Exploiting Magnetic Relaxation in x-Space Magnetic Particle Imaging

By

Daniel Westbrook Hensley

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Joint Doctor of Philosophy
with University of California, San Francisco

in

Bioengineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Steven Conolly, Chair
Professor Chris Diederich
Professor John Clarke

Summer 2017

Exploiting Magnetic Relaxation in x-Space Magnetic Particle Imaging

Copyright 2017
by
Daniel Westbrook Hensley

Abstract

Exploiting Magnetic Relaxation in x-Space Magnetic Particle Imaging

by

Daniel Westbrook Hensley

Joint Doctor of Philosophy
with University of California, San Francisco in Bioengineering

University of California, Berkeley

Professor Steven Conolly, Chair

Magnetic particle imaging (MPI) is a new and rapidly developing imaging modality. Here we explore new ways to use MPI for molecular imaging and theranostic applications. A major focus is exploring methods to leverage the magnetic relaxation dynamics associated with the tracers used in MPI. In the context of canonical MPI implementations, these dynamics can reduce resolution and signal; however, they are also a source of molecular imaging contrast. In the first part of this thesis, we experimentally characterize relaxation in the canonical sinusoidal MPI implementation and explore novel ways to exploit this relaxation for colorized imaging and theranostic applications. In the second part, we describe a new approach to signal encoding in MPI that we call pulsed MPI (pMPI). We show that pMPI can allow one to fully prevent relaxation-induced resolution loss and can even provide ways to improve resolution over traditional continuous wave encoding techniques. Furthermore, pMPI can be leveraged to directly quantify magnetic relaxation that does occur and report this information in an imaging format as *relaxation images* (in contrast to the standard tracer mass/concentration images). Such approaches may pave the way for greatly improved color MPI and related molecular imaging methods. We believe a major takeaway from this work is that there remains significant untapped potential in MPI by way of unexploited physics. We believe development of these possibilities will be an exciting part of our field's future.

Acknowledgements

My work at the Berkeley Imaging Systems Lab (BISL) has been an incredible time of learning, collaboration, and excitement. So many people have been a part of this process and I am eternally grateful to all of you.

Specifically, I would like to say thank you to my wonderful fiancée, Whitney Walsh. Above all others, she has had to deal with my idiosyncrasies, frustrations, and the stress of pursuing a PhD. But she's always been there for me in a wonderful and supportive way. Thank you.

To all my close friends and family – thank you for your support and providing the outlets needed to maintain sanity. Special thanks to my mother, Sandy Hensley, for raising me and getting me to a place where I could even pursue a PhD and for always being so confident in and supportive of me.

I've had the privilege of working with a truly fantastic group of people at BISL. I appreciate the incredible concentration of brilliance, kindness, and personality the group has provided. And I had the wonderful opportunity to work in close collaboration with so many of you over the years. Special thanks to Zhi Wei Tay, with whom I have had so much fun exploring new MPI methods over the last few years. In particular we have worked in close partnership to develop various aspects of pulsed magnetic particle imaging, both of us making significant contributions individually and together. Special thanks as well to Patrick Goodwill for getting me started and providing strong mentorship early in my tenure. Patrick's fearlessness, brilliance, and well-roundedness are an inspiration to all engineers. Furthermore, it has been a pleasure to work with BISL colleagues Bo Zheng, Kuan Lu, Justin Konkle, Laura Croft, Emine Saritas, Paul Keselman, Ryan Orendorff, Elaine Yu, Prashant Chandrasekharan, and Xinyi Zhou.

A very special thanks to my PI and mentor Steve Conolly. Steve you have provided a truly enabling environment by providing a rigorous and foundational engineering culture, you care deeply about your students and seek to help us find a great place in the world, and you have provided an intellectually free environment to pursue our graduate studies. Thank you, and I will be forever in your debt for this.

Finally, I have dedicated this work to my father Mark Hensley. I can think of no bigger influencer in my life. You were larger than life and your creativity, brilliance, and sheer energy were incredible. I can only aspire to your unique synthesis of these qualities. I am forever in your debt for the worldview, intellectual framework, and moral compass you provided. Thank you.

"We consider it a good principle to explain the phenomena by the simplest hypothesis possible."

Claudius Ptolemy

Preface

Magnetic particle imaging is at a pivotal stage of development. After a little over a decade of international development, we now have multiple stable implementations. Much of this first decade of development concerned building, refining, and characterizing the first generation of MPI hardware and implementations. More recently, we have been able to focus on all important applications work. We that work in this field are certainly fascinated and delighted by the engineering challenges and elegance we see in MPI. But in the end, MPI must demonstrate clear and compelling advantages in specific applications. Fortunately, we believe we have already started building a very strong case. Accordingly, initial commercialization of the technology is playing out at this very moment. As a researcher and engineer, there is no more exciting and dynamic a time to be involved.

In this work, we set out to vigorously explore an area of MPI that holds enormous potential: exploitation of our tracers' magnetic dynamics, or *relaxation* behavior. We believe we are able to make a very compelling case for two important themes: 1) magnetic relaxation can enable powerful molecular imaging capabilities and 2) magnetic relaxation can be leveraged to fundamentally improve the encoding of our signal, with the potential to greatly improve fundamental imaging parameters of merit such as resolution.

In Chapter 1, we will provide a broad and high level introduction to MPI including its motivation and position in medical imaging, the features of its basic implementation, some of the important applications work and directions, and an intuitive discussion of current challenges. We then organize our exploration of new methods of exploiting relaxation in two parts:

In **Part I** we will explore magnetic relaxation in the current continuous wave implementation of MPI. In Chapter 2 we provide a basic introduction to the physical mechanisms of magnetic relaxation encountered in MPI and spend time cataloging and analyzing experimental data demonstrating relaxation-induced behavior in MPI over a broad parameter set.

Chapters 3 and 4 focus on using magnetic relaxation in MPI. One of the most intriguing aspects of magnetic relaxation is the encoding of physiologic information, and in Chapter 3 we exploit this by describing an approach to x-space 'color' MPI including the first *in vivo* demonstration of such an approach. We switch gears in Chapter 4 to explore theranostic MPI where we describe building the first system that combines MPI and magnetic fluid hyperthermia (MFH) in a single device. We show the ability for on demand spatial selection of heating and the first simultaneous heating and imaging.

We hope that one takeaway from Part I is a sense of opportunity by solidifying the notion that MPI physics allow for significant encoding of environmental state and physiology of interest – a key requirement for any aspiring molecular imaging modality. Future developments in this young field will go well beyond that which is contemplated here.

In **Part II**, we expand the possibilities of using relaxation in MPI, and MPI more generally, by rethinking how we excite our tracer. We explore moving beyond the narrow scope of sinusoidal continuous wave excitation and call these new approaches pulsed magnetic particle imaging (pMPI). In Chapter 5, we discuss our initial pMPI methods using square and trapezoidal wave excitation, providing a detailed and general theoretical analysis along with supporting experimental data for the one dimensional (1D) case. We show that this square wave approach allows us to remove deleterious effects of relaxation and achieve steady-state Langevin conditions. With this basic harness equipped, we then discuss several ways of actively using and exploiting relaxation with pulsed MPI techniques in Chapter 6. We focus on three approaches, first providing detailed theoretical analysis in 1D and then demonstrating these methods in proof-of-concept experiments. We show that these approaches can provide direct, quantitative measures of magnetic relaxation impulse responses, robust and quantitative measurement of local microenvironmental state, and greatly improve resolution *beyond* what was previously thought fundamental via steady-state Langevin theory. We also show that many of these abilities can be provided *simultaneously* (using only a single acquisition) in pMPI. Finally, in a short Chapter 7, we describe implementation of pMPI in the multidimensional context, including numerous simulations and one comparison of the simulator with basic proof-of-concept data. We also discuss the flexibility of pMPI in terms of generalized pulse sequencing and highlight some intriguing, but unimplemented, approaches that could be good targets for future work.

We believe this work lays an important groundwork and framework for a compelling new family of MPI approaches. But it is only a start; we hope future developments will greatly expand on these initial methods.

Daniel Hensley
Berkeley, CA
August 8, 2017

Contents

List of Figures

List of Tables

List of Abbreviations

1D	**1 D**imensional
2D	**2 D**imensional
3D	**3 D**imensional
AWR	**A**rbitrary **W**aveform **R**elaxometer
BISL	**B**erkeley **I**maging **S**ystems **L**ab
DI	**D**e-**I**onized
DSP	**D**igital **S**ignal **P**rocessing
EMF	**E**lectro**m**otive **F**orce
EPR	**E**hanced **P**ermeability (and) **R**etention
FFL	**F**ield-**F**ree **L**ine
FFP	**F**ield-**F**ree **P**oint
FFR	**F**ield-**F**ree **R**egion
FOV	**F**ield **o**f **V**iew
FWHM	**F**ull **W**idth (at) **H**alf **M**aximum
LSI	**L**inear (and) **S**hift **I**nvariant
MAA	**M**acro**a**ggregated **A**lbumin
MFH	**M**agnetic **F**luid **H**yperthermia
MNP	**M**agnetic **N**ano**p**article
MPI	**M**agnetic **P**article **I**maging
MRI	**M**agnetic **R**esonance **I**maging
PEG	**P**oly**e**thylene **G**lycol
PET	**P**ositron **E**mission **T**omography
pFOV	**p**artial **F**ield **o**f **V**iew
pMPI	**p**ulsed **M**agnetic **P**article **I**maging
PSF	**P**oint **S**pread **F**unction
ROI	**R**egion **o**f **I**nterest
Rx	**R**eceive
SAR	**S**pecific **A**bsorption **R**ate
SNR	**S**ignal (to) **N**oise **R**atio
SPECT	**S**ingle **P**hoton **E**mission **C**omputed **T**omography
SPIO	**S**uper**p**aramagnetic **I**ron **O**xide
Tx	**T**ransmit or excitation

Physical Constants

Boltzmann constant	$k = 1.380\,648\,527\,9 \times 10^{-23}\,\mathrm{J\,K^{-1}}$
Vacuum permeability	$\mu_0 = 4\pi \times 10^{-7}\,\mathrm{H\,m^{-1}}$

List of Symbols

B	induction field magnitude	T
B_1	coil sensitivity	$\mathrm{T\,A^{-1}}$
f	frequency	Hz
f_0	fundamental excitation frequency	Hz
G	linear magnetic gradient magnitude	$\mathrm{T\,m^{-1}}$
H	magnetic field magnitude	$\mathrm{A\,m^{-1}}$
I	moment of inertia	$\mathrm{kg\,m^{2}}$
k	magnetic Langevin constant	$\mathrm{m\,A^{-1}}$
$\mathcal{L}$	Langevin function	none
m	magnetic moment	$\mathrm{A\,m^{2}}$
M	magnetization magnitude	$\mathrm{A\,m^{-1}}$
x_s	location of MPI field-free region center	m
η	viscosity	$\mathrm{Pa\,s}$
γ	gyromagnetic ratio	$\mathrm{rad\,s^{-1}\,T^{-1}}$
ρ	particle density	$\mathrm{particles\,m^{-3}}$
ω	angular frequency	rad
τ	time constant	s

Dedicated to Mark Hensley

Chapter 1

Introduction and Background

Magnetic particle imaging (MPI) is an emerging tracer-based and molecular imaging modality. First publicly described in 2005 by Gleich and Weizenecker (Gleich and Weizenecker, 2005), MPI endeavors to be the next big clinical imaging modality, following in the footsteps of X-ray computed tomography (CT), magnetic resonance imaging (MRI), nuclear medicine techniques, ultrasound, and other clinical modalities of the 20th century.

As a young modality, MPI is a fascinating and fertile area for research. I have greatly enjoyed my chance to engage with this technology as a graduate student researcher. The work herein encompasses a variety of topics all with the common goal of improving and maturing MPI.

In this introductory chapter, we will introduce the overarching context of MPI – medical imaging – and discuss the basic physical principles that govern MPI and constitute the x-space MPI process. We will conclude by discussing some of the current and ongoing challenges in the field. This work takes aim at a few of these challenges.

1.1 Medical Imaging

Medical imaging is concerned with whole-body imaging of human subjects for medical analysis and diagnostic value. Although no clinical MPI device currently exists, this is the clear long term goal of the field. Current research and commercial efforts are focused on preclinical applications. Preclinical imaging is typically concerned with whole-body imaging of animal subjects for the purpose of basic scientific research or for analysis of eventual clinical applications. In this sense, preclinical imaging is distinct from microscopy and other forms of small scale basic science imaging but does straddle application areas of interest to both basic science and clinical medicine. For example, preclinical functional brain imaging can benefit basic neuroscience research and provide valuable information as we continue

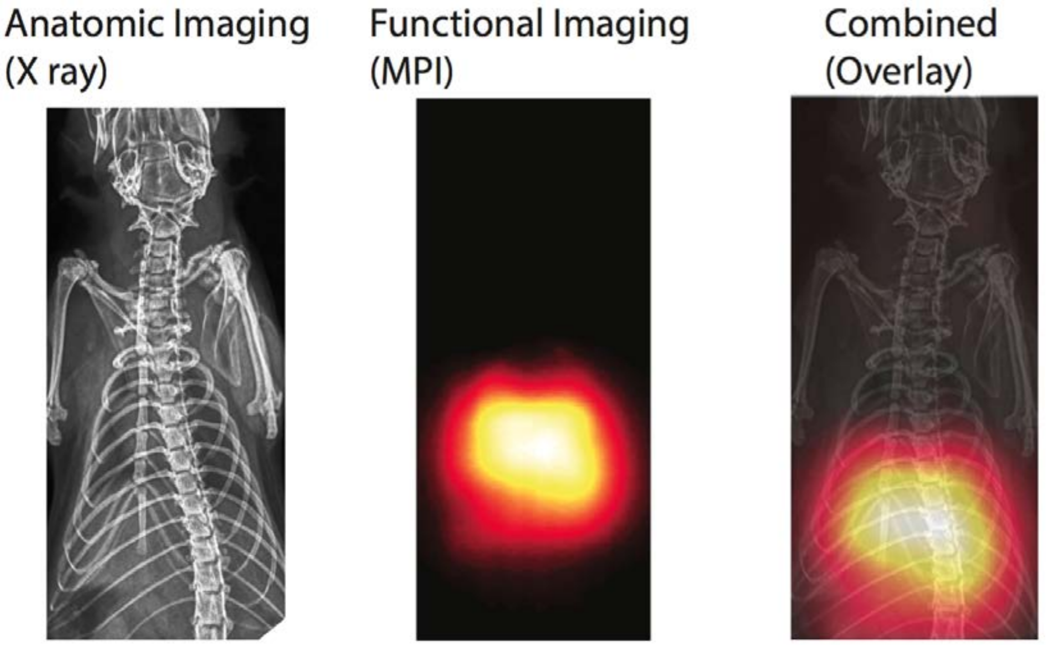

FIGURE 1.1: Categorizing imaging as anatomic vs. functional. Anatomic imaging (left) can provide high resolution images of structure but with little to no contrast relevant to physiology or pathophysiology of interest. Functional or molecular imaging techniques, which often leverage contrast agents or tracers (center) specifically target the illumination of contrast related to physiology of interest but perhaps without high resolution, or any, structural information. (Right) We can often get the best of both worlds by combining techniques such as in this overlay. These images are from a 'color MPI' dataset we will explore in Chapter 3.

to understand the mechanisms of the brain (Park and Friston, 2013), while other practitioners of functional brain imaging may more exclusively target clinical application. Research domains such as immune cell tracking and stem cell tracking also explicitly straddle both basic science research and possible clinical diagnostic/therapeutic methods (Weissleder, Nahrendorf, and Pittet, 2014; Ahrens et al., 2005; Ahrens and Bulte, 2013; Zheng et al., 2015; Zheng et al., 2016). Preclinical imaging methods may also be attractive in the drug development space providing, for example, analysis of drug delivery and pharmacokinetics *in vivo* prior to full human trials.

As Fig. 1.1 illustrates, a conceptual distinction between 'anatomic' imaging and 'functional' or related 'molecular' imaging can be useful for discussing various imaging modalities and their goals (James and Gambhir, 2012). In this context, MPI is being developed primarily as a new functional and molecular imaging modality.

1.1.1 Anatomical Imaging Techniques

Anatomical imaging modalities derive signal from all or most tissue in the subject, possibly with some contrast that distinguishes between different tissue types. Typical modalities include magnetic resonance imaging (MRI), X-ray projection imaging, X-ray computed tomography (CT), and ultrasound. These modalities can provide exquisite high resolution information (spatial and/or temporal) about the anatomic and structural makeup of a subject. The development of early anatomic techniques such as X-ray projection imaging in the 20th century greatly altered and improved the practice of medicine by noninvasively providing a visual window into the body.

While anatomic techniques are invaluable, native contrast between tissues or contrast associated with a specific physiology or pathology of interest may be poor. Furthermore, the strong background signal in the tissue may dominate small amplitude signals carrying this contrast information.

1.1.2 Functional Imaging Techniques

In contrast to purely anatomic imaging, functional and molecular imaging methods and modalities seek to illuminate specific physiologic, functional, or pathological information in the image. Often these methods leverage contrast agents or tracers. Modalities that are fundamentally anatomical in nature can provide functional information with the use of contrast agents, as in the case of gadolinium-based agents in MRI or microbubbles in ultrasound (James and Gambhir, 2012).

Contrast agents modify the native or background signal, adding desirable contrast. Positive contrast agents increase the local intensity of the image (*e.g.*, gadolinium T1 agent in MRI) while negative contrast agents decrease the local intensity (*e.g.*, T2* iron oxide tracers in MRI). Typically, positive contrast agents are preferred due to the ambiguity associated with negative contrast. In MRI, for example, myriad conditions such as the presence of air pockets may induce a low signal region, complicating identification of the agent.

Tracer modalities such as MPI and nuclear medicine modalities positron emission tomography (PET), scintigraphy, and single photon emission computed tomography (SPECT) are different. In tracer modalities, there is no background signal to modulate. Instead, the only signal detected comes from the tracer itself. This is powerful for molecular imaging purposes because one can theoretically obtain *ideal* contrast for a target physiology or pathology of interest. Furthermore, typically much lower (orders of magnitude) concentrations of tracer are required than in the case of contrast agents. One associated downside is that tracer modalities can lack contextualizing anatomical information. For this reason, combined anatomic and molecular imaging overlays are common as in PET/CT and MPI/CT (Beyer et al., 2000; Zheng et al., 2016).

1.2 LSI Imaging Systems and LSI Analysis

Linear systems theory has provided a crucial foundation for developments in physics and engineering. Even for the case of explicitly nonlinear systems, it is the basis for most analytical approaches (Gaskill, 1978; Prince and Links, 2006). Medical imaging is no different. All of the major medical imaging modalities, such as MRI, projection X-ray and X-ray CT, and ultrasound operate in linear or approximately linear modes. This is desirable for several reasons. First, linear systems theory is well understood and mathematically tractable. This translates to reliable and robust imaging systems that can be easily modeled, their behavior easily predicted, noise properties easily understood, and artifacts minimized. Importantly, the nature of images reconstructed from systems that are linear and shift invariant (LSI) do not depend on the nature of the truth image itself or any *a priori* assumptions built into the reconstruction scheme beyond the linear model. While these features of LSI systems are powerful, we must caveat this discussion by acknowledging that no systems are perfectly LSI and nonlinear approaches that inject *a priori* information can be leveraged to great effect (Lustig, Donoho, and Pauly, 2007; Silva et al., 2010). We note, however, that clinical medicine is an exceedingly high bar. Even if nonlinear artifacts are somewhat rare, and perhaps well tolerated in many other application areas, this can impact a patient diagnosis in medical imaging.

A linear system in real space, described by the linear operator H, exhibits the following superposition property:

$$H(\alpha\rho(x) + \beta\rho(x)) = \alpha H(\rho(x)) + \beta H(\rho(x)) \quad \forall \alpha, \beta \in \mathbb{R} \tag{1.1}$$

If $\tilde{\rho}(x) = H(\rho(x))$, then a shift invariant system exhibits the following behavior:

$$H(\rho(x - x_0)) = \tilde{\rho}(x - x_0) \tag{1.2}$$

Powerful linear algebra techniques can be brought to bear on LSI systems. At risk of brushing all-important details aside, any LSI system can be elegantly described mathematically by:

$$y = Ax \tag{1.3}$$

where A (*e.g.*, represented as a matrix) is a linear operator that acts on the input x to produce the output y. The operator A describes everything that occurs in the linear transformation. For example, in MPI, the operator A describes the blurring effect of the fundamental Langevin physics, the action of inductive reciever coils to collapse spatial information to temporal information, and the specifics of our image encoding trajectory (Goodwill and Conolly, 2011; Rahmer et al., 2009). In this context, we note

$$A = A_N A_{N-1} \dots A_1$$

that A may be broken down into a composition of various conceptually distinct linear components of the LSI system. This is important, because solving a linear system, or in the case of MPI, reconstructing an image from the raw data, involves inverting this linear system (whether from explicitly inverting a linear operator or performing this implicitly in a procedural algorithm). If a linear process is well-posed, then this is straightforward. But inversion of a linear system can be ill-posed (*e.g.*, deconvolution of a modality's fundamental physical blur), meaning that noise and any sources of coherent error may be greatly magnified in the operation (Shahram and Milanfar, 2004). Inversion may also be impossible if information is destroyed in the forward operation, a case that often motivates the use of *a priori* assertions. This analysis carries over to the individual components of a linear operator. For example:

$$\begin{aligned} \mathrm{inv}(A) &\rightarrow \text{ill-posed} \\ \mathrm{inv}(A_N A_{N-1} \dots A_2) &\rightarrow \text{well-posed} \\ \mathrm{inv}(A_1) &\rightarrow \text{ill-posed} \end{aligned}$$

In the context of MPI, we can choose to reconstruct by inverting only the well-posed portion of the linear process (MPI trajectory spatio-temporal encoding) while avoiding deconvolution of the fundamental physics blur:

$$
\begin{aligned}
y &= \underbrace{(A_N A_{N-1} \dots A_2)}_{\tilde{A}} \cdot \underbrace{(A_1 x)}_{x_{\mathrm{bl}}} = \tilde{A} \cdot x_{\mathrm{bl}} \\
\hat{x}_{\mathrm{bl}} &= \tilde{A}^{-1} y \rightarrow \text{well-posed}
\end{aligned}
\tag{1.4}
$$

where we reconstruct an estimate ($\hat{x}_{\mathrm{bl}}$) of the truth image, blurred by the fundamental physics, and do not attempt to deconvolve the fundamental blur. This is the approach that we take in x-space MPI reconstruction (Goodwill and Conolly, 2010; Goodwill and Conolly, 2011; Goodwill et al., 2012b; Lu et al., 2013).

We do not deconvolve for at least two reasons: (a) avoidance of stark signal-to-noise ratio (SNR) loss inherent to deconvolution and artifact introduction due to imperfect knowledge of the blurring function, and (b) the relative youth of MPI means that there are many avenues to improve resolution and other aspects of the MPI process through advances in the way the *fundamental imaging physics* are leveraged. This may be through novel encoding schemes and/or hardware implementations. In this manner, we may be able to fundamentally modify A_1 in the above examples, including its posedness. This is an approach taken in part II of this work. When one has exhausted improvements at the fundamental physics level, then reconstruction and postprocessing methods may become the only targets for improvement. But until then, we find that improving the fundamentals typically yield much higher rewards.

As we have harped on the ill-posed deconvolution of blurring functions, we close out this section introducing the nomenclature for this concept in imaging. All LSI imaging systems are associated with some fundamental physics-induced blur that limits the ability to resolve fine details. This blur can be expressed mathematically as a convolution relation. A convolution operator is a type of linear operator that shows up routinely in physics.

Convolution of distribution ρ with blur h in 1D:

$$
\rho_{\mathrm{blur}}(x) = (\rho * h)(x) = \int \rho(u) h(x-u) \, \mathrm{d}u \tag{1.5}
$$

The important and powerful connection between convolution and LSI imaging systems is that the blur is fully and totally described by a 'point-spread function' (PSF). The PSF can be analytically derived or empirically measured by subjecting the system to a point source input (insert $\delta(x)$ for $\rho(x)$ in (1.5)). The δ function is the identify function for convolution. With the PSF in hand, the effect of subjecting any possible input to the system is known. For an LSI system H, with PSF $h(x)$:

$$\begin{aligned} H(\delta(x)) &\triangleq h(x) = \delta(x) * h(x) = h(x) \\ \therefore H(\rho(x)) &= \rho(x) * h(x) \quad \forall \rho(x) \end{aligned}$$

With knowledge of the PSF, one can derive many of the important parameters of merit for an imaging system, such as resolution and signal scaling.

1.3 Magnetic Particle Imaging

Here we introduce the field of MPI broadly, with an emphasis on describing the MPI process as a cohesive and intuitive whole. We begin by discussing the physical principles governing MPI, then describe the traditional methods of encoding imaging information and the reconstruction approaches that follow, move on to the state-of-the-art in terms of MPI applications, and finish with a discussion of some open problems in the field.

1.3.1 Basic Principles

MPI is a tracer-based modality, and the governing physics are those associated with the tracer. At the most fundamental level, MPI measures the spatial distribution of a magnetic tracer as revealed by magnetic flux changes when the tracer responds to translating gradient fields. MPI hardware then consists of systems that apply magnetic excitation in a specific manner to this distribution (transmit or Tx) and systems that perceive the magnetic response of the tracer (receive or Rx).

The core of any imaging modality is how spatial information is encoded in the recieved signal. In MPI, this is accomplished by the use of magnetic gradient fields and tracers that exhibit magnetic saturation (and most often associated with superparamagnetism). In this context, most MPI tracers are superparamagnetic iron oxides (SPIOs). As illustrated in Fig. 1.2(a), the steady-state physics of SPIOs, and superparamagnetic materials in general, is described by the Langevin function. The salient feature of a Langevin-shaped magnetization curve (magnetization as a function of externally applied field) is a rapidly varying, approximately linear region at low applied

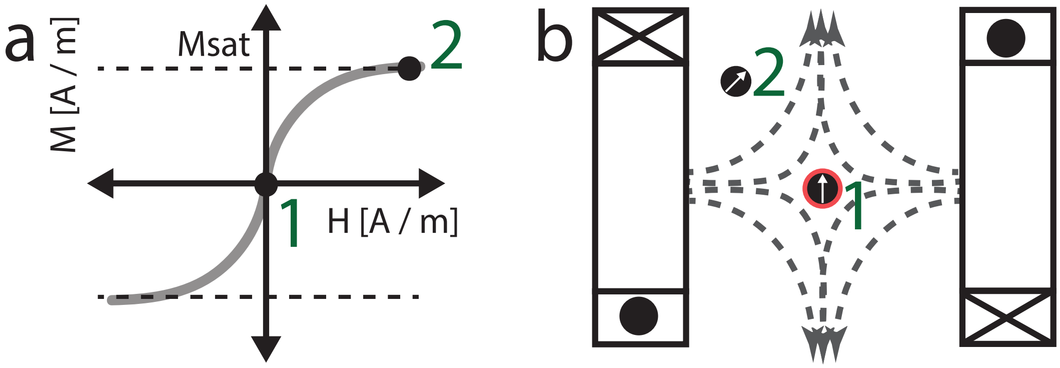

MPI tracers are chracterized by Langevin M-H curves exhibiting saturation phenomena.

This allows for spatial localization of MPI signals using gradient fields. Tracer far from the 'field-free' region are saturated and unable to respond significantly to an AC excitation.

FIGURE 1.2: Spatial localization of the MPI signal via Langevin physics. (a) A depiction of a typical magnetization vs. applied field curve for a superparamagnetic MPI tracer showing saturation phenomena. (b) Spatial gradient fields can harness this physics to create a sensitive region. This is the crux of spatial localization in MPI.

field magnitudes surrounded by regions of saturation at high field magnitude. This structure is exploited in MPI to produce a sensitive region as illustrated in Fig. 1.2(b). Only tracer located near the center of this 'field-free' region (FFR) is unsaturated and thus capable of responding to a superposed AC excitation.

The theory of superparamagnetism is fundamental to MPI and sufficient for a complete derivation of the MPI signal and imaging equations, as traditionally performed (Rahmer et al., 2009; Goodwill and Conolly, 2011). However, a major assumption in doing so is that dynamic magnetic processes (called *relaxation* in the field) are negligible. Strictly speaking, these derivations are only valid for a tracer that instantly responds to the excitation fields. In some cases, the response times associated with magnetic dynamics are much faster than the MPI excitation slew rates, in which case the steady-state assumption is accurate. However, this is known to be violated to some degree generally and especially when using tracers with larger magnetic cores. And by virtue of the steady-state imaging and signal equations, use of larger tracers is highly motivated (Goodwill and Conolly, 2010).

Relaxation in MPI and the ramifications of non-negligible relaxation will be discussed at length in this thesis and are a primary concern of part I and the major motivation behind part II. A more detailed discussion of steady-state Langevin physics can be found in the literature (Bean and Livingston, 1959; Gleich and Weizenecker, 2005; Rahmer et al., 2009; Goodwill and Conolly, 2010; Goodwill and Conolly, 2011; Lu et al., 2013).

1.3.2 Signal Encoding and Reconstruction

MPI may be understood as a sensitive point (region) imaging modality. At the highest level, all MPI systems involve creation of a field-free region and translation of this FFR over space during an imaging scan. As illustrated in Fig. 1.2, any SPIO tracer located far from the FFR is saturated while tracer coincident with the FFR can respond strongly to an excitation field. An MPI scan is nothing more than moving the FFR structure around an imaging volume to query the imaging field-of-view (FOV) for the presence of tracer.

Strong magnetic field gradients are used to create a sensitive FFR delineated from an encompassing region of high magnetic field magnitude as illustrated in Fig. 1.2. To scan the FOV, spatially homogeneous but temporally varying excitation and shift fields are superposed on the gradient field and/or mechanical means of shifting the excitation magnets and sample are employed. The result is to shift the position of the FFR relative to the object of interest as desired. A third magnetic field component, a spatially homogeneous but 'fast' AC excitation (*e.g.*, $20\ \mathrm{kHz}$), is then superposed. As illustrated in Fig. 1.3, these mechanisms allow us to sample our FOV and encode an image: The gradient fields provide spatial selection, slow-slewing shift mechanisms allow coverage of the large FOV over time, and the AC excitation induces the bulk of the received signal at any given instant. In some approaches, fast excitation concomitantly in multiple directions can allow encoding of data in a small FOV quickly without the use of distinct slow shift mechanisms (Gleich, Weizenecker, and Borgert, 2008). As in other modalities, inherent tradeoffs between parameters such as image acquisition speed, total FOV size, native resolution, and SNR lead to opportunities to optimize the excitation and reconstruction strategies as desired.

Regardless of the details of the trajectory, due to the nonlinear saturation magnetization that defines SPIO tracers, only when the FFR is coincident with a region containing tracer will a significant change in the tracer magnetization be induced by the AC excitation field. When the FFR is farther away, the tracer is saturated and no appreciable signal is detected because there is no appreciable change in SPIO magnetization.

We term our approach 'x-space' because signal acquisition is intuitively just a direct sampling of the spatial domain ('x-space') using a translating

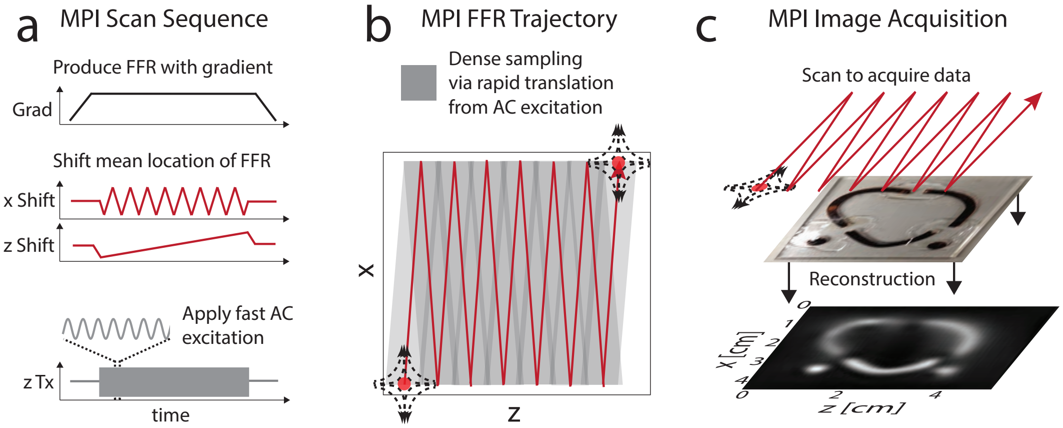

FIGURE 1.3: The canonical x-space MPI approach. (a,b) During an MPI scan, gradient waveforms create the FFR, shift waveforms move the mean location of the FFR to cover the entire field-of-view (FOV), and an AC (transmit or Tx) waveform provides the temporally fast excitation. The superposition of shift and excitation waveforms provides an overlapping and dense sampling of the FOV. (c) MPI experimental results for a single slice from a 10 minute tomographic scan of a simple phantom using our in house FFP scanner. In reconstruction, the received time-domain data is correlated with the trajectory of the FFR in the image domain to grid images.

sensitive region (in contrast to encoding directly in the Fourier domain as in the 'k-space' of MRI, or significant usage of the Fourier domain as in system matrix MPI methods). Inductive pickup coils are used to record the changes in magnetization over time. For reconstruction, a simple gridding procedure can be used to map the received time-domain voltages to an image domain grid (pixel/voxel grid) using knowledge of the FFR trajectory. Intermediate steps such as compensating for the changing velocity of the AC excitation and recovery of information lost due to direct feedthrough mitigation strategies are described in prior work (Lu et al., 2013; Konkle et al., 2015).

In contrast to this fully x-space approach, a system matrix approach is also widely used in the field (Rahmer et al., 2009; Rahmer et al., 2012). This approach differs from the x-space method in several ways. First, in reconstruction, the time-domain signal is generally transformed to the Fourier domain. Second, the reconstruction is formulated as an inverse problem by measuring system functions and constructing a suitable system matrix (Rahmer et al., 2009; Rahmer et al., 2012). This calibration matrix is constructed using numerous full MPI scans taken with a calibration point source located at each position associated with a voxel in the output image domain. The Fourier domain output of each scan becomes a column in the system matrix.

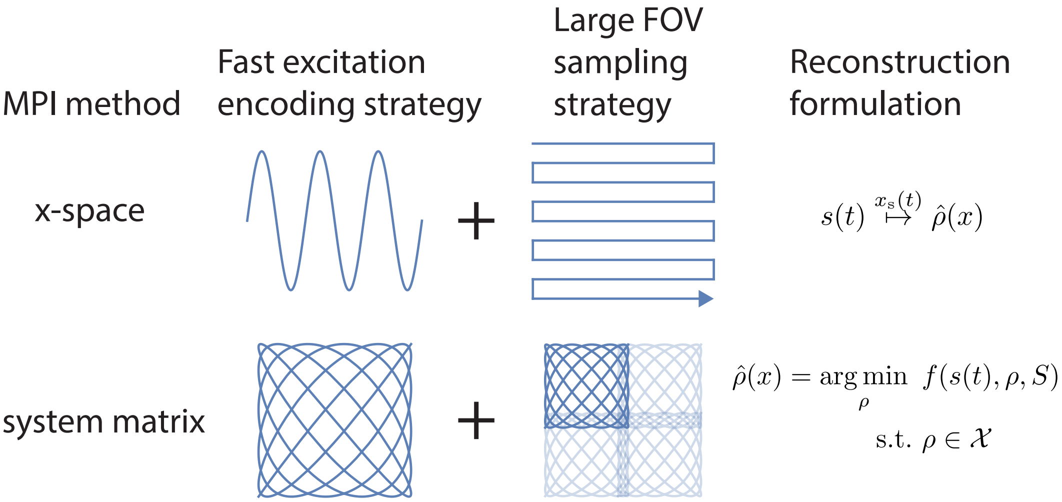

FIGURE 1.4: Major approaches to MPI signal acquisition and reconstruction. In the x-space MPI approach, a sinusoidal excitation (typically in one fixed direction per scan) is combined with slow-shift rastering to sample a large FOV. Reconstruction consists of gridding the received voltages from inductive receiver coils to images using the known FFR trajectory. In the system matrix approach, FFR trajectories often consist of (sinusoidal) Lissajous or similar fast excitation patterns in multiple dimensions. One or more such encodings are used to sample an entire FOV. Reconstruction consists of inverting a system matrix constructed by scanning a known point source located at each voxel position.

Mathematical modeling may help ease the time and space burden associated with constructing and storing this matrix (Knopp and Weber, 2015). An image acquisition then consists of applying this same scan to an unknown tracer distribution. Reconstruction is performed by regularized inversion of the system matrix applied to the Fourier-transformed receiver coil data. In effect, the signal for the unknown image scan is projected onto a space consisting of the results of point source scans associated with each voxel. Depending on the relationship between the intrinsic resolution (dictated by the steepness of a tracer's Langevin curve and the linear gradient strength) and the spacing between voxel centers, the system matrix will be more or less ill-conditioned and inversion will perform more or less implicit deconvolution of the fundamental blur.

Fig. 1.4 describes some of the high level differences between the x-space and system matrix methods, as typically performed. The sensitive point intuition applies to both methods, they simply differ in how to use FFR trajectories to reconstruct images. Per the saturation physics described in Fig. 1.2, both of these methods report resolution that scales inversely with the gradient strength (Rahmer et al., 2009; Goodwill and Conolly, 2010).

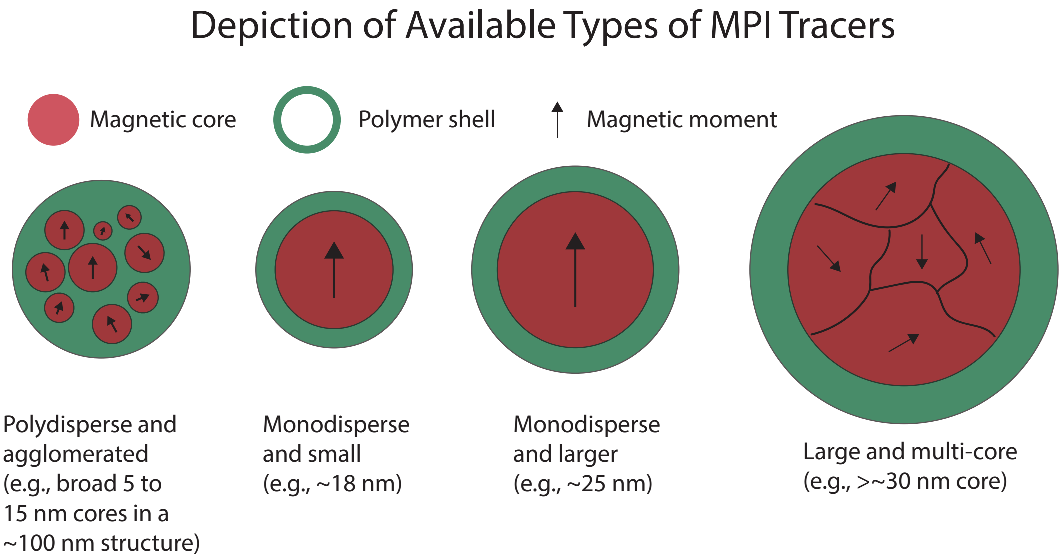

FIGURE 1.5: Depiction of types of tracers encountered in MPI. The earliest and most widely available tracers such as Resovist are composed of highly polydisperse cores embedded in a polymer (*e.g.*, dextran) shell structure (left). More recently available, and typically the highest performing, tracers are single core and monodisperse (center). If the tracers get too large (*e.g.* greater than 30–40 nm), then only multi-domain states are stable.

1.3.3 MPI Tracers

MPI tracers are dominated by SPIOs with magnetic cores in the range of approximately 10 – 30 nm (Goodwill and Conolly, 2010; Rahmer et al., 2009; Ferguson, Khandhar, and Krishnan, 2012). Fig. 1.5 depicts some of the common types of tracers encountered in MPI. They range from small and polydisperse to large, monodisperse, but multidomain.

In the context of superparamagnetism, the Langevin function, which dictates achievable resolution in canonical MPI, is a statement of the energy balance between the coherent magnetic torque imposed on the tracer by an applied field and the effect of thermal energy to oppose coherent alignment with the field. Superparamagnetic saturation occurs when the magnetic torque is sufficiently stronger than thermal effects such that, on average, the tracer is fully aligned with the applied field. Larger magnetic cores have larger magnetic moments leading to saturation at relatively lower fields.

The Langevin theory incentivizes use of larger and larger magnetic core sizes for improved resolution and SNR. However, several factors have restricted the field to the limited range of approximately 10 – 30 nm. First, as SPIO tracers increase beyond about 30 nm, individual nanoparticles are only stable in multiple domains, eventually giving rise to ferromagnetic behavior. Larger ferromagnetic tracers can aggregate, produce high amounts of localized heating, and can physically move in response to the strong gradient fields of MPI. The latter is particularly alarming due to the risk to the subject's tissues and severe imaging artifacts due to dynamic perturbation of the tracer distribution itself. Nonetheless, MPI tracers exhibiting ferromagnetic (or a similar phenomenon perhaps best termed dynamic coercivity) may be of great future interest in MPI.

The second reason for the limited MPI core size range is due to magnetic relaxation effects. As the tracer size increases, the characteristic response time associated with all physical mechanisms of tracer relaxation increases, eventually becoming non-negligible with respect to excitation slew rates. When this occurs, canonical MPI systems are adversely affected, degrading both signal intensity and achievable resolution (Tay et al., 2017). The effects of magnetic relaxation in MPI is a very active area of research (Goodwill et al., 2011; Croft et al., 2016; Dhavalikar et al., 2016; Deissler, Wu, and Martens, 2014; Biederer et al., 2009; Bauer et al., 2015). On the other hand, magnetic relaxation may also encode valuable information about a tracer's identity and local microenvironmental conditions. This information can be used to provide desirable image contrast or colorized images (Rahmer et al., 2015; Hensley et al., 2015b). The exploitation of magnetic relaxation is a major theme of this work.

The state of MPI tracers is rapidly evolving. Currently, polydisperse tracers such as Resovist and perimag (micromod Partikeltechnologie GmbH, Germany) are most often employed. This is largely for legacy and availability reasons as recent development in highly monodisperse tracers, especially those around 25 nm in core size, have shown significant increases in MPI performance (Ferguson, Khandhar, and Krishnan, 2012; Ferguson et al., 2015). Commercial entities are beginning to supply such tracers (LodeSpin Labs, Washington and Imagion Biosystems, New Mexico).

While the magnetic core dictates much of the MPI behavior, the polymer shell used to encapsulate the cores are a very important, if currently underexplored, aspect of MPI tracers. A polymer shell is necessary to imbue the tracer with hydrophilicity and provide stability in aqueous environments. Furthermore, these polymer shells can be functionalized as desired for MPI applications, *e.g.*, antibody and other targeted binding moieties or PEGylation to increase blood circulation half-life. When non-negligible magnetic relaxation is present, design of the polymer shell can be important for color MPI and related methods that attempt to exploit relaxation behavior.

We fully expect future development to produce exciting MPI tracer systems such as magnetic switches and drug delivery vehicles – the fundamental physics for such developments is clearly present. But getting there will require work from experts in areas such as chemistry and materials science. This thesis will focus largely on engineering systems, encoding methods, and reconstruction regimes that exploit MPI physics and available tracer paradigms. It is now clear that tracer design is another very important frontier for advances and improvements to the state of the art of MPI.

1.4 MPI Applications

After the initial public introduction of MPI in 2005, the field busied itself with developing and maturing hardware systems, encoding approaches, and reconstruction methods. Eventually, as these systems matured enough, we began exploring the important end goal: applications. MPI's tracer nature makes it appealing in a number of application domains, such as angiography and real-time imaging, quantitative cell tracking, lung perfusion imaging, and cancer imaging. Next we review some of these developments and provide sources in the literature.

1.4.1 Angiography and Perfusion Imaging

Much of the early focus of MPI was in angiography and real-time imaging. Fast, real-time MPI is possible and MPI tracers may have important safety advantages over the contrast agents, such as gadolinium and iodine, that are used in other contrast agent-based angiographic methods.

Real-time MPI in 3D requires explicit navigation of fundamental tradeoffs between image acquisition speed, FOV, resolution, and SNR. For example, the physical extents covered by sinusoidal oscillations in a given dimension is equal to the amplitude of the excitation divided by the gradient strength. For high resolution systems (*e.g.*, 6 $\mathrm{T\,m^{-1}}$), a typical 20 mT excitation only covers about 6.5 mm. Gradient strengths can be relaxed to improve this problem at the expense of a loss of fundamental resolution as the sensitive FFR becomes physically larger. Nonetheless, even at high resolution, 2D slice (field free point, FFP) or projection images (field free line, FFL) with a FOV in the range of 5 cm x 5 cm can be obtained at a temporal resolution of a few images per second (Goodwill et al., 2012b).

In a seminal 2009 paper, Weizenecker et al. showed real-time (21.5 ms per volume acquisition) imaging of a beating mouse heart (FOV: 20.4 $\times$ 12

$\times$ 16.8 mm^3) using a 3D Lissajous excitation trajectory and a system matrix reconstruction (Weizenecker et al., 2009). In this case, the temporal resolution was primarily achieved by encoding an entire 3D volume in a single Lissajous repetition. Clinically relevant concentrations of the Resovist tracer were used. Resovist is an SPIO originally developed as an MRI T2* contrast agent. Due to its availability and established clinical use, it has become the *de facto* gold standard tracer in MPI.

Our group typically prioritizes high SNR and high native resolution, without the use of deconvolution or related methods. Recently, we have shown the first *in vivo* traumatic brain injury imaging in MPI (Orendorff et al., 2017). We have also accumulated a number of high resolution brain perfusion scans and real-time 2D perfusion scans soon to be published. These data indicate that even when using large gradients for high resolution, real-time data acquisition is possible.

1.4.2 Cell Tracking

The first major application thrust for our group was in quantitative cell tracking using MPI. In many ways, cell tracking is an ideal application for MPI. Cell labeling with SPIOs is a fairly well understood process. And if we want to track a very small number of cells in the body over time, the zero background signal, quantitation, and high sensitivity of MPI are a perfect match. In two recent papers, our group highlighted these capabilities with the first *in vivo* MPI cell tracking studies (Zheng et al., 2015; Zheng et al., 2016).

In the first study, the fate of neural stem cells injected into the brain was tracked over the course of three months. This emphasized the longitudinal and quantitative capabilities of MPI. In this work, one of the stem cell injections went awry and, although unintended, this provided a demonstration of the utility of cell tracking with MPI as we were unambiguously able to see the movement of part of the tracer population to the ventricle.

In the second study, mesenchymal stem cells were injected through the tail vein and tracked. Due to size, the stem cells all immediately became entrapped in the lungs – the first encounter of capillary level flow for tracer introduced venously. Longitudinal scans then quantitatively tracked the physiologic clearance process as the labeled cells were slowly cleared from the lungs and the tracer processed through the liver. En route, high contrast and quantitative images of the lungs and then liver were obtained. At the time of this work, experiments established 200 labeled stem cells as an SNR = 1 threshold, demonstrating the high sensitivity MPI. Furthermore, this value is in no way fundamental, only characteristic of our device and approach at the time. These initial application papers confirmed the capabilities of MPI to provide high contrast, high sensitivity, and quantitative cell tracking.

1.4.3 Lung Perfusion

The promising results of 'lung imaging' using injected stem cells directed us to intentional lung imaging. Imaging of the lungs can be difficult with established techniques such as MRI and ultrasound (air-tissue interface issues). Traditionally, nuclear medicine 'V/Q' techniques have been used for some lung imaging applications. In this approach, separate inhalation and perfusion nuclear medicine scans are performed to assess lung function. For perfusion scans, macro-aggregated albumin entities are tagged with a radiotracer. Upon injection, as in the case of the labeled stem cells in our previous study, the aggregates get stuck in the capillaries of the lung. In this manner lung function can be assessed, *e.g.*, gross obstructions to flow can be detected. We hypothesized that MPI could be used in the same manner, but with the benefits of safety and optional longitudinal scanning. Zhou et al. were able to show just that in a recent paper from our group (Zhou et al., 2017). Subsequently we have obtained preliminary MPI ventilation scan data (currently unpublished) using nebulized MPI tracers, allowing us to provide both V and Q lung images in MPI. Future work will begin to asses MPI V/Q imaging in pathological models.

1.4.4 Cancer

One of the most important frontiers in medicine is cancer. And while we will leverage any and every tool available in this fight, molecular imaging modalities are particularly attractive. Recently, our group has shown the first *in vivo* cancer imaging in MPI (Yu et al., 2016). In this work, tumors induced in a rat model were imaged with high contrast after applying untargeted but long-circulating MPI tracer. In this case, the contrast came from the enhanced permeability and retention (EPR) effect exhibited by *some* lesions. This exciting study highlights a few of the reasons MPI can be disruptive in molecular imaging applications. First, the ideal contrast provided by the tracer nature of the modality allowed identification of the tumor without ambiguity, even in the presence of a substantially more intense signal from other regions (*e.g.*, the heart or the liver depending on the time point). Second, the availability of a long-circulating tracer allowed us to take many longitudinal time points and the quantitative nature of MPI allowed us to quantify the tracer biodistribution over these time points. In particular, we observed distinct early rim-filling of the tumor, followed by wash-in of the core, and ending in eventual wash-out and clearance of the tumor. Quantitative analysis of the results included interesting multi-compartmental modeling of the biodistribution between, *e.g.*, the tumor, circulatory volume, and liver. In future work, we will be exploring more advanced imaging methods such as active targeting, but we believe these high contrast and quantitative results speak to the potential for MPI to impact the cancer behemoth.

1.5 Open Challenges in MPI

The MPI field has made a lot of progress in the 12 years since the seminal publication by Gleich and Wiezenecker (Gleich and Weizenecker, 2005). But we are still very much a young and dynamic field. MPI 12 years from now may look nothing like what it does today. And while we believe we have established compelling groundwork to motivate taking this new modality to the clinical level already, it is clear that fundamental new developments and innovation in all directions will continue to make MPI a more attractive and compelling modality. Here we briefly describe some of the open challenges in MPI to give the interested reader some perspective on the current state of the art and provide the context that motivated the work in this thesis.

1.5.1 Hardware Optimization

MPI hardware is a ripe area for development and innovation. While a few systems worldwide have matured to a state capable of providing robust and consistent imaging, designs are hardly unassailable. Here we describe some of the challenging areas of hardware design in MPI:

Tx Coil and Resonant Systems Design

In MPI we need to provide high power excitation ($\approx 20\ \mathrm{kHz}$ and up to 30 mT) homogeneously over a large imaging FOV. Furthermore, canonical MPI operates in continuous mode (Tx concomitant with Rx) and the tracer signal lives in the space of harmonics of the fundamental excitation frequency. The tracer signal is orders of magnitude lower than the excitation, and so operating in continuous Tx/Rx mode means that we typically must abandon the tracer signal at the fundamental frequency due to a practical inability of separating the extremely high energy fundamental feedthrough from the smaller tracer signal. Additionally, any higher harmonic distortion in the excitation waveform is highly problematic. The orders of magnitude difference between Tx and Rx make distortion that may be negligible in other applications highly problematic in MPI.

Unlike MRI, MPI is quite robust to imperfect field homogeneity (*e.g.*, $\approx$ 95% typically being adequate). Instead, however, MPI is highly sensitive to Tx tonal purity. This has motived the design of unique resonant systems that provide high field while extremely attenuating sources of higher harmonic distortion. Such transmit chains involve distinct impediance matching, low-pass filter, and resonant stages. Since all of these stages generally interact, and due to the sensitivity of resonant systems (*e.g.* to system thermal fluctuations, etc.), design of these systems is challenging. Finally, it is generally desirable to be able to transmit in multiple (*e.g.*, orthogonal) directions, possibly simultaneously. In the general case, this requires multiple Tx coils, each with its own resonant chain, in close proximity. And although coil designs for coils with principal axis down the bore of the device can be designed as simple solenoids, orthogonal coils require more complex geometries such as saddle coil designs. Coupling between the coils can lead to issues such as detuning or dynamically modulated feedthrough, further complicating the Tx system design space.

Rx Coil and Rx Chain Design

Design of receiver systems is always incredibly important. Per our use of continuous Tx/Rx in canonical MPI, a major feature of MPI receiver coil systems is a gradiometric configuration. Gradiometer coil designs are composed, for example, of distinct components of the coil wound in opposite directions such that (in the limit) zero net area is presented to the homogeneous Tx field. The sample is placed only within one of these components. The result is significant cancellation of Tx feedthrough while fully coupling with the tracer signal.

An analog receive chain typically contains a notch filter at the fundamental excitation frequency to attenuate the strong feedthrough that is still present even after gradiometric coil attenuation. The Rx chain will also include various important amplification steps. Because MPI is not 'body-noise' dominant yet as in current MRI systems, these analog receive chain components generally dictate fundamental performance characteristics of an MPI system such as Signal-to-Noise ratio (SNR). Recent work in our group has explored pre-amplifier designs in MPI (Zhang et al., 2015), which is a challenging proposition due, for example, to the conflicting desire to optimally noise match and the need for broad band detection (a typical receive bandwidth of $\approx 30\ \mathrm{kHz} - 1\ \mathrm{MHz}$). As in the case of the Tx coils, the desire for multichannel Tx/Rx further complicates the design and implementation of MPI receiver systems.

Main Gradient Magnet Design

The main gradient magnets are a key feature of any MPI scanner. The strength of the linear gradient field establishes one limit on achievable resolution. Establishing large gradients on the order of 6 T/m over sizable FOVs is a challenging proposition. Some systems make full or partial use of strong permanent magnets to provide such gradient fields. Permanent magnets have the advantage of requiring no active power to establish the gradient field, with the downside of always being active and providing a static strength. Electromagnet systems, on the other hand, allow one to dynamically set the gradient strength or even modulate the gradient strength within a single scan. Designing these electromagnet systems, however, requires powerful and expensive amplifiers and mechanisms to dissipate heat generated in the coil. Power limits typically dictate the strength of the gradient. Recently our group has built electromagnets with iron return systems that tighten the return path of magnetic flux in the system compared to an air core magnet. One ramification is that we can attain the same gradient strength with much less power or, alternatively, provide higher gradient strengths with the same power.

As we seek to scale to clinical scanners in the future, a major question mark is how to scale these gradient systems. While advances to MPI tracer performance can provide much needed relief, optimal magnet design is an important open challenge for scaling up MPI. No doubt the field will see the evolution of innovative new magnet systems as we push forward. Perhaps we can learn from fields such as supercollider systems designers that have been solving similar challenges but in very different fields. And while superconducting magnets may be an option, this would add significant cost, complexity, and development time to MPI scanners. Perhaps the easiest and most prudent path is in improved tracer design.

Shielding, Power, and Bore Space

Assembling all of these components (and those not discussed, such as distinct shift field magnets) into a unified design is itself a major challenge. For example, it is critical to shield the Tx and Rx subsystems from the noisy gradient magnet and shift magnet systems and the saturating iron of the flux guide. But this must be done in an efficient manner that, for example, does not lead to dumping too much power from excitation coils into shielding subcomponents. Furthermore, efficient use of the bore space is crucial. We want our excitation coils to be as close to the FOV as possible for efficiency, but especially in the context of multiple Tx and Rx coils aligned with different directions, we must avoid geometric coupling between the various coils. This requires careful design and implementation of many complex resonant systems in a small space.

