

Temporal Regularization Use in Dynamic Contrast-Enhanced MRI

by

Kimberly A. Khalsa

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biomedical Engineering)
in The University of Michigan
2011

Doctoral Committee:

Professor Jeffrey A. Fessler, Chair
Professor Thomas L. Chenevert
Professor Douglas C. Noll
Associate Professor Anna C. Gilbert

© Kimberly A. Khalsa 2011

All Rights Reserved

To my mama. You are the strongest woman I know.

ACKNOWLEDGEMENTS

I am incredibly grateful to so many people for all of the love, support, and guidance I have received over the years. Of course, I'd like to thank my advisor, Jeff Fessler, who, for better or for worse, convinced me to go to grad school in the first place. He inspired me as an undergrad, and he inspires me still today. I think it's safe to say that this dissertation wouldn't exist without Jeff's endless guidance, patience, and support. Thank you.

I also want to thank the other members of my committee: Doug Noll, who first taught me the nuts and bolts of MRI, and whose group meetings and group members were an important part of my graduate education; Tom Chenevert, whose insights and knowledge regarding "real life" DCE-MRI were invaluable to this research project; and Anna Gilbert, for her support and insights. A big thank you also goes to my funding sources: the National Science Foundation Graduate Research Fellowship Program, the Department of Defense Breast Cancer Research Predoctoral Traineeship Program, and the Rollin M. Gerstacker Foundation.

I could not have survived grad school without the endless love and support of my amazing family and friends. To my mom, Mary St. Clair, you are truly an inspiration. From making soup when I'm sick, to being my favorite shopping companion, to providing hugs and encouragement and chocolate cake (when necessary), you are an extraordinary mother and friend. Thank you for everything. To my papa, Gurudarshan Khalsa, thank you for always believing in me (often more than I believed in myself), and for regularly reminding me to "Party on, Wayne". Your perspective on life is always helpful and appreciated. Thank you to my unci, Gurucharn Khalsa, for visiting often, taking me to dinner and the movies, and for still spoiling me in many small ways, just as you always did as I was growing up.

Last, but certainly not least, a HUGE thank you to my incredible friends, without whom I never would have made it through the last N+ years. Of the many things I am taking away from grad school, including the degree itself, your friendships are what I treasure most. There are so many people who have helped make the non-academic parts of my grad school years great, a few of whom I want to acknowledge personally.

Mike Diaz, your steadfast support, especially during the final few months, has been incredible. You have been there for me when I needed you most, always ready to listen, to provide an encouraging word, or simply to make me laugh. Emine Cagin, Awlok Josan, Danese Joiner, Kiran Pandey, and Tzeno Galchev- you guys have been there since the beginning, and have seen me through the end. From coffee breaks to football games to world travels, we've managed to squeeze some fun out of these years. Each of you has added so much to my life- thank you for being there with me. I couldn't have done it without you.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	viii
LIST OF TABLES	xii
ABSTRACT	xiii
CHAPTER	
I. Overview	1
1.1 Motivation: Application to Breast Cancer	1
1.2 Tradeoff between Spatial and Temporal Resolution	2
1.3 Contributions of this Thesis	5
1.4 Outline of the Thesis	6
II. Introduction to Dynamic Contrast-Enhanced MRI	7
2.1 MR Basics	7
2.1.1 Nuclear Magnetic Resonance	7
2.1.2 K-space	8
2.1.3 Signal Equation	10
2.2 Dynamic Contrast-Enhanced MRI	11
2.2.1 Contrast Agents	11
2.2.2 DCE-MRI of the Breast	12
2.2.3 Limitations in Dynamic MRI	14
2.2.4 Parallel Imaging	15
2.3 Existing Reconstruction Methods	16
III. TRUIR: Temporal Regularization Use In Reconstruction	19
3.1 Measurement Model	19

3.1.1	Static MRI	19
3.1.2	Dynamic MRI	21
3.2	TRUIR Cost Function	22
3.3	Preliminary Simulation Results	24
3.4	Extension to Parallel Imaging	27
3.5	Algorithm Acceleration	29
IV.	Phase Encode Sampling Schemes	31
4.1	Original Trajectory	31
4.2	Reordered Trajectory 1	34
4.3	Reordered Trajectory 2	39
V.	Regularization Parameters	42
5.1	Initial DCE-MRI Simulation	42
5.2	Local Impulse Response	45
5.2.1	Accelerated Computation for the Single Coil Case	46
5.2.2	Single Coil Results	47
5.3	Multi-coil Case	50
5.3.1	Spatial PSF vs β	50
5.3.2	Spatial PSF vs α	52
5.3.3	Temporal PSF	64
VI.	Kinetic Parameters in DCE-MRI	66
6.1	Modeling Tissue Physiology	66
6.1.1	Definition of Kinetic Parameters	67
6.2	Estimating Kinetic Parameters from an Image Sequence	71
6.2.1	Computing Contrast Agent Concentration	71
6.2.2	Concentration to Kinetic Parameter Estimates	72
6.3	Variable Projection	72
6.3.1	Using VP to Estimate Kinetic Parameters	75
VII.	Simulation Study and Results	77
7.1	Experiment Setup	77
7.1.1	Modeling the Dynamic Object	77
7.1.2	Coil sensitivity	80
7.1.3	Acquired PE Trajectories	81
7.1.4	Reconstruction Methods	82
7.1.5	Assigning Frame Times	84
7.1.6	Regularization Parameters	85
7.2	Results	86
7.2.1	Small Rapidly Enhancing Lesion	87

7.2.2	General Results	102
7.2.3	Kinetic Parameters	106
7.2.4	Summary	110
VIII. Conclusions and Future Work		113
8.1	Summary	113
8.2	Future Work	114
BIBLIOGRAPHY		117

LIST OF FIGURES

<u>Figure</u>		
1.1	Night Photography	3
2.1	Representation of an object in object space and in k-space	11
2.2	Kuhl et al.’s time intensity curves from [1]	13
2.3	Traditional Keyhole reconstruction.	17
3.1	Illustration of the effect of TRUIR’s spatial and temporal regularization	23
3.2	Enhancement curve used in simulations	24
3.3	K-space trajectory used in simulations.	25
3.4	Reconstruction error, entire object.	26
3.5	Reconstruction error in ROI surrounding lesion.	26
3.6	5th frame of reconstructed image sequence	27
3.7	The measurement model for parallel imaging includes coil sensitivity maps	28
4.1	Clinical PE sample locations, and full sampling	32
4.2	Original Trajectory	33
4.3	PE locations sampled in the first and second half of a “full” frame acquisition	35
4.4	Ring-based reordering of PE locations	36
4.5	Reordered Trajectory 1	37

4.6	Radial distribution of sample locations over time	38
4.7	Creation of Reordered Trajectory 2	39
4.8	Reordered Trajectory 2	40
5.1	Enhancement curve for simulated lesion and reconstructions	43
5.2	True image and various reconstructions	44
5.3	FWHM _s vs β for single-coil TRUIR	48
5.4	FWHM _t vs α for single-coil TRUIR	49
5.5	Spatial PSF measures vs β for multi-coil TRUIR	51
5.6	Spatial PSF measures vs α , for multi-coil TRUIR	54
5.7	PSF for 3 α s.	55
5.8	PSF features as a function of α for the Original Trajectory with 24- frame TRUIR reconstruction	58
5.9	PSF features as a function of α for Reordered Trajectory 1 with 24- frame TRUIR reconstruction	60
5.10	PSF features as a function of α for Reordered Trajectory 2 with 24- frame TRUIR reconstruction	61
5.11	PSF tail energy vs α	63
6.1	The two-compartment model relates the contrast agent concentration in the blood plasma, $C_p(t)$, to the tissue concentration, $\mathbf{C}_t(t)$	68
6.2	Modeled contrast agent concentration in the aorta ($C_p^{Aorta}(t)$), blood plasma ($C_p(t)$), and tissue of interest ($C_t(t)$).	69
6.3	Transport function, $h_p(t)$	69
6.4	Tissue transfer function, $h_t(t)$	70
7.1	True dynamic object used in our simulations, shown at a single point in time (1 minute).	78

7.2	True enhancement curves for the slow, moderate, and rapidly enhancing lesions in our simulated object.	79
7.3	Sensitivity maps used in simulations.	81
7.4	Three phase encode trajectories used in our simulations.	82
7.5	Flow Chart of the Homodyne+SENSE reconstruction method . . .	83
7.6	Example of how TRUIR and HS group the available data when reconstructing at an increased frame rate.	85
7.7	Frames from 3 reconstructed dynamic image sequences	87
7.8	Effect of α on enhancement curves of the small, rapid lesion for different PE orderings.	88
7.9	Enhancement curves of HS reconstructions of the small, rapidly enhancing lesion	90
7.10	Enhancement curves of TRUIR reconstructions of the small, rapidly enhancing lesion	91
7.11	HS reconstructions of the small, rapidly enhancing lesion and the normalized absolute error of each reconstruction	93
7.12	TRUIR reconstructions of the small, rapidly enhancing lesion and the normalized absolute error of each reconstruction	94
7.13	Profiles through HS reconstructions of the small, rapidly enhancing lesion shown in Fig. 7.11	96
7.14	Profiles through TRUIR reconstructions of the small, rapidly enhancing lesion shown in Fig. 7.12	97
7.15	SENSE-only, HS, and TRUIR 12-frame reconstructions of the small, rapidly enhancing lesion	98
7.16	Profiles through SENSE, HS, and TRUIR reconstructions of the small, rapidly enhancing lesion shown in Fig. 7.15	99
7.17	Enhancement error and ROI error for TRUIR and HS reconstructions	101
7.18	Mean normalized absolute enhancement error for HS reconstructions	103

7.19	Enhancement error as a function of temporal regularization parameter, α	104
7.20	Mean normalized absolute enhancement error for TRUIR reconstructions with $\log_2 \alpha = 6$ and $\log_2 \beta = 6$	106
7.21	K^{trans} estimation error as a function of α	107
7.22	k_{ep} estimation error as a function of α	108
7.23	v_e estimation error as a function of α	109

LIST OF TABLES

Table

7.1	Values used in simulations.	80
7.2	Kinetic Parameter values for simulated lesions.	80

ABSTRACT

Temporal Regularization Use in Dynamic Contrast-Enhanced MRI

by

Kimberly A. Khalsa

Chair: Jeffrey A. Fessler

Dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) studies demand both high spatial and high temporal resolution. We need high spatial resolution to accurately visualize tissue morphology, and we need high temporal resolution to accurately follow the contrast kinetics of the tissue, which provide clinically important physiological information. We can only acquire so many measurements per unit time, however, and this limited data implies an inherent tradeoff between spatial and temporal resolution in the reconstructed image sequence. Most existing methods undersample the data and then employ some sort of data sharing technique in the k-space domain to recover the ‘missing’ data points. These data sharing schemes are based on an implicit assumption that the dynamic object varies smoothly in time.

We present an image reconstruction scheme based on an object domain model that does not attempt any k-space data recovery, but rather explicitly uses the assumption of temporal smoothness in the image domain to estimate the image sequence that best fits the available data. Our proposed method is called Temporal Regularization Use in Image Reconstruction (TRUIR), and is a penalized likelihood formulation that includes spatial and temporal regularization terms in addition to the data fidelity term. This work presents our TRUIR formulation for both single coil and parallel imaging, and explores various aspects of TRUIR reconstructed image sequences. We evaluate the effect of the spatial and temporal regularization parameters on the resolution properties of TRUIR reconstructed image sequences, and present our work toward establishing selection criteria for these parameters. In evaluating our proposed TRUIR

method, we focus on the application of DCE-MRI in the characterization and assessment of breast cancer. Our simulation studies model contrast uptake in a dynamic digital breast phantom. Results show that TRUIR reconstructions offer improved temporal dynamics when compared to more traditional frame-by-frame reconstructions, as well as more accurate estimates of kinetic parameters, particularly when TRUIR is used in conjunction with two new proposed k-space sampling trajectories, which are also presented in this work. These new trajectories are also shown to be more robust to regularization parameter choice. Further work is needed to improve TRUIR's spatial resolution.

CHAPTER I

Overview

1.1 Motivation: Application to Breast Cancer

Breast cancer is increasingly prevalent in the United States and throughout the world. It is the second leading cause of cancer deaths among women in the United States (after lung cancer). Women have a 1 in 8 chance of being diagnosed with breast cancer in their lifetime [2]. While regular screening using mammography greatly increases the chances of early detection, and thus increases chances of survival, this is not the end of the story for the role of medical imaging in breast cancer detection and management.

Magnetic Resonance Imaging (MRI) is highly sensitive, and, unlike mammography, its sensitivity is independent of breast density. It has been reported that MRI is more sensitive than mammography for detecting tumor size and additional lesions [3]. Dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) is being increasingly used for detection, diagnosis and treatment of breast cancer. In DCE-MRI, a contrast agent is injected into the patient's blood stream and we perform a dynamic MRI scan of the patient as the contrast agent travels throughout the body. The result is a series of MR images that represent the patient's internal anatomy over time, and which provide radiologists with clinically important information both in terms of physical features and temporal behavior of the tissues.

Tumor morphology is an important factor in detection and diagnosis of breast cancer [4–6]. There are a number of morphological indicators; for example, spiculated margins are indicative of a malignant lesion, while smooth edges are indicative of a benign one. Therefore, we want to have high spatial resolution in our MR images to best examine the tissue morphology. In addition to morphological features, the temporal behavior of a tumor in response to injection of a contrast agent is useful in tumor evaluation. [1, 7, 8]. Slow, persistent enhancement is characteristic of benign

lesions, while fast enhancement followed by washout is indicative of malignancy. To accurately assess the temporal kinetics of the tissue, we want to have a series of MR images with high temporal resolution, i.e., where the images are closely spaced in time. To recap, high spatial resolution and high temporal resolution are both desirable for detection, diagnosis, and evaluation of breast cancer. The challenges associated with simultaneously achieving high spatial and temporal resolution are discussed briefly in the next section and in more detail in Section 2.2.3.

The overall goal of this project is to develop, implement, and evaluate methods for improving image quality in dynamic MR imaging. While our methods may have several potential applications, we focus on areas related to breast cancer. In addition to detection and diagnosis, DCE-MRI is being investigated as a method of early prediction of tumor response to neoadjuvant chemotherapy. Chemotherapy is an intense and extremely physically taxing treatment, and unfortunately not all tumors are responsive to it. Neoadjuvant chemotherapy refers to cases when a patient undergoes chemotherapy prior to surgery in an attempt to shrink the tumor before removal. Because some tumors are non-responsive, it is desirable to be able to predict a tumor's response to chemotherapy before treatment is complete. We hope the image sequences created using our proposed reconstruction technique will be able to assist clinicians in evaluating a tumor's initial response to chemotherapy, and thus predict its response (or lack thereof) to an entire course of chemotherapy treatment. For those patients whose tumors are determined to be non-responsive early in the chemotherapy treatment, this mode of therapy could be discontinued. This would save huge amounts of time, emotional strain, and physical distress for these patients, as well as time and money for health care providers.

1.2 Tradeoff between Spatial and Temporal Resolution

If you have ever tried to take a picture at night, you are likely familiar with the challenge of capturing a high quality image in a situations where the available data, or measurements, are limited. In low light, there are fewer photons striking the camera's detector, meaning we have fewer incoming measurements representing the scene we want to capture, and often the resulting pictures are of poor quality.

A common solution for taking pictures in low light is to use a longer exposure time, which means that we leave the camera's shutter open longer, allowing more total photons to strike the detector. That is, we collect more measurements of the scene by accepting incoming photons over a longer period of time, and more measurements




Figure 1.1: Night Photography. Two pictures of the same scene taken with different exposure times. (Image courtesy of <http://www.photoxels.com/tutorial-night-photography.html>)

of the scene results in improved image quality. Figure 1.1 shows an example of two pictures of the same scene taken in low light with different shutter speeds. The picture on the left uses a short exposure time of 0.25 seconds, which is clearly not long enough to collect enough measurements for a good photograph. The picture on the right uses a much longer exposure time of 3.0 seconds, during which time enough measurements are recorded to create a clear, spatially resolved photograph.

