inc., Stony Brook, USA). Institutional animal care and use guidelines were followed throughout.

3.2.3 Calibration Testing - Results

Figure 3.12 shows images of both the phantom and rat, using a single echo acquisition, as well as before and after the calibration procedure to align echoes during an eight-echo train. Without the calibration, eight-echo images are prone to ghosting and significant distortions. When the gradient amplitude and pulse phase are optimized, the result is an image which is comparable in quality to single-echo acquisitions but which is obtained in a fraction of the time (8 minutes instead of 64). Some blurring of the multi-echo image is unavoidable, due to T_2 modulation of the PSF in the echo train direction, but the overall artifact level is low.

The images in Figure 3.12 correspond to the nominal echo time of TE = 8 ms but, as shown in Figure 3.13, the image quality remains constant throughout the rise and fall of the spin echo. There is no significant difference between the calibrated eight-echo images and the single-echo images, nor between their time courses. This is true even in areas containing significant off-resonance signal due to fat, such as the highlighted voxel whose time course is shown in Figure 3.13(k), as well as in voxels containing SPIO (not shown). These effects produce modulations of the time course, as well as altered contrast in the corresponding images (see, for example, Figure 3.13(b,d), where the signal in the marked voxel is far lower than at t=TE), but without the detrimental effects on image quality that would be seen in frequency-encoded data.

3.3 Discussion and Summary

TurboSPI has demonstrated robust image quality, both *in vitro* and *in vivo*, once stimulated echoes are properly aligned with the fast and straightforward calibration procedure described above. As with all SPI-type sequences, the series of images produced by TurboSPI is insensitive to a variety of artifacts that affect frequency-encoded techniques. To our knowledge this is the first implementation of a spin-echo

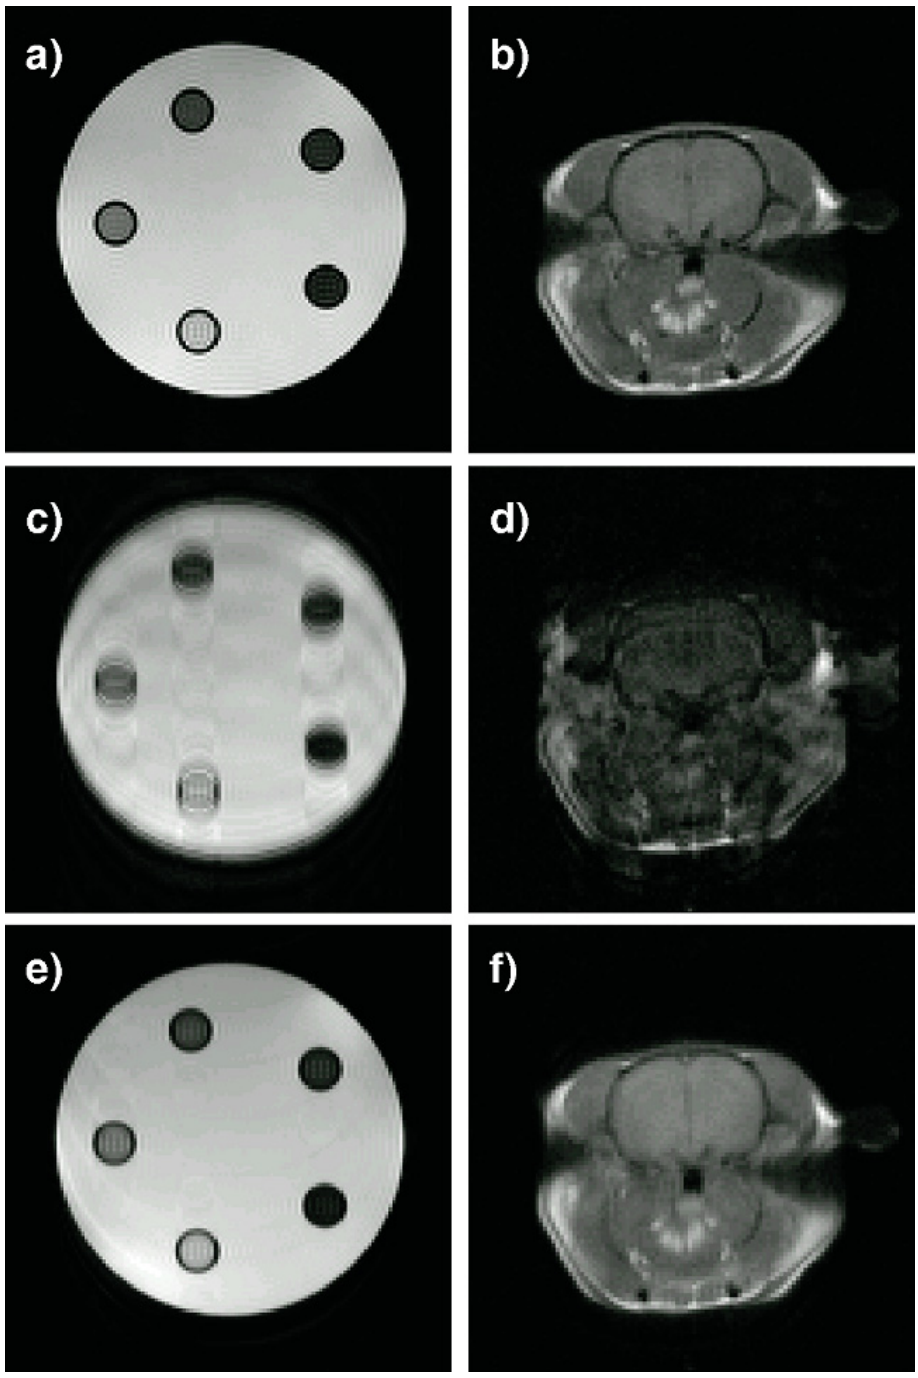

Figure 3.12: TurboSPI images are significantly improved by the automatic calibration procedure. (a,b) Images of a phantom (a) and *in vivo* rat brain (b) acquired in 64 minutes using a single echo are artefact-free. (c,d) Images acquired using an eight-echo train in 8 minutes without adjustment to gradient amplitude or RF pulse phase; significant artifacts are present in the echo direction (top to bottom in these acquisitions). (e,f) Eight-echo train images with gradient amplitude adjusted and pulse phase optimized. Some R_2 blurring remains but ghosting and distortion are reduced.

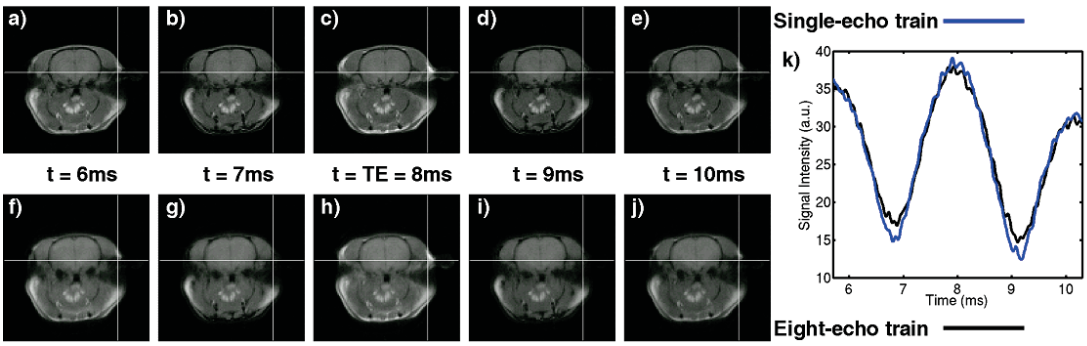

Figure 3.13: The quality of TurboSPI images remains constant through the time course. (a-e) Images of an *in vivo* rat brain obtained from different points in the time course, using a single echo. (f-j) Images from the same time points, using an eight-echo train. Aside from R_2 blurring in the multi-echo phase encode (up/down) direction, quality is similar throughout. (k) Time courses from a typical pixel (marked with crosshairs), showing that the time course behavior is unaffected by the use of a multi-echo train; the modulation is due to fat signal within the voxel.

SPI sequence which retains the signal from stimulated echoes instead of suppressing it, potentially improving signal intensity while increasing acquisition speed and avoiding artefacts due to improper echo alignment.

The calibration procedure described in this chapter allows image quality to be considerably improved after a short, two-stage prescan. In practice, the adjustment of the refocusing pulse phase is more critical to artifact levels in the image than the refocusing gradient amplitude. This is because phase differences between adjacent echoes tend to produce ghosting along the phase encoding direction, while an improperly adjusted refocusing gradient generally leads only to an overall reduction in signal amplitude. However, it is beneficial to adjust the refocusing gradient amplitude first, since this maximizes the available signal and improves the reliability of the phase adjustment.

All of the acquisitions described in this chapter were 2D slice-selective, with a relatively low in-plane spatial resolution (400 μm) and a relatively large slice thickness (5 mm). For most cellular imaging purposes a high-resolution 3D image is desirable, with resolution less than 200 μm in all three spatial directions. While the image and time course SNR are more than sufficient to attain these dimensions, the practical limitation on TurboSPI's performance is acquisition time. Since TurboSPI requires several minutes to produce single slice images at moderate resolution, high-resolution

3D *in vivo* studies would be impractically long if performed with the sequence as described in this chapter.

While lengthening the echo train from 8 echoes (as used in this chapter) to 12 or 16 echoes would seem to be a straightforward method for accelerating the acquisition, this approach has limitations. For a fixed repetition time, increasing the train length will decrease the amount of signal available for the next excitation, reducing overall SNR. Long echo trains also introduce more T_2-related blurring due to stronger modulation of the acquired k-space. Fortunately, TurboSPI is highly amenable to a number of more recently developed techniques for acceleration of MR imaging; the process of implementing these techniques for use with TurboSPI will be the subject of Chapter 6.

First, however, we will consider the ability of TurboSPI to image concentrations of SPIO such as those present in cellular imaging studies. The large number of images and the short time interval between them makes TurboSPI datasets an excellent candidate for high temporal resolution relaxometry, which is a key requirement for the accurate quantification of SPIO-labeled cells.

Chapter 4

In Vitro Quantification of SPIO with TurboSPI

As mentioned in Chapter 1, SPIO particles are readily taken up into a number of different cell types, and their effect on MRI relaxation rates allows the resulting concentration of iron to be quantified. In this chapter we will outline the physical basis of this relaxation effect and present a simple relationship between R_2' and iron concentration. Because concentrations of SPIO within cells produce a much larger increase in the reversible R_2' relaxation than the irreversible R_2 relaxation, the TurboSPI sequence is a natural candidate for imaging labeled cells. TurboSPI's ability to reliably image and quantify iron using the established theoretical relationship will then be demonstrated using samples of micron-sized iron oxide particles and SPIO-labeled cells *in vitro*.

4.1 Relaxation Behaviour of SPIO

The relaxation of the MRI signal is physically based on molecular motion, as discussed in Chapter 2. For SPIO and other similar contrast agents, numerous studies have shown that the R_2 and R_2' relaxivity have a strong dependence on the particle size [88, 89, 90]. For small particles, both relaxivities increase with particle diameter, but beyond a certain diameter the R_2 relaxivity begins to decrease, while R_2' continues to increase and eventually reaches a maximum value, as shown in Figure 4.1.

This phenomenon can be understood by considering the diffusion of water molecules in the vicinity of an SPIO particle. The portion of the particle's inhomogeneous field which is strong enough to cause significant dephasing of signal occupies a certain volume in space, which is proportional to the particle's radius as well as its magnetic susceptibility. If a water molecule's diffusion between the excitation and acquisition of signal is sufficient to move it completely through the majority of the particle's magnetic field, the loss of coherence will be unrecoverable even after a refocusing RF pulse, leading to relaxation dominated by R_2. This is called the diffusion-narrowing

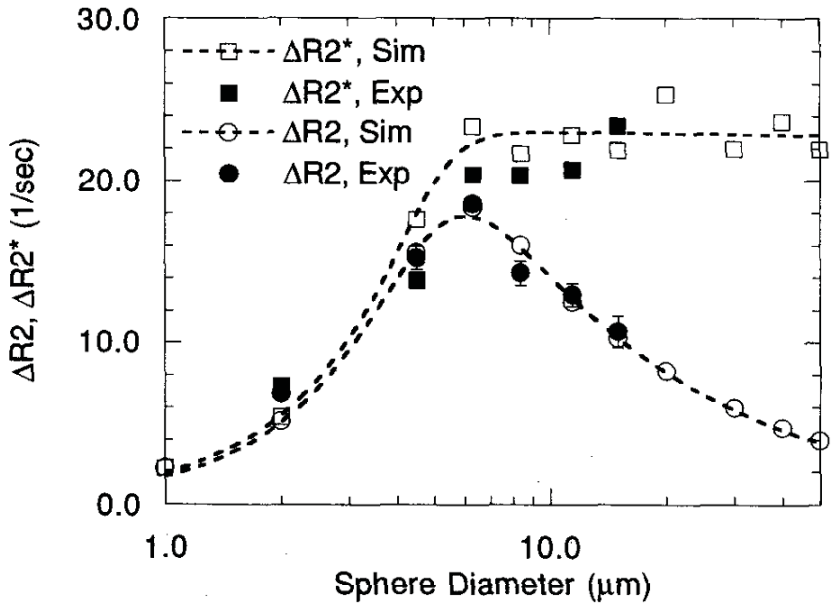

Figure 4.1: R_2 and R_2^* relaxivity as a function of particle size, for microspheres containing Dysprosium-DTPA, a paramagnetic contrast agent. The change in R_2 produced by the contrast agent reaches a peak between 5 and 10 μm, and then decreases with particle size, while the R_2^* relaxivity reaches a maximum value at approximately the same size. Figure reproduced from [88].

or motional narrowing regime[1].

Alternatively, if the particle diffuses through only a small portion of the field, the inhomogeneity is effectively time-independent and the dephasing can be reversed. This is known as the static dephasing regime (SDR), in which R_2' is the dominant relaxation mechanism. Once a system has reached the SDR, further increases in particle size will have no impact on relaxivity, and R_2' reaches a plateau as observed experimentally. (The precise conditions which must be satisfied for a system to be considered in the static dephasing regime will be addressed in Section 4.1.4.)

In the case of cellular imaging, because the individual SPIO particles are compartmentalized within the cell, it is not the particle size which is of importance, but the size of the cell. When modeling the relaxation behavior we can then consider an SPIO-loaded cell as a single large SPIO particle with an appropriate susceptibility.

[1]The motional narrowing regime is so named because, in inhomogeneous systems with very broad NMR linewidths, diffusion can actually have the effect of reducing the apparent field inhomogeneities through motional averaging, leading to a narrower linewidth than expected.

The theory describing signal behavior in the static dephasing regime can then be used to predict the relationship between the total amount of iron present in an area and the R_2' relaxivity. This prediction, which will be derived below, is the basis of SPIO quantification with relaxometry.

4.1.1 Relaxation in the Static Dephasing Regime

The first comprehensive theory describing MRI signal behavior in the static dephasing regime was that of Yablonskiy and Haacke [91]. This model completely neglects the effects of diffusion and, in its original formulation, describes the signal in the vicinity of magnetized perturbers immediately following excitation, i.e. during an FID experiment.

Consider an ensemble of N perturbers with radii R_n (but geometry otherwise unspecified for now), which occupy a total volume v in a background medium of size V_0 such that the total volume $V = V_0 + v$. When placed within an external field B_0 these perturbers will create inhomogeneous magnetic fields in their vicinity, affecting the local Larmor frequency $\omega(\vec{r})$ about each perturber. Regardless of the precise form of $\omega(\vec{r})$, the NMR signal immediately after RF excitation should be of the form

$$s(t) = \frac{1}{V}\rho \int_{V_0} e^{-i\omega(\vec{r})t} d\vec{r} \tag{4.1}$$

with the local frequency $\omega(\vec{r})$ given by a sum over all N perturbers,

$$\omega(\vec{r}) = \sum_{n=1}^{N} w_n(\vec{r}) \tag{4.2}$$

Equation 4.1 is similar in form to equation 2.37 for the signal in the presence of a field gradient, though the integral is now over V_0 and normalized by the total volume. Indeed, the scaling factor ρ is essentially $\rho(\vec{r})$ integrated over this volume, and as such it incorporates effects due to flip angle and hardware sensitivity that are unrelated to relaxation. For the purposes of this discussion, any T_2 relaxation due to the background medium can also be subsumed into ρ.

Instead of an integral over a sum of perturbers, equation 4.1 can be expressed as a product of integrals, each of which represents the contribution of a single perturber. This is valid so long as the volume occupied by the perturbers is small and their inhomogeneous fields do not significantly overlap. Assuming that the perturbers are

uniformly and independently distributed, i.e. $P(\vec{r}_n)d\vec{r}_n = d\vec{r}_n/(V - v_n)$, the signal becomes

$$\begin{aligned} s(t) &= \rho\frac{V_0}{V}\prod_n \frac{1}{V - v_n}\int_{V-v_n} e^{-i\omega_n(\vec{r})t}d\vec{r} \\ &= \rho(1-\zeta)\prod_n \frac{1}{V - v_n}\int_{V-v_n} e^{-i\omega_n(\vec{r})t}d\vec{r} \end{aligned} \tag{4.3}$$

with the volume fraction $\zeta = \frac{v}{V}$ introduced for convenience. The $(1-\zeta)$ term reflects the fact that we are only considering signal from the background medium, i.e. that the perturbers themselves are impermeable and contain no MR-sensitive nuclei.

In the simplest case, the perturbers all have the same radius R, and the volume v_n is $v(R)$, the volume of a single perturber. The product over all perturbers can be rewritten as [91]

$$\begin{aligned} s(t) &= \rho(1-\zeta)\prod_n \left(1 - \frac{1}{V - v_n}\int_{V-v_n}(1 - e^{-i\omega_n(\vec{r})t})d\vec{r}\right) \\ &= \rho(1-\zeta)\left(1 - \frac{v(R)}{V - v(R)}f\right)^N \end{aligned} \tag{4.4}$$

with the function f defined here as

$$f = \frac{1}{v(R)}\int(1 - e^{-i\omega(\vec{r})t})d\vec{r} \tag{4.5}$$

The benefit of re-expressing $s(t)$ in this form is that only the f function depends on the geometry of the perturbers through $\omega(\vec{r})$ and $d\vec{r}$.[2] In the statistical limit $N \to \infty$ and using $\frac{v(R)}{V-v(R)} \approx \frac{v(R)}{V} = \zeta/N$, $s(t)$ can be further simplified to

$$s(t) = \rho(1-\zeta)\left(1 - \frac{\zeta f}{N}\right)^N \tag{4.6}$$

$$= \rho(1-\zeta)e^{-\zeta f} \tag{4.7}$$

In the more general case that the perturbers have a distribution of radii, with perturbers of radius R_n occupying a volume fraction $\zeta(R_n)$, it is straightforward to demonstrate that

$$s(t) = \rho(1-\zeta)e^{-\int \zeta(R)f(R)dR} \tag{4.8}$$

[2]The f function also has a useful geometric interpretation which will be revisited in Chapter 5; briefly, it relates to the effective volume throughout which a perturber is able to dephase signal.

4.1.2 FID Experiment with Spherical Perturbers

The expression just derived is true for perturbers of any geometry; to proceed further we need to compute f by substituting appropriate expressions for $\omega(\vec{r})$ and $\mathrm{d}\vec{r}$ into equation (4.5). Though this can be done for any geometry, including arrangements of parallel or randomly oriented cylinders [91], for our purposes we consider only the case of spherical perturbers.

Regardless of whether it represents a single SPIO particle or an SPIO-loaded cell, a spherical perturber of radius R and with a magnetic susceptibility χ will produce an inhomogeneous magnetic field in the form of a magnetic dipole, such that

$$B(r,\theta,\phi) = \frac{4\pi}{3}\Delta M \left(\frac{R}{r}\right)^3 (3\cos^2\theta - 1) \tag{4.9}$$

The derivation of this expression is given in Appendix A.1.

The change in magnetization is $\Delta M = (\chi - \chi_0)B_0$, where χ is the susceptibility of the SPIO particles and χ_0 is that of the surrounding medium.[3] Using the Larmor equation (equation (2.6)) the frequency shift experienced by a water molecule in the vicinity of this SPIO particle can be expressed as

$$w(r,\theta,\phi) = \delta\omega \left(\frac{R}{r}\right)^3 (3\cos^2\theta - 1) \tag{4.10}$$

with the characteristic frequency $\delta\omega$ defined as the shift at the sphere's equator,

$$\delta\omega = \omega(r = R, \theta = \pi/2) = \frac{4\pi}{3}\gamma\Delta M \tag{4.11}$$

This characteristic frequency will be of importance when defining the static dephasing condition.

Substituting equation (4.10) into equation (4.5) and using the volume element $d\vec{r} = r^2 dr \sin\theta d\theta d\phi$, we obtain

$$f = \frac{3}{4\pi R^3}\int_0^{2\pi} d\phi \int_0^{\pi} \sin\theta d\theta \int_R^{\infty} r^2 dr (1 - e^{-i\delta\omega t\left(\frac{R}{r}\right)^3 (3\cos^2\theta - 1)}) \tag{4.12}$$

which can be simplified using the substitutions $U = \cos\theta$ and $u = R^3/r^3$, such that $r^2 dr = \frac{-R^3 du}{3u^2}$. Then

$$f = \frac{1}{2}\int_{-1}^{1} dU \int_0^1 \frac{du}{u^2}(1 - e^{-i\delta\omega t u(3U^2 - 1)}) \tag{4.13}$$

[3] At many field strengths used for MRI, the magnetization of the SPIO will *saturate* or reach a maximum, when all magnetic domains are aligned. In this case ΔM is the saturation magnetization.

Though not integrable in general, this expression does have asymptotic forms for small and large values of $\delta\omega t$ [91]. During the short time scale $t \leq 1.5/\delta\omega$, f has a quadratic form,

$$f = 0.4(\delta\omega t)^2 \tag{4.14}$$

while on long time scales ($t \geq 1.5/\delta\omega$), f is linear:

$$f = (C + iC')\delta\omega t - 1 \tag{4.15}$$

The constants C and C' can be computed to be

$$C = \frac{2\pi}{3\sqrt{3}} = 1.2092$$
$$C' = \frac{2}{3}\left(\frac{1}{\sqrt{3}}\ln\frac{\sqrt{3}+1}{\sqrt{3}-1} - 1)\right) = -0.1598$$

In most practical applications $1.5/\delta\omega$ is very small, and signal can only be collected during the long time scale. In this case, f from equation (4.15) can be substituted into equation (4.7) to give the evolution of the signal as

$$s(t) = \rho(1-\zeta)e^{-C\zeta\delta\omega t}e^{\zeta}e^{-i\zeta C'\delta\omega t} \tag{4.16}$$

The e^{ζ} term is close to 1 as long as ζ is small, and the oscillatory term will only affect the phase of the signal, not the decay of its amplitude; both are often neglected from further consideration. We therefore identify the relaxation rate R_2' as

$$R_2' = C\zeta\delta\omega = \frac{8\pi^2}{9\sqrt{3}}\zeta\gamma\Delta M \tag{4.17}$$

This defines the maximum relaxivity for any set of spherical magnetic perturbers in the static dephasing regime, regardless of their size.

It is convenient to express this relaxivity in terms of the quantity of magnetic material present per voxel. While units of s^{-1}/mM are often used in the context of contrast agents, knowledge of the molar concentration of SPIO is not always readily available. An alternative is to express R_2' in terms of the local magnetic dose, which is defined as the magnetic moment per unit volume, or [5]

$$LMD = \zeta\Delta M \tag{4.18}$$

In these units, the static dephasing relaxivity becomes

$$R_2' = \frac{8\pi^2}{9\sqrt{3}}\gamma \cdot LMD \simeq 10.78 \cdot LMD \tag{4.19}$$

4.1.3 Spin-Echo Experiment

In systems with high concentrations of SPIO, and therefore large R_2', it can be difficult to sample the FID before it decays completely, making accurate relaxometry challenging. It may be beneficial to refocus the R_2' decay with a 180 degree pulse and sample the signal during the resulting spin echo. Though the treatment above has been based on an FID experiment, in the static dephasing regime it is assumed that the signal decay due to magnetic perturbers should be mostly reversible, and $s(t)$ should have the same general form but with a dependence on the echo time TE.

Specifically, Yablonskiy [92] has modeled the signal during the spin echo as follows. For clarity we will also include R_2 relaxation from the background medium, which had been neglected to this point. The rising portion of the spin echo prior to t = TE - $1.5/\delta\omega$ will have the form

$$s(t) = \rho(1-\zeta)e^{\zeta}e^{R_2'(t-TE)-R_2t} \tag{4.20}$$

and similarly, after t = TE + $1.5/\delta\omega$,

$$s(t) = \rho(1-\zeta)e^{\zeta}e^{-R_2'(t-TE)-R_2t} \tag{4.21}$$

The only difference in these forms is the sign of the R_2' term, since this portion of the signal is rephasing before the spin echo and dephasing afterwards. Between t = TE - $1.5/\delta\omega$ and t = TE + $1.5/\delta\omega$ the signal will evolve non-linearly using the f given in equation (4.14).

Sampling of the spin echo has been demonstrated with multiple-gradient-echo sequences [58, 59] and carries a number of advantages. For example, if both sides of the spin echo are sampled, it becomes possible to separate the reversible (R_2') and the irreversible (R_2) contributions to the overall relaxation rate R_2^*. Furthermore, though the region of non-linear behavior around the echo peak is predicted to be quite short, it has been experimentally demonstrated that there is measureable deviation from linear behavior around the peak of the FID [92]. This can be used to estimate the volume fraction ζ occupied by the perturbers, by comparing the extrapolated signal $s(TE)$ from equations (4.20) and (4.21) to the actual signal measured at t=TE.

4.1.4 Conditions for the Static Dephasing Regime

As described at the beginning of this section, for a system to be considered in the static dephasing regime, diffusion must be sufficiently slow that a water molecule will not move significantly through a perturber's inhomogeneous field during the experiment. A condition to describe such systems can be formulated in terms of the perturber radius R, the diffusion coefficient D of the background medium, and the strength of the inhomogeneous magnetic field around the perturber, which is proportional to the characteristic frequency $\delta\omega$.

The time for a particle to diffuse through a distance R in a system of d dimensions ($d = 1, 2, 3$) is given by

$$t_D = \frac{R^2}{2 \cdot d \cdot D} \tag{4.22}$$

and we can also define a characteristic time as $t_c = 1/\delta\omega$. The simplest criterion for defining the static dephasing regime is then

$$t_c \ll t_D \tag{4.23}$$

In other words, the characteristic time t_c for signal from a proton near a perturber to dephase should be much less than the time needed for that proton to diffuse away from the particle. If this condition holds, the field should appear effectively static.

However, Yablonskiy [91] observed that this criterion is not appropriate for all system geometries, since it does not consider the volume fraction ζ occupied by the magnetic perturbers. He instead proposed a condition based on the NMR signal decay which, as we have seen, occurs on a time scale proportional to $(\zeta\delta\omega)^{-1}$. The static dephasing condition would then be

$$\frac{1}{\zeta \cdot \delta\omega} \ll \frac{(\bar{r}/2)^2}{2 \cdot d \cdot D} \tag{4.24}$$

Another modification is the use of $\bar{r}$, the average distance between neighboring perturbers, which is a more relevant distance scale for loss of NMR signal phase coherence.

For systems of perturbers such as cylinders, where the volume fraction varies as $(R/\bar{r})^2$, this condition is equivalent to the simpler form given in equation (4.23). However, since spheres of radius R occupy a volume fraction of approximately

$$\zeta \approx \left(\frac{R}{\bar{r}/2}\right)^3 \tag{4.25}$$

it can be shown that the simple form is not recovered, and instead we find that

$$\begin{aligned}\left(\frac{\bar{r}/2}{R}\right)^3 &\ll \frac{(\bar{r}/2)^2}{6\cdot D}\cdot\delta\omega \\ \left(\frac{\bar{r}/2}{R}\right) &\ll \frac{R^2}{6\cdot D}\cdot\delta\omega \\ \zeta^{-1/3} &\ll \frac{R^2}{6\cdot D}\cdot\delta\omega \end{aligned} \tag{4.26}$$

as stated in equation (24) of [91]. Since ζ must be less than 1, this condition is more stringent than the simpler expression above, and if perturber concentrations are very small, a system which would be otherwise expected to be in the SD regime might not actually satisfy this condition.

In the case of a spin-echo experiment, different considerations apply. As the echo time TE increases, the diffusive attenuation of the signal grows accordingly, and it is this attenuation which must be small compared to the phase dispersion introduced by the presence of the perturbers. Majumdar and Gore [93] determined that, if the field in the vicinity of a diffusing molecule can be approximated by a gradient G, the expected attenuation of signal at TE is

$$\exp\left[-\frac{1}{12}D\cdot\left\langle(\gamma G)^2\right\rangle TE^3\right] \tag{4.27}$$

Based on this assumption, Yablonskiy [92] computed that, for an arrangement of parallel cylinders, the diffusive attenuation would be[4]

$$\exp\left[\frac{1}{24}\zeta\frac{TE^3}{t_c^2 t_D}\right] \tag{4.28}$$

To satisfy the static dephasing regime condition we require the argument of the exponential to be small, and therefore

$$\begin{aligned}\frac{1}{24}\zeta\frac{TE^3}{t_c^2 t_D} &\ll 1 \\ TE &\ll 2.9(T_2' t_c t_D)^{1/3}\end{aligned} \tag{4.29}$$

This result has been referenced in the literature [62] in the context of spherical perturbers, though it is technically only valid for parallel cylinders. However, a calculation for the case of spherical perturbers yields a similar result:

$$TE \ll 2.4(T_2' t_c t_D)^{1/3} \tag{4.30}$$

[4]In this expression the characteristic time t_c has been modified to include the coefficient C from equation (4.15), such that $t_c = \frac{\zeta}{R_2'}$.

For details of these calculations, see Appendix A.2.

Interestingly, despite the differences in system geometry, the condition differs from the cylindrical case only in the numerical coefficient, and even then, the difference is minimal. This means that authors who have cited equation (4.29) in the context of a spherical system were not greatly mistaken, since a system which fulfills that criterion will fulfill equation (4.30) as well.

It should also be be noted that the parameters best suited to meeting the static dephasing regime conditions differ significantly between FID and spin-echo experiments. While large particles and slow diffusion (hence large t_D) are always required for static dephasing, in FID experiments the particles should have large characteristic frequency offsets $\delta\omega$ and occupy a large volume fraction ζ, such that dephasing due to field gradients is more rapid than dephasing due to diffusion. However, in a spin-echo experiment, particles with small $\delta\omega$ occupying a small volume fraction are more likely to meet the criterion stated in equation (4.30) since these conditions reduce diffusive attenuation of the signal. In practice a compromise between these factors is necessary to meet both criteria.

4.2 Methods

With a theoretical understanding of the relationship between the relaxation rate R_2' and the local iron concentration, we can now explore the potential of using a TurboSPI time series to quantify concentrations of SPIO with R_2^* relaxometry. To demonstrate the accuracy and range of this quantification, a number of samples were prepared for *in vitro* imaging, which is more reproducible and is limited by fewer confounding factors than would be present if attempting to characterize TurboSPI *in vivo*. Other quantification techniques were also employed for comparison with TurboSPI and to allow verification of the theoretically described relaxation behavior.

4.2.1 Sample Preparation

Three series of samples were prepared to test the ability of TurboSPI to accurately quantify intra-cellular and freely suspended iron. The first two series of samples consisted of micron-sized particles of iron oxide (MPIO, Bangs Laboratories, Fishers, IN) which had a mean diameter of 0.96 μm in the first group, and 1.63 μm in the

second. Various concentrations of MPIO were used; calculated iron concentrations in the 0.96 μm group ranged from 0 to 24 μg/mL, and from 0 to 120 μg/mL in the 1.63 μm group.

A third series of samples was produced which contained HPV C3 cervical cancer cells [94], which were loaded with Molday ION Rhodamine B, a 35 nm SPIO particle obtained from BioPal Inc (Worcester, MA). The C3 cell line was maintained in Iscove Modified Dulbeccos Medium (IMDM; Sigma, St. Louis, MO) supplemented with 10% heat-inactivated fetal calf serum (Sigma, St. Louis, MO), 2 mM l-glutamine (Gibco, Burlington, ON), 50 mM 2-mercaptoethanol (Gibco, Burlington, ON), 100 U/mL penicillin and 100 μg/mL streptomycin (Gibco, Burlington, ON). Following incubation of the cells in 5% CO_2 at 37°C for 48 hours, Molday ION Rhodamine B was added to the cells at a concentration of 0.1 mg/mL and incubated for a further 22 hours.

The cells were washed to remove excess iron oxide, trypsinized, centrifuged and re-suspended in a 1X HBSS (Sigma, St. Louis, MO) solution with 200 μL/mL HEPES (Gibco, Burlington, ON). The suspended cells were divided into groups with concentrations ranging from 1 million to 5 million cells/mL. Cellular iron uptake was assessed with UV/VIS spectrophotometry and was determined to be approximately 8 pg Fe/cell. The average cell size was calculated using the Countess Automated Cell Counter (Invitrogen, Burlington, ON) and accompanying software; the average diameter was determined to be 13 μm.

To prepare each sample, an appropriate amount of MPIO particles or labeled cells were suspended in a heated solution of gelatin (4% by weight). This mixture was doped with 168 μM $MnCl_2$, a paramagnetic contrast agent, to give the samples a background T_1 and T_2 comparable to that which would typically be found in biological samples. 2 mL of each iron concentration was placed in a 5 mm NMR tube and the gelatin was set with an icewater bath to avoid settling.

For imaging with TurboSPI, these NMR tubes were placed in groups of five within a cylindrical phantom (see Figure 4.2). The rest of this phantom was filled with distilled water that was also doped with 168 μM $MnCl_2$.

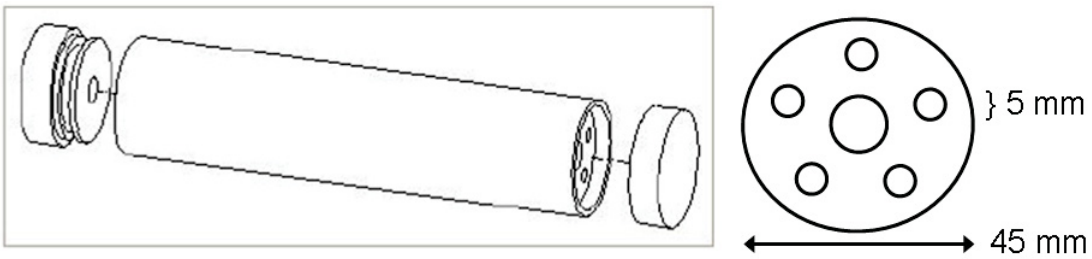

Figure 4.2: Cylindrical phantom for SPIO quantification with TurboSPI. The phantom can hold up to five 5mm NMR tubes in a doped water bath, and is constructed of ABS plastic. Schematic reproduced from [60].

4.2.2 Imaging Parameters

All TurboSPI acquisitions were 2D slice-selective, with the same parameters as described in the previous section (5 mm slice thickness, 128x128 matrix size, 50x50 mm field of view, train length of 8, repetition time of 250 ms, TE = 8 ms, 512 time points collected at 100 kHz). The home-built quadrature RF coil was used for all TurboSPI image acquisitions. RF pulse durations were 250 μs for the shaped excitation pulse and 50 μs for the hard refocusing pulses.

For comparison, bulk measurements of the R_2^* relaxation rate were obtained using a slice-selective pulse-acquire sequence to collect the FID of a 10 mm slice through each tube, and a slice-selective single-echo spin-echo sequence to collect the decaying signal immediately following a spin echo. These sequences were implemented by modifying spectroscopy sequences included on the system console (the voxel-selection gradients were replaced with a single slice-select gradient). For the spin-echo sequence, the echo time was set at TE = 8 ms to match the TurboSPI acquisitions. Tubes were placed individually into the single-channel solenoid coil with no water bath. The R_2 relaxation rate of each tube was also measured with a nonselective CPMG sequence (see Figure 2.10). In each case the relaxation rate was determined by fitting a single exponential decay curve to the acquired data with `nlinfit` in Matlab 2010a.

As discussed in the previous section, the relaxivity of a sample in the static dephasing regime can be expressed in terms of the local magnetic dose (LMD) of that sample, which can be measured using MR susceptometry. The complete procedure for performing this measurement is outlined in [5], and will only be summarized here; see Figure 4.3 for an illustration of the procedure.