1.5.2 Resolution

Of all imaging parameters of merit, resolution is probably the most lacking in MPI. The fundamental Langevin physics, gradient strength limitations, and empirical limits of available tracers mean that ‘native’ resolution (without deconvolution) in current systems is around 1 mm. In a pre-clinical environment, we need sub-millimeter resolution to be compelling and fully competitive with other modalities, in terms of resolution. We know that tracer design will be a large part of this story, as early MPI-specific tracer design efforts have produced monodisperse tracers capable of resolution down to 600 microns (Ferguson et al., 2015) with our systems. This is only the tip of the iceberg, and as we will discuss, we now know there are tracer physics and new encoding techniques that we can harness to provide large improvements to resolution, well beyond our current capabilities.

1.5.3 Relaxation

Non-instantaneous magnetic response times of MPI tracers will be explored extensively in this work. At a basic science level, we still lack a fundamental understanding of tracer relaxation physics and behavior in MPI. Partly this is due to a unique intersection of high absolute fields, high field gradients, and tracers at the interface between multiple distinct physics. Currently, relaxation is known to limit achievable resolution in MPI but has also been recently leveraged to provide molecular imaging contrast. We will have much more to say about relaxation in subsequent chapters. Although tracer physics is somewhat complicated and incompletely understood, the upside is a huge potential font of rich physics to improve aspects of MPI in the future.

1.5.4 Image Contrast

A major proposition of MPI is to be a new and exciting molecular imaging modality. At first glance, use of a tracer that can be sensitive to local microenvironmental conditions means there should be many ways to encode physiologic and molecular imaging information into the modality. Initial work has provided glimpses of this, and in this work we seek to show just how rich the tailorable image contrast story is in MPI. Our efforts will primarily look at exploiting magnetic relaxation. But we note that our tracer may be functionalized for active targeting and long circulation, may be used in complex nanocarrier composites, and be used as a drug delivery vehicle. As a field, we need to reach out to tracer developers and materials scientists to help explore this hugely important part of MPI, especially as we are now mature enough as a field to have robust scanner implementations capable of exploring many different application areas.

1.5.5 Reconstruction Techniques

Currently, two image reconstruction paradigms dominate the field: the x-space approach and the system matrix approach. Both of these approaches are intuitively quite simple, and a basic implementation of each is straightforward. There remains, however, much room for improvement with these approaches. For example, all canonical MPI systems to date more or less ignore the possibility of temporally- and spatially-varying relaxation (system matrix) or relaxation altogether (x-space). First and foremost, significant relaxation currently degrades the quality of all canonical reconstruction methods. In a typical x-space approach, we reconstruct distinct images associated with each direction of the excitation waveform (negative and positive for the case of a constant direction excitation). Only one of these images is ever used, throwing out 50% of our information.

Because MPI is so young, the question of improving reconstruction is highly tied to acquisition and encoding strategies. We will show in part II of this work how changing the encoding strategy can completely change the story about how relaxation affects reconstruction. More generally, MPI has a rather elegant physics: Until you specify an MPI pulse sequence (or FFR trajectory), MPI is described by a tensor PSF. Recent work from our group has shown how serial scans with orthogonal excitation directions can sample different components of the base tensor PSF. The choice of receiver coil geometry and instantaneous excitation direction collapse the tensor PSF to a scalar PSF when scanning. We can combine data from orthogonal scans to provide isotropic resolution (which is unattainable with just one single direction excitation scan in typical x-space MPI). Furthermore, we can perform 'equalization' and other well-motivated postprocessing steps to improve image conspicuity based on a knowlege of the tensor PSF structure. We typically use only the most basic Cartesian scanning methods in x-space MPI today, but we believe innovative scanning trajectories in the future can improve imaging parameters of merit such as SNR-efficiency, physiologic contrast, and signal locality.

Part I

Molecular Imaging and Theranostics in Continuous Wave Magnetic Particle Imaging

Chapter 2

Relaxation Effects in Sinusoidal MPI

Magnetic relaxation dynamics influence current canonical MPI approaches by limiting achievable theoretical signal intensity and resolution and providing intriguing mechanisms for encoding contrast. In this chapter we explore relaxation in standard sinusoidal MPI primarily from an empirical perspective. We report results from testing a wide variety of MPI tracers across a large excitation parameter space. Subsequent analysis confirms and quantifies existing experiences in MPI, highlights a few novel behaviors, and allows us to explore possibilities for improving current MPI methods. In particular, this work provides a prelude for color MPI and relaxation imaging, discussed in detail in following chapters, which aim to exploit the unique information encoded in relaxation behavior.

2.1 Introduction and Background

The canonical MPI signal and imaging equations are derived with the assumption of instantaneously-achieved, steady-state physics (Goodwill and Conolly, 2011; Rahmer et al., 2012; Lu et al., 2013). Concretely, this is equivalent to assuming that some temporal parameter of merit that quantifies or bounds tracer magnetic dynamics is always negligibly small compared to the slew rates used in MPI excitation. Magnetic excitation is thus 'slow' with respect to magnetic relaxation. For a typical sinusoidal excitation frequency of around $25\ \mathrm{kHz}$, the excitation period, T, is $40\ \mu\mathrm{s}$. For steady-state models to accurately describe the evolution of the signal during the MPI scan, we must have $\tau_{\max} \ll 40\ \mu\mathrm{s}$ where $\tau_{\max}$ is a conservative upper bound time constant that characterizes the magnetic dynamics of the tracer.

In practice, this ideal condition is not generally met, even considering the most commonly used MPI tracers. Nonetheless, if the ratio of $\tau_{\max}$:

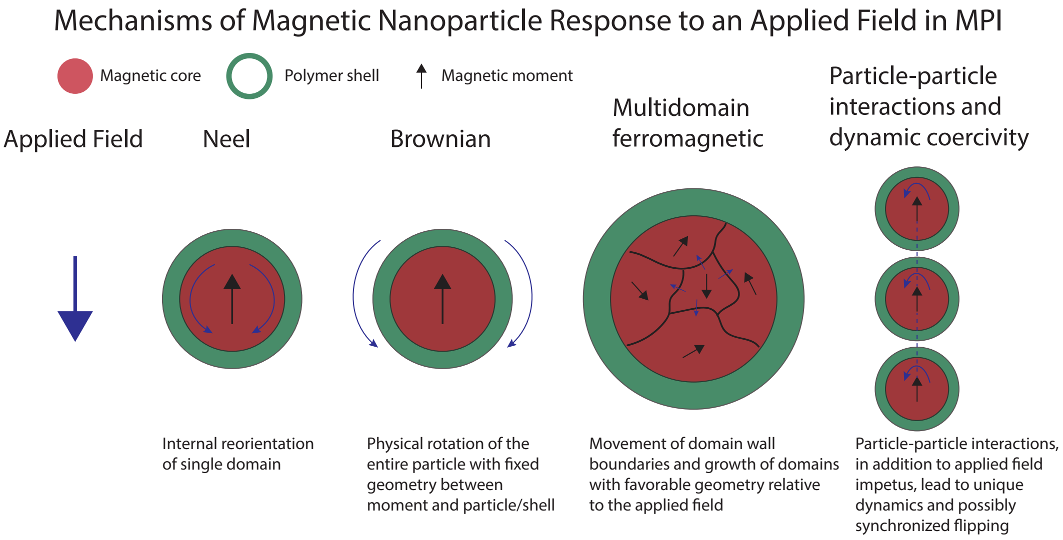

FIGURE 2.1: Physical mechanisms for magnetic nanoparticles to respond to applied fields in the context of MPI. These include Néel relaxation where a single magnetic domain internally rotates with respect to the particle shell, Brownian relaxation where the entire particle physically rotates maintaining a fixed geometry between the magnetic moment and the shell, multidomain ferromagnetism that involves internal domain wall movement and domain growth, and the case where particle-particle interactions affect relaxation, *e.g.*, particles closely adjacent along a line potentiate the response to an applied field.

T is not vanishingly small but still below a less conservative limit, there may be a modest and tolerable effect on the MPI signal and image domain resolution. Indeed, this may be desired to encode valuable physiologic information. However, as τ_{max} increases further, catastrophic signal loss and image degradation will eventually occur. Magnetic relaxation is a strong function of tracer core size, thus, while steady-state Langevin theory predicts ever improving resolution with core size, the influence of dynamic processes mean that relaxation will become a road block at some point. In this chapter we explore these and other effects of magnetic relaxation on the canonical continuous wave (sinusoidal excitation) MPI process from a phenomenological and experimental perspective.

In the remaining sections of this chapter, our objective is to quantify magnetic relaxation and the impact of this relaxation on MPI image domain parameters of interest such as signal intensity and resolution. We accomplish this by performing a large number of MPI experiments, in which we vary excitation parameters such as amplitude and frequency, across a large number of MPI tracer types with different properties such as magnetic core size and polydispersity, magnetic core material, and shell properties. This allows us to quantify MPI relaxation over a broad parameter space and better informs how we may best exploit MPI relaxation behavior, as explored in subsequent chapters.

2.1.1 Origins of Magnetic Relaxation

Assessing magnetic relaxation in MPI is challenging due to significant complexity in the dynamic physics of magnetic nanoparticles generally, and in the context of MPI specifically. There are multiple distinct mechanisms a given tracer may employ to align with an applied field, including Néel relaxation, Brownian relaxation, ferromagnetic relaxation, and relaxation affected by particle-particle interactions as depicted in Fig. 2.1 (Neel, 1949; Brown Jr, 1963; Raible and Engel, 2004; Spaldin, 2010). The specific quantitative nature of these mechanisms, as well as the dominant mechanism observed at a given point in time, depends on a large number of factors including properties of the tracer such as magnetic core size, shape, and shell thickness, and external factors such as the instantaneous strength of the local magnetic field, environmental viscosity, and temperature. The ferrofluids literature has explored very similar questions and is generally a good reference point for the interested reader (Shliomis, 2002).

Several aspects of MPI further complicate complex parametric dependencies in magnetic tracer dynamics. First, all tracers used in MPI are characterized by some degree of polydispersity in their properties (the canonical MPI tracer Resovist is notably polydisperse in core size (Eberbeck et al., 2013)), although recent progress has been made in monodisperse MPI tracers (Ferguson et al., 2015). Second, the strong dependence of tracer dynamics on applied field strength and the strong spatial gradient fields employed in MPI mean that the dominant physics and its character depend on the relative geometry between the MPI FFR and the tracer under consideration at a given point in time. Our use of inductive receiver coils then spatially integrates these spatially-varying dynamics at all times. Finally, tracer at a given location and at a given time may be responding through a complex combination of one or more of these relaxation phenomena.

Models employed to describe magnetic relaxation are often simplified and ignore many of the aforementioned details, assess averaged or expectation variables, or are explicitly phenomenological (Croft et al., 2016; Dhavalikar and Rinaldi, 2016; Shliomis, 2002; Deissler, Wu, and Martens, 2014). Sometimes this is deemed a practical requirement, but as engineers often leverage, even greatly simplified models may well describe observed phenomena and have predictive power in certain constrained parameter spaces.

MPI is in an early stage of development, however, and many novel methods are yet undiscovered. To fully exploit tracer physics in this context, we must keep these complexities in mind and endeavor to bring them into our models when and as appropriate. New tools such as our recently developed arbitrary waveform relaxometer (AWR) (Tay et al., 2016) and continual theoretical and empirical work are making this more and more possible (Deissler, Wu, and Martens, 2014; Croft et al., 2016; Dhavalikar and Rinaldi, 2016).

2.1.2 First Order Model

The simplest model of magnetic relaxation is to ignore it and assume instantaneous physics. We state this case for two reasons: (1) it is the basis for all 'non-zero' relaxation models and captures the basic spatial encoding of information in MPI and (2) it is the current basis for fully developed signal and imaging equations in MPI.

Ignoring time-dependent relaxation, we arrive in the Langevin physics regime. The Langevin equation is an expression of steady-state energy balance, describing how the interplay between magnetic torque and thermal agitation leads to average magnetization behavior (of an ensemble of particles, or any number of particles observed over a 'long' period of time) as a function of applied field (Neel, 1949; Brown Jr, 1963). This theory is extremely accurate for steady-state conditions:

$$\mathbf{M}_{\text{ss}}(t, \mathbf{x}) = m\rho(\mathbf{x})\mathcal{L}(k\|\mathbf{H}(t, \mathbf{x})\|) \cdot \hat{\mathbf{H}}(t, \mathbf{x}) \tag{2.1}$$

where $\mathbf{M}_{\text{ss}}$ is the steady-state, dynamic equilibrium magnetization in $[\mathrm{A\,m^{-1}}]$ of an ensemble of magnetic nanoparticles distributed over space $\mathbf{x}$, m is the magnetic moment of a single particle in $[\mathrm{A\,m^{2}}]$, ρ is the density of magnetic nanoparticles in $[\mathrm{particles\,m^{-3}}]$, $\mathcal{L}$ is the Langevin equation, k in $[\mathrm{m\,A^{-1}}]$ is a tracer-specific parameter that characterizes the relationship between magnetic and thermal energies, and $\hat{\mathbf{H}}$ is the unit vector pointing in the direction of the applied field $\mathbf{H}$ in $[\mathrm{A\,m^{-1}}]$.

Accounting for the details of non-equilibrium dynamics is non-trivial from a first-principles perspective. An alternative approach is to develop phenomenological equations that well approximate the real dynamics under certain circumstances. One such option is the family of first-order models, as in a Debye process (Debye, 1954) or the Bloch equations (Bloch, 1946) of MRI. A Debye model has been described by Goodwill and Croft (Goodwill et al., 2011; Croft et al., 2016), similar to the following relation:

$$\frac{\mathrm{d}\mathbf{M}(t,\mathbf{x})}{\mathrm{d}t} = \frac{\mathbf{M}_{\mathrm{ss}}(t,\mathbf{x}) - \mathbf{M}(t,\mathbf{x})}{\tau(t,\mathbf{x})} \tag{2.2}$$

where deviation from the steady-state, ensemble magnetization $\mathbf{M}_{\mathrm{ss}}(t,\mathbf{x})$ described by equation (2.1) is captured by a first order relation and a time constant $\tau(t,\mathbf{x})$. In equation (2.2), we have left the possibility of a temporally and spatially varying τ explicit. We can further simplify this equation if we assume that $\tau(t,\mathbf{x}) \approx \tau_{\mathrm{eff}}(\mathbf{x})$ is stationary in time or both time and space (alternatively one can consider such a term to describe behavior in aggregate): $\tau(t,\mathbf{x}) \approx \tau_{\mathrm{eff}}$. In this case of a single scalar time constant, equation (2.2) would have a simple exponential solution if the steady-state magnetization were also constant in time. However, in traditional x-space Magnetic particle imaging (MPI), this term is not stationary in time as our excitation is varying sinusoidally. In the case of such a 'moving steady-state', solving for the magnetization yields a temporal convolution:

$$\mathbf{M}(t,\mathbf{x}) = \mathbf{M}_{\mathrm{ss}}(t,\mathbf{x}) \overset{t}{*} \tau_{\mathrm{eff}}^{-1}\exp(-t/\tau_{\mathrm{eff}}) \tag{2.3}$$

where we can fit τ_{eff} from the data. In Section II, chapter 6, we describe the case of spatially-dependent relaxation times in detail. There we employ specific models of the field-dependence of various relaxation mechanisms, which maps to a spatial dependence in MPI due to the presence of strong gradient fields. In this chapter, however, we will focus on experimental data and not consider more sophisticated models.

The phenomenological expression (2.3) can be useful for modeling some effects of magnetic relaxation in MPI, but is of limited utility for uncovering the specific physical nature of the relaxation. The simplified Debye model states that magnetic relaxation can be described as an exponential kernel that blurs the ideal magnetization signal over time. Due to the direct mapping of the time domain signal to the image domain in x-space MPI, such a temporal blur maps directly to an image domain blur. Thus, in addition to time delays detectable in the received time domain signal, we would expect image domain implications such as shifts, increased FWHM values, and decreases in peak signal/SNR. Here, we will empirically measure and explore these effects across a wide variety of MPI tracers.

2.2 Materials and Methods

Numerous tracer samples were tested under a variety of conditions on our arbitrary waveform relaxometer (AWR) (Tay et al., 2016). Just under 1500 individual experiments were run, emphasizing the high throughput and flexibility provided by the AWR.

2.2.1 Tracer Collection and Attributes

To empirically catalog the effect of magnetic relaxation in MPI, a wide ranging collection of 12 MPI-compatible tracers was tested with our AWR. The collection included 'traditional' tracers such as nanomag MIP (micromod Partikeltechnologie GmbH, Germany), University of Washington monodisperse tracers, and the gold standard Resovist; monodisperse tracers from the same commercial vendor but varying in core size (18, 21.6, 23.6, 24.4, 25, 27.4, and 32.1 nm – Imagion Biosystems, Inc., Albuquerque, NM); and nonstandard tracers such as cobalt-ferrite SPIOs from the University of Florida, elemental iron tracers from the University of California, Davis, and polydisperse tracers embedded in aggregated albumin from Universidade Federal de Goiás. Table 2.1 summarizes the properties of the tracers we tested. All tested samples consisted of 40 μL of tracer at the supplied concentration (typically in the range [5, 20] $\mathrm{mg\,mL^{-1}}$). This diverse set of tracers provides a sampling of the tracer-specific factors that affect magnetic relaxation including magnetic core size, polymer shell thickness and composition, magnetic material, and polydispersity.

2.2.2 The AWR and Excitation Conditions

For this work, we heavily leveraged our x-space MPI AWR, pictured in Fig. 2.2(a). The AWR is the most recent addition to our family of x-space relaxometers. These tabletop MPI devices do not generally contain gradient field-inducing components. Instead, they are typically constructed with an AC excitation system, a bias field system, and a gradiometric receiver coil all aligned with sensitivity in the same principal axis. Lacking a gradient field, the sample is effectively collapsed to a point source, as all components of the sample feel the same applied field at any point in time. Accordingly, we can acquire data to reconstruct a 1D point-spread function (PSF) by applying the AC excitation while the bias field slowly ramps across a magnetic field of view (FOV). Because we use linear gradient fields in MPI, the PSF

Tracer	**Core**	**Comments**
Ferucarbotran (Resovist)	≈ 17 nm (polydisperse)	De facto standard MPI tracer.
Nanomag MIP	≈ 21 nm (polydisperse)	Commercially available MPI-optimized polydisperse tracer similar to Resovist, with noticeably improved resolution.
UW	25 nm (monodisperse)	Highest performing monodisperse tracer in MPI.
Imagion	21.6, 23.6, 24.4, 25.0, 27.4, 32.1 (monodisperse)	Commercial monodisperse tracers available with various functionalized moieties such as carboxylic acid, oleic acid (in hydrophobic solvent), PEG, and streptavidin.
UF CoFe	21 nm (monodisperse)	Cobalt-ferrite tracers, guaranteed to have only Brownian magnetic behavior (no Néel).
UC Davis Fe	≈ 20 nm (monodisperse)	Prepared with elemental iron core and maintained under nitrogen atmosphere until testing.
MMA-SPIOs	polydisperse	Polydisperse SPIOs embedded in macro-aggregated albumin.

TABLE 2.1: Summary of experimentally tested MPI tracers. The list includes canonical tracers such as Resovist, nanomag MIP, and tracer from the University of Washington, monodisperse tracers that sample a range of core sizes from Imagion Bioysystems, and non-standard tracers such as cobalt-ferrite particles from the University of Florida, elemental iron core tracers from UC Davis, and macro-aggregated albumin, polydisperse tracer from Universidade Federal de Goiás.

reconstructed over this magnetic FOV is equivalent to a 1D imaging PSF after dividing by the appropriate gradient strength (Tay et al., 2016). One way of seeing this is to consider the signal from the AWR when the slowly varying bias field is equal to -6 mT and 0 mT. Let us assume a scanner gradient strength of 6 $\mathrm{T\,m^{-1}}$, then the former is equivalent to an imaging scenario in which an MPI field-free region (FFR) is shifted to $-6\,[\mathrm{mT}]/6\,[\mathrm{mT\,mm^{-1}}]$ = $1\,[\mathrm{mm}]$ and $0\,[\mathrm{mT}]/6\,[\mathrm{mT\,mm^{-1}}]$ = $0\,[\mathrm{mm}]$ away from a point source, along the axis of the excitation direction. Our previous work has experimentally verified this equivalence (Tay et al., 2016).

In contrast to the effort and time required to acquire an image of a PSF on a scanner, it takes on the order of a second to acquire and reconstruct the data associated with a single AWR PSF scan. Furthermore, the AWR, unlike our previous x-space relaxometers and related MPI spectrometers (Biederer et al., 2009), is constructed with an untuned excitation system. This allows us to choose any arbitrary excitation frequency and amplitude we desire, subject to our power constraints. We have built a modular software stack around this system that allows us to exploit these properties to automate the collection of parametrically varying scans. The result is an incredibly fast, high throughput, and flexible system for characterizing MPI tracers. The ability to directly measure 1D MPI PSFs in this manner is critically enabling for MPI tracer development and testing new signal encoding and reconstruction strategies.

Here we demonstrate these capabilities by describing the results from approximately 1500 individual experiments. To assess the effect of MPI excitation parameters such as excitation frequency and amplitude, our collection of tracers was subjected to AWR experiments at multiple frequencies (3, 10, 17, 25, 32, 40, 65, 85, 100, 150) kHz and multiple amplitudes (10, 20, 30) mT. Due to power limitations, permutations of the highest amplitudes and highest frequencies were not possible, and each amplitude supported a different number of the frequencies in the listed range. Three replicate trials were performed per experiment to assess statistical variation in the scanning and reconstruction process.

2.2.3 Reconstruction and Data Analysis

Standard x-space reconstruction was applied to the acquired data (Goodwill and Conolly, 2010; Tay et al., 2016) using in-house Python scripts to reconstruct 1D point-spread functions (PSFs) for each sample. Peak signal and FWHM resolution can be directly measured from the resulting PSFs. As depicted in Fig. 2.2(c,d), in the standard x-space reconstruction, distinct

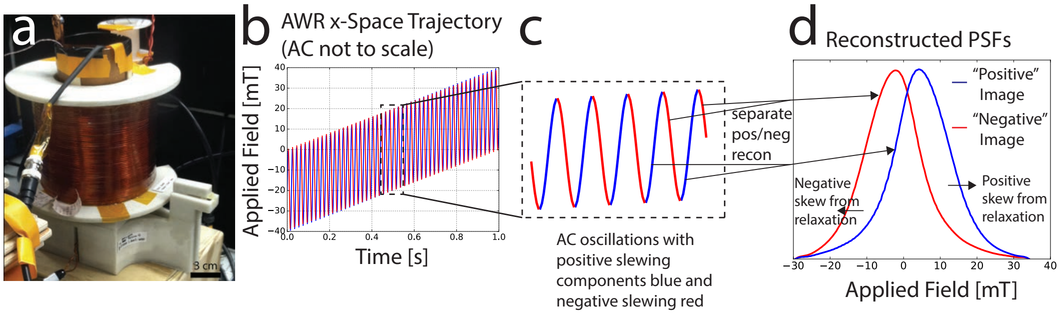

FIGURE 2.2: Relaxation effect is directional with Tx direction. The delay, skew, or blur associated with finite magnetic relaxation in MPI tracers occurs in the direction of the excitation. (a) Photograph of our MPI AWR. (b) Depiction of a typical 1D AWR x-space scanning trajectory. (c) A region of the trajectory showing successive negative and positive slewing components. The signal associated with these successive components of the excitation are reconstructed separately yielding 'negative' and 'positive' images. (d) Experimental AWR data showing the directional effect of relaxation on the positive and negative PSFs.

images are reconstructed from positive-slewing and negative-slewing components of the excitation trajectory. This is done to prevent mixing of slew direction-dependent relaxation effects. As equation (2.3) implies and experimentation has borne out, relaxation leads to an asymmetric skew/blur in the direction of the excitation. Fig. 2.2 illustrates the nature of this asymmetric skew. Without separating data associated with different excitation directions, the shifts and directional blurs would average when gridding data to the image domain. The effect would be to average the two PSFs shown in Fig. 2.2(d), yielding much worse resolution than that obtained by simply separating the data.

In addition to PSF reconstruction, we constructed relaxation maps to quantify magnetic relaxation in more detail. This was accomplished by calculating the relative time delay between the signal associated with adjacent pairs of negative- and positive-slewing regions of the excitation trajectory. We call these regions partial fields-of-view (pFOVs). Standard cross-correlation was performed on successive pFOV pairs after time-reversing and negating one of the pFOVs. Due to the symmetry and alternating polarity of sinusoidal excitation, time-reversed and negated pFOV pairs should be identical in the absence of relaxation or other symmetry-breaking and delaying effects. In this manner a single time delay value can be computed for each pair and gridded to the mean location of the trajectory (field, or position in the case of an MPI imaging system with gradients). This value is an average (in a nonlinear manner) value of the relaxation delays across the pFOV and is thus a coarse indication of time-domain lag caused by magnetic relaxation as a function of applied field strength. A Savitzky-Golay filter was applied to raw relaxation maps to smooth out regions of rapid variation, typically observed at the high field edges of some of the relaxation maps where the SNR is poor. Due to analysis on adjacent pFOVs, each with a total time equal to half of the excitation period T, and the cross-correlation method used, the technique can record a maximum time delay value of $T/4$, or a maximum phase delay of $\pi/2$.

The coarseness of the relaxation map is a direct function of the size of the pFOV which is determined precisely by the amplitude of the excitation sinusoid. Lower amplitudes provide more resolved relaxation maps. Furthermore, we typically divide the calculated time delays by the excitation period to report relaxation maps in units of phase delay. Directly reporting time delays can be misleading if it is interpreted as a fundamental map of relaxation time constants vs. field strength or location, for example, because the specific time delay is highly dependent on the excitation frequency in these non-steady-state excitation conditions. Nonetheless, we believe that these phase delay relaxation maps can be thought of, in some senses, as a magnetic relaxation PSF that report the underlying field-dependence of magnetic relaxation. Relaxation map construction and the data analysis required to produce plots of MPI data were performed with Python scripts.

2.3 Results

We first show empirical demonstration of the MPI 'relaxation wall' preventing Langevin improvement in resolution with larger tracer core sizes. We then provide figures summarizing the effects of magnetic relaxation on the image domain parameters of merit FWHM resolution, peak signal, and image shift across the spectrum of tested tracers and excitation condition. We conclude our results with quantification and comparisons of relaxation maps.

2.3.1 The Relaxation Wall

Fig. 2.3 shows resolution results as a function of magnetic core size, using samples from Imagion Biosystems. These results quantify and summarize a trend we have seen in many tracers and which we refer to as the 'relaxation wall'. Smaller tracers have relatively worse resolution, as predicted by Langevin physics. As one increases size from approximately 18–25 nm,

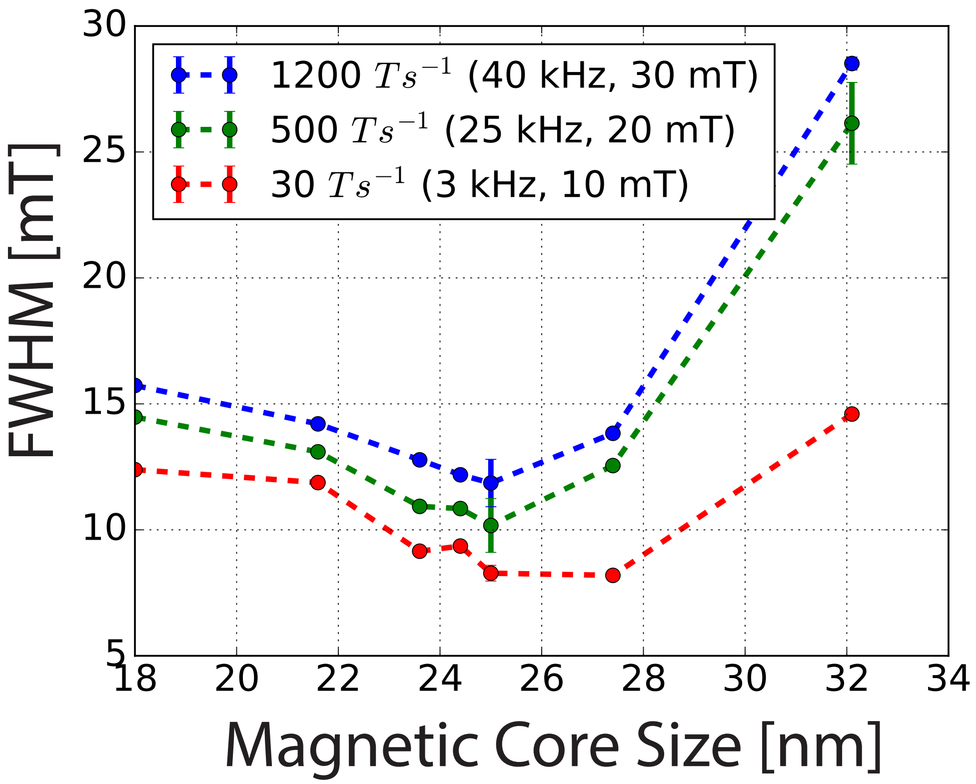

FIGURE 2.3: Magnetic relaxation limits MPI resolution. The steady-state Langevin theory teaches that increasing magnetic core size will improve MPI resolution (as well as peak signal). Experimentally, we observe this relationship in the range of approximately 18 – 25 nm. Larger tracers have significant magnetic relaxation dynamics that induce a secondary blur. The data shows the effect of modifying the excitation slew rate. Even a very low slew rate (compared to canonical systems) of 30 $\mathrm{T\,s^{-1}}$ (3 kHz, 10 mT) shows loss in resolution with larger core sizes, although the effect is somewhat mitigated and delayed.

resolution improves but then degrades again after about 25 nm. The latter effect, in conflict with Langevin predictions, is the result of secondary relaxation blurring effects. The three curves in Fig. 2.3 were taken with excitation slew rates of 30 $\mathrm{T\,s^{-1}}$, 500 $\mathrm{T\,s^{-1}}$, and 1200 $\mathrm{T\,s^{-1}}$. While overall resolution is improved with lower slew rates generally, the same trend with a noticeable minimum resolution at intermediate core sizes is apparent in all three cases. The lowest 30 $\mathrm{T\,s^{-1}}$ case appears to also shift the minimum point to a slightly higher core size.

2.3.2 Parametric Evaluation of Relaxation Image Domain Effects

FWHM Resolution

Fig. 2.4 shows 1D image domain FWHM resolution as a function of both excitation frequency and amplitude for a selection of 9 of the tracers used in this study. In the top row (Fig. 2.4(a)), well known MPI tracers are shown, in the middle row (Fig. 2.4(b)), monodisperse tracers from Imagion Biosystems at three different core sizes are shown, and in the bottom row (Fig. 2.4(c)), non-standard MPI tracers are shown.

The standard MPI tracers show typical dependencies in Fig.!2.4(a): worsening resolution as both frequency and amplitude increase, with a greater relative sensitivity to amplitude in terms of the typical parametric ranges available in canonical MPI. Minor curvature differences between the three tracers are apparent and, as has been observed previously, the UW tracer shows the best overall resolution, followed by Nanomag MIP, and then Resovist. Nanomag MIP appears to show a greater sensitivity to excitation amplitude than Resovist or UW.

The monodisperse tracers from Imagion Biosystems show behavior in Fig. 2.4(b) quite distinct from any other tracers. This unique behavior is characterized by an initial increase in FWHM resolution as a function of frequency as typically observed, followed by a decrease and plateauing of FWHM vs. frequency after an inflection point located around 5 – 20 Hz. This pattern is observed for all core sizes and all amplitudes. As in the case of the other tracers, FWHM resolution does worsen with increasing amplitude, however.

The non-standard tracers M2 macro-aggregated albumin and Davis elemental iron tracers also exhibit the typical relationships of generally worsening resolution with increased amplitude and frequency. The M2 tracer generally performs very well in terms of resolution, second only to UW,

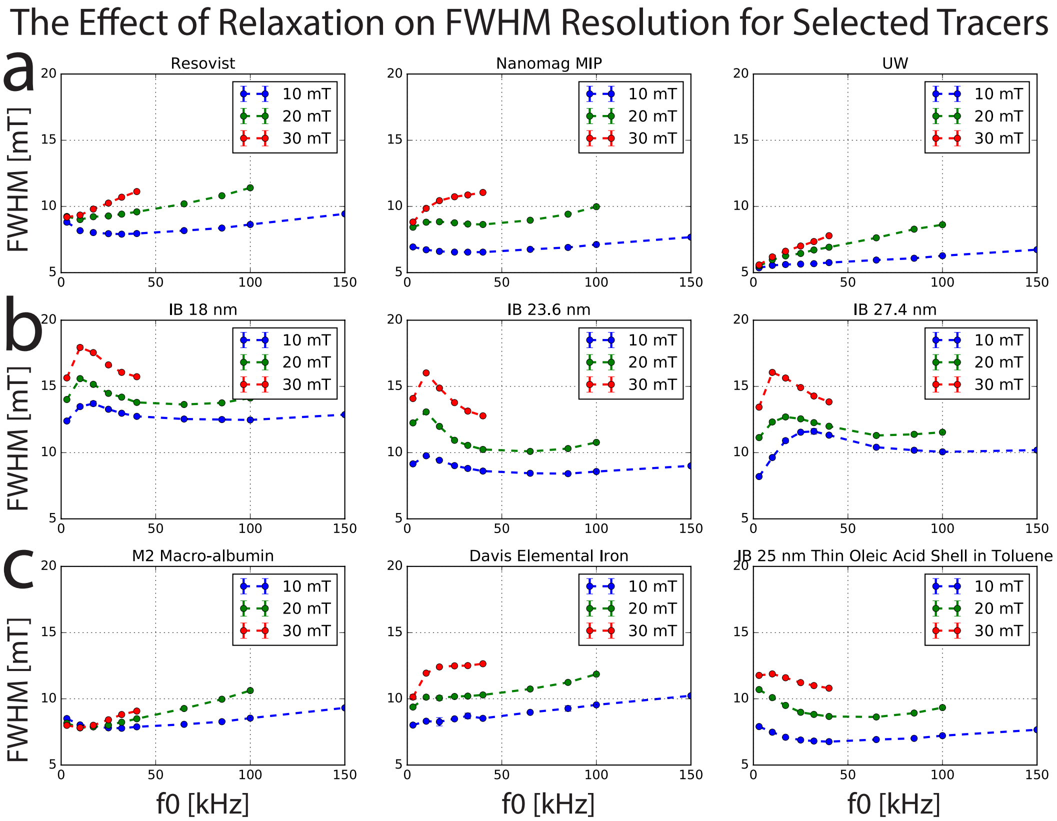

FIGURE 2.4: MPI experimental FWHM as a function of excitation frequency and amplitude for selected tracers. (a) Canonical MPI tracers (Resovist, nanomag MIP, and tracers from the University of Washington) show mostly monotonically worsening resolution with respect to both frequency and amplitude. (b) Three different core sizes of monodisperse tracers from Imagion Biosystems show a unique inflection point in FWHM vs. frequency at around 10 kHz, after which FWHM improves or plateaus with increasing frequency. FWHM resolution does degrade monotonically with amplitude. (c) Non-standard MPI tracers including polydisperse iron oxides macro-aggregated with albumin, a tracer with partial elemental iron core, and monodisperse tracer from Imagion Biosystems without a large hydrophilic shell and suspended in toluene.

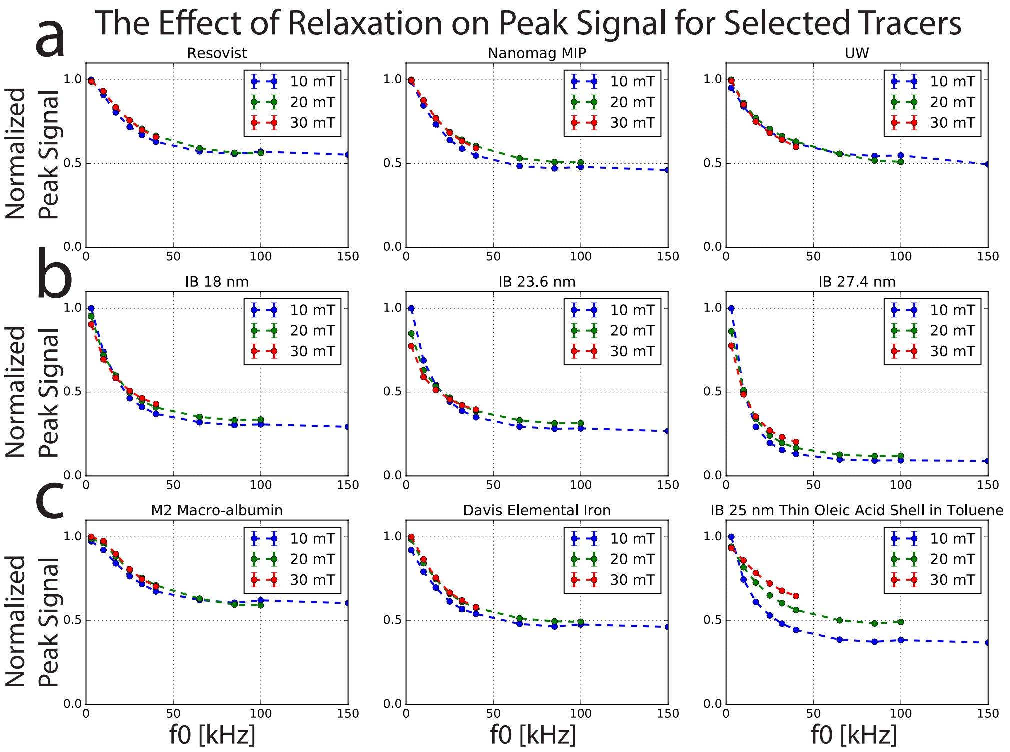

FIGURE 2.5: Relative peak signal as a function of excitation frequency and amplitude for selected tracers. All tracers exhibited monotonic loss of peak signal with increasing frequency but negligible, minor, or mixed relationships with amplitude. (a) Canonical MPI tracers (Resovist, nanomag MIP, and tracers from the University of Washington). (b) Monodisperse tracers from Imagion Biosystems generally show more loss in peak signal as a function of frequency. Larger core tracer shows significantly more loss and at lower frequencies than the smaller core tracers. (c) Non-standard MPI tracers (iron oxide macro-aggregated with albumin, partial elemental iron core, monodisperse core with thin oleic acid shell in toluene).

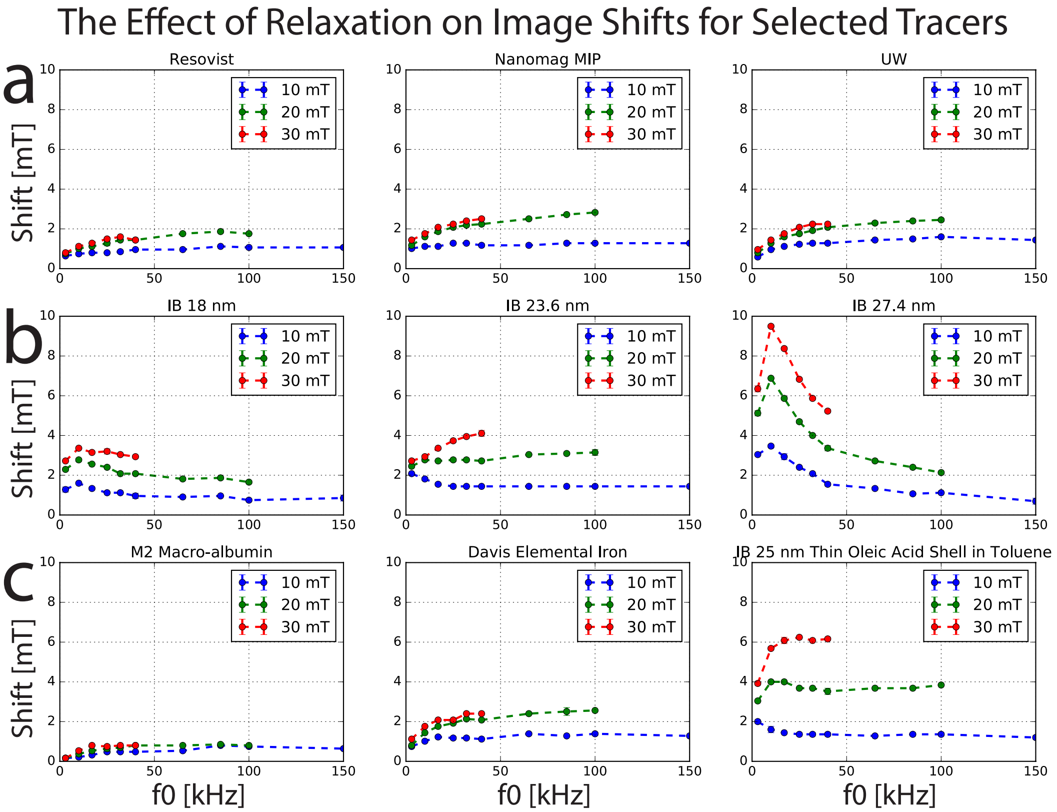

FIGURE 2.6: MPI image domain shifts as a function of excitation frequency and amplitude for selected tracers. Similar to FWHM variation, most tracers show monotonically larger shifts with increasing frequency and amplitude. (a) The canonical tracers (Resovist, nanomag MIP, and tracers from the University of Washington) show similar results with less than 3 mT shifts. (b) Monodisperse tracers from Imagion Biosystems generally show larger shifts and inflection points at low frequencies, after which the shift decreases with increasing frequency. (c) Two of the non-standard tracers (iron oxide macro-aggregated with albumin and tracer with partially elemental iron core) show typical behavior while the monodisperse tracer with only a thin oleic acid shell in toluene shows a very large sensitivity to amplitude.

and shows relatively small variations in resolution with respect to amplitude and frequency. The Davis elemental iron tracer, like nanomag MIP, shows a relatively strong sensitivity to excitation amplitude. Finally, the oleic acid-only, toluene-solvated 25 nm monodisperse tracer from Imagion Biosystems exhibits similar improving resolution with increasing frequency as in the hydrophilic tracers of Fig. 2.4(b), although no clear inflection point is seen, at least for the cases of 20 and 30 mT excitation. It is quite possible that an inflection point is apparent at a lower frequency than that tested in these experiments.

Peak Signal Intensity

Fig. 2.5 shows 1D image domain peak signal intensity as a function of both excitation frequency and amplitude for the same selection of 9 tracers. The data is first normalized for the excitation slew rate associated with each data point (amplitude * f_0) to normalize for the expected scaling due to Faraday's law of induction and our use of inductive receiver coils. The normalized values are then normalized to the maximum signal intensity for straightforward relative analysis of peak signal changes.

All tracers exhibit gradual loss in relative peak signal with increasing frequencies. The effect of amplitude is relatively much smaller and has no clear trend. Some tracers exhibit minor increases in peak signal with increasing amplitude (*e.g.*, Resovist and Nanomag MIP), others the opposite (*e.g.*, UW), and the monodisperse tracers from Imagion Biosystems show a transition between these two conditions at different frequencies: At lower frequencies, larger amplitudes are associated with lower signal while the opposite is true at higher frequencies.

Most tracers show a reduction in Faraday-normalized peak signal intensity of about 50% at 100-150 kHz from a peak at 3 kHz. The M2 macroaggregated albumin and Resovist tracers show the smallest overall reductions in peak signal, both remaining above 50% loss from the maximum in the tested range. As shown in Fig. 2.5(b) monodisperse tracers from Imagion Biosystems, however, generally exhibit higher signal losses and clear core size-dependent behavior. The 27.4 nm tracer shows large drops in peak signal intensity, below approximately 75% of the 3 kHz maximum by 25 kHz. The 18 nm tracer shows the lowest overall loss in peak signal of the group and the least sensitive rate of loss with increasing frequency. The 23.6 nm tracer is intermediate, but very close to 18 nm in behavior such that a very large change between 23.6 and 27.4 nm is apparent. This is consistent with greatly increased lossiness due to increasingly prominent relaxation behavior as core size increases.

Image Domain Shifts

Fig. 2.6 quantifies image domain shifts as a function of excitation amplitude and frequency. Magnetic relaxation induces a temporal blurring of the ideal signal which maps to a spatial blur in the direction of excitation. This induces asymmetric skew in the image domain, and one way of quantifying this phenomena is through the apparent peak shift. Almost all tracers show increasing shift with both increasing frequency and amplitude. The amount of shift and sensitivity of the shift value with respect to the excitation parameters, however, varies largely between tracers. The M2 macroaggregated albumin and Resovist tracers are associated with the smallest shift values, similar to the case with peak signal loss see in Fig. 2.5.

As shown in Fig. 2.6(b,c), the monodisperse tracers from Imagion Biosystems generally show more significant shifts than the other tracers. The 27.4 nm tracer in particular, shows a unique trend of increasing shift with frequency at low frequencies followed by decreasing shift values at larger frequencies, with an inflection point around 10 kHz. These unique results mirror the same type of relationship observed in the FWHM relationships of Fig. 2.4.

2.3.3 Relaxation Maps

Fig. 2.7 quantifies and explores magnetic relaxation effects in detail through relaxation maps for two tracers: 18 and 32 nm Imagion Biosystems monodisperse tracers. Fig. 2.7(a,d) plot phase relaxation maps for the 18 nm and 32 nm tracers, respectively, for the three different amplitudes tested, all at a fixed frequency of 3 kHz. In all cases, we see symmetric peaks centered around 0 mT applied field. The 18 nm relaxation maps are significantly lower in magnitude than the 32 nm relaxation maps, both in peak phase delay and at all other points. Furthermore, the minimum phase delay observed at the highest applied field values for the 32 nm relaxation map is larger in magnitude than the peak phase delay in the 18 nm relaxation map. The peak phase delay is generally higher for lower excitation amplitudes, although this may be solely due to the increased averaging/blurring of the relaxation maps associated with larger amplitudes due to the larger pFOVs.

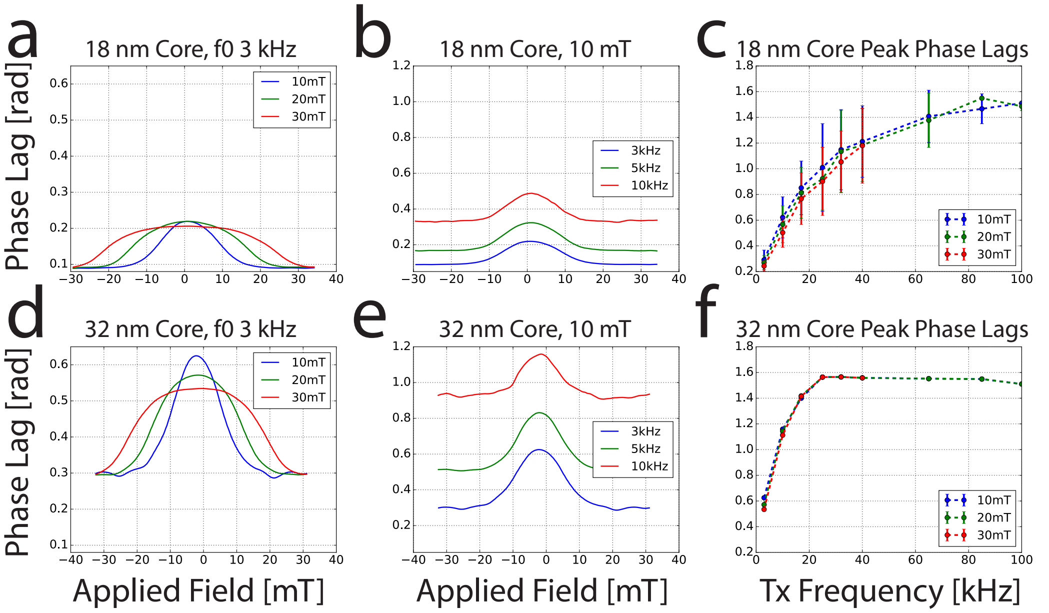

FIGURE 2.7: Quantifying time domain magnetic relaxation effects in MPI. (Top) Results for monodisperse 18 nm tracer. (a) Phase relaxation maps – phase delay as a function of applied field for excitation at 3 kHz and 10, 20, and 30 mT. (b) Phase relaxation maps for excitation at 3, 5, and 10 kHz and 10 mT. (c) Peak phase lag as a function of frequency and amplitude. (Bottom) The same results for monodisperse 32 nm tracer. Per the pFOV-specific nature of the algorithm, increasing amplitude acts like a rectangular blur applied to the phase map. In both cases, increasing frequency leads to increased phase lag, with the rate of increase and total phase lag much larger for the larger 32 nm tracer. Excitation amplitude has relatively little effect on peak phase lag.

Fig. 2.7(b,e) plot phase relaxation maps for a fixed excitation amplitude of 10 mT and at the frequences 3, 5, and 10 kHz. As before, the 32 nm relaxation maps are high in magnitude across the board compared to the 18 nm case. In both cases, increased frequency leads to significantly higher phase delay. Furthermore, the 32 nm tracer shows larger increases in phase delay with increases in frequency compared to the 18 nm tracer.

Fig. 2.7(c,f) plot peak phase lag as a function of frequency and amplitude for both tracers. The 32 nm tracer clearly shows much more significant phase lag that rises to the $\pi/2$ maximum (per our cross correlation method) by approximately 20 kHz. The 18 nm tracer, in contrast, shows a much gentler rate of phase lag change with excitation frequency, reaching $\pi/2$ peak phase lag only near 100 kHz. These data also show how the phase lag changes very little with excitation amplitude, in contrast to the dependence on frequency.

2.4 Discussion

2.4.1 The Relaxation Wall

The availability of highly monodisperse tracer at seven distinct core sizes between 18 and 32 nm, all from the same supplier and subject to the same synthesis techniques, allowed us to gather robust data demonstrating and quantifying a 'relaxation wall' effect in standard MPI. This phenomenon, whereby resolution degrades after approximately 25 nm in core size, directly contradicts the canonical Langevin theory. Furthermore, our analysis across different excitation conditions and slew rates indicates that this phenomenon cannot be easily prevented by simply manipulating two of the chief excitation sequence design levers: amplitude and frequency. Our data show that lower slew rates can shift the curve and generally improve resolution across the board, but the relative worsening of resolution with increasing core size is observed even for an excitation slew rate of 30 $\mathrm{T\,s^{-1}}$ – over an order of magnitude reduced from our typical scanner operating conditions. Reducing the excitation slew rate in such a manner also greatly reduces achievable SNR due to our inductive signal reception. Furthermore, reducing the excitation amplitude increases the relative amount of signal energy at the fundamental harmonic. Because this information is lost in all canonical scanner designs, it must be recovered in the reconstruction process. Reducing the amplitude is thus limited due to eventual ill-conditioning and destabilization of MPI reconstruction algorithms.

The relaxation wall of MPI is an important topic for the field to understand, as imaging resolution is arguably one of the weakest attributes of MPI. We fully expect this relaxation wall is in no way fundamental. We believe our current exploitation of the physics of MPI tracers is in its infancy and that future methods of signal encoding and excitation trajectories as well as tracer design will move us beyond this current limitation. A better understanding of tracer relaxation is crucial to these endeavors. Previous MPI literature has often focused on the distinction between Néel and Brownian relaxation mechanisms. We believe expanding this to include other mechanisms such as ferromagnetism, and particle-particle interactions is important. Furthermore, while theoretical models of Néel and Brownian physics have been explored significantly in the literature, fully understanding MPI dynamics will require placing and vetting these models within the complex MPI context in which strong total applied fields and gradient fields are rapidly varying.

So far, it has been difficult to fully explain the raw MPI signal using first-principles models, and understanding how the multiplicity of possible relaxation dynamics interact is a part of this. Even if we have complete and robust models for field-dependent Néel and Brownian relaxation separately, it is unlikely that the result of tracers exhibiting both mechanisms manifests as a simple orthogonal or parallel combination of the individual models. Finally, the heterogeneous conditions of any individual MPI scan and the polydispersity in tracer properties such as core and shell size further complicate this picture.

In this context, the challenges to fully understanding magnetic relaxation in MPI are legion. But we find this an exhilarating part of MPI's future. The sheer options and complexity imply an incredibly rich and fertile ground to fundamentally improve MPI performance and add robust molecular imaging capabilities in the future. As relaxation leads to a relaxation wall on the one hand, it provides contrast-enhancing information on the other. Additionally, en route to exploring these dynamics to improve MPI, we may be able to provide important contributions to the basic sciences and nanophysics. Recent developments in the field such as highly monodisperse tracers and flexible tools such as the AWR that allow us to probe parametric dependencies in much finer detail will be crucial components of this process as we believe robust empirical studies must be leveraged intensely along with theoretical work to get to the bottom of MPI physics.

2.4.2 Relaxation Effects Across the Tracer Spectrum

Testing a variety of MPI tracers across a wide spectrum of excitation conditions both further documented known behavior and demonstrated that some tracers behave very differently from others.

The FWHM data show that for the most part, MPI tracers exhibit monotonically worsening resolution with increasing excitation frequency and amplitude. And as we observed, the excitation amplitude has a stronger relative impact across the typical parametric ranges encountered in MPI. The monodisperse tracers from Imagion Biosystems show an intriguing inflection point and reversal of this trend in terms of frequency dependence. Multiple explanations may account for this behavior, including a unique physics for these tracers or that our x-space reconstruction implementation induces this as an artifact. Although the x-space reconstruction method is extremely general (the core of the algorithm is gridding from the time domain to the image domain), the analytical results in the literature are derived assuming no relaxation. Multiple steps in the x-space reconstruction process such as velocity compensation and gridding are phase sensitive. Phase lags from relaxation lead to skews and shifts but can also lead to more severe artifacts when phase shifts are very large. Additionally, while monotonically worsening resolution with increasing excitation frequency is typical, the convexity of the relationship differs across tracers.

The peak signal data show that reduced signal intensity with increasing frequency is very conserved across all tracers tested. The rate of reduced signal with frequency, however, varies across the tracers. The tracers from Imagion Biosystems all show greater rates of signal reduction than other tracers, regardless of core size. Interestingly, even the 18 nm tracer, a size typically associated with negligible relaxation behavior, shows above average signal loss. This could be more evidence supporting the idea that these tracers differ from more typical MPI tracers in terms of relaxation physics. The largest core tracers show significantly greater peak signal loss, corroborating the association of larger core sizes with increased relaxation dynamic time constants (regardless of mechanism) and therefore relatively lossier responses to MPI excitation.

Finally, the peak image domain shift data paints a very similar picture as the FWHM data. Shifts generally increase in severity with increasing excitation slew rate, with the Imagion Biosystems tracers uniquely showing a very different behavior characterized by an inflection point around 10 kHz after which the shifts plateau or decrease.

Taking all of these results together, we can make clear associations between tracers that consistently show less relaxation effects and those that show more. For example, the UW (monodisperse) and M2 macroaggregated albumin (polydisperse) tracers consistently perform the best with respect to raw MPI metrics. These two tracers have the best resolution, the smallest changes of FWHM resolution with excitation parameters, the smallest losses in peak signal intensity with varying excitation parameters, and the smallest image domain shifts. The Nanomag MIP and Resovist tracers, somewhat similar in their core polydispersity and dextran coating, also perform well overall and quite similarly to each other. Interestingly, Nanomag MIP performs slightly better in overall FWHM resolution, but tends to have noticeably higher parametric variation in FWHM resolution, peak signal intensity loss, and image domain shifts compared to Resovist. This aligns well with what we do know of the differences between the two – Nanomag MIP is composed of polydisperse cores with a larger average core size. This implies Nanomag MIP should have somewhat more detectable relaxation effects and that the improved resolution is due to the mean core size being closer to the optimal value in terms of the relaxation wall.