Similar to night photography, the available data in MRI is also limited in that we can only collect so many measurements per unit time, and this can result in poor quality images. And, as with night photography, a common solution in MRI is to collect data over a longer period of time. The increased number of measurements enable us to reconstruct a higher quality image.

The focus of this dissertation is dynamic MRI, in which our goal is to create image sequences that accurately represent a dynamic object as it changes in time. Returning to the analogy of night photography, consider what happens if the scene is changing while we are taking the picture. For example, what if a person walks across the sidewalk in Figure 1.1 during the 3 seconds of exposure? Furthermore, what if the person's location at different times is a key element of what we want to capture in our picture(s)? One option is to stick with the 3 second exposure, which will allow collection of a sufficient number of measurements of the scene to create a spatially resolved picture. The caveat, of course, is that the person is walking during those 3 seconds, and will appear as blurred version of himself in the picture. If we want to

determine the person’s location over those 3 seconds based on the photograph, we are essentially limited to stating that during those 3 seconds, he was somewhere within the area of the blurred representation of himself.

In our attempt to accurately capture this dynamic scene, another option is to reduce the length of exposure. This means that we are collecting measurements over a shorter period of time, which will have the positive effect of reducing the blur of the walker in the resulting photograph. However, using fewer measurements also reduces the spatial resolution of the photographic image, as we saw in Figure 1.1.

We can think about our predicament thus: we are collecting measurements of the scene over some length of time, and each individual measurement reflects the scene at the precise time at which it was collected. In grouping the measurements collected over the full 3 second exposure to create a single photograph, we are implicitly assuming that all of the combined measurements represent the same scene, and we consider the resulting photographic image to be an accurate representation of the scene over those 3 seconds. Now that a person is walking across the sidewalk, that scene is different at different points in time during data collection, and the resulting photograph is not a completely accurate representation of the scene, but rather some sort of time averaged approximation of the scene during those 3 seconds.

Because we have limited measurements of a dynamic scene, there is an inherent tradeoff between choosing to group measurements over a larger range of time (i.e. use a long exposure), which will yield photos with high spatial quality (except for the dynamic features), or choosing to group measurements over a smaller time range (i.e. use a shorter exposure), which will result in photos with more accurate representation of the dynamic features, but poorer spatial quality.

This is essentially the same tradeoff we are faced with in dynamic MRI. We want to create a series of magnetic resonance images over time that has good spatial resolution, and that also accurately represents the dynamic features of the object being imaged. And, like with night photography, we can collect only a limited number of measurements per unit time. Using large groups of measurements to reconstruct each image frame results in a dynamic image sequence with high spatial resolution, but poor temporal resolution. Using small groups of measurements to reconstruct each image frame results in a sequence of images with high temporal resolution, but poor spatial resolution. This is the classic spatial vs temporal resolution tradeoff of dynamic MRI.

This thesis proposes a new approach to dynamic MR image reconstruction, which attempts to alleviate this tradeoff. Our reconstruction formulation combines infor-

mation provided by the (limited) measured data with *a priori* information about the dynamic object under study in determining the reconstructed image sequence. We explore various properties of our proposed method, and present results from simulation studies.

1.3 Contributions of this Thesis

The overall goal of this project is to develop, implement, and evaluate methods for improving image quality in dynamic MR imaging. We focus specifically on dynamic contrast-enhanced (DCE) imaging of breast cancer patients. The fundamental challenge in dynamic MRI is the tradeoff between spatial resolution and temporal resolution. In addressing this problem, dynamic acquisition methods usually collect undersampled dynamic data, and most traditional reconstruction methods have been based on operations in the data domain, implicitly assuming that the object varies smoothly in time. We present a reconstruction approach that instead uses explicit temporal models in object space. We use iterative methods to fit these models to the measured data using regularized estimators without attempting to synthesize any of the ‘missing’ data.

The main contributions of this work are:

- An image-domain based reconstruction method for dynamic MRI with a penalized likelihood formulation that includes explicit temporal regularization in object space. We refer to our method as Temporal Regularization Use in Image Reconstruction (TRUIR).
- Acceleration of TRUIR via exploitation of Toeplitz matrices in our reconstruction process.
- Extension of the TRUIR algorithm to incorporate parallel imaging techniques.
- Presentation of two new phase encode sampling trajectories that are similar to a current clinical sampling scheme, but offer significantly more flexibility for reconstructing at higher frame rates.
- Evaluation of the spatial and temporal resolution properties of our system, and their dependence on chosen regularization parameters.
- An accelerated method of computing the local impulse response of our system for the single-coil case, which is an important tool in resolution analysis.

- Improved enhancement curves and kinetic parameter estimates from image sequences reconstructed using TRUIR.

1.4 Outline of the Thesis

This thesis first presents an introduction to dynamic contrast-enhanced magnetic resonance imaging in Chapter II, including its application to breast imaging, its limitations, and discussion of existing reconstruction methods. In Chapter III we present our proposed TRUIR reconstruction approach in detail. We present one existing and two new phase encode sampling schemes in Chapter IV, which we use in conjunction with TRUIR. The challenge of choosing appropriate spatial and temporal regularization parameters is discussed in detail in Chapter V, along with a proposed selection method based on resolution analysis. Chapter VI gives a detailed presentation of three common kinetic parameters used to quantify the temporal behavior of tissues in DCE-MRI, and discusses the relationship of these parameters to the underlying tissue physiology. Here we also provide the mathematical models used to represent this relationship, and establish how to derive kinetic parameters from a reconstructed dynamic image sequence. Our simulation study incorporating our proposed TRUIR reconstruction method with phase encode sampling strategies, regularization parameter selection, and kinetic parameter estimation is presented in Chapter VII. Chapter VIII concludes the thesis and discusses possible areas for future work.

CHAPTER II

Introduction to Dynamic Contrast-Enhanced MRI

2.1 MR Basics

2.1.1 Nuclear Magnetic Resonance

Magnetic Resonance Imaging is based on a phenomenon called nuclear magnetic resonance (NMR) exhibited by the nucleus of certain atoms. The nuclei of atoms with an odd number of protons and/or neutrons, such as ^1H , ^{13}C and ^{19}F , have an intrinsic angular momentum that produces a small magnetic moment. Often visualized as tiny spinning charged spheres, these nuclei are commonly referred to as “spins” throughout MR literature [9, 10].

In the presence of an external magnetic field B_0 , these spins can be excited out of equilibrium by a second applied magnetic field B_1 , that is tuned to the resonant frequency of the spins. As they relax back to equilibrium, the spins will precess around the the axis of the main field B_0 at a rate proportional to the strength of that magnetic field. The rate of precession is called the Larmor Frequency, and depends on both the type of excited nucleus and the strength of the magnetic field experienced by the nucleus. This relationship is described by

$$\omega = \gamma B, \tag{2.1}$$

where ω is the frequency of precession, γ is the gyromagnetic ratio of the nucleus in question, and B is the magnetic field experienced by the nucleus. In MRI of humans, we are mostly interested in the behavior of the hydrogen nucleus ($\underline{\gamma} = \frac{\gamma}{2\pi} = 42.58$ MHz/T), since the human body contains an abundance of them (we’re mostly water). Precessing spins induce current in a nearby receiver coil, and we measure this current to determine a spin’s frequency of precession, or equivalently, the strength of the magnetic field acting on the spin.

After excitation, spins eventually relax back to equilibrium. The return to equilibrium is governed by the spin-lattice time constant, T1, and the spin-spin time constant, T2. T1 characterizes the recovery of the longitudinal component of the magnetization, and is determined by interactions of spins with the surrounding lattice. T2 defines the decay of the transverse component of the magnetization, and is determined by interactions of spins with other spins. Different tissues exhibit different inherent relaxation rates, and these differences are often exploited in MRI acquisitions to produce images with high contrast between certain tissues.

The phenomenon of nuclear magnetic resonance was discovered by physicists long before the advent of Magnetic Resonance Imaging in the 1970's. However it wasn't until 1973 that Paul Lauterbur and Peter Mansfield figured out how to take advantage of this naturally occurring physical phenomenon to create an image of an object [11, 12]. They each proposed the application of a third, spatially-varying magnetic field, G , over the object, in the same direction as the main magnetic field B_0 . Because these additional magnetic fields vary in space (and time), they are called gradient magnetic fields.

Use of additional gradient fields enables spatial encoding of the object (and hence creation of an image of the object), because the location of a spin can be deduced from its rate of precession, and that rate of precession is determined by the strength of the known spatially-varying magnetic field. This concept of spatial encoding via gradient magnetic fields, which Lauterbur initially proposed be called zeugmatography, is the essence of what we now know as Magnetic Resonance Imaging. In 2003, Lauterbur and Mansfield were awarded the Nobel Prize in Medicine for its discovery.

2.1.2 K-space

We know that a spin will precess at a rate proportional to the magnetic field it experiences, as expressed in Eq. (2.1). When we use gradient fields to vary the rate of precession, a spin at spatial location (x, y) accumulates phase according to

$$\phi(x, y, t) = \gamma \int_0^t B(x, y, \tau), d\tau, \quad (2.2)$$

where

$$B(x, y, t) = B_0 + G_x(t)x + G_y(t)y, \quad (2.3)$$

and $G_x(t)$ and $G_y(t)$ represent the time-varying gradients applied in the x and y direction, respectively, and have units of T/m. If we let

$$k_x(t) = \frac{\gamma}{2\pi} \int_0^t G_x(\tau) d\tau, \quad (2.4)$$

$$k_y(t) = \frac{\gamma}{2\pi} \int_0^t G_y(\tau) d\tau, \quad (2.5)$$

and combine with Eqs. (2.2) and (2.3), we can rewrite the accumulated phase for a spin at location (x, y) as,

$$\phi(x, y, t) = \omega_0 + 2\pi k_x(t)x + 2\pi k_y(t)y \quad (2.6)$$

[10]. From their definition, we see that $k_x(t)$ and $k_y(t)$ have units 1/distance. The pair $(k_x(t), k_y(t))$ represent a single point in the (2D) spatial frequency domain, which in MRI we call k-space. We vary the gradients over time to traverse different points in k-space, but it is important to remember that any particular instance in time corresponds to a single k-space location.

To gain a better understanding of the meaning of k_x and k_y and their meaning as spatial frequencies, we can examine how the 2π periodicity of phase relates spins at different spatial locations. For simplicity, we let $G_y(t) = 0$, i.e., apply no gradient in the y -direction, and evaluate the problem in 1D. We know that spins at two locations x_1 and x_2 will have the same phase at time t if

$$\phi(x_2, t) = \phi(x_1, t) + 2\pi N, \quad (2.7)$$

where N is any integer. If we want to find the two closest points x_1 and x_2 that have the same phase, i.e., we want to know the spatial period of the phase, we let $N = 1$. Combining the phase relationship in Eq. (2.7) and the expression of phase in terms of $k_x(t)$ from Eq. (2.6), we can figure out the spatial distance between x_1 and x_2 :

$$\begin{aligned} \phi(x_2, t) &= \phi(x_1, t) + 2\pi \\ \omega_0 + 2\pi k_x(t)x_2 &= \omega_0 + 2\pi(k_x(t)x_1 + 1) \\ k_x(t)x_2 &= k_x(t)(x_1 + 1/k_x(t)) \\ x_2 &= x_1 + \frac{1}{k_x(t)}. \end{aligned}$$

We now clearly see that at time t , the phase is periodic in the x direction with spatial

period $1/k_x(t)$, and therefore $k_x(t)$ is the spatial frequency in x at time t . Similar analysis holds for $k_y(t)$, which is the spatial frequency in y at time t .

2.1.3 Signal Equation

In MRI we use one or more receiver coils to collect data. The object contains a multitude of precessing spins, each one of which contributes to an induced time varying current in the receiver coil that we measure. This leads us to the basic signal equation for MRI, which relates the object of interest to the received signal in the coil.

For simplicity, we first present the signal equation for MR imaging that uses a single receiver coil, which was standard practice for the first $\tilde{20}$ years of MRI. Parallel imaging, which uses multiple receiver coils, is discussed in Section 2.2.4.

Each individual spin over the entire object contributes to the received signal. A spin at location \vec{r} has accumulated a spatially dependent phase at time t according to the applied gradient fields. The signal received by the coil at time t is a sum of the signals produced by all spins throughout the object and therefore is represented as an integral over space. Ignoring relaxation and field inhomogeneity, the signal equation for a MR system with a single receiver coil is given by

$$\begin{aligned} s(t) &= \int f(\vec{r}, t) e^{-i\phi(\vec{r}, t)} d\vec{r} \\ &= \int f(\vec{r}, t) e^{-i2\pi(\vec{k}(t) \cdot \vec{r})} d\vec{r}, \end{aligned} \quad (2.8)$$

where $s(t)$ is the signal received by the coil at time t , $f(\vec{r}, t)$ is the magnetization of the object at location \vec{r} at time t , and $\vec{k}(t) = [k_x(t), k_y(t)]$ is the k-space location at time t , which is a function of the applied magnetic field gradients.

To recover an image that is an estimate of the true magnetization of the object, the relationship in Eq. (2.8) suggests a Fourier relationship between the object, $f(\vec{r}, t)$, and the received signal, $s(t)$. Indeed, taking an inverse Fast Fourier Transform (IFFT) of the k-space data is the traditional image reconstruction method used in MRI. Figure 2.1 shows an object in the image or spatial domain, and the object's equivalent representation in the spatial frequency domain, or k-space.




Figure 2.1: Representation of an object in object space and in k-space. Left: object represented in the spatial/image domain, Right: the object's representation in k-space or the spatial frequency domain.

2.2 Dynamic Contrast-Enhanced MRI

We are interested in dynamic MRI, in which we acquire a series of images over time. The added temporal component can provide additional important information for a variety of applications. As discussed in Section 1.1, our application of interest is the use of dynamic MRI in the study of breast cancer. In most dynamic MRI studies of the breast, a contrast agent is injected into the patient's bloodstream, because the presence of the contrast agent improves the resulting image sequence. Hence, this method of imaging is called dynamic contrast-enhanced MRI (DCE-MRI). The most prevalent contrast agent in DCE-MRI studies of the breast is gadolinium diethylenetriaminepentaacetic acid (Gd-DTPA).

2.2.1 Contrast Agents

Contrast agents (CA) have a T1 shortening affect, which means that tissues where contrast agent is present will exhibit a shorter apparent T1 than their inherent T1 value, and the longitudinal component of the spins located in those areas will recover faster than they would in the absence of contrast agent [9, 10, 13]. In a steady-state gradient echo sequence, an RF pulse is applied at regular intervals and faster recovery of the longitudinal magnetization during each interval (due to T1 shortening) results in larger transverse magnetization after the next excitation. This increased transverse magnetization results in the receiver coils picking up a stronger signal from the spins in those areas, which in turn translates to higher intensity in the reconstructed images. The degree of T1 shortening is proportional to the concentration of the contrast agent present in the tissue, and also depends on the contrast agent of choice as well as

magnetic field strength. The details of this relationship are discussed in Chapter VI. Factors affecting CA concentration in a tissue include blood flow and permeability of the tissue [14, 15].

For any cell to grow and divide, it must have sufficient blood supply to deliver the nutrients needed for cell survival and growth. Tumor cells grow and divide more rapidly than healthy tissue cells, and thus require even higher blood flow and nutrient supply. To ensure this increased supply of blood, many tumor cells produce proteins that promote angiogenesis (growth of new blood vessels) and vascular permeability. This same mechanism that promotes rapid tumor growth also results in faster delivery of an injected contrast agent to tumor tissue, and so we often see higher concentration of contrast agent in tumor tissue compared to healthy tissue [15].

2.2.2 DCE-MRI of the Breast

While x-ray mammography is still the standard method used for breast cancer detection and diagnosis, DCE-MRI of the breast offers some important advantages over mammography. Mammographic images are two-dimensional projections of the three-dimensional object, and therefore offer less overall resolution than MR images, which are generally a collection of 2D profiles of the 3D object. MRI also offers better soft-tissue contrast than mammography, and has been shown to be better at detection of multi-focal, multi-centric cancer in dense breasts, compared to mammography [3]. DCE-MRI can also provide a measure of vascularity, due to the behavior of the contrast agent in the body and its affect on the dynamic image sequence [16].