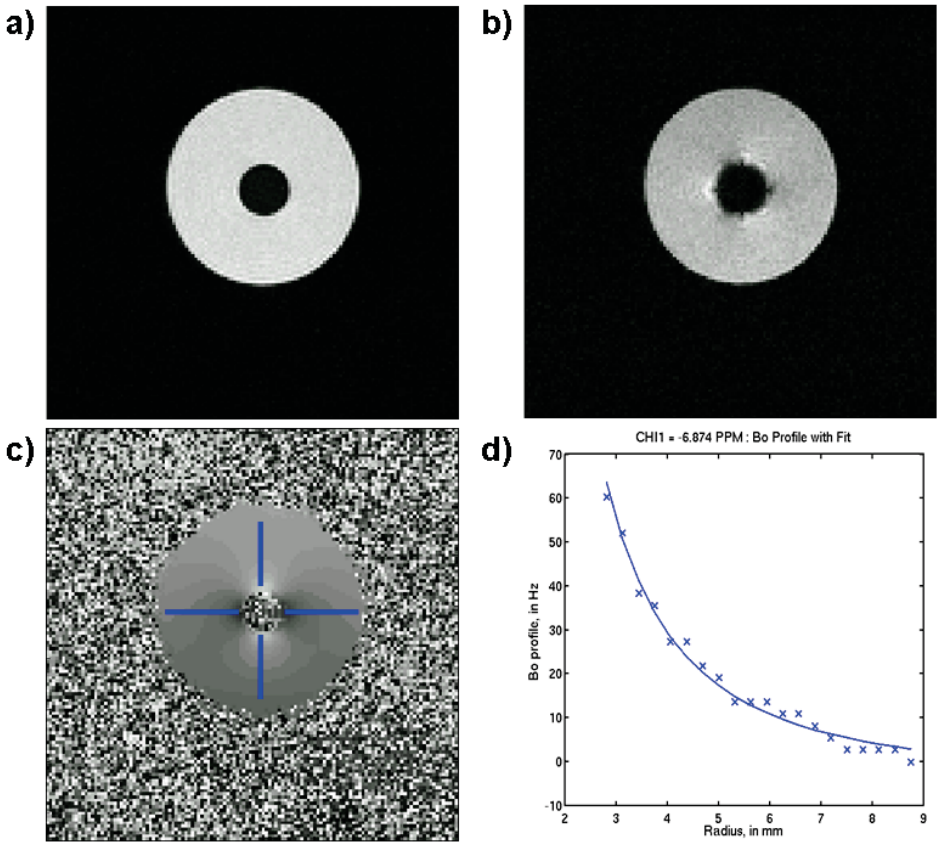

Figure 4.3: Overview of the procedure for performing MR susceptometry and computing the local magnetic dose of the contents of an NMR tube. (a) Image of an NMR tube in a water bath, TE = 10 ms. (b) Same image but with TE = 100 ms; note the loss of signal around the tube. (c) Calculated field map based on all 10 images, and lines along which profiles are extracted for fitting. (d) An appropriate model is fit to the average of the field map along all four profiles, which yields the susceptibility χ.

A single NMR tube is placed within a water-filled cylindrical phantom oriented perpendicular to B_0, which in turn is placed within the solenoid RF coil. A 2D multiple-gradient-echo sequence is used to acquire a field map of the water bath (sequence parameters: repetition time 200 ms, 10 echoes with 10 ms spacing, 15 degree flip angle, one 5 mm slice, 128x128 matrix, 40x40 mm^2 FOV). The field map can be compared to a theoretical description of the inhomogeneity pattern resulting from an arrangement of concentric cylinders, from which the magnetic susceptibility of the sample can be calculated using a script written in Matlab. The LMD is the susceptibility difference, relative to a reference tube containing no magnetic material, multiplied by the field strength B_0 in appropriate units.

4.2.3 Quantification Procedure

As was shown in Figure 3.6, the result of a TurboSPI acquisition is a series of images which are well-suited to relaxometry. At each pixel in the image, the relaxation rate R_2^* is determined by a fit to the linearly decaying exponential region of the spin-echo time course. These fits typically use 200 to 250 of the collected time points, corresponding to a duration of 2 to 2.5 ms. For reasons to be discussed, the complete time course including the peak of the spin echo was not used. Fits to the bulk spin-echo signal are performed similarly, while fits to the bulk FID data use the entire collected time course.

In all cases a two-parameter single exponential decay function (with a baseline of zero) is fitted to the time course with a non-linear least-squares optimization in Matlab 2010a. In the case of TurboSPI, R_2^* mapping of the complete image set requires less than one minute. Fits of background noise pixels are suppressed by only fitting those pixels whose peak intensity is above a given threshold, typically 5% of the image maximum.

To assess the variability of the determined TurboSPI relaxation rates within a sample, time courses are fitted for each pixel, such that a standard deviation can be computed. However, R_2^* values used for comparisons with bulk measurements and for determination of particle relaxivity were obtained from fits to the average of a 5x5 pixel ROI centered on each sample. Obtaining relaxation rates from an average of pixels in this manner has been shown to improve accuracy and precision, as compared

to the average of rates from individual pixels [95]. R_2' for a given sample is calculated by subtracting R_2 (obtained by a CPMG measurement) from the fitted R_2^*.

The relaxivity for each group of samples is determined with a nonlinear regression (using Matlab's `nlinfit` and `nlparci` to obtain a 95% confidence interval for the parameters). All values are reported as the relaxivity ± the half-width of the confidence interval.

4.3 Results

A typical TurboSPI image of five MPIO tubes within the cylindrical holder is shown in Figure 4.4, along with a single-pixel time course and the corresponding fit used to determine R_2^*. As can be seen in the time course, there is significant non-linear-exponential behavior around the peak of the spin echo, and a visible shift of the echo peak away from the nominal echo time. These effects are seen in all of the MPIO samples and, as shown in Figure 4.5, are also present in the samples containing SPIO-loaded cells, albeit to a much smaller degree. The presence of these effects make fits to the entire time course difficult.

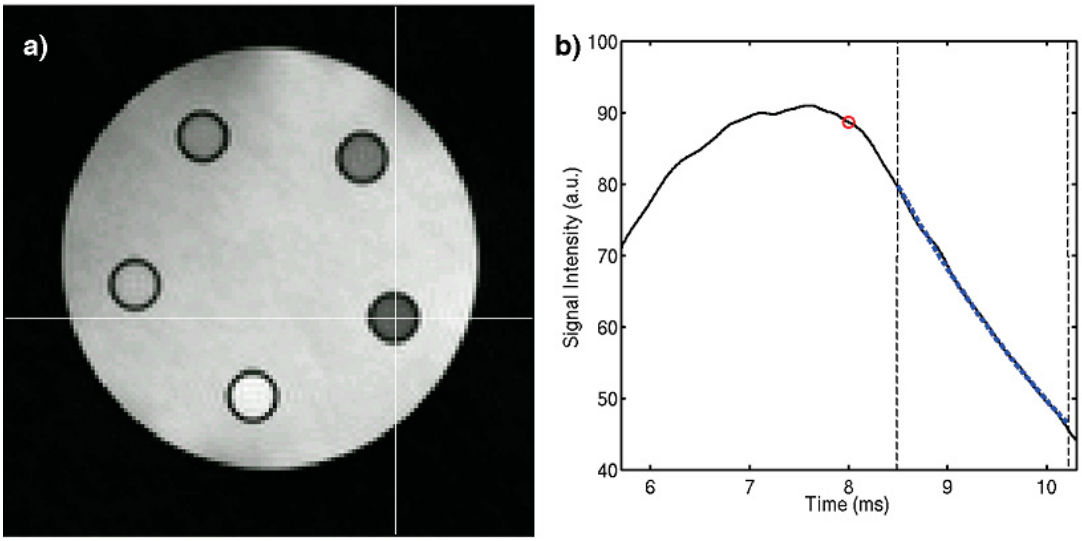

Figure 4.4: R_2^* quantification using TurboSPI. (a) TurboSPI image from t = TE = 8 ms, with crosshairs indicating the pixel of interest. The sample holder contains five tubes of 1.63 μm MPIO in varying concentrations. (b) Time course from indicated pixel, with fitted relaxation curve superimposed (blue line). Thin dashed lines bracket the data used for the fit, and the red circle marks the echo time TE. Note the significant shift in the echo peak away from the nominal echo time.

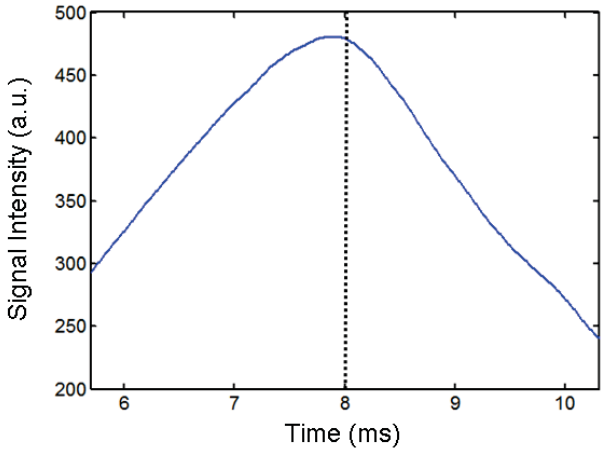

Figure 4.5: TurboSPI time course from a representative sample containing SPIO-loaded C3 cells. Though the echo is not as significantly shifted, the peak still forms before TE = 8 ms (dotted line), and there is some non-linear-exponential behavior.

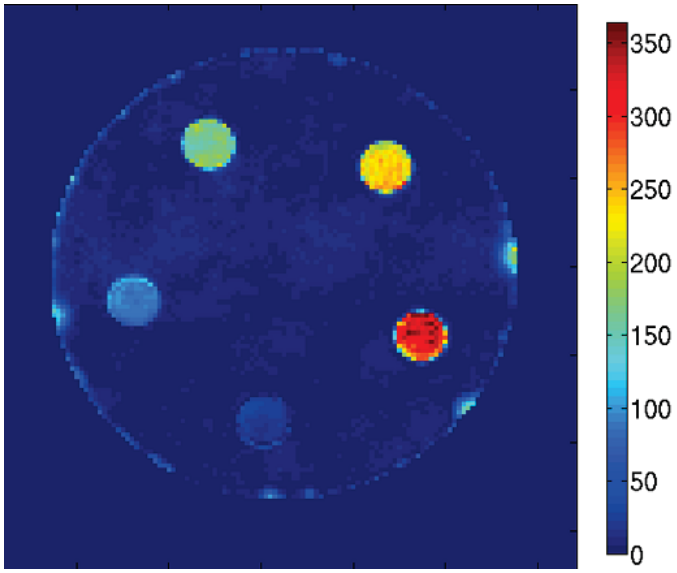

Figure 4.6: Example of an R_2^* map produced by analysis of TurboSPI time course data. In this set of five tubes of 1.63 μm MPIO particles in gel suspension, MPIO concentration increases in the clockwise direction starting with the tube nearest the bottom.

A typical R_2^* map for a group of 1.63 μm MPIO tubes is shown in Figure 4.6. The fitted R_2^* values for each of these tubes are fairly homogeneous, even at moderate iron concentrations. The homogeneity of the R_2^* maps for the loaded cells is also similar.

Figure 4.7 demonstrates the relaxation behavior of the SPIO-loaded C3 cells. As

expected for these relatively large perturbers, the R_2' to R_2 ratio is high, approximately 40:1, indicating that the static dephasing conditions should be met. The bulk R_2' relaxivity as measured by the slice-selective FID experiment was 10.73 $\pm$ 1.6 s^{-1}/mG, while the relaxivity based on the decaying portion of the spin echo was measured as 10.36 $\pm$ 2.9 s^{-1}/mG by the bulk measurement and 10.47 $\pm$ 2.3 s^{-1}/mG by TurboSPI. None of these values are significantly different from the theoretically predicted value of 10.78 s^{-1}/mG for a static dephasing system.

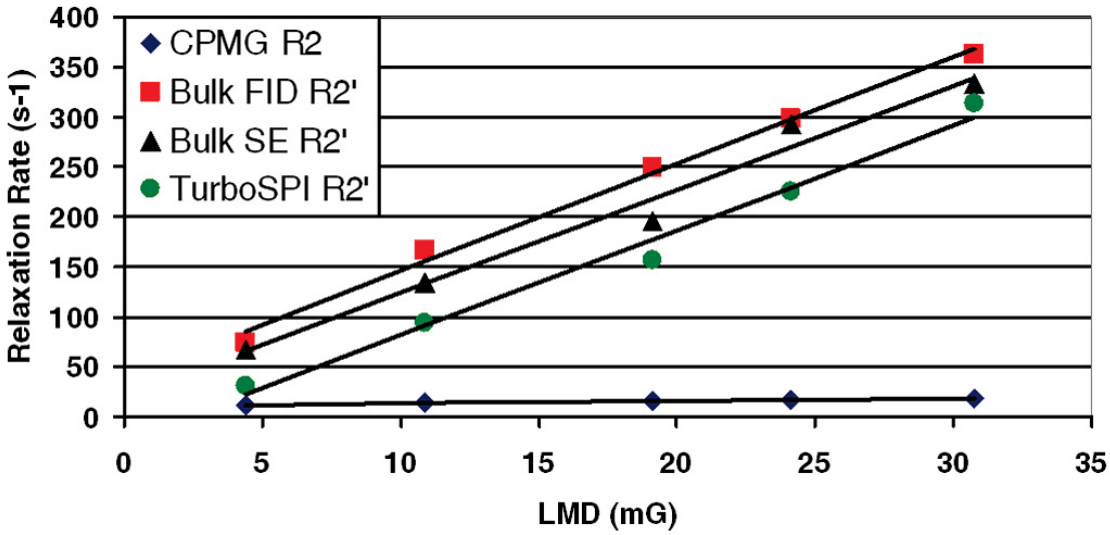

Figure 4.7: Relaxation rates for C3 cells loaded with 35 nm SPIO, as determined by a bulk CPMG measurement (diamonds), slice-selective bulk measurements of the FID (squares) and spin echo (triangles), and 2D TurboSPI (circles). All relaxivities are expressed as a function of the local magnetic dose (LMD) of magnetic material.

The relaxation behavior of the MPIO particles in gel suspension is illustrated in Figure 4.8. As with the loaded cells, the change in the R_2' relaxation rate with iron concentration is greater than the change in R_2, though only by a factor of 10 for the 0.96 μm particles, and by a factor of 13 for the 1.63 μm particles. The relaxivity R_2' from the bulk FID measurement was 10.68 $\pm$ 0.38 s^{-1}/mG for the 0.96 μm particles, which is not significantly different from the theoretically predicted value, and 11.00 $\pm$ 0.18 s^{-1}/mG for the 1.63μm particles, which is different from 10.78 at the $\alpha = 0.05$ level but not at $\alpha = 0.1$.

However, unlike the SPIO-loaded cells, the R_2' relaxivities measured with the bulk spin-echo experiment and with TurboSPI are significantly less than the static

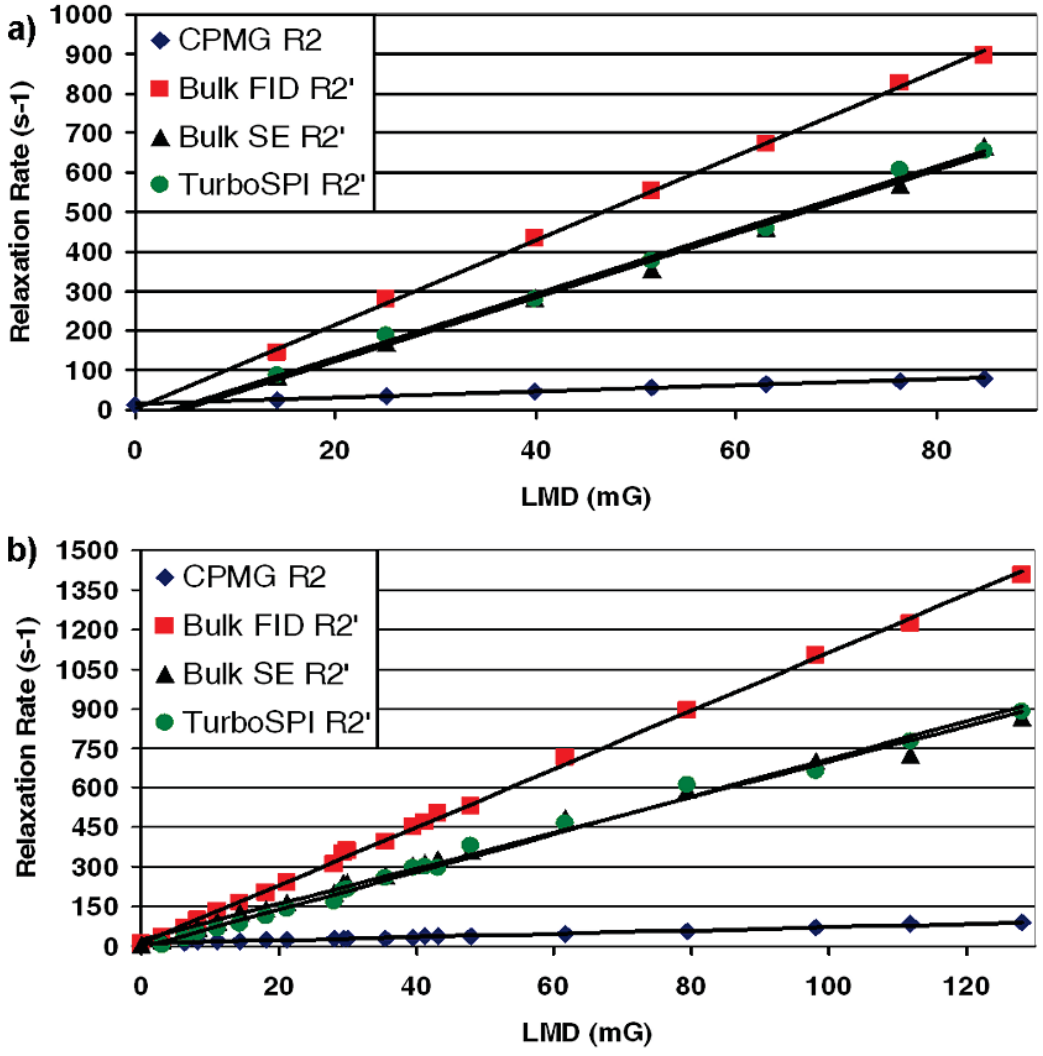

Figure 4.8: Relaxation rates for (a) 0.96 μm and (b) 1.63 μm MPIO particles in gel suspension, as determined by a bulk CPMG measurement (diamonds), by slice-selective bulk measurements during the FID (squares) and following a spin echo (triangles), and by 2D TurboSPI (circles). All relaxivities are expressed as a function of LMD.

dephasing prediction for both sizes of particles. The relaxivity of the 0.96 μm particles was determined to be 8.03 $\pm$ 0.65 s^{-1}/mG by the bulk measurement and 8.06 $\pm$ 0.68 s^{-1}/mG by TurboSPI, while the 1.63 μm particles were found to have a relaxivity of 6.74 $\pm$ 0.30 s^{-1}/mG with the bulk measurement and 7.13 $\pm$ 0.31 s^{-1}/mG with TurboSPI. In both cases the bulk spin-echo relaxivity is not significantly different from that obtained by TurboSPI.

The relaxivities obtained from each measurement and for all three types of perturbers are summarized in Table 4.1. Table 4.1 also lists the local magnetic dose per unit iron concentration for each particle, with cellular iron concentration based on UV/VIS spectrophotometry and MPIO concentrations based on the volume of particles used in preparing the samples.

	0.96 μm MPIO	1.63 μm MPIO	Loaded cells
R_2 from CPMG	0.79 $\pm$ 0.08	0.56 $\pm$ 0.03	0.25 $\pm$ 0.04
R_2' during FID (Bulk)	10.68 $\pm$ 0.38	11.00 $\pm$ 0.18	10.73 $\pm$ 1.6
R_2' after Spin Echo (Bulk)	8.03 $\pm$ 0.65	6.74 $\pm$ 0.30	10.36 $\pm$ 2.9
R_2' after Spin Echo (TurboSPI)	8.06 $\pm$ 0.68	7.13 $\pm$ 0.31	10.47 $\pm$ 2.3
LMD per unit iron concentration	3.01	1.00	0.82

Table 4.1: Relaxivities and local magnetic dose (LMD) per unit iron concentration for all sample groups. All relaxivities are in units of s^{-1}/mG and are reported with 95% confidence interval bounds. Local magnetic doses are in units of mG / (μg/mL).

The variability of R_2^* values obtained from fits to individual TurboSPI pixels is shown in Figure 4.9, which plots the standard deviation of the 25 fitted R_2^* values within a particular ROI as a function of sample T_2 for all of the MPIO samples imaged. There is a significant increase in standard deviation as T_2 approaches the echo time TE, underscoring the requirement that R_2 be small to ensure acceptable quantification. Note that it is not sufficient that R_2 be much less than R_2', since TurboSPI's high sampling rate around the spin echo allows accurate fits to very large R_2^*. However, this assumes that there is sufficient signal at the echo peak, which is only the case if R_2 is not too large; beyond a certain threshold, the SNR is too low to guarantee accurate fits, especially if there is some residual ghosting or blurring of background signal.

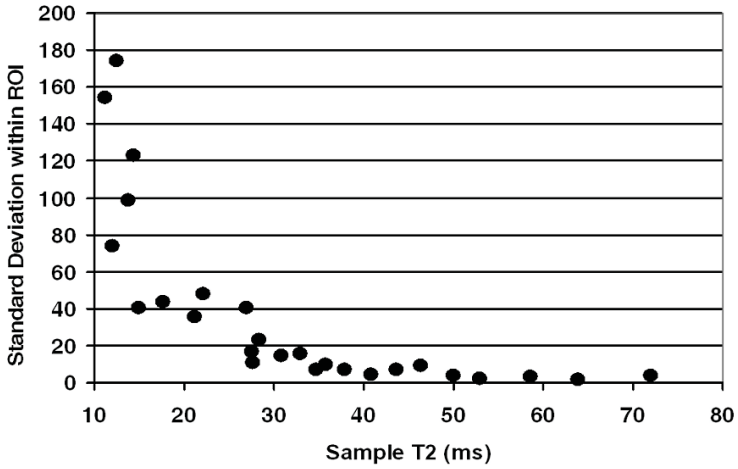

Figure 4.9: Standard deviation of R_2^* values obtained from single-pixel fits to TurboSPI data (TE = 8 ms) within a 25-pixel ROI centered on each MPIO sample.

4.4 Discussion and Summary

Reliable *in vitro* quantification of both MPIO particle and SPIO-loaded cells over a wide dynamic range has been demonstrated with TurboSPI. In all cases the measured R_2^* relaxation rates are linearly proportional to the local magnetic dose (LMD) as measured by MR susceptometry, and are comparable to the rates obtained using a bulk spin-echo measurement. This is true even at high iron concentrations with relaxation rates in excess of $R_2^* = 1000s^{-1}$, so long as R_2 is sufficiently small.

For MPIO particles, fit quality begins to degrade significantly around $T_2 = 2TE =$ 16 ms, or $R_2 > 60$. Assuming a similar maximum R_2 for SPIO-loaded cells such as those used in this study, which have an $R_2' : R_2$ ratio of 40, it should be possible to quantify relaxation rates up to $2400s^{-1}$, corresponding to an iron concentration of 270 μg/mL Fe, or 27 million cells/mL at a loading level of 10 pg/cell. This represents a reasonable loading level and cell density that could be used in longitudinal studies of cellular therapy [7].

While the bulk FID relaxivities of all three sample groups satisfy the prediction of the static dephasing regime theory ($R_2' = 10.78 \cdot LMD$), only in the case of SPIO-loaded cells was this prediction also satisfied for the spin-echo case. For both types of MPIO particles, the R_2' relaxation measured after a spin echo was significantly lower than predicted. This indicates that these particles did not satisfy the conditions of

the SDR for the spin-echo experiment, as outlined in section 4.1.4. Indeed, Table 4.2 summarizes the relevant properties of all three particle types and shows that, while the FID criterion is met in all cases, only the SPIO-loaded cells also meet the SE criterion. While data are only listed for one sample from each group, these conclusions should remain valid for all other samples in that group due to the small effect of changes in ζ.

	0.96 μm MPIO	1.63 μm MPIO	C3 cells (13 μm)
Iron concentration	20 μg/mL	30 μg/mL	8 μg/mL
Volume fraction, ζ	1.5 x 10^{-5}	1.9 x 10^{-5}	1.44 x 10^{-3}
Characteristic freq, $\delta\omega$	8.48 x 10^6 s^{-1}	3.88 x 10^6 s^{-1}	2.72 x 10^4 s^{-1}
Characteristic time, t_c	0.097 μs	0.213 μs	0.030 ms
Diffusion time, t_D	0.0384 ms	0.111 ms	16.7 ms
FID SD criterion: $\delta\omega\zeta^{1/3}t_D$	8.03	11.49	25.1
$\delta\omega\zeta^{1/3}t_D \gg 1$?	YES	YES	YES
FID relaxivity (Bulk)	10.68 s^{-1}/mG	11.00 s^{-1}/mG	10.73 s^{-1}/mG
$TE_c = 2.9 \cdot (t_c \cdot t_D/R_2')^{1/3}$	71 μs	160 μs	4.24 ms
$TE = 8\text{ms} \approx TE_c$?	NO	NO	YES
Spin-echo relaxivity (TurboSPI)	8.06 s^{-1}/mG	7.43 s^{-1}/mG	10.47 s^{-1}/mG

Table 4.2: Relevant properties for representative tubes from each group of samples. Volume fractions for MPIO samples are calculated based on iron concentration, particle mass as stated by the manufacturer, and particle mean volume. Volume fractions for cells are based upon measured concentration of cells and the average cell size. The characteristic frequency depends on the susceptibility of an individual perturber, obtained by dividing the tube susceptibility by the volume fraction.

These differences in relaxation behavior between MPIO particles and SPIO-loaded cells are primarily due to the particle size, which leads to a significantly longer diffusion time (since t_D varies as R^2). As well, the smaller characteristic frequency $\delta\omega$ of the loaded cells reduces diffusive attenuation in the spin-echo case. It should be noted that, even though the critical echo time for the cell sample listed in Table 4.2 is much larger than that for the MPIO samples, it is still less than the nominal echo time TE. However, static dephasing behavior is still observed, suggesting that the condition derived above, which requires TE $\ll$ TE$_c$, might be too restrictive. In practice it may be sufficient that TE be comparable to TE$_c$.

The results of this study are consistent with the findings of Seevinck and colleagues [62] in which relaxation rates of holmium microspheres were measured with GESFIDE and GESSE, and significantly lower relaxivity was observed in the spin-echo experiment. They also observed that the holmium microspheres satisfied the SDR condition

in the FID case but not the spin-echo case, and concluded that quantification using FID signal was more reliable. We contend that spin-echo-based quantification such as that allowed by TurboSPI is no less reliable than FID based quantification, as evidenced by the strong linear relationship between R_2' and local magnetic dose in all three sample groups; it is simply necessary to measure the reduced relaxivity beforehand under controlled conditions.

The differences between the relaxivity of loaded cells and MPIO can also be used to support one of the main assumptions of this work; namely, that the SPIO-loaded C3 cells can be treated as a single magnetic perturber with a radius equal to that of the cell. While this seems a reasonable assumption in the event that individual SPIO particles within a cell are evenly distributed throughout the cytoplasm, it might also be possible for particles to be sequestered into one or more endosomes, in which case a more appropriate effective size for the perturber would be that of the endosome. However, if this were true for the C3 cells tested in this experiment, we would anticipate behavior similar to that of the MPIOs, which have physical sizes comparable to a single endosome. Since this was not observed, our assumption of a cell-sized perturber seems valid. This can be further explored with more advanced models of signal relaxation.

All of the spin-echo relaxometry described in this chapter used only the decaying portion of the time course, as shown in Figure 4.4. The significant shifting of the echo peak (seen in Figure 4.4 and observed in all of the MPIO samples) meant that the rising portion of the spin echo was largely outside the data acquisition window, and the portion which remained was unsuitable for curve fitting due to non-linear exponential behavior around the peak. While increasing the acquisition window would permit more of the rising portion to be acquired, this would also increase the echo time of the sequence and impede quantification of large R_2^* samples. The non-linear behavior would also remain, and it cannot be fitted to the existing model of relaxation described in this chapter. The methods for dealing with this unanticipated behavior, and the opportunities it presents, are the subject of the next chapter.

Chapter 5

Extended Analytical Description of SPIO Relaxation

The peak shifting and non-linear behavior observed by TurboSPI measurements represent more than impediments to quantification. The model presented in the previous chapter that describes relaxation in the static dephasing regime does not account for these effects, which are even observed (albeit to a much smaller degree) in the SPIO-loaded cells that satisfy both static dephasing conditions. The ability of TurboSPI to monitor spin echo relaxation at very high temporal resolution provides an opportunity to explore these effects and gain a greater understanding of the underlying physical mechanisms, which in turn will provide opportunities to improve the quantification of SPIO.

Spurred by the experimental observations provided by TurboSPI, an expanded analytical description of the MRI signal in the presence of SPIO can be developed, which may allow additional information to be obtained from TurboSPI data. This section will describe some of the existing theoretical models that apply outside the static dephasing regime, outline their predictions and applicability, and develop a related model suitable for fitting to TurboSPI data. The results of this model will then be compared to simulations and experimental data, and their application to a number of practical outcomes will be considered.

5.1 Relaxation Outside of the Static Dephasing Regime

Yablonskiy and Haacke's description of relaxation in the static dephasing regime [91] is appealing for studies of SPIO-induced relaxation because of its simple prediction of a relaxivity that varies linearly with iron concentration. Its relevance to the case of SPIO-loaded cells in an FID experiment has been well documented, both in previous studies [5, 55] and in the current work. However, its inability to fully describe the behavior observed in TurboSPI measurements limits its applicability to this technique.

The static dephasing model described in the previous chapter does allow for non-linear exponential behavior near the peak of the spin echo, but only for a duration proportional to $1/\delta\omega$, which in the case of MPIO particles is on the order of 10^{-6} s (see Table 4.2). The extent of non-linear behavior observed in TurboSPI measurements is far greater, usually on the order of 10^{-3} s, and this extent increases with $\delta\omega$; it is far less evident in the case of SPIO-loaded cells which have $1/\delta\omega \approx 10^{-4}$ s. The static dephasing model also provides no explicit mechanism for the peak shifting observed in all measured samples. While T_2 relaxation will lead to a spin echo that peaks earlier in time, the T_2 of the samples tested in Chapter 4 is not short enough to account for all of the observed shift.

The static dephasing model relies on the assumption that diffusion can be neglected entirely, which may not be the case even with large effective particle sizes. This is especially true in the case of a spin echo, where dephasing effects are refocusing and diffusion becomes the dominant mechanism of signal attenuation. Indeed, the presence of diffusion is likely a central factor in the differences between this model's predictions and the observations of TurboSPI. The logical course of action is then to consider a model which describes the MRI signal near magnetic perturbations and in the presence of diffusion, and a number of such theoretical models have been proposed in the literature.

5.1.1 Predicted Behavior

While an exact analytical description of relaxation behavior will obviously require a realistic mathematical model, it is actually possible to derive the broad functional form of some key results using only geometrical arguments. This process will yield predictions of which variables contribute to signal decay and in what proportion.

The formulation of the static dephasing model outlined in Chapter 4 used a function f to contain all of the terms which depended on the perturber geometry, and which plays a central role in determining the relaxation rate R_2'. This function was defined such that

$$s(t) \propto \left(1 - \frac{v(R)}{V - v(R)} f\right)^N \tag{5.1}$$

If $f = 0$, each of the N perturbers contributes a factor of 1 to the overall signal. For non-zero f, this contribution is reduced by the fraction of the total volume $V(R)$

which does not contribute signal. This in turn depends on the volume $fv(R)$ throughout which signal is fully dephased [65]. f therefore represents a scaling factor that quantifies how the perturber affects a region larger than its physical size.

From symmetry considerations we can assume this region has the same geometry as the perturber, and define r to be its radius. For the case of spherical perturbers, f can then be estimated as

$$f = \frac{4\pi r^3/3}{4\pi R^3/3} = \frac{r^3}{R^3} \tag{5.2}$$

Ignoring numerical coefficients and the angular dependence, the frequency offset experienced by a diffusing proton a distance r from a perturber of radius R is

$$\omega(r) = \delta\omega \frac{R^3}{r^3} \tag{5.3}$$

During an FID experiment, if the particle's position does not change significantly, then the phase it accumulates during a time t is

$$\phi = \omega(r)t \tag{5.4}$$

For the proton to be significantly dephased, such that it does not contribute to the MRI signal, we require that the accumulated phase be on the order of 1. (We could equivalently assume a phase of 2π; this would only add an overall scaling factor without changing the functional form.) This means

$$\omega(r)t \sim 1$$
$$\delta\omega \frac{R^3}{r^3} t \sim 1$$
$$\delta\omega R^3 t \sim r^3$$

Comparing this to our estimate for f gives

$$f \propto \delta\omega t \tag{5.5}$$

This matches the asymptotic behavior of f for large values of t, as given by equation (4.15), verifying the usefulness of this geometrical argument.

The more interesting result is the case of a spin echo, in which the dominant effect is not dephasing, but instead the attenuation of signal due to diffusion through local field gradients. Around a spherical perturber these gradients are proportional to

$$g = \frac{d\omega(r)}{dr} \propto \delta\omega \frac{R^3}{r^4} \tag{5.6}$$

and the attenuation due to diffusion processes becomes significant under the condition $Dg^2TE^3 \sim 1$ [96]. Substituting, and using the dimensionless variable $\lambda = \frac{D}{R^2\delta\omega}$,

$$\begin{aligned} D\left(\delta\omega\frac{R^3}{r^4}\right)^2 TE^3 &\sim 1 \\ D\delta\omega^2 R^6 TE^3 &\sim r^8 \\ \lambda(\delta\omega TE)^3 R^8 &\sim r^8 \\ \lambda^{3/8}(\delta\omega TE)^{9/8} R^3 &\sim r^3 \end{aligned}$$

The predicted form of f at the peak of a spin echo is therefore

$$f = \frac{r^3}{R^3} = \lambda^{3/8}(\delta\omega TE)^{9/8} \tag{5.7}$$

Now there is a dependence on both the diffusion coefficient and the particle size (through λ), as well as a nonlinear dependence on the echo time TE. It is anticipated that this functional form will reappear in any satisfactory model of spin-echo relaxation near the static dephasing regime.

5.1.2 Overview of Alternate Models

Numerous models have been proposed which incorporate the effects of diffusion on relaxation, but for our purposes, only a small number are of interest. Models such as the Gaussian phase approximation of Sukstanskii and Yablonskiy [97] and several other models [98, 99] are valid only in the motional narrowing regime, in which diffusion is the dominant mechanism of signal loss even in the FID case. While these have applications in the study of free SPIO particles, such systems are not well suited to imaging with TurboSPI because of their low $R_2' : R_2$ ratio.

Jensen and Chandra [100] developed a model using a strong field approximation which, if satisfied, ensures that the magnetic field gradient experienced by a diffusing particle before its signal decays is approximately linear. This assumption leads to a function G describing the decay of signal which, in general, must be evaluated numerically. However, from this function a number of useful asymptotic forms can be derived, including predictions of signal behavior during FID, spin-echo and multi-echo experiments, for both spherical and cylindrical geometries. For example, for an FID experiment or far from the peak of a spin echo, the signal in a system of spherical

perturbers (ignoring the scaling factor of ρ) is

$$\begin{aligned} s(T) &= e^{-4pi\gamma\chi B_0\zeta(0.4031-0.0533i)T} \\ &= e^{-\zeta(1.21-0.16i)\delta\omega T} \end{aligned} \tag{5.8}$$

in agreement with the static dephasing model (equation (4.15)). At the peak of a spin echo, the signal in this model is

$$\begin{aligned} s(T) &= e^{-2.2184\zeta(3DTE^3\delta\omega^2/4R^2)^{3/8}} \\ &= e^{-1.99\zeta(\delta\omega TE)^{9/8}\lambda^{3/8}} \end{aligned} \tag{5.9}$$

As expected from geometrical considerations, this expression is non-linear in TE and has contributions from diffusion and particle size with the correct exponents. However, these predictions are only valid far from the echo peak or at the echo time t=TE; a complete form of the echo shape close to TE is not provided by this model. A generalization of Jensen and Chandra's original model is given in [101] which extends the results to multi-echo sequences with non-uniform echo spacing, but still only defines the signal at echo peaks.

A model based on a strong collision approximation has been used by Bauer and colleagues [89, 102] to analytically describe relaxation for a wide range of diffusion values. While significantly more complex, this model does have the advantage of providing descriptions of the signal in the intermediate motion regime, which many other models do not describe; a Laplace transform is used to translate static dephasing results into other diffusion regimes. Good agreement has been shown between this model and experimental or simulated data [103]; however, as with most other models, only the observed spin-echo relaxation rate is predicted, not the complete time course as would be ideal for use with TurboSPI [104].

One model which does provide a complete analytical description of the MRI signal during a spin echo is that of Kiselev and Posse [65]. This model is valid in the slow diffusion regime where diffusion effects are limited, though still present, and also assumes a linear field variation in the vicinity of a diffusing particle. Asymptotic forms derived from this model match those given by Yablonskiy and Jensen, but include correction terms to account for the effects of diffusion during the FID. As well, the model describes shifting of the echo peak similar to that observed in TurboSPI measurements.

There are two drawbacks of Kiselev's model in the context of this work. First, it was formulated to model signal changes in microvasculature, and as such is only valid for systems of randomly oriented cylinders; it is not immediately applicable to the case of SPIO-labeled cells. Second, the full time course is described as an integral over the random distribution of perturbers, which is difficult to use in practice. However, these limitations are not insurmountable, and otherwise this model is an excellent candidate for exploring the relaxation behavior observed by TurboSPI.

5.2 Extension of the Slow Diffusion Model

This section will describe in more detail the slow diffusion model of Kiselev and Posse, first described in [65], which is extended here for the first time to the case of spherical perturbers. An approximate form is also derived for the spin-echo signal time course which will be more suitable for fitting to experimental data than the full analytical expression.