One might be tempted to expect that excitation slew rate is the operative parameter in magnetic dynamics, predicting a symmetry in the effects of varying excitation amplitude and frequency. However, parsing the effects, we see that the excitation amplitude has a much stronger impact on FWHM resolution and image domain shift but little to no effect on peak signal intensity. The impact of excitation frequency is largely reversed. This difference may speak both to the nature of magnetic relaxation and how the x-space gridding algorithm maps relaxation effects to the image. For example, it may be that excitation frequency determines the lossiness or energy efficiency of the particle response relatively independently of the amplitude. In contrast, the amplitude is responsible for the degree to which laggy signal is spread over the pFOV in the gridding scheme of reconstruction. For example, larger amplitudes mean larger pFOVs and thus a relatively larger spreading of lagging signal over space during time domain to image domain gridding.

In general, the amplitude and pFOV size have a relaxationship with x-space gridding error completely distinct from the excitation frequency: The excitation amplitude and pFOV size fundamentally limit the maximal extent of any gridding error in reconstruction. If we assume that magnetic relaxation dynamics are very fast with respect to the slow shifting bias or shift fields (typically approximately 10 Hz, so this is a good assumption), then in the limiting case of vanishingly small excitation amplitudes, relaxation blur and shift effects must also be vanishingly small in the image domain *regardless* of the severity of magnetic relaxation phase lag. This is because, no matter how severe lags are, they cannot locally be gridded to physical space larger than the extents of the pFOV. Scaling up beyond vanishingly small amplitudes, we expect the pFOV size to still control and scale the effects of relaxation signal spread and blur in the image domain.

2.4.3 Relaxation Maps and Image Contrast

It is known that all expected magnetic relaxation phenomena, *e.g.*, Néel, Brownian, ferromagnetic, are strong functions of the applied magnetic field. Furthermore, the MPI process consists of rapidly varying large magnetic fields with strong gradients. While various models have been put forth in the literature, the field-dependence of MPI tracer relaxation is complex and incompletely understood, especially considering simultaneously active mechanisms. The relaxation maps constructed in this work provide the first, if coarse, empirical quantification of this dependence. As expected, the data show phase lag vs. applied field relationships that are symmetric about a peak centered around zero field. We also see very large differences in relaxation maps of an 18 nm core tracer and a 32 nm core tracer, with the latter showing much more pronounced relaxation behavior, both in terms of peak phase delay and phase delays at any applied field strength.

These relaxation maps also make clear one of the ways in which magnetic relaxation leads to a loss of FWHM resolution, in a process we call 'peak suppression'. As the various theoretical models assume and these data confirm, relaxation phase delay (and underlying relaxation time constants) is increased markedly near zero field. This implies that the inability to achieve steady-state magnetization conditions is that much more pronounced under low field conditions. This relative reduction in signal due to increased lossiness at low field should lead to artificially reduced MPI signal at low field, precisely the location of the peak MPI signal due to the nature of the Langevin function and the tracer's M-H curve. The result is for the PSF peak to be artificially reduced with respect to the rest of the PSF. This will directly lead to relatively wider PSF with a worse FWHM value compared to the ideal.

These relaxation maps may provide a powerful mechanism to improve current x-space reconstruction. x-Space reconstruction is very sensitive to incorrect phase, and it is likely that many of the untoward effects of relaxation can be mitigated by taking into account tracer-specific and trajectory-localized phase information provided by relaxation maps. For example, by simply delaying the signal of each pFOV by the locally calculated phase delay, some of the resolution and peak signal-damaging effects during the x-space gridding procedure may be mitigated.

Finally, relaxation maps like this may represent a great way to report tracer identity or physiologic contrast. In a most straightforward way, construction of relaxation maps as presented here, when applied to imaging scanner data, can provide a *new type* of image. Such a relaxation image is related to (gross tracer distribution) but distinct from the standard iron mass/concentration image and would have intensity in units of phase or time. As we have shown here, different tracers have very distinct relaxation maps, and a tracer with significant Brownian behavior should produce a relaxation image with contrast associated with changing physiologic parameters such as viscosity, pH, or binding state. Distinct relaxation images or some combination of current iron mass/concentration images and relaxation images could provide powerful new 'colorized' molecular imaging applications for MPI. We will look at color MPI in the next chapter and revisit relaxation maps/imaging in detail in Section II.

2.5 Conclusions

We have empirically documented the relaxation-induced behavior of MPI tracers over a large parametric space, including magnetic material, core size, polydispersity, excitation amplitude, and excitation frequency. The analysis of approximately 1500 experiments over such a broad space is a testament to the flexibility of our recently developed untuned AWR device. These results reinforce a number of previously observed phenomena and provide detailed and specific quantification of these phenomena that was previously unavailable. We have also documented some surprising results for some tracers, such as the FWHM vs. frequency relationship of the Imagion Biosystems tracers. We were able to rigorously document and quantify the 'relaxation wall' phenomenon in which magnetic relaxation prevents realization of the Langevin promise of increasing resolution with increasing magnetic core size beyond approximately 25 nm. We also show a quantification of magnetic relaxation time delays and phase lag as a function of applied field by constructing relaxation maps.

The richness of the dataset and analytical methods allowed us to hypothesize and describe aspects of MPI relaxation effects such as how markedly more significant relaxation phase lag at low field leads to PSF 'peak suppression' and reduced resolution. The relaxation maps we have described appear to have great future utility: The scan-specific and localized relaxation information provided by relaxation maps may be used to greatly improve x-space reconstruction for MPI tracers with significant relaxation. Moreover, we believe these relaxation maps could be the tip of the iceberg for new molecular imaging contrast in MPI. Whether by directly providing new 'relaxation images' or informing new algorithms, it is clear that we can take advantage of magnetic relaxation in 'color' MPI applications where information about tracer identity and physiologic conditions can be extracted and reported in an imaging format.

In future work, we will continue to characterize and test the growing family of MPI tracers and build a robust database of empirical quantification of MPI relaxation behavior. A major thrust should be unifying these empiricial results with theoretical models. This will likely involve both using and vetting existing models as well as development of new or modified models. Finally, we will be actively exploring how these results can inform improved reconstruction techniques and new molecular imaging approaches.

2.6 Acknowledgements

I would like to thank Laura Taylor for her help acquiring the large dataset discussed in this chapter and Zhi Wei Tay for developing and building the AWR system that made acquiring and analyzing these data possible. Also, I would like to acknowledge the numerous collaborators that provided us with the wide range of tracer samples used in this work: Professor Kannan Krishnan at the University of Washington, Dr. Matt Ferguson and colleagues at LodeSpin Labs, Professor Carlos Rinaldi and Rohan Dhavalikar at the University of Florida, Gainesville, Dr. Erika Vreeland and associates at Imagion Biosystems for providing custom particles, Professor Angie Louie at the University of California, Davis, and Professor Andris Bakuzis at Universidade Federal de Goiás.

Chapter 3

x-Space Color MPI

Non-negligible magnetic relaxation generally degrades signal intensity and resolution in standard MPI systems. However, magnetic relaxation can also report information such as local microenvironmental conditions, tracer identity, or tracer binding state. Here we explore the use of MPI magnetic relaxation to provide multi-color images in MPI. We describe an encoding and reconstruction scheme to accomplish this task, show proof-of-concept results and quantification imaging a test apparatus, and finally show the first *in vivo* application of color MPI.

3.1 Introduction and Background

As described in previous chapters, non-negligible magnetic relaxation generally degrades signal intensity and resolution in standard MPI systems. However, there is another very important side to the story. Magnetic relaxation may also encode useful imaging contrast that can be exploited and revealed through new acquisition and reconstruction strategies.

Consider the case of a relatively large MPI tracer, *e.g.*, a sample consisting of monodisperse 30 nm particles, whose relaxation is dominated by the Brownian mechanism. This tracer must physically rotate to align with applied excitation fields. The physical rotation is driven by the magnetic torque imposed by the applied field but is also affected by local microenvironmental factors such as viscosity, pH, and binding states. Differing microenvironmental conditions at different locations, different states of the tracer at different locations (*e.g.*, bound to some biological target), or different tracer identities will lead to different magnetic relaxation behavior. Thus, the magnetic relaxation dynamics may encode the state of various physiologic variables of interest.

The ability to distinguish tracer state and report this information in a multi-color imaging format (color MPI) has recently been explored in the literature (Rahmer et al., 2015; Hensley et al., 2015b; Stehning, Gleich, and Rahmer, 2016). Here, we describe our approach to color MPI. We focus on use of two different MPI tracers with different magnetic relaxation characteristics and assess an approach to recognize and disambiguate signal originating from one or the other. Ideally, this allows us to quantitatively unmix and quantify populations of these two tracers, even when co-localized in space, and report the results as multi-color images. We describe an encoding and reconstruction scheme to accomplish this task, show that our method is capable of disambiguating co-localized tracer species in an imaging format with a test vial apparatus, and finally show the first *in vivo* application of color MPI.

3.1.1 Magnetic Relaxation in MPI

As we explored in the previous chapter, Langevin theory only describes steady-state conditions and says nothing about the dynamics of how tracer populations evolve toward a new steady-state magnetization in response to a change in applied field. The latter is the domain of magnetic relaxation. Models that include relaxation will be indistinguishable from models that assume instantaneous achievement of Langevin steady-state when the MPI excitation slew rates are much faster than the time constants associated with the dynamic response. As these time constants approach the same order of magnitude as the excitation components and beyond, however, we must include specific dynamic models to correctly understand how the MPI signal will evolve. We begin by recapitulating the relationship between FFR trajectory and applied field in the canonical x-space approach to MPI (Goodwill and Conolly, 2011):

$$\mathbf{H}(t, \mathbf{x}) = \mathbf{H}_{\mathrm{s}}(t) + \mathbf{G} \cdot \mathbf{x} \tag{3.1}$$

$$\mathbf{x}_{\mathrm{s}}(t) = \mathbf{G}^{-1} \cdot \mathbf{H}_{\mathrm{s}}(t) \tag{3.2}$$

where $\mathbf{H}(t, \mathbf{x})$ is the total applied field in space over time in $[\mathrm{A\,m^{-1}}]$, $\mathbf{H}_{\mathrm{s}}(t) = [H_x(t)\ H_y(t)\ H_z(t)]^\top$ is the spatially homogeneous but time varying excitation field, $\mathbf{G} \in \mathbb{R}^{3\times3}$, $\mathrm{trace}(\mathbf{G}) = 0$ is the 3D linear gradient field ($[\mathrm{A\,m^{-2}}]$) employed during the MPI scan, and $\mathbf{x}_{\mathrm{s}}(t)$ is the instantaneous location of the FFR center.

The steady-state magnetization can be expressed with the fundamental spatial encoding mechanism in MPI, the Langevin function:

$$\mathbf{M}_{\text{ss}}(t,\mathbf{x}) = m\rho(\mathbf{x})\mathcal{L}(k\|\mathbf{H}(t,\mathbf{x})\|) \cdot \hat{\mathbf{H}}(t,\mathbf{x}) \tag{3.3}$$

where m in $[\mathrm{A\,m^2}]$ is the magnetic moment of the tracer, $\rho(\mathbf{x})$ is the tracer density over space in $[\mathrm{particles\,m^{-3}}]$, $\mathcal{L}$ is the Langevin function, and $\mathbf{M}_{\text{ss}}(t,\mathbf{x})$ is the steady-state magnetization with magnitude in $[\mathrm{A\,m^{-1}}]$.

As described in Chapter 2, we can phenomenologically describe magnetic relaxation in MPI as a Debye-like relationship, yielding:

$$\mathbf{M}(t,\mathbf{x}) = \mathbf{M}_{\text{ss}}(t,\mathbf{x}) \overset{t}{*} \frac{1}{\tau_{\text{eff}}(t,\mathbf{x})} \mathrm{e}^{-t/\tau_{\text{eff}}(t,\mathbf{x})} \tag{3.4}$$

where $\mathbf{M}(t,\mathbf{x})$ is the total magnetization in the imaging FOV over time and space and $\tau_{\text{eff}}(t,\mathbf{x})$ is the 'effective' relaxation time constant of the tracer which in the general case also varies with time and position.

3.1.2 Image Domain Effects of Relaxation

In our x-space approach to MPI data acquisition and reconstruction, we simply grid the voltage signal on our inductive receiver coil(s) to the known instantaneous location of the FFR as described by equation (3.2) (Goodwill and Conolly, 2011; Lu et al., 2013). Thus, there is a direct linear mapping between time and space in these equations. This implies that if magnetic relaxation is described as a temporal convolution with an exponential kernel as in equation (3.4), then this low-pass filter effect will also translate to the gridded image domain. We would expect a temporal filter to lead to image domain skewing, blurring, and translation.

To develop a simplified model for the impact of temporal relaxation in the final output image domain, we follow (Croft et al., 2016) but stay in full 3D space. First we assume that $\tau_{\text{eff}}(t,\mathbf{x}) = \tau_{\text{eff}}(v(t))$ where we assert that the effective relaxation time constant depends only on the magnitude of the instantaneous excitation velocity $v(t) = ||\dot{\mathbf{x}}_{\text{s}}(t)||$. We also assume the ability to obtain the portion of the signal collinear to any arbitrary excitation direction (Goodwill and Conolly, 2011) (*e.g.*, via projection from the combined signal of receiver coils aligned with all three principal axes). This allows us to express the following relationship for our estimated tracer density as a

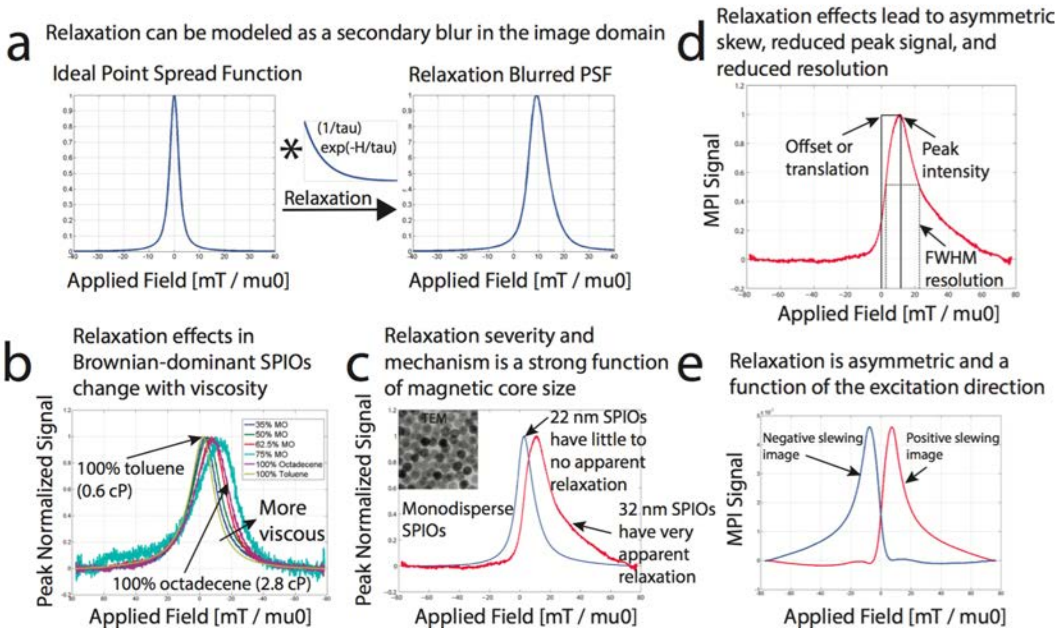

FIGURE 3.1: MPI relaxation effects in the image domain. (a) Relaxation can be phenomenologically modeled as a convolution of the ideal image with an exponential kernel in the direction of excitation. (b) For Brownian relaxation-dominant cobalt-ferrite SPIOs, increasing viscosity leads to increased relaxation effects. (c) Magnetic relaxation is a strong function of magnetic core size. (d) Image domain parameters that characterize relaxation include peak offset or translation, changes in peak intensity, and broadening of the PSF. (e) Separately reconstructed positive- and negative-slewing images make clear the scan direction dependence of relaxation effects.

function of the FFR trajectory:

$$\hat{\rho}(v(t), \mathbf{x}_\mathrm{s}(t)) \approx \rho_{\dot{\mathcal{L}}}(\mathbf{x}) \overset{\hat{\mathbf{x}}_\mathrm{s}(t)}{*} \frac{1}{v(t)\tau_\mathrm{eff}(v(t))} \exp\left(\frac{-x}{v(t)\tau_\mathrm{eff}(v(t))}\right) u(x)\bigg|_{\mathbf{x}=\mathbf{x}_\mathrm{s}(t)} \tag{3.5}$$

where $\hat{\rho}(v(t), \mathbf{x}_\mathrm{s}(t))$ describes our reconstructed image including relaxation effects which therefore depends on the nature of the excitation trajectory specified by the excitation velocity waveform $v(t)$ and FFR spatial trajectory (characterized by positional trajectory $\mathbf{x}_\mathrm{s}(t)$ and instantaneous excitation direction $\hat{\mathbf{x}}_\mathrm{s}(t)$. $\rho_{\dot{\mathcal{L}}}(\mathbf{x})$ is the ideal steady-state reconstructed image, blurred by the Langevin function but without any relaxation effects and $u(x)$ is the Heaviside step function. Equation (3.5) states that an arbitrary MPI acquisition trajectory will sample an image that results from a 1D spatial convolution, in the direction of the instantaneous excitation direction, between the steady-state Langevin-blurred MPI image and a spatial Debye relaxation kernel.

As is typically the case in our systems, we now further assume excitation in a single, constant direction during a scan and reconstruction using a receiver coil collinear with this excitation direction. As in (Croft et al., 2016), we also approximate the excitation velocity in a given n-th scan as a constant: $v_n = A_n f_0 / G$. A_n is the excitation amplitude of the n-th scan in $[\mathrm{T}]$, f_0 is the fixed excitation frequency in $[\mathrm{Hz}]$, and G is the linear gradient strength in the excitation direction in $[\mathrm{T\,m^{-1}}]$. This allows us to express our estimated image as a function only of the tracer density over space and the excitation amplitude of a given scan:

$$\hat{\rho}(A_n, \mathbf{x}) \approx \rho_{\dot{\mathcal{L}}}(\mathbf{x}) \overset{\hat{\mathbf{x}}_n}{*} \frac{G}{A_n f_0 \tau_n} \exp\left(\frac{-Gx}{A_n f_0 \tau_n}\right) \tag{3.6}$$

where τ_n is the exponential constant for the n-th excitation scan as determined by the amplitude A_n and $\hat{\mathbf{x}}_n$ is a unit vector pointing in the constant excitation direction of the n-th scan. We have thus expressed the impact of our time domain relaxation blur in the image domain.

Inspecting equation (3.6), we see that our model predicts that our images will be blurred in the excitation direction in an amplitude-dependent manner. By varying the excitation amplitude, both the intensity of the image and the severity of the blur/skew/translation effects should change.

As illustrated in Fig. 3.1, reconstructed images from our x-space scanners and AWR indicate that this model well captures the image domain behavior when magnetic relaxation is present. In addition to a reduction in the signal intensity, there is specific structure to the blurring effect – asymmetric skew and translation in the direction of the excitation. Because these translation and signal intensity effects are a function of the excitation amplitude, this implies that we can modulate and control the effect of magnetic relaxation by modifying the amplitude. This relationship has been studied and experimentally quantified by Goodwill, Croft, and others (Goodwill et al., 2011; Croft et al., 2016).

While perhaps capturing salient image domain behavior, equation 3.6 highly simplifies the process of relaxation. In reality, relaxation for tracer at any point and at any time depends strongly on the magnitude of the applied field at that location, which varies sharply in space due to the presence of gradient fields. Moreover, our inductive reception integrates over the spatially varying relaxation. In a multidimensional context, significant signal from vector rotation of saturated tracer leads to somewhat long hazy tails in

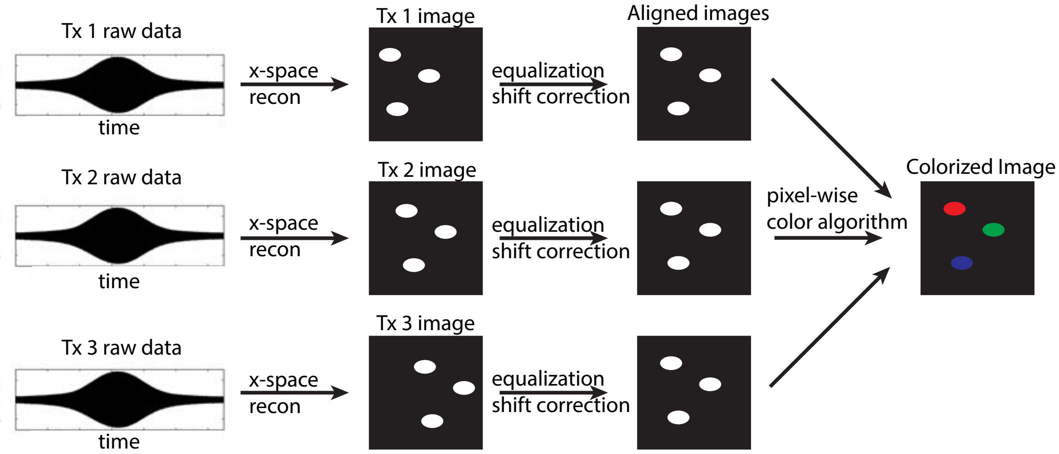

FIGURE 3.2: x-Space color MPI algorithm. Raw data from scans using different excitation amplitudes are reconstructed individually using standard x-space methods. The resulting series of images may undergo color pre-processing steps including application of an MPI equalization filter and shift correction. The pixel-wise color algorithm is then applied.

the multidimensional PSF. The relaxation associated with this long tail component will differ, perhaps significantly, from the tracer signal near the FFR center due to the strong dependence on applied field. Additionally, our x-space reconstruction algorithms do not account for tracer phase lag, which affect velocity compensation and gridding steps, leading to artifacts relative to the ideal. However, this expression is useful in capturing the general effects of non-negligible relaxation and explains the approach to relaxation encoding in our method of color MPI.

In the remainder of this chapter, we focus on how we may leverage these changes to produce colorized images.

3.2 Materials and Methods

We have developed an encoding and reconstruction method for x-space color MPI, described at a high level in Fig. 3.2. Experimentally, we tested the ability of this approach to disambiguate two tracers distributed in the imaging FOV: chemicell (chemicell GmbH, Germany) and nanoMag MIP (micromod Partikeltechnologie GmbH, Germany). We tested the method first on 2D image slices of test vials with known amounts of each tracer either physically separated from each other or co-located as a mixture. We then tested the ability to use this algorithm for 3D *in vivo* color MPI in a rat.

3.2.1 x-Space Color MPI Encoding

All imaging was performed with our in house [3.5 x 3.5 x 7] $\mathrm{T\,m^{-1}}$ FFP scanner (Goodwill et al., 2012a). To encode magnetic relaxation information, we sequentially scanned the object under study multiple times, each using a different excitation amplitude (typically, $3-5$ distinct scans with amplitude in the range of 10–20 mT). According to equation (3.6), the different excitation slew rate used in each scan should elicit different relaxation effects. Furthermore, we assume that the relaxation effects will be different for each of the tracer types present in the image. This is shown experimentally in Fig. 3.3(b,c). The 1D PSFs, obtained using an x-space relaxometer (Croft et al., 2016), show different behaviors for two tracers when the excitation amplitude is changed. In particular, a smaller 22 nm core tracer (Imagion Biosystems, Inc., NM USA) exhibits little to no detectable scaling with excitation amplitude after accounting for the expected linear scaling from changing slew rates per Faraday's Law. In contrast, a larger 27 nm tracer from the same manufacturer shows significant relaxation-induced changes and asymmetries.

Curiously, and in a very repeatable fashion, the higher excitations induce a *larger* signal, after Faraday normalization, than smaller excitations. Our simplified equation (3.6) predicts the opposite behavior. Aside from the many simplifications used in generating equation (3.6), the effects of increased relaxation delays on the x-space algorithm itself may lead to this discrepancy. We will consider this more in our discussion section. Nonetheless, we do observe strong amplitude-dependent changes in image domain intensity.

3.2.2 x-Space Color MPI Reconstruction Algorithm

The color MPI algorithm is applied to the set of images after x-space reconstruction is applied to the raw data (Lu et al., 2013) as indicated in Fig. 3.2. Additionally, combining orthogonal direction excitation scans and an MPI equalization algorithm can provide x-space images with isotropic resolution and reduced haze (Lu et al., 2015). The background haze and component of the MPI PSF that is attenuated using this and similar equalization filters comes from tracer that is relatively far from the FFR. Being far from the FFR, this signal originates from tracer in the presence of relatively high field. Even tracer with significant apparent relaxation at the low fields associated with the FFR center will typically have negligible or weakly apparent

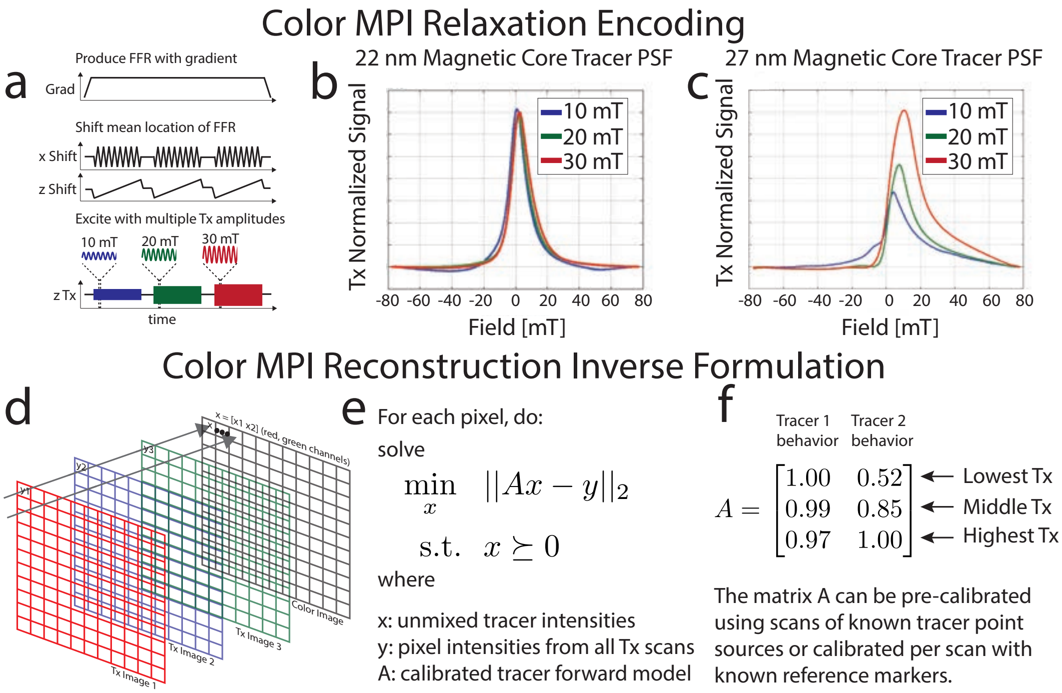

FIGURE 3.3: Signal encoding and reconstruction in x-space color MPI. (a) Standard x-space acquisitions with different excitation amplitudes are taken sequentially. (b) A small 22 nm core tracer shows no detectable relaxation-induced changes with amplitude while (c) a larger 27 nm core tracer shows considerable changes. (d) Depiction of the process of converting a stack of M x-space color MPI acquisitions to a colorized image with N colors. (e) The colorizing algorithm solves a convex nonlinear least squares optimization problem per image domain pixel/voxel. (c) The forward model matrix A describes how various tracers map to different Tx amplitude images. The n-th column of A describe the (relaxation) behavior of the n-th tracer over the excitation amplitude space.

relaxation in such a context due to the strong dependence all magnetic relaxation mechanisms have on the applied field strength. From the perspective of the color algorithm, in this scenario, the distinguishing relaxation of the tracers may be erased. We expect that application of an equalization filter will improve the accuracy of the color MPI algorithm.

We can develop our relaxation encoding into an N-compartment model reconstruction. In particular, the excitation amplitude encoding admits an image-domain forward model that may be solved by standard inverse problem techniques as depicted in Fig. 3.3(d–f). The algorithm compares the same pixel location across different excitation images and is thus sensitive to any relative shifts that may occur during reconstruction. Therefore, a simple shift correction, using the known location of reference point sources, may be employed to rectify any minor shifts that arise in reconstruction.

The following color algorithm is applied in a pixel/voxel-wise manner:

$$\begin{aligned} \min_{x} \quad & ||Ax - y||_2 \\ \text{subject to} \quad & x \succeq 0 \end{aligned} \tag{3.7}$$

where $x \in \mathbb{R}^N$ contains the image intensity associated with each of the N tracers under consideration, $y \in \mathbb{R}^M$ contains the image intensity from each of the M input images (each acquired with a distinct excitation amplitude), and $A \in \mathbb{R}^{M \times N}$ is the forward model mapping tracer behavior across the excitation amplitude space. The columns of A contain the expected (peak-normalized) change in pixel intensity as a function of excitation amplitude for a given tracer type and after normalizing for the expected linear change in signal due to inductive signal reception. An MPI tracer that exhibits negligible relaxation would lead to a column of unity values. The matrix of Fig. 3.3(f) is taken from experimental data; the left column corresponds to the chemicell tracer which shows negligible relaxation effects while the right column corresponds to the nanoMag MIP tracer which shows clear relaxation-induced changes. Considering (3.6), this implies that the *effective* observed time constant decreases with increasing amplitude for this tracer and using this specific implementation of x-space reconstruction methods.

The only *a priori* information injected into the inverse problem is a non-negativity constraint imposed on x, which is well-motivated because negative mass or concentration of a tracer is not physically possible. We note that there are no other imposed constraints, tradeoff terms, or regularization parameters in the cost function of this convex optimization problem. One consequence of this fact is that this pixel-wise algorithm, applied to each voxel in a tomographic dataset is equivalent to solving a single global optimization problem with the same form as equation (3.7), but with the A matrix used to construct a much larger block diagonal forward operator and the x and y vectors accordingly concatenated to form much larger vectors.

Since we are combining a measurement of relaxation (amplitude changes) with an N-compartment model, we do require specification of our forward model matrix A. Typically, we can do this using fiducials placed in the image FOV during a scan, in which we isolate the amplitude effects for each tracer. We note, however, that this is only required to fit to the N-compartment model, which provides explicit quantitative reporting of each compartment (*e.g.*, each tracer). We can also directly report 'relaxation images' by gridding the amplitude variance for each pixel. This could be combined with a standard MPI image via overlay, for example. Such an approach is 'calibration-free', but we will focus on the explicit N-compartment model that requires calibrating the simple forward model A in this chapter.

3.2.3 Initial Experiments

For initial testing of the color MPI algorithm, we used a test apparatus consisting of Eppendorf tubes arranged along a line as shown in Fig. 3.4(a). Each Eppendorf tube was filled with 50 μL of tracer at 25 $\mathrm{mg\,mL^{-1}}$. The vial on one end contained only the nanomag tracer, the central vial contained only the chemicell tracer, and the vial on the other end contained a 50/50 mixture of chemicell and nanomag by volume. The vials were then serially scanned using 10, 20, and 30 mT excitation (4 x 4 x 8.5 cm FOV, approximately 10 min per scan). Each scan was reconstructed using our standard x-space MPI technique followed by application of a high pass equalization filter (Lu et al., 2015) The color algorithm was then applied to the set of scans to produce tracer-specific images that can be superposed using colorized contrast. The maximum intensity pixel in each of the single tracer-only regions was used to construct the forward model matrix A.

3.2.4 In-vivo Experiments

A 0.2 kg female Fisher rat was injected with 0.229 mg of Chemicell SPIOs via tail vein injection. This was followed by an injection of 875,000 (0.08 mg Fe) macroaggregated albumin SPIO (MAA-SPIO) conjugates as described in Zhou et al. (Zhou et al., 2016) using the nanomag MPI tracer. The physically large MAA-SPIOs injected in this manner will immediately get sterically stuck in the lungs, the first place where they encounter capillary level blood flow, followed by slow degradation and transport to the liver over the course of hours. Two reference markers (2 μL each) corresponding to the two tracers (chemicell at 4.16 $\mathrm{mg\,mL^{-1}}$ and MAA-nanomag at 0.67 $\mathrm{mg\,mL^{-1}}$) were placed in the FOV above the chest. The rat was then serially scanned at four excitation amplitudes of 10, 15, 17.5, and 20 mT (4 x 3.75 x 10 cm FOV, approximately 10 min per scan, with respiratory gating). X-ray imaging was performed with a Kubtec Xpert 40. Each scan was reconstructed using our standard x-space MPI techniques. The colorizing algorithm was then applied to the set of scans to produce distinct 3D data sets for each tracer, allowing for creation of colorized images. The two reference markers

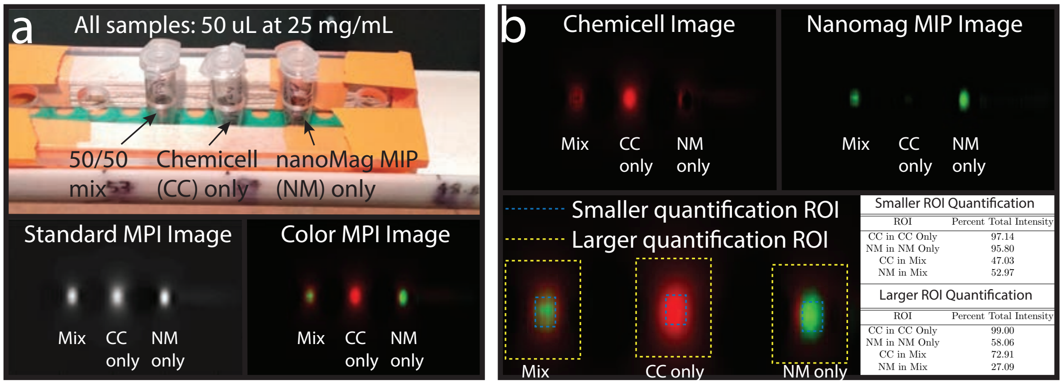

Smaller ROI Quantification

ROI	Percent Total Intensity
CC in CC Only	97.14
NM in NM Only	95.80
CC in Mix	47.03
NM in Mix	52.97

Larger ROI Quantification

ROI	Percent Total Intensity
CC in CC Only	99.00
NM in NM Only	58.06
CC in Mix	72.91
NM in Mix	27.09

FIGURE 3.4: Color MPI results with test vial apparatus. (a) Multi-amplitude acquisition for color MPI encoding was performed at 10, 20, and 30 mT on a series of Eppendorf tube samples. 50 μL samples (25 $mg\, mL^{-1}$) of chemicell only, nanomag only, and a 50/50 mixture by volume were placed in Eppendorf tubes and aligned in the imaging bore. A single standard MPI (20 mT acquisition) image and the constructed color MPI image are shown. (b) A view of the chemicell and nanomag distinct tracer images and quantification of the results. The color algorithm is generally successful in disambiguating the two tracers, even when co-localized, except for visually apparent error in classifying all low intensity background MPI haze and peripheral signal as chemicell (even in the nanomag-only ROI).

were used to construct the forward model matrix A. These animal procedures were conducted in accordance with the National Research Council Guide for the Care and Use of Laboratory Animals and approved by the UC Berkeley Animal Care and Use Committee.

3.3 Results

We first report proof-of-concept x-space color MPI using the Eppendorf tube apparatus along with quantitative analysis of the results. We then report the first *in vivo* MPI applied to a tomographic rat dataset.

3.3.1 Test Vial Imaging and Quantification

Fig. 3.4 shows imaging results with some quantitative analysis of the color algorithm using our test vial setup. In Fig. 3.4(a), the Eppendorf tubes are visually colorized as expected: The nanomag-only and chemicell-only vials show up mostly as their respective colors, with the mixture vial showing up colored by both. However, there is a noticeable degree of chemicell color in the periphery of the nanomag-only vial.

In Fig. 3.4(b), each tracer image is shown separately along with region of interest (ROI) quantitative analysis. The individual tracer images show more clearly how the nanomag-only vial does show up to a non-trivial degree in the chemicell image. Visually, an outer halo is classified as chemicell.

A small ROI, smaller than the total signal area for a given Eppendorf tube but which captures the high intensity core, as well as a larger ROI fully encompassing each vial signal area were assessed. The quantitative analysis of the small ROI matches well with what is expected. In terms of integrated signal, 97% of the total signal associated with the chemicell-only vial ROI is labeled chemicell (*i.e.*, shows up only in the chemicell image), 95% of the signal associated with the nanomag-only vial ROI is labeled nanomag, and the 50/50 mixture vial is split 47% chemicell and 53% nanomag. However, when we analyze the larger fully-encompassing ROIs, the bias toward chemicell in terms of the low intensity haze and background that we see upon qualitative inspection is clear. While 99% of the chemicell-only vial ROI signal is labeled chemicell, only 58% of the nanomag vial ROI signal is labeled nanomag. Furthermore, 73% of the total signal of the mixture vial was assigned to chemicell while only 27% was assigned to nanomag.

In terms of the fully encompassing ROIs, which include the long tails and haze in MPI, these numbers indicate that the color algorithm misclassifies roughly 50% of the nanomag tracer as chemicell. The smaller ROI results, however, which align much better with our expectations, indicate that the misclassification is almost exclusively confined to the outer haze and periphery.

3.3.2 In-vivo Results

Figure 3.5 shows the first *in vivo* application of color MPI. On the left of Fig. 3.5(a), a standard MPI image from a 20 mT excitation scan is overlaid on an x-ray projection of the animal and indicates a large signal mass in the abdomen with no clear distinction between the lung and liver. As shown in the right of Fig. 3.5(a), the color algorithm parses the MAA-nanomag tracer predominately in the lungs from the chemicell tracer in the liver. Because we use an explicit compartmental model and the standard MPI signal is linear with mass or concentration, we can report this information quantitatively in terms of the mass of iron tracer per voxel in $[\mu\mathrm{g}\,\mathrm{mm}^{-3}]$. Finally, we can visualize the combined colorized image in the right of Fig. 3.5(a) where we have overlaid the color image on the X-ray projection of the rat. The anatomical location and identification of the lungs and liver is clear with the color contrast.

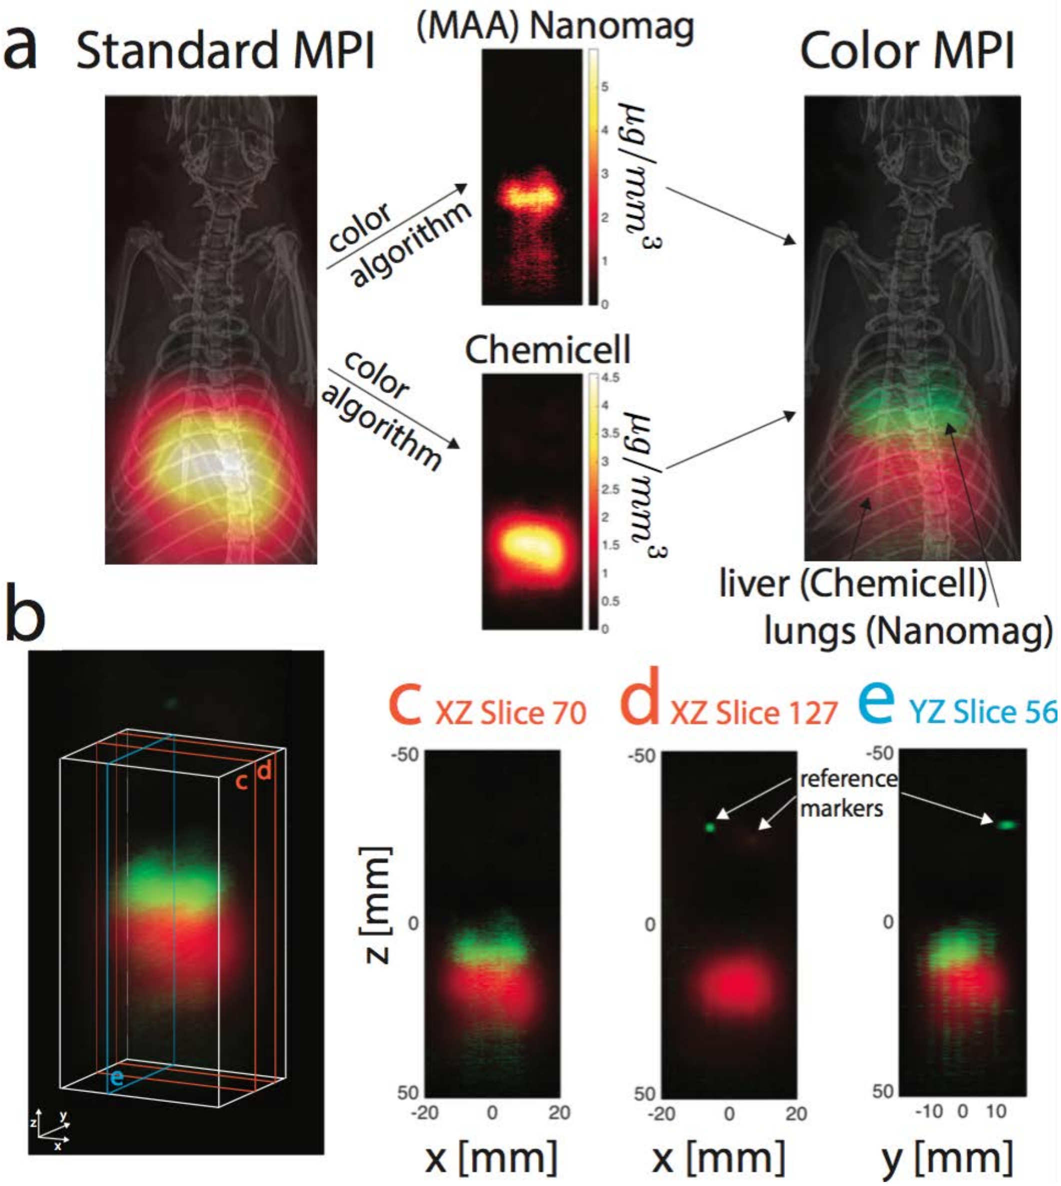

FIGURE 3.5: First *In vivo* color MPI results. (a) A standard MPI scan following IV injection of the Chemicell MPI tracer and MAA-SPIOs (constructed with the nanomag tracer). Our color algorithm quantitatively unmixes the MAA-SPIOs (trapped in the lungs) from the Chemicell tracer (cleared straight to the liver). We can use the results to generate a colorized image that clearly distinguishes the two tracer populations and organs, which is not the case in the standard MPI scan. (b) The entire 3D tomographic data set was colorized. (b-e) Here we show select colorized slices in the xz and yz planes where the smaller MAA-SPIO tagged lung is further distinguished from the Chemicell-tagged liver.

Figure 3.5(b) illustrates the tomographic nature of the color algorithm. On the left, an image is shown after rotating the tomographic dataset and then taking a projection at the new angle. The coordinate system and bounding box indicate the relative positions of three slices shown in Fig. 3.5(c–e). (c) In the first xz plane slice, we can see the lungs and the liver as before. (d) In a second xz slice plane, the lungs are no longer in the plane, but the larger liver still is. We can also see the presence of the reference markers (point sources of the two tracers) used to construct the color algorithm's forward model. (e) Finally in a yz slice plane we can see the lung and liver relative geometry from a different vantage point. The nanomag reference marker also shows up in this slice.

3.4 Discussion

3.4.1 Accuracy of the Color MPI Algorithm

In this chapter we have demonstrated our x-space color MPI algorithm. The test vial results and associated quantitative analysis indicate that the algorithm is not perfect in unmixing tracer. In particular, we noticed a bias for the chemicell tracer (little to no relaxation changes) in terms of the low intensity peripheral regions of signal. We believe this is easily understood in the context of the field dependence of magnetic relaxation in MPI. Peripheral signal originates from a high field context that erases the relaxation distinction between the two tracers that we use to fit the two compartment model. In particular, in a lower field context, our nanomag tracer exhibits detectable relaxation while chemicell exhibits none. In the high field context, however, even nanomag appears to have little to no relaxation and thus is classified as chemicell by the algorithm.

Our algorithm is based on a gross or averaged difference between tracer relaxation behavior in the image domain. The reality of relaxation as discussed in 3.1.1 is more nuanced. Ultimately, the final intensity of image domain pixels superposes contributions from all non-zero signal sources when the FFR is located in the given pixel region. These contributions have spatially varying relaxation governed by their position relative to the FFR. Another curiosity is that the nanomag tracer, with noticeable relaxationd-dependent behavior, shows a higher peak signal (after Faraday normalization) for higher excitation amplitudes. Our simplified model (3.6) predicts the opposite. On the other hand, if we consider that the largest signal contributions occur when the mean location of the FFR (*i.e.*, from the slowly-varying shift coils) is coincident with a region of signal, the larger amplitudes lead to a much larger excursion of the FFR with each RF ocsillation (pFOV). This means that the signal associated with the mean location will spend more time exposed to much higher field strengths in the oscillation. In this scenario, a larger amplitude could lead to a greater mean change in magnetization per unit time due to the very strong field-dependence of magnetic relaxation. In Chapter 5 and 6, we will explore field-dependent relaxation in much more detail and observe a related 'peak depression' phenomenon.

Considering the MPI PSF (Goodwill and Conolly, 2011), the most peripheral components of the MPI signal come from tracer that is either in a location of higher fields, but not completely saturated, or is fully saturated and generating a small signal by coordinated vector rotation following the local field lines. These physics give rise to the unique tensor nature and asymmetries of the MPI PSF, a dependence of the PSF on the excitation trajectory, and lead to characteristic "long tails" (Goodwill and Conolly, 2011). Use of orthogonal excitation scans and a rigorously developed equalization filter such as that described in (Lu et al., 2015) should help mitigate the presence of the long peripheral tails. This is both desirable generally and should help mitigate the classification error of the color MPI algorithm described herein that is due to a loss of contrast between the relaxation behavior of the tracers in a high field context. Future experimental work should rigorously quantify the performance of the colorizing algorithm over a large number of known test conditions and phantom designs (*e.g.*, geometry and concentration) and with and without the use of equalization filters.

3.4.2 First In Vivo Results

The *in vivo* results shown here are the first such results using color MPI algorithms. Qualitatively, the algorithm is clearly able to distinguish the lungs from the liver, whereby each is marked with a different tracer and in which there is significant overlap in the signal. We believe this is a powerful result that proves the concept of color MPI in the all important *in vivo* context. In addition to the *in vivo* aspect, we showed that our x-space color MPI can be used to generate tomographic colorized data sets that can be analyzed as desired. We also showed that the color MPI algorithm reports distinct tracer images that are quantitative, although we did not independently verify the accuracy of this *in vivo* quantitation.

Per the test vial analysis, it is likely that the quantitation of the distinct tracer images is affected by a bias toward the chemicell (liver) tracer, especially in peripheral signal regions. Again, we expect orthogonal excitation and application of an equalization filter to help mitigate this possibility. Follow up *in vivo* work should establish the repeatability and performance of the algorithm generally and in various different model contexts specifically.

3.4.3 Conditioning of the Algorithm

The forward model operator A employed in this work, as described in equation (3.7) has a condition number typically in the range of approximately 5–15. For example, the operator A used in the *in vivo* results has a condition number of 8.7. The variation occurs because, typically, the matrix A is automatically calculated for a given dataset from reference markers, with known concentrations of tracer in the relaxation states to be unmixed. It is possible, however, to pre-calibrate this matrix and apply it to an imaging dataset without the need for such reference markers. In all cases one is constructing or calibrating the forward matrix, care should be taken to ensure conditions of the calibration match those in the final imaging context as much as possible.

The relatively low condition numbers in our data mean that we were able to solve the straightforward inverse model without requiring use of regularization or injection of significant *a priori* knowledge. The only *a priori* knowledge built into the problem is a non-negativity constraint, which is very well motivated because the output image intensity, corresponding to tracer concentration, cannot physically be negative. We believe that not requiring regularization or other 'twiddle' parameters is highly desired, leading to an algorithm with more straighforward properties and that does not broadly inject nonlinearities whose impact across the input space is difficult to understand. We note that the structure of this problem generally disallows a forward operator with the ideal condition number of 1. Because standard MPI image intensities in a region of high signal are always positive and non-zero, the columns of A will always have entries that are all positive and non-zero. Therefore, no two columns can be orthogonal. If, however, an MPI tracer "switch" is developed whereby signal is zero for one or more of the color encoding scans, orthogonality is possible.

3.4.4 Future Directions

Future x-space color MPI work should focus on validating the algorithm(s) for repeatability, robustness, quantitativeness. Additionally, exploring models beyond the disambiguation of two tracer species is highly desirable (*e.g.*, a single tracer in different binding states).

The algorithm described in this chapter is very simple and ignores known relaxation information. For example, as indicated in Fig. 3.1(c), non-trivial magnetic relaxation leads to image domain translation and other asymmetries in addition to the intensity modifications we leveraged in this chapter. An even better route could be to incorporate time-domain information (before gridding to the image domain), where the effect of magnetic relaxation is fundamentally observed. For example, the relaxation maps we showed in Chapter 2 directly represent a measure of relaxation with resolution at the level of adjacent pFOVs throughout the image. This information can be directly reported as a contrasting relaxation image but can also be brought into the type of N-compartment model demonstrated here. In general, incorporating these other sources of relaxation information will motivate the use of larger and more complex color MPI algorithms.

One of the most important directions to explore in the future is changing how the relaxation information is encoded in the MPI signal in the first place. In this way, one can directly improve the conditioning of the problem prior to the formulation of any specific reconstruction or postprocessing scheme. A simple improvement to the approach herein is to modulate the amplitude of the sinusoidal excitation such that the color MPI excitation encoding across the amplitude space is obtained with a single MPI scan (and image domain information is co-localized in time) rather than a series of sequentially gathered scans. More generally, because relaxation (and MPI in general) is sensitive to the exact FFR trajectory employed, pulse sequence design should be a fantastically rich area to improve color MPI algorithms. The pulsed MPI methods described in part II of this thesis may be the most promising way forward for color MPI methods.

3.5 Conclusions

In this chapter we described the first x-space color MPI algorithm. The algorithm leverages differences in relaxation states between tracer populations to disambiguate and unmix a series of standard MPI images into a single colorized image dataset. We first quantitatively assessed the accuracy of the algorithm to unmix two distinct tracers in a test vial model with known tracer concentrations. We then showed the first *in vivo* color MPI results showing the ability to disambiguate a tracer targeting the lungs and a tracer in the liver. The standard MPI imaging results yield only a single distribution with no contrast or resolved features capable of distinguishing the two organs. We believe these first proof-of-concept color MPI results indicate the great potential in leveraging MPI relaxation physics to provide physiologic contrast. Although we have only explored discrimination of two distinct tracer species, discriminating a single species in distinct relaxation states (*e.g.*, bound vs. unbound) should be possible in the same manner, provided the distinguishing behavior admits a suitably conditioned algorithm. Future developments that improve the encoding of relaxation information in the raw signal and include more relaxation and *a priori* information in the colorizing algorithm formulation are expected to greatly improve the conditioning, applicability, and robustness of color MPI.

3.6 Acknowledgements

I would like to thank Patrick Goodwill and Laura Croft for assistance in getting this project started and Patrick for many great discussions while I was developing the approaches described in this chapter. I would also like to thank both Zhi Wei Tay and Xinyi Zhou for assistance with the *in vivo* experiment described in this chapter. Xinyi also supplied the MAA tracer used in the *in vivo* experiment which she created via her own custom procedure.

Chapter 4

Combining MPI and Magnetic Fluid Hyperthermia

Magnetic relaxation in MPI is a thermodynamically lossy process and this can be explicitly leveraged for therapeutic purposes, such as in targeted heating and drug delivery. The field of magnetic fluid hyperthermia (MFH) has been exploring these applications for many years. In this chapter, we will look at combining MPI and MFH in a single device. In doing so, we can build a new theranostic platform where MPI gains an interventional aspect and MFH gains potentially real-time and seamless image guidance as well as a powerful new ability for spatially localized heating/actuation. We will describe construction of the first combined MPI-MFH device and provide proof-of-concept data showing simultaneous MPI and MFH as well as spatially selective MFH targeting tracer samples anywhere in our device field of view with approximately 3 mm resolution.

4.1 Introduction and Background

While anatomical imaging modalities such as X-ray computed tomography (CT) and magnetic resonance imaging (MRI) are staples of clinical medicine, many pathologies such as cancer may not be readily perceived in such images, especially in early and critical stages of the disease (Etzioni et al., 2003; Serres et al., 2012; Li, 2014). Physiologic contrast between healthy and pathological tissue is needed in these cases. Molecular imaging modalities such as MPI can provide this valuable physiologic and functional information in a clinical context (James and Gambhir, 2012; Pablico-Lansigan, Situ, and Samia, 2013).

Imaging and other diagnostics are ultimately combined with some interventional therapy to treat disease. Combining both aspects into a single "theranostic" platform can deliver faster, more flexible, and more precise treatment and ultimately improve patient outcomes. Real-time interactions and feedback between the diagnostic imaging and therapeutic components can lead to enhanced functionality.

An ideal theranostic imaging platform would be noninvasive, safe, and provide fully real-time feedback. We believe such a platform can be provided by the union of MPI and magnetic fluid hyperthermia (MFH). MFH is a method of heating nanoparticles by driving them through their nonlinear magnetization curve. In this context, MPI provides high contrast and diagnostic molecular imaging of tracer distributions that may be targeted for actuation by MFH. Actuation, in this context, may be varying degrees of localized and macroscopic heating or other non-macroscopic or even non-thermal effects such as induced drug release from mechanically- or thermally-labile bonds. By using MPI hardware, modified slightly from what is typically used, we have the potential to provide *both* imaging and MFH treatment in real-time.

The magnetic nanoparticles used in MPI and MFH significantly overlap in character and physical properties. Both techniques leverage SPIO nanoparticles with magnetization curves characterized by nonlinear saturation at high field magnitudes. In the saturated regions, the particles are effectively locked in place and neither an appreciable MPI signal nor heating occur in response to an AC excitation field. We and other groups have been actively exploring the ramifications of MPI and gradient fields in MFH (Tasci et al., 2009; Khandhar et al., 2012; Murase et al., 2013; Murase et al., 2015; Kuboyabu et al., 2016; Maruyama et al., 2016; Bauer et al., 2016; Behrends, Buzug, and Neumann, 2016).

Here we describe the construction of a combined MPI-MFH scanning system, which we use to explore and quantify the possibilities of an MPI theranostic system. We show that standard MPI hardware can guide MFH treatments to any desired region within a subject, providing heating with high spatial resolution as well as seamless switching between imaging and therapy while the subject remains in the scanner (Hensley et al., 2017).