Tumor morphology is an important indicator in breast cancer evaluation, as is the dynamic behavior of a tumor in the presence of contrast agent. In 1999, Kuhl et al. established the importance of a tumor's time-course in lesion diagnosis [1]. They classified tumors into 3 types, according to their enhancement behavior (see Figure. 2.2). They found that rapid enhancement followed by washout (Type III) was a strong indicator for malignancy, while tumors that exhibited steady enhancement over time (Type Ia and Ib) were more likely to be benign. Several additional DCE-MRI studies have focused on diagnostic characteristics of breast cancer, in terms of both architectural features and dynamic features, and have found that both types of features are useful in lesion classification [17–20]. This finding underscores our need for image sequences with both high spatial resolution (for accurate evaluation of architectural features), and high temporal resolution (for accurate evaluation of dynamic features).

Three standard parameters have emerged to quantify the dynamic behavior of




Figure 2.2: Kuhl et al.'s time intensity curves from [1]. Kuhl et al. characterized tumors according to 3 types of enhancement. Types Ia and Ib represent persistent enhancement, Type II plateaus after initial enhancement, and Type III exhibits washout after initial enhancement.

tissue in DCE-MRI [14]. These parameters are the transfer constant, K^{trans} , the rate constant, k_{ep} , and the volume of extravascular extracellular space (EES) per unit volume of tissue, v_e . These origin and meaning of these parameters are discussed in detail in Chapter VI.

A major area of focus in the DCE-MRI research community has been investigating the use of these kinetic parameters in the breast cancer diagnosis and evaluation [21, 22], as well as their utility in predicting and evaluating patients' response to treatment. The patients in treatment response studies are undergoing chemotherapy prior to surgery, with the aim of decreasing the tumor size and better surgical outcome. However, not all patients respond well to neoadjuvant chemotherapy, meaning the size and extent of the tumor are not significantly decreased by the treatment. Thus, a method of early prediction of treatment response is desirable, so non-responders can be identified early and their treatment course altered. Several studies have focused on evaluating the use of kinetic parameters in early prediction of a patient's response to neoadjuvant chemotherapy. [23–26]. In these studies, several DCE-MRIs are acquired from each patient: one baseline scan is performed before treatment begins, a second scan is performed early during the treatment (typically after 1 cycle of chemotherapy), and often further scans are performed at later points during and/or after treatment.

The results of these studies have been mixed. Pickles et al. found that eventual responders had significantly higher K^{trans} in a lesion “hot spot” during the initial scan than non-responders, and responders also showed a significant decrease in both K^{trans} and k_{ep} between the pre-treatment scan and early treatment scan. They also found a significant difference in the v_e of responders vs non-responders in the early treatment scan [23]. Meanwhile, a similar study by Thukral et al. found that none of the parameters were good predictors of treatment response after a single cycle of chemotherapy [25]. Studies of the use of MRI for response prediction continue, including one here at the University of Michigan.

The goal of this work is to provide dynamic contrast-enhanced MR image sequences that are more accurate than those currently available, thereby enabling better accuracy in the clinical assessments that are based on these images.

2.2.3 Limitations in Dynamic MRI

The classic trade-off in dynamic magnetic resonance imaging is between spatial resolution and temporal resolution. For an MR image to have high spatial resolution, one must collect many measurements in the data domain, called k-space. These

measurements are then used in a reconstruction algorithm to create an image. As discussed in Section 2.1.2, acquiring each k-space measurement takes a finite amount of time. Thus acquiring enough measurements to create an image or volume with acceptable spatial resolution is far from instantaneous.

We want our dynamic image sequence to have both high spatial resolution and high temporal resolution. The ideal sequence would be comprised of a series of high spatial resolution images that are closely spaced in time. That is, a sequence in which each image frame represents the object at a time shortly after the previous frame. We would like to be able to increase the frame rate as high as we want, while maintaining high quality images at each frame. Fixing the frame rate, however, means that we only have time to collect so many measurements for each frame. This, in turn, means that the images produced from these measurements using traditional inverse FFT reconstruction methods will suffer in spatial resolution. Therefore, the physical limitations of the system combined with traditional reconstruction methods will not allow us to simultaneously achieve both high spatial resolution and high temporal resolution. To address this, many alternative sampling and reconstruction approaches have been proposed. These methods are discussed in Section 2.3.

2.2.4 Parallel Imaging

Traditional MRI uses a single coil to transmit and receive the signal. In the past decade the idea of using multiple receiver coils has been introduced and the practice of multi-coil imaging is now widespread [27–32]. The method of using multiple receiver coils is known as *parallel imaging*, because each of the coils is receiving signal from the object at the same time, i.e., in parallel. Because the coils in multi-coil arrays are distributed in space around the object, each coil has non-uniform spatial sensitivity to spins generated throughout the object. A coil positioned near the left of the object will be more sensitive to spins emanating from the left side of the object than those originating from the right side. Thus, each coil in a multi-coil array has its own spatial sensitivity pattern or profile, also known as a *sensitivity map*.

The measurements from an individual coil in a multi-coil array are a product of the true signal from the object and the spatial sensitivity of that particular coil. Using $s_l(t)$ to denote the signal received by the l -th coil and $c_l(\vec{r})$ to denote the spatial sensitivity of the l -th coil, the signal equation becomes

$$s_l(t) = \int f(\vec{r}, t) c_l(\vec{r}) \exp(-i2\pi(\vec{k}(t) \cdot \vec{r}) d\vec{r}. \quad (2.9)$$

When the sensitivity maps are known, we have an added dimension of information for the reconstruction problem. This allows us to reconstruct images from fewer measurements, thereby decreasing the required acquisition time. The need for speed in dynamic imaging applications means that most current MR imaging approaches utilize parallel imaging, although its use is certainly not limited to the dynamic case.

Sensitivity Encoding (SENSE) is one popular method for reconstructing images from multiple coils [28]. The SENSE method (for Cartesian sampling) calls for undersampling in the phase encode direction. With standard reconstruction, increasing the sample spacing in k-space (via undersampling) would result in a reduced Field of View (FOV) in image space and the resulting image would be folded over onto itself. The SENSE reconstruction method uses the known sensitivity maps to derive an unfolding matrix, which is then incorporated into the reconstruction process, resulting in an unaliased final image.

2.3 Existing Reconstruction Methods

A variety of methods have been developed to address the space-time resolution tradeoff in MRI. Many existing methods rely on k-space interpolation or other k-space operations, an approach referred to as k-space data sharing. When sharing k-space data between time frames, one is implicitly using the assumption that the object varies slowly in time. The Keyhole was one of the first such methods, and exemplifies this type of data sharing [33, 34]. In a Keyhole acquisition, one acquires a fully sampled reference frame before the dynamic portion of the study begins. Sometimes a second fully sampled reference frame is acquired after completion of the dynamic acquisition. During the dynamic portion of the study, one collects only a subset of the desired samples for each frame, so as to achieve a faster frame rate (i.e., higher temporal resolution). To reconstruct these undersampled frames, one uses the data from the fully sampled reference frame to first fill in the missing data from the undersampled dynamic frames. The complete image sequence is then reconstructed, frame-by-frame, by taking the inverse FFT of the now “complete” data for each time frame.

Since the introduction of Keyhole imaging, many adaptations and improvements have been made to data sharing techniques. A plethora of methods have been developed for improved recovery of the missing data points. These range from using a simple sliding window over a set of fully sampled data frames to increase temporal resolution, to more sophisticated methods of determining the values assigned to missing k-space locations [30, 35–41].




Figure 2.3: Traditional Keyhole reconstruction.

Another method proposed for dynamic reconstruction is reduced-encoding imaging by generalized-series reconstruction (RIGR) [42–44]. RIGR also fills in missing data for each frame, but is not considered a data sharing method because much of the analysis occurs in the image domain. RIGR recovers the missing samples in under-sampled dynamic frames using a generalized series model, where the basis functions are based on the fully sampled reference image(s). While the RIGR method promises better data consistency than Keyhole, the final step of both methods is the same: apply an inverse FFT to each frame of the imputed k-space data set. In fact, this is the final reconstruction step for all data sharing reconstruction methods. We say this to emphasize the major difference between many existing reconstruction techniques and our proposed reconstruction method. While data sharing techniques propose to recover the missing information in the data domain, we propose to use object-based models to recover the desired information in the image domain.

Another approach to the dynamic imaging problem is Partial Fourier imaging. Partial Fourier imaging with homodyne reconstruction was first proposed to reduce the acquisition time in static MRI, and pre-dates the Keyhole method by several years [45]. Partial Fourier imaging and associated reconstruction methods exploit k-space symmetry, allowing us to reconstruct an image from just over half of the prescribed k-space sample locations, thereby reducing the required scan time by almost half [46]. As with many data sharing methods, Partial Fourier imaging can be combined with parallel imaging [47, 48]. As a means of comparison, we implemented the method of King and Angelos [47], which we will refer to as Homodyne + SENSE (HS) reconstruction. In Chapters V–VI, we compare results from our TRUIR reconstruction to results from HS reconstruction.

Another approach to dealing with the tradeoff between spatial and temporal resolution is to examine the problem in a combined spatiotemporal domain. Krishnan and Chenevert developed a spatio-temporal bandwidth based method (STBB) based

on maximizing the energy captured by the acquired samples, for a given class of objects [49, 50]. Methods such as k-t BLAST and k-t SENSE are based on analyzing spatiotemporal correlations in the dynamic imaging sequence [51–54]. These methods rely on temporal periodicity of the dynamic object and are therefore well suited for dynamic cardiac imaging, but ill suited for DCE-MRI.

While a majority of existing dynamic MRI reconstruction methods are based in the data domain, there are some, such as RIGR, that include image domain models. Block et al. uses an image domain based formulation in [55], and Chen et al. explores the use of Projection onto Convex Sets for reconstruction of DCE-MR images with a temporal total variation constraint in the image domain [56].

Our proposed TRUIR method is based in the image domain and uses only the measured data, undersampled as it may be, in reconstruction. Instead of data sharing, TRUIR explicitly enforces our assumption that the object varies smoothly in time by using an image-domain temporal regularizer.

CHAPTER III

TRUIR: Temporal Regularization Use In Reconstruction

3.1 Measurement Model

Before we can present our proposed TRUIR reconstruction method, it is important to establish the measurement model that we use.

In Chapter II we discussed the MRI signal equation that relates the object to the signal we receive from the MR system. Sections 3.1.1 and 3.1.2 describe how we get from the signal equation to the matrix-vector form of the measurement model that is used throughout this thesis. For simplicity, we first describe the measurement model for the case of a static object and a single receiver coil, and then expand on the model to incorporate dynamic objects and multiple coils.

3.1.1 Static MRI

Ignoring relaxation and field inhomogeneity, the signal equation for a MR system with a single receiver coil and a static object becomes

$$s(t) = \int f(\vec{r}) \exp(-i2\pi(\vec{k}(t) \cdot \vec{r})) d\vec{r}, \quad (3.1)$$

where $s(t)$ is the signal received by the coil at time t , $f(\vec{r})$ is the magnetization of the object at location \vec{r} , and $\vec{k}(t)$ is the k-space sample location at time t , which is a function of the applied magnetic field gradients [10].

In MRI, the measurements y_i are noisy samples of the received signal,

$$y_i = s(t_i) + \varepsilon_i, \quad (3.2)$$

where ε_i is random white Gaussian noise. The combination of Eqs. (3.1) and (3.2) form a continuous to discrete mapping of the object to the MR measurements. Reconstructing the continuous object from our discrete measurements is quite an ill-posed problem; there are infinitely many continuous objects, $f(\vec{r})$, that would perfectly match our N discrete samples.

We want to establish a discrete to discrete model to make the reconstruction problem more tractable. One approach is to create a discrete estimate of the continuous space object by parameterizing it using a discrete number of spatial basis functions. A nice presentation of this approach is given in [57], and we use the same approach, which is outlined below.

The continuous object $f(\vec{r})$ is parameterized using n_p basis functions,

$$f(\vec{r}, t_m) \approx \sum_{j=1}^{n_p} x_j p_s(\vec{r} - \vec{r}_j), \quad (3.3)$$

where $p_s(\vec{r})$ are the spatial basis functions. Because each image is displayed using pixels, a natural choice for the spatial basis functions is $p_s(\vec{r}) = \text{rect}(\vec{r})$, with n_p representing the number of pixels in the reconstructed image. With this formulation, the reconstruction problem can be reduced to estimating the coefficients, x_j , of the spatial basis functions.

Plugging this discrete object estimate into the signal equation (3.1), we get

$$\begin{aligned} s(t) &\approx \int \sum_{j=1}^{n_p} x_j p_s(\vec{r} - \vec{r}_j) \exp(-i2\pi(\vec{k}(t) \cdot \vec{r})) d\vec{r} \\ &\approx \sum_{j=1}^{n_p} x_j \int p_s(\vec{r} - \vec{r}_j) \exp(-i2\pi(\vec{k}(t) \cdot \vec{r})) d\vec{r} \\ &\approx P(\vec{\nu}) \sum_{j=1}^{n_p} x_j \exp(-i2\pi(\vec{k}(t) \cdot \vec{r}_j)), \end{aligned} \quad (3.4)$$

where $P(\vec{\nu})$ is the Fourier Transform of $p_s(\vec{r})$.

Combining Eq. (3.4) with (3.2), we now have a discrete to discrete form of our measurement model:

$$y_i = P(\vec{\nu}_i) \sum_{j=1}^{n_p} x_j e^{-i2\pi(\vec{\nu}_i \cdot \vec{r}_j)} + \varepsilon_i,$$

where $\vec{\nu}_i = \vec{k}(t_i)$ and ε_i is random white Gaussian noise. We stack our n_d measurements collected during the scan and establish the matrix vector form of the measure-

ment model:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\varepsilon},$$

where

$$\mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_{n_d} \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} P(\vec{\nu}_1) \exp(-i2\pi(\vec{\nu}_1 \cdot \vec{r}_1)) & \dots & P(\vec{\nu}_1) \exp(-i2\pi(\vec{\nu}_1 \cdot \vec{r}_{n_p}) \\ \vdots & \ddots & \vdots \\ P(\vec{\nu}_{n_d}) \exp(-i2\pi(\vec{\nu}_{n_d} \cdot \vec{r}_1)) & \dots & P(\vec{\nu}_{n_d}) \exp(-i2\pi(\vec{\nu}_{n_d} \cdot \vec{r}_{n_p}) \end{pmatrix},$$

with \mathbf{x} and $\boldsymbol{\varepsilon}$ stacked similarly to \mathbf{y} . So our system matrix \mathbf{A} is an $n_d \times n_p$ matrix with elements $a_{ij} = P(\vec{\nu}_i) \exp(-i2\pi\vec{\nu}_i \cdot \vec{r}_j)$.

3.1.2 Dynamic MRI

In dynamic MRI, the measured data is a collection of M scans $\mathbf{y}_1, \dots, \mathbf{y}_M$, where $\mathbf{y}_m = (y_{m1}, \dots, y_{m, N_m})$ and N_m denotes the number of k-space samples acquired for the m th scan. To create a discrete representation a dynamic object $f(\vec{r}, t)$ we add a temporal basis function, $p_t(t_m)$, to Eq. (3.3). We parameterize the dynamic object $f(\vec{r}, t_m)$ during the m th scan as:

$$f(\vec{r}, t_m) \approx \sum_{j=1}^{n_p} \sum_{l=1}^M x_{lj} p_s(\vec{r} - \vec{r}_j) p_t(t_m - \tau_l). \quad (3.5)$$

In this work, we use rect functions as our temporal basis functions, i.e., $p_t(t) = \text{rect}(\frac{t}{\Delta t})$. This choice simplifies some of our analysis, although other basis functions may eventually prove more suitable for modeling the temporal component.