5.2.1 Description of the Model

In many ways the model of Kiselev and Posse is similar to that of Yablonskiy and Haacke, but simply generalized to account for cases when diffusion is slow but not negligible. The fundamental difference is that, instead of starting with the signal

$$s(t) = \frac{1}{V}\rho \int_{V_0} e^{-i\omega(\vec{r})t} d\vec{r} \tag{5.10}$$

we require a more general form that incorporates the effects of diffusion. This can be easily achieved by noticing that $e^{-i\omega(\vec{r})t}$ is simply the solution to the Bloch equations (equation (2.33)) for signal in the transverse plane and in the absence of relaxation. To add the effects of diffusion we replace the Bloch equation with the Bloch-Torrey equation, which has an additional term containing D [105]. If ψ represents the transverse magnetization, the Bloch-Torrey equation is

$$\begin{aligned} \frac{d\psi}{dt} &= D\nabla^2\psi - i\gamma\vec{r}\cdot\vec{G}\psi \\ &= D\nabla^2\psi - i\omega(\vec{r})\psi \end{aligned} \tag{5.11}$$

This form assumes D is isotropic; in the case of anisotropic diffusion D becomes a tensor and the solution is more complicated.

The solution to the Bloch-Torrey equation can be obtained by various methods, one of which is outlined in Appendix A.3. In the case of an FID experiment it can be shown that this solution is

$$\psi(\vec{r},t) = exp\left(-i\omega(\vec{r})t - \frac{1}{3}D[\nabla\omega(r)]^2t^3\right) \tag{5.12}$$

while the solution for a spin-echo experiment is

$$\psi(\vec{r},t) = exp\left(-i\omega(\vec{r})(t-TE) - \frac{1}{3}D[\nabla\omega(r)]^2t^3F(TE/t)\right) \tag{5.13}$$

The F function in the second term is defined as

$$F(z) = 1 - \frac{3}{2}z^2 + \frac{3}{4}z^3 \tag{5.14}$$

and its derivation is also discussed in Appendix A.3. This expression is valid for t $>$ TE/2; before this time the FID solution applies. In both cases, when $D = 0$ the solution reduces to the expression used in the static dephasing model as expected.

The signal can now be written as a function of Ψ, the solution for the magnetization in the presence of all N perturbers:

$$s(t) = \rho(1-\zeta)\int \Psi(\vec{r},t)\prod_n^N d\Omega_n \tag{5.15}$$

As before, ρ incorporates all non-relaxation effects, and $(1-\zeta)$ reflects the contribution of protons in the volume not occupied by the perturbers. The $d\Omega_n$ encompasses all of the integrations necessary to average over the position of the n^{th} object in the network. If the volume fraction ζ occupied by the magnetic perturbers is small, we are justified in factoring Ψ into separate contributions due to each object,

$$\Psi(\vec{r},t) = \prod_n^N \psi_n(\vec{r},t) \tag{5.16}$$

such that the average over the entire network is simply the product of integrals due to a single object:

$$\begin{aligned} s(t) &= \rho(1-\zeta)\prod_n^N \frac{1}{V}\int \psi_n(\vec{r},t)d\Omega_n \\ &= \rho(1-\zeta)\prod_n^N \left[1 - \frac{1}{V}\int(1-\psi_n(\vec{r},t))d\Omega_n\right] \end{aligned} \tag{5.17}$$

At this point we can define an f function as was done in Yablonskiy and Haacke's model, though in this case

$$f = \frac{1}{v(R)} \int d\Omega (1 - \psi(\vec{r}, t)) \tag{5.18}$$

But the signal equation in terms of f is exactly the same as in the static dephasing model,

$$s(t) = \rho(1-\zeta)\left(1 - \frac{\zeta f}{N}\right)^N = \rho(1-\zeta)e^{-\zeta f} \tag{5.19}$$

As before, to determine a form for f we must use our knowledge of the system geometry and the type of experiment (FID or spin-echo).

FID Experiment

As described in [65], to obtain f for the FID case we substitute the expression (4.10) for the frequency shift into equation (5.12) and integrate ψ over the network of perturbers. We will first need to compute the gradient of ω for substitution in equation (5.12).

$$\begin{aligned}
\nabla\omega &= \frac{\partial\omega}{\partial r}\hat{r} + \frac{1}{r}\frac{\partial\omega}{\partial\theta}\hat{\theta} \\
&= -3\delta\omega\frac{R^3}{r^4}(3\cos^2\theta - 1)\hat{r} - 6\delta\omega\frac{R^3}{r^4}\cos\theta\sin\theta\hat{\theta} \\
&= -3\delta\omega\frac{R^3}{r^4}((3\cos^2\theta - 1)\hat{r} + 2\cos\theta\sin\theta\hat{\theta})
\end{aligned} \tag{5.20}$$

Squaring this expression gives

$$\begin{aligned}
[\nabla\omega]^2 &= 9\delta\omega^2\frac{R^6}{r^8}((3\cos^2\theta - 1)^2 + (2\cos\theta\sin\theta)^2) \\
&= 9\delta\omega^2\frac{R^6}{r^8}(9\cos^4\theta - 6\cos^2\theta + 1 + 4\cos^2\theta\sin^2\theta) \\
&= 9\delta\omega^2\frac{R^6}{r^8}(5\cos^4\theta + 4\cos^2\theta(\cos^2\theta + \sin^2\theta) - 6\cos^2\theta + 1) \\
&= 9\delta\omega^2\frac{R^6}{r^8}(5\cos^4\theta - 2\cos^2\theta + 1)
\end{aligned} \tag{5.21}$$

This can now be substituted into equation (5.12) and integrated using the procedure described in [65]. However, some minor modifications are necessary due to the change in geometry from cylinders to spheres. As in the static dephasing case we use

the volume element $r^2 dr \sin\theta d\theta d\phi$ to obtain

$$f^{FID} = \frac{3}{4\pi R^3}\int_0^{\pi} d\theta \sin\theta \int_R^{\infty} r^2 dr \int_0^{2\pi} 1 - \exp\left(-i\delta\omega\frac{R^3}{r^3}(3\cos^2\theta - 1)t \right.$$
$$\left. -3D\delta\omega^2\frac{R^6}{r^8}(5\cos^4\theta - 2\cos^2\theta + 1)t^3\right) d\phi \qquad (5.22)$$

We will make several variable substitutions to bring this function into alignment with the notation of Kiselev. First, the time t is replaced with the dimensionless quantity $\tau = \delta\omega t$. The diffusion coefficient is also modified into a dimensionless parameter,

$$\lambda = \frac{D}{R^2\delta\omega} \qquad (5.23)$$

Finally, the variable of integration is changed from r to $u = R^3/r^3$, and $r^2 dr$ becomes $\frac{-R^3 du}{3u^2}$. The result is

$$f^{FID} = \int_0^{\pi} d\theta \frac{\sin\theta}{2}\int_0^1 \frac{du}{u^2}\, 1 - \exp\left(-iu(3\cos^2\theta - 1)\tau - 3\lambda\tau^3 u^{8/3}(5\cos^4\theta - 2\cos^2\theta + 1)\right) \qquad (5.24)$$

This resembles the form of f^{FID} in Kiselev's original model for cylinders (equation (25) in [65]), though it is no longer possible to express the imaginary term as a Bessel function. It can also be seen that, in the case $\lambda = 0$ this formula becomes equivalent to the static dephasing case given in equation (4.13), and it will have the same asymptotic forms.

In the case $\tau \ll 1$ a λ-dependent correction to the static dephasing result can be calculated, such that

$$f^{FID} = \frac{2}{5}\tau^2 + \frac{12}{5}\lambda\tau^3 \qquad (5.25)$$

In principle a similar correction to the $\tau \gg 1$ version of f should be feasible, but the resulting integral is more challenging and at this time has not been evaluated.

Spin-Echo Experiment

By comparing the differences between equations (5.12) and (5.13), the form of f for the spin-echo case can be determined straightforwardly. Since these differences involve only the presence of the function $F(TE/t)$ in the decaying exponential and the $(t - TE)$ dependence of the oscillatory term, we can simply alter the resulting terms in f^{FID} by inspection. Introducing the dimensionless variable $\tau_E = \delta\omega TE$ to

replace the echo time, we find

$$f^{SE} = \int_0^{\pi} d\theta \frac{\sin\theta}{2} \int_0^1 \frac{du}{u^2} \left(1 - \exp\left(-iu(3\cos^2\theta - 1)(\tau_E - \tau) \right.\right. \\ \left.\left. -3\lambda\tau^3 F(\tau_E/\tau) u^{8/3}(5\cos^4\theta - 2\cos^2\theta + 1)\right)\right) \tag{5.26}$$

As in the FID case, asymptotic forms for this expression can be computed. Far from the echo peak, t $\gg$ TE and $F \to 1$, which means that the forms computed in the FID case will be recovered. The other case of interest is t=TE (or equivalently, $\tau = \tau_E$), in which case the oscillatory exponential term vanishes and the integral is simplified.

Though not of significant practical use, the value of f at $\tau = \tau_E$ with $\tau_E \ll 1$ can be found to be

$$\begin{aligned} f^{SE} &= \int_0^{\pi} d\theta \frac{\sin\theta}{2} \int_0^1 \frac{du}{u^2} 1 - exp\left(-\frac{3}{4}\lambda\tau_E^3 u^{8/3}(5\cos^4\theta - 2\cos^2\theta + 1)\right) \\ &= \frac{1}{2}\int_{-1}^1 dU \int_0^1 \frac{du}{u^2}\left(\frac{3}{4}\lambda\tau_E^3 u^{8/3}(5U^4 - 2U^2 + 1)\right) \\ &= \frac{1}{2}\int_{-1}^1 \frac{9}{20}\lambda\tau_E^3(5U^4 - 2U^2 + 1)dU \\ &= \frac{9}{20}\lambda\tau_E^3(1 - \frac{2}{3} + 1) \\ &= \frac{3}{5}\lambda\tau_E^3 \end{aligned} \tag{5.27}$$

which, as in the case of the cylinder model [65] has a form similar to the correction on f^{FID} but differs by a factor of $F(1) = \frac{1}{4}$.

Of more interest is the form for f^{SE} at $\tau = \tau_E$ when $\tau_E \gg 1$, which is generally the case for SPIO particles imaged at most experimentally relevant echo times. To evaluate this integral we make the substitution $U = \cos\theta$ as before, and also $v = u\tau_E^{9/8}$, such that $\tau_E^3 u^{8/3} = v^{8/3}$. This substitution is suggested by the geometrical considerations outlined in section 5.1.1.

$$f^{SE}(\tau_E) = \frac{1}{2}\int_{-1}^1 dx \int_0^{\infty} \tau_E^{9/8}\frac{dv}{v^2} 1 - exp\left(-\frac{3}{4}\lambda v^{8/3}(5x^4 - 2x^2 + 1)\right) \tag{5.28}$$

The upper limit of the integral over v has been extended to infinity to assist in evaluation; if desired for completeness, a correction term to account for this can be computed, which is of order $\tau_E^{1/8}$.

Though it may appear that this integral diverges, it can be evaluated by defining the integral

$$A = \int_0^\infty dv v^{-2+\epsilon} e^{-cv^{8/3}} \tag{5.29}$$

such that substituting $\epsilon = 0, c = \frac{3}{4}\lambda(5x^4 - 2x^2 + 1)$ will recover the second term in the equation above, while $\epsilon = 0, c = 0$ gives the first term. Using the definition of the gamma function,

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt \tag{5.30}$$

we can determine a form for A in terms of c and ϵ by using the substitution $t = cv^{8/3}$ and $dt = \frac{8}{3}cv^{5/3}dv$. Then $\frac{dv}{v^{2+\epsilon}} = \frac{3}{8}dtc^{3/8+3\epsilon/8}t^{-11/8+3\epsilon/8}$, and

$$\begin{aligned} A &= \int_0^\infty dv v^{-2+\epsilon} e^{-cv^{8/3}} \\ &= \frac{3}{8} c^{3/8+3\epsilon/8} \int_0^\infty t^{-11/8+3\epsilon/8} e^{-t} dt \\ &= \frac{3}{8} c^{3/8+3\epsilon/8} \Gamma(-3/8 + 3\epsilon/8) \\ &= c^{3/8+3\epsilon/8} \Gamma(5/8 + 3\epsilon/8) \end{aligned} \tag{5.31}$$

Here we have also made use of the property $\Gamma(z+1) = z\Gamma(z)$. Setting $\epsilon = 0$ and substituting for c,

$$f^{SE}(\tau_E) = \frac{1}{2}\tau_E^{9/8}\Gamma(5/8)(\frac{3}{4}\lambda)^{3/8} \int_{-1}^{1} (5x^4 - 2x^2 + 1)^{3/8} dx \tag{5.32}$$

The integral over x can be evaluated numerically and yields a result of 2.162. Gathering coefficients,

$$f^{SE}(\tau_E) = 1.39\tau_E^{9/8}\lambda^{3/8} \tag{5.33}$$

This has the functional form predicted by geometrical argument in section 5.1.1, and also agrees with the form given by Jensen and Chandra [100] with a slightly different numerical coefficient.

5.2.2 Approximate Forms

The complete expression for the evolution of the signal during the spin echo (equation (5.26)) is valuable for comparison with simulated and experimental data, as will be described in the next section. However, for practical purposes a simpler expression is required that can be easily fitted to experimental data, in order to extract

estimates of key parameters such as $\delta\omega$ and the volume fraction ζ. While regions far from the echo peak can be fitted by the asymptotic form given in equation (4.15), an approximate form that remains valid near the peak is also desirable.

The most straightforward way to approximate a complex function like f^{SE} is with a Taylor series,

$$f(\tau) \approx f(\tau_0) + \frac{df}{d\tau}|_{\tau=t_0}(\tau-\tau_0) + \frac{1}{2}\frac{d^2f}{d\tau^2}|_{\tau=\tau_0}(\tau-\tau_0)^2 + ... \tag{5.34}$$

with the echo time τ_E a logical choice for the point τ_0 about which to expand the function. The first term in this Taylor series is, indeed, the expression just derived for $f^{SE}(\tau_E)$. Further terms in the expansion can be obtained by differentiation of equation (5.26). Since neither integral is time-dependent, this involves differentiating the integrand and then evaluating the resulting integral, which in many cases is similar to the form A defined in the previous section.

For example, the linear-order Taylor series term is found by computing

$$\begin{aligned}\frac{df^{SE}}{d\tau} &= \frac{d}{d\tau}\int_0^\pi d\theta\frac{\sin\theta}{2}\int_0^1\frac{du}{u^2}\left(1-\exp\left(-iu(3\cos^2\theta-1)(\tau_E-\tau)\right.\right.\\ &\qquad \left.\left.-3\lambda(\tau^3-\frac{3}{2}\tau_E^2\tau+\frac{3}{4}\tau_E^3)u^{8/3}(5\cos^4\theta-2\cos^2\theta+1)\right)\right)\\ &= \int_0^\pi d\theta\frac{\sin\theta}{2}\int_0^1\frac{du}{u^2}\left(3\lambda(3\tau^2-\frac{3}{2}\tau_E^2)u^{8/3}(5\cos^4\theta-2\cos^2\theta+1)+iu(3\cos^2\theta-1)(\tau_E-\tau)\right)\cdot\\ &\qquad \exp\left(-iu(3\cos^2\theta-1)(\tau_E-\tau)-3\lambda(\tau^3-\frac{3}{2}\tau_E^2\tau+\frac{3}{4}\tau_E^3)u^{8/3}(5\cos^4\theta-2\cos^2\theta+1)\right)\end{aligned}$$

Setting $\tau=\tau_E$ and taking only the real part of the resulting expression yields

$$\begin{aligned}\frac{df^{SE}}{d\tau} &= \int_0^\pi d\theta\frac{\sin\theta}{2}\int_0^1\frac{du}{u^2}\left(\frac{9}{2}\lambda\tau_E^2u^{8/3}(5\cos^4\theta-2\cos^2\theta+1)\right)\cdot\\ &\qquad \exp\left(-\frac{3}{4}\lambda\tau_E^3u^{8/3}(5\cos^4\theta-2\cos^2\theta+1)\right)\\ &= \frac{1}{2}\int_{-1}^1 dx\int_0^\infty dv\left(\frac{9}{2}\lambda\tau_E^{1/8}v^{2/3}(5x^4-2x^2+1)\right)\exp\left(-\frac{3}{4}\lambda v^{8/3}(5x^4-2x^2+1)\right)\end{aligned}$$

This integral is of the same form as equation (5.29), and can be solved by a similar method, or with a computer algebra system. The final result is

$$\frac{df^{SE}}{d\tau}|_{\tau=\tau_E} = 3.132\lambda^{3/8}\tau_E^{1/8} \tag{5.35}$$

Further terms can be computed in a similar manner, though these calculations become increasingly complex. For our purposes we shall stop at the second-order Taylor series,

$$f^{SE}(\tau) \approx 1.39\tau_E^{9/8}\lambda^{3/8}+3.132\lambda^{3/8}\tau_E^{1/8}(\tau-\tau_E)+(0.3915\lambda^{3/8}\tau_E^{-7/8}+0.326\lambda^{-3/8}\tau_E^{-9/8})(\tau-\tau_E)^2 \tag{5.36}$$

Figure 5.1 shows that this series expansion is a fair approximation to f^{SE} near the peak, for a number of different choices of parameters (listed in Table 5.1). In all cases where a "full" f^{SE} is shown, it is obtained by numerically integrating equation (5.26) in Matlab. The accuracy of the Taylor expansion will eventually degrade away from τ_E, but as it does the linear asymptotic form (equation (4.13)) becomes valid, meaning that the full spin-echo time course can now be modeled in some way.

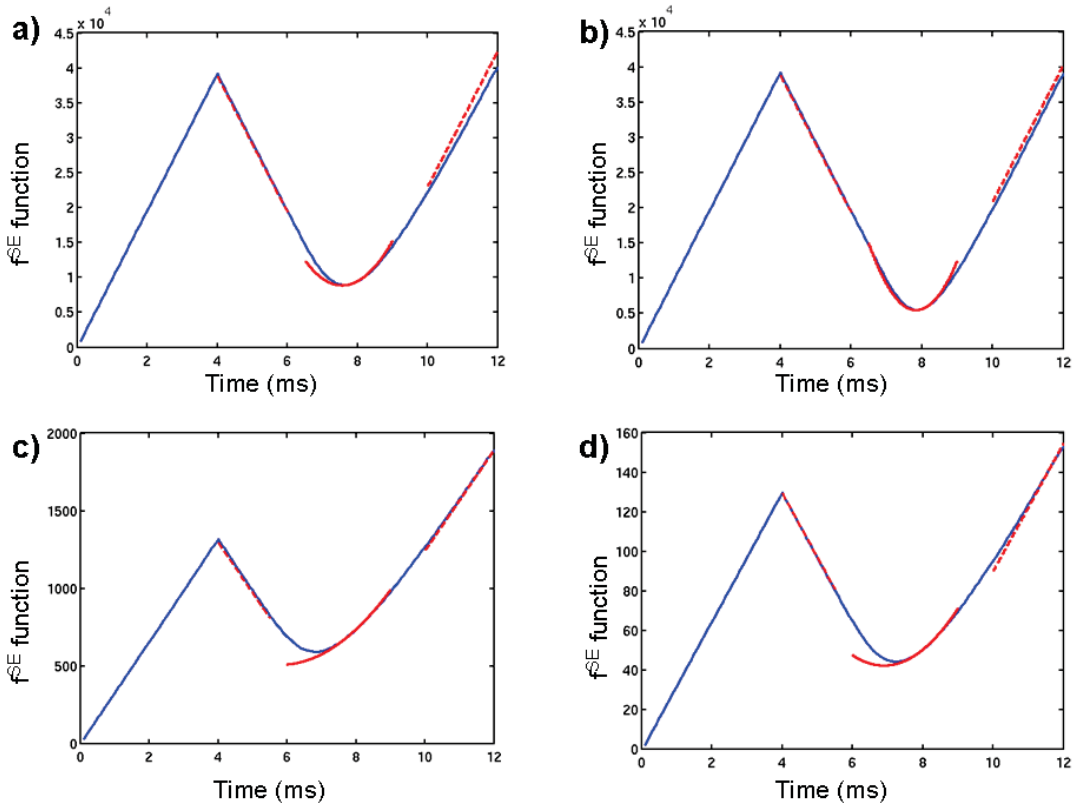

Figure 5.1: Comparison of the full (numerically integrated, blue) and approximate forms (red) for f^{SE} for four representative sets of parameters, which are listed in Table 5.1. Red dashed lines are the asymptotic form (equation (4.13)), while the red solid line is the series expansion (equation (5.36)). (a) Small, strongly magnetized particles with slow diffusion. (b) Large, strongly magnetized particles with fast diffusion. (c) Small, weakly magnetized particles with slow diffusion. (d) Large, weakly magnetized particles with fast diffusion.

In cases near the static dephasing regime (top two panels of Figure 5.1) the agreement is excellent throughout the time course. For cases with more prominent peak shifting the fit to the peak is somewhat less accurate, but far from the echo peak all time courses converge to the asymptotic form predicted by static dephasing theory.

	D (μm^2/ms)	R (μm)	$\delta\omega$ (s^{-1})	λ
Plot 5.1(a)	0.5	1	8.0 x 10^6	6.23 x 10^{-5}
Plot 5.1(b)	2.0	4	8.0 x 10^6	1.56 x 10^{-5}
Plot 5.1(c)	0.5	1	2.7 x 10^5	1.9 x 10^{-3}
Plot 5.1(d)	2.0	7	2.7 x 10^4	1.5 x 10^{-3}

Table 5.1: Model parameters used to generate plots in Figures 5.1, 5.4 and 5.5.

Modification for Peak Shifting

Based on the Taylor expansion given in equation (5.36) it is straightforward to calculate the position of the echo peak. In the absence of T_2 relaxation, this is simply the minimum of f^{SE}, which occurs at

$$\frac{df^{SE}}{d\tau} = 3.132\lambda^{3/8}\tau_E^{1/8} + 2(0.3915\lambda^{3/8}\tau_E^{-7/8} + 0.326\lambda^{-3/8}\tau_E^{-9/8})(\tau - \tau_E) = 0 \quad (5.37)$$

This gives a peak shift $\Delta\tau = \tau_E - \tau$ of

$$\Delta\tau = \frac{3.132}{0.783\tau_E^{-1} + 0.652\lambda^{-3/4}\tau_E^{-5/4}} \quad (5.38)$$

This expression has the desired property of $\Delta\tau \to 0$ at $\lambda = 0$.

To verify this expression, the actual minimum of the full, numerically integrated f^{SE} was calculated for a range of parameters using the `fminsearch` function in Matlab, and compared with equation (5.38). Figure 5.2 shows a plot of the predicted versus calculated peak shift, indicating that this expression is not suitable for most of the parameters tested. This is likely due to the contribution of higher-order Taylor series terms that were neglected.

To generate a more suitable form for predicting the peak shift, a model of the form

$$\Delta\tau = \frac{a}{b\tau_E^{-1} + c\lambda^{-3/4}\tau_E^{-5/4}} \quad (5.39)$$

was fitted to the peak shifts calculated from each combination of $\delta\omega$ and TE, and the parameters a, b and c were obtained by a non-linear least squares optimization. The values obtained from this process (expressed as mean $\pm$ standard deviation), were

$$a = 3.125 \pm 0.010$$

$$b = 10.382 \pm 0.094$$

$$c = 0.654 \pm 0.002$$

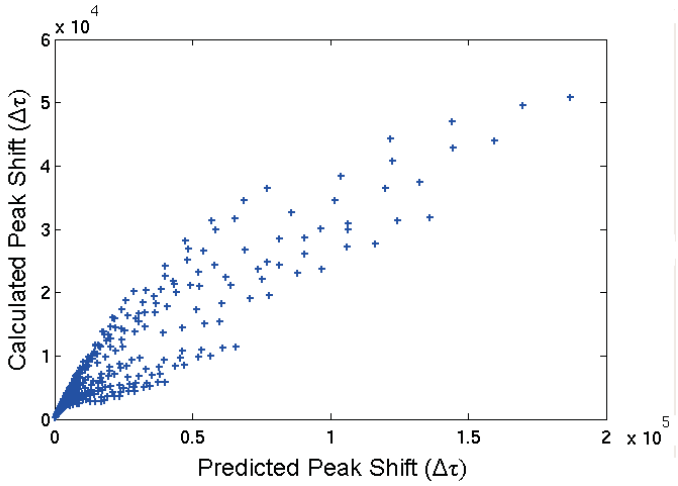

Figure 5.2: Comparison of peak shifts predicted by the Taylor series form for a particular set of parameters (x axis) and the peak shift calculated by minimizing the full f^{SE} (y axis). All calculated peak shifts are smaller than predicted, in some cases by factors of 5 or more.

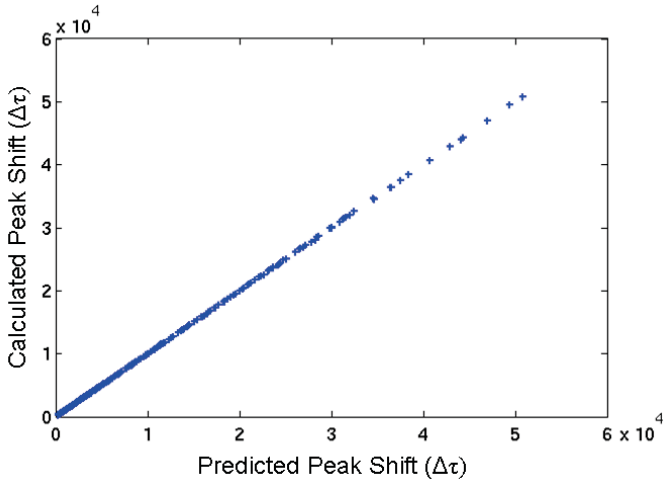

Figure 5.3: Comparison of peak shifts predicted by the empirical model (x axis) and the peak shift found by minimizing the full f^{SE} (y axis).

As seen in Figure 5.3, this fitted expression provides a far better fit to the calculated peak shifts. Interestingly, the values a and c are in excellent agreement with the prediction based on the Taylor series, but b differs by more than an order of magnitude. No errors were found in the Taylor series expansion of f^{SE} outlined in the previous section, so the reason for this discrepancy is not clear. Nevertheless, this suggests that the Taylor series expansion of f^{SE} should be

$$f^{SE}(\tau) \approx 1.39\tau_E^{9/8}\lambda^{3/8} + 3.132\lambda^{3/8}\tau_E^{1/8}(\tau-\tau_E) + (5.191\lambda^{3/8}\tau_E^{-7/8} + 0.326\lambda^{-3/8}\tau_E^{-9/8})(\tau-\tau_E)^2 \tag{5.40}$$

Re-plotting Figure 5.1 with this new approximation used around the peak also results in slightly increased performance in cases further from the static dephasing regime, as shown in Figure 5.4. We therefore propose the use of this approximation in cases where use of the full model is difficult, such as curve fitting to experimental data.

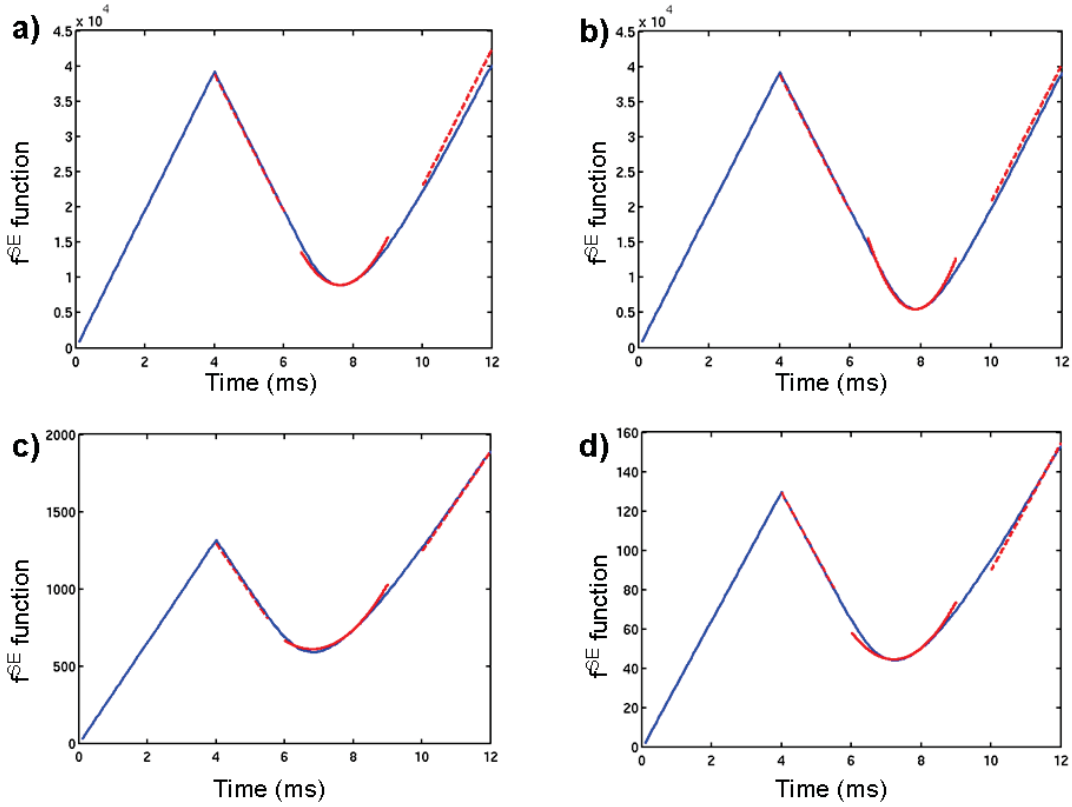

Figure 5.4: Comparison of the full (blue) and revised approximate forms (red) for f^{SE}. Parameters are the same as in Figure 5.1. Little change is observed in the near-static dephasing cases (top panels), but the modified approximation does perform slightly better in cases further from the SDR (bottom panels).

For completeness, we also note that in the case where R_2 relaxation cannot be neglected, the expression for the peak shift must be modified to include R_2 using

$$\begin{aligned} \frac{ds(t)}{dt} &= s(t)(-R_2 - \zeta\frac{df^{SE}}{dt}) = 0 \\ \frac{df^{SE}}{dt} &= -\frac{R_2}{\zeta} \\ \frac{df^{SE}}{d\tau} &= -\frac{R_2}{\zeta\delta\omega} \end{aligned} \tag{5.41}$$

which gives a peak shift of

$$\Delta\tau = \frac{3.132 + R_2\zeta^{-1}\lambda^{-3/8}\tau_E^{-1/8}}{0.783\tau_E^{-1} + 0.652\lambda^{-3/4}\tau_E^{-5/4}} \tag{5.42}$$

5.3 Validation of the Model

If the model developed in the previous section is to be of practical use, its accuracy and applicability must be tested against experimental data. As we have demonstrated, TurboSPI is able to provide such data at high temporal resolution, though in practice obtaining this data over a wide range of parameters (such as particle radius and susceptibility, volume fraction and diffusion) is challenging. MPIO particles with arbitrary properties are not available and the iron content of SPIO-loaded cells can be difficult to control. We will therefore supplement the available experimental data with simulated relaxation curves.

5.3.1 Monte Carlo Simulations

The use of computer simulations is a valuable tool for characterizing the behavior of systems that would otherwise remain inscrutable due to their complexity or the difficulty associated with observing them experimentally. Systems containing random processes such as diffusion, or which involve random distributions and orientations of particles, are not easily modeled by deterministic algorithms, but can be simulated by Monte Carlo methods, which involve repeated random sampling of an otherwise deterministic system.

A wide variety of Monte Carlo studies have been performed to examine relaxation behavior around magnetic perturbers. Some early examples are that of Hardy and

Henkelman [106], who attempted to derive an empirical description of transverse relaxation near iron oxide particles, as did Fisel [107] for the case of susceptibility in brain tissue. Kennan [108] and Muller [90] performed further simulations on cylindrical and spherical geometries respectively, which were later reproduced and augmented with experimental data by Boxerman, Weisskoff and colleagues [88, 109]. More recently, Gillis [110] explored relaxation in the case of multiple-echo sequences, and Matsumoto [111] simulated the effects of particle clustering.

All of these studies share common elements in terms of how the simulation is carried out, and most of these elements are reproduced in the current set of simulations.

Simulation Details

To simulate the MRI signal in the presence of a network of spherical perturbers, a number of impermeable spheres of radius R and susceptibility $\Delta\chi$ are distributed randomly within a cubical universe of side length L, filling it to a specified volume fraction ζ. The side length L is chosen based on the sphere radius and the mean diffusion path length [109], as well as a scaling factor based on ζ to ensure that the number of spheres generated is sufficiently large (typically at least 1000). Spheres are not allowed to overlap. At each iteration of the simulation, a single proton is placed within the center of the cube; if the network of perturbers is random, this is equivalent to placing the proton at a random position within the cube.

Diffusion is simulated by computing a random path for each particle to follow, with a step every $\Delta t = 100\mu s$. Each step is a randomly oriented vector in three dimensions, with magnitude $\sqrt{6D\Delta t}$ and direction chosen from a uniform distribution. Steps which would move the proton inside a sphere are disallowed, and periodic boundary conditions are enforced. A total of 5 x 10^4 random walks are generated for each simulation, with a new network of spheres generated every 100 iterations.

At each time step, the magnetic field at the proton's position is calculated by superposing the fields from all of the spheres inside the cube. The resulting change in the phase of the proton's magnetization is tracked throughout the random walk, with a phase reversal applied at time TE/2 to simulate a 180 degree pulse if desired. For the purposes of these simulations, both T_1 and T_2 decay are neglected. The sum of all N proton magnetizations becomes the simulated MRI signal.

These simulations were performed on a Dell PowerEdge 2900 server (8 Intel Xeon 3.16 GHz CPUs, 50 GB RAM) in Matlab, using the Distributed Computing Toolbox to simulate multiple protons in parallel. The total simulation time for each set of parameters depends on the number of iterations, particle radius and volume fraction (which affect the number of perturbers and the universe size L) but is typically between 3 and 10 minutes.

Comparison with Analytical Model

Monte Carlo simulations were run for a variety of combinations of the parameters R, D, $\delta\omega$ and ζ, and the MRI signal generated by these simulations was compared

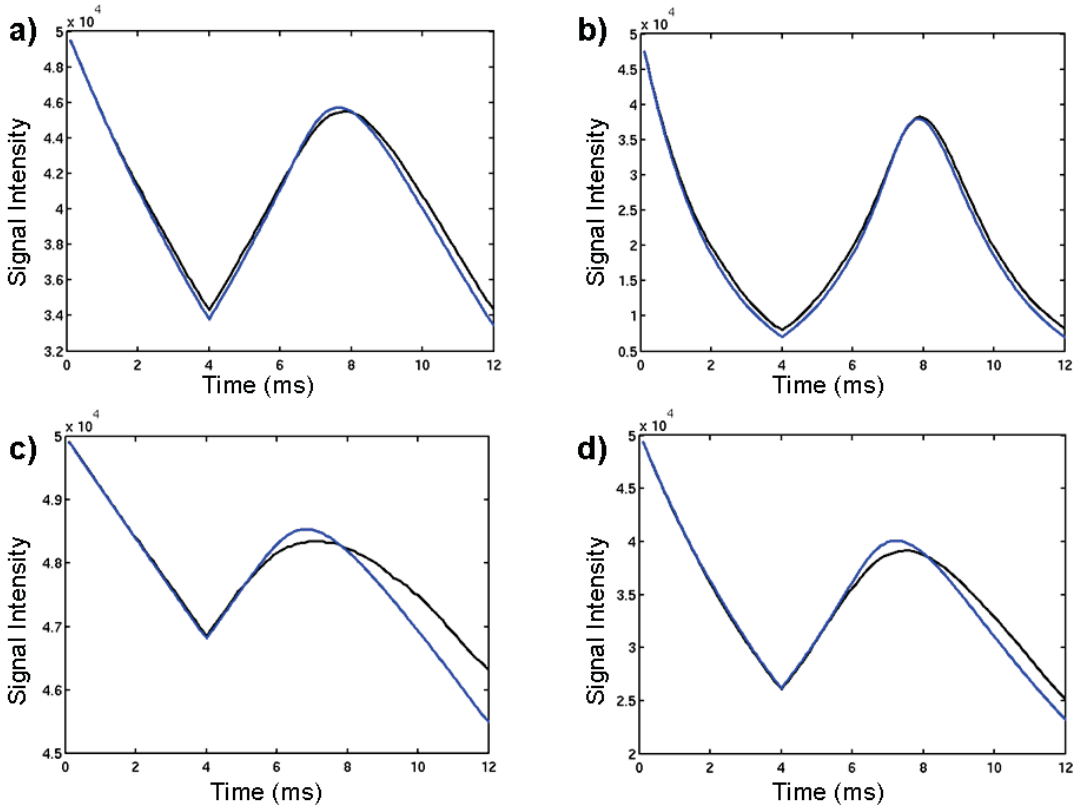

Figure 5.5: Comparison of the analytical signal ($Ne^{-\zeta f}$, blue) with the signal generated by Monte Carlo simulation (black). Parameters are the same as in Figure 5.1, with the following volume fractions used for each simulation: (a) $\zeta = 1$ x 10^{-5}, (b,c) $\zeta = 5$ x 10^{-5}, (d) $\zeta = 5$ x 10^{-3}. Agreement is excellent throughout the time course in the two cases close to the SDR (top panels). Further from the SDR (bottom panels), the curves diverge somewhat after the echo peak.

with the analytical result obtained by numerically integrating equation (5.26) with appropriate parameters and then calculating the signal $s = Ne^{-\zeta f}$, where N was the number of particles in the simulation. Representative examples of these signals are given in Figure 5.5, and show good overall agreement.