4.1.1 Magnetic Fluid Hyperthermia

MFH encompasses an array of therapeutic approaches that use magnetic nanoparticles (MNPs) to couple magnetic energy into the body in the form of heat to treat diseased tissue (Jordan et al., 1999; Rosensweig, 2002; Jordan et al., 2006; Thiesen and Jordan, 2008). Ultimately, this heating is derived from non-adiabatic, lossy magnetic dynamics, or ‘relaxation’ in the language of the MPI field. Direct hyperthermia and tissue ablation that

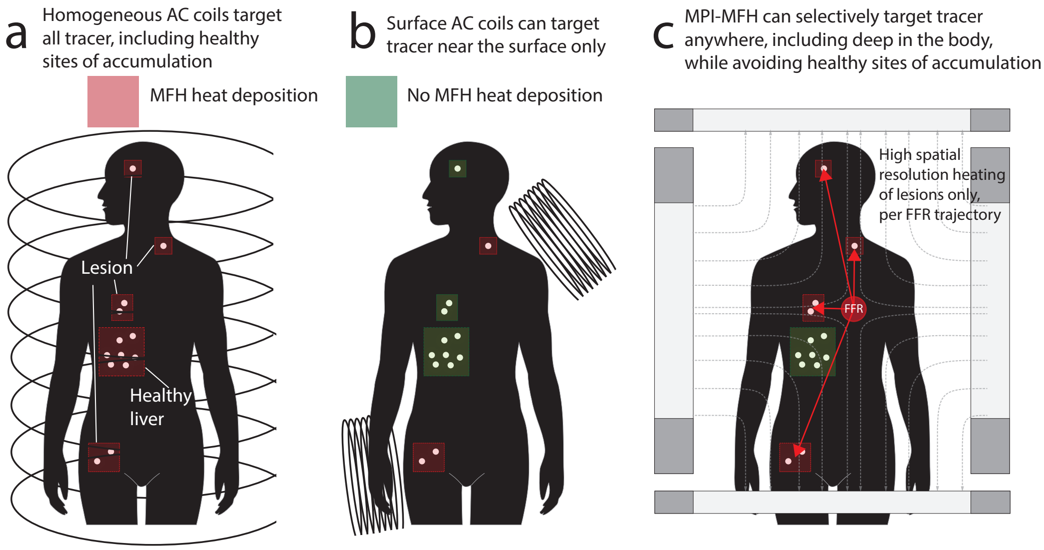

FIGURE 4.1: Localizing magnetic fluid hyperthermia. (a) Homogeneous AC coil designs apply therapy to all locations where the SPIOs have accumulated, including healthy sites. (b) Surface AC coils can be used to target sites of accumulation near the surface of the subject but not sites deep in the body. (c) MPI-MFH as proposed herein can target any site desired using the FFR, including those deep in the body, with high resolution. Healthy sites of accumulation can be avoided.

rely on large temperature changes have been studied in applications such as cancer treatment (Hilger et al., 2002; Johannsen et al., 2010; Branquinho et al., 2013). For tissue ablation therapies, a typical goal is to heat a target tissue region to above 48–50 °C to achieve the desired effect. However, heat-induced cell death is generally a complex function of time and temperature, motivating the use of more heuristic assessments. For example, the non-linear metric cumulative equivalent minutes at 43 °C (CEM_{43}) establishes a thermal dose for any time-temperature trajectory that tissue is subjected to (Sapareto and Dewey, 1984; Dewhirst et al., 2003; Dewey, 2009; Hensley et al., 2015a). In this context, a goal such as $\mathrm{CEM}_{43} = 240$ may be sought in the region of interest.

Other therapeutic techniques require increasing temperature by only a few degrees. For example, limited or mild *hyperthermia* therapy in conjunction with chemotherapy or radiation therapy has been explored as a way to enhance therapeutic potency (Torres-Lugo and Rinaldi, 2013; Maier-Hauff et al., 2011). More recently, research has indicated that no macroscopic temperature change may be necessary to induce cell death in targeted tissues (Creixell et al., 2011; Domenech et al., 2013). Additionally, MFH may be leveraged for drug delivery through, for example, selective breaking of thermally labile bonds or use of thermally-dependent permeabilities without detectable macroscopic temperature changes (Kumar and Mohammad, 2011; Zhang and Misra, 2007).

4.1.2 Spatial targeting of MNPs and MFH

Spatially targeting MFH therapies is a primary concern of the technique. There are two ways to spatially target the effects of MFH: targeting the tracer and targeting the energy deposition. We discuss both approaches separately.

Targeting the Tracer

MFH intrinsically provides spatial localization of the therapy via the distribution of the MNPs. A major challenge in MFH is ensuring high enough concentrations of MNPs at target sites to effect significant macroscopic temperature changes. Accordingly, many *in vivo* pre-clinical and clinical studies to date have used direct injection of tracer (*e.g.*, intratumoral) in lieu of effective active targeting mechanisms following systemic injection (Johannsen et al., 2010). In general, the required concentration is a function of the MNP properties, heat transfer conditions, and magnetic field parameters. It has been reported that 5 milligrams ferrite per gram of tumor, or approximately 5 $\mathrm{mg\,mL^{-1}}$ MNP concentration, is suitable for MFH using clinically acceptable magnetic field excitation conditions (Jordan et al., 2009). In one *in vivo* human trial to treat prostate cancer, the authors report direct injection of 12.5 mL of 120 $\mathrm{mg\,mL^{-1}}$ MNP solution (injected in smaller volumes at 24 different locations) into a 35 mL prostate (Johannsen et al., 2005). Assuming uniform distribution, that all of the injected MNP fluid remained in the prostate, and summing the injected and prostate volumes, we calculate that up to approximately 30 $\mathrm{mg\,mL^{-1}}$ MNP concentration was achieved in the target lesion. With these concentrations, temperatures in excess of 48 °C were achieved in the lesion using MFH.

A longer term goal is to leverage various *in vivo* MNP targeting mechanisms such that the therapy may be localized by the distribution of the particles following systemic introduction. However, high specificity targeting of nanoparticles and small molecules is an incredibly challenging unsolved problem. And regardless of the method of introduction, MNPs will to some degree make their way into non-target regions such as the liver. An ability to robustly and reliably target only those MNPs associated with a therapeutic region is an important ultimate requirement for any clinical MFH approach.

Targeting the Energy Deposition

Our ability to focus electromagnetic fields at the frequencies used in MFH (around 300 kHz) is fundamentally limited by diffraction, regardless of the number of external coils employed, to about half the *in vivo* wavelength. Unfortunately, at 300 kHz, this value is roughly 50 meters and much larger than the subject. Many MFH methods use large homogeneous coils as illustrated in Fig. 4.1(a). This leads to significant heat deposition at all sites where the MNP concentration is high, including healthy sites of accumulation. Alternatively, surface coils which do not provide a homogeneous excitation field can be used to target the heating, as depicted in Fig. 4.1(b). Due to the rapid decay of the magnetic field with distance from the coil, this approach can selectively target lesions near the surface of the subject, but it cannot deliver energy to particle distributions deeper in the body. There is currently no way in MFH to target specific regions of accumulated magnetic nanoparticles arbitrarily deep in the body. This is a serious clinical challenge since the non-specific uptake of even targeted MNPs is far higher in the excretory organs (liver, spleen, or kidneys) than in the targeted region (Wilhelm et al., 2016).

The use of a strong magnetic field gradient can potentially solve MFH spatial targeting, as illustrated in Fig. 4.1(c). A strong DC field fully saturates MNPs, locking them into alignment. A superposed AC excitation field is then unable to induce a rotation, which is required to generate heating. With an MPI gradient system in place, heating will be isolated to the FFR – the only location where MNPs are unsaturated. The spatial resolution of heating should scale similarly to the MPI spatial resolution since the localization mechanism is identical. In this manner, an MPI-MFH system can provide targeted heating with high resolution arbitrarily deep in the body.

4.1.3 Combining MPI-MFH

Combining MPI and MFH could lead to even greater gains than spatially targeted MFH by providing real-time feedback for more refined and safe therapy. Murase et al. recently showed that MPI can be used to predict the effect of subsequent MFH due to the common physics that generates the MPI signal and MFH heating (Murase et al., 2015). Combined MPI-MFH could enable continuous monitoring of tumor position, real-time quantitation of SAR deposition or temperature (Weaver, Rauwerdink, and Hansen, 2009; Stehning, Gleich, and Rahmer, 2016), and real-time assessment of treatment success. In this context, rigorous quantitation may require accounting for the effects of magnetic relaxation on the MPI signal and MFH therapy. For example, it is known that the MPI signal is affected by the excitation amplitude due to magnetic relaxation dynamics (Croft et al., 2016), and binding events when using targeted tracers may also change the nature of the signal.

Such real-time or simultaneous approaches are in contrast to current MFH approaches that are open loop or require invasive temperature probes, complex pre-planning heating simulations, or provide only limited imaging feedback (Laurent et al., 2011). Here we demonstrate a unified MPI-MFH system that may one day solve the problem of high resolution localization of MFH while providing seamless and non-invasive feedback *in vivo*.

4.2 Materials and Methods

In this work, a combined MPI-MFH system was built and several MPI systems previously built by our group were used. Simulations based on recent MFH models, modified to include the presence of gradient fields, were also performed.

4.2.1 MPI-MFH System

A field-free line (FFL) based combined MPI-MFH system was constructed with a 2.35 $\mathrm{T\,m^{-1}}$ gradient provided by a quadrupole NdFeB magnet array oriented around a cylindrical imaging bore as shown in Fig. 4.2(a) and Fig. 4.3(a) (Goodwill et al., 2012b; Konkle et al., 2013). A solid copper shield defines the imaging bore in which concentric solenoidal excitation (Tx) and MPI receive (Rx) coils were placed. A simple resonant filter chain was implemented as shown in Fig. 4.2(b), consisting of an initial impedance matching and low pass filter component and final resonant stage connected to the Tx coil. This system is resonant at f_0 = 353 kHz with an optimal real input impedance of 50 Ω seen by the power amplifier at f_0. A Tomco BT00500-AlphaS-CW amplifier (Tomco Technologies, South Australia) was used to power the system. This setup allowed for up to 20 mT excitation across a 95% homogeneity region of approximately 2 cm. A sample delivery and movement system was constructed using linear stages driven by stepper

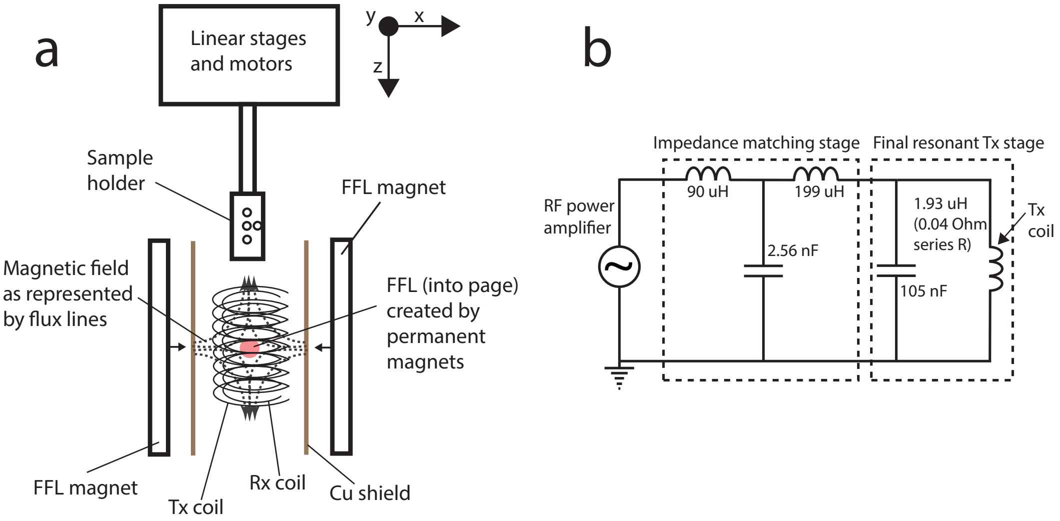

FIGURE 4.2: MPI-MFH system diagrams. (a) Diagram of the major structural components: permanent magnet array for creation of a field-free line (FFL), imaging bore with concentric transmit excitation (Tx) and MPI receive (Rx) coils, and sample movement system with linear stages and motors. (b) Circuit diagram for the 353 kHz resonant Tx chain with major components: impedance matching and low pass filter, final resonant Tx stage, and Tx coil.

motors in all three principal axes. In this manner, the position of a 3D-printed sample holder can be controlled precisely to realize desired MFH or MPI-MFH scan trajectories coded in MATLAB (MathWorks, MA, USA) scripts. All MFH and MPI-MFH scans using this system were taken in a one dimensional format, with relative movement between the sample and FFL confined along the axis of the cylindrical imaging bore and solenoidal Tx/Rx coils. This is the ‘z’ axis by our convention.

As in our previous MPI systems, the Tx coil was wound using hollow-core copper wire. A room temperature hydraulic circuit for cooling was installed inline with the coil to allow Joule heating of the coil to be removed without confounding the MPI-MFH experiments. A Gaymar T/Pump controlled temperature and pump system (Stryker Corporation, Kalamazoo, MI, USA) powered the cooling circuit, providing 0.5 liters per minute flow of room temperature water.

For MFH data acquisition, temperature data was collected using optical temperature probes designed for safe use in magnetic systems (Neoptix, Canada). These small probes were placed directly in contact with the tracer fluid in phantom vials containing MNPs in aqueous solution. Temperature data from these probes was sampled at 10 Hz using a MATLAB script and a standard serial communication interface.

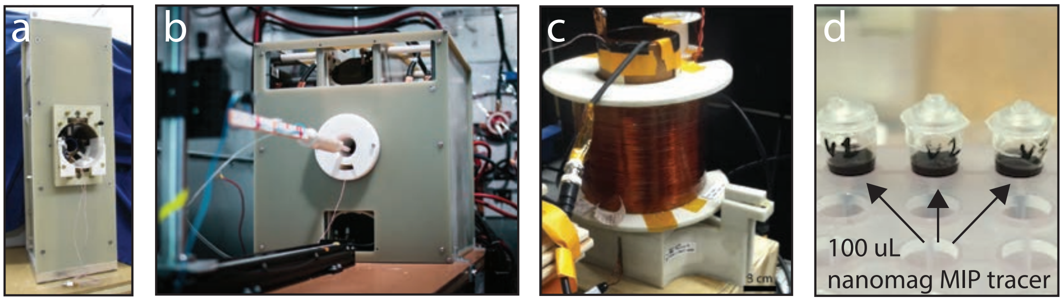

FIGURE 4.3: Hardware systems and vial phantoms used in this work. (a) Combined MPI-MFH system. (b) FFP scanning system used for multi-dimensional imaging. This system produces a 3.5 x 3.5 x 7 $[\mathrm{T\,m^{-1}}]$ gradient field with excitation at 20 kHz. (c) Arbitrary waveform relaxometer (AWR) table top MPI characterization system with no gradients. The AWR is capable of characterizing the 1D MPI PSF for a sample over a wide range of frequencies and field amplitudes. (d) Vial samples used in phantoms. 100 μL of nanomag MIP magnetic nanoparticles were placed in small PCR tubes.

For MPI data acquisition, an MPI Rx coil wound in a gradiometric fashion was placed internally concentric with the transmit coil as depicted in Fig. 4.2(a). As is standard in MPI, to maximize dynamic range in the receive chain, this coil was physically tuned for maximum cancellation of the transmit feedthrough that couples into the Rx coil due to the mutual inductance of the Tx and Rx coils. A second phase sniffer coil was also placed in the imaging bore to calibrate relative phase between the transmit and receive systems and ensure phase coherence in MPI analysis and reconstruction. The basic receive chain consisted of an SR560 low noise voltage preamplifier (Stanford Research Systems, Inc., CA, USA), and the raw data was sampled at 10 MHz using a PCI-6115 12-bit ADC (National Instruments, TX, USA). A Python script was used to analyze 1D MPI data.

To characterize the Tx coil and calibrate the input voltage to output field relationship, a Fluxtrol Magnetic AC field probe (Fluxtrol Inc., MI, USA) was used. To characterize the gradient field, a Lake Shore DSP 475 DC field probe (Lake Shore DSP, OH, USA) was used. These probes were placed in the sample holder and moved through the imaging bore using the linear stage and motor system for precise calibration.

4.2.2 MPI-MFH Experiments

Several types of MPI-MFH experiments were performed, all with phantoms constructed using samples of nanomag MIP SPIOs (78-00-102, micromod

Partikeltechnologie GmbH, Rostock, Germany) as shown in Fig. 4.3(d) (Eberbeck et al., 2013). These particles are polydisperse SPIOs coated with dextran (total hydrodynamic diameter in the range of 20 – 100 nm) and are similar to Ferucarbotran (Resovist) used widely in the MPI field. In some experiments, a single vial of nanomag MIP tracer (100 μL at 10.6 $\mathrm{mg\,mL^{-1}}$ iron) was placed in the sample holder, and an MFH sequence was applied using 20 mT excitation applied for 10–60 s with the sample fixed at a single position relative to the FFL. The 10.6 $\mathrm{mg\,mL^{-1}}$ is similar to clinically relevant MFH concentrations described in the literature (Jordan et al., 2009; Johannsen et al., 2005). This procedure was repeated after the sample was moved along the axis of the bore to sample the MFH signal as a function of position in one dimension. Replicate experiments were carried out to assess repeatability and report statistical variation.

In other experiments, the ability for repeatable and precise spatial selection in MFH was tested directly with a phantom containing three nanomag MIP vials (each 100 μL at 10.6 $\mathrm{mg\,mL^{-1}}$ iron) separated from each other by 3 mm (7 mm center-to-center distance). In some actuation sequences, a specific vial was individually targeted for heating while in others, each vial was targeted for heating in rapid succession. Replicate experiments were carried out to assess repeatability and report statistical variation.

In a third set of experiments, simultaneous MFH and MPI was performed. The setup and trajectory were the same as reported above (single nanomag MIP vial actuated at different points along the imaging bore), but MPI data was recorded. Data was also taken for a PBS control subjected to the same scan trajectory. In analyzing the MPI signal, baseline removal was performed using the control data and linear baseline subtraction.

The imaging phantom used in MPI-MFH experiments was also separately imaged in MPI-only devices to assess the quality of the phantoms in standard MPI. The field-free point (FFP) scanner shown in Fig. 4.3(b) was used for multi-dimensional MPI imaging of the phantoms. This scanner has a 3.5 x 3.5 x 7 $\mathrm{T\,m^{-1}}$ gradient system with excitation at $f_0 = 20\ \mathrm{kHz}$ and up to 30 mT. An amplitude of 20 mT was used in this work. Full 3D tomographic data sets took approximately 10 minutes to acquire and reconstruct.

In addition to the MPI-MFH and MPI-only experiments, the nanomag MIP sample vials were also tested on our arbitrary waveform relaxometer (AWR) system. This table-top system has no gradient fields. Instead, a sinusoidal excitation and linear bias field are superposed to test the aggregate response of a sample in the applied magnetic field space (Tay et al., 2016). This system can reconstruct 1D MPI point-spread functions (PSFs) (Croft et al., 2016; Tay et al., 2016). The AWR was used to test the quality of MPI at the higher frequencies used in MFH. Experiments were run with f_0 = 353 kHz to match the MPI-MFH system and an excitation amplitude of 2 mT.

4.2.3 Simulation

Simulations were carried out to predict the response of particles to AC fields in the presence of a static magnetic field gradient. The simulations were based on the theoretical model developed by Dhavalikar and Rinaldi (Dhavalikar and Rinaldi, 2016) which was constructed using the magnetization relaxation equation described by Martsenyuk et al. (Martsenyuk, Raikher, and Shliomis, 1974) to calculate heat dissipation as specific absorption rate (SAR). SAR values are calculated directly from a thermodynamic model wherein the work done by an applied magnetic field on the particles is dissipated as heat.

The nanomag-MIP particles primarily consist of 19 nm core particles. These cores are arranged as clusters in dextran with hydrodynamic diameters in the range of 20–100 nm (Eberbeck et al., 2013). In simulation, an iron oxide domain magnetization of 446 $\mathrm{kA\,m^{-1}}$ (Rosensweig, 2013) and density of 5.18 $\mathrm{g\,cm^{-1}}$ were used (Narasimha, Brennan, and Holtham, 2006; Rosensweig, 2002). The thickness of the dextran shell was varied to obtain SAR values comparable to those in the experiments. Interparticle interactions were not included in the model. For an aqueous solution at room temperature, as in the experiments, a viscosity of 0.89 $\mathrm{mPa\,s}$ and temperature of 298 K were assumed. Simulations were carried out an excitation amplitude of 20 mT and an excitation frequency of 353 kHz to match the experimental setup. To assess the effect of field gradient on the spatial distribution of SAR, simulations were performed with MPI field gradients in the range of 1–7 $\mathrm{T\,m^{-1}}$ based on gradient strengths typically encountered in existing MPI scanners.

4.2.4 SAR Calculations

To calculate experimental SAR, we used a common approximate method based on a linear fit to the initial temperature rise (Murase et al., 2013):

$$SAR = \left(\frac{\Delta T}{\Delta t}\right)_{\mathrm{i}} \cdot \frac{C_{\mathrm{w}} m_{\mathrm{w}} + C_{\mathrm{p}} m_{\mathrm{p}}}{m_{\mathrm{p}}} \tag{4.1}$$

where ΔT is the change in temperature in units of K or $^{\circ}\mathrm{C}$ in time Δt $[\mathrm{s}]$ over some interval i. $C_{\mathrm{w}} = 4190\ \mathrm{J\,kg^{-1}\,K^{-1}}$ is the specific heat capacity of water, m_w is the dimensionless relative mass of water in the sample, $C_p = 670\ \mathrm{J\,kg^{-1}\,K^{-1}}$ is the specific heat capacity of the magnetic nanoparticle (Maenosono and Saita, 2006), and m_p is the dimensionless relative mass of the magnetic nanoparticle. The units of SAR reported herein are $\mathrm{W\,g^{-1}}$. The relative masses were calculated from the following absolute values in each phantom vial: a $9.92\mathrm{e}^{-5}\ \mathrm{kg}$ water mass and $1.06\mathrm{e}^{-6}\ \mathrm{kg}$ tracer mass.

4.3 Results

4.3.1 Quantification of Spatial Localization of Heating

Fig. 4.4 shows 1D simulation and experimental MPI-MFH results. In Fig. 4.4(a), simulations of an MFH theoretical model including a spatial gradient field component are provided for f_0 = 353 kHz and at gradient strengths of 2.35, 4, and 7 $\mathrm{T\,m^{-1}}$. Qualitative agreement between the simulated SAR profile (heating PSF) and experimental SAR data were obtained by modeling an 8 nm dextran shell surrounding the 19 nm core. The predicted SAR FWHM for a thermal point source using a 2.35 $\mathrm{T\,m^{-1}}$ gradient is approximately 10 mm.

Fig. 4.4(b) shows experimental SAR and initial rate-of-change temperature data as a function of position for a single vial phantom and a PBS control using the MPI-MFH system. Data was recorded for excitation across a 2 cm FOV in steps of 1 mm with three trials per location. MFH was applied for 10 seconds at each location. A peak SAR deposition of about 150 $\mathrm{W\,g^{-1}}$ and a rate of temperature change just under 0.4 $\mathrm{{}^\circ C\,s^{-1}}$ were observed when the sample was coincident with the center of the FFL. The SAR full-width-at-half-maximum (FWHM) is approximately 7.5 mm.

Fig. 4.4(c) plots experimental SAR and MPI signal using the MPI-MFH system to scan a single vial sample. These data were obtained using the MPI-MFH device at f_0 = 353 kHz and 20 mT excitation as in all other heating scans. The MPI signal from the MPI-MFH system is reported as the baseline subtracted magnitude of the 3^rd harmonic of the excitation frequency in the time domain data associated with the discrete mean locations of the FFL per the MPI-MFH trajectory. Both data sets are normalized to their peak values for comparison.

Fig. 4.4(d) plots a 1D PSF of one of the sample vials constructed from data taken with the AWR. The PSF FWHM is approximately 3 mm and was taken using a lower 2 mT excitation amplitude at 353 kHz, in contrast to the 20 mT excitation used in the MPI-MFH device.

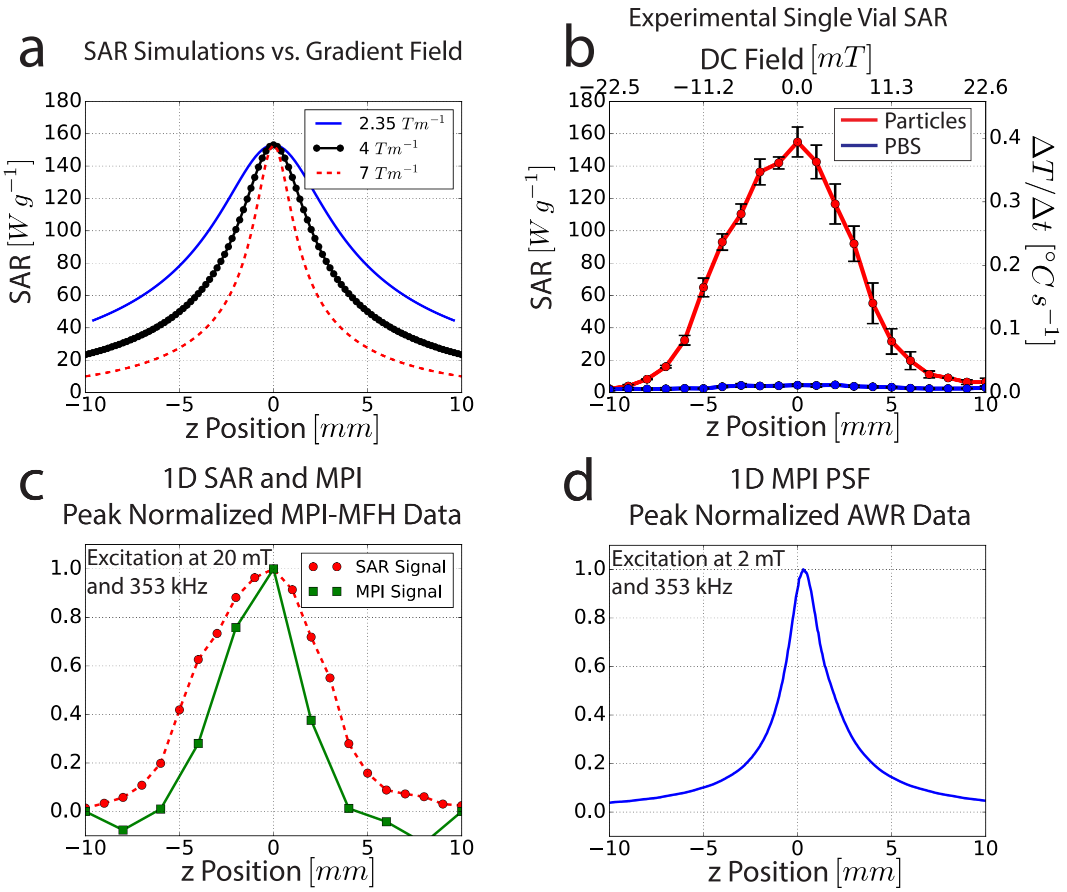

FIGURE 4.4: Spatial localization of heating with MPI-MFH in 1D. (a) Simulations of specific absorption rate (SAR) for the SPIOs used in this work at various gradient fields. (b) Experimental SAR data using the MPI-MFH system and applied to both a single vial with SPIOs and a PBS control. (c) Normalized plots of experimental SAR and experimental MPI signal (magnitude of the 3rd harmonic) from the MPI-MFH system. (d) Experimental 1D MPI PSF from the AWR. All experimental data and simulations use 353 kHz excitation.

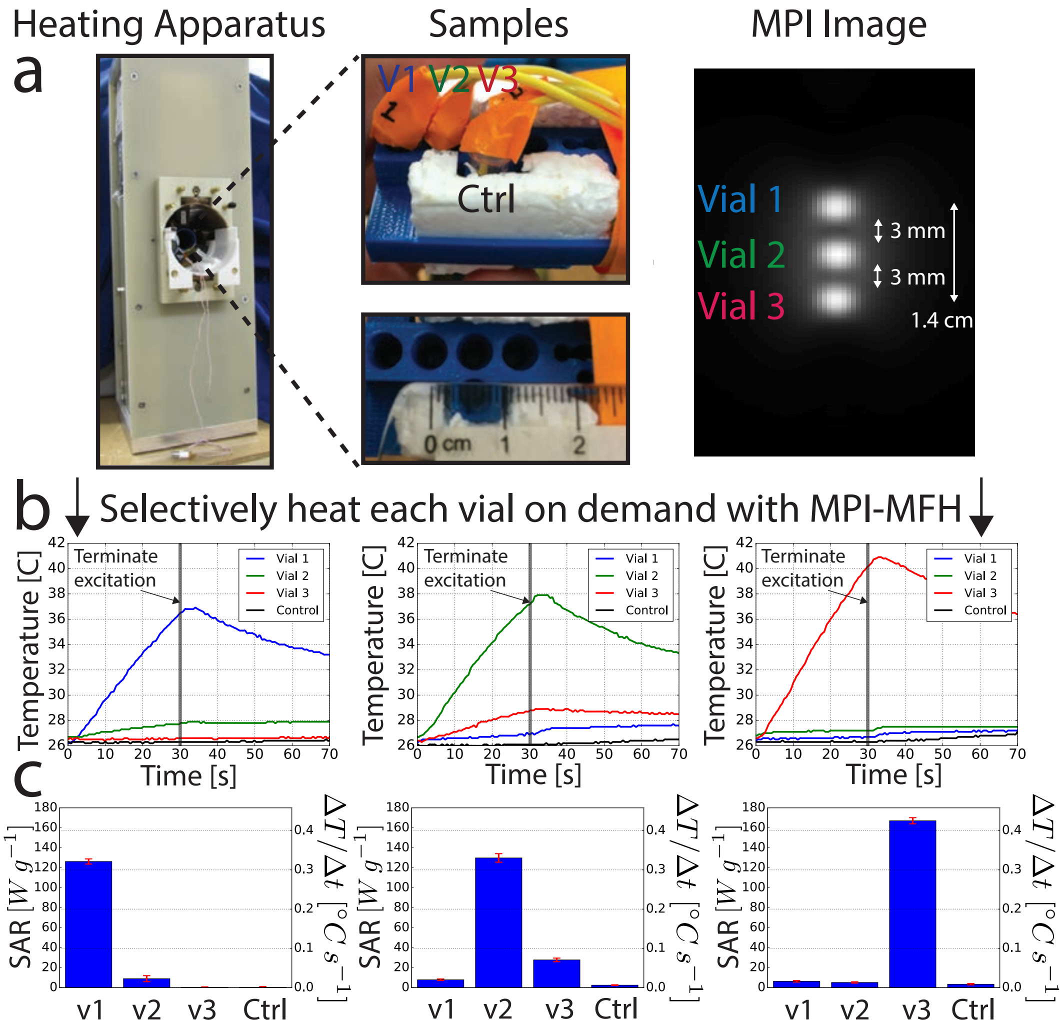

FIGURE 4.5: Selective heating with MPI-MFH. (a) MPI-MFH device, SPIO phantom, and MPI image taken with a FFP scanner. The phantom vials were 3 mm apart (7 mm center-to-center). (b) MPI-MFH temperature data for trials when each of the three phantom components was targeted for individual heating. (c) SAR and temperature data quantified for all 5 trials associated with targeting each of the 3 phantom vials.

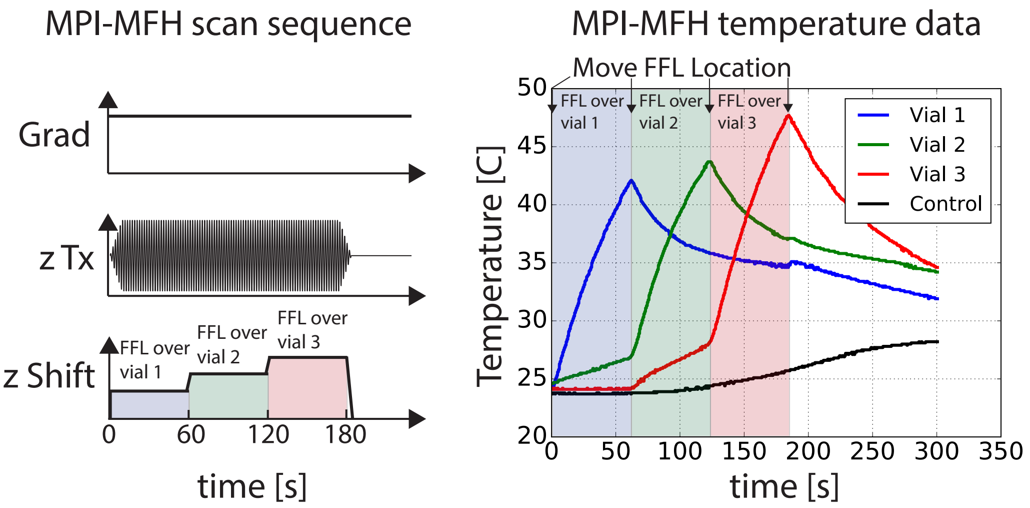

FIGURE 4.6: Implementation of an MPI-MFH scanning sequence to actuate multiple distinct locations in a single therapeutic scan. While AC excitation is applied, the localizing FFL is sequentially aligned with each of the three vials in the phantom to realize temporo-spatially selective heating. The MPI-MFH method provides high resolution, flexible, and dynamically configurable therapeutic targeting.

4.3.2 Selective Heating and Imaging of a Phantom

Fig. 4.5 shows MPI and MPI-MFH data using a triple-vial phantom constructed with the sample vials in Fig. 4.3(d). The phantom vials were 3 mm apart from each other (7 mm center-to-center distance). Fig. 4.5(a) shows the phantom imaged at 20 kHz and 20 mT excitation using the FFP scanner of Fig. 4.3(b). Data acquisition took about 10 minutes and a maximum intensity projection is shown.

Fig. 4.5(b,c) show data using the MPI-MFH system to individually target each of the three vials of the phantom for heating in separate trials. Fig. 4.5(b) shows temperature data from individual trials. In all cases, the vial that was targeted rose in temperature significantly, around 0.3 – 0.4 $^\circ\mathrm{C}\,\mathrm{s}^{-1}$, while the other vials showed little or no rise. As indicated in Fig. 4.5(b), actuation was performed for an initial 30 second period during which a total temperature increase of 10 – 15 $^\circ\mathrm{C}$ was observed in the targeted vial each time. After termination of the AC excitation, heating ceased. Total temperature increases in non-targeted vials was in the range of 0 – 2.5 $^\circ\mathrm{C}$, generally less than 1 $^\circ\mathrm{C}$.

Fig. 4.5(c) shows temperature and SAR data quantified and with statistical standard deviation reported for all 5 trials associated with the targeting of each of the 3 vials. The rates of heating of the targeted vials are similar to the SAR experiments of Fig. 4.4(b), in the range of 130 – 170 $\mathrm{W\,g^{-1}}$ and $(\Delta T/\Delta t)_{\mathrm{i}}$ in the range of 0.3 – 0.4 $^{\circ}\mathrm{C\,s^{-1}}$.

Fig. 4.6 shows the use of a longer MPI-MFH scanning sequence targeting each of the three vials sequentially for 60 seconds along with the resulting temperature data. Each vial was heated approximately 20 $^{\circ}\mathrm{C}$ over a 60 second actuation period.

4.4 Discussion

4.4.1 MFH Localization

In this chapter, we explored the use of MPI magnetic gradient fields to localize SAR heating. This approach enables spatially localized MFH arbitrarily deep in the body with high resolution. Fig. 4.4 and Fig. 4.5 show data that quantify localized heating and demonstrate that SAR spatial resolution can be modeled, with a strong dependence on the strength of the gradient field. Although the nanomag MIP particles are polydisperse and have a mixed Brownian and Néel character in terms of magnetic relaxation, the simulation results with our method qualitatively agree with the experimental data. In the future, simulations may be valuable as a pre-planning tool in MPI-MFH therapeutic applications.

The shape of the simulated SAR profiles is very similar to that of the 1D experimental MPI PSF (Langevin-derivative or Lorentzian-like), which is expected given that the same nonlinear saturation physics mediate both MPI and MFH. The shape of the experimental SAR profile, however, is more rounded. This is also to be expected because the 2 mm-wide vials used in the experimental phantom are not good approximations of point sources. For both the simulated SAR and 1D MPI PSF data, however, a point source was assumed and realized, respectively.

The measured FWHM for the SAR deposition data is about 7.5 mm, and because the sample vial is 2 mm wide, we estimate a heating FWHM of approximately 5.5 mm. The phantom experiments shown in Fig. 4.5 provide a direct empirical measure of the available thermal resolution in MPI-MFH. When targeting the center vial of the triple-vial phantom, with each vial placed 3 mm away from the adjacent vials, there was modest heating of one of the adjacent vials, indicating we are nearing the thermal resolution limit. In that specific case, the adjacent vial rose in temperature by about 2 $^{\circ}\mathrm{C}$ during the 30 second actuation, while the the targeted vial rose over 10 $^{\circ}\mathrm{C}$.

The data also shows that, when actuating a vial at an end of the phantom, the vial at the other end (10 mm away, 14 mm center-to-center) does not heat.

These data confirms our estimate of approximately a 5.5 mm thermal resolution. If we extrapolate from a 2.35 $\mathrm{T\,m^{-1}}$ gradient to the 7 $\mathrm{T\,m^{-1}}$ gradient in our FFP MPI scanner, we estimate a SAR FWHM of 1.8 mm. In the context of future clinical therapies, a 'thermal resolution' in the range of 5.5 mm, as demonstrated in this work, should be more than adequate to avoid heating sites such as the liver while targeting various lesions. If electromagnet gradients are used, the gradient strength can be increased when higher thermal resolution is needed and decreased when spatial separation is large. Lower gradients will allow for more efficient heating of larger lesions as there is a direct relationship between heating efficiency and FFR size.

We believe that the consistently higher heating of vial 3 when compared to vials 1 and 2 (approximately 30%) in the experiments in Fig. 4.5 are primarily due to the sensitivity of alignment that comes with MPI-based spatial localization in MPI-MFH. In general, there is a tradeoff between thermal resolution and alignment sensitivity (as well as heating efficiency) due to the spatial localizing effect of the FFR. In these data, there were degrees of freedom in both the x and z axes. Imperfect FFL-vial alignment in either or both axes will lead to less than ideal heating of the vial, and this was observed in initial experimentation. Possible sources of relative misalignment include imperfect shift trajectories in z and imperfect vial positioning or tilt in the phantom, a possible confound in both the x and z axes. In general, this motivates the use of a combined MPI-MFH system capable of parallel or simultaneous and real-time imaging to ensure proper alignment during therapy. This is especially true for high gradient and therefore high resolution localized MFH. Differing heat transfer conditions, thermal probe placement in the vial fluid, and minor pipetting differences during sample creation could also have contributed to the observed variation.

For MFH actuation of duration on the order of 10 – 30 s, as performed in Fig. 4.4 and Fig. 4.5, control vials containing PBS show essentially no heating and confirm that Joule heating in the transmit coil system and subsequent conduction or convection heat transfer is not responsible for the heating observed within the tracer vials. However, in Fig. 4.6, we note heating in the PBS control vial beginning at 150 s. This is due to the hardware limitations of the setup which eventually leads to heat transfer from the Tx coil and/or thermal cross-talk between the neighboring vials.

The small standard deviations in our experimental heating data show that the system performs consistently. While the sequentially actuated data shown in Fig. 4.6 represents a basic MPI-MFH sequence, it foreshadows how precise x-space trajectories may be leveraged in this theranostics platform.

4.4.2 Simultaneous MPI-MFH

The 1D MPI PSF data of Fig. 4.4(d) was taken with a tabletop AWR at the same frequency $f_0 = 353\ \mathrm{kHz}$ as the MPI-MFH system but at an amplitude of 2 mT. A 2 mT excitation amplitude was used due to power limitations with the device but is also relevant because it represents how one might seamlessly mode switch between heating and imaging with this platform. At 2 mT, detectable macroscopic heating is generally not possible even with high tracer concentrations because of the two orders of magnitude reduction in heating compared to MFH at 20 mT. Thus dialing up and down the excitation amplitude, in addition to varying FFR trajectories, represents a simple and seamless way of switching between heating and imaging modes.

The narrower shape of the MPI PSF compared to the SAR data could be due to some fundamental differences between the two types of data but may also be explained by the large difference in excitation amplitude. We have demonstrated a strong dependence of the MPI PSF FWHM on the excitation amplitude in the past (Croft et al., 2016) due to the impact of magnetic relaxation, changing slew rates, and the differing size of the partial FOV covered by each oscillation of the excitation field. In general, smaller amplitudes lead to higher resolution with all other variables being constant (Croft et al., 2016). We therefore posit that changing the experimental drive field from 2 to 20 mT would likewise result in broader MPI PSFs, but this experiment is outside the hardware capabilities of the AWR. Improved MPI capabilities of the combined MPI-MFH device in the future will allow for this analysis directly. It may also be of interest to explore if the shape of the SAR deposition changes analogously with changing excitation amplitude, independent of magnitude scaling.

The 3^{rd} harmonic MPI signal from the MPI-MFH device shown in Fig. 4.4(c), taken while applying a scanning heating sequence to a single vial, is a proof-of-concept demonstration of simultaneous MPI-MFH. These data align qualitatively well with the SAR data and suggest real-time SAR quantitation and thermal dose quantification via the MPI signal may be obtained in a straightforward manner. These data are an important first step en route to using real-time, quantitative feedback in MPI-MFH.

4.4.3 MNP Concentrations in Relation to Clinical Applications

In the literature, a ‘moderate’ concentration for clinical MFH of 5 $\mathrm{mg\,mL^{-1}}$ has been reported (Jordan et al., 2009), and in one clinical MFH pilot study, we calculated that a prostate tissue MNP concentration of up to 30 $\mathrm{mg\,mL^{-1}}$ was established during the therapy (Johannsen et al., 2005). In comparison, a MNP concentration of 10.6 $\mathrm{mg\,mL^{-1}}$ was used in this work, which fits within this range of clinical MFH concentrations. Using this concentration, we showed a repeated ability to achieve a 10–15 $^{\circ}\mathrm{C}$ temperature rise within 30 seconds. Given these results, we believe heating up to and beyond the 43 $^{\circ}\mathrm{C}$ desired in ablation applications is possible with this approach in an *in vivo* clinical setting. However, active blood perfusion will reduce achievable heating rates in this context. Potentially much lower concentrations could be used in MPI-MFH applications where only mild hyperthermia therapy is desired or in applications where no macroscopic temperature change is required (*e.g.*, drug delivery applications).

We can also compare the MNP concentration used in this work to clinical intravenous iron therapy, as relevant for systemically introduced tracer in diagnostic MPI and MPI-MFH. A 2010 study reported on the safety of Ferumoxytol iron therapy introduced systemically by intravenous injection (Lu et al., 2010). The Ferumoxytol was administered as two 510 mg injections separated by 3 – 8 days. If we consider direct injection into a 35 mL target lesion, as in a prostate treated in the previously mentioned MFH clinical pilot study, then the 10.6 $\mathrm{mg\,mL^{-1}}$ concentration used in this work would be equivalent to approximately 371 mg of total iron.

We conclude that the use of MNPs in this work is within the clinically relevant range, both in terms of concentration for MFH therapy and total mass in systemic delivery. However, future studies should elaborate on these and other clinical concerns.

4.5 Conclusions and Future Work

Here we have described the theoretical foundation, physical construction, and testing of an MPI-MFH theranostics platform. We showed that the SAR predicted by theoretical models agrees well with experimental data and may be the basis of future treatment planning optimization strategies. We showed the first experimental data using MPI gradients to deliver targeted heating on demand to components of a phantom with an ability to selectively heat targets separated by as little as 3 mm. We also showed the ability to serially target the component vials using an MPI-MFH sequential scanning sequence. We demonstrated that MPI imaging still works at the high frequencies used in MFH through the use of the AWR. Last, we demonstrated simultaneous MPI-MFH by measuring the MPI signal when using the MPI-MFH device in a heating scan applied to a single vial phantom. Together, these data represent an important step in the development of a theranostics platform for combined MPI-MFH.

Work on combining MPI and MFH has just begun, and we expect development to accelerate in the coming years. An important next step is to explore combined MPI-MFH with *in vivo* applications work such as targeted cancer therapy. Exploring novel therapeutic approaches that do not require macroscopic temperature changes, such as activation of lysosomal pathways and thermal drug delivery, are also of great interest.

Although this work demonstrated the MPI-MFH principle in the context of inducing large macroscopic temperature changes, we believe this may not be the best avenue for MPI-MFH in the future. The high frequency and large amplitudes required to accomplish such a therapy at the scale of a human body are confounded by the dependence of non-specific SAR deposition which scales strongly with increasing size. Furthermore, achieving the high concentrations required, even with such high frequency and high amplitude excitation, is not generally possible without direct injection of tracer. In contrast, we believe therapies that do not reqruire macroscopic temperature changes may be much more promising. For example, consider the case of a drug placed on a long-circulating SPIO tracer via thermally- (microscopic scale) or mechanically-labile bonds. Myriad alternative systems such as the case where actuation increases the porosity and releases entrapped drug may also be used. If we know a desired region or lesion location where we would like to deposit drug, we can move the FFR to this location and excite. Actuation, and in this case, drug delivery will only occur when the functionalized SPIOs pass through the region through normal circulation. Importantly, drug may be released in a highly targeted and spatially resolved manner but without *any* form of biochemical active targeting. Instead, the onus here is on a drug entrapment scheme that releases drug only with MPI-MFH actuation with high fidelity. We believe these type of opportunities, especially as made possibly by spatially selective MPI-MFH actuation, are particularly interesting to explore.

4.6 Acknowledgements

I would like to thank Patrick Goodwill for assistance in getting this project started, especially for helping with the design and building of early MPI-MFH prototypes not featured in this chapter, but from which I learned a great deal and enabled me with respect to this work. I would also like to thank Zhi Wei Tay for helping with the construction and testing of some of the MPI-MFH hardware featured in this chapter. Finally, I would like to thank Rohan Dhavalikar and Professor Carlos Rinaldi for helpful discussions and for providing all of the simulated data described in section 4.2.3 and shown in Fig. 4.4(a).

Part II

Pulsed Magnetic Particle Imaging

Chapter 5

Square Wave pulsed MPI in 1D

Engineering efforts are rapidly improving current methods and probing new implementations in MPI. The Langevin physics most fundamental to MPI predicts improving resolution with increasing tracer magnetic core size. As described in Section I, for larger tracers, and given typical AC excitation slew rates used in the field, this steady-state prediction is compromised by dynamic processes. These dynamics induce a secondary blur that leads to an optimal core size and prevents realization of the Langevin promise for larger tracers. Here we propose, investigate, and evaluate a new method of excitation and signal encoding in MPI we call pulsed MPI (pMPI) to overcome this phenomenon. pMPI allows us to directly encode the steady-state Langevin physics into the time-domain signal. This in turn gives rise to a simple reconstruction algorithm to realize images free of secondary relaxation-induced blur. The core feature of the pMPI approach is the use of excitation waveforms containing periods of temporally static applied field, separated by potentially rapidly varying periods as in the case of a mathematical square wave. Rapid transitions between these field states can allow for significant temporal separation of the time-varying excitation and the tracer signal, mitigating the feedthrough issues that plague current continuous-wave, sinusoidal MPI. Here we provide a detailed description of the theory behind our pMPI approach in 1D and show experimental data demonstrating better than 500 μm resolution using 'large' magnetic core tracers.

5.1 Introduction

Beginning in 2005, published work has described the physical and mathematical nature of the MPI signal (Rahmer et al., 2009; Lu et al., 2013), the system matrix encoding and reconstruction approach (Rahmer et al., 2009; Weizenecker et al., 2009), the x-space encoding and reconstruction approach

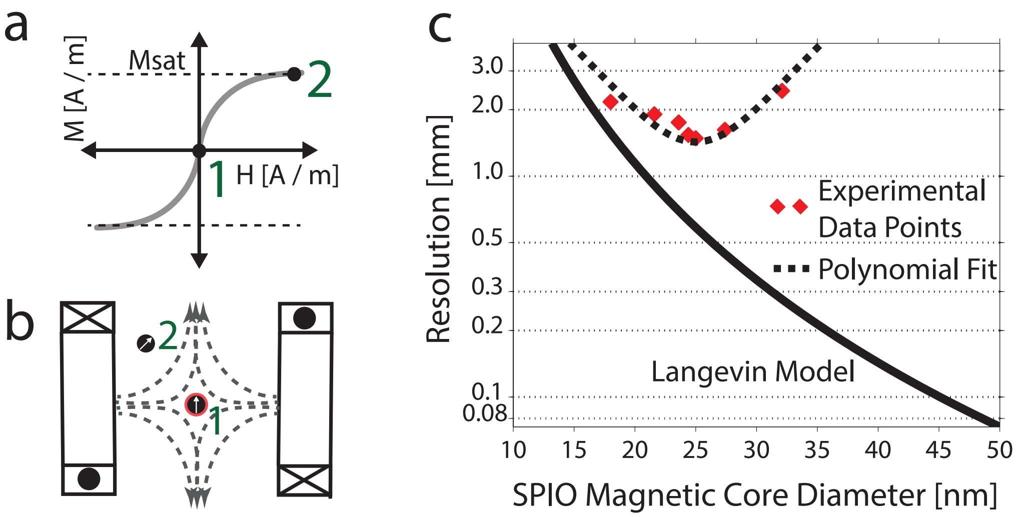

FIGURE 5.1: Langevin theory and limitations in experimental continuous wave MPI. (a) SPIO tracers used in MPI are characterized by saturation. (b) In MPI this provides spatial localization of our signal via sensitive 'field-free' regions (FFRs) constructed using gradient fields. (c) The Langevin theory is a steady-state theory that predicts monotonically improving resolution with magnetic core size. Experimentally, this is observed up to an inflection point around 25 nm, past which resolution degrades due to effects of magnetic dynamics or relaxation. Data was taken with our arbitrary waveform relaxometer (AWR) using $f_0 = 20\ \mathrm{kHz}$ excitation frequency and $20\ \mathrm{mT}$ amplitude.

(Goodwill and Conolly, 2010; Goodwill and Conolly, 2011), and other distinct implementations of MPI (Vogel et al., 2014; Goodwill et al., 2009). All are governed by the Langevin physics illustrated in Fig. 5.1(a,b) that describe the steepness of the tracer's M-H curve, which translates directly to imaging resolution (Goodwill and Conolly, 2010). The M-H curve steepness monotonically increases and thus imaging resolution monotonically improves as a function of magnetic core size according to Langevin theory. Increasingly, however, the impact of tracer magnetic dynamics on the imaging process as explored in Chapter 2 has become a focus of the field (Ferguson, Minard, and Krishnan, 2009; Biederer et al., 2009; Rauwerdink and Weaver, 2010; Goodwill et al., 2011; Croft et al., 2016; Dhavalikar and Rinaldi, 2016). Langevin physics is strictly a steady-state model and can only rigorously model the MPI process if the dynamic response of the tracer is always much faster than the slew rates used in excitation. For small Néel-dominant tracers, this may be very well approximated. However, as the magnetic core size of a tracer increases, the magnetic response time increases as well, regardless of the mechanism (*e.g.*, Néel or Brownian) (Bauer et al., 2015; Deissler, Wu, and Martens, 2014).

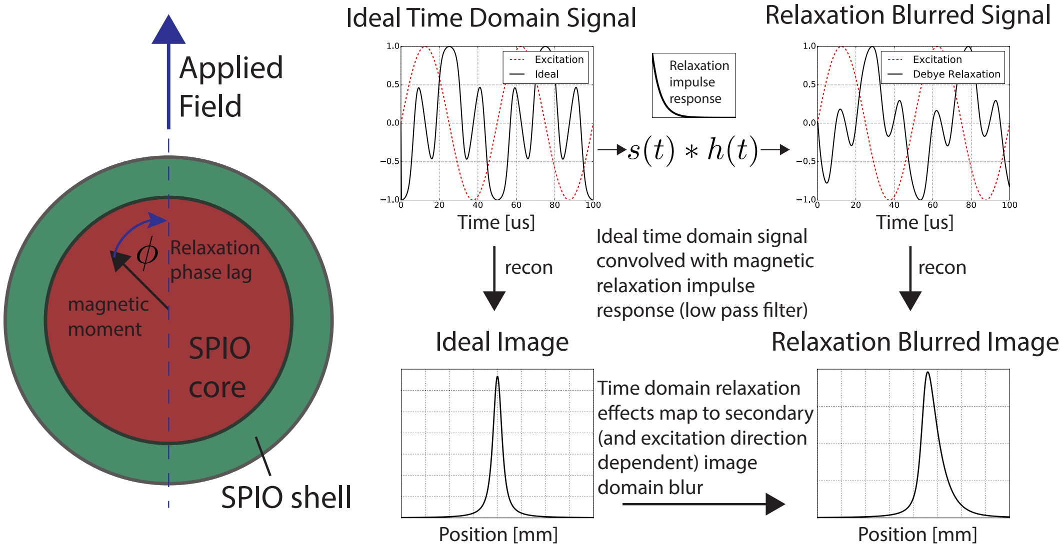

FIGURE 5.2: Magnetic relaxation effects on the MPI signal. The MPI tracer cannot respond instantaneously to the continuously varying AC excitation in canonical MPI. When the time constant(s) associated with the magnetic relaxation impulse response is nontrivial compared to the excitation period, non-negligible relaxation delays and phase lags manifest. We can model this as a convolution of the ideal steady-state tracer response with a relaxation kernel. Because of the direct mapping of time domain signal to space in MPI reconstruction, this time-domain low pass filter maps directly to a low pass filter in the image domain. As a result, image domain translations and reduced resolution manifest.

Fig. 5.2 describes how finite and non-negligible relaxation time constants induce secondary relaxation blurs in MPI. Intuitively, relaxation can be understood as a relaxation phase lag as the tracer struggles to keep up with a continuously varying applied field. Mathematically, this can be modeled as a time domain convolution between the ideal steady-state signal and a relaxation kernel (*e.g.*, a Debye exponential kernel). Regardless of the details, the relaxation kernel will generally act as a low-pass filter. From an x-space MPI perspective, this temporal low-pass filter translates to an image domain low-pass filter due to the direct mapping between the two domains. From the Fourier domain or system matrix perspective, the higher order harmonics that contain high spatial resolution information are attenuated. As a result, relaxation manifests as a secondary blur in reconstructed images, leading to compromised resolution.

With typical excitation parameters in standard MPI, we have found that improved resolution with increasing magnetic core size follows the Langevin prediction up to approximately 25 nm. After this inflection point, we see reduced performance and worsening resolution as shown in Figs. 2.3 and 5.1(c) and which is in direct conflict with the Langevin theory. At this point, relaxation dynamics start to play an increasingly important roll, breaking the steady-state assumption and leading to significant secondary blurring as depicted in Fig. 5.2. The precise location of this inflection point is dependent on various tracer properties and the magnetic slew rates used in excitation as illustrated in Fig. 2.3, but accounting for the effect of dynamic physics in MPI data acquisition as currently performed generally establishes a ‘relaxation wall’ that limits the achievable resolution in MPI (Tay et al., 2017).