We must also make some notes regarding the temporal argument in Eq. (3.5), t_m . This argument refers to the time point we associate with the m th frame in our reconstructed image sequence. Although we assign a single time point t_m to the m th reconstructed frame, acquisition of the data used to reconstruct that frame does not happen instantaneously, and the dynamic object is certainly changing during that finite acquisition time. Therefore, there is some model mismatch in our measurement model for dynamic imaging.

Because the predominant noise in MRI is white Gaussian, the measurement model for the m th frame becomes

$$\begin{aligned} \mathbb{E}[\mathbf{y}_m] &= \mathbf{A}_m \mathbf{x}_m \\ [\mathbf{A}_m]_{ij} &= P(\vec{\nu}_{m,i}) e^{-i2\pi\vec{\nu}_{m,i} \cdot \vec{r}_j}, \end{aligned}$$

where \mathbf{y}_m are the measurements collected during the m th frame, \mathbf{A}_m is the system matrix associated with that frame, and \mathbf{x}_m is the object at frame m . Note that because we are using a rect basis function in time, each \mathbf{y}_m is modeled as a function of the object at a single frame time, t_m . Stacking the data for all M time frames, we can write the measurement model for the entire dynamic sequence as

$$\mathbf{E}[\mathbf{y}] = \mathbf{A}\mathbf{x},$$

$$\mathbf{y} = \begin{pmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_M \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} \mathbf{A}_1 & 0 & \dots & 0 \\ 0 & \mathbf{A}_2 & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{A}_M \end{pmatrix},$$

with \mathbf{x} stacked similarly to \mathbf{y} .

For complex Gaussian measurements, the negative log-likelihood is

$$L(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2,$$

assuming $\sigma^2 = 1$. Estimating $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_M)^T$ by minimizing the negative log-likelihood would be equivalent to reconstructing each scan independently. However, because each \mathbf{y}_m is undersampled in general, the ML estimator would be poorly conditioned and the reconstructions would have low spatial resolution. Instead, we propose the TRUIR cost function, which includes penalty terms that better condition the reconstruction problem, and encourage spatial regularity and temporal smoothness.

3.2 TRUIR Cost Function

Our proposed reconstruction method is called Temporal Regularization Use in Reconstruction (TRUIR). Our reconstruction is the dynamic image sequence, $\hat{\mathbf{x}}$, that is the minimizer of our cost function, $\Psi(\mathbf{x})$:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \Psi(\mathbf{x}),$$

$$\Psi(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \alpha R_t(\mathbf{x}) + \beta R_s(\mathbf{x}). \quad (3.6)$$

The first term in the cost function is the data fidelity term, which ensures that the image estimate is consistent with the measurements and the system model. The second and third terms in the cost function are the temporal and spatial regularizers,




Figure 3.1: Illustration of the effect of TRUIR’s spatial and temporal regularization. The temporal regularization term in Eq. (3.6) penalizes differences between pixels at the same spatial location in adjacent frames.

respectively. We use these terms to incorporate our *a priori* knowledge about the object, namely that there is a certain smoothness expected in both space and time. The regularization parameters α and β determine the relative weighting of the temporal and spatial regularization terms. We use a quadratic penalty in both time and space, with

$$R_t(\mathbf{x}) = \frac{1}{2} \|\mathbf{C}_t \mathbf{x}\|^2$$

$$R_s(\mathbf{x}) = \frac{1}{2} \|\mathbf{C}_s \mathbf{x}\|^2 = \sum_{m=1}^M \frac{1}{2} \|\mathbf{C}_{s0} \mathbf{x}_m\|^2,$$

where \mathbf{C}_{s0} is a spatial differencing matrix that compares neighboring pixels within an image frame and $\mathbf{C}_s \triangleq \mathbf{I}_M \otimes \mathbf{C}'_{s0} \mathbf{C}_{s0}$. The temporal differencing matrix \mathbf{C}_t compares pixels at the same spatial location in adjacent time frames, as illustrated in Figure 3.1.

A first-order temporal penalty function looks like

$$R_t(\mathbf{x}) = \sum_{m=2}^M \frac{1}{2} \|\mathbf{x}_m - \mathbf{x}_{m-1}\|^2,$$

and higher order penalty functions may also be reasonable. We currently use a quadratic penalty in space, but certainly think that using edge preserving penalties in space, such as total variation, will be worth examining in the future.

Letting $R(\mathbf{x}) = \alpha R_t(\mathbf{x}) + \beta R_s(\mathbf{x})$, we can write the minimizer of our TRUIR cost function as

$$\hat{\mathbf{x}} = [\mathbf{A}'\mathbf{A} + \mathbf{R}]^{-1} \mathbf{A}'\mathbf{y},$$

where \mathbf{R} is the Hessian of $R(\mathbf{x})$. We use the conjugate gradient (CG) algorithm to compute this estimate.




Figure 3.2: Enhancement curve used in simulations. True curve (solid) and reconstructions.

3.3 Preliminary Simulation Results

As a first validation step, we performed simulations using a digital phantom in which we generated k-space data and reconstructed an image sequence using our method as well as some existing methods. The digital phantom was comprised of a real, bilateral breast image with an inserted (simulated) circular lesion. The lesion exhibited enhancement according to the curve in Fig. 3.2, while the rest of the image remained static. We generated k-space data for one fully-sampled reference frame, prior to contrast agent injection/enhancement, and eight subsequent under-sampled dynamic frames. The under-sampled frames were sampled using a rotating blade trajectory. Each blade contains 1/16 the samples of the fully sampled frame and is centered in k-space. Essentially, we alternated the phase encode and read-out directions for each dynamic frame. The k-space sampling trajectory is shown in Figure 3.3). For these initial simulations, we did not use multi-coil data, but rather simulated data for a single coil. We reconstructed the data using the methods of simple zero-padding, Keyhole, and RIGR, and compared to reconstructions using TRUIR.




Figure 3.3: K-space trajectory used in simulations.

Because we used a digital phantom, we can compare the spatial resolution and temporal resolution of our reconstructed image sequences to the known true image sequence. For our TRUIR approach, the reconstructions are highly dependent on the choice of the temporal and spatial regularization parameters, α and β . The choice of these parameters is discussed in detail in Chapter V. For the results presented here, α is 2^{14} and β is 2^4 .

The true enhancement curve, which reflects temporal resolution for the true object, along with enhancement curves from various reconstructions are shown in Figure 3.2. The values for the tumor intensity were determined by averaging the pixel values for a region within the lesion boundaries. The TRUIR reconstruction achieves an enhancement curve much closer to the true enhancement curve than the other methods.

As a quantitative measure, we computed the normalized root mean square error (NRMSE) for each frame of the reconstructed sequences over the entire image, as well as within a region of interest (ROI) surrounding the lesion. The error over the object at each frame is shown in Figure 3.4. The error of the zero-padded reconstruction is huge, as expected, due to the severe undersampling of the dynamic frames. The error in the Keyhole and RIGR methods is significantly smaller, although both are outperformed by our iterative, TRUIR reconstruction.

We defined a second, larger circular ROI surrounding the lesion with a diameter twice that of the true lesion. Because accurate reconstruction of morphological features is a primary goal of our method, defining the ROI to surround the lesion allows us to potentially capture poor reconstruction of the lesion edges/margins. The NRMSE within this larger ROI is shown in Figure 3.5 across time. Again, our iterative TRUIR method has lower error than the other reconstructions.

The reason TRUIR has lower error than the competitors can be seen by examining the images reconstructed for the 5th time frame in Figure 3.6. From left to right, top to bottom, the images are: true image, zero-padded reconstruction, Key-




Figure 3.4: Reconstruction error, entire object.




Figure 3.5: Reconstruction error in ROI surrounding lesion.




Figure 3.6: 5th frame of reconstructed image sequence. Top row (left to right): True image, zero-padded reconstruction, Keyhole reconstruction. Bottom row (left to right): RIGR reconstruction, our TRUIR reconstruction.

hole reconstruction, RIGR reconstruction, and our TRUIR reconstruction. The effect of undersampling is best seen in the zero-padded reconstruction, which shows significant blur in the undersampled direction. Note that the direction of the blur alternates in each time frame (not shown), as expected from the alternating blade trajectory (see Fig. 3.3). While both Keyhole and RIGR produce sharper reconstructions than zero-padding, the dynamic feature, the lesion, still shows significant blur in the undersampled direction. Both of these methods use the pre-enhancement, fully sampled frame as a reference, but because the lesion is not visible in the reference frame, neither method is able to adequately reconstruct this dynamic feature. Because our TRUIR method jointly estimates the entire image sequence at once, it is better able to take advantage of the information in the reference frame as well as the alternating blade trajectory to produce an image sequence with significantly reduced blur of the dynamic feature. This can be seen in Figure 3.6 and is also reflected in the reduced NRMSE in the ROI for the iterative method in Figure 3.5.

3.4 Extension to Parallel Imaging

In parallel imaging, we use multiple coils to simultaneously record the measurements. The signal equation for multiple receiver coils is shown in Eq. (2.9), and indicates that the measurements recorded on a particular coil are now a function of both the the object and that coil’s sensitivity pattern (see Fig. 3.7).




Figure 3.7: The measurement model for parallel imaging includes coil sensitivity maps. Each coil records the signal from the object, scaled by that coil's spatial sensitivity profile.

The use of multiple coils of course changes our measurement model. The measurements recorded by the l th coil for the m th frame are given by

$$\begin{aligned}
 y_{m,l} &= \mathbf{A}_m \begin{pmatrix} c_{l,1}x_{m,1} \\ \vdots \\ c_{l,n_p}x_{m,n_p} \end{pmatrix} + \varepsilon \\
 &= \mathbf{A}_m \mathbf{C}_l \mathbf{x}_m + \varepsilon,
 \end{aligned}$$

where the sensitivity of coil l is represented by \mathbf{C}_l . The sensitivity of the l th coil to position \vec{r}_j is given by $c_{lj} = c_l(\vec{r}_j)$, and $\mathbf{C}_l = \text{diag}\{c_{lj}\}$.

Combining the measurements from all L coils, we have the multi-coil measurement model for the m th frame:

$$\mathbf{y}_m = \mathbf{B}_m \mathbf{x}_m + \varepsilon, \quad \mathbf{y}_m = \begin{bmatrix} \mathbf{y}_{m,1} \\ \vdots \\ \mathbf{y}_{m,L} \end{bmatrix}, \quad \mathbf{B}_m = \begin{bmatrix} \mathbf{A}_m \mathbf{C}_1 \\ \vdots \\ \mathbf{A}_m \mathbf{C}_L \end{bmatrix},$$

where \mathbf{y}_m are the stacked measurements from all L coils for the m -th frame, and \mathbf{B}_m incorporates the system matrix for the m -th frame as well as the coil sensitivities.

Finally, for dynamic parallel imaging, we can write the multi-coil measurement model for the entire image sequence as

$$\mathbf{y} = \mathbf{B} \mathbf{x} + \varepsilon, \quad \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_M \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{B}_1 & 0 & \dots & 0 \\ 0 & \mathbf{B}_2 & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{B}_M \end{bmatrix},$$

and \mathbf{x} stacked similarly to \mathbf{y} . In our experiments, we assume that the sensitivity maps are static throughout the scan, although this measurement model does allow the use of updated sensitivity maps for each frame.

For parallel imaging our TRUIR cost-function minimization becomes,

$$\begin{aligned}\hat{\mathbf{x}} &= \underset{\mathbf{x}}{\operatorname{argmin}} \Psi(\mathbf{x}), \\ \Psi(\mathbf{x}) &= \frac{1}{2} \|\mathbf{y} - \mathbf{B}\mathbf{x}\|^2 + \alpha R_t(\mathbf{x}) + \beta R_s(\mathbf{x}),\end{aligned}$$

and the minimizer is given by

$$\hat{\mathbf{x}} = [\mathbf{B}'\mathbf{B} + \mathbf{R}]^{-1} \mathbf{B}'\mathbf{y}, \quad (3.7)$$

where \mathbf{R} is again the Hessian of $R(\mathbf{x})$. Most of the results presented in the following chapters are for the multi-coil case.

3.5 Algorithm Acceleration

A drawback of iterative reconstruction methods, compared to conventional methods, is increased computation time. The most computationally expensive step in using the conjugate gradient algorithm to minimize our cost function, $\Psi(\mathbf{x})$, is computing its gradient:

$$\nabla \Psi(\mathbf{x}) = -\mathbf{A}'(\mathbf{y} - \mathbf{A}\mathbf{x}) + \alpha \nabla R_t(\mathbf{x}) + \beta \nabla R_s(\mathbf{x}). \quad (3.8)$$

If we let $\mathbf{T} = \mathbf{A}'\mathbf{A}$ and $\mathbf{b} = \mathbf{A}'\mathbf{y}$, we can rewrite the equation as

$$\nabla \Psi(\mathbf{x}) = \mathbf{T}\mathbf{x} - \mathbf{b} + \alpha \nabla R_t(\mathbf{x}) + \beta \nabla R_s(\mathbf{x}). \quad (3.9)$$

Because \mathbf{b} does not depend on \mathbf{x} , it needs to be computed only once per minimization, and not once per CG iteration. Additionally, for equally spaced (spatial) basis functions, \mathbf{T} is block Toeplitz with Toeplitz blocks. Recall that we are using $\operatorname{rect}(\vec{r} - \vec{r}_j)$ as our spatial basis functions, chosen because we are reconstructing image pixels. This choice also means that the spatial coordinates \vec{r}_j are equally spaced.

Calculating the gradient in each iteration using Eq. (3.9) now involves multiplying \mathbf{T} by the current guess of \mathbf{x} , which can be done efficiently by embedding \mathbf{T} into a block circulant matrix and doing an FFT [58]. This acceleration technique has previously been investigated for use in static, field-corrected MR image reconstruction [59], but to our knowledge, we are the first to apply it to dynamic MRI.

To validate the speed-up in minimizing our cost function via Eq. (3.9) compared to Eq. (3.8), we performed (single-coil) simulations using each formulation. We found that the Toeplitz-modified conjugate gradient algorithm was 1.7 times faster than the original CG algorithm [60]. While a speed-up factor of 1.7 hardly brings the computation time of our iterative algorithm down to the level of non-iterative reconstruction methods, it is a significant improvement and is certainly a step in the right direction.

CHAPTER IV

Phase Encode Sampling Schemes

Any reconstructed MR image is dependent on the k-space locations that were sampled during data acquisition, and our TRUIR reconstruction method is no exception. Because the dynamic object is changing continuously in time, each individual measurement represents the object at a different point in time. However one cannot reconstruct an entire image of the object from a single measurement, so measurements are grouped into frames for reconstruction. We think of a single frame of the reconstructed image sequence as representing the dynamic object at a single point in time, but in actuality, that image frame is created from measurements taken over a range in time, and is thus only an approximation of the dynamic object during that time.

We employ a flexible definition of an image frame in our TRUIR model, and therefore the acquired k-space locations, as well as their order, are especially important. We explored three k-space sampling schemes and examined the resulting reconstructed image sequences created using both TRUIR and conventional reconstructions. Because sampling in the readout direction happens fairly quickly, we focus only on rearranging/redistributing the phase encode sample locations in our studies. In all of our studies, the phase encode directions are k_y and k_z , and we ignore the readout direction k_x .

4.1 Original Trajectory

The first phase encode (PE) sampling scheme we studied is one that is currently in use for DCE MRI studies of the breast at the University of Michigan hospital, on a Philips 3T scanner. We refer to this as the Original Trajectory. It is an elliptically shuttered, partial Fourier acquisition, that is also undersampled by a factor of 2 in the SENSE direction, k_y . The complete set of sample locations acquired for each frame is




Figure 4.1: Clinical PE sample locations, and full sampling. PE sample locations used in current clinical acquisitions (blue), and sample locations prescribed for full Nyquist sampling (cyan). The original (clinical) pattern is an elliptically shuttered, partial-Fourier acquisition. It also has a SENSE undersample factor of 2 in the k_y direction, which can be seen in the zoomed in portion of k -space pictured on the right.

shown in Figure 4.1 (blue), along with the sample locations that would be prescribed for a fully sampled acquisition (cyan), based on the reconstructed image size of 384×166 and Nyquist criteria. This acquisition has a TR of 4.6 msec, and includes 4766 PE sample locations, which means the nominal frame rate is 22 sec/frame.