In the case of systems further from the static dephasing regime, such as those illustrated in Figure 5.5(c,d), the simulated and calculated signals diverge after the echo peak, to a degree depending primarily on the degree of diffusion in the system. It should be noted that, in either case, the calculated and simulated signals coincide at t=TE=8ms, and this is generally true for all combinations of parameters tested.

5.3.2 Experimental Data

TurboSPI time courses obtained using the methods described in the previous chapter were also compared with both the calculated signal predicted by the slow diffusion model, and the signal generated by the Monte Carlo simulations. Comparisons were made with samples from each group described in section 4.2.1: 0.96 μm MPIO, 1.63 μm MPIO, and SPIO-labeled C3 cells.

In each case, for a valid comparison with experimental data, the parameters input to the model and Monte Carlo simulations must be representative of the acquired data. Fortunately, most of these parameters are already known or can be calculated easily. For loaded cells, the cell density and radius (and therefore the volume fraction) can be obtained by cell counting techniques, which are often performed as a matter of course during the preparation of cultures. For particles such as MPIO, the mean radius and particle density are typically provided by the manufacturer, and the volume fraction can be therefore calculated based on the concentration used to prepare each sample. In both cases the susceptibility of each perturber can be computed from the sample susceptibility once the volume fraction is known.

The only variable not known beforehand was the diffusion coefficient of the samples, which must be obtained experimentally. Since all samples of a particular particle size are prepared at the same time using the same gelatin suspension, it can be assumed that all samples within a group will have the same diffusion coefficient. Since the presence of cells or MPIO should not significantly affect diffusion, a single measurement performed on a reference sample without any SPIO should suffice.

The diffusion coefficient can then be obtained with an MRI measurement based on the Stejskal-Tanner pulsed gradient experiment [112]. Such measurements were performed, and yielded an average diffusion coefficient of 2.02 ± 0.31 μm^2/ms which was used as input to all calculations.

Data were compared for several samples in each group. Once analytical forms and simulations were obtained for each set of parameters, T_2 decay was added to the resulting signals, with T_2 relaxation times for each sample measured as described in Section 4.2.2. The experimental data was scaled so that its intensity at t=TE matched the analytical signal at t=TE, and all three results were compared. The results of this comparison for representative samples are shown in Figure 5.6, with parameters used to generate the model and simulation in Table 5.2.

	D (μm^2/ms)	R (μm)	$\delta\omega$ (s^{-1})	λ	ζ	R_2 (s^{-1})
Plot 5.6(a)	2.0	0.5	8.01 x 10^6	4.98 x 10^{-4}	5 x 10^{-5}	46
Plot 5.6(b)	2.0	0.8	4.03 x 10^6	2.49 x 10^{-4}	5 x 10^{-5}	23
Plot 5.6(c)	2.0	7	2.67 x 10^5	7.6 x 10^{-5}	1 x 10^{-3}	17

Table 5.2: Parameters used to generate model and simulated signals for comparison to experimental data, as plotted in Figure 5.6.

In all cases the experimental data agree well with the Monte Carlo simulations, and follow the same trends observed in the previous section; near the static dephasing regime there is good agreement with the analytical model, while the data and the model prediction do tend to diverge slightly as the effect of diffusion in the system increases.

5.4 Applications

With the accuracy of the model validated against simulations and experimental data, we can now consider some of the ways in which this improved description of signal relaxation could be applied for practical benefit.

5.4.1 Improved Quantification

In the static dephasing regime, relaxation behavior in the presence of magnetic perturbers depends only on four parameters: the characteristic frequency of the perturber $\delta\omega$ (which is proportional to the susceptibility $\Delta\chi$), the volume fraction ζ

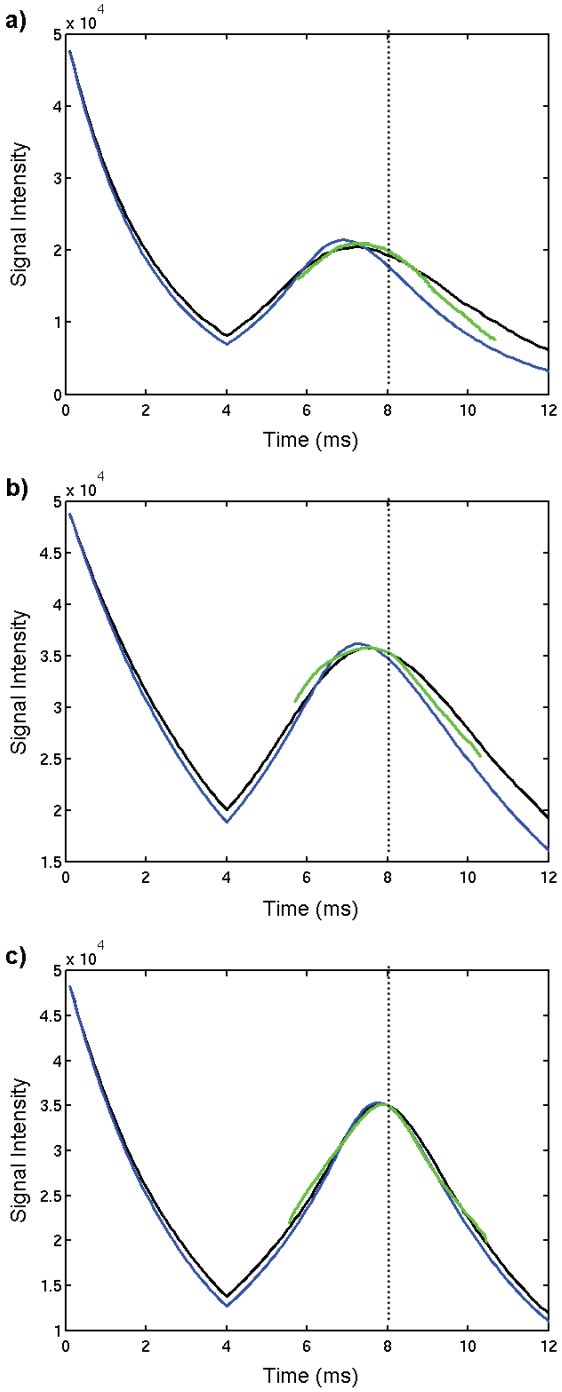

Figure 5.6: Comparison of the analytical (blue) and simulated (black) signal with experimental data (green), acquired using TurboSPI. Representative samples are shown, with parameters listed in Table 5.2. (a) Tube of 0.96 μm MPIO particles, (b) Tube of 1.63 μm MPIO particles, (c) Tube of SPIO-loaded C3 cells.

occupied by the perturbers, the background R_2 relaxation rate, and the scaling factor ρ which reflects the overall signal intensity. It is not possible to obtain independent estimates of all of these parameters with an FID experiment, since the effective relaxation rate R_2^* cannot be separated into its constituents R_2 and $R_2' \propto \zeta \cdot \delta\omega$.

Experiments which sample both the rise and fall of the spin echo (or the FID and the rising half of the spin echo) provide the opportunity to separate R_2 from R_2' by comparing the rephasing and dephasing relaxation rates [59]. Also, as demonstrated by Yablonskiy [92], the deviation from linear exponential behavior at the peak of the echo allows estimation of ζ by comparing the actual signal intensity at the echo peak to that extrapolated from the R_2^* fit. As noted in reference to equation (4.16), the presence of an e^{ζ} term in the linear-exponential region allows the volume fraction to be computed as

$$\zeta = log \frac{S_{extrapolated}(TE)}{S(TE)} \tag{5.43}$$

Because this approach assumes the signal is symmetric about the echo time TE, is it not directly applicable to the analytical model developed in this chapter. However, it is possible to fit the Taylor-series approximation (equation (5.40)) to the echo peak and the asymptotic form (equation (4.15)) to the rising and falling portions. This would enable estimation of all of the parameters included in the model, including the diffusion time R^2/D contained within the parameter λ.

There would be clear benefits to obtaining this additional information in applications involving SPIO-labeled cells. While R_2^* relaxometry can be used to quantify the total iron load of a voxel, in the case of a population of labeled cells which is dividing and distributing the SPIO among daughter cells, this quantity will remain constant (assuming no migration of cells or loss of SPIO) and is therefore not useful for monitoring the behavior of cells over time. However, accurate and separate estimates of the number of cells and the loading level of each cell, obtained through ζ and $\delta\omega$ respectively, would allow cellular propagation to be observed more quantitatively. Estimation of the background R_2 could also be incorporated to further improve quantification accuracy.

However, early attempts to perform such quantification using simulated and experimental data have been unsuccessful. While a non-linear least-squares curve fitting routine (`lsqcurvefit` in Matlab) can be used to fit the Taylor series (equation (5.40))

and the asymptotic form (equation (4.15)) to the entire time course, the resulting parameter estimates are strongly dependent on the initial conditions supplied to the fitting routine, which will not be known *a priori* in experimental applications. This reflects the difficulty in correctly separating the contributions of ζ, R^2/D and $\delta\omega$, even in the ideal case of simulated data.

In particular, the linear asymptotic form of the time course depends on $\zeta \cdot \delta\omega$ while the quadratic approximation near the peak depends on $\zeta\delta\omega^{9/8}\lambda^{3/8}$. Because of the similarity of these forms it is possible for changes in $\delta\omega$ to be mistakenly assigned to λ, such that ζ and $\delta\omega$ remain hard to separate. More sophisticated approaches than simple curve fitting may be required if reliable parameter estimation is to become practically feasible. For example, it may be possible to use Yablonskiy's method described above to provide a rough estimate of ζ which can then be refined with curve fitting.

5.4.2 Robustness to Field Inhomogeneity

To this point, we have always assumed that the external magnetic field, or at least the magnetic field experienced by a particular imaging voxel, is homogeneous. Under most experimental conditions this is not the case, due to imperfect hardware or gradients induced by anatomical structures, and the observed signal will be affected by these inhomogeneities. These can be generally regarded as macroscopic field variations, as compared with the mesoscopic fields around the perturbers, or the microscopic fields which lead to background R_2 relaxation.

Assuming that the gradient across an imaging voxel is a constant, G, such that the range of frequencies across the width of the voxel Δx is

$$\Delta\omega = \gamma G \Delta x \tag{5.44}$$

then it can be shown [92] that the signal will be modulated by a function

$$F(t) = \frac{\sin(\Delta\omega t/2)}{\Delta\omega t/2} \tag{5.45}$$

Examples of this modulation applied to Monte Carlo data are shown in Figure 5.7. These were generated by adding a field offset based on the particle's position at each step of the random walk, to simulate the presence of an applied gradient centered on

the imaging voxel. Note that the modulation shifts the echo peak back towards the nominal echo time, and that the relaxation behavior far from the peak is no longer linear-exponential.

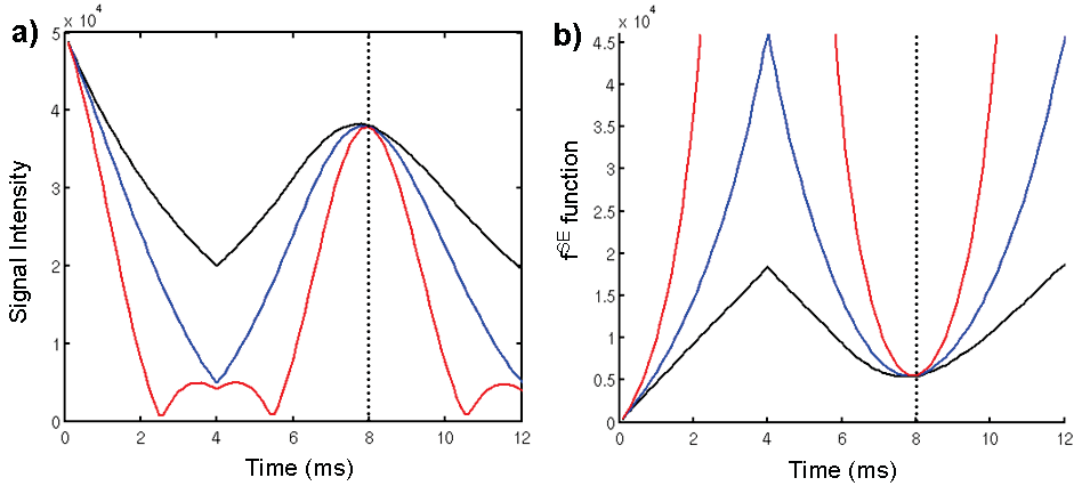

Figure 5.7: Examples of modulation introduced by macroscopic field gradients. Data are from a Monte Carlo simulation with TE = 8ms, D = 2 μm^2/ms, R = 2 μm, $\delta\omega$ = 4 x 10^6 s^{-1}, ζ = 5 x 10^{-5}. Black line: No external gradient. Blue line: External gradient of 0.5 G/cm across the imaging voxel. Red line: External gradient of 1 G/cm. (a) Computed signal intensity. (b) Equivalent f function, obtained from the logarithm of signal intensity.

Such a modulation can be added to the model used for fitting experimental data, to attempt to separate the effects of macroscopic inhomogeneities from the mesoscopic effects of interest. As long as the argument $\Delta\omega t$ is small, $F(t)$ can be approximated by a quadratic exponential,

$$F(t) \approx e^{\Delta\omega^2 t^2/24} \tag{5.46}$$

A basic example showing demodulation of Monte Carlo data using such an approximation is presented in Figure 5.8. Fits to the rising and falling portions of the spin echo are computed, and the quadratic term of the fit is used to correct for the modulation, leading to a curve that better matches the simulation performed without macroscopic inhomogeneity. Furthermore, the coefficient of the quadratic term can be used, in conjunction with equations (5.46) and (5.44), to estimate the macroscopic gradient. In this example, the computed gradient is 0.286 G/cm, which compares well with the actual simulated gradient value of 0.25 G/cm.

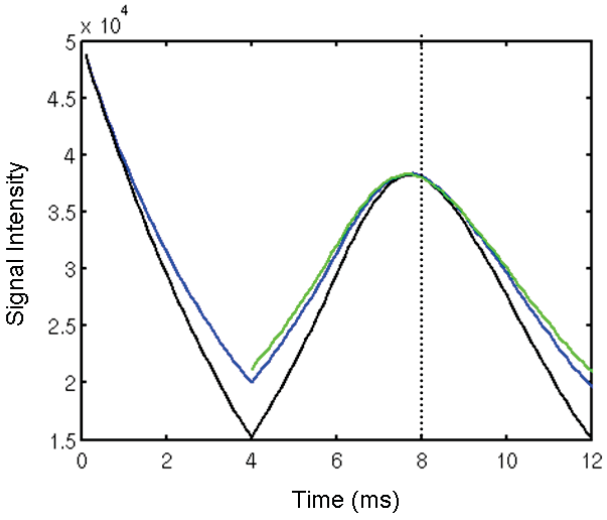

Figure 5.8: Removal of modulation by macroscopic gradients, using data from a Monte Carlo simulation with parameters matching those of Figure 5.7. Blue line: Signal from simulation with no macroscopic gradient. Black line: Signal from simulation with external gradient of 0.25 G/cm. Green line: Signal with effects of external gradient demodulated.

This extra term to account for macroscopic inhomogeneity could easily be added to the model described in this section. TurboSPI data should be particularly well-suited to fitting with such an expanded model due to the high number and density of time samples. This has yet to be tested rigorously, but should become possible once a suitable procedure is devised for reliably fitting the basic model to experimental data. One potential challenge will be the similarity of the modulation term and the Taylor expansion near the peak, which are both quadratic in t. However, as long as the acquisition window is wide enough, non-linear behavior away from the echo peak should be evident and distinguishable from the linear effects of relaxation.

5.5 Summary

In this chapter we have developed an analytical description of MRI signal behavior throughout an FID or a spin echo, in the presence of spherical magnetic perturbers and slow but non-negligible diffusion. The analytical model demonstrates excellent agreement with both Monte Carlo simulations and experimentally acquired data. This model, and its approximate forms, may allow the full TurboSPI time course to be employed for more accurate quantification of systems of SPIO-labeled cells.

Chapter 6

Acceleration of TurboSPI

To this point, the use of TurboSPI for quantification of SPIO has only been demonstrated *in vitro*, and its other applications described in the literature [63, 113] all involve the imaging of non-living systems. In order to be of practical use for the detection and quantification of labeled cells, TurboSPI acquisitions must be obtained *in vivo*. This is challenging for a variety of reasons, the most important of which is TurboSPI's relatively slow imaging speed, which must be improved significantly if 3D images are to be acquired of animals under anesthesia. This section will explain how TurboSPI can be accelerated to the extent that it can successfully image animal models *in vivo*.

6.1 Theory - MRI with Compressed Sensing

The speed at which any MR image can be acquired depends on a variety of factors, but is ultimately limited by the fact that enough data must be collected in k-space to faithfully represent the object being imaged. Increases in acquisition speed can be achieved by omitting portions of k-space, but this always comes at a cost in terms of image resolution or quality, as measured by the level of artifact or by the signal-to-noise ratio (SNR). For example, omitting the corners of k-space reduces image resolution, since some high spatial frequencies are not sampled. Sampling k-space asymmetrically and reconstructing the remainder using conjugate symmetry [79] increases the noise level of the image since there are fewer independent measurements. Instead of changing the extent of coverage in k-space, it is possible to reduce the density of sampling (i.e. increasing the step size Δk), but as discussed in Chapter 2, this leads to aliasing artifacts in which multiple copies of the image overlap.

Over the past decade a number of techniques [114, 115, 116] have been proposed for recovering un-aliased images from such undersampled datasets by using what has come to be known as parallel imaging. A full discussion of parallel imaging is beyond

the scope of this work, since the equipment available did not permit its application to TurboSPI. As a brief summary, an object is imaged using an array of RF receiver coils, each of which is sensitive over a different portion of the FOV. Instead of simply combining the resulting images, which would still contain aliasing, the different spatial sensitivities of the coils allow un-aliased images to be reconstructed, with the specific procedure varying from one imaging method to another. In theory the image can be undersampled by a factor equal to the number of coils in the array, though in practice this optimal undersampling is difficult to achieve even with a suitably constructed RF coil array. In the absence of such an array, other approaches must be considered.

Rather than uniformly undersampling k-space, we can instead consider the case of random undersampling. This will still impact the quality of the resulting image, but instead of coherent aliasing, the resulting artifacts will be incoherent and noise-like, as shown in Figure 6.1. Recent advances in signal processing [117] have demonstrated that, given certain assumptions, the original un-aliased signal can be recovered from such randomly undersampled data. This technique is known as *compressed sensing*, and its application to MRI [118] is a topic that has generated considerable interest, since it allows significant acceleration of MRI acquisitions without the need for specialized hardware, and without corresponding sacrifices in image quality.

Three conditions must be met in order to perform compressed sensing on an MRI dataset. First, the data must be sparse or *compressible*, meaning that only a small fraction of the total number of voxels have signal which differs significantly from zero. This sparsity can be in the image domain (as in MR angiography where only blood vessels appear bright [119]) or in some other transform domain (for example, the majority of medical images are made sparse by a wavelet transform or discrete cosine transform [118]). The theory of compressed sensing states that the number of samples required to properly reconstruct such data can be far less than the Nyquist criterion would normally allow, assuming the underlying data is sufficiently sparse [117].

The second requirement for compressed sensing is that k-space must be undersampled randomly, such that the resulting point spread function creates incoherent aliasing artifacts in whichever domain sparsifies the image. In practice the undersampling need not be uniformly random; variable-density undersampling has advantages as well, which will be discussed shortly. The third requirement is that the data

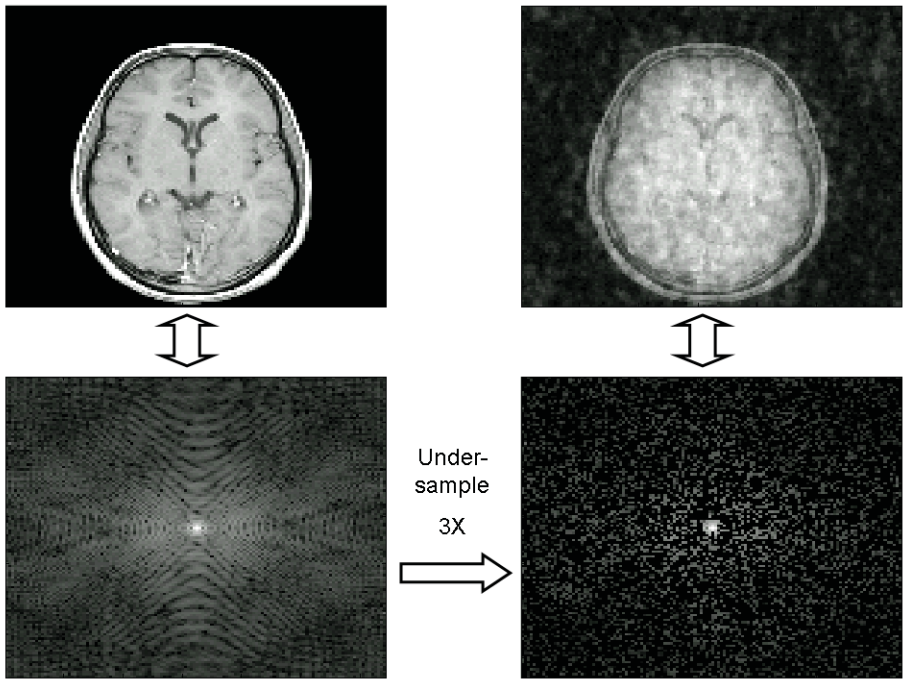

Figure 6.1: If the k-space representation of an MRI image (left) is randomly undersampled, the result is incoherent artifacts in the corresponding image (right). In this example, one-third of the total k-space is sampled.

must be reconstructed with a non-linear optimization that enforces sparsity and data consistency. The resulting recovery of aliased data is illustrated in Figure 6.2.

More rigorously, the problem we wish to solve can can be expressed mathematically in the form

$$min||\Psi x||_1 \quad \text{such that} \quad ||\mathcal{F}_u[x] - y||_2 < \epsilon \tag{6.1}$$

where x is the reconstructed image, Ψ is the sparsifying transform, y is the acquired k-space data and $\mathcal{F}_u$ is the Fourier transform operator followed by undersampling. $||\cdot||_1$ denotes the L_1 norm, the sum of the absolute values of each element, while $||\cdot||_2$ is the L_2 norm, the sum of their squares. This has been shown [117] to promote sparsity of the solution while ensuring that the result is still consistent with the acquired data, to within the noise level ϵ.

A variety of algorithms can be used to solve this problem, such as Orthogonal Matching Pursuit [120] and steepest descent [121]. Of particular interest for MRI applications is SparseMRI, a free Matlab toolbox used in [118] and elsewhere, which solves equation (6.1) with a conjugate gradient descent algorithm. Extra terms can

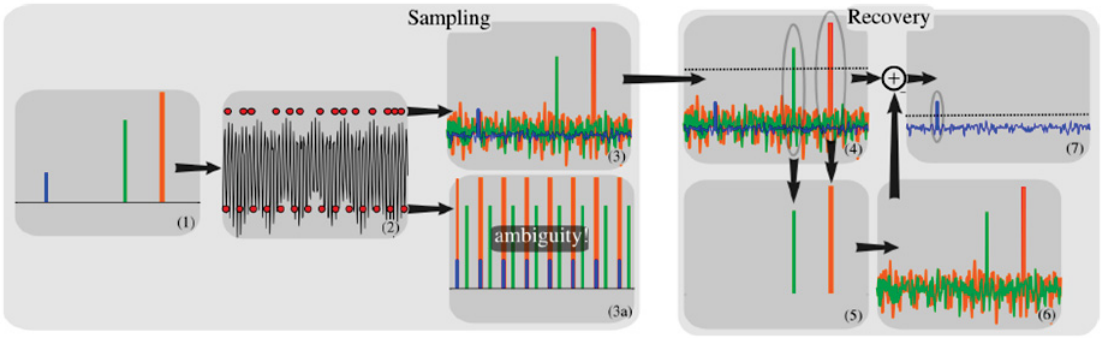

Figure 6.2: In this conceptual description of compressed sensing, a sparse signal (1) is undersampled either randomly or uniformly (2). In the case of uniform undersampling the aliased copies have intensities equal to the true signal and cannot be distinguished (3a), but random undersampling leads to incoherent artifacts (3). These artifacts obscure some of the signal, but a thresholding operation (4) recovers the dominant components (5). Knowledge of the sampling pattern allows the incoherent artifacts generated by the recovered components to be iteratively calculated (6) and subtracted from the data, allowing more components to be recovered (7). Figure from [118].

be added to equation (6.1) to provide further constraints on the reconstruction, and potentially improve the quality of the resulting image. For example, a total variation (TV) penalty can be added [118], which reduces noise by selecting images which are sparse under a finite-difference transform in addition to the sparsifying transform Ψ. The L1-SPIRiT algorithm [122] combines compressed sensing with parallel imaging and includes a constraint based on calibration data acquired by the RF coil array.

The use of prior information as an additional constraint has also been explored. If the geometry of the desired object is known beforehand, only those solutions to equation (6.1) which match that geometry need to be considered, i.e. the difference between the reconstruction and the prior information should be sparse. This is particularly useful when a series of images are acquired to study dynamic systems, as in functional MRI, since compressed sensing can be used to accelerate image acquisition and improve the temporal resolution [123].

Such prior information can be a static reference image acquired before the onset of the dynamic process [124] or a composite image made up of the data acquired at all time points [125, 126, 127]. In the first case, the prior information can also be used to prescribe an optimal undersampling pattern [128], while in the second case, no additional scans are needed if the undersampling pattern is changed for each

acquisition in the time series. Another use of prior information is to sort the image voxels in order of intensity, greatly increasing the sparsity under certain transforms [129, 130] by minimizing the differences between neighboring voxels.

6.2 Compressed Sensing with TurboSPI - Methods

In its original implementation, TurboSPI is too slow for most *in vivo* applications because of the use of pure 3D phase encoding. However, for this same reason, TurboSPI is particularly well suited to acceleration by compressed sensing. Frequency-encoded sequences can be undersampled in at most two spatial dimensions, since complete lines of k-space are acquired in the frequency-encode direction. A purely phase-encoded sequence like TurboSPI can be undersampled in all three dimensions, increasing the incoherence of the resulting aliasing artifacts and improving the quality of the resulting reconstruction [118].

It has been shown that compressed sensing can be favorably applied to ^{19}F Chemical Shift Imaging *in vivo* [121] and to the quantitative dynamic imaging of moisture content with SPI [124]. Both studies demonstrate the acquisition of information along a fourth dimension, either spectral (peaks corresponding to different ^{19}F markers) or temporal (changes in moisture content over time), which remains comparable to that acquired by conventional acquisitions. We shall show in this section that TurboSPI can be similarly accelerated with compressed sensing, and that TurboSPI presents a unique opportunity to acquire prior information, increasing the maximum acceleration factor that can be achieved.

6.2.1 Reconstruction with Prior Information

Unlike many other kinds of time-resolved acquisitions, TurboSPI does not acquire individual time points with separate acquisitions, but obtains all of the temporal information for a single location at once. As a result, it is not possible to generate prior information retrospectively by combining a series of differently undersampled time points, as reported in [125] and elsewhere. We must instead use a separate acquisition to provide prior information, similar to that described in [124]. In that study the prior information was generated using a scan of the same type as the subsequent accelerated scans, which was permissible because the dynamic process

being observed (water ingress into a material) had not been initiated, and faster imaging was unnecessary. For TurboSPI this approach is also unfeasible, and a slightly different approach must be taken to obtain a suitable image.

TurboSPI acquires temporal information at each pixel because of the lack of readout gradient. If a readout gradient is introduced to the sequence but all other parameters are unchanged, as shown in Figure 6.3, the sequence becomes identical to the typical Fast Spin-Echo (FSE) sequence used for rapid T_2-weighted imaging. A single image is produced whose contrast is determined by the signal at t=TE (see Figure 6.4), but which is obtained in a fraction of the time needed to acquire the full TurboSPI dataset. For example, to acquire a complete 128x128x16 TurboSPI dataset with ETL = 8 and TR = 250 ms requires 136 minutes; to acquire an FSE image with the same parameters takes just over 1 minute. The refocusing pulse phase must also be recalibrated to maximize image quality, since the presence of the readout gradient will affect the optimal value, but this requires little extra time. This provides a means of quickly obtaining information about the sample geometry, which can be used both to prescribe an undersampling pattern and to constrain the reconstruction.

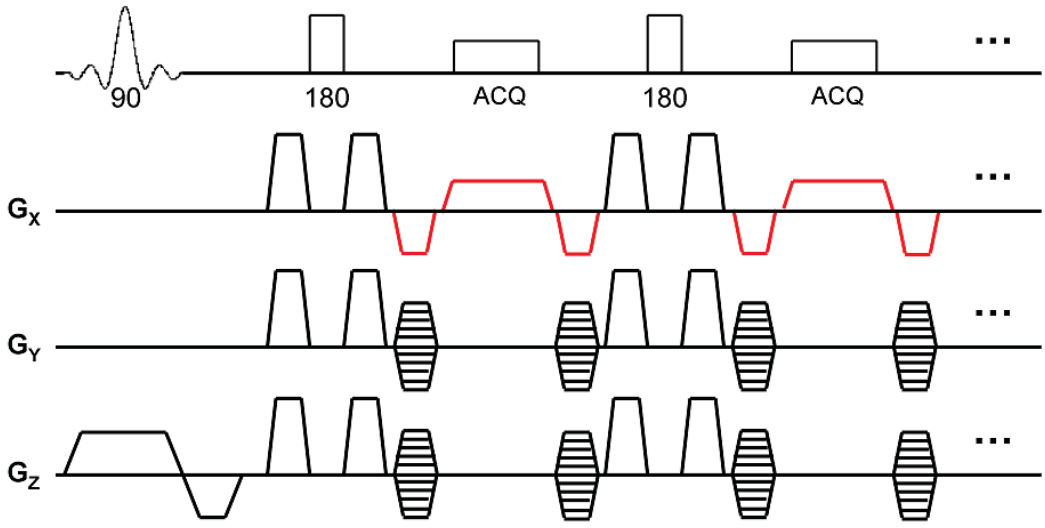

Figure 6.3: TurboSPI pulse sequence with a readout gradient added (red), effectively converting it into a Fast Spin-Echo sequence with all other parameters identical.

Clearly a single FSE image cannot provide an accurate representation of the entire TurboSPI image series; as we have seen, the contrast changes significantly throughout the rise and fall of the spin echo. However, the FSE image will match the TurboSPI

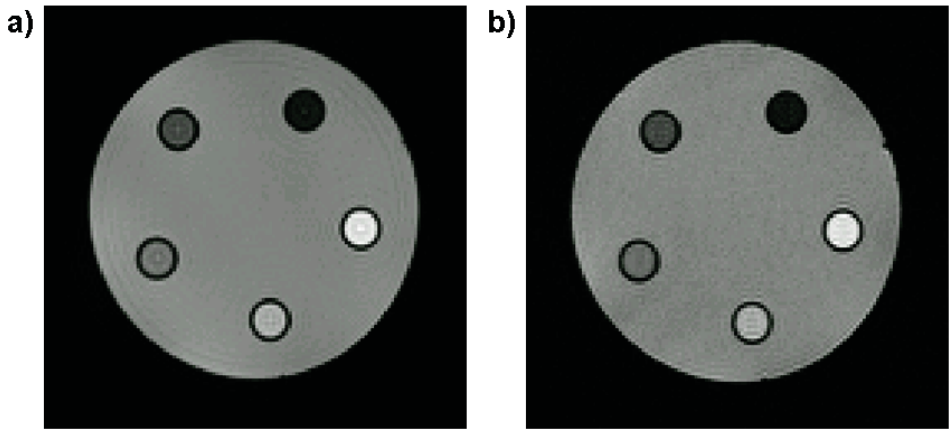

Figure 6.4: Comparison of (a) a TurboSPI image from t=TE and (b) the FSE image with matched parameters. The sample is a group of five MPIO samples in the cylindrical phantom (see Chapter 3).

image contrast at one critical point in the time course - the echo time, t=TE. Since the spacing between time points is so small (typically tens of microseconds), each image should be similar to those immediately adjacent to it. Therefore, we can use the FSE image to provide prior information for the reconstruction of the TurboSPI image at t=TE, then use that image to assist in the reconstruction of the adjacent images. Each image reconstructed can be used to constrain or guide the one next to it, until the entire time course is processed. This is illustrated in Figure 6.5.

There are many ways to make use of the available prior information (or guide image) for a particular TurboSPI image, regardless of whether it is the FSE image or a previously reconstructed time point. The simplest is to provide the guide image as an initial condition for an algorithm such as SparseMRI, so the optimization process will converge more rapidly to a solution that resembles the guide image but is still wavelet-sparse and consistent with the sampled k-space. Alternatively, the guide image can be used to sparsify the TurboSPI data, assuming that the differences between the guide and desired image are sparse. These differences become the target of the reconstruction and, once computed, can be added to the guide image to produce the desired output. Finally, Vaswani and Liu [66] have recently shown that knowledge of the support (or most significant components) of the image to be reconstructed can further improve reconstruction quality, by ensuring that voxels within this support are not discarded during attempts to find a sparse solution.

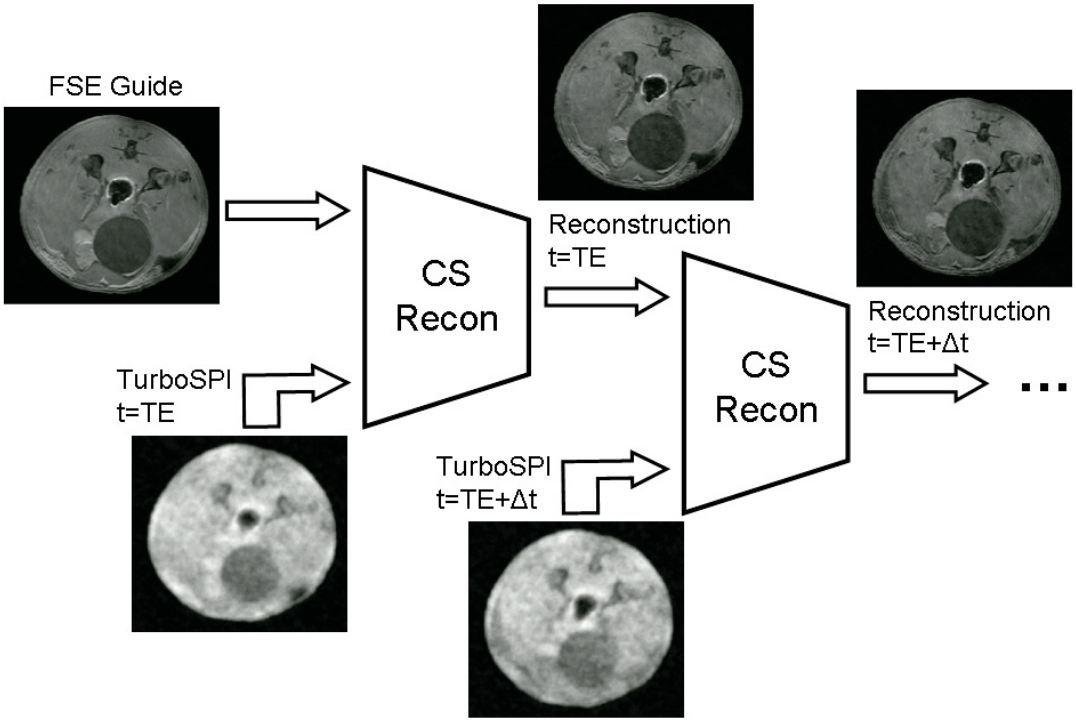

Figure 6.5: Compressed sensing reconstruction of a TurboSPI time series with guide images. The reconstruction at each time point is guided by either the matched FSE image (for the time point at t=TE) or a previously reconstructed adjacent image (for all other time points). For details of the image reconstruction see Figure 6.6.

In the current implementation of 3D CS-accelerated TurboSPI, a combination of these last two approaches is used. This process, illustrated schematically in Figure 6.6, is performed for each point of the time course as indicated in Figure 6.5.

Given a guide image and undersampled k-space, the guide image is first intensity-scaled, such that its overall intensity matches that of the undersampled data. The guide is then phase corrected, using the central points of k-space to obtain a broad estimate of the target phase; it is suggested in [118] that this increases image sparsity and improves reconstruction accuracy. The corrected guide is Fourier transformed and artificially undersampled to match the acquired data, and this k-space is subtracted from the acquired k-space data. The k-space resulting from this subtraction corresponds to the residual, or the differences between the guide and target image, which is assumed to be sparse.