While magnetic relaxation typically has a deleterious effect on achievable resolution in canonical MPI, it can also be leveraged to provide powerful contrast and molecular imaging capabilities. For example, in Chapter 3 we showed the ability to provide colorized contrast and unmixing when multiplexing multiple tracer species, and in Chapter 4 we showed the ability to combine MPI with magnetic fluid hyperthermia to build a new theranostics platform that relies on the lossy, non-instantaneous nature of the tracer response. These and related concepts such as sensing temperature have recently been described in the literature (Cho, Cho, and Cho, 2014; Perreard et al., 2014; Hensley et al., 2015b; Rahmer et al., 2015; Stehning, Gleich, and Rahmer, 2016; Murase et al., 2013; Hensley et al., 2017).

Many of these newer developments as well as the distinction between x-space and system matrix MPI can be described in terms of differences in signal encoding. This is possible because there is great flexibility in how the FFR can be used to sample an imaging field-of-view (FOV). And paramount for the work in the next few chapters, the way in which relaxation information is encoded in the raw signal is a strong function of the excitation waveform.

Here we propose a new excitation approach for MPI which allows direct encoding of steady-state Langevin information into the acquired time domain signal. This enables reconstruction of images that asymptotically approach the theoretical Langevin resolution, even for large particles with significant magnetic response times. It also provides an ability to temporally separate excitation feedthrough from the MPI signal.

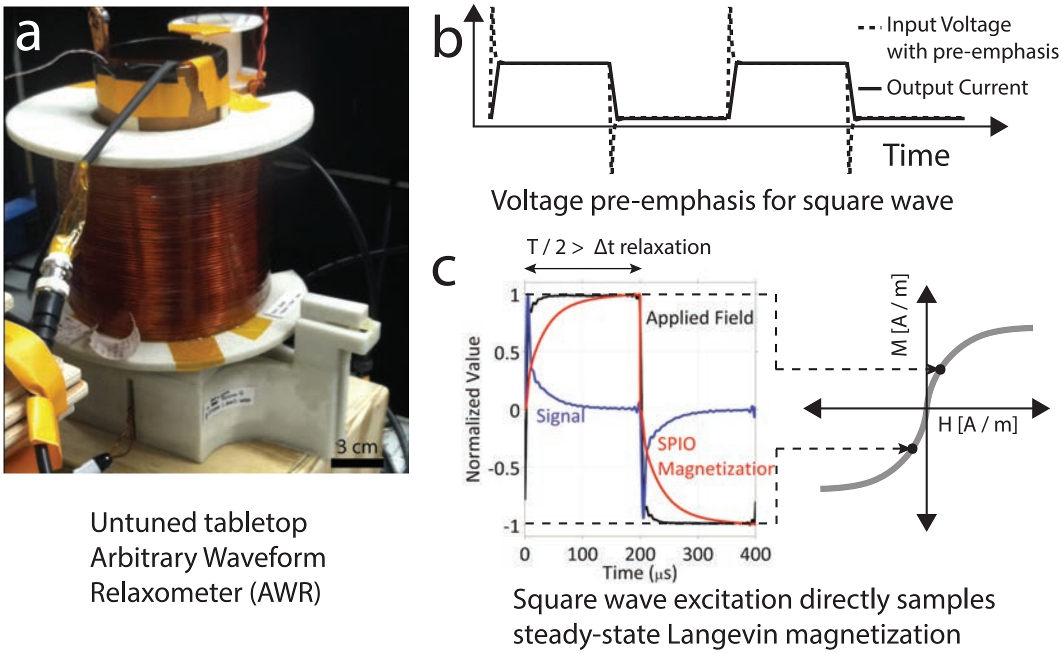

Untuned tabletop
Arbitrary Waveform
Relaxometer (AWR)

Voltage pre-emphasis for square wave

Square wave excitation directly samples
steady-state Langevin magnetization

FIGURE 5.3: Square wave excitation in pulsed MPI. (a) Arbitrary waveform relaxometer (AWR) used in this work. The AWR contains an untuned excitation coil and DC bias coil that can realize arbitrary waveforms for 1D pMPI pulse sequence testing and analysis. (b) Pre-emphasized voltage waveforms sent to linear amplifiers in voltage-control mode are used to create square and trapezoidal waveforms. (c) The period of the square wave is chosen such that the tracer achieves steady-state magnetization by the end of each half-period. In this manner, we sample points on the Langevin curve of the tracer, encoding this steady-state information in our raw signal. Experimental raw signal from AWR square wave excitation applied to 32 nm core tracer is shown.

5.2 Pulsed MPI Theory and Implementation in 1D

The crux of pMPI is use of non-sinusoidal waveform components in MPI excitation trajectories. We can construct pulsed waveforms from two primitive waveform components: 1) periods of temporally static applied field and 2) periods of rapidly varying applied field. The former allows encoding of steady-state information in the raw signal while the latter provides high peak time domain SNR as well as temporal separability between the tracer signal and excitation feedthrough. These components can be assembled into larger periodic excitation waveforms or pulse sequences. One of the mathematically simplest cases is a square wave as depicted in Figs. 5.3 and 5.4.

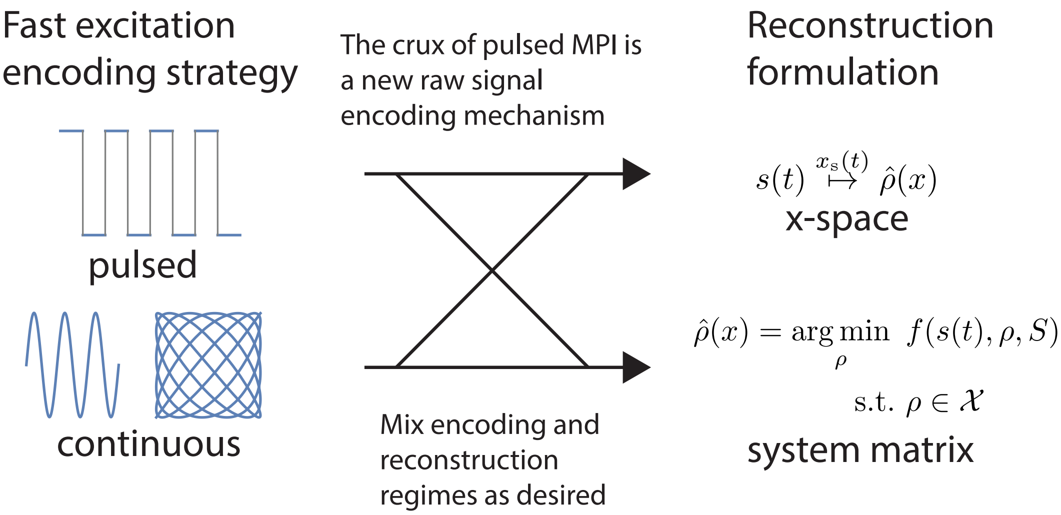

FIGURE 5.4: Pulsed encoding in MPI acquisition. The fast excitation strategy is responsible for how the fundamental tracer signal is captured. pMPI replaces the canonical continuous wave sinusoidal fast excitation waveform with pulsed waveforms. Subsequent aspects of the acquisition scheme such as large FOV sampling strategy and reconstruction method may be chosen as desired as long as the pMPI constraints described herein are respected. For example, slew rates involved in FOV sampling cannot compromise steady-state induction. As an encoding back-end, pMPI may be incorporated into any of the MPI imaging paradigms, such as the popular x-space or system matrix methods.

Pulsed MPI is a complete departure from typical continuous wave excitation sequences used in the field, requiring new transmit and receive electronics. Other aspects of the typical MPI system infrastructure can be retained; however, such as the gradient-producing magnets and shift mechanisms that allow sampling of a large FOV over time in a canonical x-space acquisition (Goodwill et al., 2012a). Furthermore, while pMPI herein is described in an x-space context, as Fig. 5.4 indicates, pMPI can be thought of as an encoding back-end that can be incorporated into other acquisition and reconstruction paradigms such as the system matrix methods.

In this chapter, we again leverage our AWR system shown in Fig. 5.3(a) to demonstrate square and trapezoidal wave pMPI. The AWR is an untuned tabletop MPI system capable of producing a broad array of excitation waveforms (Tay et al., 2016). We can superpose a slowly ramping homogeneous bias field to measure 1D point-spread functions (PSFs).

5.2.1 Magnetic Relaxation

Much of Section I of this thesis is concerned with MPI magnetic relaxation, and generally this is a very active area of research (Ferguson, Minard, and Krishnan, 2009; Biederer et al., 2009; Rauwerdink and Weaver, 2010; Goodwill et al., 2011; Croft et al., 2016; Dhavalikar and Rinaldi, 2016). As illustrated in Fig. 2.1, multiple relaxation processes may be active simultaneously, including Néel, Brownian, and ferromagnetic processes. All of these processes have different dependencies on applied field conditions, different processes may be coupled in interesting ways, and nonlinear interactions between particle domains may occur.

In continuous-wave MPI, significant relaxation leads to signal and image degradation in both x-space and system matrix methods. Relaxation becomes more prominent, and therefore the negative effect on resolution becomes relatively worse, the larger the tracer. The effect is stark when comparing predictions of Langevin theory with experimental results as shown in Fig. 5.1(c). In the following chapter we will explore some detailed theoretical models of magnetic relaxation in MPI.

In this chapter, however, we emphasize how pMPI encoding mitigates the deleterious effect of relaxation on resolution in a *general way*: To realize the benefits of pMPI, we require only 1) the assumption that magnetic relaxation processes have a finite step response, which is guaranteed for all real physical processes, and 2) superparamagnetism, which is typically assumed for MPI in general. More specifically, for the latter, we assume no asymmetry in steady-state magnetization trajectories (hysteresis) as exists in coercive ferromagnetic systems. On this point, we note that, in the context of large populations of nanoparticles free to rotate independently on long time scales, *steady-state* behavior will appear superparamagnetic and symmetric in aggregate even if some (*e.g.*, large tracers) exhibit some individual ferromagnetic behavior.

The experimental data of Fig. 5.3(c) shows the decaying signal elicited after the rapid transitions of a typical square wave excitation. In this context, and as we will show below, the ramifications of pMPI encoding should pertain for a great variety of possible MPI tracers *regardless* of the specific nature of their magnetic relaxation physics. We believe this is a powerful statement given the complexity and clearly incomplete understanding we currently have of magnetic relaxation in MPI.

5.2.2 Square Wave Signal and Imaging Equations

Here we develop the 1D signal and imaging equations for pMPI with ideal square wave excitation. We begin by defining the applied field in terms of a field-free point (FFP) as described in canonical x-space theory (Goodwill and Conolly, 2010):

$$H(t,x) = H_0(t) - Gx \tag{5.1}$$

where $H(t,x)$ is the total applied field in $[\mathrm{A\,m^{-1}}]$ composed of separable superposed field contributions: a time-varying but spatially homogeneous field $H_0(t)$ and a temporally static but spatially varying linear gradient field Gx parameterized by G, the linear gradient strength in $[\mathrm{A\,m^{-2}}]$.

We then express the field as a function of the time varying FFP isocenter, $x_\mathrm{s}(t)$ in $[\mathrm{m}]$, as follows:

$$\begin{aligned} H(t,x) = 0 &= H_0(t) - Gx_\mathrm{s}(t) \\ H_0(t) &= Gx_\mathrm{s}(t) \\ x_\mathrm{s}(t) &= G^{-1}H_0(t) \end{aligned} \tag{5.2}$$

followed by substitution of expression (5.2) into equation (5.1):

$$H(t,x) = H_0(t) - Gx = G(G^{-1}H_0(t) - x) = G(x_\mathrm{s}(t) - x) \tag{5.3}$$

The steady-state magnetization of a 1D tracer distribution exposed to this field is given by:

$$M_\mathrm{ss}(t,x) = m\rho(x)\mathcal{L}(kH(t,x)) \tag{5.4}$$

where m is the magnetic moment of the tracer in $[\mathrm{A\,m^2}]$, $\rho(x)$ is the tracer distribution number density in $[\mathrm{particles\,m^{-3}}]$, $\mathcal{L}$ is the Langevin function, and k is a tracer-specific parameter with units of $[\mathrm{m\,A^{-1}}]$ that characterizes the relationship between magnetic and thermal energies:

$$k = \frac{\mu_0 M_\mathrm{sat}\pi d^3}{6k_\mathrm{B}T} \tag{5.5}$$

where d is the diameter of the tracer magnetic core in $[\mathrm{m}]$ and M_sat is the saturation magnetization of the tracer in $[\mathrm{A\,m^{-1}}]$.

We now consider excitation with an ideal square wave:

$$x_{\mathrm{s}}(t) = \sum_{n=0}^{k} (-1)^n \, \Delta x \, \sqcap \left(\frac{t - n\Delta t}{\Delta t} \right) \tag{5.6}$$

with x-space displacement $\Delta x = AG^{-1}\mu_0^{-1}$ where A is the square wave applied magnetic field amplitude in [T], Δt is the square wave half-period in [s], and $\sqcap$ is the rectangular function (Gaskill, 1978).

To derive the signal equation, we assume use of an inductive receive coil with homogeneous sensitivity B_1 [T A^{-1}] in the imaging FOV which leads to a volume integral (Hoult and Richards, 1976). Considering only the case of 1D and leaving B_1 to carry the proper sign/polarity relationship, the voltage on the receiver coil is:

$$\begin{aligned} V(t) = B_1 \frac{\mathrm{d}\Phi(t)}{\mathrm{d}t} &= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \iiint_V M(t, \mathbf{u}) \, \mathrm{d}V \\ &= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \iiint_V M(t, u)\delta(v)\delta(w) \, \mathrm{d}u \, \mathrm{d}v \, \mathrm{d}w \\ &= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \int_u M(t, u) \, \mathrm{d}u \end{aligned} \tag{5.7}$$

To proceed, we assume that the tracer magnetic relaxation can be described as a linear time-invariant (LTI) process with impulse response $h_{\mathrm{m}}(t, x)$ and related step response $r_{\mathrm{m}}(t, x)$. We let these responses vary as a function of x because magnetic relaxation, regardless of the specific mechanism, is typically a strong function of applied field which in the context of MPI varies over space due to our use of gradient fields. We do note, however, that these expressions are normalized such that $\bar{h}_{\mathrm{m}} = \int_{-\infty}^{\infty} h_{\mathrm{m}}(t, x) \, \mathrm{d}t = 1 \; \forall x$. We keep the appropriate scalars separate and explicit. We now consider the voltage on the receiver coil over one half-period of the square wave.

Let $t_{n-1} = (n-1)\Delta t$, $t_n = n\Delta t$, $M_{\mathrm{ss}}(t_{n-1} < t < t_n, x) = M_{n-1}(x)$, $M_{\mathrm{ss}}(t_n < t < t_{n+1}, x) = M_n(x)$, and $\Delta M_n(x) = M_n(x) - M_{n-1}(x)$. Assuming that steady-state conditions are achieved at the end of each previous $(n-1)$-th square wave half-period, the evolution of the magnetization over the next n-th half-period is mathematically identical to a step response:

$$M_{\mathrm{ss}}(t, x) \Big|_{t_n}^{t_{n+1}} = M_{n-1}(x) + \mathrm{hv}(t - t_n)\Delta M_n(x) \tag{5.8}$$

where $\mathrm{hv}(t)$ is the temporal heaviside function. We now specify a signal equation $s_n(t)$, evaluated for each square wave half-period per the index n:

$$\begin{aligned}
s_n(t) &= V(t)\Big|_{t_n}^{t_{n+1}} \\
&= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \int_u M(t,u)\,\mathrm{d}u \Big|_{t_n}^{t_{n+1}} \\
&= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \int_u M_{\mathrm{ss}}(t,u) \overset{t}{*} h_{\mathrm{m}}(t, u-u_n)\,\mathrm{d}u \Big|_{t_n}^{t_{n+1}} \\
&= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \int_u (M_{n-1}(u) + \mathrm{hv}(t-t_n)\Delta M_n(u)) \overset{t}{*} h_m(t, u-u_n)\,\mathrm{d}u \Big|_{t_n}^{t_{n+1}} \\
&= B_1 \frac{\mathrm{d}}{\mathrm{d}t} \int_u \bar{h}_{\mathrm{m}}(u-u_n) M_{n-1}(u) + \Delta M_n(u) r_m(t-t_n, u-u_n)\,\mathrm{d}u \Big|_{t_n}^{t_{n+1}} \\
&= B_1 \int_u \frac{\mathrm{d}}{\mathrm{d}t} \left(1 \cdot M_{n-1}(u) + \Delta M_n(u) r_{\mathrm{m}}(t-t_n, u-u_n)\right)\,\mathrm{d}u \Big|_{t_n}^{t_{n+1}} \\
&= B_1 \int_u \Delta M_n(u) h_{\mathrm{m}}(t-t_n, u-u_n)\,\mathrm{d}u \Big|_{t_n}^{t_{n+1}}
\end{aligned} \tag{5.9}$$

where $h_{\mathrm{m}}(x-x_n)$ indicates that the field-dependent relaxation impulse response is centered about the FFR location x_n during the n-th square wave exctiation. The expansion of the magnetization in terms of the steady-state magnetization distributions as shown assumes that the square wave half-period is chosen to be longer than a parameter of merit that quantifies the convergence of $h_{\mathrm{m}}(t,x)$. For example, in the case of a Debye process (Croft et al., 2016), let the half-period be longer than 5 * τ_{max} where $\tau_{\mathrm{max}} = \max_x \tau(x)$ is the maximum observed time constant for the MPI tracer under MPI excitation conditions (*i.e.*, the zero field relaxation time constant).

The signal equation is composed of chunks with a very specific interpretation: The signal associated with each half-period n is a weighted form of the tracer magnetic impulse responses integrated over space and scaled by the change in steady-state magnetization distribution. This steady-state magnetization change depends on the underlying tracer distribution and the square wave amplitude or step size. For tracer at any given location, and comparing the signal at the beginning of the half period with that at the end (in which steady-state is established by definition), we are sampling two distinct points on the Langevin curve as illustrated in Fig. 5.3(c) and the direct encoding of the steady-state Langevin magnetization is explicit.

With this fundamental signal equation, we now consider an approach to achieve an imaging equation where we integrate the time-domain signal associated with each half-period. This value is then associated with the center location for each half period (half way between the initial and final location of each square wave step):

$$\begin{aligned}
s(n) &= \int_{t_n}^{t_{n+1}} s_n(t)\,\mathrm{d}t \\
&= B_1 \int_{t_n}^{t_{n+1}} \left(\int_u \Delta M_n(u) h_\mathrm{m}(t - t_n, u - u_n)\,\mathrm{d}u \right) \mathrm{d}t \\
&= B_1 \int_u \Delta M_n(u) \left(\int_{t_n}^{t_{n+1}} h_\mathrm{m}(t - t_n, u - u_n)\,\mathrm{d}t \right) \mathrm{d}u \\
&= B_1 \int_u \Delta M_n(u)\,\mathrm{d}u \\
&= B_1 \int_u \left(M_n(u) - M_{n-1}(u) \right) \mathrm{d}u \\
&= B_1 \int_u \left(m\rho(u)\mathcal{L}(kG(u - u_n)) - m\rho(u)\mathcal{L}(kG(u - u_{n-1})) \right) \mathrm{d}u \\
&= B_1 m\rho(x) \overset{x}{*} \mathcal{L}(kGx)|_{x=x_n} - B_1 m\rho(x) \overset{x}{*} \mathcal{L}(kGx)|_{x=x_{n-1}} \\
&= B_1 m\rho(x) \overset{x}{*} \left(\mathcal{L}(kGx)|_{x=x_n} - \mathcal{L}(kGx)|_{x=x_{n-1}} \right) \\
&= B_1 m\rho(x) \overset{x}{*} \left(\dot{\mathcal{L}}(kGx) \overset{x}{*} \sqcap\left(\frac{x}{2\Delta x}\right)\Big|_{x=x_n-\Delta x} \right) \\
&= B_1 m\rho(x) \overset{x}{*} \dot{\mathcal{L}}(kGx) \overset{x}{*} \sqcap\left(\frac{x}{2\Delta x}\right)\Big|_{x=x_n-\Delta x} \qquad (5.10)
\end{aligned}$$

where $2\Delta x = x_n - x_{n-1}$ and in the last three lines we have used the fundamental theorem of calculus and definition of $\sqcap$ to formulate a spatial convolution from the difference of a function evaluated at two distinct points. Finally, we can use this discrete signal equation to form a 1D square wave pMPI imaging equation:

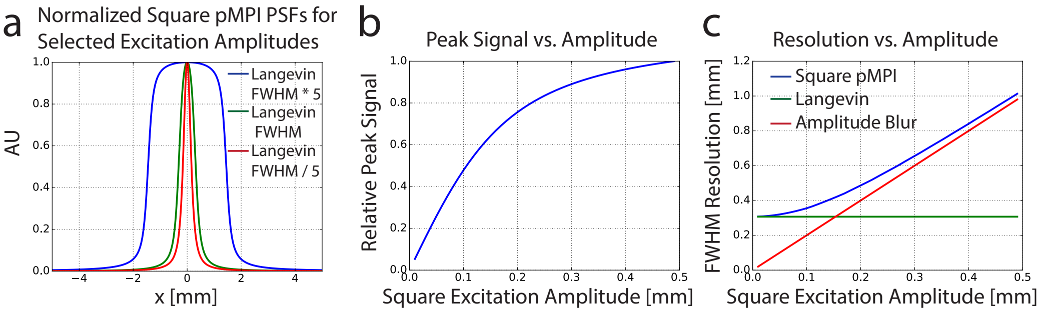

FIGURE 5.5: Theoretical pMPI square wave PSF analysis assuming a 32 nm magnetic core tracer. (a) Peak-normalized PSFs for the cases of square excitation with an amplitude of 5 times the Langevin full-width-at-half-maximum (FWHM), equal to the Langevin FWHM, and one-fifth of the Langevin FWHM. The rectangular function's blurring effect is apparent in the large excitation. Larger excitations also lead to increased signal. (b) Relative peak signal as a function of the excitation amplitude. (c) PSF FWHM resolution for square wave pMPI as formulated herein plotted along with the Langevin resolution and the rectangular function width.

$$\begin{aligned}\hat{\rho}(x) &= \frac{s(n)}{B_1 m}\\ &= \rho(x) \overset{x}{*} \dot{\mathcal{L}}(kGx) \overset{x}{*} \sqcap\left(\frac{x}{2\Delta x}\right)\Bigg|_{x=x_n-\Delta x}\\ &= \rho(x) \overset{x}{*} h_{\dot{\mathcal{L}}}(x) \overset{x}{*} h_r(x) = \rho(x) \overset{x}{*} h(x)\end{aligned} \tag{5.11}$$

Our 1D square wave pMPI PSF, $h(x)$, is then the steady-state Langevin PSF, $h_{\dot{\mathcal{L}}}(x)$, convolved with a rectangular function, $h_r(x)$, defined by the square wave excitation amplitude Δx. Fig. 5.5(a) shows peak-normalized PSFs as a function of various excitation amplitudes. We summarize some ramifications of this new imaging equation:

- This result is general for any steady-state superparamagnetic tracer with LTI relaxation physics characterized by an impulse response that converges to zero, such as the Debye model (Goodwill et al., 2011; Croft et al., 2016).
- The results are predicated on choosing an excitation half-period longer than some measure of the maximum observable impulse response time, such as $5 \cdot \tau_{\max}$ in the case of an exponential response with time constant τ.

- The results of (5.10) and (5.11) hold for non-square waveforms, such as a trapezoidal waveform, that contain a flat, time-invariant section longer than the aforementioned measure of the maximum observable impulse response time. In this case, however, (5.9) does not hold for the signal equation, but a more general temporal convolution form does.
- We obtain resolution related directly to the steady-state Langevin function and an excitation parameter under our direct control: the excitation amplitude.
- In this formulation, we can only asymptotically approach the Langevin resolution in the limit of small excitation amplitudes.
- The choice of excitation amplitude provides a convenient and continuous tradeoff between SNR and resolution.

5.2.3 Resolution

Rigorously, we can solve for the full-width-at-half-maximum (FWHM) of the square wave pMPI PSF in equation (5.11) by solving for x in the following expression:

$$\mathcal{L}(kG(x+\Delta x)) - \mathcal{L}(kG(x-\Delta x)) = \frac{\mathcal{L}(kG\Delta x) - \mathcal{L}(-kG\Delta x)}{2}$$

This expression is difficult to solve analytically, although it is easily numerically visualized as in Fig. 5.5(c). The square wave pMPI PSF is composed of two components: the steady-state derivative of the tracer Langevin function and a rectangular function with width equal to twice the square wave amplitude (the peak-to-peak amplitude). The FWHM of the derivative of the Langevin function has been reported (Goodwill and Conolly, 2010) and the FWHM of a rectangular function is just its width. Fig. 5.5(a) shows how the PSF changes as a function of the excitation amplitude and Fig. 5.5(c) quantifies how the square wave pMPI FWHM resolution varies with excitation amplitude. The Langevin FWHM and rectangular function blur are also shown. Qualitatively we note:

$$\mathrm{FWHM} \approx \begin{cases} \mathrm{FWHM}_{\dot{\mathcal{L}}} & \text{if } \Delta x \ll \mathrm{FWHM}_{\dot{\mathcal{L}}} \\ 2\Delta x & \text{if } \Delta x \gg \mathrm{FWHM}_{\dot{\mathcal{L}}} \\ \mathrm{FWHM}_{\dot{\mathcal{L}}} + \Delta x & \text{otherwise} \end{cases} \tag{5.12}$$

where $\Delta x = AG^{-1}\mu_0^{-1}$ is the square wave amplitude in x-space.

When using this, to significantly realize the gains in resolution that larger tracers promise via the steady-state Langevin curve, we must limit the square wave excitation amplitude.

5.2.4 SNR and Bandwidth

The signal equation and receive bandwidth of square wave pMPI are both explicitly related to the tracer impulse response. Lacking a generic and closed form expression for tracer impulse responses, we do not provide explicit expressions for SNR. However, we can qualitatively assess the ramifications of equations (5.9), (5.10), and (5.11). Furthermore, we will assess experimental SNR and bandwidth relations in section 5.4 and in the following chapter.

The time-domain signal equation (5.9) states that SNR in square wave pMPI at any given time is directly proportional to the tracer magnetic moment, mass of the tracer, and the difference between two points (separated by the square wave excitation amplitude) along a tracer's Langevin curve. This leads to similar conclusions as in traditional (relaxation-free) MPI analysis: tracers with high magnetic moment and steeper M-H curves lead to higher SNR (Goodwill and Conolly, 2010).

Unlike most treatments of MPI, we see that the magnetic impulse response of the tracer is incorporated explicitly into the raw signal equation. Tracers with faster relaxation will exhibit a higher peak time domain SNR. This also has ramifications for the required bandwidth. In the case of an ideal square wave excitation, just as in the case of sinusoidal MPI, and due to the combination of excitation periodicity and the non-linear nature of the Langevin function, Fourier analysis reveals that the received signal should be confined to harmonics of the fundamental square wave frequency. In this case, however, we are exciting at more than one tone – at multiple harmonics of a fundamental tone. The received bandwidth of interest, as in sinusoidal is then composed of small bandpass regions around some N harmonics, where the number of harmonics with support above the noise floor depends on relaxation dynamics and system parameters. Additionally, the distribution of energy across the harmonic bands will depend on the nature of the magnetic relaxation. In Chapter 6, we will explore more about emergent spatio-temporal and spatio-frequency relationships that are exploitable in pMPI. As in canonical MPI, the size of the bandpass regions required around each harmonic depends on the slew rates used in the slow shift or bias coils that translate the FFR to cover a large imaging FOV in the general case.

SNR Efficiency

The discrete signal and imaging equations (5.10) and (5.11) integrate the time-domain signal over each half-period, collapsing this information to a single value in which the impulse response integrates to unity given our assumptions. Thus the image domain signal intensity is directly proportional to the difference in magnetization between the two points sampled by a step excitation, regardless of the relaxation dynamics. However, long tracer responses associated with the larger tracers that have better Langevin resolution will require a relatively long half-period to reach steady-state. Furthermore, to achieve high resolution with this approach, we must keep the excitation amplitude small to keep the rectangular blur component of the PSF small. This combination of small amplitude excitation and long fundamental excitation period can lead to reduced SNR efficiency compared with canonical sinusoidal MPI using smaller core tracers. For optimal SNR and SNR-efficiency in this context, one should carefully choose the excitation half-period, making it no longer than necessary. Optimization of tracer shell properties (*e.g.*, to minimize time constants for large Brownian tracers) may play a crucial roll in square wave pMPI performance. Furthermore, fully establishing steady-state (*e.g.*, in a Debye $5 * \tau$ sense) may not be necessary, and, as in the case of standard MPI, resolution-bandwidth tradeoffs are manifest. We will explore experimentally some of the ramifications of breaking our assumption that the quiescent half-period be longer than the relaxation response. Fig. 5.5(b) quantifies the relationship between peak image domain signal and excitation amplitude and Fig. 5.5 (b) and (c) together illustrate how the choice of excitation amplitude represents a continuous SNR-resolution tradeoff.

We summarize some considerations when comparing SNR efficiency between continuous wave MPI and square wave pMPI as presented here, especially keeping ΔM per Tx oscillation per unit time in mind:

Square wave pMPI advantages:

- No loss of signal associated with the fundamental harmonic due to separability of Tx feedthrough and tracer signal.
- No need to trim Rx signal near sinusoidal transition regions (this is needed to avoid noise gain in the x-space velocity compensation step when the excitation sinusoid transitions through zero velocity).
- No need to separately reconstruct and inability to combine images associated with each excitation direction (*e.g.*, negative and positive slewing directions in single direction Tx). This improves SNR by $\sqrt{2}$.

Possible square wave pMPI disadvantages:

- SNR efficiency loss when steady-state hold times lead to square wave f_0 much slower than traditional sinusoidal f_0. Note this is not actually pMPI specific, as the same is true for sinusoidal or other methods that must also slow the excitation slew rates to cope with relaxation physics when using large tracers and targeting high resolution.
- Using amplitudes much lower than sinusoidal amplitudes for high resolution can greatly reduce SNR efficiency (ΔM per oscillation).

Comparing square wave pMPI and sinusoidal MPI depends greatly on the tracer regime one considers, *e.g.*, comparisons for very large tracers (square wave will be superior in resolution and SNR), comparisons at tracer sizes ideal for sinusoidal MPI (square wave will likely be approximately equivalent in resolution but much worse in SNR). To bridge the gap, perhaps one useful metric for capturing the utility of square wave pMPI over sinusoidal MPI is SNR efficiency change per unit resolution improvement when each is considered using its 'ideal' tracer. We will explore aspects of square wave pMPI SNR experimentally in this chapter but leave a rigorous exploration of SNR and SNR efficiency comparison for future work.

5.3 Materials and Methods

To experimentally assess our pMPI approach, we tested various MPI tracers with different magnetic core sizes using our AWR system. We assessed the relationship between encoding parameters such as excitation amplitude, pulsed rise time, and pulsed hold time as well as image domain parameters of merit such as FWHM resolution and peak signal. We plot and compare experimental results with the theoretical predictions presented in section 5.2. We also empirically explore the ramifications of breaking the basic pMPI steady-state constraint by using quiescent periods not long enough to fully establish steady-state.

5.3.1 Implementing pMPI Excitation Pulse Sequences

Implementation of a square wave, or pMPI pulse sequences in general, requires transmit and receive systems that differ from that used in standard sinusoidal MPI. The latter are typically built using resonant systems in both the transmit and receive electronics. The use of resonance improves available reactive power at the single excitation frequency while relatively attenuating noise and interference contributions at other frequencies. Because the excitation is continuous with reception, direct transmit feedthrough usually necessitates a receive side filter to remove energy at the fundamental frequency, in addition to broadband common mode interference and feedthrough mitigation strategies such as use of a gradiometer receive coil(s).

pMPI pulse sequences, however, require broadband excitation with relatively high frequency components. Given magnetic relaxation times encountered with available MPI tracers, a square wave rise time in the range of 1–5 μs has been typical for us. As indicated in equation (5.11), smaller excitation amplitudes are desired for high resolution in the approach outlined in this chapter, which helps constrain the required slew rates.

This type of high power pulse sequence design has been addressed in many previous application spaces (Scott et al., 2001). Our AWR provides high bandwidth excitation by using a low inductance receive coil and benefits from a small sample bore (Tay et al., 2016). With this system, we can use a linear amplifier in voltage-controlled mode with pre-emphasized waveforms to achieve the magnetic slew rates we need in pMPI. Other approaches, such as use of timed resonant switcher circuits are also possible.

The low-pass filter effect of an electromagnet system is well understood through Maxwell's equations. To implement pMPI pulse sequences, we can use this knowledge to design an input voltage waveform that will achieve the desired field, subject to our constraints such as the power amplifier maximum output voltage and maximum voltage slew rate. In the context of these constraints, the coil inductance and series resistance are important parameters to optimize.

For many of our square wave and trapezoidal excitation experiments, we have computed the voltage waveforms by solving the following convex optimization:

$$\begin{aligned} \min_{x} \quad & ||Ax - y||_2^2 \\ \text{s.t.} \quad & x \in \mathcal{X} \end{aligned} \tag{5.13}$$

where the matrix A is the physics forward model (including RLC circuit analysis and knowledge of the coil sensitivity) relating the magnetic field produced by the coil (inductor) in response to an input voltage, x is the voltage waveform, y is the desired or target field waveform, such as a square wave, and $\mathcal{X}$ is the constraint set imposed on x, including the maximum voltage and maximum voltage slew rate.

Solutions to (5.13) are guaranteed optimal in an ℓ_2-norm sense and subject to uncertainty and error in construction of A. This general formulation can be used to build any desired pMPI pulse sequences. Typically, our excitation waveforms are periodic, as we apply the same excitation as we raster across a FOV. This means we only need to solve the pre-emphasis convex optimization for a single period, making the computation simple and fast, even for long scans. To minimize error that does result from imperfect modeling of A, output from the optimization problem is sometimes corrected using a small number of iterations of a simple online negative feedback control script.

5.3.2 Data Acquisition

In our experiments, we used a set of monodisperse SPIO tracers all purchased from the same supplier (Imagion Biosystems, Ltd., Albuquerque, NM, USA) with the following core sizes: 18.5, 21.6, 24.4, 27.4, and 32.1 nm. We subjected samples of each tracer to a number of different encoding and excitation conditions to build parametric data sets. These included excitation amplitude (range: $[1, 20]$ mT), steady-state hold time (range: $[25, 750]$ μs), and excitation rise time (range: $[3, 380]$ μs)

All experimental samples consisted of 40 μL of stock tracer solution placed in small PCR vials. While stock solutions contained tracer at different concentrations, all data is reported with signal normalized by the mass of iron in the sample for proper comparison.

Our AWR system shown in Fig. 5.3 is described in detail in our recent work (Tay et al., 2016). Briefly, an inner solenoidal transmit coil is separated from an outer solenoidal bias coil with a cylindrical copper shield, all arranged coaxially. The transmit and bias coils provide the fast excitation (pMPI or sinusoidal waveforms) and a slowly varying DC offset field, respectively. The biasing field allows us to query the response of a sample across a large magnetic FOV and thus fully characterize the 1D PSF corresponding to the sample. An inductive receive coil wound in a gradiometric fashion is placed coaxial and internal to the transmit coil. A mechanical tuning mechanism allows for on-demand adjustments to the relative geometry between the transmit and receive coils to maximize cancellation of direct feedthrough. Unlike continuous wave MPI, there is no notch filter associated with a fundamental frequency. Instead the received signal passes through a standard pre-amplifier (SR560, Stanford Research Systems, Inc., California, USA) in a broad band-pass setting. The pre-amplifier begins to attenuate frequencies beyond $1\,\mathrm{MHz}$; we therefore constrained the received bandwidth to less than $500\,\mathrm{kHz}$. MATLAB (Mathworks, Inc., Natick, Massachusetts, USA) scripts were used to drive data acquisition. The output of the SR560 pre-amplifier was sampled at $10\,\mathrm{MHz}$ with a PCI-6115 12-bit ADC (National Instruments, TX, USA).

All 1D excitation trajectories consisted of an (approximate) square wave, trapezoidal, or sinusoidal excitation superposed with a linear bias field to cover a magnetic FOV from -40 to 40 mT over a total single acquisition time of 0.5–0.7 seconds. To improve SNR, up to 25 single acquisitions were taken serially and averaged together, for a maximum of 17.5 seconds total acquisition time per 1D trial. Additionally, void baseline scans with no sample present, were taken for time domain baseline subtraction in reconstruction. In general, three replicate trials were performed in each experiment to report statistical variation in device performance.

5.3.3 Reconstruction

To reconstruct images in our pMPI x-space implementation, we first perform basic digital signal processing (DSP) on the digitized received voltage waveform. This includes application of a digital filter with a $200 - 500\,\mathrm{Hz}$ passband around each harmonic of the fundamental frequency (depending on the bias field slew rate), up to a maximum absolute frequency cutoff of 500 kHz. With a fast slewing square wave excitation, the transmit feedthrough, in a $\mathrm{d}/\mathrm{d}t$ sense, is limited to a few $\mu\mathrm{s}$ after each step jump. Due to the significant relaxation of larger core tracers, this can be separated from the majority of the tracer signal. In such cases, this allows us to retain most of the tracer signal without need for explicit feedthrough removal or even baseline subtraction. Time-domain baseline subtraction is a robust process with our AWR and so is generally performed to remove

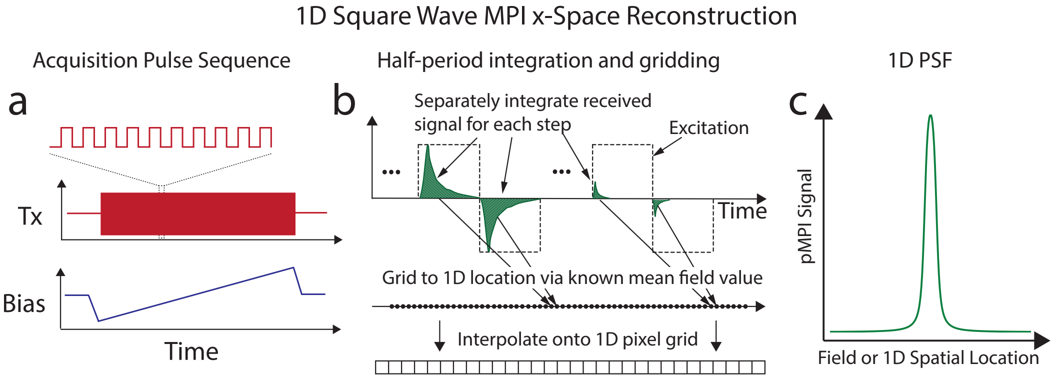

FIGURE 5.6: Square Wave data acquisition and image reconstruction using the AWR. (a) In acquisition, a slowly-varying bias field allows us to sample a large magnetic FOV over time with the faster square wave excitation (Tx). This pulse sequence leads to a dense sampling of points in the applied magnetic field domain. (b) In a simple x-space reconstruction, we can integrate the signal for each square wave half-period and grid this value to the mean field location. In a final step we can interpolate the gridded data onto a desired 1D output pixel grid. (c) The result is a 1D square wave PSF of the AWR sample in aggregate. If we divide the applied field by an assumed linear gradient strength, we arrive at a 1D spatial PSF.

the amount of feedthrough that does overlap along with any remaining sources of interference. However, without time-domain baseline subtraction, feedthrough that remains in the signal translates to a global DC offset in the reconstructed image domain. Using robust *a priori* assumptions of zero signal at the edges of the imaging FOV, we can remove this global DC without requiring a more complex partial FOV DC recovery as required in sinusoidal x-space MPI (Lu et al., 2013).

Fig. 5.6 graphically illustrates the reconstruction process. In the AWR, we sample the 1D applied field space during acquisition using a pulsed excitation waveform superposed with a bias field ramp to cover a large magnetic FOV. After DSP, we fold the signal into half-period chunks, integrate each chunk, and grid this value to a location in the magnetic field domain: $H_n - \Delta H$ where H_n is the value of the total applied field (including the bias field contribution) immediately prior to a step jump and ΔH is the square wave amplitude. The result is a 1D PSF of the MPI signal gridded against applied magnetic field space. ΔH may be positive or negative depending on the slew direction. We also note that the nature of pMPI excitation (abrupt dynamic changes interspersed with quiescent periods) as well as the integration step simplify phase coherence methods required in x-space implementations. Precise identification of the initiation of each half-period excitation is typically easily obtained from the raw data itself.

Because no gradient fields are present in these 1D scans, the system collapses the sample geometry to a point and the resulting PSF is representative for the sample in aggregate. We have previously described theoretically and shown experimentally (Croft et al., 2016; Tay et al., 2016) that the 1D PSF obtained using an x-space relaxometer corresponds directly to the 1D PSF obtained in an imaging scanner after division by the gradient strength of the scanner. We therefore plot some of our experimental 1D PSFs as a function of position by assuming a 7 $\mathrm{T\,m^{-1}}$ gradient, the gradient strength of our highest resolution in house scanner (Goodwill et al., 2012a).

We have implemented our reconstruction algorithms in a variety of languages: MATLAB, Python, and Rust. MATLAB and Python are high level interpreted languages with established numerical, I/O, and visualization libraries. Rust is a lower level compiled language, in this sense similar to C, that we sometimes leverage as a computationally efficient backend.

5.4 Experimental Results

5.4.1 Square Wave Signal, Resolution, and Comparison with Sinusoidal MPI

Fig. 5.7 shows a direct comparison between sinusoidal MPI and pMPI in terms of relaxation-induced and excitation direction-dependent asymmetric blurring. Not only does relaxation reduce resolution in canonical MPI, but it leads to skews in the direction of the excitation direction. This typically necessitates reconstructing separate images for each excitation direction used in acquisition using the x-space method. For constant excitation with a sinusoid in one axis, this reduces to separate 'positive' and 'negative' images that are not easy to combine without harming resolution. Because pMPI removes these effects of relaxation in the image domain, these asymmetries do not appear. We achieve improved resolution, symmetric PSFs, and no need to separately handle or reconstruct received signal based on the excitation direction.

Fig. 5.8 shows experimental square wave MPI data and a comparison with conventional sinusoidal MPI. Fig. 5.8(a) shows experimental square wave pMPI PSFs for the 27.4 nm tracer excited with different excitation amplitudes. As predicted by the theory, the PSF shape, SNR, and resolution are strongly affected by the choice of excitation amplitude, with the characteristic rectangular blur prominent in the case of the largest excitation.

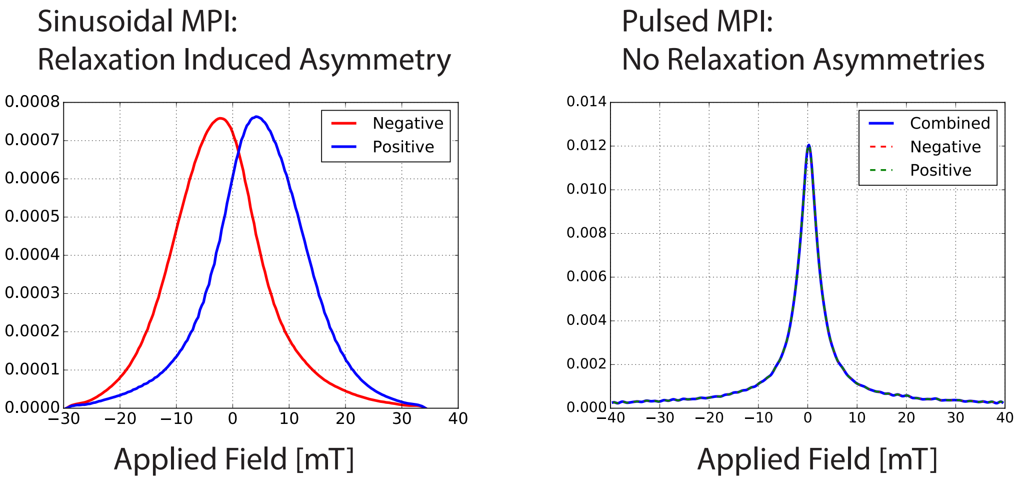

FIGURE 5.7: pMPI removes relaxation-induced asymmetries in the image. In canonical sinusoidal MPI, the presence of significant relaxation reduces resolution due to a blur that originates from time domain relaxation and that is asymmetric and directional with excitation direction. In a typical x-space reconstruction, this will necessitate separate reconstruction for the positive and negative slewing components of the excitation waveform to avoid worst-case naïve averaging of the two. With this approach we get asymmetric skew and image domain translations in the excitation direction along with reduced overall resolution. In pMPI, on the other hand, relaxation effects are removed, yielding native resolution, a symmetric PSF, and no need to distinguish excitation directions in reconstruction.

Fig. 5.8(b) experimentally quantifies the relationship between FWHM resolution and excitation amplitude, showing excellent agreement with the theory, and Fig. 5.8(c) quantifies the experimental peak signal as a function of excitation amplitude, also agreeing well with the theory.

Fig. 5.8(d) shows 1D PSFs obtained with the AWR for each of the tracers but using conventional sinusoidal MPI excitation. Here we used the standard $f_0 = 25\,\mathrm{kHz}$ and $20\,\mathrm{mT}$ excitation amplitude used in our imaging systems. The effects of magnetic relaxation are apparent in the data: blur and translation in the scanning direction that increases in severity as the tracers get large. In the the most severe cases (the 32 nm sample), distortion artifacts are present caused by large phase lags disrupting the reconstruction process. Fig. 5.8(e) shows pMPI PSFs obtained with the same samples using $1\ \mathrm{mT}$ square wave excitation. The steady-state encoding of pMPI leads to PSFs without blurring and translation from magnetic relaxation effects, and larger tracers produce narrower PSFs. Finally, Fig. 5.8(f) plots the FWHM resolution for sinusoidal MPI, square wave pMPI, and the theoretical Langevin model. We see the characteristic inflection point around 25 nm for sinusoidal excitation, leading to worse resolution for the larger tracers.

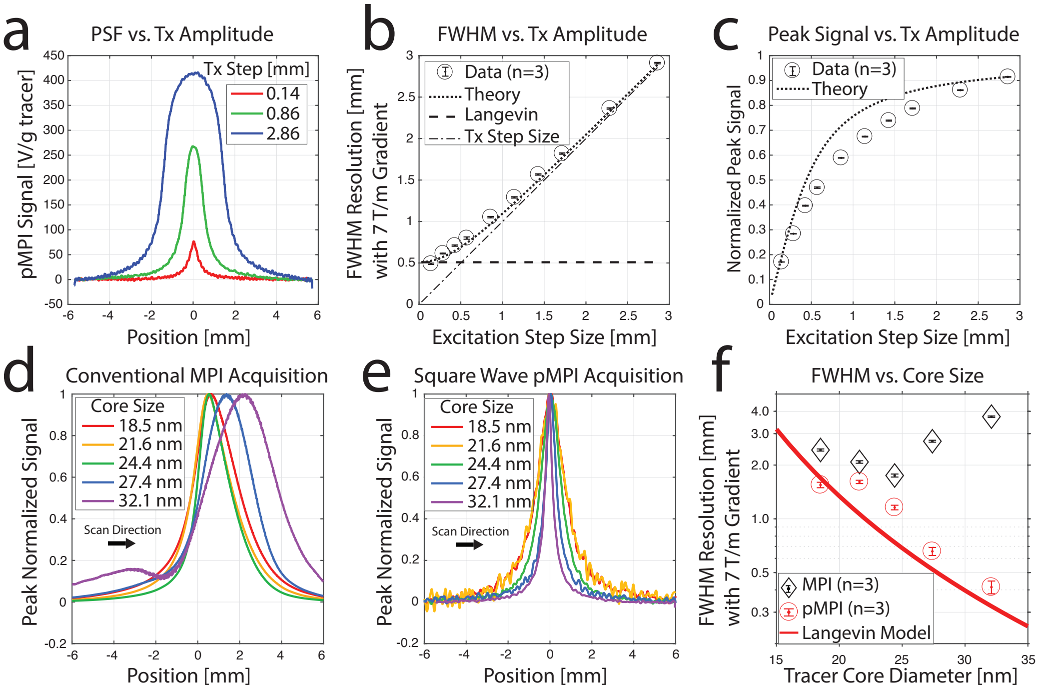

FIGURE 5.8: Experimental square wave MPI and comparison with conventional methods. (Top) Square wave MPI experimental results as a function of excitation amplitude and comparison with theory for 27 nm particles. (a) 1D PSFs using different exciation amplitudes and normalized by tracer mass. Experimental data points overlaid on theoretical predictions for (b) FWHM resolution and (c) peak signal as functions of excitation amplitude. (Bottom) Experimental resolution comparison between conventional sinusoidal MPI (25 kHz and 20 mT) and square wave pulsed MPI (1 mT and 5 kHz fundamental square wave frequency). (d) Sinusoidal data, (e) square wave data, (f) comparison of experimental sinusoidal MPI FWHM resolution, square wave pMPI FWHM resolution, and theoretical Langevin prediction as a function of tracer core size.

In contrast, the square wave pMPI data points continue to follow the theoretical Langevin curve with a slight offset. The effect of the excitation amplitude means we only asymptotically approach the theoretical Langevin resolution, leading to the observed minor offset.

5.4.2 Effects of Hold Time and Rise time in Square Wave pMPI

Fig. 5.9 shows the effects of hold time and rise time in experimental square and trapezoidal wave pMPI. Fig. 5.9(a) shows square wave pMPI PSFs constructed from acquisitions with square wave half-periods of 25 μs, 100 μs,

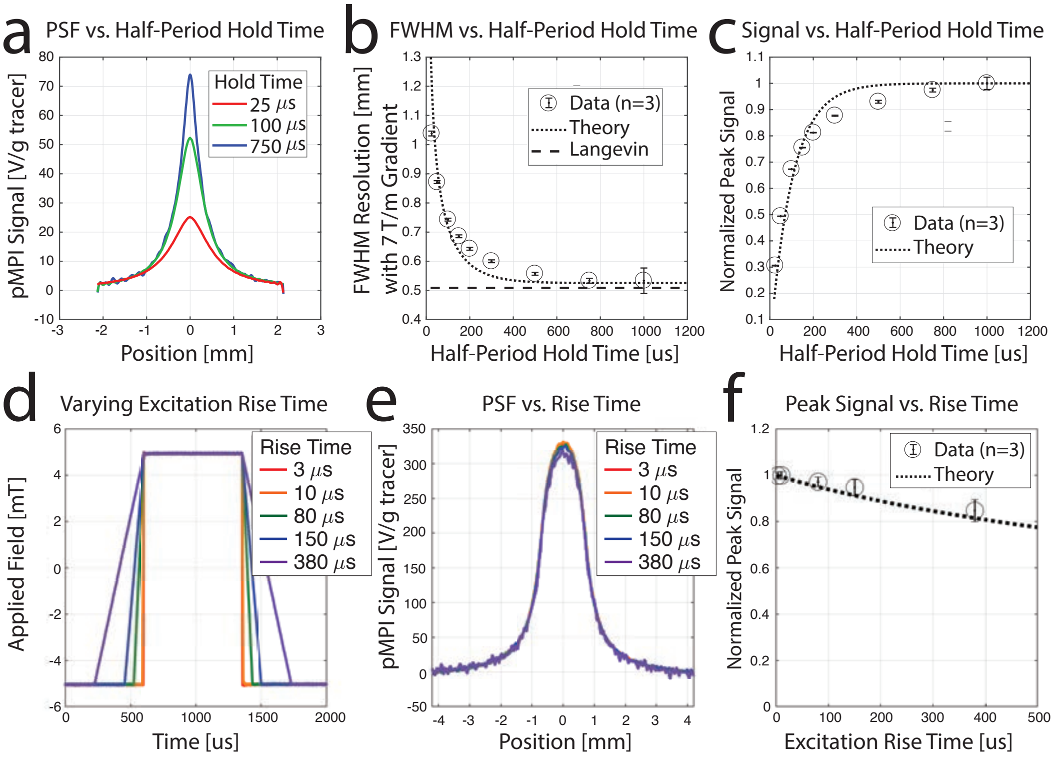

FIGURE 5.9: Effects of hold time and rise time in experimental square wave MPI. (Top) Square wave MPI experimental results as a function of hold time – the length of time at a constant applied field to establish steady-state – and comparison with theory for 27 nm particles. (a) 1D PSFs normalized by tracer mass. Experimental data points overlaid on theoretical predictions for (b) FWHM resolution and (c) peak signal as a function of excitation amplitude. The theoretical curves were constructed using a simple Debye model of MPI relaxation. (Bottom) Experimental analysis of the effect of rise time using trapezoidal excitation waveforms with different rise times. (d) Experimental excitation waveforms, (e) pMPI PSFs, (f) Peak SNR as a function of the excitation rise time.

and 750 μs. A phenomenon of ‘peak depression’ is apparent for shorter half-periods which reduces the FWHM resolution of the PSF as shown in Fig. 5.9(b). The magnetic relaxation is a strong function of the applied field strength. At low field strengths the relaxation time is longest and unable to achieve steady-state conditions at, *e.g.*, 25 or 100 μs. This reduces the signal at low fields, although steady-state may be achieved at the higher field strengths where the higher magnetic torque leads to faster relaxation processes. Fig. 5.9(c) quantifies the peak signal loss as a function of hold time due to this peak depression effect. The theoretical curves of Fig. 5.9(b,c), were constructed using a model of relaxation as a simple Debye process.

Fig. 5.9(d) shows experimental pMPI excitation waveforms with varying linear rise times. The shortest duration rise time, $3\,\mu\mathrm{s}$ is effectively a square wave. Trapezoidal waveforms with rise times of 10, 80, 150, and 380 $\mu\mathrm{s}$ are also shown. Fig. 5.9(e) shows pMPI PSFs reconstructed from acquisitions using excitation waveforms with the different rise times while Fig. 5.9(f) shows how the peak signal varies with the rise time. There is essentially no effect on PSF shape and FWHM resolution as a function of rise time. PSF shape depends only on steady-state induction with a long enough hold time. As long as this parameter is chosen to achieve steady-state as indicated in Fig. 5.9(b), the desired FWHM resolution will be achieved. Longer rise times do lead to a minor decrease in peak image domain signal due to reduced SNR efficiency and our use of a fixed total scan time. In this case, a longer rise time means we achieve fewer excitations per unit time. The theoretical curve of Fig. 5.9(f) takes this into account.

5.5 Discussion

5.5.1 Breaking the Relaxation Wall with a New Encoding Strategy

The theoretical results from derivation (5.11) as illustrated in Fig. 5.5 and our experimental data in Fig. 5.8 show excellent agreement. We believe this rigorously establishes a new encoding strategy in MPI that has many new compelling features when compared with standard sinusoidal MPI. Perhaps most importantly, for the first time, we have been able to follow the Langevin steady-state predictions of resolution, even for large particles with core sizes over 30 nm. In this respect, we have broken the ‘relaxation wall’ that has plagued canonical MPI. Accordingly, we experimentally report the best ‘native’ (*i.e.* without explicit or implicit deconvolution) resolution we have ever achieved in x-space MPI.

The ability of pMPI encoding and reconstruction to remove the deleterious effects of relaxation not only improves raw resolution metrics, but also improves robustness of MPI in at least two more ways. First, pMPI removes the asymmetry and skew imparted by relaxation on the PSF as illustrated in Fig. 5.2 and Fig. 5.7. As further illustrated in Fig. 5.7, the asymmetric skew in canonical MPI makes combining information associated with negative and positive slewing components of the excitation waveform difficult. In practice, we often just use an image reconstructed from one or the other as our final output, throwing away 50% of our data. While methods (*e.g.*, forward model approaches) may be used to combine the images, this has its own costs and is sensitive to model error. As illustrated in Fig. 5.7, we can effortlessly combine these data in pMPI due to the removal of the offending relaxation asymmetry.