One can see that the Original Trajectory is severely undersampled; in fact it samples only 4766 PE locations, compared to the 63,744 locations that would constitute full sampling. The Philips setup at the UM hospital uses a 7-coil receiver array for acquisition, bringing the total number of measurements per frame up to 33,362, which is still barely half of the number required for full sampling. Therefore we can expect imperfect reconstruction of this undersampled data set, regardless of reconstruction method. In all of the alternative sampling schemes we present, we hold constant the total number of samples in a complete trajectory at 4766 to ensure a consistent total scan time. Thus all of the presented PE trajectories are severely undersampled, so when examining various reconstruction schemes, we are not looking for perfect reconstructions, but rather dynamic image sequences that represent an improvement over those created with conventional reconstruction of the original PE sample locations and order.

We are interested in utilizing a flexible definition of frame in our reconstruction methods, in particular, looking at assigning fewer samples per frame, thereby increasing the frame rate, or, equivalently, reconstructing more frames from a given set of dynamic data. For a given set of measurements, if we increase the frame rate by a factor of 2, we reduce the number of measurements per frame by half. The Original




Figure 4.2: Original Trajectory. Dark circles represent PE locations that are sampled earlier in the acquisition, and light circles represent locations that are sampled later during acquisition. The Original Trajectory starts sampling near DC and works its way out, with the highest frequency sample locations being sampled latest in the acquisition.

Trajectory acquires samples roughly from the inside out. That is, it samples low frequency locations earlier in the acquisition and high frequency locations later in the acquisition. Figure 4.2 shows the sample acquisition order of the Original Trajectory. The color of the circle at each PE location indicates when during the acquisition the data at that location is measured. Dark circles indicate locations that are sampled earlier in the acquisition, with the circles getting lighter as time progresses.

The low to high frequency ordering of the original PE acquisition means that if we break the data into 2 subframes, the first half-frame of PE locations will all be from lower spatial frequencies, while the second half-frame of data will all be from higher spatial frequencies. This is illustrated in the top row of Figure 4.3. Without any data sharing or temporal regularization, reconstructing at this higher frame rate would result in an image sequence where every odd frame has only low frequency information, and therefore increased spatial blur, and every even frame has only high frequency information and therefore shows only edges and has very little contrast. Clearly, this is an undesirable result. A 24-frame reconstruction of data collected according to the Original Trajectory would not be a recommended strategy, and is somewhat of a straw man. That is, no one would plan to reconstruct 24 image frames from 12 frames of data acquired with the Original Trajectory, due to the obvious

negative results. But, we are in part considering the situation in which one has only retrospective access to 12 frames of collected data, and one wants to reconstruct 24 image frames; therefore we include the case of a 24-frame reconstruction with the Original Trajectory for completeness.

4.2 Reordered Trajectory 1

We would like the samples within each subframe to be more balanced than those from the Original Trajectory in terms of information content, so that we have the option to increase the frame rate used for reconstruction without suffering the detrimental effects described above. This train of thought led us to examine a trajectory that samples the same k-space locations as the original one, but in a different order. Our first reordered trajectory implements a reordering based on the samples' radial distance from the center of k-space. First we take the range of k-space covered by the Original Trajectory and break it into NR annuli or rings. The number of samples that fall within each annulus determines how frequently we sample from that annulus. That is, each annulus is sampled with a period that is inversely proportional to the number of samples in that annulus. For a 20 second acquisition, a ring containing 2 samples should be sampled every 10 seconds, and our ring-based reordering scheme will assign integer sample time indices to these samples that correspond to 0 sec and 10 sec. A ring containing 4 samples should be sampled every 5 seconds, and those samples will be assigned integer sample time indices

that correspond to 0, 5, 10, and 15 sec. It is easy to see that conflicts will quickly arise with this approach, because prescribed sample times will overlap for rings containing the same number of samples as each other, or for rings that contain a number of samples that is a multiple of the number contained in another ring. And we can only sample one PE location at a time (1 PE location per TR). We solve this overlap problem by adding a small random perturbation to the integer time index assigned to each sample. Sorting the PE locations according to their perturbed time indices (in ascending order) gives us our ring-based reordering. Limiting the perturbations to the range (-0.4 to 0.4) ensures that all PE locations that were prescribed to be sampled at a particular time index will be sampled before any samples that were prescribed to the next time index.

Figure 4.4 shows the Original Trajectory as well as a few ring-based reorderings for different values of NR . As in Fig. 4.2, darker circles represent PE locations that are sampled earlier in the acquisition, and lighter circles represent locations that are

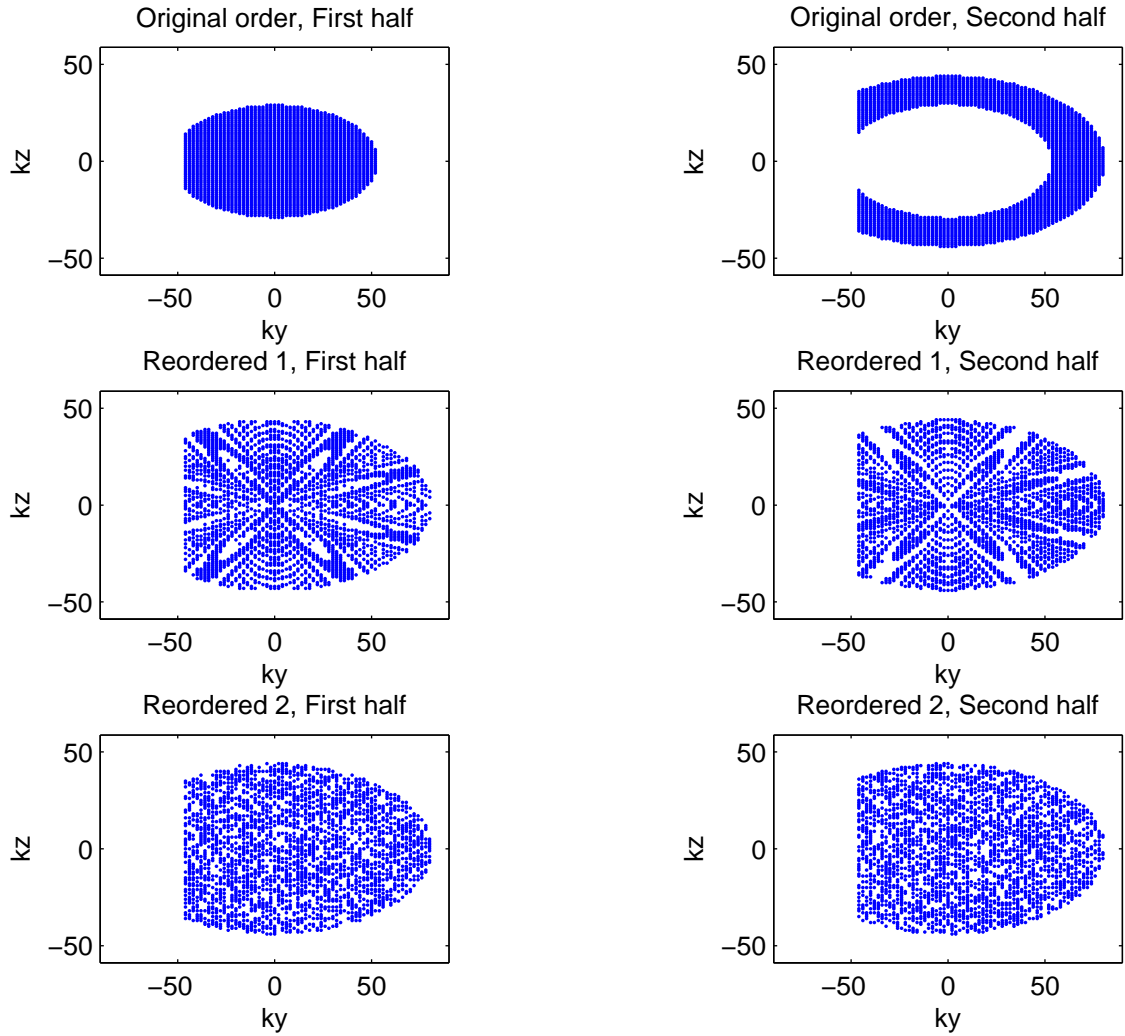


Figure 4.3: PE locations sampled in the first and second half of a “full” frame acquisition. Sample locations acquired during the first half of acquisition are shown in the left column and those acquired during the second half are shown on the right. The top row shows half frame sample locations from the Original Trajectory, the middle row shows those from Reordered Trajectory 1, and the bottom row shows results from Reordered Trajectory 2.




Figure 4.4: Ring-based reordering of PE locations. PE sampling patterns resulting from our initial reordering scheme using varying number of rings (NR). Sample order is indicated by the darkness of circles. PE locations sampled early in the acquisition are darkest, and the markers get lighter as time progresses. The Original PE order is included for reference (top left) and ring-based reordered sampling patterns are shown for NR = 5 (top right), NR = 10 (bottom left), and NR = 100 (bottom right).

sampled later. Reorderings using 5 and 10 rings exhibit a clear circular structure, while the samples reordered using 100 rings look more distributed. Therefore we used the reordering resulting from using 100 rings as our Reordered Trajectory 1, which is shown in Figure 4.5.

The intended (and actual) result of the ring-based reordering is that if we take any given subframe in time, the samples associated with that subframe in the reordered trajectory will cover a wider range of k -space locations than those acquired for the same subframe in the Original Trajectory. This feature is illustrated in Figure 4.6, which plots each sample location of the acquisition in terms of its radial distance from the center of k -space vs the time at which that that sample is acquired. In the plot




Figure 4.5: Reordered Trajectory 1. This trajectory is composed of the same sample locations as the Original Trajectory, but with a different sample order. PE locations indicated by dark circles are sampled earlier in the acquisition, and those indicated by lighter circles are sampled later. Reordered Trajectory 1 was created using our ring-based reordering scheme with 100 rings.

for the Original Trajectory (top), the samples' increasing radial distance over time indicates the inside to out ordering of that trajectory. If we look at samples for only a subframe of the acquisition, say the first 4 seconds, we see that with the Original Trajectory, the samples only cover a small range of radial distances. In contrast, the samples acquired during the first 4 seconds with the Reordered Trajectory 1 (middle plot) essentially cover the same range of radial distances as a full frame acquisition.

Another view of the improved distribution of sample locations in Reordered Trajectory 1 can be seen in Figure 4.3 (second row), which shows that if we break the samples into two half-frames, each half-frame will cover the same *range* of k-space as a full frame. Figure 4.3 also shows that, even using 100 rings, our Reordered Trajectory 1 exhibits some structure in terms of the distribution of samples, such that taking samples from only the first or second half-frame leaves some undesirable holes in k-space coverage. This led us to reconsider our initial reordering scheme, and come up with a second reordering scheme that more uniformly samples k-space throughout the scan.

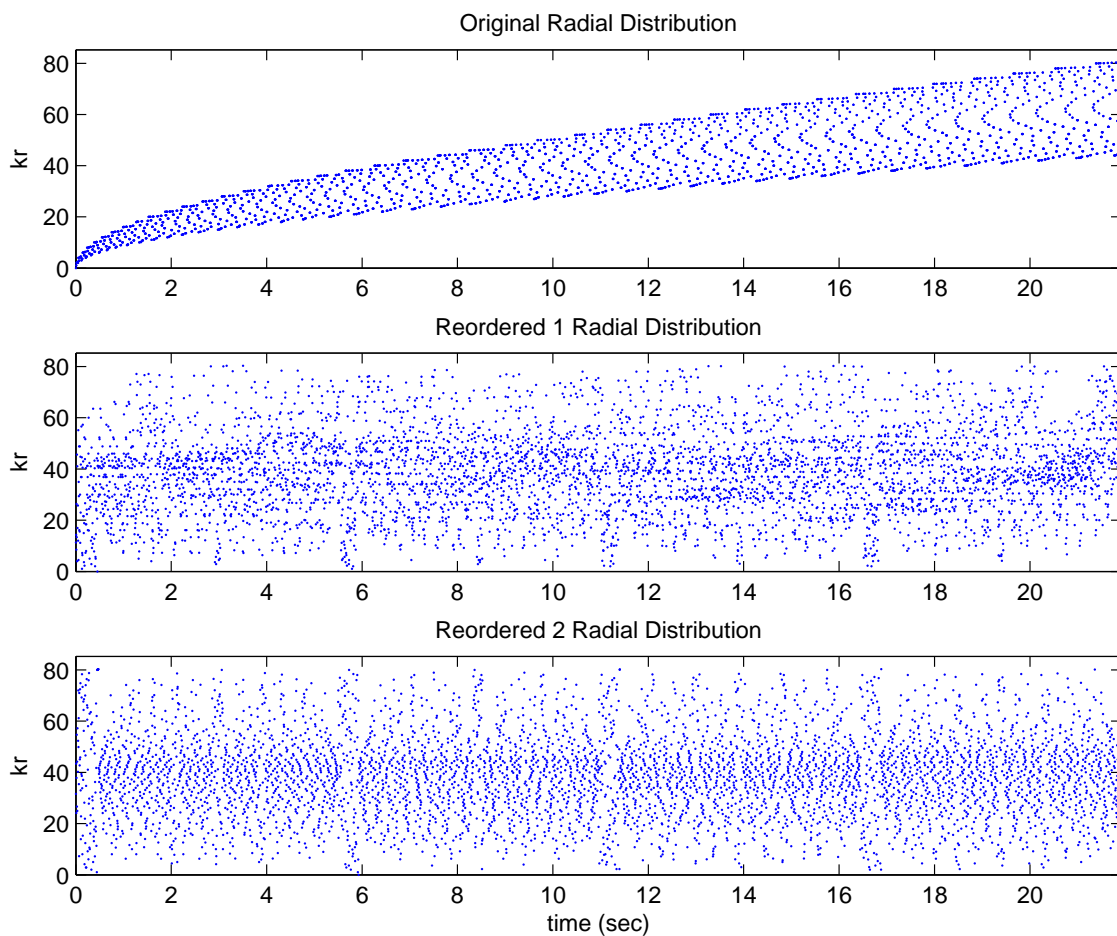


Figure 4.6: Radial distribution of sample locations over time. Original Trajectory (top), Reordered Trajectory 1 (middle), and Reordered Trajectory 2 (bottom). Both reordered trajectories show a more uniform distribution of sample locations over time, in terms of the samples' radial distance from the center of k-space (shown on the y-axis), than the Original Trajectory.




Figure 4.7: Creation of Reordered Trajectory 2. In the first step, we break the original PE sample locations into equally sized random subframes. This figure shows the PE locations for 4 subframes: sample locations for the first subframe are indicated in red, and those in the second, third and fourth subframes and indicated with cyan, blue, and yellow, respectively. After division into subframes, the samples from each subframe are reordered using the ring-based method from Reordered Trajectory 1.

4.3 Reordered Trajectory 2

Our Reordered Trajectory 2 is based on the same premise as the Reordered Trajectory 1, the goal being a more uniform distribution of PE sample locations over time. Our second reordering scheme is similar to our first, but involves one extra step. The Reordered Trajectory 2 is a ring-based reordering of n_{sub} random subsets of the original PE locations. To create our Reordered Trajectory 2, we first divide the original sample locations into n_{sub} distinct random subsets of even size. This first step is depicted in Figure 4.7, where $n_{sub} = 4$ and PE locations belonging to the first, second, third and fourth subsets are represented by red, cyan, blue and yellow, respectively.