The guide image is also sparsified with a wavelet transform, and thresholded to retain only the highest-intensity components (typically the top 5% sorted by intensity). The matrix T denoting the indices of these components is passed to the non-linear reconstruction algorithm along with the k-space of the residual. During the reconstruction, the sparse components corresponding to T are held fixed. If, in the course of attempting to find a sparser solution to equation (6.1), the algorithm attempts to alter these components, those changes are disallowed unless they are required to maintain consistency with the supplied k-space data. The result of the reconstruction is a sparsified residual, which can be converted back to image space and added to the corrected guide data to yield the desired image.

The optimization algorithm used in this implementation is the nonlinear conjugate gradient descent algorithm provided by the Sparse MRI toolbox [118], modified to account for the prior knowledge T [66]. The minimization problem becomes

$$min\left(||\mathcal{F}_u[x] - r||_2 + \alpha||(\Psi x)_T||_1 + \beta TV(x)\right) \tag{6.2}$$

The first term enforces consistency with the residual, the second enforces wavelet sparsity with support T, and the third enforces total variation (TV) sparsity. The weights α and β control the importance of the sparsity penalties relative to data consistency, and can be specified independently to determine the impact of these constraints on the final image. Typical, empirically determined values for these weights are $\alpha = \beta = 0.01$.

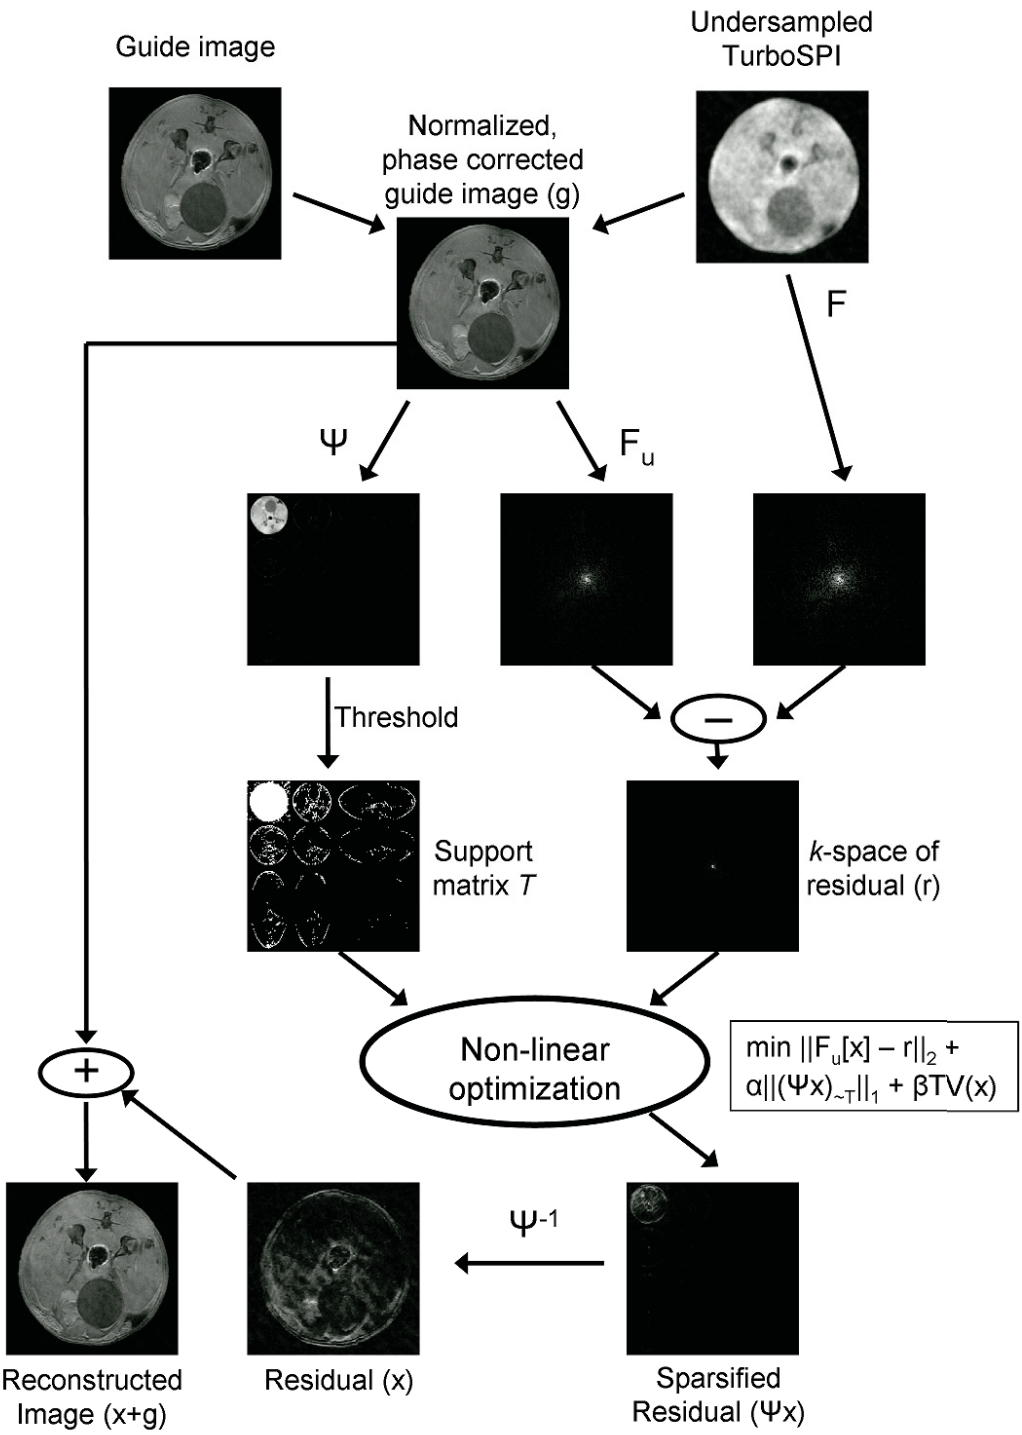

Figure 6.6: Reconstruction algorithm for an individual point in the TurboSPI time series. Given the undersampled k-space data and a guide image, the output of the non-linear optimization is the residual or difference between the guide and the acquired data, which is added back to the guide to yield the desired image.

6.2.2 Phase Encode Ordering and Undersampling

A key component of compressed sensing is a randomly undersampled acquisition of k-space, which is straightforward to perform with a purely phase-encoded sequence such as TurboSPI, requiring only some modifications to the order in which k-space locations are encoded.

The implementation of TurboSPI discussed in Chapter 3 used a linear phase encode ordering, but this assumes a full and uniform sampling of a rectilinear k-space. Random sampling patterns or those which sample the center more densely than the periphery are not as amenable to such partitioning because each k-space line will contain different numbers of points [131]. A centric ordering, such as that shown in Figure 6.7, is a better choice in general. Points acquired at earlier echoes are always near the center of k-space in both directions, with later echoes collecting points radially outward; this changes the resulting artifact patterns so they create circular ghosts instead of linear ones.

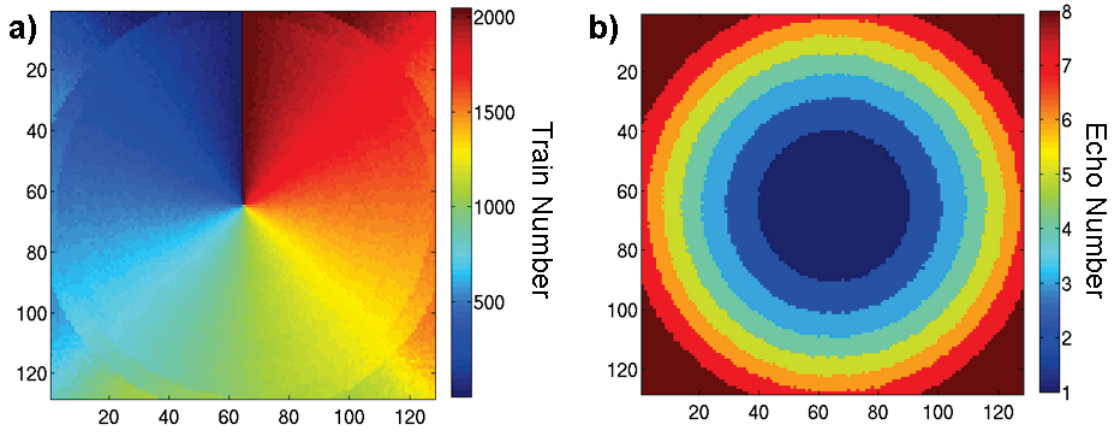

Figure 6.7: Centric phase encode ordering for a 2D TurboSPI acquisition. As in the linear case, there are 2048 trains of eight echoes. (a) Train number of each sampled point; trains are acquired beginning along $k_x = 0$ and proceeding counterclockwise around the $k_x - k_y$ plane. (b) Echo number of each sampled point; the central region of k-space is sampled early in the train, with the corners sampled at the end.

The algorithm to sort k-space locations into centric ordering is straightforward. The distance of each location to be sampled relative to the k-space center is calculated, along with the azimuthal angle relative to the k_y axis. Sampled locations are sorted into groups based on distance, with the number of groups equal to the number of

echoes in the TurboSPI echo train. If the total number of samples N is not evenly divisible by the train length, highest-distance samples are discarded as needed. Each group is then assigned to trains by azimuthal angle, with samples closer to the center of k-space acquired earlier in the train.

The algorithm is easily extended to three dimensions; before sorting by distance the samples are sorted by zenith angle into a number of conical shells, similar to those used in Conical-SPRITE [132]. The number of shells is chosen to be the largest integer divisor of N which is smaller than $\sqrt{N}$; if N is prime, samples are again discarded to make N more suitable. Each shell is then assigned to trains by radius and azimuthal angle as in the 2D case. This leads to a sampling pattern such as that shown in Figure 6.8, in which the acquisition starts along the $+k_z$ axis and traces a roughly helical path through the k-space cube, finishing along $-k_z$.

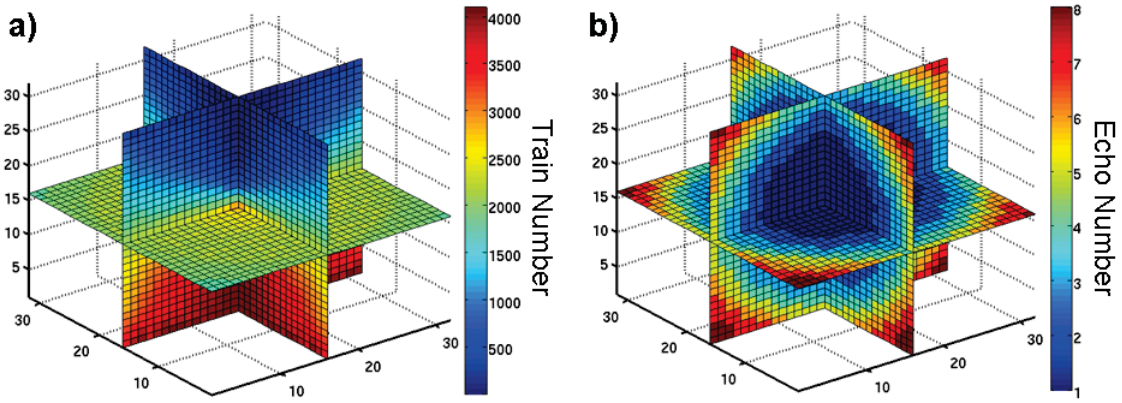

Figure 6.8: Centric ordering can also be extended to three dimensions, by assigning the sampled k-space points to conical shells. In this case a 32x32x32 matrix is fully sampled by 4096 trains of eight echoes, divided into 64 shells of 64 trains each. This small matrix size is used for illustrative purposes only; in practice, matrix sizes of 128x128x16 or 192x192x32 are used. (a) Train number during which points are sampled. (b) Echo number during which points are sampled.

As demonstrated in Figure 6.9, centric ordering is also well-suited to arbitrary undersampling factors. To generate an undersampling pattern, a subset of the full k-space is selected based on a probability density function (PDF). Though truly random undersampling (i.e. a uniform PDF) would give the most random aliasing, variable-density undersampling has the advantage of ensuring the central regions of k-space are fully sampled, and the resulting artifacts are generally still sufficiently incoherent to

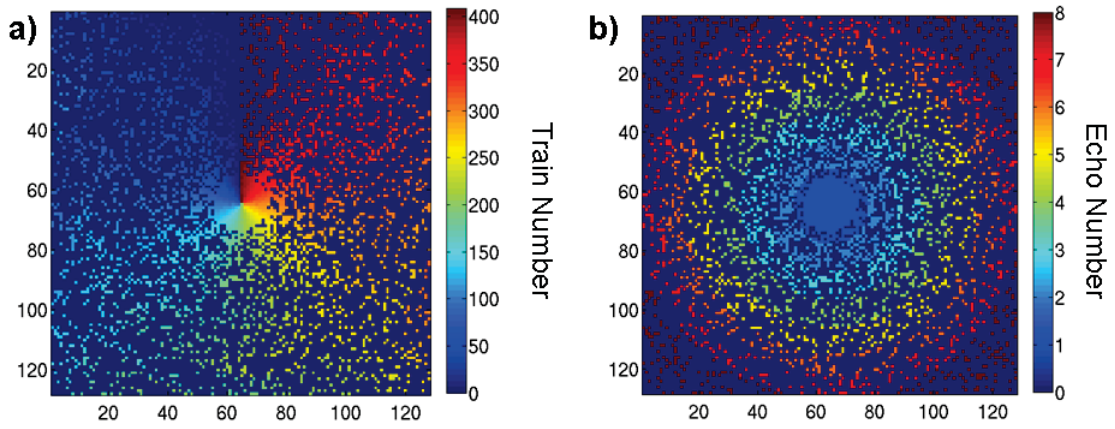

Figure 6.9: Centric order also lends itself naturally to undersampling of k-space. In this example a 128x128 matrix is undersampled by a factor of 5; only 410 trains are acquired, sampling 3280 points of a possible 16384. Points are chosen using a variable-density probability function such that the center of k-space is more densely sampled than the periphery. (a) Train number during which points are sampled. (b) Echo number during which points are sampled.

allow successful reconstruction. Indeed, variable-density undersampling outperforms purely random undersampling at high acceleration factors, even with a simple PDF that decreases as a polynomial function of distance from the origin [118].

If prior information is known concerning the object geometry, this can be used to determine a sampling pattern as well [128, 133]. Indeed, this is the method we

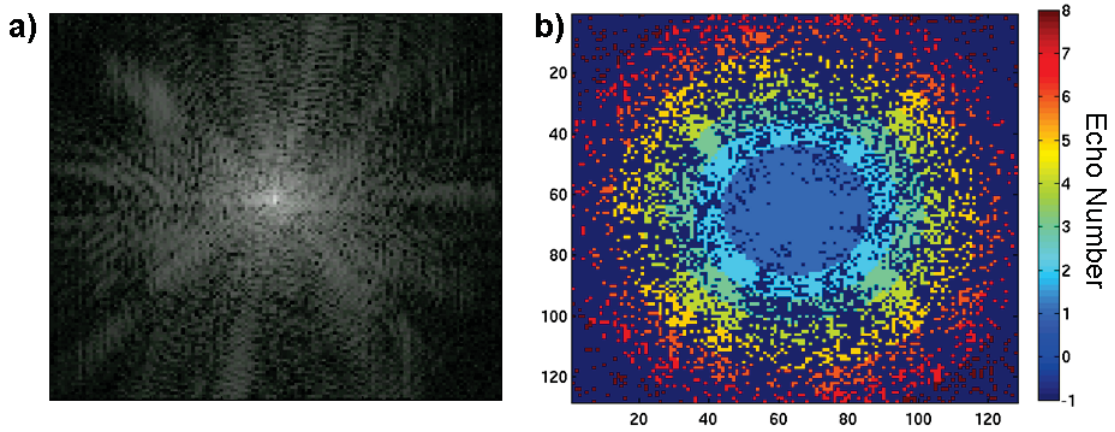

Figure 6.10: Undersampling patterns can be generated based on the FSE template image. (a) k-space of the FSE image, scaled to emphasize features. (b) Resulting phase encode pattern for 8X acceleration. Note the increased sampling density in areas where the k-space is more intense.

propose to use in conjunction with TurboSPI. The FSE image described above can be used to prescribe an undersampling pattern which will ensure that sampling will be denser in regions of k-space which are more important to the definition of the image, while remaining sufficiently random to prevent coherent artifacts. This produces a phase encoding pattern unique to the object being imaged, a step which requires little additional time to compute but which may further improve reconstruction accuracy. An example of such a pattern is given in Figure 6.10.

6.2.3 Assessment of Reconstruction Quality

To assess the quality of TurboSPI images accelerated and reconstructed with compressed sensing, it is first necessary to acquire a fully-sampled image to serve as a gold standard. Due to the slow scanning speed of the sequence this cannot be done *in vivo*, but in order to simulate many of the conditions that would be present during typical applications, an animal model is necessary.

Therefore, a single male Long-Evans rat was imaged *ex vivo* using TurboSPI. The rat was euthanized by cardiac puncture and injection of 1 mL urethane, then securely wrapped in a plastic sheet and placed within the 50 mm inner diameter home-built quadrature RF coil. Since any motion of the rat due to settling or decomposition during the fully-sampled TurboSPI acquisition would reduce image quality, the rat was left in the coil for several hours such that it would naturally settle into a stable position. The temperature of the water-cooled gradient coil was also reduced to $7°C$ during this time and for all acquisitions, to slow decomposition as much as possible.

Images were acquired using a 192x192x32 matrix and a 55x55x25 mm FOV, providing a nominal resolution of 280x280x780 μm; this is not TurboSPI's maximum resolution, but the largest image that could be acquired in the time available. A 20 mm slab was excited, and the echo time was TE=8 ms. Using an eight-echo train and a repetition time of 250 ms, acquisition of the fully sampled dataset required 10 hours and 15 minutes. Acquisition of the matched FSE image required 12 minutes with 4 averages to improve SNR, and this dataset was used to prescribe undersampling patterns for subsequent TurboSPI images. In total, eight datasets were acquired using acceleration factors ranging from 3 to 30, parameters for which are summarized in Table 6.1. Reconstruction of undersampled data required 2 hours per dataset.

Undersampling	Phase Encodes	Echo Trains	Acquisition Time
None	1179648	147456	10 hr 14 min
3X	393200	49150	3 hr 25 min
6X	196688	24586	1 hr 42 min
10X	117936	14742	61 min
15X	78624	9828	41 min
20X	59008	7376	31 min
25X	47200	5900	25 min
30X	39336	4917	21 min

Table 6.1: Undersampling parameters and corresponding acquisition times for the *ex vivo* rat images used to assess the quality of compressed sensing reconstruction. All other imaging parameters were identical.

To evaluate reconstruction accuracy, the normalized root-mean-square error was computed at each point in the time series. The nRMSE of a reconstructed image is the square root of the mean squared difference between the fully-sampled and reconstructed images, divided by the maximum intensity of the fully-sampled image.

6.2.4 Assessment of Quantification Accuracy

Since our desired application involves quantification of SPIO using the time course information provided by TurboSPI, it is critical to determine whether this information is disrupted by reconstruction with compressed sensing, and if so, to what degree. This was done by acquiring fully sampled 3D TurboSPI datasets, using the samples and cylindrical holder described in Chapter 3, followed by successive undersampling by factors of 3, 5, 8 and 10. In total, 15 different tubes of 1.63 μm MPIO were imaged in three groups of five, and R_2^* was measured for each tube and at each undersampling factor using the methods presented in Chapter 3. Fitted R_2^* values from individual voxels were averaged across a 25-voxel region of interest, such that a standard deviation could also be computed.

6.3 Compressed Sensing with TurboSPI - Results

6.3.1 Quality of Reconstructed Images

Images from a representative slice of the high-resolution *ex vivo* rat dataset at t = TE = 8 ms are shown in Figure 6.11. Even at undersampling factors up to 30, image quality remains very good overall, and agrees well with both the fully-sampled

data and the matched FSE image as expected. Some degradation of SNR and altered contrast is observed, the latter of which will be discussed shortly.

A more quantitative measure of the reconstruction error is the normalized root-mean-square error with respect to the fully-sampled data. Figure 6.12(a) plots the nRMSE as a function of time for each undersampling factor; Figure 6.12(b) plots the average nRMSE throughout the entire time series as a function of undersampling factor. The difference between images is minimized near the nominal echo time where the FSE guide and the reconstructed data match more closely, and increases as the reconstruction proceeds away from t=TE.

The differences between images tend to be maximized around t=9.5 ms, corresponding to the time when areas containing fat have reached minimum intensity due to off-resonance signal cancelation. Images at this time point and from each undersampling factor are compared in Figure 6.13, along with the fully-sampled TurboSPI image from the same time point, and the FSE image matched to t=TE. The overall quality of the images degrades as the undersampling factor increases, though many areas (such as the spinal cord, in the top middle) retain high-resolution features throughout the time course. Reconstructed images tend to differ most from the fully-sampled dataset in regions containing fat, whose contrast changes the most from t=TE.

Another feature which emerges with increasing undersampling factor is the presence of structures which are not visible in the fully-sampled image, but which are in the guide image, such as the structure highlighted by the red box in Figure 6.13. Such differences are also visible even at the t=TE time point. This behavior is not unexpected; as undersampling increases, the reconstruction algorithm will rely increasingly on the guide image to define the overall image contrast, with the result being that the reconstructed image begins to more closely resemble the guide image. This underscores the importance of ensuring that the FSE guide image matches the TurboSPI dataset as closely as possible, as well as the need to select an appropriate undersampling factor for a given experimental setup.

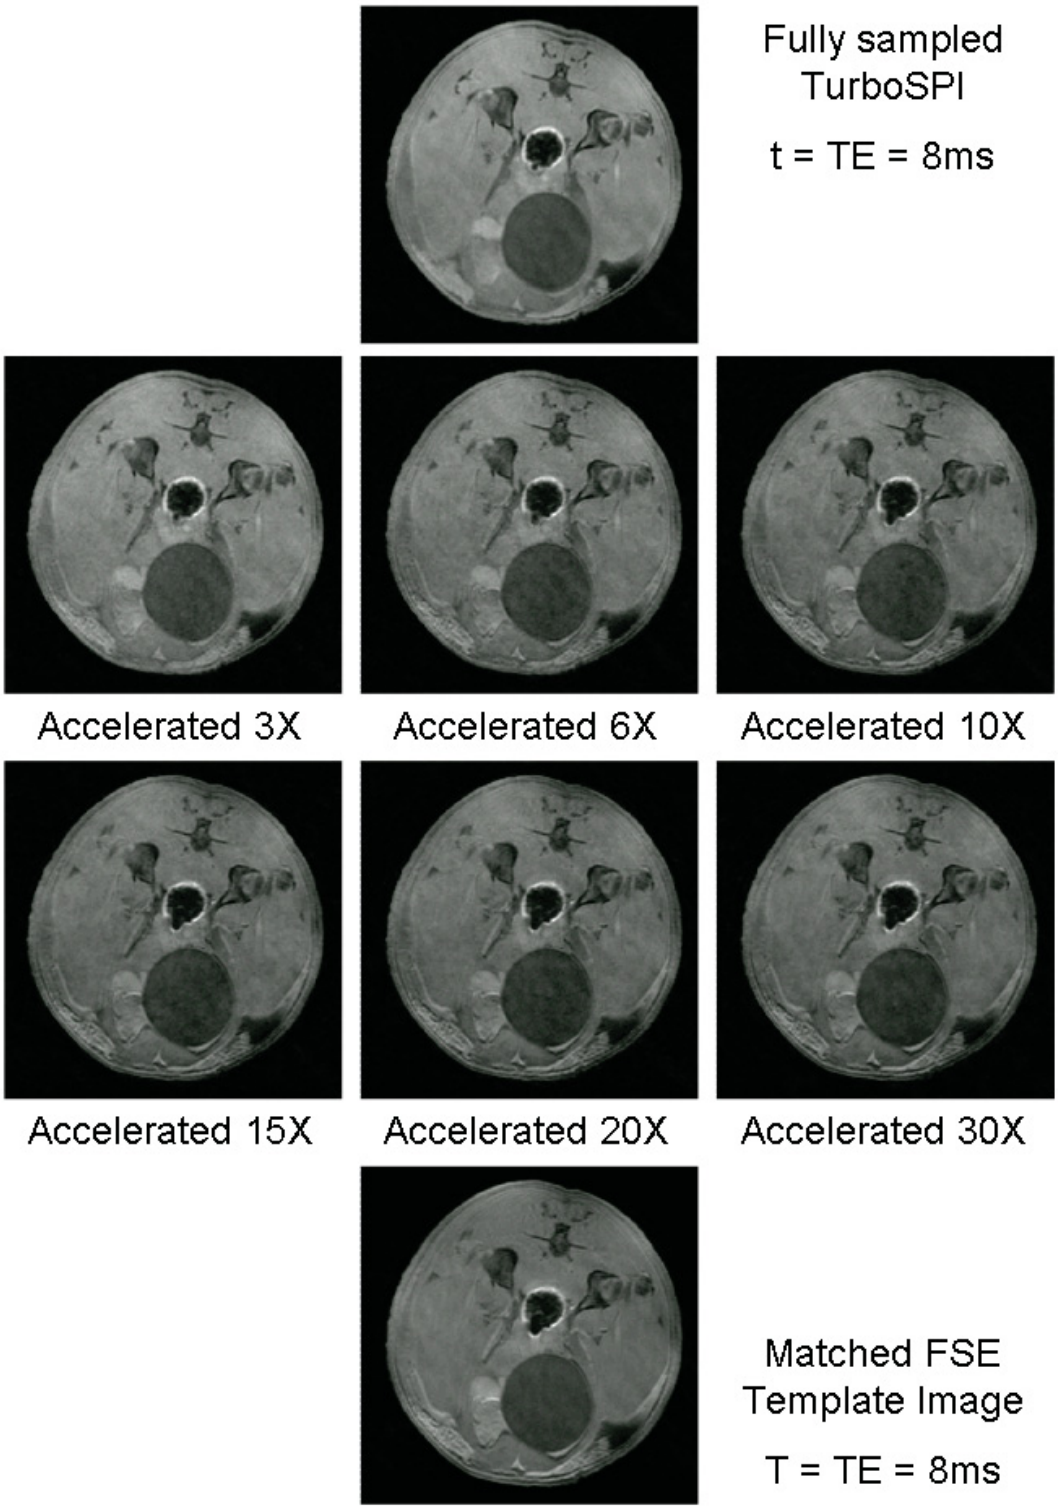

Figure 6.11: Reconstructed images from a representative slice of the high-resolution *ex vivo* rat dataset at t = TE = 8 ms and acquired with various undersampling factors are compared to fully-sampled data at t=TE and to the matched FSE image. Overall image quality remains good even at 30X undersampling.

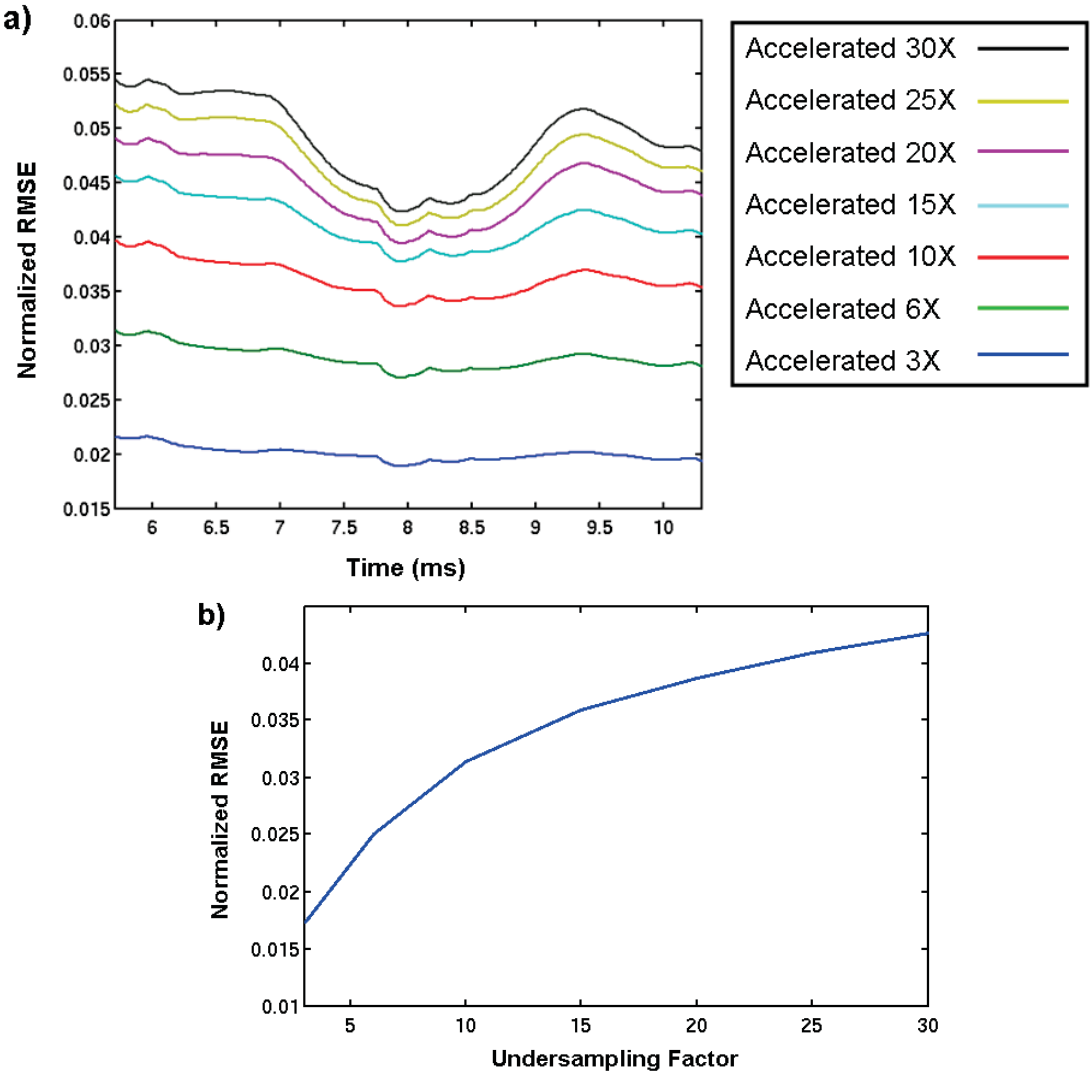

Figure 6.12: Normalized root-mean-square error of reconstructed TurboSPI datasets with respect to fully-sampled data. (a) nRMSE at each point in the time series, for undersampling factors from 3 to 30. Differences are minimal near TE = 8 ms and larger away from the echo peak. (b) Mean nRMSE throughout the time series as a function of undersampling factor.

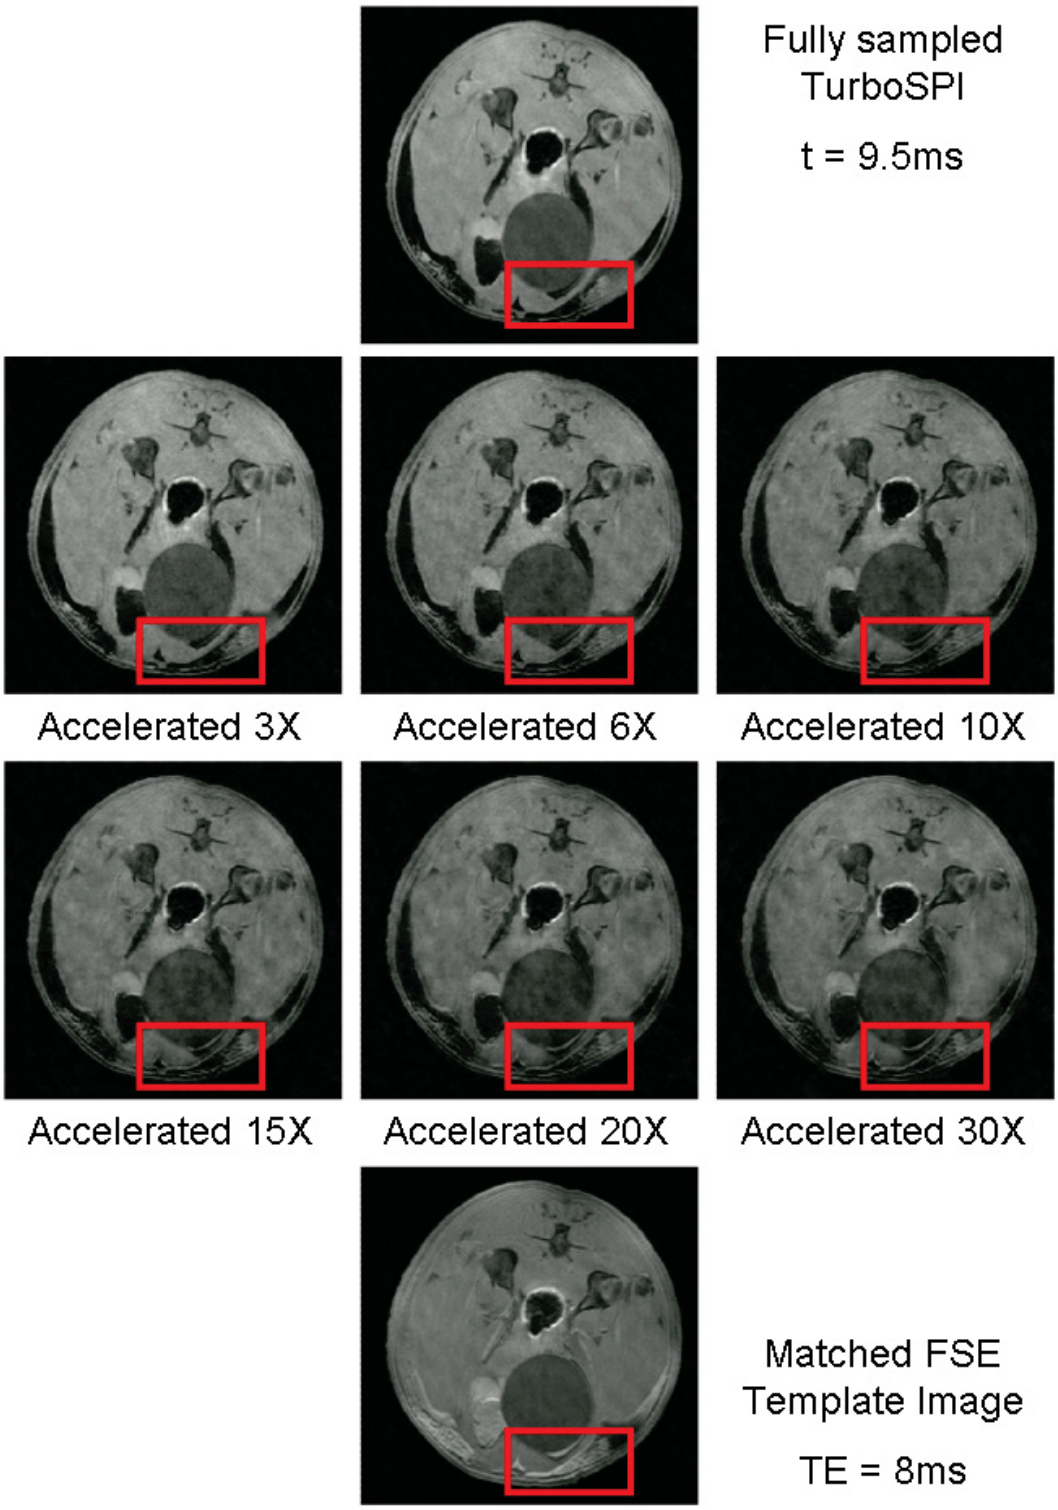

Figure 6.13: Comparison of reconstructed images at various undersampling factors with fully-sampled data at t = 9.5 ms and with FSE guide image matched to t = TE = 8 ms. Red boxes indicate a region where, as undersampling increases, contrast begins to more closely resemble the guide image than the TurboSPI data.

6.3.2 Impact on Quantification

Representative time courses from an individual voxel are shown in Figure 6.14, demonstrating that CS acceleration does not greatly alter the overall shape of the time course, including the decaying slope used to calculate R_2^*. An overall baseline offset is sometimes added, but this does not significantly affect the measured relaxation rate.

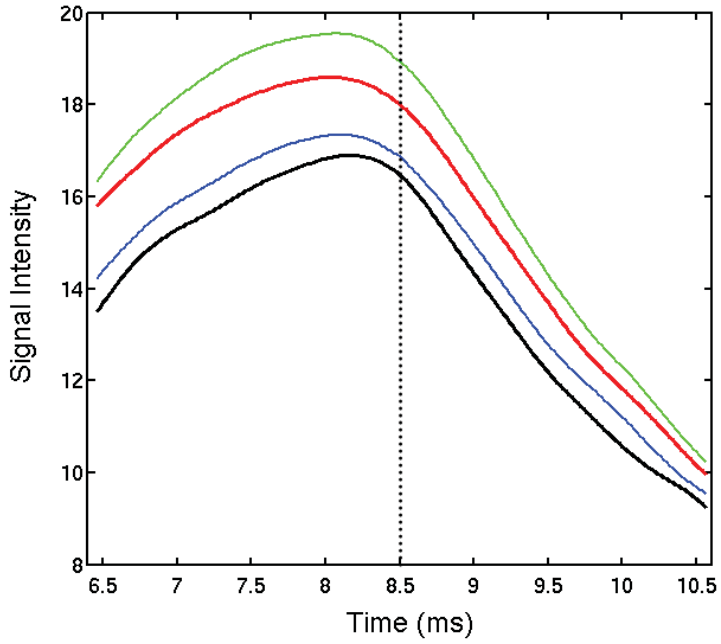

Figure 6.14: TurboSPI time courses from an individual voxel of a dataset acquired at various undersampling factors: unaccelerated (black), 5X (blue), 8X (green) and 10X (red). The voxel is in the center of a tube containing 1.63 μm MPIO and has $R_2^* \approx 300$ s^{-1}. The echo time TE=8.5 ms is indicated with the dotted line.