Second, the presence of significant relaxation in sinusoidal MPI compromises the linearity of the MPI signal, possibly in space- and condition-variant ways. For example, the presence of disparate viscous environments in an imaging FOV or use of different excitation amplitudes will augment the tracer's peak signal depending on the significance of relaxation effects. We have isolated this effect in the pMPI analysis of Fig. 5.9, where insufficiently long readout periods leads to a 'peak depression' effect. In Chapter 3, we used these effects to enable color MPI multiplexing precisely by using small changes in peak signal intensity. At the same time, it also means that nonlinearities in our signal intensity occur. When due to the excitation amplitude, we can control and calibrate this; however, when due to spatially varying environmental conditions, we cannot. In pMPI, the complete removal of relaxation effects from the final image obviates this confound.

5.5.2 Scope of Theoretical Results

We believe the general applicability of the pMPI derivations in 5.2 is compelling. The nature of magnetic relaxation behavior in MPI, especially in the presence of multiple simultaneously active physical mechanisms, can be very complex and is not well understood. However, the results of this chapter pertain across a wide swathe of such physics, without any need for explicitly accounting for them. We simply require a finite magnetic relaxation response and steady-state superparamagnetism. The only obvious candidates that can violate this broad classification are aggregated ferromagnetic tracers.

To realize the pMPI theory, we must use excitation waveforms with quiescent readout periods (*e.g.*, square wave half-periods) longer than the longest magnetic response time of the tracer. However, this is easily and rigorously measured experimentally, without requiring a first principles knowledge of the mechanics. As we will explore in more detail in the next chapter, the signal evolution over the course of each quiescent readout period is exactly a magnetic impulse response. Furthermore, as Fig. 5.9(a–c) demonstrate, we can calibrate the required readout period length in a fast 1D parametric search using the AWR or a gradient-less scanner, looking for convergence of resolution and peak signal as a function of readout time.

Understanding the tracer physics from first principles will certainly benefit pMPI and MPI at large, but we can proceed with pMPI confidently even without such a specific knowledge of a tracer's governing physics.

5.5.3 pMPI Advantages

The pMPI approach has allowed us to achieve Langevin resolution for large tracers in the limit of small excitation amplitudes. In standard x-space MPI, use of very small excitation amplitudes (*e.g.*, $1\ \mathrm{mT}$) is untenable, even if desired. This is because by shrinking down the excitation amplitude (FFR excursion per oscillation) generally, and well below the Langevin FWHM in particular, the received signal energy is pushed more and more into the first harmonic. First harmonic information is lost in all canonical sinusoidal systems due to necessary excitation feedthrough mitigation strategies. So use of smaller amplitudes in sinusoidal MPI leads to increasingly ill-posed reconstruction. In pMPI, however, our use of fast transitions can realize a temporal decoupling between excitation feedthrough and the tracer response. In the limit of an ideal square wave, ideal system response, and finite relaxing particles, we approach complete separability. In our real systems, there is some overlap between the two phenomena, but we have observed separation of the majority of the signal energy even with our initial systems.

Another advantage of the pMPI approach described herein is a new ability to continuously tradeoff resolution and SNR in a seamless way by changing the excitation amplitude. This is due to the unique relationship these two parameters have with respect to the amplitude in pMPI. Specifically, our integrated gridding reconstruction locally projects the received single associated with each step excitation across the distance spanned by the excitation. This leads to the appearance of the rectangular function blur in our new imaging equation. A similar tradeoff is possible in canonical MPI when using electromagnets by varying the linear gradient strength. However, having another tradeoff lever is useful generally, and specifically the amplitude lever pertains when using permanent magnet systems and may be more seamless to implement in practice compared to dynamically configurable electromagnets. Furthermore, the physically large, high inductance, and high power gradient electromagnets cannot be dynamically modified as fast as the excitation fields.

5.5.4 pMPI Challenges

There are challenges associated with pMPI as we've outlined so far. Most importantly, and as discussed in 5.2.4, in taking advantage of pMPI to allow us to realize very high resolution using larger tracers, we must accommodate the longer tracer relaxation with long square wave hold times that allow us to approach steady-state conditions. Because this follows from the fundamental physics of the tracer that is the basis of our signal in MPI, it will ultimately impose limits on achievable SNR and SNR efficiency. This is true both in pMPI as presented in this chapter as well as when using other excitation strategies such as sinusoidal MPI with reduced excitation slew rates.

One challenge specific to the square wave pMPI approach we described in this chapter is the tradeoff between SNR and resolution imposed by the rectangular blur that appears in our imaging equation. To achieve very high resolutions in square wave pMPI, *e.g.*, at or near the Langevin limit, we must use very small excitation amplitudes. For example, in our experimental data, we reduced the excitation amplitudes (1 mT) greatly compared to that typically used in sinusoidal MPI (20 mT). This leads to reduced signal per unit time than if we could use much higher amplitudes. We consider this effect, however, an *implementation detail* of our simple integrated gridding reconstruction method. Namely, we are locally projecting (spatial blurring) all information associated with a square wave step excursion and do not attempt to resolve information inside this range.

There are also practical, *cf.* theoretical, challenges associated with implementing pMPI. Namely, the transmit and receive systems must be redesigned from resonant chains to broadband systems capable of potentially very rapid step-like excitations. The desire for highly sensitive and large volume coils can be in conflict with, for example, the desire for low inductance coils when implementing pre-emphasis and linear amplifier based systems as presented here. Alternatively, resonant switcher circuits may be employed. The strain of step response-like feedthrough on the receive system may motivate the use of cross diodes, variable gain pre-amplifiers, and other new receive chain features. In general, pMPI leads to a different set of design parameters to optimize for when building excitation and receiver systems. These new challenges have been explored deeply in other fields and we do not see these challenges as more difficult than the canonical MPI systems, only different.

5.5.5 Other Sequence and Reconstruction Approaches

As described, our definition of pMPI is very broad, including a variety of excitation waveforms constructed by composing the basic primitives of fast transitions and constant field or readout periods. Here we only explored the simplest of such waveforms: periodic square and trapezoidal waveforms. With regards to SNR efficiency challenges, for example, it is quite possible that more sophisticated pulsed waveforms could further mitigate this challenge while preserving the benefits of pMPI. More generally, we believe that pMPI waveforms have a rich ability to manipulate and exploit the encoding of the MPI signal during acquisition, perhaps one day providing a system as rich as pulse sequence design in MRI which, among other things, can manipulate contrast encoding due to the presence of T1 and T2 relaxation parameters. Indeed, we believe there are many parallels between pMPI and MRI sequencing in terms of encoding and the mathematics even though the specific physics differ greatly in nature.

The amplitude-dependent blur of our current reconstruction method arises because we do not try to spatially disambiguate any of the signal associated with a square wave half-period. By integrating the total signal and gridding this value to a single pixel, we are locally projecting all spatial information over the size of the excitation amplitude. We expect to be able to discriminate spatial information inside of a half-period in the future, which will release us from the amplitude based SNR limitations of this chapter. For example, we may be able to use much larger amplitudes and perform projection reconstruction, locally or globally, to recover this spatial information. More exotic approaches such as excitation along the line of an FFL system may also be possible.

5.5.6 Safety

We have left the assessment of safety (*e.g.*, considerations at the human scale) outside the scope of this chapter. At a basic and intuitive level, however, we do not expect this to be an issue. Or at least we expect the flexibility in our excitation parameters will allow us to conform to the imposed limitations. The low amplitudes and long static readout periods described in this chapter may be viewed as a relatively low energy and low duty cycle excitation waveform when considering RMS safety in terms of specific absorption rate (SAR) energy deposition. Furthermore, when using very rapid step excitations, as in the case of a square wave, the high frequency nature of the step jump should generally remain outside of the frequency range known to elicit a magnetostimulation response. Nonetheless, in our future work, we will explore this in more detail both theoretically and experimentally.

5.6 Conclusions

We have described a new signal encoding and acquisition strategy, pulsed MPI. We derived fundamental signal and imaging equations in 1D and confirmed the predictions of the derivations with experimental 1D data. pMPI allows us to break the 'relaxation wall' encountered in canonical MPI by removing the deleterious effects of magnetic relaxation on the MPI image, allowing us to achieve continued resolution improvement with increase tracer core size, even for tracers larger than 30 nm. With this development we have showed native 1D MPI FWHM resolution below $500\mu\mathrm{m}$ for the first time. Additionally, our derivations show that these achievements of pMPI encoding apply generally to MPI tracers *regardless* of the specific nature of their relaxation.

pMPI provides other new advantages, such as an ability to significantly decouple excitation feedthrough from the tracer response, and an ability to seamlessly use received signal associated with any arbitrary excitation direction when producing a single output image. We have also contemplated new challenges associated with square wave pMPI such as reduced SNR efficiency in some cases and requirements for new transmit and receive system design. Here we have assessed only the simplest square and trapezoidal pMPI pulse sequences in our initial description. We expect that future developments will demonstrate myriad pMPI pulse sequences that further exploit aspects of MPI signal encoding, perhaps in ways similar to T1 and T2 relaxation contrast in MRI.

Future work will demonstrate pMPI in multidimensional imaging, explore the ability to quantify and exploit the relaxation information encoded directly in the readout periods of pMPI acquisition, rigorously analyze safety considerations, and explore novel and more sophisticated pMPI pulse sequences.

5.7 Acknowledgements

Developing our new pulsed MPI methods has been a group effort, and I want to thank Zhi Wei Tay and Steve Conolly for their assistance and helpful discussions throughout this process. I developed the theoretical analysis and mathematical results and framework presented in this chapter. Additionally, all of the discussion and analysis of results as well as all of the writing is my work. Zhi Wei Tay and I designed the experimental approach and much of the pMPI data acquisition/reconstruction/analysis code used in this chapter together, and Zhi Wei acquired data presented in Figs. 5.1(c), 5.8, and 5.9. Finally, Steve Conolly provided helpful mathematical analysis and discussions, especially when we initially launched this project.

Chapter 6

Pulsed MPI Relaxation Imaging in 1D

In this chapter we explore ways of measuring and exploiting relaxation in 1D pMPI. This is in contrast to Chapter 5 where we concerned ourselves only with removing deleterious effects of relaxation. We first describe and develop theoretical models for field-dependent Néel and Brownian relaxation. Using these models and our square wave signal and imaging equations from Chapter 5, we then propose methods to construct relaxation images and quantify magnetic relaxation. We also propose methods to combine the relaxation spatial encoding in pMPI with standard square wave imaging for improved resolution using what we call 'relaxation weighting' techniques. We describe these methods using signal and imaging equation derivations as appropriate. We then demonstrate and evaluate these proposed techniques with proof-of-concept experiments. We show that experimental data aligns well with our theoretical framework, show compelling capabilities to quantify the effect microenvironmental variables such as viscosity, pH, and binding have on magnetic relaxation, and show the ability to improve resolution to less than 300 μm and *beyond the limits of Langevin theory* using relaxation weighting methods. We conclude with a lengthy discussion of the results that contemplates various issues and unexpected behavior that we observed and propose various future directions to follow this initial work. We believe these results introduce and demonstrate powerful new methods of exploiting magnetic relaxation that pMPI can provide over canonical continuous wave methods. To us, these approaches are especially exciting because they are due to new usage of fundamental physics – the best kind of improvement. However, it is very important to follow up on these early results in more detail and, especially, in a multi-dimensional imaging format.

6.1 Introduction

In the previous chapter, we described a new approach to signal encoding and acquisition in MPI we call pulsed MPI. In particular, we focused on rigorously developing a new imaging equation that encodes only steady-state tracer magnetization in the final image, allowing us to overcome previous relaxation-induced resolution limitations in canonical continuous wave MPI. In particular, the square wave paradigm we described provides complete removal of relaxation effects from the final output image. Insofar as these effects are deleterious, such as secondary blurs and signal nonlinearities, this is desirable. On the other hand, Section I illustrated how relaxation information can be used to enhance desirable image contrast and provide colorized multiplexing. In this chapter, we will explore in detail the second half of the story told in the last chapter – how pMPI offers unprecedented new ways of encoding and exploiting magnetic relaxation.

We first show that we can produce new types of MPI images, relaxation images, that quantify magnetic relaxation on a per pixel basis using the same square wave acquisition described in chapter 5. In this manner, we can fully orthogonalize our inspection of steady-state Langevin physics, as in chapter 5, and our inspection of relaxation-based information, both derived from the same data set. We will contrast this approach to our similar approach explored for continuous wave MPI in chapter 2.

We will then describe the first non-square wave pMPI pulse sequence: Using a 'steady-state recovery' sequence that allows us to sensitively quantify parameters such as environmental viscosity and tracer binding state that affect relaxation dynamics. Our discussion of this new sequence also serves to introduce the myriad possibilities of more sophisticated pulse sequencing in pMPI.

Finally, we will show that pMPI allows us to combine relaxation phenomena and Langevin phenomena (*i.e.*, not keep them orthogonal) in ways that can improve native resolution well below the previously fundamental Langevin limit. The fundamental physics enabling this is a second spatio-temporal encoding provided by relaxation (and distinct from the steady-state spatio-temporal encoding described by the Langevin equation). By moving beyond the strictures of single tone AC excitation, pMPI is a much more capable platform for exploiting the rich physics provided by MPI tracers. The greatly improved resolution and new SNR-resolution tradeoffs we describe are accomplished through encoding and exploiting raw physics – an approach, as we discussed in Chapter 1, that we deem preferable to model-based postprocessing techniques such as deconvolution.

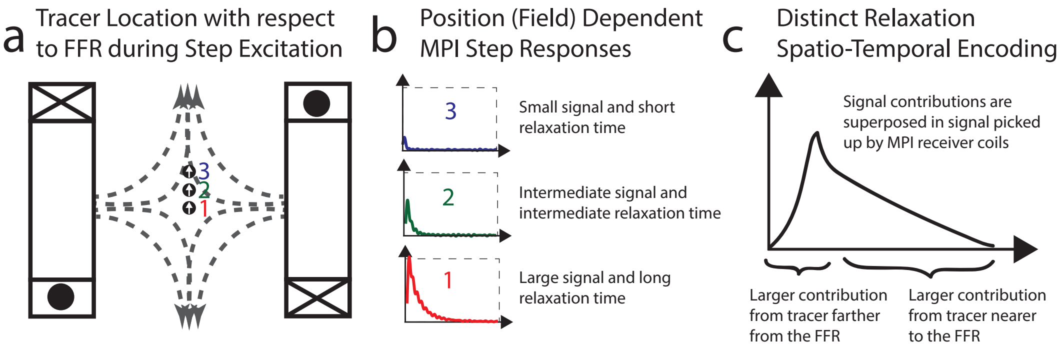

FIGURE 6.1: Field and position dependent MPI relaxation in pulsed MPI. (a) Consider tracer located at three different positions relative to the MPI FFR isocenter during a pMPI step excitation. (b) Both the total signal energy (from steady-state Langevin M-H curve) and the relaxation dynamics depend on location and thus applied field via the gradient fields of MPI. (c) In an imaging context, the inductive Rx in MPI superposes these contributions. Early stage signal will be weighted relatively more by tracer farther from the FFR while later stage signal will be composed almost solely of signal from tracer nearer the FFR. This is a second spatio-temporal encoding, independent of the Langevin spatio-temporal encoding, that we can exploit.

6.2 Field Dependent Relaxation Models

In Chapter 5 we showed that the implications of establishing steady-state during periods of temporally constant applied field in pMPI were applicable generally, regardless of the nature of the magnetic relaxation physics. In this chapter, we consider quantifying and exploiting magnetic relaxation. While it is still possible to measure and report relaxation information without specific knowledge of relaxation physics, optimizing pulse sequence design to most effectively achieve these goals, choosing proper quantification metrics, and properly interpreting the results motivate a more explicit understanding or modeling of relaxation phenomena.

Perhaps the most important feature of magnetic relaxation physics that underly the methods in this chapter is a strong dependence on the applied field strength. In Chapter 2, we showed a first attempt at ‘relaxation maps’ that capture this field dependence by gridding adjacent pFOV cross correlation delays against the mean applied field. In this chapter, we will explore field dependent relaxation in much more detail, including both theoretical and empirical analysis. Much of the literature of magnetic nanoparticle physics is concerned with activity in the absence of an applied field, but MPI involves high total applied magnetic field magnitudes, large magnetic field gradients, and rapid temporal shifts in these fields – a complex applied field space. One of the reasons it has been difficult to fully understand MPI relaxation phenomena in general is this complexity of the applied field.

Fig. 6.1 describes how field-dependent relaxation manifests in pMPI. In the previous chapter, we showed mathematically that the raw time domain MPI signal after an ideal step excitation is exactly the magnetic relaxation impulse response. As we will show, we can directly measure or fit a relaxation parameter of merit for each half-period of our square wave excitation and grid this value to the mean location of the total field. This is a much better approach than the relaxation maps constructed from sinusoidal excitation because we can directly evaluate the relaxation impulse response at a specific location of the FFR. In contrast, the cross-correlation technique of Chapter 2 is a relatively indirect measure that is averaged over an entire excitation oscillation.

Fig. 6.1(b) shows experimental data where the impulse response and its field-dependent nature are apparent. Tracer farther away from the FFR during a square half-period step excitation will contribute a relatively small (*e.g.*, integrated) total signal over the duration of the half-period due to near Langevin saturation but will be characterized by the fastest relaxation times due to the strong field-dependence of magnetic relaxation. Signal closer to the FFR will exhibit the opposite behavior – large total signal magnitude evolving more slowly over time. Furthermore, as Fig. 6.1(c) depicts, this field-dependent relaxation (and position-dependent in the presence of an MPI gradient field) leads to a spatio-temporal encoding independent of the core Langevin spatio-temporal encoding. We will take advantage of this phenomenon in this chapter.

First we consider how we might model the field dependence of the MPI signal. We will build models of the two relaxation mechanisms that typically dominate MPI tracers: Néel and Brownian. Both of these mechanisms are known to have a strong dependence on the applied field (Arami et al., 2013; Deissler, Wu, and Martens, 2014).

6.2.1 Néel Relaxation

The literature describes various Néel models, often derived from an initial Fokker-Planck starting point. We begin with the zero field expression for the Néel time constant (Deissler, Wu, and Martens, 2014):

$$\tau_{\mathrm{N}}(0) = \frac{\sqrt{\pi}\beta(1+\alpha^2)M_{\mathrm{sat}}}{4\gamma\alpha(\beta K)^{3/2}} \cdot \exp(\beta K) \tag{6.1}$$

where V_c is the magnetic core volume of the tracer in $[\mathrm{m}^3]$, k is Boltzmann's constant in $[\mathrm{J\,K^{-1}}]$, T is the ambient temperature in $[\mathrm{K}]$, and $\beta = V_\mathrm{c}/kT$ in $[\mathrm{m^3\,J^{-1}}]$; α is a a dimensionless damping constant, M_sat is the tracer saturation magnetization in $[\mathrm{A\,m^{-1}}]$ or equivalently $[\mathrm{J\,m^{-3}\,T^{-1}}]$; γ is the electron gyromagnetic ratio in $[\mathrm{rad\,s^{-1}\,T^{-1}}]$; and K is the tracer anisotropy constant in $[\mathrm{J\,m^{-3}}]$.

We now consider Néel relaxation in the presence of a non-zero applied field. Deissler et al. (Deissler, Wu, and Martens, 2014) report that a high barrier approximation for the Néel relaxation time associated with a magnetic step excitation initially provided by Brown (Brown Jr, 1963) agrees well with numerical simulations of the base Fokker-Planck derivation:

$$\begin{aligned}&\tau_\mathrm{N}(H_\mathrm{k}) = \\ &\quad \tau_\mathrm{N0} \cdot \frac{\sqrt{\pi}(\alpha_\mathrm{K}/2)^{-3/2}}{1-H_\mathrm{k}^2} \cdot \\ &\quad \big((1+h)\mathrm{exp}((-\alpha_\mathrm{K}/2)(1+H_\mathrm{k})^2) + (1-H_\mathrm{k})\mathrm{exp}((-\alpha_\mathrm{K}/2)(1-H_\mathrm{k})^2\big)^{-1}\end{aligned} \tag{6.2}$$

where H_k is a normalized and dimensionless applied field value: $H_\mathrm{k} = M_\mathrm{sat}B/2K$, B is the applied field magnitude in $[\mathrm{T}]$, $\alpha_\mathrm{K} = 2K\beta$ is a dimensionless constant, and we have the relation

$$\tau_\mathrm{N}(0) = (\sqrt{\pi}/2)(\alpha_\mathrm{K}/2)^{-3/2}\mathrm{e}^{\alpha_\mathrm{K}/2}\tau_\mathrm{N0}.$$

An important ramification of this result is that we should expect Néel time constants to be vary strongly with the locally applied field strength.

6.2.2 Brownian Relaxation

The perhaps more intuitively accessible Brownian physics can be understood as the attempt of the magnetic nanoparticle to physically rotate in space to align with the applied field, driven by the magnetic torque from the applied field and opposed by viscous resistance and stochastic torques imparted by the Brownian motion of the surrounding solvent. Accordingly, we can model Brownian physics in one dimension with a torque balance equation (Raible and Engel, 2004):

$$I\ddot{\theta}(t) = -M_\mathrm{sat}V_\mathrm{c}B(t)\sin(\theta(t)) - 6\eta V_\mathrm{h}\dot{\theta}(t) + \sqrt{2D}w(t) \tag{6.3}$$

where I in $[\mathrm{kg\,m^2}]$ is the magnetic nanoparticle's moment of inertia, $\theta(t)$ is the angle between the moment of the particle and the applied field, η is the viscosity in $[\mathrm{Pa\,s}]$, V_{h} in $[\mathrm{m^3}]$ is the hydrodynamic volume of the nanoparticle, and the expression $\sqrt{2D}w(t)$ with $\langle w(t)\rangle = 0$ is a stochastic term that represents the torque imparted on the particle by the Brownian motion of the solvent. While rigorously a deterministic function of the trajectories of all interacting solvent molecules, our practical inability to obtain this detailed information motivates use of a stochastic term that captures the aggregate behavior.

The nonlinear and inhomogeneous equation (6.3) can be solved numerically for arbitrary magnetic field trajectory $B(t)$, but here we focus on square wave-like excitation and consider the evolution of the tracer state over a single square wave half-period such that $B(T) = B$ is constant. We further simplify this equation by assuming negligible inertia ($I \approx 0$), which is well motivated for the very small magnetic nanoparticles used in MPI. Furthermore, we make the small angle approximation such that $\sin(\theta(t)) \rightarrow \theta(t)$ and neglect the effects of stochastic Brownian motion. The small angle approximation is, in general, not well motivated to describe the dynamics generally. But our goal here it to obtain an approximate understanding of the expected behavior, and these last two steps provide a pathway to simplified analytical solution: We are left with a first order homogeneous differential equation that is readily solvable, yielding an exponential solution:

$$\tau_{\mathrm{Bo}}(B) = \left.\frac{6\eta V_{\mathrm{h}}}{M_{\mathrm{sat}} B V_{\mathrm{c}}}\right|_{B>0} \tag{6.4}$$

where the subscript 'o' indicates this is an overdamped solution. In this context, our model predicts a Brownian time constant that is inversely proportional to the applied field strength, proportional to the viscosity, and proportional to the ratio of the hydrodynamic and magnetic core volumes. These latter two predictions indicate that MPI tracers with significant Brownian character can report physiologic and microenvironmental state, and that we can easily modify this behavior by modifying the tracer shell characteristics. Equation (6.4) approaches infinity in the limit of small applied field strengths. The zero-field Brownian relaxation time is widely reported as:

$$\tau_{\mathrm{B}}(0) = \frac{3\eta V_{\mathrm{h}}}{kT} \tag{6.5}$$

which is a statement that, under no applied field conditions, the stochastic

Brownian torques will eventually drive a population of tracers that were previously polarized by a non-zero applied field to zero net magnetization. To model Brownian relaxation maps, we combine this zero-field expression with our derived field-dependent expression as follows. This effectively reintroduces the effects of the stochastic Brownian term, which we ignored earlier, to govern very low field conditions:

$$\hat{\tau}_{\mathrm{Bo}}(B) = \frac{3\eta V_{\mathrm{h}}}{(1/2)M_{\mathrm{sat}}BV_{\mathrm{c}} + kT} \tag{6.6}$$

where this equation reduces to (6.5) in the limit of small applied fields and reduces to (6.4) in the limit of large applied fields.

Finally, we also consider a simplified underdamped case for Brownian relaxation, where for very large particles, inertia is non-negligible. In this case, ignoring the stochastic term is also well motivated. If we again make the small angle approximation, we can solve the manifest second order homogeneous equation that arises with relaxation time envelope (here underdamping leads to ringing) characterized by the time constant:

$$\tau_{\mathrm{Bu}}(B) = \frac{I}{3\eta V_{\mathrm{h}}} \tag{6.7}$$

where the subscript 'u' indicates this is an underdamped solution. Interestingly, in this regime, the model predicts an inversely proportional relationship between relaxation time and viscosity, the opposite of the overdamped case. The model also predicts no dependence on the step excitation strength or core volume, only on the hydrodynamic volume. While this expression should not apply for any of the currently used MPI tracers, we will run macroscopic experiments to confirm the relationship.

6.2.3 Visualizing Model Predictions

It is common in the literature to model the observed relaxation as the parallel combination of the Néel and Brownian relaxation times (Deissler, Wu, and Martens, 2014):

$$\tau(B) = \frac{\tau_{\mathrm{N}}(B) \cdot \tau_{\mathrm{B}}(B)}{\tau_{\mathrm{N}}(B) + \tau_{\mathrm{B}}(B)} \tag{6.8}$$

The accuracy and applicability of this parallel model, however, have been questioned because, among other things, it assumes that the Néel and

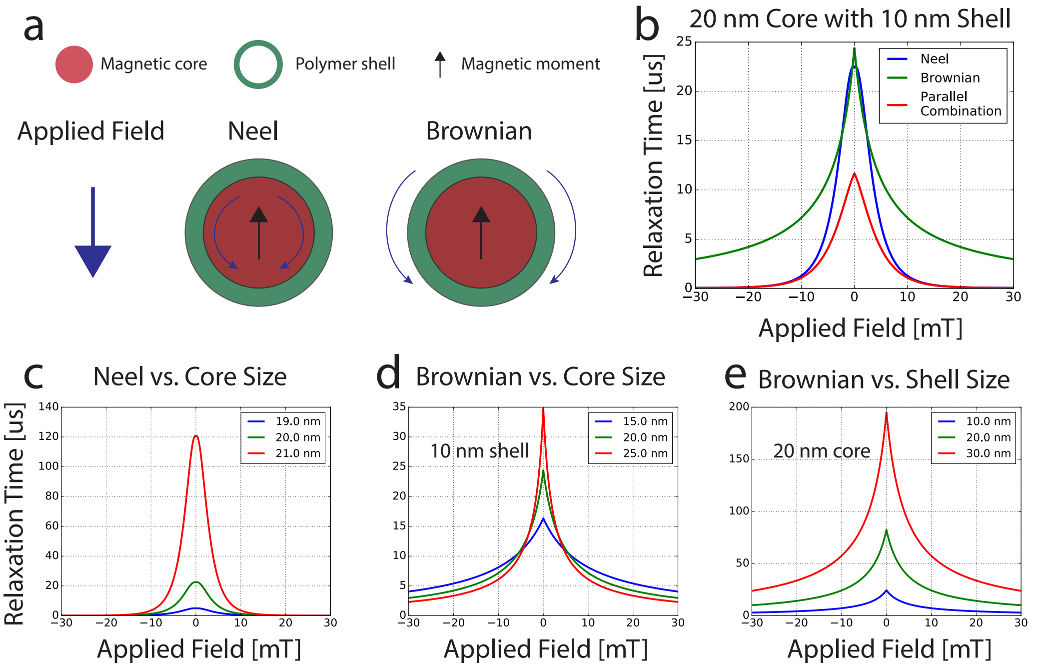

FIGURE 6.2: Magnetic relaxation field dependencies predicted by theoretical models. (a) Depiction of Néel relaxation with internal realignment of the moment cf. Brownian relaxation with physical rotation of the entire particle. (b) Model predictions using the following parameters: 20 nm magnetic core, 10 nm polymer shell, ambient temperature of 300 K, magnetic damping coefficient $\alpha = 0.1$, electronic gyromagnetic ratio, $\gamma 1.75e11$ $[\mathrm{rad\,s^{-1}\,T^{-1}}]$, magnetic anisotropy constant $K = 11000$ $[\mathrm{J\,m^{-3}}]$, magnetic saturation $M_{\mathrm{sat}} = 474000$ $[\mathrm{J\,m^{-3}\,T^{-1}}]$, and viscosity $\eta = 1.0049$ $[\mathrm{mPa\,s}]$. (c) Néel relaxation map dependence on core size. (d) Brownian relaxation map dependence on core size with fixed shell thickness. (e) Brownian relaxation map dependence on shell thickness with fixed core size.

Brownian mechanisms are fully decoupled, which is not necessarily the case (Deissler, Wu, and Martens, 2014).

Fig. 6.2 visualizes these field dependent derivations. Fig. 6.2(a) depicts the different nature of Néel vs. Brownian relaxation mechanisms (internal moment realignment vs. physical particle rotation, respectively). Fig. 6.2(b) shows the model predictions for the following base parameter values: 20 nm magnetic core, 10 nm polymer shell, ambient temperature of 300 K, magnetic damping coefficient $\alpha = 0.1$, electronic gyromagnetic ratio, $\gamma = 1.75 \times 10^{11}$ $[\mathrm{rad\,s^{-1}\,T^{-1}}]$, magnetic anisotropy constant $K = 11\,000$ $[\mathrm{J\,m^{-3}}]$, magnetic saturation $M_{\mathrm{sat}} = 474\,000$ $[\mathrm{J\,m^{-3}\,T^{-1}}]$, and viscosity $\eta = 1.0049$ $[\mathrm{mPa\,s}]$. The Néel, Brownian, and parallel combination of the two are all plotted.

Fig. 6.2(c) demonstrates the significant predicted sensitivity of the Néel relaxation maps on magnetic core size, Fig. 6.2(d) demonstrates the relatively minor predicted sensitivity of the Brownian relaxation maps on core size, and Fig. 6.2(e) demonstrates the significant sensitivity of the Brownian relaxation maps on the tracer shell thickness predicted by our model.

In addition to being derived with explicit simplifying assumptions we listed previously, these models include a number of parameters that are difficult to measure precisely for tracers of interest and the situation is further complicated by the significant parameter polydispersity and non-ideality that exists in real tracer samples. For these reasons, we are not necessarily interested in the exact absolute values of relaxation times predicted by these models. However, we believe the gross shape, patterns, and relative parameter sensitivities are instructive. In this context, we make the following observations:

- Néel relaxation and Brownian relaxation times can be very similar in the MPI tracer regime of 15–30 nm. For example, the model predicts a 'crossover' event whereby the zero field Néel time constant becomes longer than the zero field Brownian time constant in this range.

- The Néel relaxation maps have a much stronger sensitivity to the applied field strength. Thus, while for larger core sizes the zero field Néel relaxation time may be orders of magnitude larger than the zero field Brownian relaxation time, the situation may be reversed for high field strengths (*e.g.*, $\geq$ 10 mT). This implies a high degree of crossover and mixing of mechanism dominance in the complex applied field geometries encountered in every MPI excitation trajectory.

- Brownian relaxation is much less sensitive to the magnetic core size, but is very sensitive to the shell thickness, or core-to-shell ratio. As an internal process, Néel phenomena are expected to have no dependence on the nature of the shell. This means there should be great flexibility in tuning Brownian behavior by tracer shell functionalization.

6.3 Using Relaxation Information in pMPI

With this modeling apparatus in place, we now consider how to exploit magnetic relaxation information in pMPI. First we consider ways of measuring/fitting/quantifying relaxation information in pMPI and reporting this information in an imaging or sensing format. In this context, we are considering methods orthogonal to the tracer mass imaging of Chapter 5 and akin to the relaxation maps we constructed in Chapter 2. We then consider how we might incorporate relaxation information into the tracer mass imaging to benefit parameters of merit such as resolution.

6.3.1 Orthogonal Relaxation Imaging Approaches

In Chapter 2, we explored gridding time delays between adjacent sinusoidal excitation half-periods against the mean location of the FFR in 1D. Although this method is indirect and blurs together relaxation information across the whole sinusoidal oscillation, this did provide us with our first relaxation maps and an MPI image distinct from the standard tracer mass/concentration images. These relaxation maps may be valuable on their own or may be overlayed or otherwise combined with the standard tracer mass image from the same dataset (*e.g.*, used to fit an N-compartment model for color MPI as in chapter 3). Using pMPI encoding, we can produce more resolved relaxation maps that directly report information about the tracers' impulse response(s).

We will explore two distinct approaches to reconstructing relaxation maps and measuring relaxation behavior in pMPI: square wave impulse response maps and steady-state recovery analysis.

Square Wave Impulse Response Fitting

In the case of square wave pMPI, the raw time domain signal associated with each half-period is a superposition of the impulse response of the tracer distribution in space, weighted both by the Langevin function (the change in magnetization from the initial field state to the stepped-to field state) and the field dependence of relaxation. In Fig. 6.3(a), raw time domain data associated with a single half-period of a square wave excitation is shown. We can easily fit the data to a parameter of merit, such as a single exponential time constant as shown.

Beginning with the raw time domain voltage equation from chapter 5 for 1D square wave excitation, associated with each n-th half period of the square wave, we have:

$$s_n(t) = V(t)\Big|_{t_n}^{t_{n+1}} = B_1 \int_u \Delta M_n(u) h_{\mathrm{m}}(t - t_n, u - u_n)\,\mathrm{d}u\Big|_{t_n}^{t_{n+1}} \tag{6.9}$$

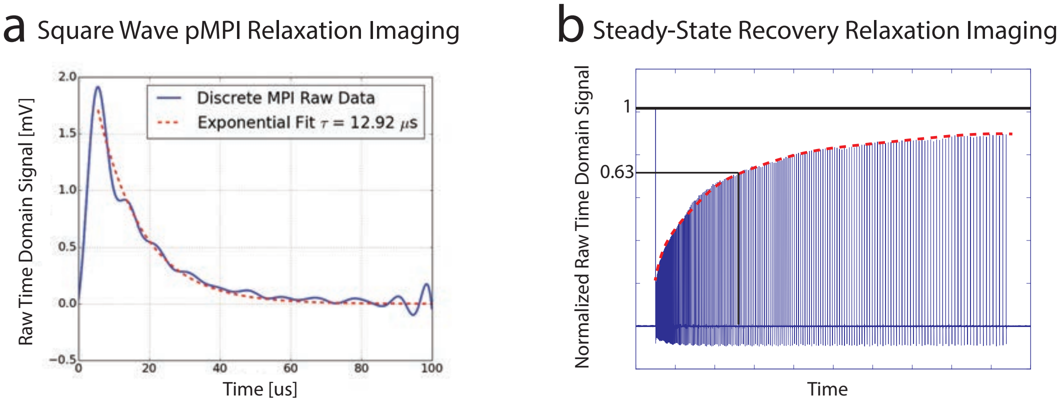

Directly measure/fit raw received data associated with each square wave half-period, which is exactly a (spatially weighted) relaxation impulse response.

Expose the tracer to a train of short pulses, each separated by increasingly longer inter-pulse periods. The quiescent time to achieve steady-state can be measured or fit from the raw data.

FIGURE 6.3: Raw time domain encoding of relaxation information in pMPI examples. (a) In square wave pMPI, the time domain signal associated with a square wave half-period is exactly the (spatially weighted) magnetic impulse response for the tracer stimulated by the step excitation. (b) Alternatively, we can consider a more complex pMPI excitation waveform in which short pulses are applied with longer and longer inter-pulse quiescent hold periods. We can then measure the inter-pulse time required to achieve steady-state.

for the square wave pMPI signal equation. Let:

$$\bar{h}_{\mathrm{m}}(t) = B_1 \int_u \Delta M_n(u) h_{\mathrm{m}}(t, u - u_n)\, \mathrm{d}u$$

Now we make two assumptions: 1) our magnetic relaxation dynamics can be modeled by field-dependent impulse responses with an exponential form (as in the prior theoretical analysis of Néel and Brownian physics) and 2) we can approximate a weighted sum of exponentials as a single aggregate exponential. Then:

$$s_n(t) = \bar{h}_{\mathrm{m}}(t - t_n)\Big|_{t_n}^{t_{n+1}} \approx B_1 \Delta M_n A_n \exp(-t/\tau_n)\Big|_0^{\Delta t} \tag{6.10}$$

where we have equated the integration of weighted exponentials to a single exponential with amplitude A and weighted time constant τ and $\Delta t = t_{n+1} - t_n$ is the square wave half-period. We can further impose $A_n = \tau_n^{-1}$ (*e.g.*, a Debye model). This is motivated because we know that after achieving steady-state, the integration of the raw time domain signal must equal the steady-state difference in magnetization before and after our step

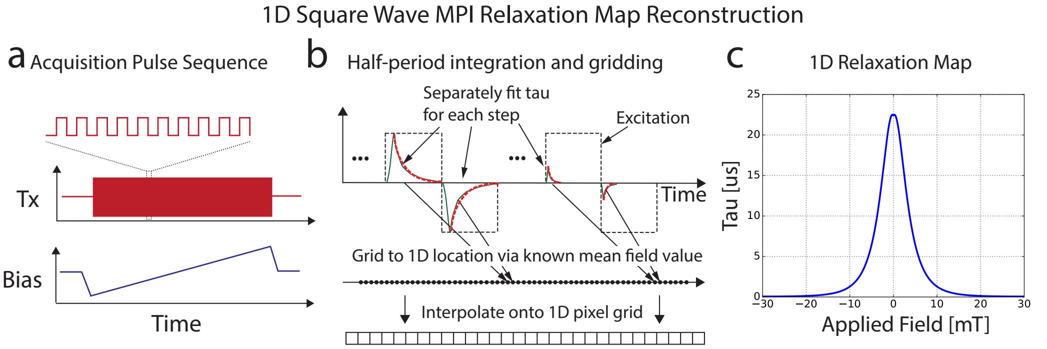

FIGURE 6.4: Relaxation imaging in 1D with square wave pMPI – relaxation maps. (a) We use a simple pulse sequence consisting of a periodic square wave excitation waveform and a linearly ramping bias field to cover the full 1D magnetic FOV. (b) The raw time domain data associated with each half-period reveals the magnetic relaxation impulse response associated with a fixed applied field value. We can fit this to a model of our choosing, *e.g.*, an exponential time constant, and grid the results against the mean location of the applied field producing a relaxation map or 1D relaxation image as shown in (c). This represents a new type of image, distinct from the standard tracer mass/concentration image and which may contain rich physiologic contrast.

excitation. This is true for tracer evolution at a given point and in aggregate. This requires $\int_0^\infty \bar{h}_m(t) = B_1 \Delta M_n$. In this context, we note that $\int_0^\infty \tau^{-1} \exp(-t/\tau) = 1$ and so we must have $A_n = \tau_n^{-1}$ in our model. Because we have access to the full steady-state data, we can integrate the area under the entire half-period to isolate $B_1 \Delta M_n$, allowing us to divide by these values before fitting against our raw data. Finally, then, we can isolate the weighted time constant through various fitting approaches, such as linear regression in $\ln$-space. Let

$$\tilde{s}_n(t) = \frac{s_n(t)}{B_1 \Delta M_n}$$

Then we estimate our 1D relaxation map as:

$$\hat{\tau}(x = x_n) = \hat{\tau}_n \tag{6.11}$$

where:

$$\hat{\tau}_n \leftarrow \min_{\tau_n} \sum_i \left(\ln(\tilde{s}_n(t_i)) - (\ln(\tau_n^{-1}) + (-t_i/\tau_n)) \right)^2 \tag{6.12}$$

In this manner, we can calculate a single value characterizing magnetic

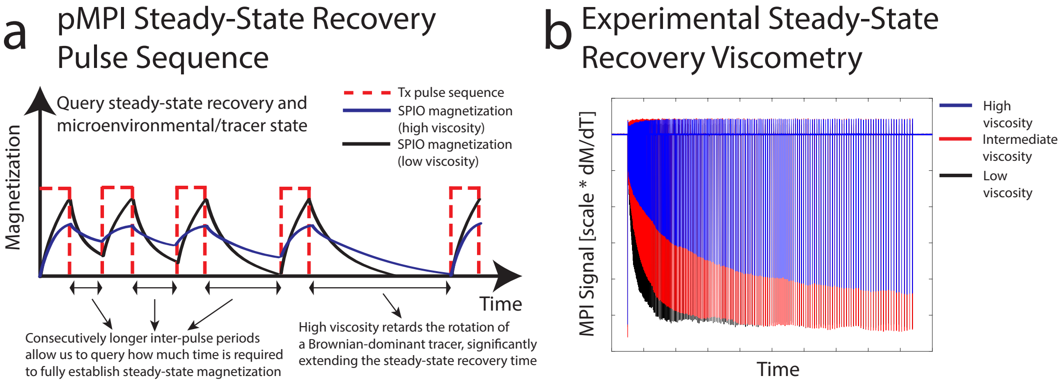

FIGURE 6.5: Relaxation analysis with pMPI steady-state recovery pulse sequence. (a) A simplified depiction of the tracer response to a train of excitation pulses separated by increasingly longer inter-pulse periods. This allows us to query how long it takes for the tracer distribution to fully establish steady-state conditions. For a Brownian-dominant tracer, any changes to microenvironmental state, such as viscosity, binding, or similar parameters, will change the time required for steady-state induction. (b) Experimental data showing a tracer's MPI response to a steady-state recovery sequence as a function of viscosity.

relaxation associated with each square wave half-period. Fig. 6.4 shows how this fits into an overall reconstruction strategy to produce a relaxation image using our AWR. In the simplified case of a relaxometer, we need only the square wave excitation waveform and a linearly ramping bias waveform as depicted in Fig. 6.4(a). We then fit a single parameter of merit for each square wave half-period as in (6.12) and grid to the mean field location as depicted in Fig. 6.4(b). After interpolating to a desired output pixel grid, the result is our relaxation image as illustrated in Fig. 6.4(c) which shows the simulation of a relaxation map from our Néel model.

Steady-State Recovery Sequence

Fig. 6.5 describes a steady-state recovery pMPI pulse sequence. This sequence is different than the square wave pulse sequence we have described thus far. Namely, instead of alternating step excitations, we expose the tracer distribution to a series of short pulses with increasing hold times in between. Fig. 6.5(a) depicts a simplified version of how the tracer magnetization might evolve when a sample with significant Brownian character is exposed to a steady-state recovery pulse sequence under high and low viscosity conditions. The inter-pulse period required to achieve steady-state magnetization in between excitation pulses is longer for the case of high viscosity. In general this time is another view of the tracer's magnetic relaxation dynamics. In the limit of simple single exponential relaxation models, this approach should yield data identical to the square wave impulse response fitting approach.

Fig. 6.5(b) shows raw time domain experimental data for an MPI tracer exposed to a steady-state recovery pulse sequence when placed in three different viscous environments. The tracer's magnetic dynamics and sensitivity to viscosity are revealed both in terms of the time required to achieve the maximum signal following a pulse (when steady-state is not achieved during an inter-pulse period, the following response to a pulsed excitation will be reduced) and the envelope of pulsed excitation responses over time.

We can analyze the MPI signal during a steady-state recovery sequence by first considering the magnetization. As in chapter 5, we assume ideal rectangular pulses and LTI magnetic relaxation characterized by impulse response h_m and step response r_m. The magnetization evolution during the application of the n-th pulse of a series of pulses, each with arbitrary duration, and of which our steady-state recovery sequence is one example, can be naturally expressed as a recurrence relation:

$$\begin{aligned}
M_0(x) &= M(t,x)\Big|_{t=t_0} \\
M_1(t,x) &= (M_0(x) + \Delta M_1(x) r_\mathrm{m}(t,x))\Big|_{t_0}^{t_1} \\
M_n(t,x) &= M_{n-1}(t=t_{n-1},x) + \Delta M_n(x) r_\mathrm{m}(t,x)\Big|_{t_{n-1}}^{t_n}
\end{aligned} \tag{6.13}$$

where $\Delta M_n(x) = M_n^\mathrm{ss}(x) - M_{n-1}(t = t_{n-1}, x)$, $M_n^\mathrm{ss}(x)$ is the steady-state magnetization associated with the n-th pulse location, and $r_\mathrm{m}(t,x)$ is the magnetic relaxation step response. In general, because no pulse necessarily achieves steady-state, $\Delta M_n(x) \neq \Delta M_n^\mathrm{ss}(x)$ where $\Delta M_n^\mathrm{ss} = M_n^\mathrm{ss}(x) - M_{n-1}^\mathrm{ss}(x)$. As in chapter 5, assuming MPI inductive signal reception, we can map these magnetization relations to received signal voltages:

$$\begin{aligned}
V_0 &= 0 \\
V_n(t) &\propto \int_u \Delta M_n(u) h_\mathrm{m}(t,x)\Big|_{t_{n-1}}^{t_n}
\end{aligned} \tag{6.14}$$

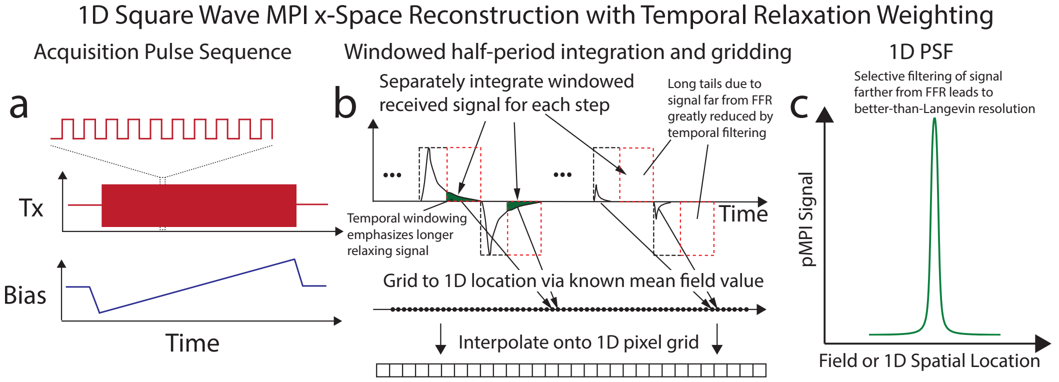

FIGURE 6.6: Temporal relaxation weighting in 1D square wave pMPI. (a) We consider a 1D square wave excitation identical to that described in Chapter 5. (b) During reconstruction, instead of integrating the full area under each half-period, we integrate only a window comprising the final portion of the signal. Here we are exploiting the relaxation spatio-temporal encoding. Signal farther from the FFR relaxes quickly due to the high applied field. We are selectively emphasizing tracer nearest to the FFR with this approach. (c) The resulting gridded image will have greatly improved resolution and attenuated tails. By throwing out a large portion of the signal, however, we will pay an SNR penalty.

Note that as the inter-pulse period tends large, we have $\Delta M_n(x) \to \Delta M_n^{\mathrm{ss}}(x)\,\forall x$ and we observe convergence to steady-state and thus the largest response to the excitation pulses as observed in Fig. 6.3(b) and Fig. 6.5. We can quantify the relaxation dynamics from this steady-state recovery data using various approaches. For example, if we assume an exponential model as we did in (6.12), then we can fit the exponential envelope of the raw data to a single time constant using linear regression.

6.3.2 Combining Tracer Mass and Relaxation Imaging

Being able to separate relaxation information and tracer mass information is very powerful, and this is an important area of difference between canonical sinusoidal MPI and pMPI. But we can also combine the two in specific ways in pMPI to greatly improve parameters of merit such as resolution. They key is exploiting the distinct spatio-temporal encoding from field-dependent relaxation that is illustrated in Fig. 6.1(b,c).

Temporal Relaxation Weighting

Fig. 6.6 illustrates our temporal relaxation weighting method, which is able to greatly improve imaging resolution and remove the 'long tails' associated with the standard Langevin PSF, all at a cost of SNR. Instead of integrating the entire received signal associated with each square wave half-period signal, we only integrate area under a specific window. In particular, we choose a window offset from the beginning of the half-period, only using some final fraction of the signal in our reconstruction. This is motivated by Fig. 6.1(c) and illustrated in Fig. 6.6(b).

From our established theory, we can express the effect of reconstructing using such a window as follows. We begin with the signal equation of chapter 5:

$$\begin{aligned} s_{\mathrm{rlx}}(n) &= \int_{t_n+t_w}^{t_{n+1}} s_n(t)\,\mathrm{d}t \\ &= B_1 \int_{t_n+t_w}^{t_{n+1}} \left(\int_u \Delta M_n(u) h_{\mathrm{m}}(t-t_n, u-u_n)\,\mathrm{d}t \right) \mathrm{d}u \\ &= B_1 \int_u \Delta M_n(u) \left(\int_{t_n+t_w}^{t_{n+1}} h_{\mathrm{m}}(t-t_n, u-u_n)\mathrm{d}t \right) \mathrm{d}u \\ &= B_1 \int_u \Delta M_n(u) \cdot w_n(u)\,\mathrm{d}u \end{aligned} \tag{6.15}$$

where, $s_{\mathrm{rlx}}(n)$ is the windowed time domain signal associated with the n-th square wave half-period. unlike in Chapter 5, we do not integrate the relaxation impulse response h_{m} to unity because our initial temporal integration bound, $t_n + t_{\mathrm{w}}$ is offset by a window threshold t_{w} where $t_{\mathrm{w}} > 0$ in the general case.

We can derive the precise expression for our weighting function w_n as follows (the subscript n indicates that the field- and position-dependent relaxation is centered about the location of the FFR isocenter specific to each n-th square wave half-period). Assessing a given half-period excitation, we let $t_n \mapsto 0$ and $t_{n+1} \mapsto \infty$ where we are assuming the half-period fully establishes steady-state conditions, as in chapter 5. As before, we also assume a general exponential form of the impulse response, but we let the time constant vary over 1D space. Then:

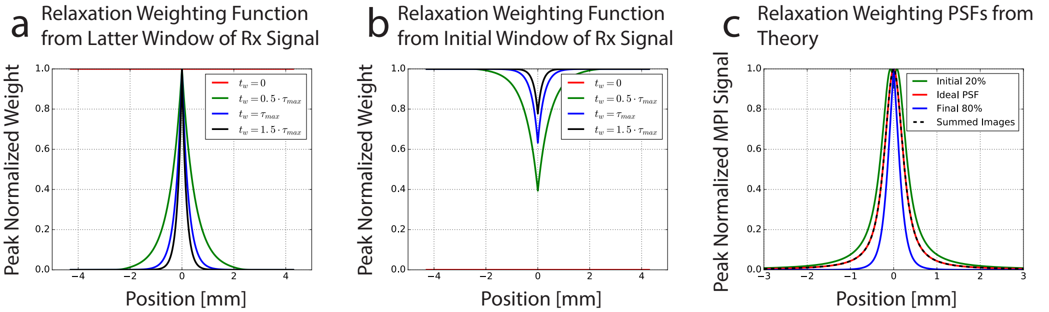

FIGURE 6.7: Temporal relaxation weighting theory and predictions. By simply windowing the received time domain signal associated with each square wave half-period in pMPI, we can exploit relaxation spatio-temporal encoding for improved resolution. The theory predicts a PSF that is the multiplication of a weighting function and the steady-state PSF. (a) The weighting function for a window covering the latter part of the Rx signal, specified by the window threshold t_{w}. (b) The weighting function for a window covering the initial part of the Rx signal. (c) Simulated PSFs for a 27 nm tracer with a 35 nm polymer shell (Brownian dominant). When using only the latter part of the signal, we get a sharper PSF, while using the initial part of the signal yields a more poorly resolved PSF than when using no window at all. The sum of the separately reconstructed PSFs is precisely equal to the image reconstructed with no windowing.

$$\begin{aligned} w(x) &= \int_{t_{\mathrm{w}}}^{\infty} h_{\mathrm{m}}(t,x)\,\mathrm{d}t \\ &\approx \int_{t_{\mathrm{w}}}^{\infty} \tau^{-1}(x)\mathrm{e}^{-t/\tau(x)}\,\mathrm{d}t \\ &= -\mathrm{e}^{-t/\tau(x)}\Big|_{t_{\mathrm{w}}}^{\infty} = -\mathrm{e}^{-\infty/\tau(x)} + \mathrm{e}^{-t_{\mathrm{w}}/\tau(x)} \\ &= \mathrm{e}^{-t_{\mathrm{w}}/\tau(x)} \end{aligned} \tag{6.16}$$

where we see that the weighting introduced by our selective integration is an exponentiated version of our magnetic relaxation map $\tau(x)$. We further note that the signal associated with the initial region of received signal, ($t_{n+1} \geq t > t_n + t_{\mathrm{w}}$), which we have excluded, is given by:

$$w_{\mathrm{init}}(x) = 1 - w(x) = 1 - \mathrm{e}^{-t_{\mathrm{w}}/\tau(x)} \tag{6.17}$$

and in general, the raw received signal may be broken up into N non-overlapping regions, each with a separate exponential weighting function.

The linearity of our system implies that the sum of all images, each reconstructed from data associated with a separate and non-overlapping window, should exactly equal a single image reconstructed from the full dataset without any windowing, *i.e.*, the windowed images represent a simple basis set.