The second step in creating our Reordered Trajectory 2 is to individually feed each subset of samples from Step 1 through the ring-based reordering scheme used for the reordered trajectory 1. Lastly, we concatenate the results from Step 2 to form our final reordered trajectory 2. We used a trajectory computed using $n_{sub} = 4$ and $NR = 100$ for all of our simulations, and this Reordered Trajectory 2 is show in Figure 4.8. As the figure shows, the Reordered Trajectory 2 seems to have less




Figure 4.8: Reordered Trajectory 2. Phase encode locations sampled earlier in the acquisition are indicated with dark circles, with the circles getting lighter for locations sampled later during the acquisition.

structure than the Reordered Trajectory 1. This is confirmed by looking at samples from the first and second half frames of our Reordered Trajectory 2 in Figure 4.3, bottom row. Each half frame shows more uniform k-space coverage and fewer large holes than Reordered Trajectory 1 (and certainly than the Original Trajectory).

The radial distribution of samples for Reordered Trajectory 2 is shown in Figure 4.6 (bottom). As with Reordered Trajectory 1, for any given subframe in time, the second reordering shows a more uniform distribution of samples' radial distance from DC than the Original Trajectory.

Changing the k-space sample trajectory is a modification made to the acquisition portion of an MRI study. Chapters V-VI discuss reconstruction related aspects of MRI, which of course also affect the final image sequence. Therefore, we postpone full presentation of results from simulations using the three PE sampling schemes described here to later chapters, where they are presented in combination with a variety of reconstruction methods. Here we simply mention that, as expected, our results show that both Reordered Trajectories do indeed result in more accurate reconstructed image sequences than the Original Trajectory when reconstructing with a higher frame rate. Additionally, reconstructions based on the Reordered Trajectories are much more robust to the choice of the temporal regularization parameter, α , than those based on the Original Trajectory. The effect of the regularization terms and

the challenges of choosing appropriate values for the regularization parameters are discussed in detail in Chapter V.

CHAPTER V

Regularization Parameters

A challenging aspect of any regularized formulation is choosing appropriate regularization terms, as well as determining the relative weights of these terms. These weights, or regularization parameters, have a significant affect on the reconstructed image sequences. Choosing good values for the regularization parameters provides the appropriate weighting between data fidelity and enforcement of our *a priori* assumptions about the object. However, determining these “good” values for the regularization parameters is not straightforward, and using poor regularization parameter values can have an extremely detrimental affect on the resulting dynamic image sequence. This chapter presents our investigations on the effect of regularization parameter choice for TRUIR reconstructed image sequences, and discusses our work towards establishing selection criteria for these parameters.

5.1 Initial DCE-MRI Simulation

To better understand how regularization parameter choice affects dynamic images, we performed a simulation of contrast agent uptake using a real bilateral breast image with an inserted (simulated) circular lesion. The lesion exhibited enhancement over time according to the true curve in Fig. 5.1 (marked with circles), while the rest of the image remained static.

We simulated a dynamic MR scan using the same alternating blade k-space trajectory from our preliminary TRUIR simulations in Chapter III (Fig. 3.3). We generated k-space data for one reference frame and eight subsequent frames, each undersampled by a factor of 1/16. For the partially sampled frames, we alternated the phase encode and readout directions with each frame.

We reconstructed the data using a spatial regularization parameter of $\beta = 2^4$ and first order temporal regularization with several values of the temporal regularization




Figure 5.1: Enhancement curve for simulated lesion and reconstructions

parameter, α . Fig. 5.2 shows the true (A) and reconstructed images from the 5th time frame; zero-padded reconstruction (B), and Keyhole reconstruction (C) are shown for reference. The “tumor” enhancement curves from the reconstructed images appear in Fig. 5.1.

Using an α that is too large degrades the temporal dynamics of the reconstructed sequence, as can be seen for $\alpha = 2^{20}$ in Fig. 5.1 (+). However, for this value of α , the reconstructed image from the middle frame of the dynamic sequence, shown in Fig. 5.2(E), has good spatial resolution. In this case the large weighting of the temporal regularization term in the TRUIR cost function (Eq. (3.7)) enforces strong correlation between image frames associated with the reference and partial datasets. The result is a sequence of reconstructed images with excellent spatial resolution, but a flattened enhancement curve, i.e., poor temporal resolution.

Choosing α too small can also degrade the quality of the reconstructed images. Fig. 5.2(D) shows the 5th frame of the reconstructed image sequence, using $\alpha = 2^5$. This image appears quite blurry because the temporal regularization term in Eq. (3.7) is not large enough to overcome the spatial blur represented by the log-likelihood term (which is due to undersampling of the k-space data in the individual frames). Note that the corresponding enhancement curve in Fig. 5.1 (x) shows good/accurate




Figure 5.2: True image and various reconstructions. True image (A), reconstruction with zero-padding (B), Keyhole (C), TRUIR with $\alpha = 2^5$ (D), TRUIR with $\alpha = 2^{20}$ (E), TRUIR with $\alpha = 2^{14}$ (F).

temporal dynamics.

Essentially, we need α large enough to provide adequate “connectivity” between the frames (especially with the reference frame), but small enough so that the reconstructed image sequence correctly reflects dynamic changes in the object. In this simulation, $\alpha = 2^{14}$ provided such a balance. The reconstructed enhancement curve in Fig. 5.1 (triangles) is a good fit to the true enhancement curve, and the reconstructed image in Fig. 5.2(F) has good spatial resolution.

While Fig. 5.2 shows only a single frame from the reconstructed image sequences (the middle frame), we did look at other image frames (e.g., the 2nd frame and the final frame) as well as the sequences as a whole (not shown). The findings and conclusions regarding spatial resolution based on this more complete examination of the results are consistent with those previously stated; essentially what we see in the 5th frame in Fig. 5.2 is representative of the spatial resolution properties for the entire image sequence.

By reconstructing with a range of regularization parameters and evaluating the resulting image sequences, we have found regularization parameters that work well in reconstruction of this particular simulated data set. However, we would prefer an analytical approach to regularization parameter selection over this ad hoc method. The following section presents our work towards establishing such an analytical method.

5.2 Local Impulse Response

We aim to address the issue of regularization parameter selection by analyzing the resolution properties of the TRUIR method, which is a penalized-likelihood reconstructor. To do this, we examine the local impulse response, which is a measure of resolution [61]. Similar analysis of penalized-likelihood reconstruction for (static) tomography was presented in [62].

In Chapter III, we established the TRUIR cost function:

$$\begin{aligned}\hat{\mathbf{x}} &= \underset{\mathbf{x}}{\operatorname{argmin}} \Psi(\mathbf{x}), \\ \Psi(\mathbf{x}) &= \frac{1}{2} \|\mathbf{y} - \mathbf{B}\mathbf{x}\|^2 + \alpha R_t(\mathbf{x}) + \beta R_s(\mathbf{x}),\end{aligned}\tag{5.1}$$

and saw that the minimizer of the TRUIR cost function is given by

$$\hat{\mathbf{x}} = [\mathbf{B}'\mathbf{B} + \mathbf{R}]^{-1} \mathbf{B}'\mathbf{y},$$

where \mathbf{R} is the Hessian of $R(\mathbf{x})$, and

$$\begin{aligned}R(\mathbf{x}) &= \alpha R_t(\mathbf{x}) + \beta R_s(\mathbf{x}) \\ &= \alpha \frac{1}{2} \|\mathbf{C}_t \mathbf{x}\|^2 + \beta \frac{1}{2} \|\mathbf{C}_s \mathbf{x}\|^2 \\ &= \alpha \frac{1}{2} \|\mathbf{C}_t \mathbf{x}\|^2 + \beta \sum_{m=1}^M \frac{1}{2} \|\mathbf{C}_{s0} \mathbf{x}_m\|^2,\end{aligned}$$

where \mathbf{C}_t is a temporal differencing matrix and \mathbf{C}_{s0} is a spatial differencing matrix.

For practical use of this estimator, one must understand how the regularization parameters α and β in the cost function (Eq. (5.1)) affect the spatial and temporal resolution of the reconstructed image sequence. One method of analyzing resolution is to look at the local impulse response [62]. Noting that $\mathbf{E}(\mathbf{y}) = \mathbf{B}\mathbf{x}$, the local impulse response near the j th pixel is given by

$$\mathbf{l}_j = [\mathbf{B}'\mathbf{B} + \mathbf{R}]^{-1} \mathbf{B}'\mathbf{B}\mathbf{e}_j.\tag{5.2}$$

This expression could be evaluated using the CG algorithm, but computation would be slow as the matrices involved are quite large and non-sparse.

Because only a partial k-space dataset is collected for each time frame, and the sampled k-space locations generally vary from frame to frame, the TRUIR method

is not locally shift-invariant in time. That is, reconstruction of an object at time frame $m + 1$ will not simply be a time-shifted version of the reconstruction of the same object in the m th frame. This complicates the analysis of the local impulse response for dynamic imaging in MRI compared to previous such analyses for Positron Emission Tomography (PET) [63]. In particular, previous methods based solely on FFT calculations are not directly applicable. But, by building on this previous work, we present an accelerated method to determine the local impulse response of the TRUIR formulation for single-coil dynamic MRI. The multi-coil case is quite a bit more complicated and is discussed in Section 5.3

5.2.1 Accelerated Computation for the Single Coil Case

For the single coil case, we do not need sensitivity maps (i.e., $\mathbf{C}_1 = \mathbf{I}$ and $\mathbf{B} = \mathbf{A}$) and we can simplify the expression for \mathbf{l}_j by using the relationship between circulant matrices and FFTs. Namely, the eigenvalues of a circulant matrix are the DFT coefficients of the first column of the matrix [64]. Because the operators $\mathbf{A}'_m \mathbf{A}_m$ and $\mathbf{C}'_{s0} \mathbf{C}_{s0}$ are locally shift invariant in the spatial dimension (but not in the temporal dimension), they can be approximated by circulant matrices and have the following orthonormal eigenvector decompositions:

$$\begin{aligned}\mathbf{A}'_m \mathbf{A}_m &\approx \mathbf{Q} \boldsymbol{\Lambda}_m \mathbf{Q}' \\ \mathbf{C}'_{s0} \mathbf{C}_{s0} &\approx \mathbf{Q} \boldsymbol{\Omega} \mathbf{Q}',\end{aligned}\tag{5.3}$$

which implies

$$\mathbf{A}' \mathbf{A} \approx (\mathbf{I}_M \otimes \mathbf{Q}) \boldsymbol{\Lambda} (\mathbf{I}_M \otimes \mathbf{Q})'\tag{5.4}$$

where

$$\boldsymbol{\Lambda} = \begin{pmatrix} \boldsymbol{\Lambda}_1 & 0 & \dots & 0 \\ 0 & \boldsymbol{\Lambda}_2 & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \dots & 0 & \boldsymbol{\Lambda}_M \end{pmatrix},$$

and \mathbf{I}_M is the $M \times M$ identity matrix. Combining (5.2), (5.3), and (5.4), we get

$$\mathbf{l}_j \approx (\mathbf{I}_M \otimes \mathbf{Q}) \hat{\mathbf{x}}_f\tag{5.5}$$

$$\hat{\mathbf{x}}_f \triangleq [\boldsymbol{\Lambda} + \beta(\mathbf{I}_M \otimes \boldsymbol{\Omega}) + \alpha \mathbf{C}'_t \mathbf{C}_t]^{-1} \boldsymbol{\Lambda} (\mathbf{I}_M \otimes \mathbf{Q}') \mathbf{e}_j.\tag{5.6}$$

While the inverse term in Eq. (5.6) is not diagonal, it is composed of two diagonal matrices, $\mathbf{\Lambda}$ and $\beta(\mathbf{I}_M \otimes \mathbf{\Omega})$, and one sparse matrix, $\alpha \mathbf{C}'_t \mathbf{C}_t$. Thus it should be relatively quick to invert. Due to the approximations and simplifications described above, computing $\hat{\mathbf{x}}_f$ in Eq. (5.6) is much faster than computing \mathbf{l}_j in Eq. (5.2). Once we have $\hat{\mathbf{x}}_f$, we obtain the local impulse response near the j -th pixel, \mathbf{l}_j , using (5.5) through a simple inverse (spatial) FFT.

5.2.2 Single Coil Results

We simulated dynamic MR scans using the same alternating blade k-space trajectory from our preliminary TRUIR simulations in Chapter III (Fig. 3.3). We generated k-space data for one reference frame and eight subsequent frames, each undersampled by a factor of 1/16. For the partially sampled frames, we alternated the phase encode and readout directions with each frame.

We computed the local impulse response using the accelerated method in Eqs. (5.5) and (5.6) for many values of α and β . For this study, using the accelerated method was over 3 times faster than evaluating the original expression for the local impulse response in Eq. (5.2).

Fig. 5.3 shows the effect of the spatial regularization parameter β on the Full Width at Half Maximum (FWHM) of the Point Spread Function (PSF) in space. The spatial PSF is 2D, and we evaluated (and show) the angularly averaged FWHM of the 2D PSF. As expected, higher values of β result in wider point spread functions, because they increase the weight of the spatial differencing penalty in the cost function. For a given sampling trajectory, we can choose the desired spatial resolution of our reconstructed image sequence and then use the established spatial FWHM vs β relationship to determine an appropriate value for β .

For these single-coil studies, we also looked at the temporal PSF and its dependence on α . We modeled an impulse in time as an image sequence that is uniformly zero, except with a single unit impulse in the middle (5th) frame. Reconstructing ‘true’ simulated data representing this temporal impulse yields the temporal point spread function. We measure the FWHM of the temporal PSF in units of time frames.

In this study, we found that the temporal regularization parameter α mainly affects the temporal PSF, and this dependence is shown in Fig. 5.4. Increasing α increases the width of the temporal point spread function, but the shape of the curve is quite different from the one representing the spatial PSF in Fig. 5.3. The shape of the temporal FWHM vs α curve is strongly dependent on the kspace trajectory in use.




Figure 5.3: FWHM_s vs β for single-coil TRUIR. Angularly averaged FWHM of spatial PSF as a function of spatial regularization parameter β for the alternating blade k-space trajectory shown in Fig. 3.3.




Figure 5.4: FWHM_t vs α for single-coil TRUIR. FWHM of the temporal PSF is shown as a function of temporal regularization parameter α for the alternating blade trajectory using first order and second order temporal penalties.

In this case, we are using a fully sampled reference frame, followed by an alternating blade trajectory in the dynamic frames. This sampling pattern is reflected in the temporal FWHM curve.

Use of a very small temporal regularization parameter ($\log_2 \alpha < -10$) results in essentially no temporal regularization, and thus the FWHM in time is near 1 frame. As we increase $\log_2 \alpha$ from -10 to 0, the TRUIR cost function enforces more regularity between time frames, and the point spread function widens accordingly. Increasing $\log_2 \alpha$ more yet, from 0 to 10, the TRUIR cost function is weighted even more toward temporal smoothness, yet we see essentially no change in the FWHM of the temporal PSF.

As we saw in our earlier results of a simulated breast lesion, using $\log_2 \alpha = 14$ yielded the best reconstructed image sequence. This value for the temporal regularization was large enough to draw important spatial information from the initial reference frame throughout the image sequence. In Figure 5.4, we see that $\log_2 \alpha = 14$ corresponds to a FWHM_t of $\tilde{2}.7$ frames for a first order temporal penalty. As with the choice of the spatial regularization parameter, the relationship of α to the temporal PSF is entirely trajectory dependent. Next we present our analysis of the effect of

regularization parameters for parallel TRUIR.

5.3 Multi-coil Case

For the multi-coil case, we must include sensitivity maps in our TRUIR reconstruction, and this significantly complicates the analysis of the affect of the regularization parameters in reconstruction. We performed a number of simulations and evaluated the results over a range of values for both α and β . Our multi-coil simulations generate simulated data from a 7-coil array, and we use real sensitivity maps to create the simulated data and also in our TRUIR reconstructions. See Chapter VII, Section 7.1.2 for details on the sensitivity maps used in our multi-coil simulations.

5.3.1 Spatial PSF vs β

We began our evaluation of regularization parameters for the multi-coil case by again looking at the local impulse response in space. Because the net sensitivity of our coils is not uniform throughout the field of view (FOV), we placed the spatial impulse off center, at $(x,y) = (289, 84)$, where we know there is good sensitivity. Sensitivity at the origin is poor for the smaps we used, so using an impulse at that location would not yield very useful results.