Relaxation rates measured from all tubes and undersampling factors are plotted in Figure 6.15. The measured R_2^* values are unaffected over a wide range of particle concentrations, with most observed differences within the standard deviation of the standard TurboSPI measurements, even at the upper range of reliable quantification. (This upper limit is somewhat smaller than that of Chapter 4 due to an increased echo time of TE=8.5 ms.)

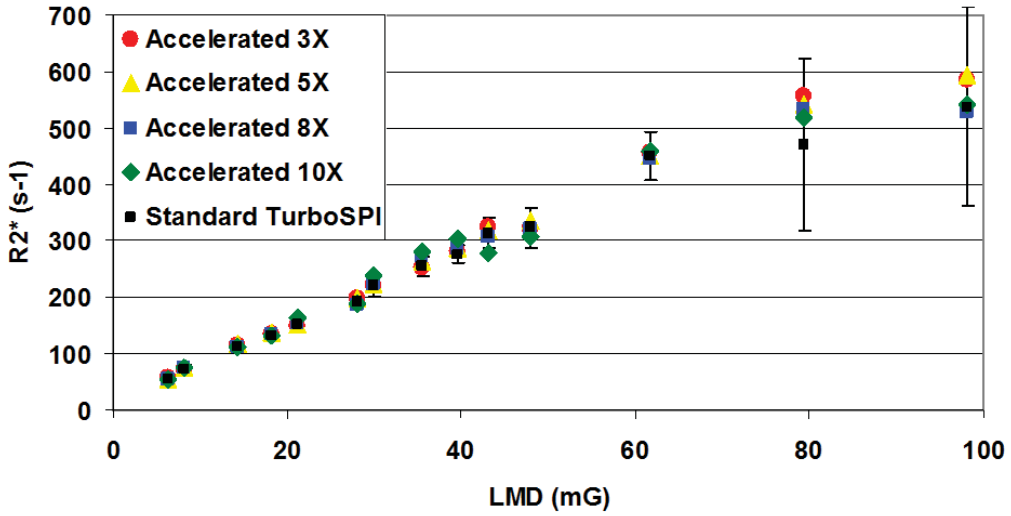

Figure 6.15: Compressed sensing acceleration (up to a factor of 10) does not significantly affect the relaxation rates measured by TurboSPI, throughout the range of 1.63 μm diameter MPIO samples tested. Error bars representing the standard deviation of the measured R_2^* with a 25-pixel ROI are only plotted for the standard (unaccelerated) TurboSPI but are comparable for all other undersampling factors.

6.4 Discussion and Summary

As a purely phase-encoded sequence with the ability to quickly acquire prior information to constrain the reconstruction, TurboSPI is especially well suited to reconstruction with compressed sensing. The reconstruction algorithm presented here allows significant acceleration of TurboSPI images without the need for special hardware and without significant impact on either image quality or the time courses used for relaxation. This is a critical requirement for implementation of TurboSPI for *in vivo* applications, which would otherwise not be possible in the time available.

As is visible in both Figures 6.11 and 6.13, as the undersampling factor increases and less TurboSPI data is available, the reconstruction must rely on the guide image to provide accurate image geometry and contrast at t=TE. In cases where the FSE and fully-sampled TurboSPI images differ, such as the structure highlighted in Figure 6.13, the contrast of the FSE image will dominate at high undersampling factors. These effects can be minimized by ensuring that the FSE image differs from the TurboSPI as little as possible, but there are practical challenges associated with this.

For example, to reduce artifacts caused by chemical shift and rapid signal decay,

which do not affect the TurboSPI image, the FSE image must acquire data very rapidly using a strong readout gradient and a correspondingly high readout bandwidth. However, unlike TurboSPI time course data, the acquired FSE data cannot be filtered without impacting the image quality, and the increased bandwidth leads to an increase in image noise. Of particular concern with the current experimental setup is a consistent noise source at approximately 50 kHz generated by the gradient amplifiers, which can be easily filtered out of TurboSPI data but which irreversibly contaminates any FSE image with a readout bandwidth larger than 50 kHz. Lowering the readout bandwidth below 50 kHz removes this source of noise but increases the prevalence of artifacts in the image. Repeated signal averaging can be used to reduce the effect of noise in the FSE image at the cost of increased imaging time.

Keeping in mind that the desired application of TurboSPI is imaging of animal models *in vivo*, the acceleration of the acquisition must only be enough to reduce the overall scan time to approximately 1-1.5 hours. Any further increases in imaging speed must be weighed against possible loss of image quality. For imaging with a matrix size of 128x128x32, an acceleration factor of 8 should be suitable for most practical applications, and for higher-resolution imaging (with matrix sizes such as 170x170x50) an acceleration factor of 15 should be sufficient.

Chapter 7

In Vivo Imaging with TurboSPI

The compressed sensing techniques presented in the previous chapter have enabled TurboSPI acquisitions to be performed at speeds suitable for *in vivo* imaging of animals, which can usually only be kept under anesthesia for 1.5-2 hours at a time. In this chapter we will consider some of the challenges specific to imaging in animal models, describe TurboSPI imaging protocols suitable for mouse imaging at 3 Tesla, and demonstrate cellular detection and relaxometry *in vivo*.

7.1 Challenges of Animal Imaging

MR imaging of animal models presents several technical obstacles which do not affect *in vitro* images such as the quantification examples given in Chapter 4. Aside from the obvious restriction on imaging time, respiratory and cardiac motion may introduce artifacts into the image depending on the area being studied. Though these motion artifacts can be reduced by gating to a cardiac or respiratory signal, this leads to longer scan times. As well, the natural background contrast of certain tissues may make detection of labeled cells difficult, and rapid transitions between tissue types may cause susceptibility gradients that affect the apparent relaxation rate in areas containing no SPIO.

Another challenge is the presence of fat, which is found throughout the bodies of most animals studied using MRI. As mentioned in Chapter 2, protons in different chemical environments have slightly different Larmor frequencies. In particular, the protons attached to lipid molecules have a frequency shifted by 3.5 parts per million (ppm) with respect to the Larmor frequency of water protons, and when the imaging region contains significant amounts of fat, a strong peak is visible in the NMR spectrum at this frequency. In an external magnetic field of 3 Tesla this shift is approximately 440 Hz.

In frequency-encoded images, the presence of fat leads to a chemical shift artifact,

in which the signal from the fat is misregistered as belonging to water protons at a slightly different spatial position. In TurboSPI acquisitions, signal from a voxel containing fat is not displaced from its true position, but the time course at that pixel is modulated at 440 Hz, with the strength of the modulation depending on the quantity of fat present. This is the source of the modulation seen in Figure 3.13, for example. Such modulation can obscure the desired relaxation effects, greatly diminishing the ability to quantify SPIO, and must be addressed in some way.

7.2 Methods

7.2.1 Fat Saturation

Since the frequency of the fat-induced modulation is known, it is theoretically possible to add an oscillatory term to the analytical model described in Chapter 5, such that the modulation can be distinguished from other effects. In practice, it is more effective to remove as much of the fat signal as possible from the image during the acquisition itself, rather than relying on post-processing of the acquired data. This can be done by applying a *fat saturation* pulse immediately before the excitation of the water signal, as illustrated in Figure 7.1. This pulse is not spatially selective, but designed to excite only a narrow range of frequencies centered 3.5 ppm away from the frequency of the water protons. The excited fat signal is then dephased with a spoiler gradient, and the remainder of the sequence proceeds normally before the fat signal can recover through T_1 relaxation.

To demonstrate the effects of fat saturation without any potential differences introduced by compressed sensing reconstructions, TurboSPI images were acquired of a mouse *ex vivo*, using a 128x128x16 matrix size covering a 30x30x15 mm field of view, for a nominal resolution of 0.25x0.25x1 mm. Eight echoes were acquired with $TE = 8$ ms and $TR = 250$ ms, sampling 512 points at 100 kHz during an acquisition window of 5 ms. The *ex vivo* acquisition allows these 136-minute scans to be collected using the same mouse in the same position within the imaging system. Images were acquired at 3 Tesla using the 30 mm inner diameter quadrature RF coil.

Because the image contrast changes significantly after the application of a fat saturation pulse, it was anticipated that the quality of images reconstructed with

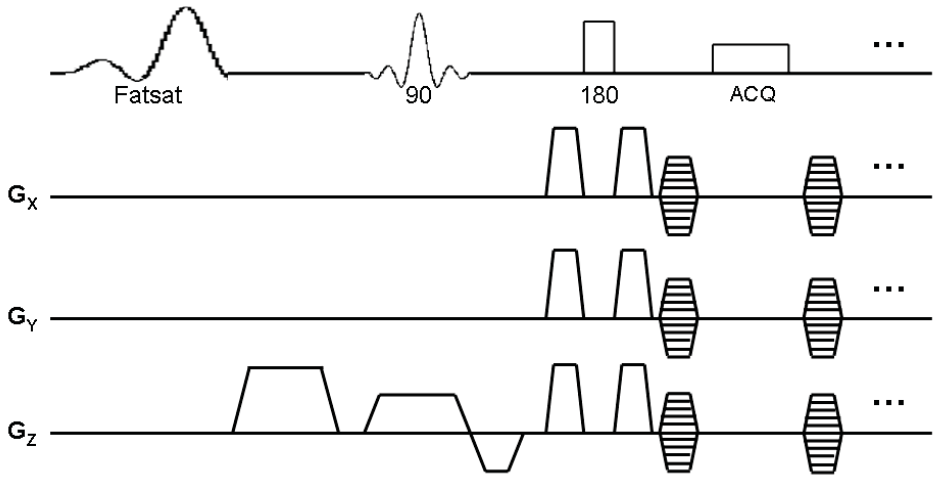

Figure 7.1: TurboSPI pulse sequence (first echo only) with a fat saturation pulse and an associated spoiler gradient. The fat saturation pulse is designed to excite a very narrow bandwidth, and therefore has a long duration of approximately 10ms.

compressed sensing could also be affected. To measure this effect, the two fully-sampled *ex vivo* datasets were retrospectively undersampled at acceleration factors between 3 and 15, with sampling patterns prescribed using an appropriate contrast-matched FSE image. For reconstruction, sparsifying transform and total-variation penalties of 0.1 each were used for all datasets. The normalized root-mean-square error with respect to the fully-sampled dataset was calculated at each point of the time series and used as a metric of reconstruction fidelity.

One of the central assumptions of the fat saturation process is that the saturation pulse will only affect the signal from fat, without having any effect on the water signal. Under most circumstances this assumption is reasonable, but in systems with broad linewidth and significant off-resonance signal, some fraction of the water protons will resonate within the frequency bandwidth of the saturation pulse, and signal from those protons will be reduced as if it were fat signal. This represents a potential problem when imaging high concentrations of SPIO, because some of the already limited signal may be further decreased by the fat saturation pulse, impacting TurboSPI's ability to perform accurate quantification.

The effect of fat saturation on fitted R_2^* relaxation rates was therefore assessed by imaging a number of samples containing 1.63 μm MPIO (as described in section 4.2.1) with and without fat saturation. R_2^* maps were computed following the procedure

described in section 4.2.3, and mean R_2^* values were obtained from identical ROIs in each dataset.

7.2.2 *In Vivo* Imaging Methods

For the *in vivo* demonstration of TurboSPI, a C57BL/6 (or "Black 6") mouse was imaged after injection with SPIO-labeled T-cells. 4 weeks prior to the experiment, 5 x 10^5 C3 (cervical cancer) cells were injected subcutaneously into the left flank, and allowed to establish a tumor. Tumor growth was monitored weekly with calipers. Effector T-cells (CD8+) were isolated from lymph nodes of naive mice and proliferated over 6 days in a medium similar to that described for the C3 cells used in Chapter 4. The cells were then incubated with Molday ION Rhodamine B at a concentration of 0.1 mg/mL for approximately 24 hours. 4 million cells were injected directly into the tumor and the mouse was immediately imaged.

After inducing anesthesia with 3% isofluorane, the mouse was placed in a nose cone that allowed administration of 1.5%-2% isofluorane throughout the scan to maintain anesthesia. The mouse was placed and secured within the RF coil, and its respiration and temperature were monitored during the scan, with a warm air blower used to maintain the mouse's temperature at 37° C. Total scan duration including setup was approximately 2.5 hours. All animal handling was done in accordance with protocols approved by the Dalhousie University Care for Laboratory Animals committee.

An initial image of the mouse was acquired using a balanced-SSFP sequence (256x170x170 matrix, 38.4x25.5x25.5 mm FOV, $\alpha = 30°$, TE = 4 ms, TR = 8 ms), which was acquired in 64 minutes with four signal averages and four acquired frequencies. The mouse was then imaged with TurboSPI using a 128x128x32 matrix, 25x25x20 mm FOV, 15 mm slab excitation, ETL=8, TR = 250 ms, and a fat saturation preparation. For both the b-SSFP and TurboSPI images, the RF coil used was the 25 mm inner diameter quadrature coil. In this particular coil, the quality of the automatic calibration was poor at typical echo times (7.5 ms), and a longer echo time of 9.5 ms was used to improve the echo alignment. Pulse durations used were 1000 μs for the shaped excitation and 500 μs for the hard refocusing pulses.

Low overall SNR in the coil necessitated extra signal averaging to ensure a high-quality FSE template; the FSE image was acquired in 16 minutes using 8 signal

averages. This was followed by a TurboSPI acquisition undersampled by a factor of 8 and acquired in 34 minutes with one average. The undersampling pattern was prescribed based on the FSE image as described in Section 6.2.1. 256 points were acquired during a 4 ms acquisition window at a sampling rate of 62.5 kHz, with the time course filtered to remove frequencies above 3 kHz. Reconstruction with the modified-CS algorithm required approximately 75 minutes for the complete time series, using transform and total-variance weights of 0.1 each.

7.3 Results

7.3.1 Fat Saturation - Effect on Image Quality

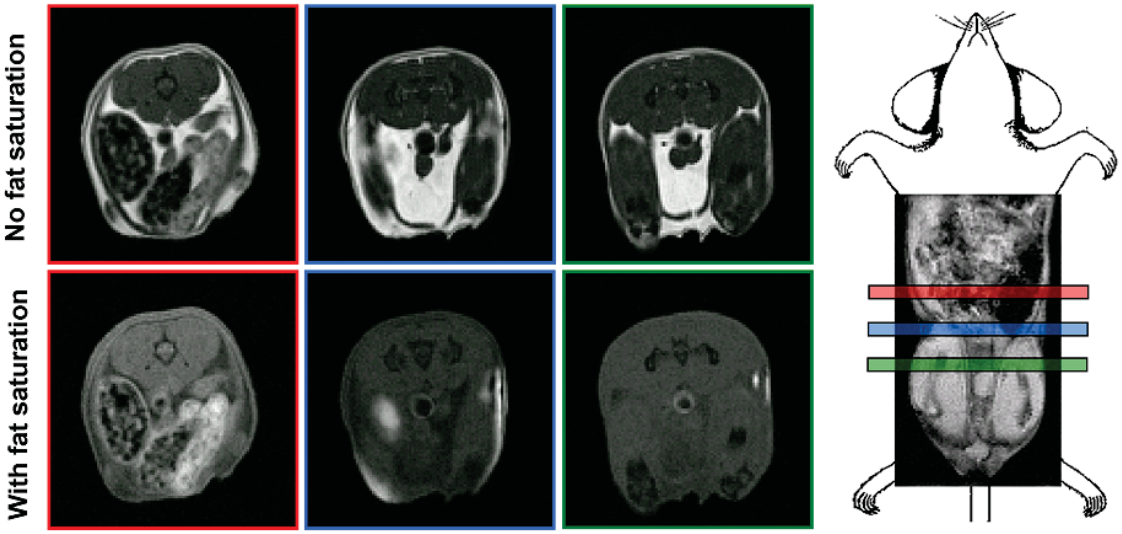

Figure 7.2: Comparison of TurboSPI images without (top row) and with (bottom row) fat saturation. The use of fat saturation greatly reduces the overall contrast in the images by suppressing signal from regions containing fat. Approximate slice positions within the mouse are shown at right; the MR data superimposed on the schematic is from a localizer scan.

Representative slices from the t = TE = 8 ms volume of *ex vivo* mouse images with and without fat saturation are shown in Figure 7.2, and demonstrate that the overall contrast is significantly altered with fat saturation enabled. Without fat saturation, areas such as the skin and other fat-rich tissues appear very bright, but with fat saturation, the bulk of the image appears more homogeneous. Figure 7.3 shows

the same slices but from a later time point (t = 9 ms) to demonstrate how the contrast changes throughout the time course. The images without fat saturation demonstrate an "India Ink" artifact [134] in which the tissues containing large quantities of fat appear to be outlined with dark contrast while remaining relatively bright (though their intensity has decreased relative to Figure 7.2). This artifact is much less pronounced with fat saturation enabled, though areas of high fat content still appear dark.

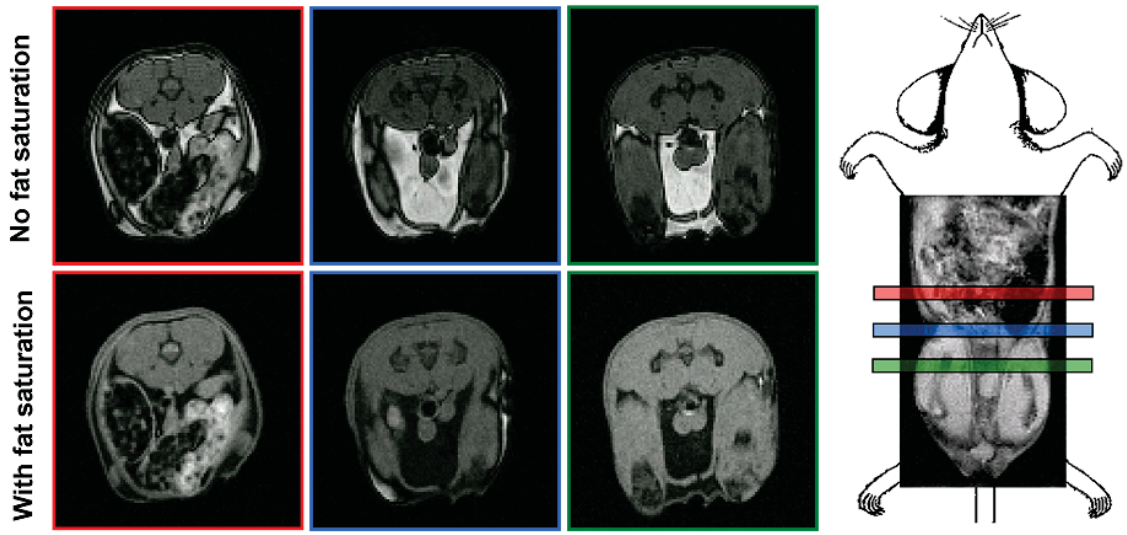

Figure 7.3: Away from t=TE, images without fat saturation (top row) are affected by darkening of the boundaries between tissues ("India Ink" artifact). Images acquired with fat saturation are affected by this artifact to a much lesser degree. Approximate slice positions within the mouse are shown at right.

These effects are reflected in the time courses obtained from the TurboSPI image series, and the impact of fat saturation depends heavily on the amount of fat present in a given location. Figure 7.4 shows that, for a time course in an area with little fat, the time course is unaffected as desired. For a time course in a region dominated by fat signal (high intensity on the T_2 weighted image at t=TE) the modulation at 440 Hz is evident, and though it cannot be removed entirely by the use of fat saturation, its intensity is greatly diminished. For regions on a boundary between these types of tissues, which contain a mixture of water and fat, the modulation is removed while the average signal level remains constant. These are the areas which produce the contrast seen in Figure 7.3, which the use of fat saturation greatly reduces.

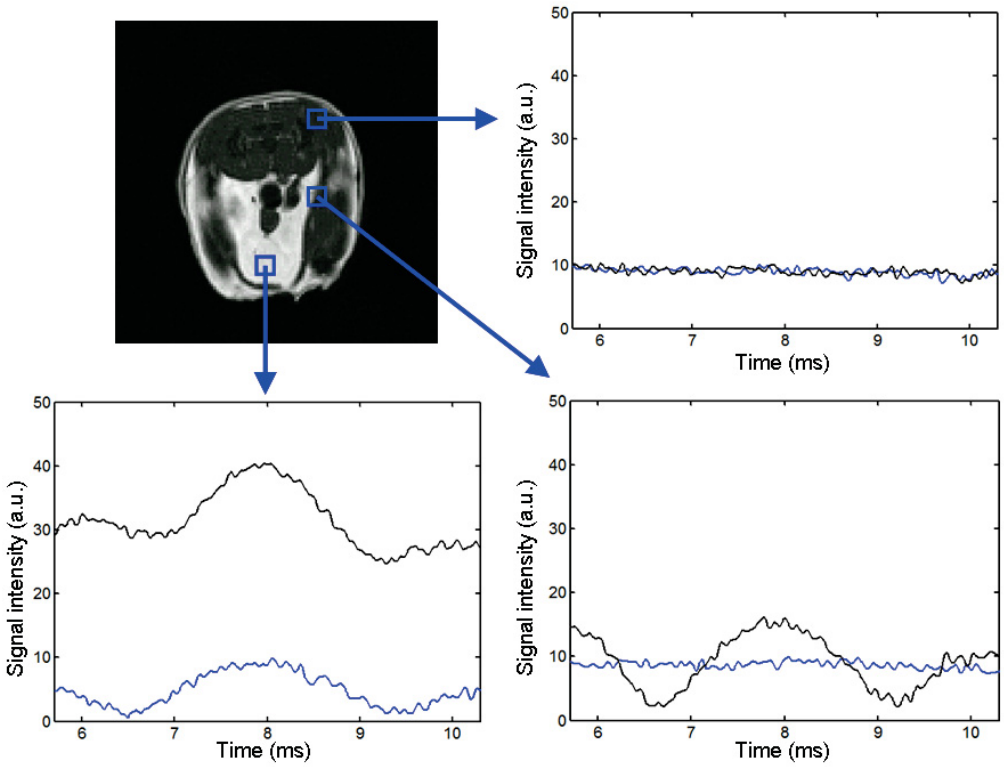

Figure 7.4: Time courses from TurboSPI images acquired with (blue lines) and without (black lines) fat saturation. Top right: time course from a region containing little fat. Bottom left: Time course from a region with large quantities of fat. Bottom right: Time course from a region between tissue types.

7.3.2 Fat Saturation - Effect on Compressed Sensing

As demonstrated in the previous chapter, TurboSPI images reconstructed with compressed sensing and initially guided by a matched FSE image differ most from the unaccelerated data when changes in contrast are highest, and the quality of the resulting reconstruction is best evaluated around this point, which for these *ex vivo* images was t = 9 ms. Results of such a comparison for a representative slice at that time point are shown in Figure 7.5.

In both cases the quality of the reconstructed image degrades with increasing acceleration, but in the images with fat saturation off, the rapidly varying contrast around tissues rich in fact leads to severe artifacts, which grow more significant at high acceleration. No such artifacts appear when fat saturation is enabled because of the altered contrast. Figure 7.6 plots the normalized RMS error as a function of

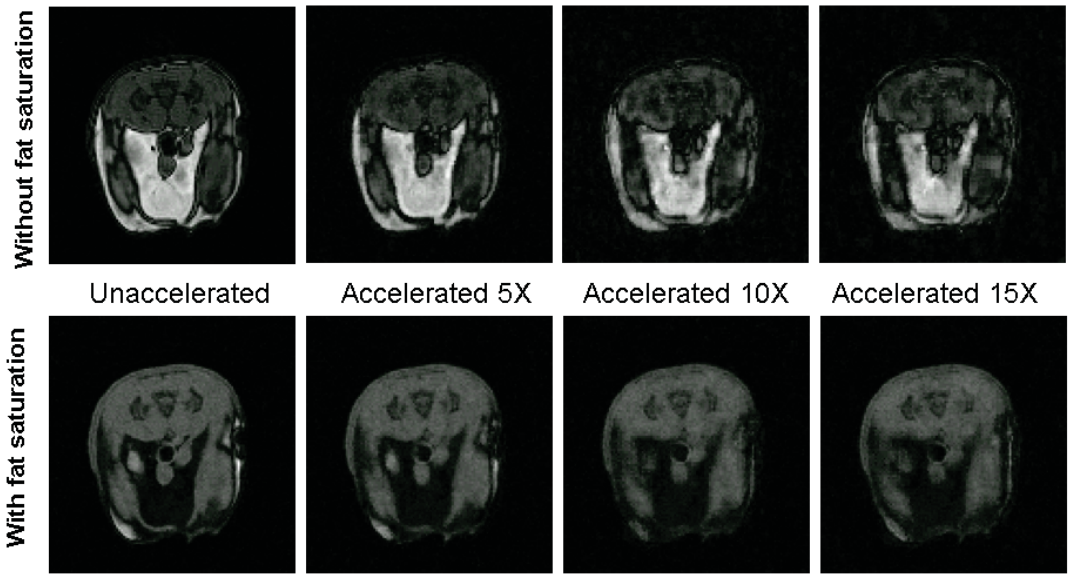

Figure 7.5: The use of fat saturation has a visible impact on the quality of images reconstructed with compressed sensing. Top row: Representative slice at t = 9 ms from the dataset with fat saturation off, reconstructed after various degrees of undersampling. Bottom row: The same slice from equivalently undersampled and reconstructed datasets with fat saturation on.

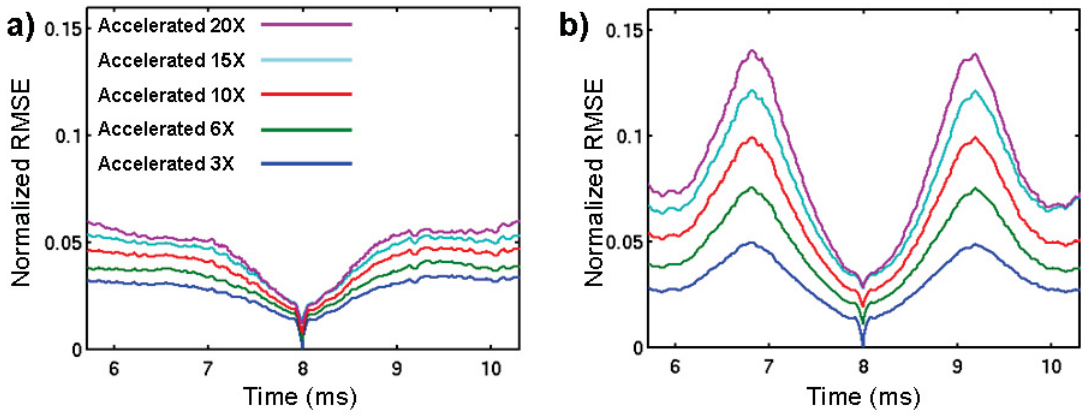

Figure 7.6: Fat saturation significantly reduces the mean squared error of reconstructed data with respect to fully-sampled data. (a) nRMSE at each point in the time series, for undersampling factors from 3 to 15, with fat saturation on. (b) nRMSE with fat saturation off. Plots use same scale and legend.

acceleration factor, further supporting the increased accuracy of reconstruction with fat saturation.

7.3.3 Fat Saturation - Effect on Quantification

The impact of fat saturation on the ability of TurboSPI to quantify SPIO is demonstrated in Figure 7.7, showing that quantification appears unaffected at low to moderate iron concentrations (corresponding to relaxation rates of $R_2^* < 400$ s^{-1}). However, as the iron concentration rises, the ability of TurboSPI to accurately measure R_2^* does appear to be affected to some degree. Even though the R_2^* values measured with fat saturation are not significantly different due to the large error bars, there is a visible trend of these rates appearing lower than those measured without fat saturation.

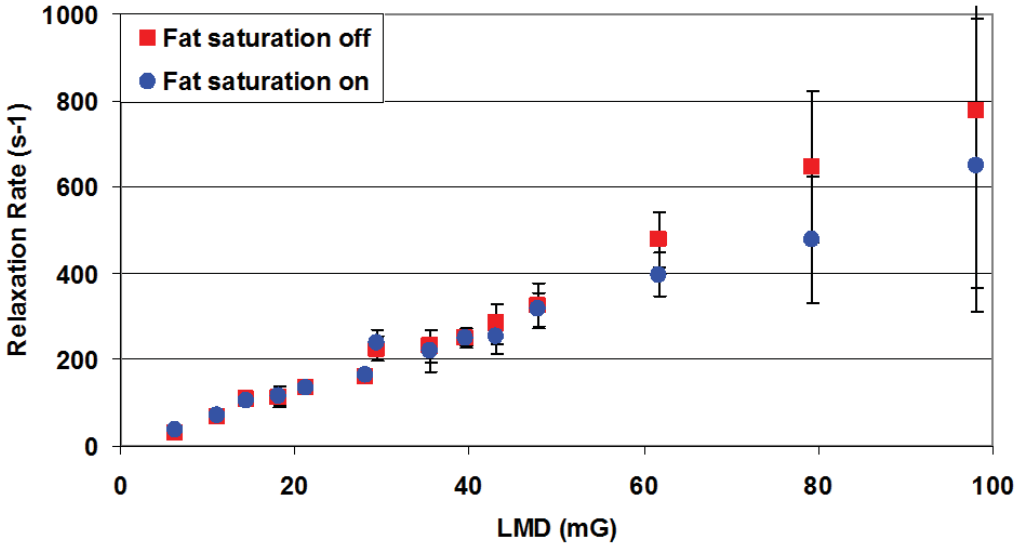

Figure 7.7: Effects of fat saturation on R_2^* quantification of MPIO samples. While the relaxations rates agree at low to moderate concentrations, there is some disagreement at higher concentrations, with R_2^* values measured with fat saturation on (blue circles) somewhat lower than those measured without fat saturation (red squares).

7.3.4 Detection of Labeled Cells

Figure 7.8 shows a mouse tumor after injection of SPIO-labeled cells, as imaged by b-SSFP and TurboSPI. The b-SSFP image (Fig. 7.8(a)) reveals a significant area of

negative contrast within the tumor as anticipated. Neither the FSE template image (Fig. 7.8(b)) nor the corresponding TurboSPI image at t=TE (Fig. 7.8(c)) show much contrast, though due to the low R_2 of SPIO-labeled cells, significant changes in contrast are not anticipated around t=TE since most of the surrounding signal should have rephased. However, TurboSPI images further away from the echo peak (Fig. 7.8(d)) show pronounced negative contrast throughout a volume similar to that observed with b-SSFP.

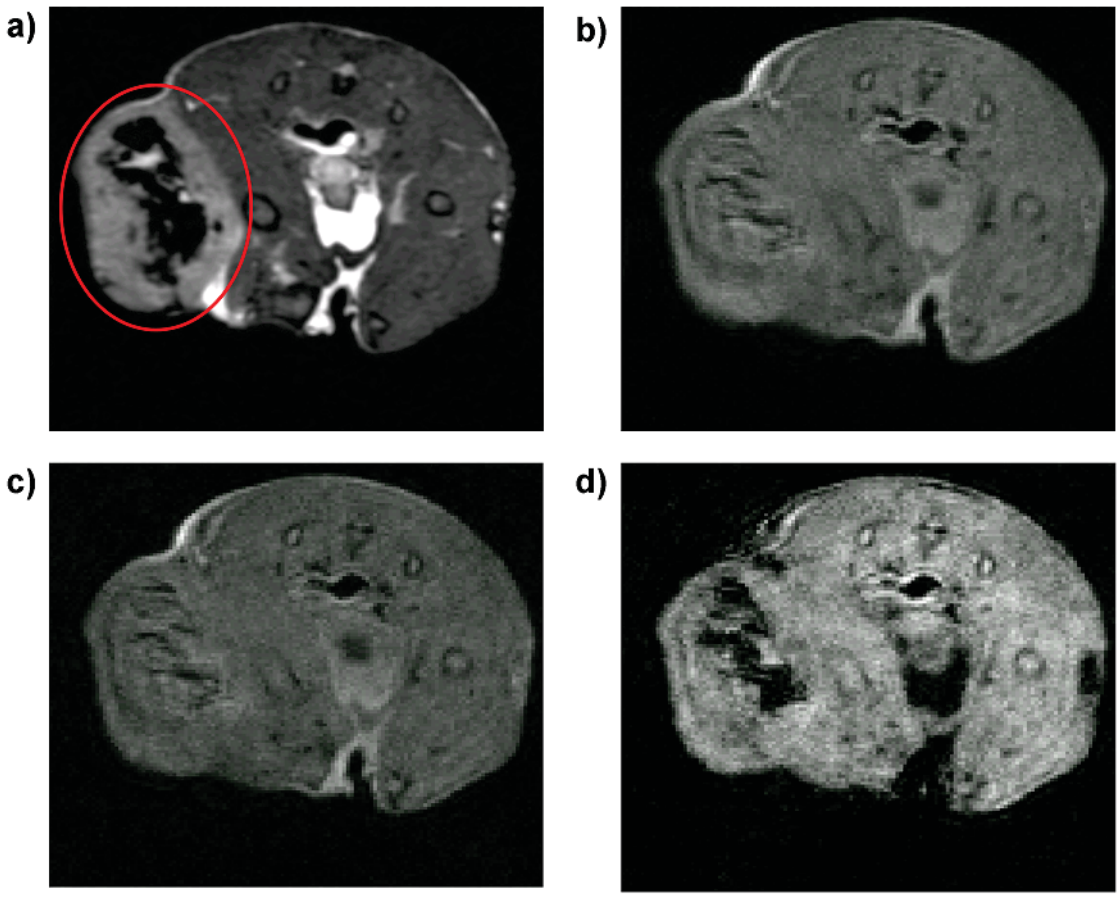

Figure 7.8: Comparison of images showing SPIO-labeled cells within a mouse tumor *in vivo*. (a) b-SSFP image, with no fat saturation. The tumor is circled in red. (b) FSE template image with fat saturation. (c) TurboSPI image at t=TE=9.5ms. (d) TurboSPI image at t=10.5ms. The significant darkening observed throughout the tumor on the b-SSFP image also appears in the TurboSPI data away from TE.

A representative time course from within the tumor is shown in Figure 7.9(a). As in the case of SPIO-loaded cells *in vitro* and as predicted by the analytical model, the

time course is fairly symmetric with a slightly rounded and shifted peak. The time courses throughout the animal can be used to generate an R_2^* map using the methods of Chapter 3, the result of which is shown in Figure 7.9(b). R_2^* relaxation rates of 1500 s^{-1} or higher are observed throughout the tumor, indicating very high iron concentrations which can nonetheless be quantified with TurboSPI. High R_2^* values are also observed in several regions outside the tumor. As seen by comparison with the b-SSFP image, these correspond to regions rich in fat, where modulation of the signal due to off-resonance effects leads to spurious fits.

An additional feature of the TurboSPI image is demonstrated in Figure 7.10, where a significant region of negative contrast seen on the b-SSFP image does not appear in the TurboSPI data, nor in the corresponding R_2^* map. It is possible that this region does not represent SPIO-labeled cells, but some other tissue-related effect such as necrosis, which appears with similar contrast on the b-SSFP image but is not detected by TurboSPI.

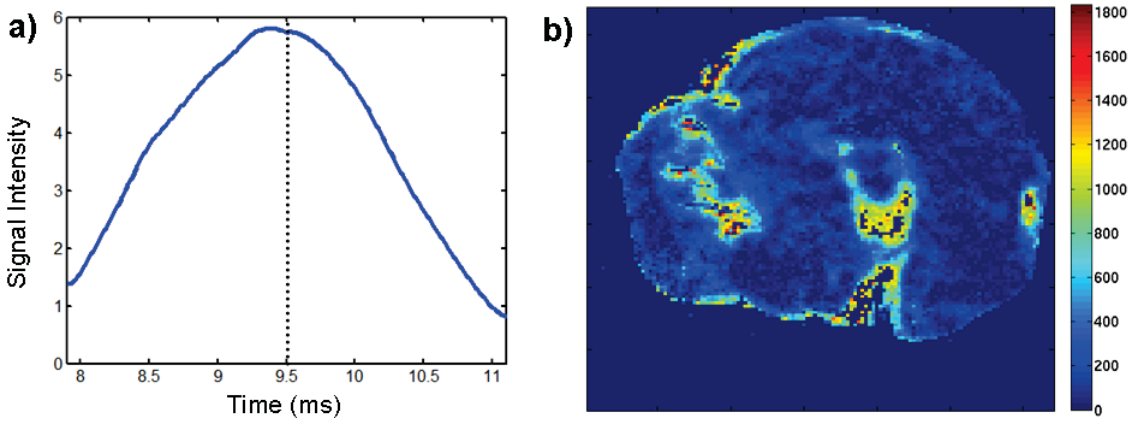

Figure 7.9: Basic relaxometry of *in vivo* TurboSPI data. (a) Time course from a representative voxel within the area of negative contrast, showing a time course characteristic of SPIO-labeled cells. (b) R_2^* map generated from a fit to the decaying portion of the spin echo throughout the entire image.