Returning to (6.15):

$$
\begin{aligned}
s_{\mathrm{rlx}}(n) &= B_1 \int_u \Delta M_n(u) \cdot w_n(u)\,\mathrm{d}u \\
&= B_1 \int_u (M_n(u) w_n(u) - M_{n-1}(u) w_n(u))\mathrm{d}u \\
&= B_1 \int_u (m\rho(u)\mathcal{L}(kG(u-u_n))w(u-u_n) - \\
&\qquad m\rho(u)\mathcal{L}(kG(u-u_{n-1}))w(u-u_n))\,\mathrm{d}u \\
&= B_1 \Bigg(m\rho(x) \overset{x}{*} (\mathcal{L}(kGx)w(x))\Big|_{x=x_n} - \\
&\qquad m\rho(x) \overset{x}{*} (\mathcal{L}(kGx)w(x+2\Delta x))\Big|_{x=x_{n-1}} \Bigg) \\
&= B_1 m\rho(x) \overset{x}{*} \big(\mathcal{L}(kG(x+\Delta x))w(x+\Delta x) - \\
&\quad \mathcal{L}(kG(x-\Delta x))w(x+\Delta x)\big)\Big|_{x=x_n-\Delta x} \\
&= B_1 m\rho(x) \overset{x}{*} (\mathcal{L}(kG(x+\Delta x)) - \mathcal{L}(kG(x-\Delta x)))\, w(x+\Delta x)\Big|_{x=x_n-\Delta x} \\
&= B_1 m\rho(x) \overset{x}{*} \left(\left(\dot{\mathcal{L}}(kGx) \overset{x}{*} \sqcap\left(\frac{x}{2\Delta x}\right)\right) w(x+\Delta x)\right)\Bigg|_{x=x_n-\Delta x} \qquad (6.18)
\end{aligned}
$$

where all definitions are the same as in chapter 5, *e.g.*, $2\Delta x = x_n - x_{n-1}$. We can express this as a new imaging equation:

$$\begin{aligned}\hat{\rho}_{\mathrm{rlx}}(x) &= \frac{s_{\mathrm{rlx}}(n)}{B_1 m} \\ &= \rho(x) \overset{x}{*} \left(\dot{\mathcal{L}}(kGx) \overset{x}{*} \sqcap\left(\frac{x}{2\Delta x}\right) w(x+\Delta x)\right)\Bigg|_{x=x_n-\Delta x} \\ &= \rho(x) \overset{x}{*} \left(h_{\mathrm{ss}}(x) w(x)\right) \\ &= \rho(x) \overset{x}{*} \left(h_{\mathrm{ss}}(x) \cdot \mathrm{e}^{-t_{\mathrm{w}}/\tau(x+\Delta x)}\right)\Bigg|_{x=x_n-\Delta x} \\ &= \rho(x) \overset{x}{*} h_{\mathrm{rlx}}(x)\Bigg|_{x=x_n-\Delta x}\end{aligned} \tag{6.19}$$

where $h_{\mathrm{ss}}(x)$ is the steady-state PSF derived in chapter 5. We have arrived at a new relaxation-weighted PSF, in which the steady-state PSF is multiplied by an exponentiated function of the tracer's relaxation map, with a window threshold value t_{w} dictating the exact sharpness of this relaxation weighted function. Fig. 6.7 visualizes aspects of this derivation. In Fig. 6.7(a) we see different peak-normalized weighting functions, $w(x)$, for various window threshold values, t_{w}, while Fig. 6.7(b) shows weighting functions for the initial part of the received signal (w_{init}). The 'one minus' relationship between the two is apparent and we can see that multiplication of the steady-state PSF by the former will yield a sharper PSF and multiplication by the latter will yield an expanded PSF. Fig. 6.7(c) shows simulated predictions for the new relaxation-weighted PSFs from these windows for a 27 nm tracer with a 35 nm shell.

We summarize some of the ramifications of this new imaging PSF:

- By leveraging relaxation spatio-temporal encoding, we are predicted to be able to achieve far better than Langevin resolution by simply windowing our received data before reconstruction.
- We will pay an SNR penalty to achieve this improved resolution as we are discarding signal to achieve this result. In general, by sliding the window threshold around, we can continuously tradeoff resolution and SNR.
- Unlike post-processing strategies such as deconvolution, we are leveraging a basic encoding of the physics. The windowing process requires no *a priori* knowledge to implement nor does it require a model (*e.g.*, PSF) in which modeling errors would lead to artifacts.

- No new acquisition changes were required – we can immediately apply this method to data acquired as described in chapter 5.
- There is an asymmetry that arises due to the steady-state PSF being centered about $x + \Delta x$ while the weighting function (and relaxation map) are centered about x. These relatively shifted components multiply. For small excitation amplitudes (small Δx), this may be negligible, but we expect significant asymmetric skew for large Δx.
- In general, and per its value in physiologic contrast applications, the relaxation maps, $\tau(x)$, that form the core of the weighting function will vary spatially depending on environmental conditions and tracer state. If severe enough, this could compromise the ideal linearity and shift invariance of this method. Néel-dominant tracer would be expected to be robust to such shift variance.

Additionally, these results allow us to concretely understand the 'peak suppresion' whereby resolution is compromised by reducing the signal near low field. This occurs when we do not wait long enough in square wave pMPI (the half-period is too short) and is generally observed in sinusoidal MPI in terms of relaxation induced blur. Namely, these cases are directly analogous to the case of relaxation-weighted reconstruction using an initial window. In this sense, Fig. 6.5(b) visually and quantitatively describes these various phenomena of relaxation-induced 'peak suppression' of resolution (and accompanying SNR loss).

We finish this section by noting that this new relaxation-weighted imaging equation and PSF provide another way of calculating a 1D relaxation image. For a given window threshold, t_{w}:

$$\hat{\tau}(x) = \frac{-t_{\mathrm{w}}}{\ln(h_{\mathrm{rlx}}(x)) - \ln(h_{\mathrm{ss}}(x))} \tag{6.20}$$

where we compute an estimate of the relaxation map by using a relaxation weighted image and a fully reconstructed steady-state image.

Frequency Domain Filtering

In the previous section, we showed explicitly how the spatio-temporal encoding illustrated in Fig. 6.1(c) can be exploited by simple windowing in the time domain. The crux of this approach is the ability to bias our image toward tracer signal originating nearest to the FFR, which has the longest relaxation times.

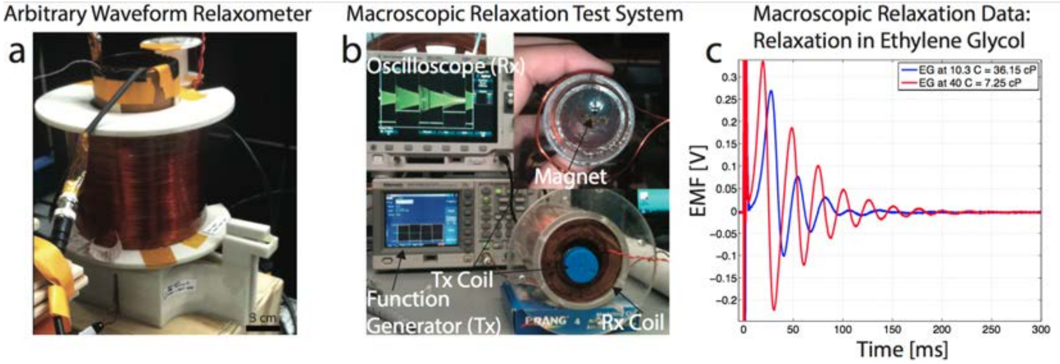

FIGURE 6.8: Hardware used to demonstrate pMPI relaxation methods. (a) Arbitrary waveform relaxometer used for all nanoscopic experiments. (b) Macroscopic relaxation system consisting of a solenoidal Tx coil surrounding a Falcon tube containing a 1 mm Neodymium magnet suspended by a string. An outer high inductance Rx coil surrounds both. The Tx coil was driven by a simple function generator using a square wave excitation while the received signal was obtained from an oscilloscope connected directly to the Rx coil. (c) To explore viscosity-impulse response time, we filled the Falcon tube with ethylene glycol and varied the temperature to vary the viscosity. We can fit the envelope of the decaying underdamped signal and compare to theory.

We can also consider approaching this problem from the frequency domain. The signal we seek that is slowest varying is composed of the lowest frequencies in the Fourier domain. If we filter out higher frequencies components of the receive bandwidth, we should be able to obtain similar effects. Moreover, there may be advantages to a frequency domain approach, such as a better SNR-resolution tradeoff due to the ability to reduce the received bandwidth and an ability to optimize receiver systems in the limit of very low bandwidth or narrow band approaches.

When using a periodic square wave excitation, we establish a pattern in the Tx and Rx signal. The ideal square wave has energy at the fundamental frequency f_0 and all odd higher harmonics, decaying $\propto k^{-1}$. The received signal is limited to small bands around these same harmonics. In pMPI, f_0 is fixed by our need to establish steady-state conditions during a square wave half-period. We will explore the signal- and resolution-bandwidth tradeoffs of frequency domain relaxation filtering experimentally and compare with temporal relaxation filtering.

6.4 Materials and Methods

Our experimental methods are similar to those described in Chapter 5 where the hardware and implementation of pMPI using our AWR (shown again in Fig. 6.8(a)) were described in detail. The primary differences in our pMPI methods in this chapter concern some pMPI encoding changes, namely the use of some non-square wave excitation waveforms, and a number of new reconstruction and fitting techniques.

In general, our experimental approach is to demonstrate each of the relaxation exploitation schemes we have described in section 6.3 with proof-of-concept experiments. We break these down into sections on relaxation mapping, steady-state recovery acquisitions, and relaxation weighted imaging in 1D. In some cases we provide detailed experimental data to characterize specific aspects of these approaches and explore the highly promising arena of microenvironmental sensitivity.

In addition to AWR nanoscopic experiments, we use a separate system shown in Fig. 6.8(b) to measure magnetic relaxation at the macroscopic scale. This allows us to compare against the underdamped and overdamped regimes of the Brownian relaxation theory.

6.4.1 AWR pMPI

When using the AWR setup, all tracer samples consisted of 20-40 $\mu\mathrm{L}$ of stock tracer placed in PCR tubes. As before, we heavily leverage monodisperse tracer samples from Imagion Biosystems (18.5, 21.6, 24.4, 27.4, and 32.1 nm). We also test the nanomag MIP tracer (micromod Partikeltechnologie GmbH, Germany) that we explored extensively in Chapter 2–4. We use some new tracers – a custom 33 nm monodisperse tracer from LodeSpin Labs (Washington, USA), and a dextran coated polydisperse tracer VivoTrax (Magnetic Insight, Inc., USA) that is essentially identical to the canonical MPI tracer Resovist. Finally, we obtained specialized streptavidin-coated versions of an Imagion Biosystems tracer (25 nm) and the nanomag MIP tracer for use in binding experiments.

Our pMPI excitation waveforms include many square wave excitations as well as the newly described steady-state recovery pulse sequence. The latter are composed of square-like pulses with intervening quiescent periods. We lengthen these quiescent periods over the course of a steady-state recovery acquisition, typically using a power law increase in interpulse period. We have found that interpulse periods that vary with a power of two provide the best balance between density of sampling and range of interpulse periods covered with our desired ≈1 second of total scan time and per typically encountered steady state recovery times.

We use the same setup and conditions as described in detail in Chapter 5 to construct our pMPI pulse sequences. For the case of steady-state recovery sequences, we use the same pre-emphasis calculating tools, only changing the duration of quiescent hold periods between our pre-emphasized pulses. Regardless of experiment type, the total scan time per experiment never exceeded 1 second. In some cases, up to 25 trials were run in series and averaged together for improved SNR. For some experiments, 3–5 replicates were performed to report statistical variation.

We test our ability to detect microenvironmental state in several ways – viscometry, pH and ionic strength, and binding. All of these approaches consisted of adding some state modifying fluid to stock solutions of tracer in our PCR sample setup. For viscometry experiments, we mixed various volume fractions of stock aqueous solution and glycerol, for pH experiments we mixed buffered acidic solutions with stock tracer solution to achieve a desired pH, and for binding experiments we used a streptavidin-biotin model in which a saturating amount of biotinylated albumin was added to stock solutions of special streptavidin-coated (shell functionalization) tracers provided by commercial suppliers Imagion Biosystems and Micromod.

6.4.2 Macroscopic Relaxation Experiments

We include data demonstrating 'macroscopic relaxation' to test our simplified Brownian model in the underdamped case. Fig. 6.8(b) shows our simple experimental setup used to obtain these data. A simple copper wire Tx coil was wound around a Falcon tube containing a ≈ 1 mm cylindrical Neodymium magnet suspended by a string which can freely rotate to align with the field applied by the Tx coil. The magnet and Tx coil were placed inside of a larger high inductance (> 1000 turns) Rx coil. For our simple experiments, we used a function generator (Tektronix AFG 3102, Tektronix, Inc., USA) to power a slow and low amplitude square wave excitation (f_0 ≈ 2 Hz at ≈ 0.15 mT). A DS0X2004A oscilloscope (Agilent Technologies, Santa Clara, United States) connected to the receive coil was used for signal reception.

This simple setup was sufficient to provide a strong magnetic relaxation impulse response signal due to use of a solid 1 mm Neodymium magnet and a very sensitive receive coil. To test the relationship between viscosity and impulse response time, the Falcon tube chamber was filled with ethylene glycol and heated while acquiring data. A magnetic safe temperature probe (Neoptix, Canada) placed in the chamber was used to measure the temperature associated with each square wave excitation half-period. By correlating the temperature with a known temperature-viscosity calibration table for ethylene glycol, we can report our data as a function of viscosity. Fig. 6.8(c) shows raw time domain signal from square wave excitation half-periods at different temperatures and viscosities. To report impulse response times, we fit an exponential to the decaying envelope of the raw signal.

6.4.3 Reconstruction

As described in Chapter 5, an initial DSP step of time domain baseline subtraction using void baseline scan data was performed in all cases. This is a robust process on our AWR and helps remove any interference and feedthrough that remains after the action of our gradiometric cancellation and SR560 preamplifier. Subsequently, we generally directly follow the mathematics outlined in 6.3 to reconstruct relaxation images, measure peak response times, and construct relaxation-weighted images. These algorithms were implemented in the Python and Rust programming languages.

A common pattern in pMPI reconstruction is calculation of a single value associated with the data for each half-period excitation. This is true whether we are reconstructing square wave mass/concentration images as in Chapter 5, or if we consider the new relaxation-based approaches in this chapter. Each value is gridded to 1D space according to the mean location of our biasing field and any pMPI-specific shifts described in our mathematical derivations. For example, one of the main methods for reconstructing relaxation maps as outlined earlier is fitting an exponential time constant to the decay associated with each square wave half-period as illustrated in Fig. 6.3. In the case of relaxation-weighted reconstruction, we perform some specific DSP (Fourier harmonic filtering or time-domain windowing) before proceeding with integrated gridding as described in Chapter 5 and illustrated in Fig. 6.6.

For steady-state recovery acquisition, we hold a constant bias field (often zero field) on top of which a long pulsed excitation train is applied. We calculate steady-state recovery times by fitting an exponential model to the envelope of pulse responses as depicted in Fig. 6.3(b) and associating each time with the bias field used in that acquisition. For many of the steady-state recovery experiments, we do not develop a full 1D image but instead

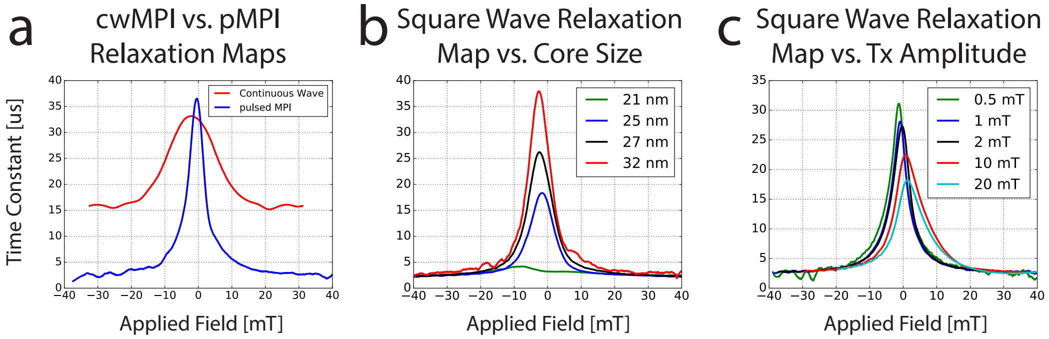

FIGURE 6.9: Square wave pMPI relaxation maps. (a) Comparison between a continuous wave relaxation map (data from Chapter 2: $f_0 = 3\,\mathrm{kHz}$ and $10\,\mathrm{mT}$) and square wave pMPI direct impulse response fitting ($f_0 = 2.5\,\mathrm{kHz}$, $2\,\mathrm{mT}$) for the 32 nm Imagion Biosystems tracer. (b) Comparison of fitted square wave relaxation maps as a function of tracer core size using the monodisperse Imagion Biosystems tracers ($f_0 = 2.5\,\mathrm{kHz}$, $5\,\mathrm{mT}$). (c) Comparison of fitted square wave relaxation maps as a function of excitation amplitude using the 27 nm tracer and a square wave $f_0 = 2.5\,\mathrm{kHz}$.

are interested in changes in zero field relaxation as a function of some microenvironmental variable of interest such as viscosity, pH, or binding state.

6.5 Experimental Results

Here we experimentally demonstrate the relaxation methods described in section 6.3. We first demonstrate relaxation maps constructed by fitting square wave excitation impulse responses. We then explore use of steady-state recovery sequences to robustly quantify aspects of the tracer microenvironment through relaxation contrast. Finally, we demonstrate and characterize relaxation-weighted imaging with square wave excitation, using both time- and frequency-domain methods.

6.5.1 Square Wave Relaxation Imaging

Fig. 6.9 shows 1D relaxation images (relaxation maps) using square wave excitation data. In Fig. 6.9(a) we compare a typical square wave pMPI relaxation map ($f_0 = 2.5\,\mathrm{kHz}$, $5\,\mathrm{mT}$), with a continuous wave relaxation map ($f_0 = 3\,\mathrm{kHz}$, $10\,\mathrm{mT}$) described in Chapter 2. Both of these relaxation maps were taken using 32 nm tracer from Imagion Biosystems. The square wave relaxation map is narrower and more resolved and also achieves much lower absolute time constant values for higher field conditions.

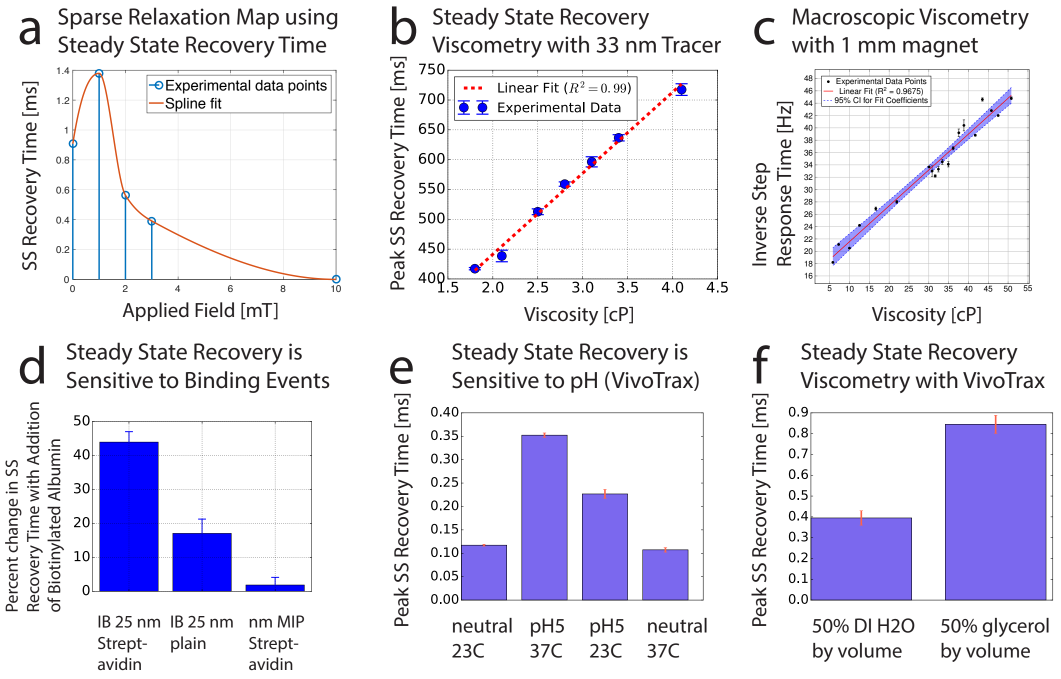

FIGURE 6.10: Pulsed MPI steady-state (SS) recovery data. (a) A sparse relaxation map constructed with SS recovery times sampled at 5 discrete bias field offsets. (b) Zero field SS recovery times change significantly across a physiologically-relevant range of viscosities for a 33 nm tracer. Viscosity was controlled by volume fraction of glycerol and water. In this overdamped nanoscopic case, a linearly varying relaxation time with viscosity is observed. (c) Step response time from viscometry with a 1 mm magnet. In this underdamped macroscopic scale, an inverse relationship between viscosity and the step response time constant is observed as predicted by theory. (d) SS recovery can detect binding events. Here, a larger monodisperse and streptavidin-coated tracer (25 nm, Imagion Biosystems) and a smaller polydisperse and streptavidin-coated tracer (nanomag MIP) were both exposed to a saturating amount of biotinylated albumin. A control non-labeled version of the monodisperse tracer was also tested. In terms of the labeled tracers, only the larger monodisperse tracer shows a non-trivial change in SS recovery time. The untagged control shows a significantly smaller but clear change as well. Even the traditional Resovist tracer (here as VivoTrax) can benefit from SS recovery pMPI, showing sensitivity to pH (e) and viscosity (f).

Fig. 6.9(b) explores the effect of tracer core size on square wave relaxation maps ($f_0 = 2.5$ kHz, 5 mT). As expected from theory, there is a strong dependence on the magnetic core size. A peak relaxation time of only ≈4 μs was observed for the small 21 nm tracer, rapidly lengthening to ≈17, 26, and 37 μs for our 25, 27, and 32 nm samples, respectively. Some minor relative peak shifts are observed with no clear pattern between the relaxation maps. The FWHM values of the relaxation maps were, however, quite conserved, all appearing in the ≈ 8–10 mT range.

Fig. 6.9(c) explores the effect of excitation amplitude on square wave relaxation maps when exciting at $f_0 = 2.5$ kHz and using the 27 nm monodisperse tracer from Imagion Biosystems. There is a clear dependence on the square wave excitation. As the excitation amplitude is increased, the relaxation map begins to spread out and shift slightly in the direction of excitation. For the smallest amplitude of 0.5 mT, a peak relaxation time of ≈ 31 μs is observed compared to a peak of ≈18 μs in the case of a 20 mT excitation. This is a 42% decrease in peak relaxation time for a 40x increase in excitation amplitude. The FWHM of the relaxation maps goes from ≈7 mT for 0.5 mT excitation to ≈10 mT for a very 43% increase in FWHM for a 40x increase in excitation amplitude. Overall, a consistent change of ≈ 2% reduction in peak relaxation time and increase in FWHM is observed per mT increase in excitation amplitude across the range of excitation amplitudes tested.

We note that the 8–10 mT values for the relaxation image PSF FWHM are equivalent to 1.1–1.4 mm 1D resolution assuming a 7 $\mathrm{T\,m^{-1}}$ gradient. This means both that pMPI relaxation images may be relatively high resolution on their own (independent of tracer mass imaging), and that we can expect relaxation weighting, which exploits this *second* spatial encoding, to greatly improve resolution.

6.5.2 Steady-State Recovery Relaxation Imaging

Fig. 6.10 explores the use of steady-state recovery (SS recovery) pMPI pulse sequences. We largely explore the ability of SS recovery sequences to provide relaxation contrast based on microenvironmental conditions.

Fig. 6.10(a) shows a very sparse SS recovery 'relaxation map' with only 5 distinct applied field data points joined by a spline fit. This serves to illustrate both that we can build relaxation maps using our steady-state recovery technique but also highlights a major scaling difference we have seen between relaxation times calculated via SS recovery methods and those calculated by direct fitting of square wave impulse responses as in Fig. 6.9. SS recovery analysis yields up to an order of magnitude larger relaxation time

– even for the same tracer. Whereas SS recovery times are typically in the 100s of $\mu\mathrm{s}$ or even greater than 1 ms, peak square wave impulse response times are typically much less than 100 $\mu\mathrm{s}$. We will have more to say about interpreting these results in our discussion section.

Fig. 6.10(b) shows very promising viscometry results using steady-state recovery sequences at a zero bias field. These results were obtained by testing the LodeSpin 33 nm tracer sample in various volume fraction mixtures of glycerol and water. These mixtures have well known viscosities, allowing us to map our measured steady state recovery times against viscosity. Here we show a linear dependence of the relaxation time with viscosity over a small and physiologically relevant range of ≈1.5–4.5 cP. Over this range, the measured tracer relaxation time varies from ≈410 $\mu\mathrm{s}$ to 710 $\mu\mathrm{s}$ for a slope of about 127 $\mu\mathrm{s}\,\mathrm{cP}^{-1}$. This linear relationship between relaxation time and viscosity matches precisely with our overdamped Brownian theory of 6.2.2. Three acquisition and fitting replicates were performed per viscosity state demonstrating very little statistical variation in the measurement process. Our results indicate great sensitivity, repeatability, and robustness of this measurement.

Fig. 6.10(c) shows macroscopic viscometry results using a 1 mm magnet and our simple macroscopic square wave system. As previously described, viscosity was varied by modifying the temperature of the measurement chamber filled with ethylene glycol and correlating this with well known temperature-viscosity calibration tables. As the theory of section 6.2.2 predicts, in this underdamped regime we observe an inverse linear relationship between measured relaxation time and viscosity. Here we see the relaxation time of the magnet vary from $\approx$ 53 ms to 22 ms over the much broader viscosity range of ≈5 – 50 cP. In terms of inverse relaxation time, this corresponds to a range of 19 – 45 Hz and a slope of approximately 0.58 $\mathrm{Hz}\,\mathrm{cP}^{-1}$. Here we also performed three acquisition and fitting replicates per viscosity (temperature) to report statistical variation in our measurement process.

Fig. 6.10(d-f) further explore sensitivity to microenvironmental conditions for other tracers. In Fig. 6.10(d) we show the relative change in measured relaxation time before and after applying a saturating amount of biotinylated albumin to stock samples of streptavidin coated 25 nm Imagion Biosystems tracer (left), uncoated 25 nm Imagion Biosystems tracer (center), and streptavidin coated nanomag MIP (right). Because albumin and aggregated albumin are physically large, we expect attaching biotinylated albumin to a streptavidin coated tracer to significantly increase the hydrodynamic diameter/size of the tracer. If a tracer's MPI response is non-trivially driven by Brownian rotation, we expect this increase in hydrodynamic size

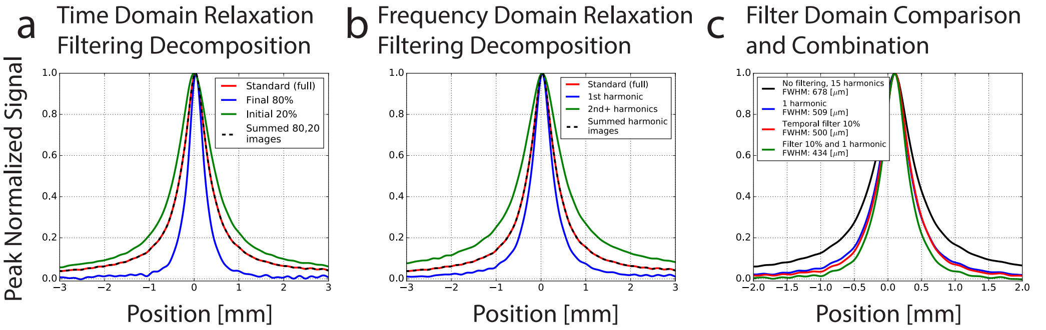

FIGURE 6.11: Exploiting relaxation spatial encoding with relaxation weighting in pMPI. (a) Time domain relaxation filtering using square wave excitation data ($f_0 = 2.5$ kHz and 2 mT) and a windowing threshold of 40 μs (20% of a half-period). We confirm the theoretical predictions: the PSF constructed from the initial window is broader than the standard square wave PSF while the PSF from the latter window is narrower with greatly reduced tails. Additionally, when we sum the separately reconstructed windows, they precisely recapitulate the standard (no windowing) square wave PSF. (b) We accomplish very similar results when reconstructing PSFs using only information from the first harmonic (narrower PSF, reduced tails) and all higher harmonics (broader PSF). The sum of these two PSFs is also equal to the standard square wave PSF. (c) Using a temporal window of 20 μs and reconstruction with only first harmonic data are almost exactly equivalent. We can also perform *both* temporal windowing and harmonic filtering for improved results.

to significantly lengthen the relaxation time per our theory developed in section 6.2.2. We observe such a change along these lines in the case of the larger monodisperse tracer coated with streptavidin but see no statistically significant change in relaxation time for the smaller multi-core nanomag MIP tracer. Interestingly, the uncoated monodisperse tracer also showed a smaller, but significant change in relaxation. This implies that bulk viscous or related changes due to addition of a large saturating amount of biotinylated albumin led to increased relaxation times independent of binding phenomena. We will further discuss sensitivity and specificity issues in our discussion.

Fig. 6.10(e) shows relaxation time results for the VivoTrax (essentially Resovist) tracer exposed to changes in pH and temperature. Stock VivoTrax at room temperature and at 37 °C shows identical relaxation times, however, by buffering the stock solution to pH 5, we see a change in relaxation time both at room temperature and 37 °C. Moreover, a difference between the two temperatures emerges when the tracer solution is buffered to pH5.

Fig. 6.10(f) shows steady-state recovery results when stock VivoTrax solution is mixed with equal parts by volume of de-ionized (DI) water and glycerol. First, we note that the DI water results are significantly longer

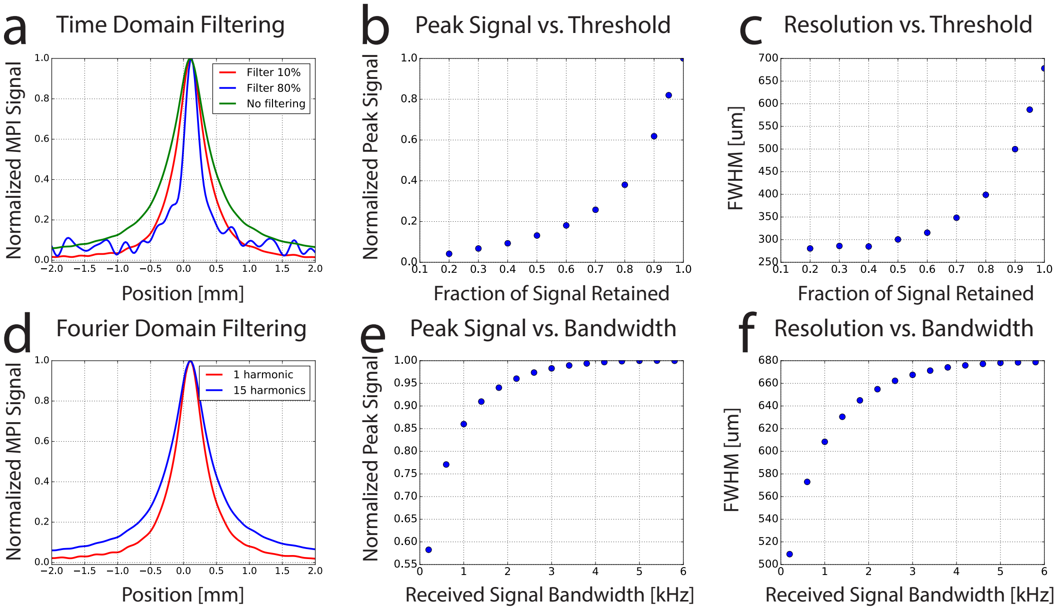

FIGURE 6.12: Signal and resolution tradeoffs using pMPI relaxation weighting methods with our 27 nm tracer ($f_0 = 2.5\ \mathrm{kHz}$, 2 mT). In (a–c) we show relationships for time domain filtering. (a) Increased window-based filtering leads to narrower PSFs but at a cost of SNR. (b) Peak signal as a function of the fraction of the half-period data retained, based on specification of the window threshold. (c) FWHM resolution as a function of the half-period data retained. In (d–f) we show relationships for frequency domain filtering. (d) We can reconstruct a relatively high SNR and modestly more resolved PSF by only using information associated with the fundamental frequency of the square wave excitation. (e) Peak signal as a function of the received bandwidth, in this case, by choosing the number of harmonics included in reconstruction. (f) Resolution as a function of the received bandwidth. Note these results (improved resolution with reduced bandwidth) are opposite that predicted by steady-state Langevin theory in continuous wave MPI.

than the relaxation times associated with unaltered stock VivoTrax solution, as shown in Fig. 6.10(e). In comparison to both of these cases, the mixed stock-glycerol condition leads to a much longer observed relaxation time. This is as expected if, for example, VivoTrax (Resovist) has some significant Brownian character. In general the results of Fig. 6.10(e,f) show that relaxation-based pMPI approaches can be used with smaller and canonical MPI tracers in MPI, *i.e.*, we are not limited to the larger core tracers.

6.5.3 Relaxation Weighted Imaging

Fig. 6.11 experimentally demonstrates the effects of our proposed relaxation weighting mechanisms in which we use the secondary spatial encoding of relaxation to improve resolution in our tracer mass images. All data was taken with square wave excitation conditions of $f_0 = 2.5\ \mathrm{kHz}$ and a Tx amplitude of 2 mT.

Fig. 6.11(a) shows that temporal relaxation filtering behaves precisely as predicted by the theory of section 6.3.2. Here we choose a relaxation window threshold of 40 $\mu\mathrm{s}$, or 20% of the 200 $\mu\mathrm{s}$ excitation half-period. The PSF constructed from the initial window is broader than the standard (no windowing) square wave PSF, while the PSF constructed from the later window is narrower with significantly reduced 'tails'. Additionally, the sum of the two windowed PSFs exactly equals the standard square wave PSF.

Fig. 6.11(b) shows frequency domain filtering behavior that is very similar to the temporal windowing behavior. The PSF constructed solely from first harmonic information is narrower and has significantly reduced tails while the PSF constructed only from higher harmonic information is broader than the standard square wave PSF. As before, the sum of the two PSFs exactly equals the standard square wave PSF.

Fig. 6.11(c) directly compares temporal and frequency domain relaxation filtering mechanisms, showing that a first harmonic PSF and a PSF constructed using a temporal window of 20 $\mu\mathrm{s}$ are almost identical, given the specific tracer and square wave excitation used. Furthermore, we can perform *both* methods during reconstruction: A 20 $\mu\mathrm{s}$ windowed, first harmonic image shows improved FWHM and tails than either method alone.

Fig. 6.12 explores the signal and resolution tradeoffs for temporal relaxation filtering and frequency domain relaxation filtering. This analysis used experimental data using the Imagion Biosystems 27 nm tracer with excitation conditions $f_0 = 2.5\ \mathrm{kHz}$ and amplitude 2 mT. In Fig. 6.12(a-c), the relationships for time domain filtering are shown. We plot the normalized peak signal (to the maximum when no windowing is used) (b) and resolution (c) as functions of the window threshold, given as a fraction of the half-period. Both have a strikingly similar exponential relationship and the range of the tradeoffs possible with time domain relaxation filtering are apparent – we can go from approximately 680 $\mu\mathrm{m}$ resolution down to approximately 280 $\mu\mathrm{m}$ when we use only 20% of the data. This is a 70% (approximately 2.5 fold) improvement in resolution over the steady-state square wave case, itself able to achieve improved resolution compared to canonical sinusoidal MPI. With a more conservative tradeoff, we can achieve 400 $\mu\mathrm{m}$ resolution while retaining 40% of our original signal. We note that this is a representative case study using Imagion Biosystems 27 nm tracer and is in no way limiting for other (*e.g.*, larger) tracers.

Fig. 6.12(d-f) shows the same relationships for frequency domain filtering where we plot the normalized peak signal and resolution as a function of the received bandwidth (*e.g.*, the number of harmonics included in the square wave case). We see a completely different relationship than that described in canonical steady-state Langevin analysis of MPI (Goodwill and Conolly, 2010), notably, we observe an inversion of the resolution-bandwidth relationship, where in this case reducing the bandwidth, and only keeping lower frequency portions in particular, improves resolution at a modest cost of signal. We observe an exponential-like relationship in both the signal vs. bandwidth and resolution vs. bandwidth curves, as in the case of temporal relaxation filtering, although with an inversion of the exponential behavior. In this square wave incarnation, we are limited to the finite number of harmonics that constitute our total signal bandwidth when choosing our signal-resolution tradeoff. Furthermore, the choice of square wave period fundamentally shifts this relationship. This is different than the case of temporal relaxation filtering where the relaxation window threshold can be varied continuously. Relatedly, in the case of the excitation and tracer used here, we are limited to improving resolution to just above 500 $\mu\mathrm{m}$ from the original 680 $\mu\mathrm{m}$ when using only the first harmonic information, retaining approximately 58% of the original peak signal in the process.

6.6 Discussion

Our main goal in this chapter was to showcase several intriguing ways of measuring and exploiting relaxation made possible by the flexibility of pMPI. Futhermore we wanted to make clear the scale of possibilities in pMPI specifically, but MPI in general, when it comes to exploiting tracer physics.

We spent time rigorously exploring the theory that suggests this is possible and then described proof-of-concept experimental results experiments for each of the major approaches we describe. In doing so, many intriguing phenomena were confirmed and others came to light. Here we will try to discuss some of the main ramifications of this work and address some of the most important questions and inconsistencies that arose. We will also make a point to suggest directions for future investigation along these lines.

6.6.1 Comparing Theoretical Analysis and Experimental Results

We spent time constructing specific relaxation models for Néel and Brownian relaxation and developing signal and imaging equations for each of the main new methods: relaxation map fitting, steady-state recovery analysis, and relaxation weighted imaging. We believe the proof-of-concept results we have shown confirm that our theoretical developments have been highly predictive and represent a strong basis for continued development of pMPI methods.

In terms of gross behavior – peak relaxation times, relaxation map shapes and FWHM values – the relaxation maps we reconstructed experimentally were very similar to those predicted by the theory. We do not draw any more specific conclusions due to the difficulty in comparing across the various parameter values easily (outside the scope of this chapter). But we also note that the models predict significant Brownian character/dominance in tracers in the range of 20–30 nm. The steady-state analysis results clearly indicate that tracer relaxation is very sensitive to local microenvironmental conditions, indicating that Brownian phenomena are significant. We were also able to confirm the ability of our simplified torque balance equation results to predict the type of relationship between relaxation time and viscosity at the two extremes of scale: nanoscopic (underdamped) and macroscopic (overdamped). We do note that some of the minor shifts in terms of peak signals relative to zero field (common to all measurements) is due in part to a known minor DC offset in our calibration of the bias field of our AWR system.

Our signal and imaging equations developed for relaxation weighted imaging are strongly supported by our experimental data. For example, Fig. 6.7 and Fig. 6.11 match extremely well. Furthermore, close inspection of Fig. 6.11(a) shows that temporal relaxation weighted results are slightly skewed relative to the standard square wave PSF due to a relative shift (2 mT in this case) between the steady-state square wave PSF and relaxation weighting functions involved, as predicted by our theory and expressed in (6.19). We have further confirmed this phenomenon with larger excitation amplitudes (data not shown) where the skew, which depends on the excitation amplitude, is much more pronounced. We note that this skew is not manifest with our frequency domain relaxation filtering approach. In general, our understanding of temporal relaxation weighting as introducing a series of weighted versions of the standard square wave PSF, with PSFs from latter windows having improved resolution and reduced tails, and in which non-overlapping windows form a basis set that can be summed to recover the full steady-state image has been confirmed.

We believe these results highlight a powerful ability to assess and describe pMPI relaxation methods generally, even when specific relaxation dynamics are incompletely understood.

6.6.2 Relaxation Maps

We described two ways of building relaxation images or maps in this work. These methods are also ways of reporting *only* relaxation information in an imaging format, although this information may easily be incorporated into tracer mass imaging via methods such as those described in Chapter 3.

The impulse response fitting method shown in Fig. 6.9 is able to measure relaxation dynamics in a much more direct way than the cross-correlation method described in Chapter 2. This is due to square wave MPI allowing us to directly observe the relaxation impulse response, untransformed, in the raw data. First of all, we believe this could be a boon for basic science researchers and our field's quest to fully understand the competing relaxation dynamic processes in the high field environment of MPI. Although we do not report detailed theoretical and experimental comparisons or significant testing of the myriad parameters that affect the relaxation map in this chapter, we believe these are important next steps in using these relaxation maps.

While we did show a dependency of impulse response relaxation maps on core size that agrees with theory, we also show that the impulse response relaxation maps have a dependency on the excitation amplitude that is not necessarily predicted by our theory. Instead, our models assume that the absolute field conditions alone determine dynamics. In this context, an excitation at 2 mT or 20 mT should yield the same relaxation maps. But as Fig. 6.9(b) clearly shows, our experimental relaxation maps have reduced peak time constants, become slightly broader, and shift slightly as the excitation amplitude is increased. Going from 0.5 mT to 20 mT excitation leads to about a 42% change in both peak relaxation time and relaxation PSF FWHM.

We believe this could be a manifestation of the angular and vector nature of magnetization that we ignored in our theory (*e.g.*, making the small angle approximation). For example, consider a 1 mT and a 10 mT square wave excitation. When we sample the relaxation at the applied field value of 3 mT in the 1 mT excitation, we are jumping from an applied field value of 1 mT to 3 mT (peak-to-peak amplitude). In this case, we are staying in the same, probably linear, region of the Langevin curve for the tracer. The tracer population will go from lightly polarized, on average, to moderately polarized in the same direction, on average, between these two states. Now consider the 10 mT excitation case where we jump from -7 mT to 3 mT when we query the relaxation at 3 mT. In this case, we go from a state of almost complete saturation aligned in the negative direction to moderately polarized in the positive direction. We might imagine that many more components of the tracer population will move through a much smaller net angle in the former than the latter. How might we expect these angular considerations to affect the relaxation maps? If the larger excitation can be thought of as introducing larger angles to move through, on average, then these larger angles would lead to relatively larger torques than smaller angles, leading to a relatively faster relaxation response, at least at the initial part of the response. Our method of fitting a single exponential to the raw time domain response is quite sensitive to the initial burst of magnetization flux. What we observe is, indeed, that larger excitation relaxation maps show faster overall fitted response times.

We do not assert that this intuitive discussion necessarily explains the excitation dependence of our relaxation maps or that we have demonstrated this to be the case in any way, but we believe thinking in this and related directions can be fruitful when exploring these unexpected behaviors in future work.

6.6.3 Discrepancy Between Impulse Response and Steady-State Relaxation Measures

We focused on steady-state recovery excitation in exploring microenvironmental sensitivity (*cf.*, impulse response relaxation fitting) for a couple of reasons First, steady-state recovery methods appeared to be much more sensitive and robust to microenvironmental-based relaxation contrast in general. Second, and importantly, in early experiments, the impulse response relaxation maps showed the *opposite* behavior as expected when we varied the microenvironmental conditions. For example, when increasing the viscosity using glycerol and water mixtures, the peak relaxation time associated with impulse response relaxation maps *decreased* with increasing viscosity. When we look at raw time domain data associated with individual excitation half-periods, what we see is that the initial part of the exponential-like response is very similar in all cases, but the latter parts of the response are greatly attenuated by higher viscosity conditions. When fitting a single exponential in this regime, of course the latter will be fit as a faster decaying exponential.

It's as if the responses generally consist of a Néel or at least a relatively static initial response in superposition with a much longer Brownian component. By selectively attenuating the Brownian component, a single fit to the combined process will register a faster response in the face of higher viscosity.

With this in mind, we also point out that peak impulse response relaxation times are typically in the range of ≈10–40 μs while the steady-state recovery method measures peak relaxation times typically in the range of 100 μs to 1 ms. This is a very large discrepancy, sometimes close to two orders of magnitude. We can consider the differences in the two approaches: The impulse response fitting fits the raw time domain data directly and is mostly determined by the high $\mathrm{d}M/\mathrm{d}t$ signal immediately after each step excitation. The steady-state recovery approach, on the other hand, queries how long it takes for the tracer to fully establish steady-state. A very slowly varying, but high integrated total area (*i.e.*, energy) component (*e.g.*, that may always be beneath the noise floor in the raw signal) that is superposed with a rapidly varying component could be completely unobserved by the direct impulse response fitting method. The steady-state method, however, will register this component and will report a total response time dictated by this slow component. The discrepency in the reported relaxation time scale could, then, be further evidence of at least two simultaneous relaxation process, one fast and Néel-like and one signficantly slower and Brownian-like.

Considering the parallel model of relaxation and the assumed independence of the various relaxation methods, how can we explain the presence of multiple relaxation mechanisms that differ in time constant by an order of magnitude or more but yet both contribute significant energy to the process? Why would the fast process not completely take care of all magnetization change? These are important questions to answer in the future. Perhaps the processes are not so independent and we could have a scenario where a 'fast Néel' process can quickly align to the nearest easy axis (or directly to a hard axis), but the much slower Brownian process must complete the rest of the change in magnetization to achieve the lowest energy, steady state in magnetization.

6.6.4 SS Recovery Sensitivity to the Microenvironment

Setting aside these intriguing discrepancies between the two methods of assessing relaxation time, we have shown in this chapter that pMPI can be very sensitive to local microenvironmental conditions such as viscosity, binding state, and pH/ion concentration state. Fig. 6.10 shows not only that we can indeed have significant relaxation contrast for relaxation imaging but that we can precisely quantify these relaxation effects. Furthermore, we believe the quantitation and sensitivity shown in these data indicate that pMPI methods can be many times more sensitive and robust for color MPI and related molecular imaging applications in MPI than that which is available in continuous wave MPI. We simply have much more flexibility to encode relaxation information into the raw signal with pMPI methods. Specific applications of interest include colorizing live vs. dead cells in MPI cell tracking applications and targeted cancer contrast via parameters such as pH, viscosity, and binding to various markers.

We note that we can perform our SS recovery analysis by operating at a static bias field rather than zero field. In general, this can improve time efficiency and the signal strength, and in some cases, we may prefer the relaxation contrast at some bias field, while at other zero or low field conditions may provide the most desirable relaxation contrast. To use SS recovery to create relaxation images, we must shorten the ≈1 second pulse sequence-per-pixel we have shown in this work. We believe optimized steady-state recovery sequences with many fewer interpulse periods can provide a much better balance between SNR for fitting and throughput.

We also point out that the data of Fig. 6.10 indicate that non-imaging biosensors could be a very interesting application area for pMPI in the future. Although it is always difficult to compete with the fantastic optical technologies for *ex vivo* analysis, biosensors applications such as glucose sensing could be attractive. For example, a small hydrogel sensitive to a target analyte could be embedded with large MPI tracers and placed subcutaneously. A small wristwatch-sized pMPI biosensor device could then report the viscosity and therefore analyte concentration locally. Advantages of such an approach could include relative insensitivity to tracer mass or concentration (we only fit the dynamics, which is independent of the amount of tracer as long as we have a certain baseline SNR) and perhaps an easier time with biofouling issues due to lack of a 2D electrode-like interface.

6.6.5 Sensitivity and Specificity in Relaxation Encoding

One consideration we must contend with in pressing forward with pMPI relaxation contrast applications is specificity. Taken together, the data of Fig. 6.10 show that we can be very sensitive to many different local environmental conditions, but we may be *too sensitive* in these proof-of-concept experiments in the sense that we have poor specificity. For example, the difference in the VivoTrax signal of the neutral, room temperature case of Fig. 6.10(e) and the case of stock VivoTrax diluted with de-ionized water in Fig. 6.10(f) indicate that simply adding DI water (and changing osmolarities and ionic concentrations) can greatly affect MPI relaxation. This calls into question how much of the results of Fig. 6.10 are due to pH and how much to other ionic states? Furthermore, the significant change in relaxation time for the Imagion Biosystems tracer not functionalized with streptavidin in Fig. 6.10(d) calls into question how much of the streptavidin functionalized change is due to binding and how much to bulk changes to viscosity, etc., from adding a large amount of biotinylated albumin? We do note, that the nanomag MPI streptavidin coated results in this case do indicate that some tracers are impervious to these microenvironmental changes (presumably smaller and Néel dominant).

It is very important that we perform careful experiments in the future to answer these questions and isolate variables. This also implies that work by chemists and material scientists to design optimal tracer shells that have only the specificity we desire will be an important part of pressing these new capabilities into the service of specific applications.

6.6.6 Relaxation Weighting Provides Sub-Langevin Resolution

The ability to use relaxation to improve resolution beyond what Langevin theory predicts is a particularly exciting feature of pMPI, as demonstrated in our relaxation weighted methods. Fig. 6.12 shows that we can achieve below 300 μm resolution using a 27 nm MPI tracer by leveraging time domain relaxation weighting. This is well below the Langevin theory and the best reported native resolution we've ever achieved in MPI. This is made possible, again, because of the *second* spatial encoding we have in MPI that is provided by the field dependence of magnetic relaxation dynamics and has been, until now, untapped. In fact, we have completely turned relaxation effects around with pMPI – from previously a resolution limiting factor in continuous wave MPI to a resolution enhancing factor in pMPI.

Aside from providing a pathway to approximately 250 μm resolution and perhaps beyond, relaxation weighting is a particularly simple and flexible approach. We have developed a relatively straightforward imaging equation and theoretical exposition in section 6.3.2 which has been demonstrated convincingly in experiments. Furthermore, relaxation weighting, whether in the time or frequency domain is implemented as a reconstruction step and requires no changes to, for example, a standard square wave pMPI excitation (although we may find that certain parameters optimize this flexibility). This is very significant – no extra up front effort beyond the initial square wave encoding is required to implement relaxation weighted improvements (as well as to construct relaxation maps).

Given that resolution is arguably the most lacking feature of MPI, having an ability to tradeoff resolution and SNR continuously to achieve previously unheard of resolution is compelling.

6.6.7 SNR Tradeoff and Challenges in Relaxation Weighted pMPI

We do want to mention a few words of caution in leveraging relaxation weighted pMPI for high resolution imaging. First of all, it does come with an SNR cost. Given that square wave pMPI itself may have reduced SNR efficiency with respect to canonical MPI methods, we need to avoid paying too steep of a penalty for the improved resolution. In terms of time domain relaxation weighting, we note that the binary window used in this work may not be ideal. Guided by theory, use of a ramping or nonlinear weighting applied before integrated gridding, could accentuate relaxation spatial encoding while retaining a better SNR tradeoff. We caution that such approaches need to be rigorously motivated and evaluated and not devolve into ‘widget’ parameters.

Ultimately, the SNR loss per unit improvement in resolution over the best sinusoidal MPI case is an important metric to keep in mind. This also implies that the frequency domain relaxation weighting can be particularly attractive. Although it appears somewhat less flexible (finite harmonic number and relationship fixed with choice of excitation period) than time domain relaxation weighting, it may have better SNR tradeoff parties. Namely, although we lose some peak signal in first harmonic-only imaging to gain resolution, our concomittant *significant* reduction in Rx bandwidth can end up being a net win in terms of total SNR. And if frequency domain relaxation filtering allows for true narrowband signal reception (*e.g.*, only at the fundamental harmonic), then we could gain even more by use of specialized Rx chain electronics optimized for narrowband conditions *cf.*, the current broadband systems.

A second aspect of relaxation weighted pMPI to note is that it may complicate the quantification and shift invariant nature of MPI. While square wave pMPI is decidedly LSI because it removes all traces of possibly shift-varying relaxation, relaxation weighted pMPI directly exploits this relaxation to improve resolution. Ultimately, a relaxation weighted image is locally scaled and shaped by local relaxation conditions. Especially when we are intentionally drawing such contrast out in relaxation imaging, if relaxation conditions are sufficiently different across a sample, this will manifest in our imaging as shift variance. The potential to disrupt quantitation is perhaps the most concerning. We will need to keep this in mind when deploying pMPI in a full imaging context in the future.

6.6.8 Future Directions

We believe pMPI is an exciting new platform in MPI, and we have only scratched the surface in this and the previous chapter. Future work should expand on the proof-of-concept experimental demonstrations of this chapter and develop a rigorous characterization of these methods. It is also important to demonstrate pMPI in a multi-dimensional imaging format by modifying existing scanners to support the new Tx/Rx modules required in pMPI. Finally, as our steady-state recovery pulse sequence has indicated, pMPI offers a huge flexibility in excitation encoding beyond that which we have explored in this work. We believe many new pulse sequences await discovery that can improve various aspects of pMPI, perhaps in application-specific ways, through modification of the basic signal encoding. For example, the math and theory indicate that relaxation weighting that we have demonstrated on the reconstruction side can apply equally to the encoding specified by signal excitation pulses. In general, a notion of dynamic 'magnetization preparation' followed by a quiescent 'readout' periods emerges when considering possible pMPI pulse sequence design. Collaborative work with tracer developers and chemists will be integral to such approaches.

6.7 Conclusions

In this chapter we have demonstrated several unique ways that pMPI can report and exploit magnetic relaxation. We described the theory and motivation behind relaxation imaging in which we can images of relaxation directly rather than tracer mass or concentration, we proposed a steady-state recovery technique as a very sensitive method of quantifying the impact of microenvironmental conditions on magnetic relaxation, and we proposed methods to exploit the relaxation spatial encoding to improve resolution in tracer mass imaging. We then verify, demonstrate, and characterize these methods with proof-of-concept experiments. We have shown the first relaxation maps in MPI drawn directly from magnetic relaxation impulse responses; show an ability to detect and quantify changes in viscosity, pH, and binding state; and show the ability to improve *native* resolution to better than 300 μm – well below the Langevin limit. We now know that the Langevin theory is not a fundamental limit in MPI and that magnetic relaxation can play a significant role in our spatial encoding. The work in this chapter is only the tip of the iceberg and shows in dramatic fashion how underexploited magnetic relaxation is in canonical MPI and how pMPI can provide an impressively flexible platform for exploiting it.

6.8 Acknowledgements

I would like to thank Zhi Wei Tay for helpful discussions and point out that Zhi Wei made the initial intuitive leap in suggesting window-based temporal relaxation weighting. I would also like to acknowledge and thank Normen Oude Booijink for stimulating and helpful conversations about the mathematics of MPI and relaxation physics. He also helped Zhi Wei and I in exploring methods such as the steady-state recovery technique during his brief but productive stay in our lab. I would also like to thank Xinyi Zhou for her help running some experiments, including the pH and viscometric experiments with VivoTrax, and Arkosnato Neogy for his help running the macroscopic relaxation experiments. Finally, I would like to thank our tracer collaborators at Imagion Biosystems, LodeSpin Labs, the University of Washington, and the University of Florida for providing custom tracers so crucial for this work. Special thanks in particular go to Erika Vreeland at Imagion Biosystems and Matt Ferguson of LodeSpin Labs for providing tracers critically featured in this chapter.

Chapter 7

Generalizing Pulsed MPI and Future Directions

In the previous two chapters, we described and demonstrated many promising features of pulsed magnetic particle imaging, but only in the 1D case. Ultimately, the most important proof-of-concept for an imaging technique is to produce compelling multidimensional images. We are very early in our exploration of pMPI and have only just begun attempts to implement it in a full MPI scanner context. In this short final chapter we consider some of our expectations for and implications of multidimensional pMPI. This discussion will extrapolate upon the intuition and mathematical structure we have explored in detail in the 1D case. We will show numerous results from a multidimensional pulsed MPI simulator we have developed and also show the first multidimensional pMPI image.

7.1 Introduction

We have shown that pulsed excitation waveforms can enable various new abilities in MPI, including the complete removal of relaxation effects in tracer mass imaging, improving resolution beyond the previously fundamental Langevin limit with relaxation weighting, and direct quantification and imaging of tracer relaxation dynamics which can in turn provide quantitative contrast based on the tracer's local microenvironmental state. We have demonstrated all of these effects in a 1D MPI format using our AWR (Tay et al., 2016). The crucial next step is implementation in an MPI scanner. We are currently working in this direction, but doing so does require significant hardware changes to the Tx and Rx systems.

In this short chapter, we will describe our current approach to and perspective on multidimensional pMPI. The main goal is to provide background and framing for future pMPI development. As we will show, we expect multidimensional pMPI to be very similar to the standard case, except being simpler to implement in terms of reconstruction. Many of our distinct pMPI results can be obtained from a single dataset (*e.g.*, steady-state images, relaxation-weighted higher resolution images, and relaxation map images). This flexibility is possible in part because the crux of pMPI is a modification of the fast excitation waveforms we apply to the tracer while the rest of our x-space MPI implementation can stay the same. The other, and fundamental, enabling factor is our ability to encode relaxation in the raw signal in very particularly ways in pMPI, as we described in detail in Chapters 5 and 6.