To get a feel for the affect of β on spatial resolution, we looked at a single frame sampled with the Original Trajectory and reconstructed using multi-coil TRUIR with a range of β values. Because all the trajectories under study are undersampled in ky by a factor of 2, reconstructing the full FOV without taking coil sensitivity into account (i.e., reconstructing using uniform sensemaps) yields two peaks, one at the true spatial location of the impulse, and one aliased peak located $ny/2$ pixels away from the true location. Note that because the modeled object (impulse) is static during the course of acquisition, our results will also hold for reconstructions using the reordered trajectories, since they sample the same PE locations as the Original Trajectory.

As in the single-coil case, we looked at the FWHM of the PSF as a function of β , which is shown in Figure 5.5. We know that for single-coil MRI, the FWHM of the PSF is useful for choosing an appropriate value for β . The existence of multiple peaks in multi-coil undersampled MRI, however, means that the FWHM alone doesn't tell us the entire story. In particular, the FWHM of the main (true) peak doesn't indicate much about the secondary (aliased) peak. Therefore, we thought it worthwhile to look at some measure of the two peaks in our analysis, in addition to FWHM.

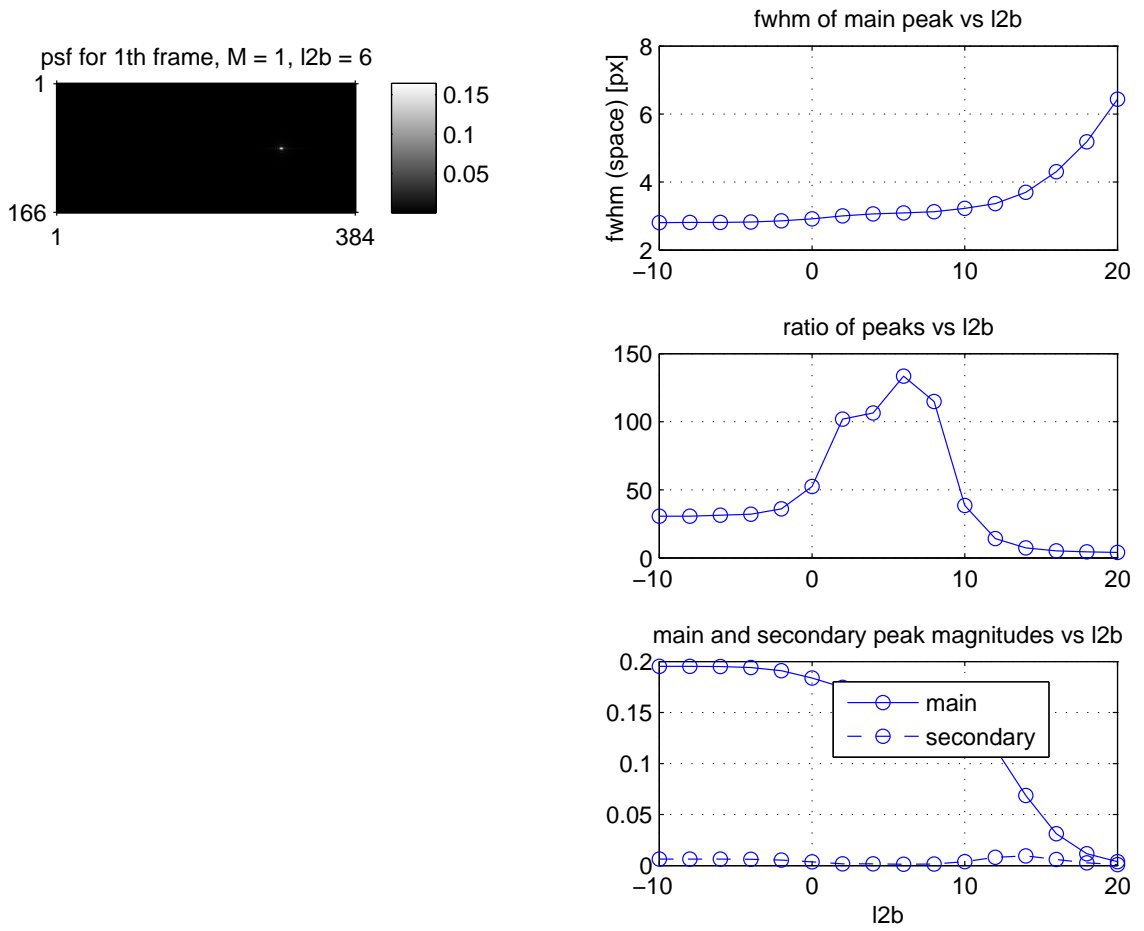


Figure 5.5: Spatial PSF measures vs β for multi-coil TRUIR. FWHM and other PSF measures are shown vs $\log_2 \beta$ ('l2b') for a single frame, fully sampled with the Original Trajectory.

We found that varying the value of β changed the magnitudes of both the main peak and the secondary peak in the reconstructed image. Figure 5.5 also shows the ratio of the peak magnitudes as a function of β (middle), as well as the individual magnitudes of the two peaks (bottom). Interestingly, while the FWHM vs β curve has the standard shape that we saw in the single-coil case, which alone would indicate that we can choose an appropriate value for β based on the desired FWHM of the reconstructed image, the plots of the peak magnitudes tell quite a different story. The height of the main peak decreases with increasing spatial regularization, and the height the secondary peak also changes with β . Although it's rather hard to see in the plot of peak magnitudes vs β , the magnitude of the secondary peak decreases for $\log_2 \beta$ between 0 and 10. The existence of the decreased secondary peak is better seen in the plot of the ratio of the peaks (Figure 5.5, middle), which is higher for β values between 0 and 10, and then decreases drastically for $\log_2 \beta$ of 10 or higher.

Based on these findings, we decided to use $\log_2 \beta = 6$ for the subsequent set of simulations which look at the affect of the temporal regularization parameter, α , on spatial resolution.

5.3.2 Spatial PSF vs α

We are interested in using TRUIR to increase the number of frames in a reconstructed dynamic image sequence, without increasing the amount of data. Doubling the number of reconstructed image frames reduces the number of measurements associated with each frame by half, and TRUIR makes up for the reduced measurements per frame through its regularization terms. In particular, adjacent frames that sample different kspace locations will contain complementary information about the object, even when the object is static. The temporal regularizer enforces smoothness between frames, and thereby allows image domain sharing of information between time frames.

We performed a series of simulations to test the effect of the temporal regularization parameter, α , on an image sequence reconstructed at an increased frame rate. For these simulations, the object is an impulse at a single spatial location that is static over time. We generated 12 “full” frames worth of dynamic data, according to each of three trajectories: the Original Trajectory, Reordered Trajectory 1, and Reordered Trajectory 2. We then used TRUIR to reconstruct 24-frame image sequences from each of the data sets, with a range of values of α ($\log_2 \alpha = [-20 : 5 : 40]$ and $\log_2 \alpha = [-10 : 2 : 20]$).

Because we double the frame rate during reconstruction, the results are no longer

independent of the order in which we acquire sample locations, as was the case for our spatial PSF v β results. Reconstructing 24 image frames from 12 full trajectory acquisitions means that every odd frame in the reconstruction is associated with PE locations measured during the first half of the trajectory’s acquisition, and every even frame in the reconstruction is associated with PE locations from the second half of acquisition. Figure 4.3 in Chapter IV shows the first and second half PE locations for each of the three simulated trajectories.

Reconstructing at this higher frame rate also means that the spatial PSF will differ from one frame to the next, because each frame is associated with a different subset of the data. Thus it is not sufficient to examine the PSF in a single frame. We looked at the spatial PSF in two adjacent image frames when evaluating the effect of α : frames 12 and 13 (of 24). We first looked at image sequences reconstructed from the Original Trajectory. The FWHM of the main peak as a function of α is shown in the top right plot of Figure 5.6. The plotted value of the FWHM at each α is the average of the measured FWHMs from the two adjacent frames. We can see that the mean FWHM is not much affected by the value of α . The mean FWHM is a constant for all values of $\log_2 \alpha$ less than 0, as well as for all values greater than 10, and these two FWHM values differ only by about 0.3 pixels.

The ratio of the main and secondary peaks is also shown in Figure 5.6 (middle right), as well as the individual magnitudes of those peaks (bottom right). Again, the plotted values are the mean of the measure values in frames 12 and 13. Unlike the mean FWHM, the mean peak ratio shows a stronger dependence on α , with values of $\log_2 \alpha$ around 10-15 yielding the best ratio. Looking at the individual peak magnitudes (bottom right), we see that the magnitude of the main peak begins to increase around $\log_2 \alpha = 5$, and the main peak magnitude is maximized for $\log_2 \alpha$ in the range of 10 to 15.

The utility of using the peak ratio (or the magnitude of the main peak) as a measure of goodness of the temporal regularization parameter remains unvalidated until we visually examine the full PSFs in space. Sample PSFs for $\log_2 \alpha = 10$ from frames 12 and 13 are shown in the left column of Figure 5.6, as well as the middle row of Figure 5.7.

To get a better idea of the affect of α on spatial PSFs, we examined two extreme cases. Figure 5.7 shows three sets of PSFs that result from using three different values of α . The low extreme of $\log_2 \alpha = -20$ is shown in the top row, a mid-range value of $\log_2 \alpha = 10$ is shown in the middle row, and the upper extreme of $\log_2 \alpha = 40$ is shown in the bottom row. The PSFs from frames 12 and 13 are pictured (stacked on top




Figure 5.6: Spatial PSF measures vs α , for multi-coil TRUIR. FWHM and other PSF measures are shown for 24-frame reconstruction of samples from the Original Trajectory.




Figure 5.7: PSF for 3 α s.

of one another) in the left column for each value of α . The right column shows the corresponding profiles through the PSFs, where the profiles for frames 12 and 13 are included on the same plot. For the very small value of α ($\log_2 \alpha = -20$), representing essentially no temporal regularization, we see that the PSFs in the two frames look quite different from each other (top left). The PSF from frame 12 is pictured directly above the PSF from frame 13 in the top left of Figure 5.7. The peak in frame 12 is weaker and more diffuse than the peak in frame 13. The profiles through the PSFs (top right) confirm the difference between the frames.

The second row of Figure 5.7 shows PSFs reconstructed using $\log_2 \alpha = 10$ (left), and their profiles (right). Here, the PSFs from adjacent frames look nearly identical, both in terms of the images and in the profile plot: profiles from both frames are plotted, but we see only a single line because the profile plots lie directly atop one another. Note that the magnitude of the PSF peak is also higher for $\log_2 \alpha = 10$ than for $\log_2 \alpha = -20$, by about two-fold.

The bottom row of Figure 5.7 shows PSFs and profiles for very large temporal regularization, with $\log_2 \alpha = 40$. As we would expect from such high temporal regularization, the PSFs from the two adjacent frames look identical in 2D space (left), and the profiles of the PSFs from the two frames lie exactly atop one another, as they did for $\log_2 \alpha = 10$. The peak magnitude in these profiles is significantly smaller, on the order of 10^{-5} , than the peaks for either $\log_2 \alpha = 10$ or $\log_2 \alpha = 20$, which are on the order of 10^{-1} . The profiles for $\log_2 \alpha = 40$ also show a significant secondary peak at the expected location for SENSE aliasing. We suspect that perhaps the TRUIR algorithm takes longer to converge with such a large regularization parameter, and that running it for more iterations may reduce the secondary peak.

Visual inspection of the PSFs from two adjacent frames reconstructed with $\log_2 \alpha = 10$ has shown that they are nearly identical, which tells us that using a temporal regularization value of $\log_2 \alpha = 10$ provides sufficient connectivity between frames to overcome each frame's undersampled data in this simulation. Because we are reconstructing a static impulse, however, we will not be able to differentiate here between sufficient temporal connectivity and oversmoothing in time.

Thus far we have looked at the effect of temporal regularization on various aspects of the spatial PSF for a large range of α s that are rather roughly sampled, with a spacing of $\log_2 \alpha = 5$ between evaluated α values. We have determined that looking at FWHM of the main peak (averaged over two adjacent frames) does not provide the same insight into choosing the temporal regularization parameter as it did for choosing the spatial regularization parameter. The ratio of the main peak to secondary peak,

or the magnitude of the main peak alone, may be useful in choosing an appropriate value for the temporal regularization parameter, or at least in narrowing the range of values to be tested.

Next we take a closer look at PSFs created from a finer sampling of α values, and evaluate yet more properties of these PSFs in terms of their predictive value for determining an appropriate weight of the temporal regularization term in our TRUIR reconstructions. We evaluated the static impulse response of reconstructions using $\log_2 \alpha = [-10 : 2 : 10]$, for each of the three phase encode trajectories, again reconstructing 24 frames from 12 frames of data.

For all tested values of α , we visually examined the PSF over all 24 frames, and plotted the value of the PSF peak vs frame, as well as the frame to frame change in FWHM of the PSF. The results for PSFs of the Original Trajectory are shown in Figure 5.8. The top portion of the figure shows the PSF at each of the 24 reconstructed frames (along the x-axis) for each of the tested values of α (along the y-axis). Here we clearly see that for $\log_2 \alpha = -10$, the PSF at every even frame looks quite different from the PSFs in the odd frames. For the Original Trajectory, all of the odd frames are associated with the low-frequency half of the PE samples, and all of the even frames are associated with the high-frequency samples in the acquisition. Therefore, the odd and even image frames initially contain complementary information (based on the data-fidelity term), and without sufficient temporal regularization this information is not adequately shared between frames. Thus, for $\log_2 \alpha = -10$ the PSFs in the odd frames are from only low-frequency samples and appear rather blurred, while the PSFs in the even frames are from only high-frequency samples, which is reflected in the reduced energy and visible sharp edges in the even frames. When we increase $\log_2 \alpha$ to 8 or 10, we start to see the effect of the temporal regularizer across the PSFs of all 24 frames. There is less variability between PSFs in adjacent frames with $\log_2 \alpha = 8$ than $\log_2 \alpha = 6$, and the interframe variability is smaller yet with $\log_2 \alpha = 10$.

In an attempt to quantify these qualitative observations, we looked at how the magnitude of the PSF peak changes from frame to frame, for the range of α values. This is the middle plot of Figure 5.8, where each value of α is represented by a different colored line. Interestingly, we found that the magnitude of the PSF peak changes very little from frame to frame, regardless of the value of α , and larger values of α s (particularly $\log_2 \alpha = 8$ and $\log_2 \alpha = 10$) result in slightly more variation of PSF peak magnitude between frames. Using $\log_2 \alpha = 8$ and 10 also increased the peak height for all frames, which is in line with our previous findings from Figure 5.6 (bottom right).




Figure 5.8: PSF features as a function of α for the Original Trajectory with 24-frame TRUIR reconstruction. The top figure shows the PSF over 24 frames (along the x-axis) for various values of α (y-axis). The middle and bottom plots are of the peak PSF value at each of the 24 frames, and (angularly averaged) FWHM of the PSF at each frame, respectively. Each value of α is represented by a different colored line.

While the peak magnitude across the dynamic sequence does provide some information, it does not fully capture the inter-frame variability we observe qualitatively by looking at the PSFs. As a second attempt to quantify those qualitative observations, we looked at the change in FWHM from frame to frame, which appears in the bottom plot of Figure 5.8. Here we see a clear distinction in interframe FWHM variability between α values that resulted in qualitatively bad and good looking PSFs. Values of α s whose PSFs show distinct differences between odd and even frames have high interframe FWHM variability. For all tested $\log_2 \alpha$ values less than 6, the FWHM for even frames was about 2.5 pixels, while the FWHM for odd frames was almost twice that, at about 4.5 pixels. However, for $\log_2 \alpha = 8$ and 10, which we have observed result in the best looking PSF sequences, the difference in FWHM between odd and even frames is much smaller, around 0.1-0.2 pixels. These findings suggest that looking at the interframe variability of the FWHM could be useful in choosing a good value for α . A temporal regularization parameter that is sufficient to reduce the interframe FWHM variability of a static PSF may also be sufficient to support adequate sharing of information in reconstructed image sequences of dynamic objects.