7.3.5 Discussion and Summary

The use of fat saturation has a significant impact on the quality of TurboSPI images, both in terms of reconstruction quality and time course fidelity. While the current implementation of compressed sensing for TurboSPI allows higher acceleration

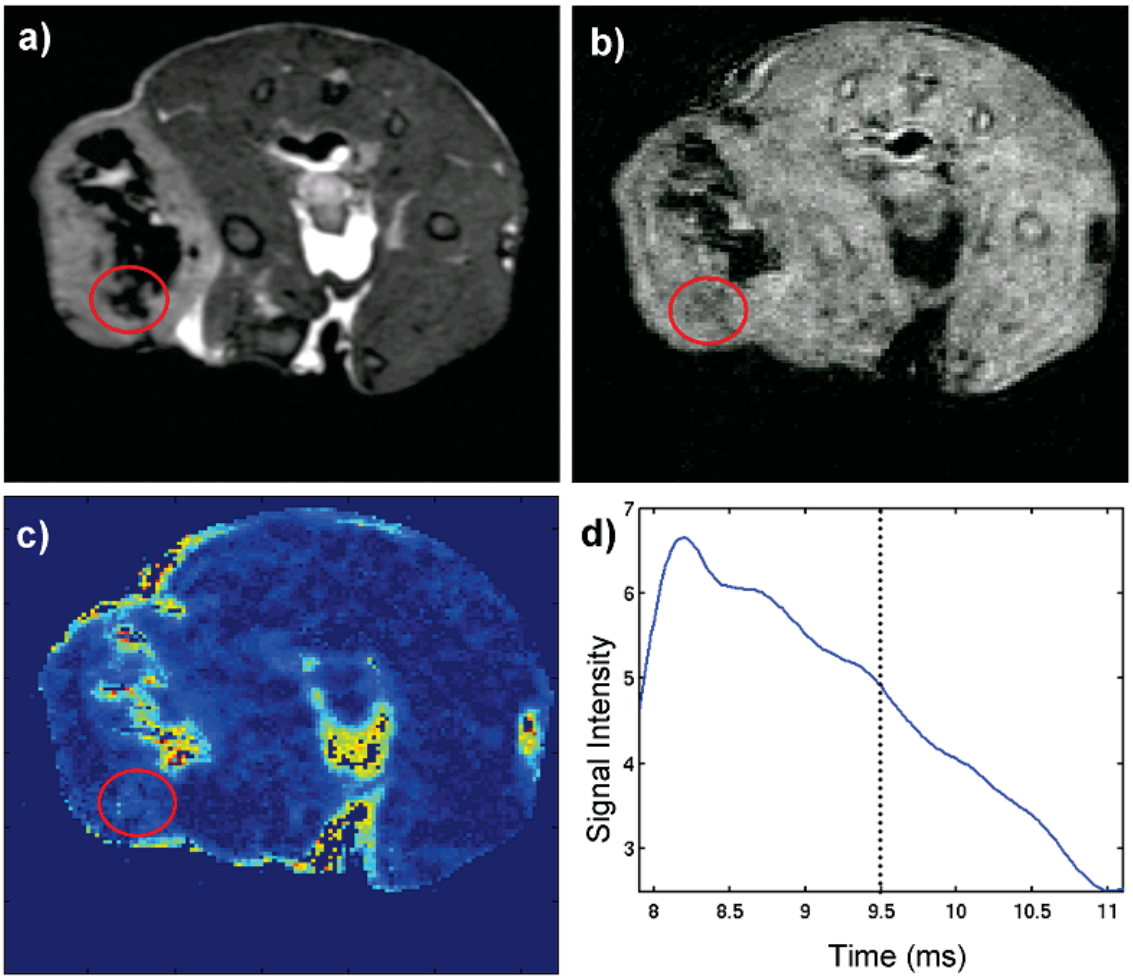

Figure 7.10: TurboSPI is potentially more specific to SPIO-labeled cells than b-SSFP. (a) b-SSFP image; circled region contains significant negative contrast. (b) A TurboSPI image at the same anatomical position shows no such contrast. (c) Most of the fitted R_2^* values in that region are not significantly different from the background, and (d) a typical time course in the region does not show evidence of relaxation consistent with labeled cells. This region may represent tumor necrosis rather than labeled cells.

of images with fat saturation, there is also evidence to suggest that quantification of very high iron loads may be impaired in such images. The question of whether to use fat saturation may therefore have to be answered on a case-by-case basis, depending on the desired resolution and anticipated SPIO concentration.

Another practical difficulty associated with the use of fat saturation is that it may render certain organs and anatomical structures more difficult to discern, by suppressing the contrast that would normally distinguish them from surrounding

tissue. If better tissue definition is required in a study where fat saturation is used, we suggest that a second FSE guide be acquired with fat saturation off, potentially using fewer signal averages to reduce the impact on overall imaging time.

An alternative to the use of fat saturation is a spectral-spatial RF pulse [135], which can be used to selectively excite water protons at a particular spatial location, effectively combining the fat saturation pulse with the standard excitation pulse. Such pulses require a time-varying gradient to be applied during the pulse, and are generally more robust against field inhomogeneity. Some preliminary testing was done with a spectral-spatial excitation preceding a TurboSPI acquisition, and though quantification did seem to be affected as in the case of fat saturation, the onset of visible effects did not occur until roughly 600-800 s^{-1} (data not shown). With additional refinements, this might become a feasible alternative to fat saturation.

Regardless of the approach used to reduce fat signal, it is challenging to completely remove its influence from TurboSPI images, and this could represent an impediment to accurate quantification in some circumstances. For example, if the acquisition window is not wide enough to encompass at least one full period of fat-induced modulation, the modulation can be confounded with R_2^* relaxation, producing spurious voxels on R_2^* maps. This is the source of the areas of high apparent R_2^* outside the tumor in Figure 7.9. It may be possible to remove these spurious fits using methods like those of section 5.4.2, if enough data is available for an accurate fit. This might entail increasing the acquisition window and therefore the effective echo time.

Though attempts have been made to use TurboSPI for detection of SPIO-labeled cells in tumors or lymph nodes following tail vein injections in mouse models, these attempts have not yet been successful. While corresponding b-SSFP images do show signs of SPIO in these tissues, TurboSPI images of the same mouse do not have time courses indicative of SPIO. This may reflect different detection limits between the two sequences, or it may reflect low-contrast information that is being lost during the reconstruction of undersampled data. Efforts to explore these differences and improve the detection capability of TurboSPI are ongoing. In the end, TurboSPI may simply be more well-suited to applications in which there are larger concentrations of cells which remain localized to a particular tissue for an extended period of time, providing significant and stable contrast.

Nonetheless, this work has demonstrated the ability of TurboSPI to acquire *in vivo* images of animal models with scan times that are feasible for preclinical studies. These images are comparable in quality to those obtained with existing protocols, and in areas known to contain large concentrations of SPIO-loaded cells, show signal time courses similar to those observed in loaded cells *in vitro*. Such time courses should be suitable for quantification by R_2^* relaxometry, and these results could be compared to existing signal-based quantification methods and later verified with histology.

Chapter 8

Future Work and Conclusions

8.1 Future Work

The results presented in this work represent the basis for *in vivo* quantification of SPIO-labeled cells with TurboSPI, a technique that had not previously been applied to biological systems. These results can be extended in a number of ways to improve the accuracy of data reconstruction and SPIO quantification, which will benefit potential applications of cellular imaging with TurboSPI.

8.1.1 Modeling Relaxation Behavior

The most significant step remaining in the practical application of the analytical model outlined in this work is the ability to use it for fitting experimental data, as described in section 5.4.1. As noted during that discussion, the best method for robustly fitting the approximate and/or asymptotic forms to the acquired data is still under investigation. The main difficulty in performing these fits is separating the contributions of the various factors involved in determining the time course shape, which include SPIO-related parameters as well as other variables such as local field inhomogeneities that may not be easily distinguishable from relaxation effects. Adjusting parameters such as the acquisition window and echo spacing may assist in producing time courses with basis functions that are more readily separated.

Regardless of the fitting method which is to be employed, experimental verification will be necessary, potentially using a series of samples like those described in Chapter 4, but with parameters chosen to test the ability of the fitting method to separate changes in ζ from changes in $\delta\omega$.

Though all of the examples tested in this work assumed that every perturber in a voxel was the same size, realistically a distribution of sizes can be expected. Using equation (4.8), the model developed in Chapter 5 can be used to describe arbitrary

distributions of particle size, and the Monte Carlo simulations can be modified accordingly to verify the model output. Though experimental measurements would not likely be sensitive enough to accurately extract much information about particle size distributions, it might be possible to distinguish the contributions of two different particle sizes, such as two mixed cell populations, or large SPIO particles that have been released from dead cells.

Another refinement which could be incorporated into the model is an extended description of signal behavior during multi-echo experiments, similar to that found in Jensen and Chandra's model [100] but which could also describe signal evolution away from the echo times nTE. The present work has assumed that the first echo of the TurboSPI train will contribute most heavily to overall signal behavior, since the center of k-space is sampled during the first echo, but the validity of this assumption could be further explored with simulations and experimental data.

While the analytical model accurately describes the behavior of systems near the static dephasing regime, as the particle's size becomes smaller or the diffusion coefficient increases, the agreement between the model and experimental data decreases. This is likely due to violations of the model's central assumption of a linear local field variation. A more sophisticated model might relax this restriction, though it would remain to be seen if the result would be too complex to be practically useful. Alternatively, a model such as the strong collision approximation of Bauer and colleagues [102] could be applied to cases outside the current model's range of applicability.

Another potential modification to the model is centered on the F function (equation (5.14)), which relates to diffusive attenuation of signal and has a significant effect on the shape of the echo peak. The propagator-based method outlined in Appendix A.3 is only one possible way to derive F; for example, Jensen and Chandra arrived at an equivalent expression by performing an integration over diffusion paths ([100] equation (24), using a "straightforward calculation" that is not elaborated upon). However, attempting to derive the F function through a direct solution of the Bloch equations (as outlined, for example, in [96]) yields a different result, which has the same asymptotic behavior at $t >> TE$ and $t = TE$ but differs elsewhere. This discrepancy has not been explained, but understanding its source might offer more insight into the underlying assumptions of the model.

8.1.2 Compressed Sensing

The modified-CS algorithm described in Chapter 6 has proven adequate for reconstruction of TurboSPI data at moderate undersampling factors, but it is certainly possible that further improvements in reconstruction speed and accuracy could be achieved. As mentioned in the initial discussion on compressed sensing, a number of alternative methods can be used to perform the non-linear optimization, each of which has different advantages and capabilities. For example, a recently proposed algorithm which exploits group sparsity in sets of images [136] could be tested with TurboSPI data.

Other improvements to the existing CS reconstruction can also be explored. Before the implementation of the modified-CS algorithm described in Section 6.2.1, the basic reconstruction provided by the Matlab SparseMRI toolbox was used for testing. When reconstructing data with that algorithm, which does not use prior information to constrain the reconstruction, it was found that processing multiple contiguous time points simultaneously was often more effective than reconstructing those time points individually, mainly due to reduced accumulation of noise or artefacts. This behavior was not observed in the current implementation of modified-CS, but it is possible that further alterations to the algorithm could yield such improvements.

With improvements in the processing hardware available, it may also become possible to reconstruct the entire 4D dataset as a single unit. Since the desired time courses are likely sparse under appropriate transforms, the ability to sparsify the data in that dimension would be valuable. Furthermore, since the approximate form of the time course is known, it might be possible to use a model-based reconstruction approach [137] to directly reconstruct parameter maps instead of deriving them from reconstructed images. As with basic R_2^* relaxometry, this approach would have to be verified against a fully-sampled dataset to ensure accurate quantification.

Alternative methods for random undersampling are also being proposed by a number of groups, such as the Poisson-disc sampling described in [138], which is claimed to produce less coherent aliasing artifacts in the wavelet domain. Implementation of such a sampling procedure, especially if it can also include the k-space intensity distribution of a template image as demonstrated in this work, might reduce image artifacts at higher undersampling factors.

Though the results of R_2^* relaxometry are not significantly altered by acceleration at factors up to 10, the effect of the reconstruction on other areas of the time course has not been demonstrated. When a more robust method for fitting the approximate form determined in Chapter 5 has been developed, such fits can be applied to accelerated data to assess the impact of compressed sensing on more advanced quantification procedures.

Finally, clinical and preclinical MRI systems are increasingly being outfitted with multiple-element RF coils and receivers which are capable of parallel imaging, as described briefly in Section 6.1. In many ways parallel imaging and compressed sensing are complementary methods of image acceleration; parallel imaging is hardware-based and requires uniform undersampling, while compressed sensing is software-based and relies on random undersampling. Combinations of these techniques have been proposed, most notably L1-SPIRiT [122] which uses the coil sensitivity information required for parallel imaging to constrain the CS reconstruction. While the MRI system used for the experiments described in this work was not equipped with the hardware to allow parallel imaging, this is a natural direction in which to continue this work, and would potentially allow even further acceleration of TurboSPI to the point where it might become practical for clinical applications.

8.1.3 *In Vivo* Imaging

As noted in the last chapter, while large quantities of SPIO-loaded cells have been detected *in vivo* using TurboSPI, detection of smaller numbers of cells remains elusive. In order to determine whether this is a limitation of the TurboSPI sequence itself or of the compressed sensing reconstruction, it would be ideal to image a mouse containing SPIO in the lymph node using b-SSFP followed by *ex vivo* TurboSPI with full k-space sampling. Experiments of this type should be possible in the near future.

The changes in contrast between b-SSFP and TurboSPI may indicate a different response to tissue-related effects such as tumor necrosis which, when present, may reduce the effectiveness of SPIO quantification through signal-based methods. TurboSPI's apparent insensitivity to tumor necrosis would have to be investigated more thoroughly before any conclusions could be drawn; the effects must be observed in multiple animals with subsequent *ex vivo* histological confirmation of necrotic tissue.

Though the chosen application of TurboSPI in this work has been the monitoring of immunotherapy, opportunities to test TurboSPI for other applications, such as the imaging of implanted stem cells, should not be neglected if and when they become available. TurboSPI may be more well-suited to imaging the very high concentrations of cells used in stem cell treatments than the relatively lower concentrations present during immunotherapies.

8.2 Conclusions

This work has demonstrated that TurboSPI, an MRI technique originally developed for imaging of materials, can also be successfully applied to the quantitative imaging of cells, even in living organisms. As a purely phase-encoded sequence, TurboSPI is less prone to several artifacts that affect traditional imaging techniques, and it retains signal even in areas of very high R_2^* relaxation so long as R_2 relaxation remains low. This permits quantification of SPIO over a large dynamic range, including the high SPIO concentrations that would be employed in some longitudinal cellular imaging studies. This range is equivalent to, or better than, that provided by most gradient-echo methods.

However, the chief advantage of TurboSPI is not necessarily its dynamic range, but its temporal resolution, which is far higher than can be achieved by traditional techniques which employ a readout gradient. This temporal resolution allows observations of relaxation behavior in unprecedented detail. On a practical level this may lead to improved quantification or robustness against confounding effects such as field inhomogeneities. As well, the high temporal resolution measurements obtained by TurboSPI have also provided an opportunity to refine the theoretical description of MRI signal evolution in the presence of SPIO.

In the current work, effects were observed that were not predicted by the existing static dephasing theory, but which could be explained to some degree by an expanded model that incorporates the effects of slow diffusion near spherical magnetic perturbers. This model, when compared against Monte Carlo simulations and experimental data, shows good accuracy over a range of parameters that describes SPIO-loaded cells. Though a robust algorithm for fitting this model to data has not yet been completed, when such a procedure is developed it should become possible

to estimate parameters that are inaccessible to existing relaxometry techniques, such as the separate contributions of SPIO susceptibility and cellular density to the R_2^* relaxation rate. Compensation for effects such as field inhomogeneities may also be possible if these effects can be reliably modeled and distinguished from relaxation behavior.

Though characterized by lengthy acquisition times, TurboSPI can be implemented for use *in vivo* with the aid of compressed sensing, which allows images of acceptable quality to be reconstructed from highly undersampled k-space data. Acquisition of a matched FSE template image, which is achieved by adding a readout gradient to the standard TurboSPI sequence, can be used to prescribe the undersampling pattern and constrain the reconstruction, allowing acceleration factors of 8-10 at moderate resolution and potentially up to 30 at high resolution, though in practice such high acceleration may introduce artifacts and may not be beneficial. The use of saturation pulses to suppress off-resonance fat signal also contributes to improved image quality, though with potential reductions in the effective range of quantification.

Imaging of SPIO-labeled cells with the current implementation of TurboSPI has been demonstrated *in vivo* with reasonable acquisition times. The resulting images are of comparable quality to those obtained by traditional methods and may show increased specificity to SPIO, allowing accumulations of labeled cells to be distinguished from other effects such as tumor necrosis that would confound signal-based analysis methods.

However, the present work has only demonstrated detection of large groups of cells after direct injection; visualization of smaller groups of cells that have migrated into tissue from a distant injection site remains challenging, and this would be necessary for most applications involving imaging of immune therapies. Modifications to the TurboSPI sequence or the reconstruction procedure will be required to improve image quality and lower the SPIO detection limit, and suitable fitting procedures will have to be developed before direct quantification of cellular density can be achieved under these experimental conditions. Alternative imaging applications such as monitoring of stem cell implants may also serve as a more effective use of TurboSPI *in vivo*.

Appendix A

A.1 Field of a Magnetized Sphere

In section 4.1.2 the following expression was given for the inhomogeneous magnetic field surrounding a uniformly magnetized sphere of radius R and susceptibility difference $\Delta\chi$ when placed within an external field of $\vec{B_0} = B_0\hat{z}$.

$$B(r,\theta,\phi) = \frac{4\pi}{3}\Delta\chi B_0 \left(\frac{R}{r}\right)^3 (3\cos^2\theta - 1) \tag{A.1}$$

This is equivalent to finding the magnetic field around a uniformly magnetized sphere with $\vec{M} = M_0\hat{z} = 4\pi\Delta\chi B_0\hat{z}$, and no background field. Such a boundary-value problem can be solved by using Laplace's equation for the scalar potential Φ.

For a three-dimensional problem with azimuthal symmetry, the scalar potential can be defined in terms of Legendre polynomials P_l as

$$\Phi(r,\theta) = \sum_{l=0}^{\infty}[A_l r^l + B_l r^{-(l+1)}]P_l(\cos\theta) \tag{A.2}$$

The coefficients A_l and B_l are determined by the boundary conditions. Since the potential must be finite at $r = 0$, inside the sphere $B_l = 0$. Similarly, $A_l = 0$ outside the sphere since the potential must remain finite at $r = \infty$. This lets us define the potential piecewise as

$$\Phi_{in}(r,\theta) = \sum_{l=0}^{\infty} A_l r^l P_l(\cos\theta) \tag{A.3a}$$

$$\Phi_{out}(r,\theta) = \sum_{l=0}^{\infty} B_l r^{-(l+1)} P_l(\cos\theta) \tag{A.3b}$$

Because of the boundary conditions on B and H at the surface of the sphere, the following conditions on Φ and its radial derivative can also be derived [139]:

$$\Phi_{in}(r = R) = \Phi_{out}(r = R)$$
$$\frac{\partial\Phi_{in}}{\partial r}\Big|_{r=R} - \frac{\partial\Phi_{out}}{\partial r}\Big|_{r=R} = \vec{M}\cdot\hat{n}$$

Since the magnetization of the sphere is in the $\hat{z}$ direction, its component normal to the surface is $\vec{M} \cdot \hat{n} = M_0 \hat{z} \cdot \hat{n} = M_0 \cos\theta$. Substituting the expressions for Φ into the above equations, we obtain

$$\begin{aligned} \Phi_{in}(r = R) &= \Phi_{out}(r = R) \\ \sum_{l=0}^{\infty} A_l R^l P_l(\cos\theta) &= \sum_{l=0}^{\infty} B_l R^{-(l+1)} P_l(\cos\theta) \\ \Rightarrow A_l &= B_l R^{-(l+2)} \end{aligned} \tag{A.4}$$

$$\begin{aligned} M_0 \cos(\theta) &= \frac{\partial \Phi_{in}}{\partial r}\Big|_{r=R} - \frac{\partial \Phi_{out}}{\partial r}\Big|_{r=R} \\ &= \sum_{l=0}^{\infty} l A_l R^{l-1} P_l(\cos\theta) + \sum_{l=0}^{\infty} (l+1) B_l R^{-(l+2)} P_l(\cos\theta) \end{aligned}$$

Since the only Legendre polynomial containing $\cos\theta$ by itself is P_1, this means that only the $l = 1$ term contributes to the expansion. Then

$$A_1 + 2B_l R^{-3} = M_0 \tag{A.5}$$

Combining (A.4) and (A.5), and substituting into (A.3), gives

$$\Phi_{in}(r, \theta) = \frac{1}{3} M_0 r \cos\theta \tag{A.6a}$$

$$\Phi_{out}(r, \theta) = \frac{1}{3} M_0 \frac{R^3}{r^2} \cos\theta \tag{A.6b}$$

Since the magnetic field is the negative gradient of the scalar potential, we can easily obtain expressions for the field inside and outside the sphere. In spherical coordinates, the gradient operator is

$$\nabla = \frac{\partial}{\partial r}\hat{r} + \frac{1}{r}\frac{\partial}{\partial \theta}\hat{\theta} + \frac{1}{r\sin\theta}\frac{\partial}{\partial \phi}\hat{\phi} \tag{A.7}$$

Applying this to (A.6) yields

$$B_{in}(r, \theta) = -\frac{1}{3} M_0 \cos\theta \hat{r} + \frac{1}{3} M_0 \sin\theta \hat{\theta} \tag{A.8a}$$

$$B_{out}(r, \theta) = \frac{2}{3} M_0 \frac{R^3}{r^3} \cos\theta \hat{r} + \frac{1}{3} M_0 \frac{R^3}{r^3} \sin\theta \hat{\theta} \tag{A.8b}$$

Since we are only concerned with the field in the direction of $\hat{z} = \cos\theta\hat{r} - \sin\theta\hat{\theta}$, the quantities of interest are

$$B_{z,in}(r,\theta) = -\frac{1}{3}M_0 \tag{A.9a}$$

$$\begin{aligned} B_{z,out}(r,\theta) &= \frac{1}{3}M_0\frac{R^3}{r^3}\left(2\cos^2\theta - \sin^2\theta\right) \\ &= \frac{1}{3}M_0\frac{R^3}{r^3}\left(3\cos^2\theta - 1\right) \end{aligned} \tag{A.9b}$$

For our purposes, the expression of interest is $B_{z,out}$. Substituting our previous definition of M_0 yields the desired expression,

$$B_{z,out}(r,\theta) = \frac{4\pi}{3}\Delta\chi B_0(\frac{R}{r})^3\left(3\cos^2\theta - 1\right) \tag{A.10}$$

A.2 Static Dephasing Condition for Spin-Echo Experiment

In this section we shall derive the condition (4.30) for a system of spherical perturbers to be in the static dephasing regime during a spin-echo experiment.

As mentioned in section 4.1.4, the diffusive attenuation during a spin-echo experiment with echo time TE was given by Majumdar and Gore [93] as

$$\exp\left[-\frac{1}{12}D\cdot\left\langle(\gamma G)^2\right\rangle TE^3\right] \tag{A.11}$$

Before proceeding with an exact evaluation, we note some general assumptions. We assume that the volume fraction ζ occupied by perturbers is small, such that field distortions from neighboring perturbers do not overlap and each can be treated independently. This allows us to work with a single object instead of a network. Also, the averaging in the above expression will take the form of a volume integral:

$$\left\langle(\gamma G)^2\right\rangle = \frac{\int(\gamma G)^2 dV}{\int dV} \tag{A.12}$$

A suitable choice of volume, given our previous assumption, is the volume around a particular perturber that is free of other perturbers. This can be written in terms of the volume fraction ζ and the volume V_p of a particular perturber as

$$V = \frac{V_p}{\zeta} \tag{A.13}$$

This will replace $\int dV$. For the upper integral we neglect the interior of the perturber, integrating outward from zero, and since the integrand goes to zero faster than r^{-2}, we can reasonably extend the upper limit to infinity:

$$\left\langle (\gamma G)^2 \right\rangle = \frac{1}{V} \int_0^\infty (\gamma G)^2 dV \tag{A.14}$$

Further evaluation requires knowledge of the gradient distribution, and hence the system geometry. As derived in Appendix A.1, the field fluctuations around a sphere of radius R and with susceptibility difference $\Delta\chi$ from the background are given by

$$B(r, \phi, \theta) = \delta B \frac{R^3}{r^3}(3\cos^2\theta - 1) \tag{A.15}$$

with the equatorial field

$$\delta B = \frac{4\pi}{3}\Delta\chi B_0 \tag{A.16}$$

We require the gradient of this expression, and since there is no dependence on ϕ we can use the form

$$\begin{aligned} \nabla B &= \frac{\partial B}{\partial r}\hat{r} + \frac{1}{r}\frac{\partial B}{\partial \theta}\hat{\theta} \\ &= -3\delta B \frac{R^3}{r^4}(3\cos^2\theta - 1)\hat{r} - 6\delta B \frac{R^3}{r^4}\cos\theta\sin\theta\hat{\theta} \\ &= -3\delta B \frac{R^3}{r^4}((3\cos^2\theta - 1)\hat{r} + 2\cos\theta\sin\theta\hat{\theta}) \end{aligned} \tag{A.17}$$

The quantity we need to integrate in the expression (A.11) above is

$$\begin{aligned} (\gamma G)^2 &= (\gamma \nabla B)^2 \\ &= 9\delta\omega^2 \frac{R^6}{r^8}((3\cos^2\theta - 1)^2 + (2\cos\theta\sin\theta)^2) \\ &= 9\delta\omega^2 \frac{R^6}{r^8}(5\cos^4\theta - 2\cos^2\theta + 1) \end{aligned} \tag{A.18}$$

with $\delta\omega = \gamma\delta B$ the characteristic frequency offset.

When performing the volume integral (A.14), the volume to be used is

$$V = \frac{4\pi R^3}{3\zeta} \tag{A.19}$$

and the expression can be evaluated as

$$\begin{aligned}\left\langle (\gamma G)^2 \right\rangle &= \frac{1}{V}\int_R^\infty 9\delta\omega^2 \frac{R^6}{r^8} r^2 dr \int_0^{2\pi} d\phi \int_0^\pi (5\cos^4\theta - 2\cos^2\theta + 1)\sin\theta d\theta \\ &= \frac{18\pi\delta\omega^2}{V}\int_R^\infty \frac{R^6}{r^6} dr \int_{-1}^1 (5x^4 - 2x^2 + 1) dx \\ &= \frac{18\pi R\delta\omega^2}{5V}\cdot\frac{8}{3} \\ &= \frac{36\zeta\delta\omega^2}{5R^2} \end{aligned} \tag{A.20}$$

We can now substitute this expression into A.11 to obtain

$$\begin{aligned}\exp\left[-\frac{1}{12}D\cdot\left\langle (\gamma G)^2\right\rangle TE^3\right] &= \exp\left[-\frac{3D\zeta\delta\omega^2}{5R^2}TE^3\right] \\ &= \exp\left[-\frac{1}{10}\frac{6D}{R^2}\zeta\delta\omega^2 TE^3\right] \\ &= \exp\left[-\frac{27}{40\pi^2}\zeta\frac{TE^3}{t_c^2 t_D}\right]\end{aligned}$$

Here we have used the definitions

$$\begin{aligned} t_D &= \frac{R^2}{6D} \\ t_c &= \left(\frac{8\pi^2}{9\sqrt{3}}\gamma\Delta\chi B_0\right)^{-1} = \left(\frac{2\pi}{3\sqrt{3}}\delta\omega\right)^{-1}\end{aligned}$$

From this we obtain the static dephasing condition

$$\begin{aligned}\frac{27}{40\pi^2}\zeta\frac{TE^3}{t_c^2 t_D} &<< 1 \\ TE^3 &<< \frac{40\pi^2}{27}\frac{t_c^2 t_D}{\zeta} \\ TE &<< 2.4(T_2' t_c t_D)^{1/3}\end{aligned} \tag{A.21}$$

with

$$T_2' = 1/R_2' = \frac{t_c}{\zeta}$$

A.3 Propagator Solution of the Bloch-Torrey Equation

Though numerous methods can be used to solve a differential equation such as the Bloch-Torrey equation (eqn. 5.11), the model originally proposed by Kiselev and

Posse [65] uses a method based on Green's functions or *propagators*, which allows solution of the differential equation once the boundary conditions are known.

In general, for any differential operator $L(x)$ there is a fundamental solution $G(x, x_0)$ defined by

$$LG(x, x_0) = \delta(x - x_0) \tag{A.22}$$

The letter G is used here since the fundamental solution is a Green's function for the differential equation. Its usefulness stems from the fact that a solution u(x) to the more general equation

$$Lu(x) = f(x) \tag{A.23}$$

can be determined by convolving the (known) function f with the fundamental solution, as

$$u(x) = \int G(x, x_0) f(x_0) dx \tag{A.24}$$

In our particular case the differential operator L corresponding to the (one-dimensional) Bloch-Torrey equation contains derivatives in both space and time,

$$L = \frac{\partial}{\partial t} - D\frac{\partial^2}{\partial x^2} - igx \tag{A.25}$$

and the resulting propagator will also have spatial and temporal dependence, and is based on the initial condition

$$LG(x, x_0, t) = \delta(t)\delta(x - x_0) \tag{A.26}$$

The form of G, which can be verified by substitution into the Bloch-Torrey equation, is [140]

$$G(x, x_0, t) = \frac{1}{(2\pi Dt)^{1/2}} \exp\left[-\frac{(x - x_0)^2}{4Dt} - igt\frac{x - x_0}{2} - \frac{1}{12}Dg^2t^3\right] e^{-igx_0t} \tag{A.27}$$

This solution is valid under the linear local field approximation, which assumes that the magnetic field variation around a given point is only due to a linear background gradient g. The desired solution Ψ to the Bloch-Torrey equation can now be found after convolution with the appropriate initial conditions, which will be described for the FID and spin-echo cases.

A.3.1 Propagator Solution for FID Experiment

In the case of an FID experiment, we need only to convolve the propagator with a delta function representing the initial position of a spin packet:

$$\Psi(t) = \int\int G(x_2, x_1, t)\delta(x_1 - x_0)dx_1 dx_2 \tag{A.28}$$

The delta function simply evaluates the inner integral at $x_1 = x_0$, and the remaining integration is performed by changing variables to $\Delta x = x_2 - x_0$, completing the square and integrating the Gaussian by inspection.

$$\begin{aligned}
\Psi &= \int_{-\infty}^{\infty} \frac{1}{(2\pi Dt)^{1/2}} \exp\left[-\frac{(x_2 - x_0)^2}{4Dt} - igt\frac{x_2 - x_0}{2} - \frac{1}{12}Dg^2t^3\right] e^{-igx_0t}dx_2 \\
&= \exp\left[-\frac{1}{12}Dg^2t^3 - igx_0t\right] \int_{-\infty}^{\infty} \frac{1}{(2\pi Dt)^{1/2}} \exp\left[-\frac{(\Delta x)^2}{4Dt} - i\frac{gt}{2}\Delta x\right] d\Delta x \\
&= \exp\left[-\frac{1}{12}Dg^2t^3 - Dt(\frac{gt}{2})^2 - igx_0t\right] \int_{-\infty}^{\infty} \frac{1}{(2\pi Dt)^{1/2}} \exp\left[-\frac{(\Delta x)^2}{4Dt} - i(\frac{gt}{2})\Delta x + Dt(\frac{gt}{2})^2\right] d\Delta x \\
&= \exp\left[-\frac{1}{12}Dg^2t^3 - \frac{1}{4}Dg^2t^3 - igx_0t\right] \int_{-\infty}^{\infty} \frac{1}{(2\pi Dt)^{1/2}} \exp\left[-\left(\frac{\Delta x}{2\sqrt{Dt}} + i\sqrt{Dt}(\frac{gt}{2})\right)^2\right] d\Delta x \\
&= \exp\left[-\frac{1}{3}Dg^2t^3 - igx_0t\right] \int_{-\infty}^{\infty} \frac{1}{(2\pi Dt)^{1/2}} \exp\left[-\frac{\left(\Delta x + 2iDt(\frac{gt}{2})\right)^2}{2\sqrt{Dt}}\right] d\Delta x \\
&= \exp\left[-\frac{1}{3}Dg^2t^3 - igx_0t\right]
\end{aligned} \tag{A.29}$$

The t^3 dependence of the attenuation function is characteristic of diffusion through a field gradient [105], and can also be derived by direct integration of the Bloch-Torrey equation. Deriving this result through the propagator framework, however, allows the approach to be verified before attempting the more relevant case of the spin-echo solution.

A.3.2 Propagator Solution for Spin-Echo Experiment

For the case of the spin echo with echo time TE, the solution to the Bloch-Torrey equation becomes

$$\Psi(t) = \int\int G(x_2, x_1, t - TE/2)G^*(x_1, x_0, TE/2)dx_1 dx_2 \tag{A.30}$$

This equation is valid only for times t after the refocusing pulse; before then, the FID solution applies. After $t = TE/2$, the new initial condition is simply the propagator prior to the refocusing pulse, with complex conjugation used to represent the reversal of the magnetization phase.

This convolution will be made easier in the Fourier representation. We therefore require the Fourier transform of G, which involves a similar integral to the one performed for the FID case, but now with a $e^{-ik\Delta x}$ term. Ignoring the oscillatory e^{-igx_0t} term since it does not depend on Δx, we find

$$\begin{aligned}
H(k,t) &= \int_{-\infty}^{\infty} \frac{d\Delta x}{\sqrt{2\pi Dt}} \exp\left[-\frac{(\Delta x)^2}{4Dt} - igt\frac{\Delta x}{2} - \frac{1}{12}Dg^2t^3\right] e^{-ik\Delta x} \\
&= \exp\left[-\frac{Dg^2t^3}{12}\right] \int_{-\infty}^{\infty} \frac{d\Delta x}{\sqrt{2\pi Dt}} \exp\left[-\frac{(\Delta x)^2}{4Dt} - i(\frac{gt}{2}+k)\Delta x\right] \\
&= \exp\left[-\frac{Dg^2t^3}{12} - Dt(\frac{gt}{2}+k)^2\right] \int_{-\infty}^{\infty} \frac{d\Delta x}{\sqrt{2\pi Dt}} \exp\left[-\frac{(\Delta x)^2}{4Dt} - i(\frac{gt}{2}+k)\Delta x + Dt(\frac{gt}{2}+k)^2\right] \\
&= \exp\left[-\frac{Dg^2t^3}{12} - Dt(\frac{gt}{2}+k)^2\right] \int_{-\infty}^{\infty} \frac{d\Delta x}{\sqrt{2\pi Dt}} \exp\left[-\left(\frac{\Delta x}{2\sqrt{Dt}} + i\sqrt{Dt}(\frac{gt}{2}+k)\right)^2\right] \\
&= \exp\left[-\frac{Dg^2t^3}{12} - Dt(\frac{gt}{2}+k)^2\right] \int_{-\infty}^{\infty} \frac{d\Delta x}{\sqrt{2\pi Dt}} \exp\left[-\frac{\left(\Delta x + 2iDt(\frac{gt}{2}+k)\right)^2}{2\sqrt{Dt}}\right] \\
&= \exp\left[-\frac{Dg^2t^3}{12} - Dt(\frac{gt}{2}+k)^2\right] \qquad \text{(A.31)}
\end{aligned}$$

The Fourier transform of the complex conjugate G^* can be found with a similar calculation to be

$$H^*(k,t) = \exp\left[-\frac{1}{12}Dg^2t^3 - Dt(\frac{gt}{2}-k)^2\right] \qquad \text{(A.32)}$$

Note the change in the sign of k, and that the oscillatory term will change sign as well. The desired solution Ψ can now be computed as

$$\begin{aligned}
\Psi(t) &= \int dk_{21} e^{ik_{21}(x_2-x_1)} \int dk_{10} e^{ik_{10}(x_1-x_0)} \int dx_2 \int dx_1 \Big(H(k_{21}, t-TE/2) e^{-igx_1(t-TE/2)} \cdot \\
&\quad H^*(k_{10}, TE/2) e^{igx_0 TE/2} \Big) \\
&= \int dk_{21} \int dk_{10} \int dx_2 e^{ix_2k_{21}} \int dx_1 e^{ix_1(k_{10}-k_{21}-g(t-TE/2))} e^{ix_0(-k_{10}+gTE/2)} \cdot \\
&\quad H(k_{21}, t-TE/2) H^*(k_{10}, TE/2)
\end{aligned}$$

After collecting terms involving x_1 and x_2, we use the definition of the delta-function

$$\delta(k) = \int dx e^{-ixk} \qquad \text{(A.33)}$$

to further simplify this expression to

$$\begin{aligned}
\Psi(t) &= \int dk_{21} dk_{10} H(k_{21}, t-TE/2) H^*(k_{10}, TE/2) \delta(k_{21}) \delta(k_{10}-k_{21}-g(t-TE/2)) e^{ix_0(-k_{10}+gTE/2)} \\
&= H(0, t-TE/2) \int dk_{10} H^*(k_{10}, TE/2) \delta(k_{10}-g(t-TE/2)) e^{ix_0(-k_{10}+gTE/2)} \\
&= H(0, t-TE/2) H^*(g(t-TE/2), TE/2) e^{ix_0(-g(t-TE/2)+gTE/2)} \qquad \text{(A.34)}
\end{aligned}$$

Into this expression we substitute

$$\begin{aligned} H(0, t - TE/2) &= \exp\left[-\frac{1}{12}Dg^2(t - TE/2)^3 - Dt(\frac{g(t - TE/2)}{2})^2\right] \\ &= \exp\left[-\frac{1}{3}Dg^2(t - TE/2)^3\right] \\ &= \exp\left[-Dg^2\left(\frac{1}{3}t^3 - \frac{t^2TE}{2} + \frac{tTE^2}{4} - \frac{TE^3}{24}\right)\right] \end{aligned} \tag{A.35}$$

and

$$\begin{aligned} H^*(g(t - TE/2), TE/2) &= \exp\left[-\frac{1}{12}Dg^2(TE/2)^3 - D(TE/2)(\frac{gTE}{4} - g(t - TE/2))^2\right] \\ &= \exp\left[-Dg^2\left(\frac{28TE^3}{96} - \frac{3tTE^2}{4} + \frac{t^2TE}{2}\right)\right] \end{aligned} \tag{A.36}$$

and finally we obtain

$$\begin{aligned} \Psi(t) &= \exp\left[-Dg^2\left(\frac{1}{3}t^3 - \frac{tTE^2}{2} + \frac{TE^3}{4}\right)\right] e^{-ix_0g(t-TE)} \\ &= \exp\left[-\frac{1}{3}Dg^2t^3\left(1 - \frac{3TE^2}{2t^2} + \frac{3TE^3}{4t^3}\right)\right] e^{-ix_0g(t-TE)} \end{aligned} \tag{A.37}$$

This matches the expression for the FID case, with a slight modification of the oscillatory term such that the phase now depends on $(t - TE)$ instead of t, as well as the addition of a factor

$$F(TE/t) = 1 - \frac{3}{2}\left(\frac{TE}{t}\right)^2 + \frac{3}{4}\left(\frac{TE}{t}\right)^3 \tag{A.38}$$

Bibliography

[1] Timm Schroeder. Imaging stem-cell-driven regeneration in mammals. *Nature*, 453(15):345–351, 2008.