In addition to describing how multidimensional MPI can be implemented, we will also take some time to describe how we might take advantage of the flexibility of pMPI. In particular, we will discuss generalized pulse sequencing which implies that myriad methods of encoding desired imaging properties in application-specific ways should be possible. This could one day lead to pulse sequence design analogous to that available in MRI. As we consider some of these approaches, some conceptual patterns emerge, such as distinct 'magnetization preparation' and signal readout components of an excitation sequence. We will also mention intriguing, if exotic, possibilities such as excitation along the line when using a field-free line (FFL) system.

We believe the new capabilities enabled by pMPI are incredibly broad and, when more completely developed, may imbue MPI with new imaging characteristics (resolution, molecular imaging contrast) that are much more powerful than that available in the relatively restricted domain of canonical continuous wave excitation. In this contexst, we hope this chapter might serve as a helpful guide in exploring multidimensional pMPI in the near future.

7.2 Multidimensional pMPI

We first describe how we can implement pMPI in a multidimensional context, namely, what multidimensional x-space pMPI acquisition trajectories and reconstruction may look like, in contrast to both continuous wave methods and the simplified AWR system used to develop pMPI in Chatpers 5 and 6. We seek to provide an intuitive description of these methods, which we are actively developing in our in-house scanners. We will show some simulations of multidimensional pMPI that confirm our intuition and proof-of-concept experimental data.

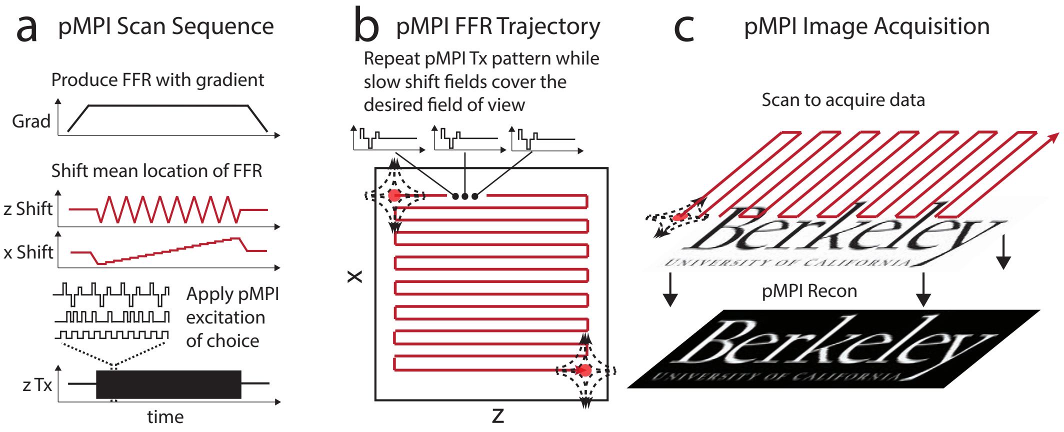

FIGURE 7.1: Depiction of multidimensional pMPI scanning. (a) Our pMPI acquisition is very similar to canonical x-space in terms of the shift field, although we do not necessarily require any overlapping of sample points (a simplification). Instead of a sinusoid, however, we can excite with a periodic pulsed waveform. (b) By repeating our core pulsed waveform as we raster across the imaging FOV, we can sample the image in full with our new encoding. (c) This x-space raster scanning procedure is followed by a simple gridded reconstruction as before to produce images, which may be much higher resolution or even consist of a measure of relaxation dynamics. Here we show a steady-state square wave pMPI simulation of a Berkeley phantom.

7.2.1 Implementing pMPI in an Imaging Format

Fig. 7.1 shows the multidimensional x-space pMPI process, similar to Fig. 1.3 but updated for the case of pulsed excitation. In general, we do not need to significantly change anything about our gradient or shift waveforms as shown in Fig. 7.1(a). These components provide spatial selectivity and allow us to query a large FOV with our sensitive region, respectively, and this paradigm applies generally to MPI, regardless of the nature of the excitation waveform. Due to our ability to temporally decouple excitation feedthrough and our received signal, we can simplify the raster trajectory because we no longer require overlapping patterns (Lu et al., 2013).

In pMPI, as in the case of continuous wave MPI, we apply a periodic excitation waveform while slowly moving our sensitive region around the FOV. This is illustrated in Fig. 7.1(a,b) and Fig. 7.2. This emphasizes how pMPI is concerned primarily with changing the encoding provided by our periodic excitation waveform. We are still just querying the position associated with our FFR over time, but our pMPI waveforms provide improved and more flexible encoding of relaxation information. Fig. 7.1(c) depicts pMPI imaging of a Berkeley phantom, using a large tracer (*e.g.*, ≈ 35 nm core) and a small excitation amplitude (1–2 mT) as described in Chapter

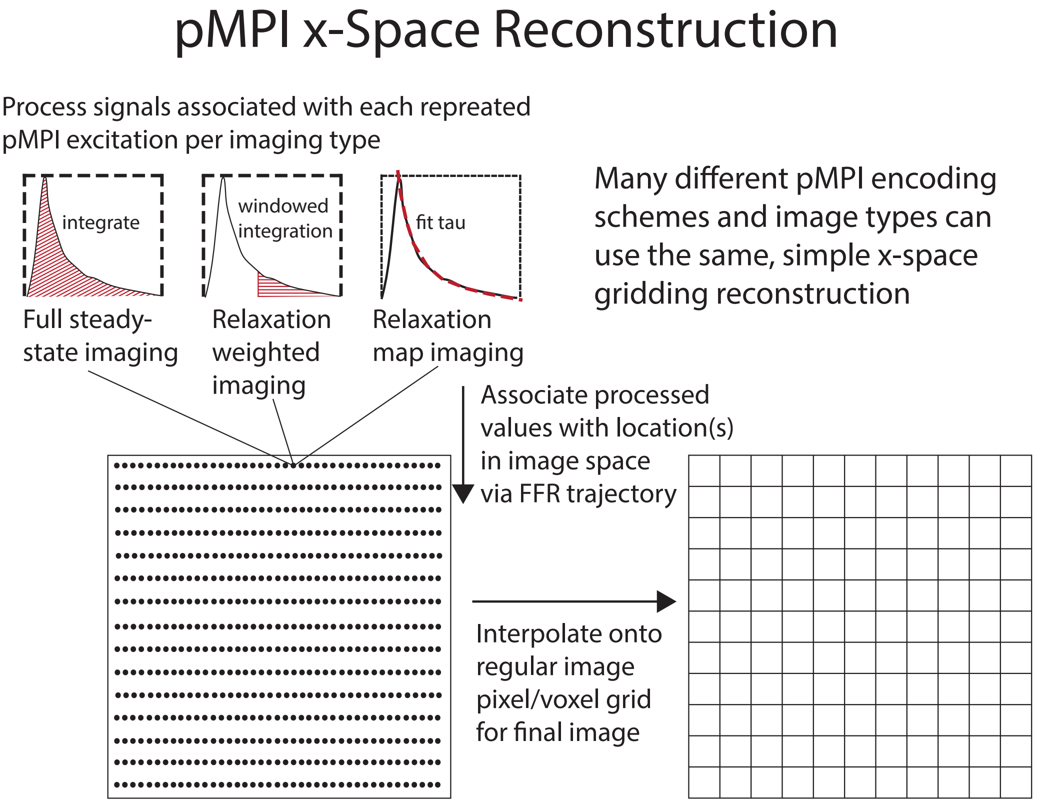

FIGURE 7.2: Depiction of simple x-space reconstruction in multidimensional pMPI. Many pMPI reconstruction approaches fall into the conceptual category of calculating a single value (integrated signal, windowed or not for relaxation weighting, or a measure of the relaxation impulse response) for each period of the periodic excitation waveform and gridding this value to the mean location of the FFR. Aside from DSP steps, the entirety of acquisition and reconstruction simply consists of slowly rastering our FFR across the FOV while applying the fast periodic excitation, gridding, and then interpolating onto a final regular pixel grid.

5. We will analyze results of our simulator in more detail in subsequent paragraphs.

Fig. 7.2 illustrates the simplicity of pMPI reconstruction and how we can obtain various *different types* of pMPI images from the same dataset. For example, as shown, a single square wave excitation acquisition can provide access to relaxation-free and high resolution images per the methods of Chapter 5, a sub-Langevin relaxation weighted image as described in Chapter 6, and a 'relaxation image' where we image mean relaxation time constants over space. All of these different images are obtained from a similar and simple process: first the signal associated with each square wave half-period is processed uniquely (integration, windowed integration, or exponential fitting) and then this single value is gridded to the image domain using knowledge of the mean FFR location associated with that readout period. In a final step, we can interpolate this possibly irregular sampling of the image domain onto a desired regular pixel grid.

Ultimately, many different pMPI encoding schemes can use this simple x-space gridding reconstruction. This is a simpler regime than the case of canonical sinusoidal MPI, where necessary loss of the first harmonic signal first requires an overlapping trajectory in acquisition and then a method of recovering the lost information, canonically using an assumption of continuity and known zero signal boundary conditions at the edges of the FOV (Lu et al., 2013). Because of the ability to temporally decouple excitation feedthrough, these methods should not be required in multidimensional pMPI.

7.2.2 Simulation Results

Our simulator calculates the integrated magnetization changes associated with each periodic pMPI readout period but does not generate a fully sampled time-domain raw signal. In this context, the simulator can be applied to steady-state square wave imaging, relaxation weighted imaging based on equations provided in Chapter 6, and can simulate integrated gridding imaging using arbitrary periodic pulse sequences. The simulator was coded in a multi-threaded manner using the low-level Rust programming language to achieve high performance.

In Fig. 7.3 we show PSF simulation results for multidimensional pMPI. Fig. 7.3(a–c) show simulations for the case of a 27.4 nm tracer, 1 mT excitation, and a $7 \times 7\ \mathrm{T\,m^{-1}}$ 2D gradient. We show the standard x-space MPI 'collinear' (b) and 'transverse' (c) images, as derived by Goodwill and Conolly (Goodwill and Conolly, 2011). A collinear image is obtained from

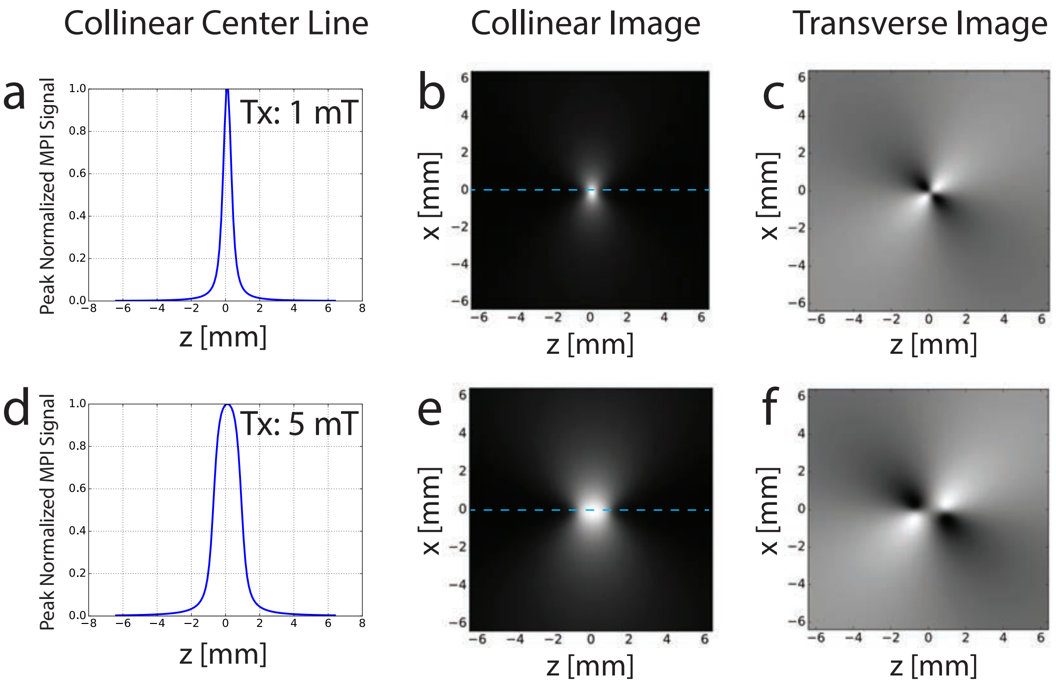

FIGURE 7.3: Pulsed MPI 2D PSF simulations. Here we simulate a 27.4 nm tracer as featured in Chapters 5 and 6 and a gradient strength of 7 $\mathrm{T\,m^{-1}}$. (a–c) Simulations of steady-state square wave pMPI with a 1 mT excitation amplitude. (a) 1D plot of the center line of the collinear PSF image. (b) The collinear image (PSF). (c) The transverse image (PSF). (d–f) Simulations with a 5 mT excitation amplitude. (d) 1D center line plot of the collinear PSF. (e) The collinear image (PSF). (f) The transverse image (PSF). As derived in the 1D case, we see the excitation amplitude leads to a rectangular function blur in the direction of excitation.

data from a receiver coil oriented along the same axis as the excitation coil, whereas the Tx and Rx coils are orthogonal in the case of a transverse image. Most x-space images reported in the literature are collinear images, in this jargon. These simulated data indicate that pMPI in multiple dimensions recapitulates the traditional theory but includes a rectangular blur, as derived in Chapter 5, that scales linearly with the amplitude of the excitation and is applied in the direction of the excitation. As in our 1D analysis, we must use small excitation amplitudes to achieve high native resolution. In terms of the transerve image, we also see an excitation amplitude-dependent blurring and separation of each pair of negative/positive lobes that constitute the transverse PSF structure.

Fig. 7.4 shows collinear simulations of a simple ‘Cal’ phantom at both 1 mT and 10 mT excitation amplitudes. These images are very similar

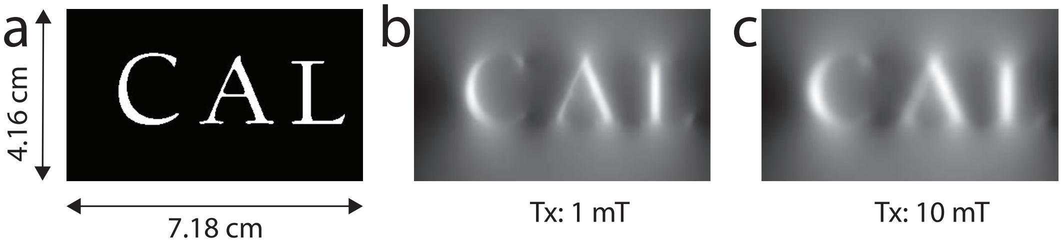

FIGURE 7.4: Cal logo phantom simulations. An approximately 7 x 4 cm phantom was simulated using a pMPI square wave excitation of 1 (b) and 10 (c) mT. The single collinear image in (b) is a higher resolution than we have been able to achieve in sinusoidal MPI thus far, while (c) is worse than the best canonical MPI results due to the large pMPI-specific amplitude blur caused by excitation with the larger 10 mT. Given our current reconstruction method, large excitations induce a secondar blur in the direction of the excitation as described in detail in Chapter 5.

to the canonical collinear image from sinusoidal excitation except for the excitation-direction blur. The 1 mT amplitude scan is higher resolution than we have ever achieved in canonical MPI since we have not been able to capitalize on the high resolution of a 27 nm tracer without suffering secondary relaxation-induced blur. Although this secondary-relaxation induced blur is completely removed when using square wave steady-state imaging as simulated here, use of the larger 10 mT excitation amplitude, as shown in Fig. 7.4(c), leads to the pMPI-specific amplitude blur that is worse than the relaxation blur of the better performing tracers in sinusoidal MPI. We must use low amplitudes to retain our resolution advantage in pMPI, at least with our current reconstruction scheme described in Chapters 5 and 6.

7.2.3 Initial Proof-of-Concept Data

My colleague, Zhi Wei Tay, has provided the first multidimensional pMPI data and image by modifying the AWR (Tay et al., 2016) to include small permanent magnets that provide a 3.5 $\mathrm{T\,m^{-1}}$ FFL over the sensitive volume of the AWR (the bottom of a standard PCR tube). With this hardware upgrade, we are able to acquire and reconstruct 2D projection images with the AWR and using an acquisition trajectory as depicted in Fig. 7.1.

Fig. 7.5 compares these first experimental results with a corresponding simulation using the simulator developed in this chapter. We used the 27.4 nm monodisperse tracer from Imagion Biosystems that was featured heavily in Chapters 5 and 6. Note that the entire phantom is approximately 4 x 3 mm in size as shown in Fig. 7.5(a). Our first experimental multidimensional

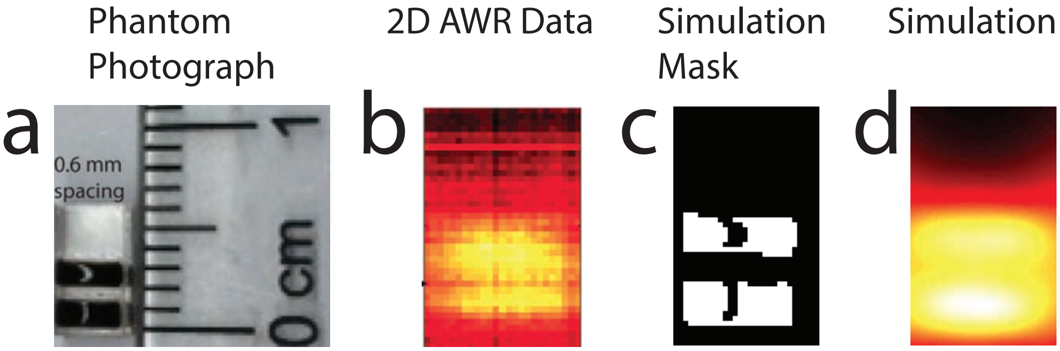

FIGURE 7.5: Comparison of proof of concept experimental 2D pMPI image and simulation. Our AWR, in 2D imaging mode provided by the addition of a 3.5 $\mathrm{T\,m^{-1}}$ permanent magnet system, was used to provide a multidimensional pMPI image (courtesy Zhi Wei Tay) (b) of a very small phantom (a). We also simulated the phantom using our pMPI simulator (c,d). Light relaxation weighting was used in reconstructing the experimental data (b) but was not modeled in the simulation (d) (case of no relaxation weighting). This explains the minor discrepancy between the two, in particular, the experimental image is more resolved and has reduced peripheral haze, as expected. Because the AWR imaging volume consists of the very small region at the base of a typical PCR tube, we had to use a particularly small phantom, and we note that the 3.5 $\mathrm{T\,m^{-1}}$ is relatively weak as far as our systems are typically constructed.

data is shown in Fig. 7.5(b) with corresponding simulation in Fig. 7.5(d). The experimental data was obtained using first harmonic-only relaxation weighted reconstruction as described in Chapter 6. For this reason, the experimental data has less haze and is better resolved than the simulation (which only models steady-state square wave reconstruction without relaxation weighting in this case). This is both a first proof-of-concept demonstration of multidimensional pMPI and of relaxation weighted approaches in multidimensional pMPI. Indeed, we see that the long tails/haze and a modest resolution improvement is obtained with these techniques in this 2D MPI data as predicted from our 1D analysis.

The extremely small volume of the AWR and the relatively weak gradient of this system mean that it is difficult to see significant features in these data, although we note that sinusoidal experiments with the same phantom (data not shown) are completely unresolved. We think this very first dataset is nonetheless important for demonstrating and proving the concept of multidimensional pMPI. Furthermore, the simulation lines up well with the experimental data, lending some degree of verification to our theory, simulations, and discussions in the rest of this chapter.

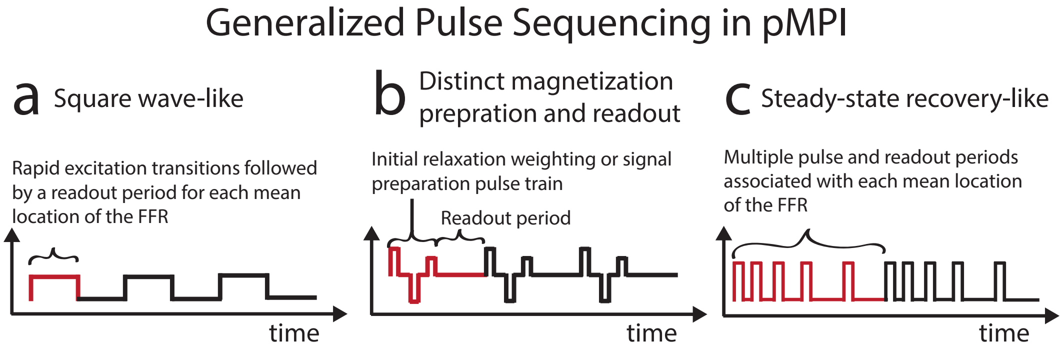

FIGURE 7.6: Generalized pulse sequencing in pMPI. (a) Square wave excitation. (b) A generic pulse sequence, whose fundamental period contains a distinct magnetization preparation/shaping train of pulses followed by a distinct readout component. (c) A steady-state recovery sequence as described in Chapter 6 contained within a repeated pMPI excitation period. Many types of pulse sequences can be composed using pulsed MPI primitives. We believe this flexibility may lead to a pulse sequence design space similar to that present in MRI.

7.3 Generalized Pulse Sequencing

One of the most intriguing aspects of pMPI is the variety of excitation waveforms available. Pulsed excitation encompasses much more than the simple square wave emphasized in Chapters 5 and 6. The steady-state recovery sequence that we described in Chapter 6 is an example of another distinct pMPI pulse sequence that we leveraged to show high sensitivity reporting of tracer microenvironmental state. We can contemplate many types of pulse sequences and we believe many are waiting to be discovered, perhaps with application- and context-specific utility as in the case of MRI pulse sequence exploitation of T1 and T2 contrast.

In addition to pMPI pulse sequencing, we can also consider other excitation methods that may improve pMPI performance, such as high amplitude excitation with projection reconstruction and excitation along the line when using an FFL scanner.

7.3.1 Magnetization Preparation and Readout Model

Fig. 7.6 depicts some of our ideas that fall under the umbrella of generalized pulse sequencing. Fig. 7.6(a) depicts the 'standard' square wave excitation that we have used extensively so far. Fig. 7.6(b) depicts a more complex, but periodic excitation pulse sequence. Here, in each fundamental period, there is an initial pulse train consisting of pulses with varying amplitudes

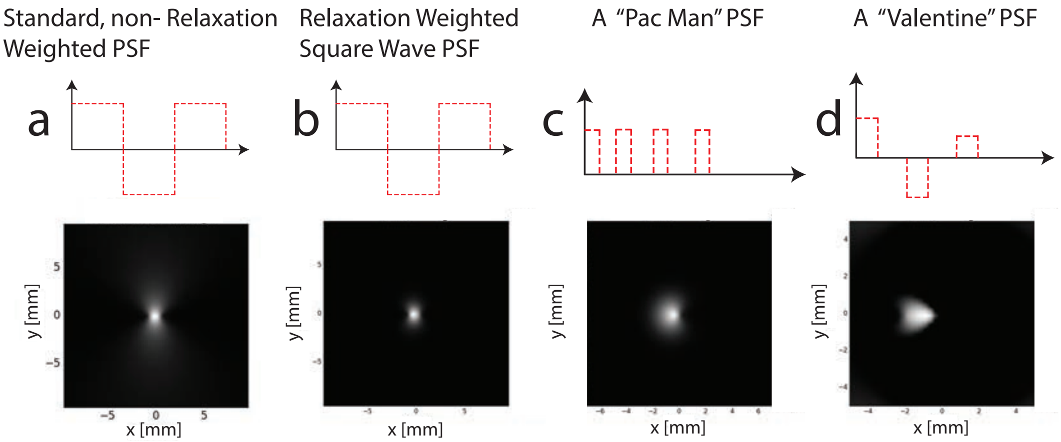
Standard, non- Relaxation Weighted PSF

Relaxation Weighted Square Wave PSF

A "Pac Man" PSF

A "Valentine" PSF

FIGURE 7.7: Simulations of relaxation weighting in generalized pulse sequencing. We can take advantage of relaxation weighting in pMPI reconstruction but also in pulse sequence encoding. In (a) we show a standard square wave PSF with no relaxation weighting in reconstruction while in (b) we show a relaxation weighted square wave PSF. Notice the removal of the long tails/haze in all directions due to selective removal of fast-relaxing signal. (c,d) Arbitrary examples of the type of generalized pulse sequences that are possible in pMPI using relaxation weighted encoding. (b) a 'Pac-Man'-like PSF and (b) a heart-shaped PSF (constructed by the author on Valentine's Day no less).

and polarities followed by a final long quiescent holding period. This is representative of a broad family of excitation sequences in which an initial 'magnetization preparation' (mag prep) or 'magnetization shaping' series of pulses excites the tracer distribution, perhaps in ways that encode desirable contrast through relaxation-weighting and careful timing of the pulsed components. This magnetization shaping component is then followed by a quiescent 'readout' period, in which we obtain our signal that has been set up in a specific manner. In this context, a square wave represents the simplest waveform in which the rapid transition between half-periods and the previous half-period readout is the magnetization preparation.

We can receive signal during a mag prep series as well, of course. But we think this framing is helpful for pulse sequence design – perhaps we can design sequences that selectively null tracer with a specific relaxation behavior, for example. Fig. 7.6(c) shows how we might employ our steady-state recovery sequence in an imaging format. We carefully choose a pulse train in terms of the interpulse periods to capture the steady-state recovery time associated with each sample point in our trajectory. In this case, our fundamental period contains a mix of 'pulse' and 'readout' sections and does not neatly fit in our sequential mag prep-then-readout model.

Fig. 7.7 shows simulated PSFs demonstrating the predicted effects of relaxation weighting in both reconstruction and encoding. Fig. 7.7(a) shows the 'standard' square wave excitation PSF that we have used extensively so far, while Fig. 7.7(b) depicts use of the same standard acquisition but with temporal relaxation weighting in reconstruction. Importantly, relaxation weighting greatly attenuates the long tails and haze in our PSF. This signal comes from when the FFR is not directly coincident with the tracer, which is thus exposed to a relatively high-field leading to fast relaxation. This fast relaxation signal is removed by our windowing based on secondary relaxation spatio-temporal encoding, as described in Chapter 6. This effect of relaxation weighting in multidimensional pMPI may obviate the need for techniques such as use of an equalization filter (Lu et al., 2015) in canonical x-space MPI.

Whereas Fig. 7.7(b) demonstrates exploitation of relaxation weighting in reconstruction, Fig. 7.6(c,d) demonstrate that we can also exploit relaxation weighting in our encoding/acquisition to produce unique PSFs. Fig. 7.7(c) shows a 'Pac-Man'-like PSF obtained from a pulse train with linearly increasing interpulse-periods. Fig. 7.7(d) shows a heart-shaped PSFs constructed with a pulse train containing pulses with both positive and negative polarities (with respect to the constant excitation direction). These arbitrary examples serve only to show the flexibility possible in pMPI pulse sequence design provided by relaxation weighting in the encoding phase of acquisition (*e.g.*, through mag prep periods).

7.3.2 Rethinking Other Aspects of Excitation and Reconstruction

We have significantly focused on the specific pulse sequences and encoding strategies available in pMPI, but these are not the only type of new methods available to us. For example, our integrated gridding reconstruction may be improved upon. While simple, it currently limits our SNR per excitation period per unit time, as discussed in Chapters 5 and 6, due to our need to use very low amplitudes to achieve high resolution.

Because temporal decoupling of excitation feedthrough and our received signal in pMPI frees us from some of the acquisition restrictions in canonical sinusoidal pMPI, we can also consider interesting approaches such as exciting along the line when using an FFL scanner. Such an approach may be able to remove the excitation amplitude barrier without requiring any changes to our simple reconstruction method.

Excitation Along an FFL

In our canonical FFL MPI systems, we always apply our excitation orthogonal to the axis of our FFL. This intuitively shifts the FFL in this orthogonal plane to sample (via projection) our FOV as depicted in Fig. 7.1. What if we consider applying our excitation waveforms (and a receiver coil) oriented along the axis of our FFL? What might this look like?

Excitation along the line will still generate signal that is projected along the dimension of the FFL long axis. However, since we are already projecting along that dimension, we can consider using much larger excitation amplitudes in the case of pMPI. We will recover the spatial information in this dimension using projection reconstruction which is well understood and standard with FFL scanners (Konkle et al., 2013). In terms of the initial production of projection images, our gridding is also simpler than the traditional case – simply grid the data to the mean location of the FFL, and since it is not translating in the orthogonal plane, this is fully determined by our shifting fields. This feature makes sinusoidal x-space MPI untenable with excitation along the line, as currently performed. This is because we require overlapping excitation trajectories which then enables us to recover lost first-harmonic information using a continuity algorithm (Lu et al., 2013). In pMPI, however, we can temporally separate excitation feedthrough and our tracer signal significantly and generally do not lose any tracer signal. This means we do not need overlapping trajectories or continuity algorithms in reconstruction, freeing us to consider implementing excitation along the line of an FFL.

One interesting aspect of excitation along the line of an FFL is that we can remove any zero-field region in the entire FOV. By applying a homogeneous field oriented along the line of the FFL, we effectively bias the FFL to some non-zero minimum field value but maintain the geometric structure of the FFL, which is important for maintaining spatial selectivity. One interesting ramification of this is that we can acquire our pMPI data in a context where the minimum field value is dictated by the excitation amplitude. For example, consider a square wave pMPI excitation along the line that switches between -5 and +5 mT. Our hold times/signal readout will occur in a context where the tracer along the line is exposed to -5 or +5 mT. One ramification is that the slowest/worse case relaxation time is now not the zero-field relaxation time constant for our tracer, but the relaxation time associated with 5 mT. An approach such as this could greatly increase the speed of signal acquisition and SNR efficiency since waiting for steady-state conditions during each signal readout can require significantly less time. We can, of course, also step between some bias field, *e.g.*, $10\ \mathrm{mT}$ and an unbiased FFL in our excitation, if low field signal readout is important for certain relaxation-weighting methods or relaxation contrast encoding. By choosing the bias field during readout, we may also be able to select the governing relaxation physics we desire in a scan (*e.g.*, the differences between a tracer's relaxation at $5\ \mathrm{mT}$ vs. $0\ \mathrm{mT}$).

Clearly, we need to rigorously explore the mathematics and other ramifications of excitation along the line in detail before making serious assertions about the technique. Perhaps, for example, vector effects will mix conventional transverse signal with conventional collinear signal in an untoward manner with this approach. Regardless, we believe this discussion implies there could be huge benefits to SNR and SNR efficiency for pMPI with such an approach and it is worth exploration. These benefits could be two-fold: from the ability to use much larger excitation amplitudes and the ability to scan much faster due to much faster achievement of steady-state during signal readout at a non-zero, biased field.

Other Reconstruction Approaches

The reason we must use very low amplitudes in our current square wave pMPI method is that our reconstruction scheme treats such an excitation in a projective manner. That is, by integrating the entire signal associated with a step excitation in pMPI, we are locally projecting all spatial information over the excursion distance of the FFR before and after a step. Mathematically, this manifests as our excitation amplitude-dependent rectangular function in our square wave pMPI PSF. In contrast, in canonical MPI, we grid the data associated with each time point to its own unique instantaneous position based on knowledge of the smoothly and continuously varying FFR location.

By using small excitation amplitudes in pMPI, we can locally project across a small distance with respect to, for example, the Langevin FWHM we expect from our tracer. This is an implementation-specific and self-imposed limitation. Surely we can devise reconstruction methods that are not so limiting in the future.

We can first consider not independently gridding the integrated signal associated with each excitation half-period. The mathematics described in Chapter 5 indicate that each successive half-period signal samples two steady-state magnetization values where the final magnetization state of each previous square wave half-period is equal to the initial magnetization state of the next square wave half-period. We can use this knowledge to avoid the rectangular blur by solving for the exact magnetization states (and not a difference between pairs). This can be performed by, for example, setting up the reconstruction problem as a series of coupled difference equations or inverting an appropriate Toeplitz matrix. Unfortunately, these methods as naïvely formulated are ill-posed. But perhaps clever sequencing and alternative methods in the future can lead to well-posed reconstruction that avoids the rectangular function blur.

Another possibility is to use as large excitation amplitudes as possible (*e.g.*, limited only by hardware constraints and magnetostimulation / nonspecific SAR) and reconstruct using projection reconstruction in the excitation direction. This would require an ability to rotate the relative geometry between the excitation coil(s) and sample arbitrarily and would be an orthogonal dimension in an FFL system. Local projection data could be taken associated with a partial FOV (pFOV) of the image corresponding to the excursion size associated with each step excitation. A set of such data, after acquiring data at many different angles (excitation directions) could be reconstructed with projection reconstruction algorithms to reconstruct a single 2D pFOV of the image. This process could be repeated for as many such pFOVs are needed to completely cover the desired full FOV, with the pFOVS stitched or averaged together to form a final image. Alternatively, large amplitude excitation could be applied as we slow shift across the entire FOV in the excitation dimension. These data can then be combined and integrated to form a single projection data point along the excitation axis. This can then be repeated for each line required to cover the full FOV in the orthogonal dimension, and for each rotation of the excitation axis needed to provide a well-posed single projection reconstruction.

While potentially inefficient and difficult to implement, incompletely described, or even ill-posed, we discuss these exotic methods to highlight possible avenues of exploration in future pMPI development. We do need to solve the problem of limited excitation amplitudes and more generally, explore the incredible new flexibility provided by pMPI. Undoubtedly, we have only scratched the surface and much more efficient and compelling implementations await discovery and exposition.

7.4 Conclusions

In this chapter we have described multidimensional pMPI, in contrast to the 1D focus of Chapters 5 and 6. We have described intuitively how multidimensional pMPI is very similar to canonical multidimensional MPI, provided many simulations illustrating the process, and compared the simulator with proof-of-concept 2D data. In this analysis, we have shown that the major features of pMPI described in Chapters 5 and 6 hold in a familiar way in the multidimensional context. For example, our PSF is the steady-state Langevin PSF as traditionally described, but blurred in an excitation amplitude-dependent manner in the direction of excitation. Additionally, our simulations and data showing relaxation weighting in multidimensional pMPI indicate that it is a very powerful effect, both improving resolution and significantly attenuating the long tails and haze associated with the canonical multidimensional Langevin PSF.

We also described the great flexibility inherent to use of arbitrary pulsed excitation waveforms in discussion of generalized pulse sequencing. We showed how we can shape and tailor our imaging PSF by changing the pulse sequence and leveraging relaxation weighting during encoding and described a potentially useful conceptual framework that breaks down a pMPI pulse sequence into magnetization preparation and signal readout periods. Finally, we describe untested and unimplemented possible approaches in pMPI that may generally improve the method but, in particular, provide a way of avoiding the low amplitude limitations of our current methods when high resolution is desired. Perhaps most intriguing is implementing pMPI excitation oriented along the axis of the FFL in an FFL MPI scanner.

We hope this chapter suggests the feasibility, flexibility, and opportunity inherent in pMPI. But the technique is very new and unexplored. We hope future work will continue to develop and mature the pMPI approach, possibly greatly increasing what is possible in MPI.

7.5 Acknowledgements

I thank Zhi Wei Tay for providing the data showing an experimental demonstration of multidimensional pMPI that was featured in Fig. 7.5(a,b) of this chapter and allowed for a comparison between data and my simulator.

Bibliography

Ahrens, Eric T and Jeff WM Bulte (2013). “Tracking immune cells in vivo using magnetic resonance imaging”. In: *Nature Reviews Immunology* 13.10, pp. 755–763.

Ahrens, Eric T et al. (2005). “In vivo imaging platform for tracking immunotherapeutic cells”. In: *Nature biotechnology* 23.8, pp. 983–987.

Arami, Hamed et al. (2013). “Size-dependent ferrohydrodynamic relaxometry of magnetic particle imaging tracers in different environments”. In: *Medical physics* 40.7.

Bauer, Lisa M et al. (2015). “Magnetic particle imaging tracers: state-of-the-art and future directions”. In: *J. Phys. Chem. Lett* 6.13, pp. 2509–2517.

— (2016). “High-performance iron oxide nanoparticles for magnetic particle imaging–guided hyperthermia (hMPI)”. In: *Nanoscale*.

Bean, CP and JD Livingston (1959). “Superparamagnetism”. In: *Journal of Applied Physics* 30.4, S120–S129.

Behrends, Andre, Thorsten M. Buzug, and Alexander Neumann (2016). “Magnetic Particle Spectrometer for the Analysis of Magnetic Particle Heating Applications”. In: *International Workshop on Magnetic Particle Imaging*, p. 47.

Beyer, Thomas et al. (2000). “A combined PET/CT scanner for clinical oncology”. In: *The Journal of nuclear medicine* 41.8, p. 1369.

Biederer, S et al. (2009). “Magnetization response spectroscopy of superparamagnetic nanoparticles for magnetic particle imaging”. In: *Journal of Physics D: Applied Physics* 42.20, p. 205007.

Bloch, Felix (1946). “Nuclear induction”. In: *Physical review* 70.7-8, p. 460.

Branquinho, Luis C et al. (2013). “Effect of magnetic dipolar interactions on nanoparticle heating efficiency: Implications for cancer hyperthermia”. In: *Scientific reports* 3, p. 2887.

Brown Jr, William Fuller (1963). “Thermal fluctuations of a single-domain particle”. In: *Physical Review* 130.5, p. 1677.

Cho, E, G Cho, H Cho, et al. (2014). “Simulational validation of color magnetic particle imaging (cMPI)”. In: *Physics in medicine and biology* 59.21, p. 6521.

Creixell, Mar et al. (2011). “EGFR-targeted magnetic nanoparticle heaters kill cancer cells without a perceptible temperature rise”. In: *Acs Nano* 5.9, pp. 7124–7129.

Croft, Laura R et al. (2016). “Low drive field amplitude for improved image resolution in magnetic particle imaging”. In: *Medical physics* 43.1, pp. 424–435.

Debye, Peter Josef William et al. (1954). “collected papers of Peter JW Debye”. In:

Deissler, Robert J, Yong Wu, and Michael A Martens (2014). “Dependence of Brownian and Néel relaxation times on magnetic field strength”. In: *Medical physics* 41.1.

Dewey, William C (2009). “Arrhenius relationships from the molecule and cell to the clinic”. In: *International journal of hyperthermia* 25.1, pp. 3–20.

Dewhirst, MW et al. (2003). “Basic principles of thermal dosimetry and thermal thresholds for tissue damage from hyperthermia”. In: *International Journal of Hyperthermia* 19.3, pp. 267–294.

Dhavalikar, R et al. (2016). “Finite magnetic relaxation in x-space magnetic particle imaging: comparison of measurements and ferrohydrodynamic models”. In: *Journal of Physics D: Applied Physics* 49.30, p. 305002.

Dhavalikar, Rohan and Carlos Rinaldi (2016). “Theoretical predictions for spatially-focused heating of magnetic nanoparticles guided by magnetic particle imaging field gradients”. In: *Journal of Magnetism and Magnetic Materials* 419, pp. 267–273.

Domenech, Maribella et al. (2013). “Lysosomal membrane permeabilization by targeted magnetic nanoparticles in alternating magnetic fields”. In: *ACS nano* 7.6, pp. 5091–5101.

Eberbeck, Dietmar et al. (2013). “Multicore magnetic nanoparticles for magnetic particle imaging”. In: *IEEE Transactions on Magnetics* 49.1, pp. 269–274.

Etzioni, Ruth et al. (2003). “The case for early detection”. In: *Nature Reviews Cancer* 3.4, pp. 243–252.

Ferguson, R Matthew, Amit P Khandhar, and Kannan M Krishnan (2012). “Tracer design for magnetic particle imaging”. In: *Journal of applied physics* 111.7, 07B318.

Ferguson, R Matthew, Kevin R Minard, and Kannan M Krishnan (2009). “Optimization of nanoparticle core size for magnetic particle imaging”. In: *Journal of magnetism and magnetic materials* 321.10, pp. 1548–1551.

Ferguson, R Matthew et al. (2015). “Magnetic particle imaging with tailored iron oxide nanoparticle tracers”. In: *IEEE Trans. Med. Imaging* 34.5, pp. 1077–1084.

Gaskill, Jack D (1978). “Linear systems, Fourier transforms, and optics”. In: *Linear Systems, Fourier Transforms, and Optics by Jack D. Gaskill New York, NY: John Wiley and Sons, 1978* 1.

Gleich, B, J Weizenecker, and J Borgert (2008). “Experimental results on fast 2D-encoded magnetic particle imaging”. In: *Physics in medicine and biology* 53.6, N81.

Gleich, Bernhard and Jurgen Weizenecker (2005). “Tomographic imaging using the nonlinear response of magnetic particles”. In: *Nature* 435, pp. 1214–1217. ISSN: 7046. DOI: `10.1038/nature03808`. URL: `http://dx.doi.org/10.1038/nature03808`.

Goodwill, Patrick W and Steven M Conolly (2010). “The X-space formulation of the magnetic particle imaging process: 1-D signal, resolution, bandwidth, SNR, SAR, and magnetostimulation”. In: *IEEE transactions on medical imaging* 29.11, pp. 1851–1859.

— (2011). “Multidimensional x-space magnetic particle imaging”. In: *IEEE transactions on medical imaging* 30.9, pp. 1581–1590.

Goodwill, Patrick W et al. (2009). “Narrowband magnetic particle imaging”. In: *IEEE transactions on medical imaging* 28.8, pp. 1231–1237.

Goodwill, Patrick W et al. (2012a). “An x-space magnetic particle imaging scanner”. In: *Review of Scientific Instruments* 83.3, p. 033708.

Goodwill, Patrick W et al. (2012b). “Projection x-space magnetic particle imaging”. In: *IEEE transactions on medical imaging* 31.5, pp. 1076–1085.

Goodwill, PW et al. (2011). “Ferrohydrodynamic relaxometry for magnetic particle imaging”. In: *Applied Physics Letters* 98.26, p. 262502.

Hensley, Daniel et al. (2015a). “Model predictive control for treating cancer with ultrasonic heating”. In: *American Control Conference (ACC), 2015*. IEEE, pp. 220–225.

Hensley, Daniel et al. (2015b). “Preliminary experimental X-space color MPI”. In: *Magnetic Particle Imaging (IWMPI), 2015 5th International Workshop on*. IEEE, pp. 1–1.

Hensley, Daniel et al. (2017). “Combining magnetic particle imaging and magnetic fluid hyperthermia in a theranostic platform”. In: *Physics in medicine and biology* 62.9, p. 3483.

Hilger, Ingrid et al. (2002). “Thermal ablation of tumors using magnetic nanoparticles: an in vivo feasibility study”. In: *Investigative radiology* 37.10, pp. 580–586.

Hoult, David I and RE Richards (1976). “The signal-to-noise ratio of the nuclear magnetic resonance experiment”. In: *Journal of Magnetic Resonance (1969)* 24.1, pp. 71–85.

James, Michelle L. and Sanjiv S. Gambhir (2012). "A Molecular Imaging Primer: Modalities, Imaging Agents, and Applications". In: *Physiological Reviews* 92.2, pp. 897–965. ISSN: 0031-9333. DOI: `10.1152/physrev.00049.2010`. eprint: `http://physrev.physiology.org/content/92/2/897.full.pdf`. URL: `http://physrev.physiology.org/content/92/2/897`.

Johannsen, M et al. (2005). "Clinical hyperthermia of prostate cancer using magnetic nanoparticles: presentation of a new interstitial technique". In: *International journal of hyperthermia* 21.7, pp. 637–647.

Johannsen, Manfred et al. (2010). "Magnetic nanoparticle hyperthermia for prostate cancer". In: *International Journal of Hyperthermia* 26.8, pp. 790–795.

Jordan, Andreas et al. (1999). "Magnetic fluid hyperthermia (MFH): Cancer treatment with AC magnetic field induced excitation of biocompatible superparamagnetic nanoparticles". In: *Journal of Magnetism and Magnetic Materials* 201.1, pp. 413–419.

Jordan, Andreas et al. (2006). "The effect of thermotherapy using magnetic nanoparticles on rat malignant glioma". In: *Journal of neuro-oncology* 78.1, pp. 7–14.

Jordan, Andreas et al. (2009). "Inductive heating of ferrimagnetic particles and magnetic fluids: physical evaluation of their potential for hyperthermia". In: *International Journal of Hyperthermia* 25.7, pp. 499–511.

Khandhar, Amit P et al. (2012). "Tailored magnetic nanoparticles for optimizing magnetic fluid hyperthermia". In: *Journal of Biomedical Materials Research Part A* 100.3, pp. 728–737.

Knopp, T and A Weber (2015). "Local system matrix compression for efficient reconstruction in magnetic particle imaging". In: *Advances in Mathematical Physics* 2015.

Konkle, Justin J et al. (2013). "Projection reconstruction magnetic particle imaging". In: *IEEE transactions on medical imaging* 32.2, pp. 338–347.

Konkle, Justin J et al. (2015). "A convex formulation for magnetic particle imaging x-space reconstruction". In: *PloS one* 10.10, e0140137.

Kuboyabu, Tomomi et al. (2016). "Usefulness of Magnetic Particle Imaging for Monitoring the Effect of Magnetic Targeting". In: *Open Journal of Medical Imaging* 6.02, p. 33.

Kumar, Challa SSR and Faruq Mohammad (2011). "Magnetic nanomaterials for hyperthermia-based therapy and controlled drug delivery". In: *Advanced drug delivery reviews* 63.9, pp. 789–808.

Laurent, Sophie et al. (2011). “Magnetic fluid hyperthermia: focus on superparamagnetic iron oxide nanoparticles”. In: *Advances in colloid and interface science* 166.1, pp. 8–23.

Li, Chun (2014). “A targeted approach to cancer imaging and therapy”. In: *Nature materials* 13.2, pp. 110–115.

Lu, Kuan et al. (2013). “Linearity and shift invariance for quantitative magnetic particle imaging”. In: *IEEE transactions on medical imaging* 32.9, pp. 1565–1575.

Lu, Kuan et al. (2015). “Reshaping the 2D MPI PSF to be isotropic and sharp using vector acquisition and equalization”. In: *Magnetic Particle Imaging (IWMPI), 2015 5th International Workshop on*. IEEE, pp. 1–1.

Lu, Min et al. (2010). “FDA report: ferumoxytol for intravenous iron therapy in adult patients with chronic kidney disease”. In: *American journal of hematology* 85.5, pp. 315–319.

Lustig, Michael, David Donoho, and John M Pauly (2007). “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic resonance in medicine* 58.6, pp. 1182–1195.

Maenosono, Shinya and Soichiro Saita (2006). “Theoretical assessment of FePt nanoparticles as heating elements for magnetic hyperthermia”. In: *IEEE transactions on magnetics* 42.6, pp. 1638–1642.

Maier-Hauff, Klaus et al. (2011). “Efficacy and safety of intratumoral thermotherapy using magnetic iron-oxide nanoparticles combined with external beam radiotherapy on patients with recurrent glioblastoma multiforme”. In: *Journal of neuro-oncology* 103.2, pp. 317–324.

Martsenyuk, MA, Yu L Raikher, and MI Shliomis (1974). “On the kinetics of magnetization of ferromagnetic particle suspensions”. In: *Sov. Phys. JETP* 38, p. 413.

Maruyama, S et al. (2016). “Development of Magnetic Nanocarriers based on Thermosensitive Liposomes and their Visualization using Magnetic Particle Imaging”. In: *Int J Nanomed Nanosurg* 2.2.

Murase, Kenya et al. (2013). “Control of the temperature rise in magnetic hyperthermia with use of an external static magnetic field”. In: *Physica Medica* 29.6, pp. 624–630.

Murase, Kenya et al. (2015). “Usefulness of Magnetic Particle Imaging for Predicting the Therapeutic Effect of Magnetic Hyperthermia”. In: *Open Journal of Medical Imaging* 5.02, p. 85.

Narasimha, M, MS Brennan, and PN Holtham (2006). “Numerical simulation of magnetite segregation in a dense medium cyclone”. In: *Minerals Engineering* 19.10, pp. 1034–1047.

Neel, L (1949). "Theorie du trainage magnetique des ferromagnetiques en grains fins". In:

Orendorff, Ryan et al. (2017). "First in vivo traumatic brain injury imaging via magnetic particle imaging". In: *Physics in medicine and biology* 62.9, p. 3501.

Pablico-Lansigan, Michele H, Shu F Situ, and Anna Cristina S Samia (2013). "Magnetic particle imaging: advancements and perspectives for real-time in vivo monitoring and image-guided therapy". In: *Nanoscale* 5.10, pp. 4040–4055.

Park, Hae-Jeong and Karl Friston (2013). "Structural and functional brain networks: from connections to cognition". In: *Science* 342.6158, p. 1238411.

Perreard, IM et al. (2014). "Temperature of the magnetic nanoparticle microenvironment: estimation from relaxation times". In: *Physics in medicine and biology* 59.5, p. 1109.

Prince, Jerry L and Jonathan M Links (2006). *Medical imaging signals and systems*. Pearson Prentice Hall Upper Saddle River, New Jersey.

Rahmer, J et al. (2015). "First experimental evidence of the feasibility of multi-color magnetic particle imaging". In: *Physics in medicine and biology* 60.5, p. 1775.

Rahmer, Jürgen et al. (2009). "Signal encoding in magnetic particle imaging: properties of the system function". In: *BMC medical imaging* 9.1, p. 1.

— (2012). "Analysis of a 3-D system function measured for magnetic particle imaging". In: *IEEE transactions on medical imaging* 31.6, pp. 1289–1299.

Raible, Martin and Andreas Engel (2004). "Langevin equation for the rotation of a magnetic particle". In: *Applied organometallic chemistry* 18.10, pp. 536–541.

Rauwerdink, Adam M and John B Weaver (2010). "Viscous effects on nanoparticle magnetization harmonics". In: *Journal of Magnetism and Magnetic Materials* 322.6, pp. 609–613.

Rosensweig, Ronald E (2002). "Heating magnetic fluid with alternating magnetic field". In: *Journal of magnetism and magnetic materials* 252, pp. 370–374.

— (2013). *Ferrohydrodynamics*. Courier Corporation.

Sapareto, Stephen A and William C Dewey (1984). "Thermal dose determination in cancer therapy". In: *International Journal of Radiation Oncology* Biology* Physics* 10.6, pp. 787–800.

Scott, Greig et al. (2001). "A prepolarized MRI scanner". In: *Proc. Int. Soc. Magn. Reson. Med.(Glasgow, Scotland)* 9.

Serres, Sébastien et al. (2012). "Molecular MRI enables early and sensitive detection of brain metastases". In: *Proceedings of the National Academy of Sciences* 109.17, pp. 6674–6679.

Shahram, Morteza and Peyman Milanfar (2004). "Imaging below the diffraction limit: A statistical analysis". In: *IEEE Transactions on image processing* 13.5, pp. 677–689.

Shliomis, Mark I (2002). "Ferrohydrodynamics: Retrospective and issues". In: *Ferrofluids*. Springer, pp. 85–111.

Silva, Alvin C et al. (2010). "Innovations in CT dose reduction strategy: application of the adaptive statistical iterative reconstruction algorithm". In: *American Journal of Roentgenology* 194.1, pp. 191–199.

Spaldin, Nicola A (2010). *Magnetic materials: fundamentals and applications*. Cambridge University Press.

Stehning, Christian, Bernhard Gleich, and Jürgen Rahmer (2016). "Simultaneous magnetic particle imaging (MPI) and temperature mapping using multi-color MPI". In: *International Journal on Magnetic Particle Imaging* 2.2.

Tasci, T Onur et al. (2009). "Focused RF hyperthermia using magnetic fluids". In: *Medical physics* 36.5, pp. 1906–1912.

Tay, Zhi Wei et al. (2016). "A High-Throughput, Arbitrary-Waveform, MPI Spectrometer and Relaxometer for Comprehensive Magnetic Particle Optimization and Characterization." In: *Scientific reports* 6.

Tay, Zhi Wei et al. (2017). "The relaxation wall: experimental limits to improving MPI spatial resolution by increasing nanoparticle core size". In: *Biomedical Physics & Engineering Express*.

Thiesen, Burghard and Andreas Jordan (2008). "Clinical applications of magnetic nanoparticles for hyperthermia". In: *International Journal of Hyperthermia* 24.6, pp. 467–474. DOI: `10.1080/02656730802104757`. eprint: `http://dx.doi.org/10.1080/02656730802104757`. URL: `http://dx.doi.org/10.1080/02656730802104757`.

Torres-Lugo, Madeline and Carlos Rinaldi (2013). "Thermal potentiation of chemotherapy by magnetic nanoparticles". In: *Nanomedicine* 8.10, pp. 1689–1707.

Vogel, Patrick et al. (2014). "Traveling wave magnetic particle imaging". In: *IEEE transactions on medical imaging* 33.2, pp. 400–407.

Weaver, John B, Adam M Rauwerdink, and Eric W Hansen (2009). "Magnetic nanoparticle temperature estimation". In: *Medical physics* 36.5, pp. 1822–1829.

Weissleder, Ralph, Matthias Nahrendorf, and Mikael J Pittet (2014). “Imaging macrophages with nanoparticles”. In: *Nature materials* 13.2, pp. 125–138.

Weizenecker, J et al. (2009). “Three-dimensional real-time in vivo magnetic particle imaging”. In: *Physics in medicine and biology* 54.5, p. L1.

Wilhelm, Stefan et al. (2016). “Analysis of nanoparticle delivery to tumours”. en. In: *Nature Reviews Materials* 1, p. 16014.

Yu, Elaine et al. (2016). “First demonstration of in vivo Cancer Magnetic Particle Imaging with IV-administered Passive Long-circulating SPIOs”. In: *2016 World Molecular Imaging Congress (WMIC 2016): Imaging Biology... Improving Therapy*. World Molecular Imaging Society (WMIS).

Zhang, J and RDK Misra (2007). “Magnetic drug-targeting carrier encapsulated with thermosensitive smart polymer: core–shell nanoparticle carrier and drug release response”. In: *Acta Biomaterialia* 3.6, pp. 838–850.

Zhang, Wencong et al. (2015). “A custom low-noise preamplifier for magnetic particle imaging”. In: *Magnetic particle imaging (IWMPI), 2015 5th international workshop on*. IEEE, pp. 1–1.

Zheng, Bo et al. (2015). “Magnetic Particle Imaging tracks the long-term fate of in vivo neural cell implants with high image contrast”. In: *Scientific reports* 5.

Zheng, Bo et al. (2016). “Quantitative Magnetic Particle Imaging Monitors the Transplantation, Biodistribution, and Clearance of Stem Cells In Vivo”. In: *Theranostics* 6, pp. 291–301. ISSN: 3. DOI: `10.7150/thno.13728`. URL: `http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4737718/`.

Zhou, Xinyi et al. (2016). “In Vivo Magnetic Particle Imaging of Lung Perfusion in Rats - Towards High Sensitivity and High Contrast Non-Radiation Based Diagnostics for Pulmonary Embolism”. In: *2016 World Molecular Imaging Congress (WMIC 2016): Imaging Biology... Improving Therapy*. World Molecular Imaging Society (WMIS).

Zhou, Xinyi Y et al. (2017). “First in vivo magnetic particle imaging of lung perfusion in rats”. In: *Physics in Medicine and Biology* 62.9, p. 3510.