[2] Jeff W.M. Bulte and Dara L. Kraitchman. Iron oxide MR contrast agents for molecular and cellular imaging. *NMR in Biomedicine*, 17:484–499, 2004.

[3] Moritz F. Kircher, Jennifer R. Allport, Edward E. Graves, Victoria Love, Lee Josephson, Andrew H. Lichtman, and Ralph Weissleder. In vivo high resolution three-dimensional imaging of antigen-specific cytotoxic T-lymphocyte trafficking to tumors. *Cancer Research*, 63:6838–6846, 2003.

[4] Chris Heyn, John A. Ronald, Lisa T. Mackenzie, Ian C. MacDonald, Ann F. Chambers, Brian K. Rutt, and Paula J. Foster. In vivo magnetic resonance imaging of single cells in mouse brain with optical validation. *Magnetic Resonance in Medicine*, 55:23–29, 2006.

[5] Chris V. Bowen, Xiaowei Zhang, George Saab, Paula J. Gareau, and Brian K. Rutt. Application of the static dephasing regime theory to superparamagnetic iron-oxide loaded cells. *Magnetic Resonance in Medicine*, 48:52–61, 2002.

[6] P. C. Lauterbur, M. L. Bernardo, Jr., M. H. Mendonca Dias, and A. W. Heldman. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 229, 1986.

[7] Joseph A. Frank, Brad R. Miller, Ali S. Arbab, Holly A. Zywicke, E. Kay Jordan, Bobbi K. Lewis, L. Henry Bryant, Jr, and Jeff W. M. Bulte. Clinically applicable labeling of mammalian and stem cells by combining superparamagnetic iron oxides and transfection agents. *Radiology*, 228(2):480–487, 2003.

[8] Daniel L. J. Thorek, Antony K. Chen, Julie Czupryna, and Andrew Tsourkas. Superparamagnetic iron oxide nanoparticle probes for molecular imaging. *Annals of Biomedical Engineering*, 34(1):23–28, 2006.

[9] Claire Corot, Philippe Robert, Jean-Marc Ide, and Marc Port. Recent advances in iron oxide nanocrystal technology for medical imaging. *Advanced Drug Delivery Reviews*, 58:1471–1504, 2006.

[10] Erik M. Shapiro, Stanko Skrtic, and Alan P. Koretsky. Sizing it up: Cellular MRI using micron-sized iron oxide particles. *Magnetic Resonance in Medicine*, 53:329–338, 2005.

[11] Michel Modo, Mathias Hoehn, and Jeff W.M. Bulte. Cellular MR imaging. *Molecular Imaging*, 4(3):143–164, 2005.

[12] Daniel L. J. Thorek and Andrew Tsourkas. Size, charge and concentration dependent uptake of iron oxide particles by non-phagocytic cells. *Biomaterials*, 29:3583–3590, 2008.

[13] Mat Lewin, Nadia Carlesso, Ching-Hsuan Tung, Xiao-Wu Tang, David Cory, David T. Scadden, and Ralph Weissleder. Tat peptide-derivatized magnetic nanoparticles allow in vivo tracking and recovery of progenitor cells. *Nature Biotechnology*, 18:410–414, 2000.

[14] Ralph Weissleder, Kimberly Kelly, Eric Yi Sun, Timur Shtatland, and Lee Josephson. Cell-specific targeting of nanoparticles by multivalent attachment of small molecules. *Nature Biotechnology*, 23(11):1418–1423, 2005.

[15] E. M. Shapiro, S. Skrtic, and A. P. Koretsky. Long term cellular MR imaging using micron sized iron oxide particles. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 11, page 166, 2004.

[16] Heike E. Daldrup-Link, Martina Rudelius, Robert A. J. Oostendorp, Marcus Settles, Guido Piontek, Stefan Metz, Hilkea Rosenbrock, Ulrich Keller, Ulrich Heinzmann, Ernst J. Rummeny, Jurgen Schlegel, and Thomas M. Link. Targeting of hematopoietic progenitor cells with MR contrast agents. *Radiology*, 228:760–767, 2003.

[17] Scott D. Swanson, Jolanta F. Kukowska-Latallo, Anil K. Patri, Chunyan Chen, Song Ge, Zhengyi Cao, Alina Kotlyar, Andrea T. East, , and James R. Baker. Targeted gadolinium-loaded dendrimer nanoparticles for tumor-specific magnetic resonance contrast enhancement. *International Journal of Nanomedicine*, 3:201–210, 2008.

[18] Luc Faucher, Andre-Anne Guay-Bgin, Jean Lagueux, Marie-France Ct, ric Petitclerc, and Marc-Andr Fortin. Ultra-small gadolinium oxide nanoparticles to image brain cancer cells in vivo with MRI. *Contrast Media and Molecular Imaging*, 6:209–218, 2011.

[19] Luisa Ottobrini, Cristina Martelli, Daria Lucia Trabattoni, Mario Clerici, and Giovanni Lucignani. In vivo imaging of immune cell trafficking in cancer. *European Journal of Nuclear Medicine and Molecular Imaging*, 38:949–968, 2011.

[20] Stephen J. Dodd, Mangay Williams, Joseph P. Suhan, Donald S. Williams, Alan P. Koretsky, and Chien Ho. Detection of single mammalian cells by high-resolution magnetic resonance imaging. *Biophysical Journal*, 76:103–109, 1999.

[21] Erik M. Shapiro, Kathryn Sharer, Stanko Skrtic, and Alan P. Koretsky. In vivo detection of single cells by MRI. *Magnetic Resonance in Medicine*, 55:242–249, 2006.

[22] R. Marc Lebel, Ravi S. Menon, and Chris V. Bowen. Relaxometry model of strong dipolar perturbers for Balanced-SSFP: Application to quantification of SPIO loaded cells. *Magnetic Resonance in Medicine*, 55:583–591, 2006.

[23] Charles H. Cunningham, Takayasu Arai, Phillip C. Yang, Michael V. McConnell, John M. Pauly, and Steven M. Conolly. Positive contrast magnetic resonance imaging of cells labeled with magnetic nanoparticles. *Magnetic Resonance in Medicine*, 53:999–1005, 2005.

[24] Matthias Stuber, Wesley D. Gilson, Michael Schar, Dorota A. Kedziorek, Lawrence V. Hofmann, Saurabh Shah, Evert-Jan Vonken, Jeff W.M. Bulte, and Dara L. Kraitchman. Positive contrast visualization of iron oxide-labeled stem cells using inversion-recovery with on-resonant water suppression (IRON). *Magnetic Resonance in Medicine*, 58:1072–1077, 2007.

[25] Venkatesh Mani, Karen C. Briley-Saebo, Vitalii V. Itskovich, Daniel D. Samber, and Zahi A. Fayad. GRadient echo acquisition for superparamagnetic particles with positive contrast (GRASP): Sequence characterization in membrane and glass superparamagnetic iron oxide phantoms at 1.5T and 3T. *Magnetic Resonance in Medicine*, 55:126–135, 2006.

[26] Hannes Dahnke, Wei Liu, Daniel Herzka, Joseph A. Frank, and Tobias Schaeffter. Susceptibility gradient mapping (SGM): A new postprocessing method for positive contrast generation applied to superparamagnetic iron oxide particle (SPIO)-labeled cells. *Magnetic Resonance in Medicine*, 60:595–603, 2008.

[27] Qun Zhao, Jason Langley, Sunbok Lee, and Wei Liu. Positive contrast technique for the detection and quantification of superparamagnetic iron oxide nanoparticles in MRI. *NMR in Biomedicine*, 24:464–472, 2011.

[28] Garry E. Gold, John M. Pauly, Albert Macovski, and Robert J. Herfkens. MR spectroscopic imaging of collagen: Tendons and knee menisci. *Magnetic Resonance in Medicine*, 34:647–654, 1995.

[29] L. A. Crowe, F. Ris, S. Nielles-Vallespin, P. Speier, S. Masson, M. Armanet, P. Morel, C. Toso, D. Bosco, T. Berney, and J.-P. Vallee. A novel method for quantitative monitoring of transplanted islets of langerhans by positive contrast magnetic resonance imaging. *American Journal of Transplantation*, 11:1158–1168, 2011.

[30] Eric T. Ahrens, Rafael Flores, Hongyan Xu, and Penelope A. Morel. In vivo imaging platform for tracking immunotherapeutic cells. *Nature Biotechnology*, 23:983–987, 2005.

[31] Ron McKay. Stem cells: Hype and hope. *Nature*, 406:361–364, 2000.

[32] Stacey M. Cromer Berman, Piotr Walczak, and Jeff W. M. Bulte. Tracing stem cells using magnetic nanoparticles. *WIREs Nanomedicine and Nanobiotechnology*, 3:343–355, 2011.

[33] J. W. M. Bulte, S.-C. Zhang, P. van Gelderen, V. Herynek, E. K. Jordan, I. D. Duncan, and J. A. Frank. Neurotransplantation of magnetically labeled oligodendrocyte progenitors: Magnetic resonance tracking of cell migration and myelination. *Proceedings of the National Academy of Sciences*, 96(26):15256–15261, 1999.

[34] Daehong Kim, Byoung gi Chun, Yeon-Kyung Kim, Yong Hyun Lee, Cheong-Soo Park, Iksoo Jeon, Chaejoon Cheong, Tae-Sun Hwang, Hyungmin Chung, Byoung Joo Gwag, Kwan Soo Hong, and Jihwan Song. In vivo tracking of human mesenchymal stem cells in experimental stroke. *Cell Transplantation*, 16:1007–1012, 2007.

[35] Jonathan M. Hill, Alexander J. Dick, Venkatesh K. Raman, Richard B. Thompson, Zu-Xi Yu, K. Allison Hinds, Breno S.S. Pessanha, Michael A. Guttman, Timothy R. Varney, Bradley J. Martin, Cynthia E. Dunbar, Elliot R. McVeigh, and Robert J. Lederman. Serial cardiac magnetic resonance imaging of injected mesenchymal stem cells. *Circulation*, 108:1009–1014, 2003.

[36] Dara L. Kraitchman, Alan W. Heldman, Ergin Atalar, Luciano C. Amado, Bradley Martin, Mark F. Pittenger, Joshua M. Hare, and Jeff W.M. Bulte. In vivo magnetic resonance imaging of mesenchymal stem cells in myocardial infarction. *Circulation*, 107:2290–2293, 2003.

[37] Venkatesh Mani, Eric Adler, Karen C. Briley-Saebo, Anne Bystrup, Valentin Fuster, Gorden Keller, and Zahi A. Fayad. Serial in vivo positive contrast MRI of iron oxide-labeled embryonic stem cell-derived cardiac precursor cells in a mouse model of myocardial infarction. *Magnetic Resonance in Medicine*, 60: 73–81, 2008.

[38] Clemens Bos, Yahsou Delmas, Alexis Desmouliere, Anne Solanilla, Olivier Hauger, Christophe Grosset, Isabelle Dubus, Zoran Ivanovic, Jean Rosenbaum, Pierre Charbord, Christian Combe, Jeff W. M. Bulte, Chrit T. W. Moonen, Jean Ripoche, and Nicolas Grenier. In vivo MR imaging of intravascularly injected magnetically labeled mesenchymal stem cells in rat kidney and liver. *Radiology*, 233:781–789, 2004.

[39] Suzanne L. Topalian, George J. Weiner, and Drew M. Pardoll. Cancer immunotherapy comes of age. *Journal of Clinical Oncology*, 29:4828–4836, 2011.

[40] Ira Mellman, George Coukos, and Glenn Dranoff. Cancer immunotherapy comes of age. *Nature*, 480:480–489, 2011.

[41] Mohan Karkada, Genevieve M. Weir, Tara Quinton, Leeladhar Sammatur, Lisa D. MacDonald, Alecia Grant, Robert Liwski, Ridas Juskevicius, Gomathinayagam Sinnathamby, Ramila Philip, and Marc Mansour. A novel breast/ovarian cancer peptide vaccine platform that promotes specific type-1 but not treg/tr1-type responses. *Journal of Immunotherapy*, 33:250–261, 2010.

[42] W.Joost Lesterhuis, John B.A.G.Haanen, and Cornelis J.A.Punt. Cancer immunotherapy - revisited. *Nature Reviews Drug Discovery*, 10:591–600, 2011.

[43] Dirk Baumjohann, Andreas Hess, Lubos Budinsky, Kay Brune, Gerold Schuler, and Manfred B. Lutz. In vivo magnetic resonance imaging of dendritic cell migration into the draining lymph nodes of mice. *European Journal of Immunology*, 36:2544–2555, 2006.

[44] Roja Rohani, Sonali N. de Chickera, Christy Willert, Yuhua Chen, Gregory A. Dekaban, and Paula J. Foster. In vivo cellular MRI of dendritic cell migration using micrometer-sized iron oxide (MPIO) particles. *Molecular Imaging and Biology*, 13:679–694, 2011.

[45] Richard Tavare, Pervinder Sagoo, Gopal Varama, Yakup Tanriver, Alice Warely, Sandra S. Diebold, Richard Southworth, Tobias Schaeffter, Robert I. Lechler, Reza Razavi, Giovanna Lombardi, and Gregory E. D. Mullen. Monitoring of in vivo function of superparamagnetic iron oxide labelled murine dendritic cells during anti-tumour vaccination. *PLoS ONE*, 6:e19662, 2011.

[46] I. Jolanda M. de Vries, W. Joost Lesterhuis, Jelle O. Barentsz, Pauline Verdijk, J. Han van Krieken, Otto C. Boerman, Wim J. G. Oyen, Johannes J. Bonenkamp, Jan. B Boezeman, Gosse J. Adema, Jeff W. M. Bulte, Tom W. J. Scheenen, Cornelis J. A. Punt, Arend Heerschap, and Carl G. Figdor. Magnetic resonance tracking of dendritic cells in melanoma patients for monitoring of cellular therapy. *Nature Biotechnology*, 23:1407–1413, 2005.

[47] Li Liu, Qing Ye, Yijen Wu, Wen-Yuan Hsieh, Chih-Lung Chen, Hsin-Hsin Shen, Shian-Jy Wang, Haosen Zhang, T. Kevin Hitchens, and Chien Ho. Tracking t-cells in vivo with a new nano-sized MRI contrast agent. *Nanomedicine: Nanotechnology, Biology and Medicine*, 2012.

[48] Mangala Srinivas, Michael S. Turner, Jelena M. Janjic, Penelope A. Morel, David H. Laidlaw, and Eric T. Ahrens. In vivo cytometry of antigen-specific t cells using 19F MRI. *Magnetic Resonance in Medicine*, 62:747–753, 2009.

[49] Sebastien Serres, Silvy Mardiguian, Sandra J. Campbell, Martina A. McAteer, Asim Akhtar, Alexandre Krapitchev, Robin P. Choudhury, Daniel C. Anthony, and Nicola R. Sibson. VCAM-1-targeted magnetic resonance imaging reveals subclinical disease in a mouse model of multiple sclerosis. *The FASEB Journal*, 25:4415–4422, 2011.

[50] Asim M. Akhtar, Jurgen E. Schneider, Stephanie J. Chapman, Andrew Jefferson, Janet E. Digby, Kulveer Mankia, Ye Chen, Martina A. McAteer, Kathryn J. Wood, and Robin P. Choudhury. In vivo quantification of VCAM-1 expression in renal ischemia reperfusion injury using non-invasive magnetic resonance molecular imaging. *PLoS ONE*, 5:e12800, 2010.

[51] Jerry S. Cheung, April M. Chow, Edward S. Hui, Jian Yang, Hung-Fat Tse, and Ed X. Wu. Cell number quantification of USPIO-labeled stem cells by MRI: An in vitro study. In *Proceedings of the 28th IEEE EMBS Annual International Conference*, pages 476–479, September 2006.

[52] Kim Brewer, Kerry Lake, Iulia Dude, Nicole Pelot, Drew Debay, Andrea Penwell, Genevieve Weir, Chris Bowen, Mohan Karkada, and Marc Mansour. Visualizing cancer vaccine clearance in vivo using magnetic resonance imaging. In *Proceedings of the American Association for Cancer Research*, 2012.

[53] Jean-Christophe Brisset, Monica Sigovan, Fabien Chauveau, Adrien Riou, Emilie Devillard, Virginie Desestret, Monique Touret, Serge Nataf, J. Honnorat, Emmanuelle Canet-Soulas, Norbert Nighoghossian, Yves Berthezene, and Marlene Wiart. Quantification of iron-labeled cells with positive contrast in mouse brains. *Molecular Imaging and Biology*, 13:672–678, 2011.

[54] Longjiang Zhang, Xiaodong Zhong, Liya Wang, Hongwei Chen, Y. Andrew Wang, Julie Yeh, Lily Yang, and Hui Mao. T1-weighted ultrashort echo time method for positive contrast imaging of magnetic nanoparticles and cancer cells bound with the targeted nanoparticles. *Journal of Magnetic Resonance Imaging*, 33:194–202, 2011.

[55] Rebecca Kuhlpeter, Hannes Dahnke, Lars Matuszewski, Thorsten Persigehl, Angelika von Wallbrunn, Thomas Allkemper, Walter L. Heindel, Tobias Schaeffter, and Christoph Bremer. R2 and R2* mapping for sensing cell-bound superparamagnetic nanoparticles: In vitro and murine in vivo testing. *Radiology*, 245(2):449–457, 2007.

[56] Gyula Kotek, Sandra T. van Tiel, Piotr A. Wielopolski, Gavin C. Houston, Gabriel P. Krestin, and Monique R. Bernsen. Cell quantification: evolution of compartmentalization and distribution of iron-oxide particles and labeled cells. *Contrast Media and Molecular Imaging*, 7:195–203, 2012.

[57] Gerben M. van Buul, Gyula Kotek, Piotr A. Wielopolski, Eric Farrell, P. Koen Bos, Harrie Weinans, Anja U. Grohnert, Holger Jahr, Jan A. N. Verhaar, Gabriel P. Krestin, Gerjo J. V. M. van Osch, and Monique R. Bernsen. Clinically translatable cell tracking and quantification by MRI in cartilage repair using superparamagnetic iron oxides. *PLoS ONE*, 6:e17001, 2011.

[58] Jingfei Ma and Felix W. Wehrli. Method for image-based measurement of the reversible and irreversible contribution to the transverse-relaxation rate. *Journal of Magnetic Resonance*, 111:61–69, 1996.

[59] Dmitriy A. Yablonskiy and E. Mark Haacke. An MRI method for measuring T_2 in the presence of static and RF magnetic field inhomogeneities. *Magnetic Resonance in Medicine*, 37:872–876, 1997.

[60] Ahmed Elkady. MRI cellular density quantification using balanced steady state free precession. Master's thesis, Dalhousie University, Halifax, Nova Scotia, 2009.

[61] Wei Liu, Hannes Dahnke, Juergen Rahmer, E. Kay Jordan, and Joseph A. Frank. Ultrashort T_2^* relaxometry for quantitation of highly concentrated superparamagnetic iron oxide (SPIO) nanoparticle labeled cells. *Magnetic Resonance in Medicine*, 61:761–766, 2009.

[62] Peter R. Seevinck, Jan-Henry Seppenwoolde, Jaco J.M. Zwanenburg, Johannes F.W. Nijsen, and Chris J.G. Bakker. FID sampling superior to spin-echo sampling for T_2^*-based quantification of holmium-loaded microspheres: Theory and experiment. *Magnetic Resonance in Medicine*, 60:1466–1476, 2008.

[63] Steven D. Beyea, Bruce J. Balcom, Igor V. Mastikhin, Theodore W. Bremner, Robin L. Armstrong, and Patrick E. Grattan-Bellew. Imaging of heterogeneous materials with a turbo spin echo single-point imaging technique. *Journal of Magnetic Resonance*, 144:255–265, 2000.

[64] Chris V. Bowen, Joseph S. Gati, and Ravi S. Menon. Robust prescan calibration for multiple spin-echo sequences: application to FSE and b-SSFP. *Magnetic Resonance Imaging*, 24:857–867, 2006.

[65] V.G. Kiselev and S. Posse. Analytical model of susceptibility-induced MR signal dephasing: Effect of diffusion in a microvascular network. *Magnetic Resonance in Medicine*, 41:499–509, 1999.

[66] Namrata Vaswani and Wei Lu. Modified-CS: Modifying compressive sensing for problems with partially known support. *IEEE Transactions on Signal Processing*, 58:4595–4607, 2010.

[67] James A. Rioux, Kimberly D. Brewer, Steven D. Beyea, and Chris V. Bowen. Quantification of superparamagnetic iron oxide with large dynamic range using TurboSPI. *Journal of Magnetic Resonance*, 216, 2012.

[68] David J. Griffiths. *Introduction to Electrodynamics, Third Edition.* Prentice Hall, 1999.

[69] William Fuller Brown. Thermal fluctuations of a single-domain particle. *Physical Review*, 130:1677–1686, 1963.

[70] Felix Bloch. Nuclear induction. *Physical Review*, 70:460–474, 1946.

[71] E.M. Purcell, H.C. Torrey, and R.V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69:37–38, 1946.

[72] J.J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, 1994.

[73] H.Y. Carr and E.M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Physical Review*, 94:630–638, 1954.

[74] Saul Meiboom and David Gill. Modified spin-echo method for measuring nuclear relaxation times. *Review of Scientific Instruments*, 29:688–691, 1958.

[75] E. Mark Haacke, Robert W. Brown, Michael R. Thompson, and Ramesh Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. John Wiley and Sons, 1999.

[76] Jurgen Hennig. Echoes - how to generate, recognize, use or avoid them in MR-imaging sequences part II: Echoes in imaging sequences. *Concepts in Magnetic Resonance*, 3:179–192, 1991.

[77] Peter Mansfield and P.K. Grannell. NMR 'diffraction' in solids? *Journal of Physics C*, 6:L422–L426, 1973.

[78] John Pauly, Patrick LeRoux, Dwight Nishimura, and Albert Macovski. Parameter relations for the shinnar-le roux selective excitation pulse design algorithm. *IEEE Transactions on Medical Imaging*, 10(1):53–65, 1991.

[79] Douglas C. Noll, Dwight G. Nishimura, and Albert Macovski. Homodyne detection in magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 10(2):154–163, 1991.

[80] S. Emid and J.H.N. Creyghton. High resolution NMR imaging in solids. *Physica*, 128B:81–83, 1985.

[81] Samuel Gravina and D.G. Cory. Sensitivity and resolution of constant-time imaging. *Journal of Magnetic Resonance*, 104:53–61, 1994.

[82] Steven D. Beyea, Bruce J. Balcom, Theodore W. Bremner, Pablo J. Prado, Albert R. Cross, Robin L. Armstrong, and Paul E. Grattan-Bellew. The influence of shrinkage-cracking on the drying behaviour of white portland cement using single-point imaging (SPI). *Solid State Nuclear Magnetic Resonance*, 13: 93–100, 1998.

[83] Pablo J. Prado, Bruce J. Balcom, Igor V. Mastikhin, Albert R. Cross, Chris B. Kennedy, Robin L. Armstrong, and A. Logan. Magnetic resonance imaging of gases. a SPRITE study. *Journal of Magnetic Resonance*, 137:342–352, 1999.

[84] Alexei V. Ouriadov, Rodney P. MacGregor, and Bruce J. Balcom. Thin film MRIhigh resolution depth imaging with a local surface coil and spin echo SPI. *Journal of Magnetic Resonance*, 169:174–186, 2004.

[85] Linqing Li, Hui Han, and Bruce J. Balcom. Spin echo SPI methods for quantitative analysis of fluids in porous media. *198*, 2009:252–260, Journal of Magnetic Resonance.

[86] J. Hennig, A. Nauerth, and H. Friedburg. Rare imaging: A fast imaging method for clinical MR. *Magnetic Resonance in Medicine*, 3:823–833, 1986.

[87] Colin S. Poon and R. Mark Henkelman. Practical T2 quantitation for clinical applications. *Journal of Magnetic Resonance Imaging*, 2:541–553, 1992.

[88] Robert M. Weisskoff, Chun S. Zuo, Jerrold L. Boxerman, and Bruce R. Rosen. Microscopic susceptibility variation and transverse relaxation: Theory and experiment. *Magnetic Resonance in Medicine*, 31:601–610, 1994.

[89] C. H. Ziener, W. R. Bauer, and P. M. Jakob. Transverse relaxation of cells labeled with magnetic nanoparticles. *Magnetic Resonance in Medicine*, 54:702–706, 2005.

[90] Robert N. Muller, Pierre Gillis, Francis Moiny, and Alain Roch. Transverse relaxivity of particulate MRI contrast media: from theories to experiments. *Magnetic Resonance in Medicine*, 22:178–182, 1991.

[91] Dmitriy A. Yablonskiy and E. Mark Haacke. Theory of NMR signal behaviour in magnetically inhomogeneous tissues: The static dephasing regime. *Magnetic Resonance in Medicine*, 32:749–763, 1994.

[92] Dmitriy A. Yablonskiy. Quantitation of intrinsic magnetic susceptibility-related effects in a tissue matrix: Phantom study. *Magnetic Resonance in Medicine*, 39:417–428, 1998.

[93] S. Majumdar and J.C. Gore. Studies of diffusion in random fields produced by variations in susceptibility. *Journal of Magnetic Resonance*, 78:41–55, 1988.

[94] Mariet C. W. Feltkamp, Henk L. Smits, Michel P. M. Vierboom, Rene P. Mmnaar, Barteld M. de Jongh, Jan Wouter Drijfhout, Jan ter Schegget, Cornelis J. M. Melief, and W. Martin Kast. Vaccination with cytotoxic t lymphocyte epitope-containing peptide protects against a tumor induced by human papillomavirus type 16-transformed cells. *European Journal of Immunology*, 23: 2242–2249, 1993.

[95] Kenneth Marro, Randolph Otto, Orpheus Kolokythas, Akiko Shimamura, Jean E. Sanders, George B. McDonald, and Seth D. Friedman. A simulation-based comparison of two methods for determining relaxation rates from relaxometry images. *Magnetic Resonance Imaging*, 29:497–506, 2011.

[96] Paul T. Callaghan. *Principles of Nuclear Magnetic Resonance Microscopy*. Oxford University Press, 2001.

[97] Alexander L. Sukstanskii and Dmitriy A. Yablonskiy. Gaussian approximation in the theory of MR signal formation in the presence of structure-specific magnetic field inhomogeneities. *Journal of Magnetic Resonance*, 163:236–247, 2003.

[98] D. S. Novikov and V. G. Kiselev. Transverse NMR relaxation in magnetically heterogeneous media. *Journal of Magnetic Resonance*, 195:33–39, 2008.

[99] A. F. Frhlich, L. stergaard, and V. G. Kiselev. Theory of susceptibility-induced transverse relaxation in the capillary network in the diffusion narrowing regime. *Magnetic Resonance in Medicine*, 53:564–573, 2005.

[100] J.H. Jensen and R. Chandra. Strong field behavior of the NMR signal from magnetically heterogeneous tissues. *Magnetic Resonance in Medicine*, 43:226–236, 2000.

[101] J.H. Jensen and R. Chandra. Theory of nonexponential NMR signal decay in liver with iron overload or superparamagnetic iron oxide particles. *Magnetic Resonance in Medicine*, 47:1131–1138, 2002.

[102] Wolfgang R. Bauer and Walter Nadler. Spin dephasing in the extended strong collision approximation. *Physical Review E*, 65:066123, 2002.

[103] C. H. Ziener, T. Kampf, G. Melkus, V. Herold, T. Weber, G. Reents, P. M. Jakob, and W. R. Bauer. Local frequency density of states around field inhomogeneities in magnetic resonance imaging: Effects of diffusion. *Physical Review E*, 76:031915, 2007.

[104] John D. Dickson, Tom W.J. Ash, Guy B. Williams, Alexander L. Sukstanskii, Richard E. Ansorge, and Dmitriy A. Yablonskiy. Quantitative phenomenological model of the BOLD contrast mechanism. *Journal of Magnetic Resonance*, 212:17–25, 2011.

[105] H. C. Torrey. Bloch equations with diffusion terms. *Physical Review*, 104: 563–565, 1956.

[106] P.A. Hardy and R.M. Henkelman. Transverse relaxation rate enhancement caused by magnetic particulates. *Magnetic Resonance Imaging*, 7:265–275, 1989.

[107] C. Richard Fisel, Jerome L. Ackerman, Richard B. Buxton, Leoncio Garrido, John W. Belliveau, Bruce R. Rosen, and Thomas J. Brady. MR contrast due to microscopically heterogeneous magnetic susceptibility: Numerical simulations and applications to cerebral physiology. *Magnetic Resonance in Medicine*, 17: 336–347, 1991.

[108] Richard P. Kennan, Jianhui Zhong, and John C. Gore. Intravascular susceptibility contrast mechanisms in tissues. *Magnetic Resonance in Medicine*, 31: 9–21, 1994.

[109] Jerrold L. Boxerman, Leena M. Hamberg, Bruce R. Rosen, and Robert M. Weisskoff. MR contrast due to intravascular magnetic susceptibility perturbations. *Magnetic Resonance in Medicine*, 34:555–566, 1995.

[110] Pierre Gillis, Francis Moiny, and Rodney A. Brooks. On T2-shortening by strongly magnetized spheres: A partial refocusing model. *Magnetic Resonance in Medicine*, 47:257–263, 2002.

[111] Yuri Matsumotoa and Alan Jasanoff. T2 relaxation induced by clusters of superparamagnetic nanoparticles: Monte carlo simulations. *Magnetic Resonance Imaging*, 26:994–998, 2008.

[112] E.O. Stejskal and J.E¿ Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42:288–292, 1965.

[113] Steven D. Beyea, Theodore W. Bremner, and Bruce J. Balcom. Minimization of diffusive attenuation in t2-weighted NMR images of porous solids using turboSPI. *Solid State Nuclear Magnetic Resonance*, 29:267–271, 2006.

[114] D.K. Sodickson and W.J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38:591–603, 1997.

[115] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, and Peter Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42:952–962, 1999.

[116] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47:1202–1210, 2002.

[117] David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, 2006.

[118] Michael Lustig, David Donoho, and John M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58:1182–1195, 2007.

[119] Tolga Cukur, Michael Lustig, and Dwight G. Nishimura. Improving non-contrast-enhanced steady-state free precession angiography with compressed sensing. *Magnetic Resonance in Medicine*, 61:1122–1131, 2009.

[120] M. Usman, C. Prieto, F. Odille, D. Atkinson, T. Schaeffter, and P.G. Batchelor. A computationally efficient OMP-based compressed sensing reconstruction for dynamic MRI. *Physics in Medicine and Biology*, 56:N99–107, 2011.

[121] T. Kampf, A. Fischer, T.C. Basse-Lusebrink, G. Ladewig, F. Breuer, G. Stoll, P.M. Jakob, and W.R. Bauer. Application of compressed sensing to in vivo 3D 19F CSI. *Journal of Magnetic Resonance*, 207:262–273, 2010.

[122] Michael Lustig and John M. Pauly. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic Resonance in Medicine*, 64:457–471, 2010.

[123] Daniel Holland, Careesa Liu, Chris V. Bowen, Andy Sederman, Lyn Gladden, and Steven. D. Beyea. Compressed sensing reconstruction improves variable density spiral functional MRI. In *Proceedings of the 19th Meeting of the International Society for Magnetic Resonance in Medicine*, 2011.

[124] P. Parasoglou, D. Malioutov, A.J. Sederman, J. Rasburn, H. Powell, L.F. Gladden, A. Blake, and M.L. Johns. Quantitative single point imaging with compressed sensing. *Journal of Magnetic Resonance*, 201:72–80, 2009.

[125] A. Samsonov, Y. Jung, A. L. Alexander, W. F. Block, and A. S. Field. MRI compressed sensing via sparsifying images. In *Proceedings of the 16th Meeting of the International Socety for Magnetic Resonance in Medicine*, Toronto, Canada, May 2008.

[126] A. Fischer, F. Breuer, M. Blaimer, N. Seiberlich, and P. M. Jakob. Accelerated dynamic imaging by reconstructing sparse differences using compressed sensing. In *Proceedings of the 16th Meeting of the International Socety for Magnetic Resonance in Medicine*, Toronto, Canada, May 2008.

[127] Guang-Hong Chen, Jie Tang, and Shuai Leng. Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. *Medical Physics*, 35: 660–663, 2008.

[128] Prodromos Parasoglou, A.J. Sederman, J. Rasburn, H. Powell, and M.L. Johns. Optimal k-space sampling for single point imaging of transient systemsg. *Journal of Magnetic Resonance*, 194:99–107, 2008.

[129] Ganesh Adluru and Edward V. R. DiBella. Reordering for improved constrained reconstruction from undersampled k-space data. *International Journal of Biomedical Imaging*, page 341684, 2008.

[130] Bing Wu, Rick P. Millane, Richard Watts, and Philip J. Bones. Prior estimate-based compressed sensing in parallel MRI. *Magnetic Resonance in Medicine*, 65:83–95, 2011.

[131] Reed F. Busse, Anja C.S. Brau, Anthony Vu, Charles R. Michelich, Ersin Bayram, Richard Kijowski, Scott B. Reeder, and Howard A. Rowley. Effects of refocusing flip angle modulation and view ordering in 3D fast spin echo. *Magnetic Resonance in Medicine*, 60:640–649, 2008.

[132] Meghan Halse, David J. Goodyear, Bryce MacMillan, Pavol Szomolanyi, David Matheson, and Bruce J. Balcom. Centric scan SPRITE magnetic resonance imaging. *Journal of Magnetic Resonance*, 165:219–229, 2003.

[133] Florian Knoll, Christian Clason, Clemens Diwoky, and Rudolf Stollberger. Adapted random sampling patterns for accelerated MRI. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 24:43–50, 2011.

[134] Gary M. Israel, Nicole Hindman, Elizabeth Hecht, and Glenn Krinsky. The use of opposed-phase chemical shift MRI in the diagnosis of renal angiomyolipomas. *American Journal of Roentgenology*, 184:1868–1872, 2005.

[135] Yuval Zur. Design of improved spectral-spatial pulses for routine clinical use. *Magnetic Resonance in Medicine*, 43:410–420, 2000.

[136] Angshul Majumdar and Rabab K. Ward. Accelerating multi-echo t2 weighted MR imaging: Analysis prior group-sparse optimization. *Journal of Magnetic Resonance*, 210:90–97, 2011.

[137] Mariya Doneva, Peter Brnert, Holger Eggers, Christian Stehning, Julien Sngas, and Alfred Mertins. Compressed sensing reconstruction for magnetic resonance parameter mapping. *Magnetic Resonance in Medicine*, 64:1114–1120, 2010.

[138] P. Lai, T. Zhang, M. Lustig, S. S. Vasanawala, and A. C. Brau. Improving compressed sensing parallel imaging using autocalibrating parallel imaging initialization with variable density tiled random k-space sampling. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 68, 2011.

[139] John David Jackson. *Classical Electrodynamics, Third Edition*. John Wiley and Sons, 1998.

[140] Valerij G. Kiselev. personal communication, 2010.