Figure 5.9 shows the PSF images, PSF peak and FWHM over 24 frames for data acquired using Reordered Trajectory 1, and $\log_2 \alpha = [-10 : 2 : 10]$. Figure 5.10 shows the same for Reordered Trajectory 2. The results for both reordered trajectories are largely similar, and the following observations hold for both. Because the samples in the reordered trajectories are more evenly distributed in kspace over time, we do not see any strong visible differences between PSFs of even and odd frames, even for the smallest tested temporal regularization parameter value of $\log_2 \alpha = -10$. In fact, visual inspection does not show a significant difference in the PSF over 24 frames for $\log_2 \alpha = -10$ compared to the PSF for $\log_2 \alpha = 10$. Furthermore, the PSFs for the reordered trajectories with any of the tested values for α look visually superior to the best PSF for the Original Trajectory. The plot of peak magnitude across frames for both Reordered Trajectories shows very little dependence on the value of alpha, with all peak magnitudes falling within 10% of each other. Similarly the difference in FWHM between frames does not show much dependence on α , at least within the range of α s included in this simulation. The maximum change in FWHM between frames is about 2%, over all values of α .

These results suggest that reconstructions of samples acquired using Reordered Trajectory 1 or Reordered Trajectory 2 are more robust to the choice of the temporal regularization parameter, α , than are reconstructions of samples acquired using the Original Trajectory.

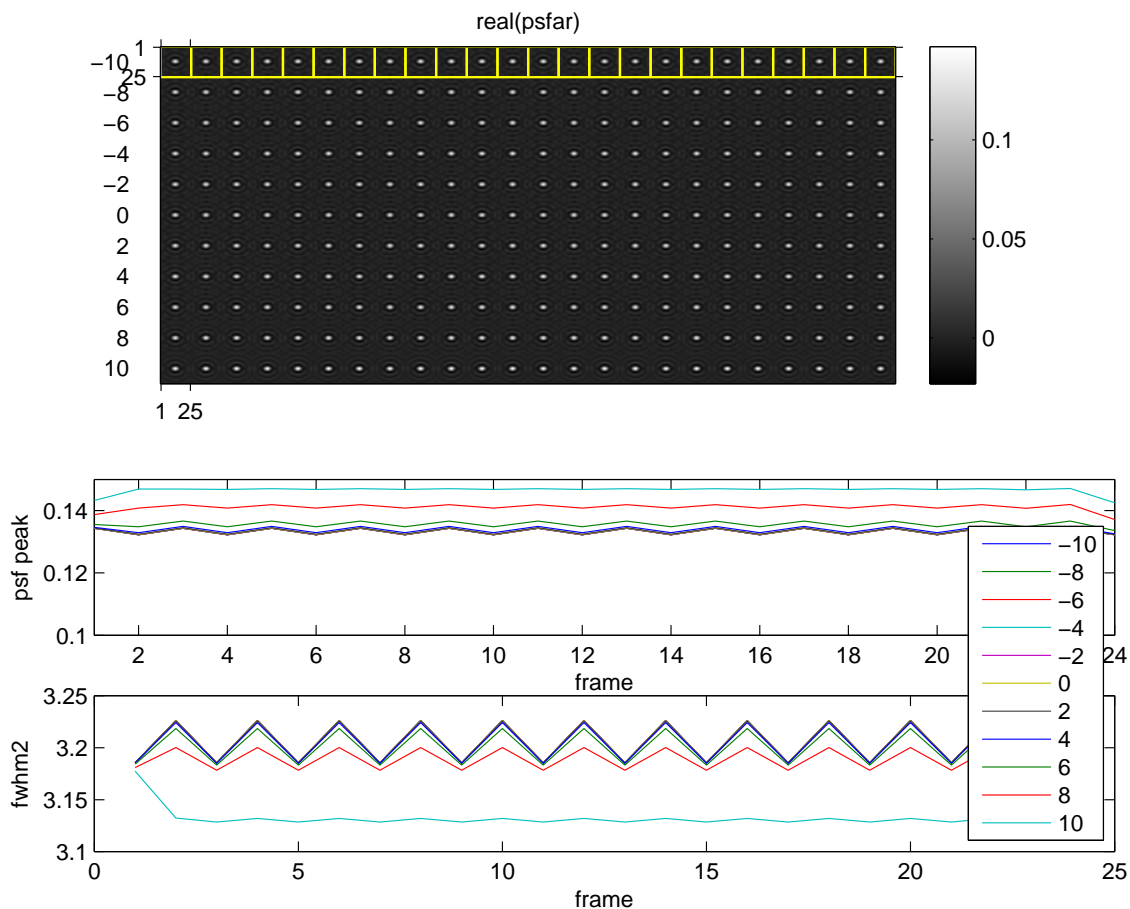


Figure 5.9: PSF features as a function of α for Reordered Trajectory 1 with 24-frame TRUIR reconstruction. The top figure shows the PSF over 24 frames (along the x-axis) for various values of α (y-axis). The middle and bottom plots are of the peak PSF value at each of the 24 frames, and (angularly averaged) FWHM of the PSF at each frame, respectively. Each value of α is represented by a different colored line.

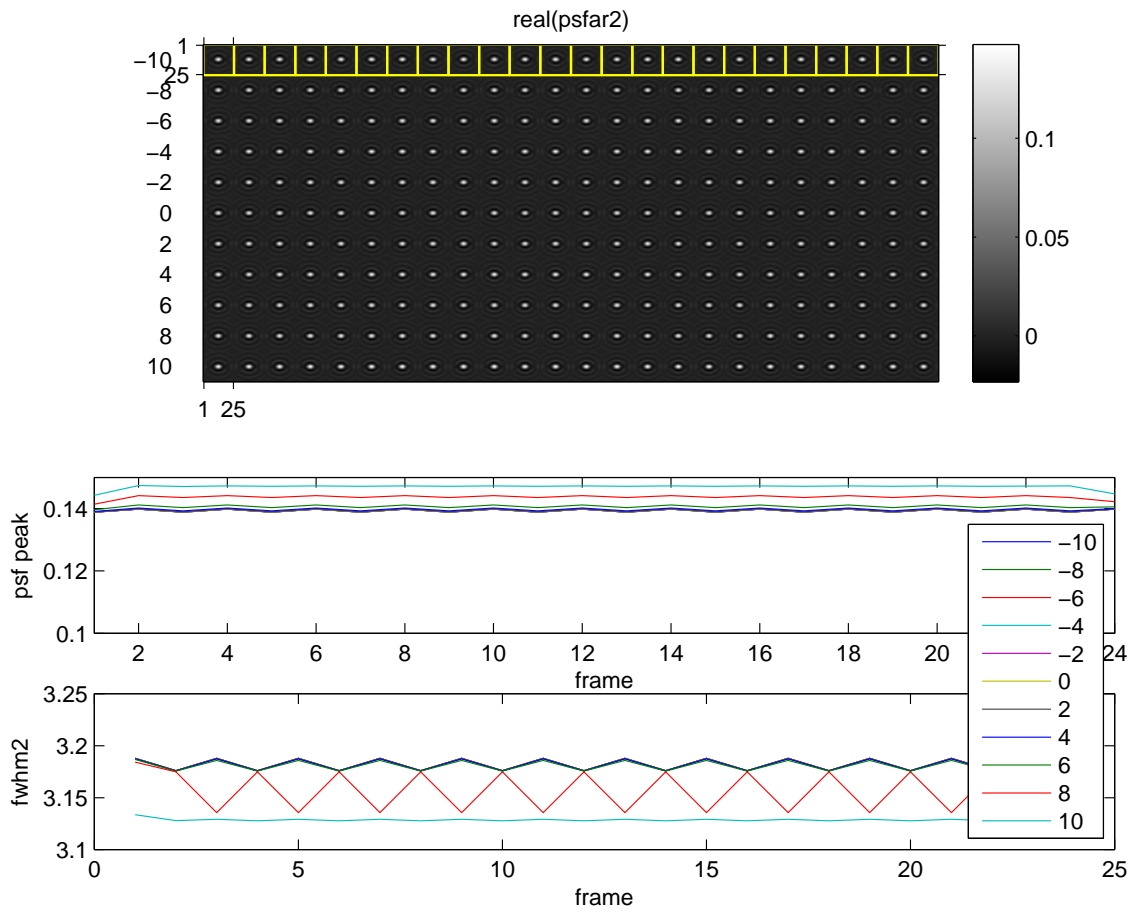


Figure 5.10: PSF features as a function of α for Reordered Trajectory 2 with 24-frame TRUIR reconstruction. The top figure shows the PSF over 24 frames (along the x-axis) for various values of α (y-axis). The middle and bottom plots are of the peak PSF value at each of the 24 frames, and (angularly averaged) FWHM of the PSF at each frame, respectively. Each value of α is represented by a different colored line.

We looked at one final measure of the effect of α on the static impulse response: the energy contained in the tails of the PSF. For each PE ordering, we computed the energy in the PSF tails, for frame 12 and frame 13. We defined the tails to consist of any signal located outside of a $r_y \times r_z$ rectangle, centered at the location of the impulse, where the dimensions of the rectangle were chosen to be equal to 3 FWHMs of the PSF reconstructed from a full frame’s worth of samples. In this study, $r_y = 13$ pixels and $r_z = 9$ pixels. The results are shown as a function of α in Figure 5.11. In the Original Trajectory, the odd frames are associated with low frequency data, and the even frames are associated with high frequency data, and the energy we see in the tails for the Original Trajectory is consistent with this. Frame 12 (high frequencies) has much higher energy in the tails than frame 13 (low frequencies), when the reconstructed image sequence is under-regularized in time. Figure 5.11 also plots the mean of the tail energy in frames 12 and 13, as well as the tail energy that results from “full” sampling, i.e., a 12-frame reconstruction in which each frame is associated with a full trajectory’s worth of samples. Curiously, for larger values of α , the energy seen in the tails for frames 12 and 13 of all three 24-frame reconstructions is lower than the tail energy of a “fully sampled” 12-frame reconstruction.

Figure 5.11 supports our previous findings that the choice of temporal regularization parameter has a much larger effect on PSFs of the Original Trajectory than on those from the Reordered Trajectories. Again this suggests relative robustness of the Reordered Trajectories to temporal regularization parameter choice.

The tail energy plots are similar for all three trajectories in that tail energy remains constant (within a frame) for α values of $\log_2 \alpha$ from -10 to 2, and the tail energy is also relatively constant for values of $\log_2 \alpha$ greater than 10. Within the range of $\log_2 \alpha$ between 4 and 10, we see a clear shift in the amount of energy in the tails of the spatial PSF. For the Reordered Trajectories, the most pronounced decrease in tail energy occurs between $\log_2 \alpha = 4$ and $\log_2 \alpha = 8$. Interestingly, in our dynamic simulations (results of which are presented in Chapter VII), TRUIR reconstructions exhibited the best temporal dynamics, i.e., the most accurate lesion enhancement curves, for values of α in this same range (see Figure 7.19). This suggests that evaluation of the tail energy of the spatial PSF may be useful in determining appropriate values for the temporal regularization parameter in TRUIR reconstructions of data acquired using the Reordered Trajectories.

The conclusions from our studies on the selection of the temporal regularization parameter for multi-coil TRUIR are as follows. For 24-frame reconstructions of samples acquired according to the Original Trajectory, we need to use a relatively large




Figure 5.11: PSF tail energy vs α . The energy in the tails of the spatial PSF is plotted vs $\log_2 \alpha$ ('l2a'), for frame 12 (blue circles) and frame 13 (red squares) of 24-frame reconstructions using the Original Trajectory (top), Reordered Trajectory 1 (middle), and Reordered Trajectory 2 (bottom). The mean tail energy seen in frames 12 and 13 is also plotted (black dashed line), as well as the tail energy that results from “full” sampling, i.e., a 12-frame reconstruction in which each frame is associated with a full trajectory’s worth of samples (as opposed to the 24-frame reconstructions in which each frame is associated with only half of the samples that comprise the full trajectory).

value for α , in the range of $\log_2 \alpha = 10$, to enforce enough temporal smoothness for adjacent frames to adequately share their complementary information, thereby compensating for the undersampling in each frame. For 24-frame reconstructions of samples acquired using Reordered Trajectory 1 or Reordered Trajectory 2, we examined many measures, and we found all of the PSF features that we investigated to be relatively insensitive to the choice of α (within the range that we tested). This is good news, as it suggests a relative robustness to temporal regularization parameter choice for the Reordered Trajectories, which means that we should be able to reconstruct quality image sequences without spending too much time looking for an α that is “just right”. Examining the energy in the tails of the spatial PSF may be useful in predicting “good” values of α for reconstructions of the Reordered Trajectories.

Examining simulation results of a realistic dynamic phantom will shed further light on the topic of temporal regularization parameter selection, and these results are presented in Chapter VII.

5.3.3 Temporal PSF

For single coil TRUIR, we examined the temporal impulse response as a function of α (see Figure 5.4). After some reflection, we decided not to look at this relationship for the multi-coil case, for a number of reasons discussed below. We feel that the spatial PSF measures examined in the previous section are a sufficient, or perhaps better, measure of the intended purpose of our temporal regularizer.

To measure the true temporal resolution of the system, we would need to measure the system’s response to an impulse in time. In our single-coil simulations, we modeled this temporal impulse as a one-pixel object that appeared in a single image frame. However, acquisition of the data that is reconstructed to create that single image frame is not instantaneous, it takes several (tens of) seconds. A true impulse in time would only be reflected in a single k-space sample, and the resulting reconstructed image would depend on which k-space location happened to be sampled at the instant the temporal impulse occurred. The image frame would exhibit a sinusoidal pattern, and the image sequence would certainly not resemble the true object. Needless to say, the system is absolutely not temporally shift invariant, even locally, since shifting the temporal impulse by just milliseconds would alter the resulting image sequence. This is a major reason we decided not to examine the temporal PSF in our later (i.e., multi-coil) studies.

Additionally, our TRUIR formulation is based on an assumption of an object that varies smoothly in time. A temporal impulse is the antithesis of a smoothly varying

dynamic object, and therefore the response of our system to a temporal impulse is not necessarily a good measure of the general response of our system for reconstructing the intended class of dynamic objects for which it was designed.

Furthermore one of the main goals of the temporal regularizer is to provide a means of sharing complementary information between undersampled image frames. The degree of information sharing between frames can be evaluated by examining the variability of the spatial PSF between adjacent frames, which we measured in various ways in the previous section. Therefore, we feel that the simulations performed and results presented in Section 5.3.2 adequately establish the degree of information sharing between image frames as a function of the temporal regularization parameter, α , and evaluation of the somewhat hazy concept of the temporal PSF is not necessary. Examining the effect of the temporal regularization parameter on a smoothly varying dynamic object will provide more information, and these results are presented in Chapter VII.

CHAPTER VI

Kinetic Parameters in DCE-MRI

In dynamic contrast-enhanced (DCE) MRI, a contrast agent is used to enhance the MR images. The contrast agent is injected into the subject's blood stream and as it travels through the body, it affects the underlying physical mechanisms of MRI contrast and thereby alters the appearance of vascularized tissues in the MR image. This section describes the mathematical models we use to represent the physical processes that occur during DCE-MRI, and how they affect MR images. These mathematical representations are useful for both accurate modeling the system to create realistic simulations, as well as for interpreting reconstructed DCE-MR image sequences to determine important underlying physiological processes.

6.1 Modeling Tissue Physiology

In DCE-MRI studies, a contrast agent (CA) is injected into the patient's blood stream at a particular rate for a particular amount of time. The concentration of CA in the blood stream and various tissues affects the resulting MR image sequence. We need to know how the CA concentration varies over time in order to accurately interpret these images. Commonly, CA concentration in a large artery is measured to establish an arterial input function (AIF), which serves as a starting point for determining the CA concentrations in other areas of the body and ultimately inferring information about the tissue(s) of interest. When an accurate AIF is not available, researchers have proposed use of reference region methods to determine CA concentrations in the tissue of interest [65].

Parker et al. established a model for the concentration of an injected contrast agent in the aorta based on a population-averaged AIF [66], and we use Yang et al.'s modified version of this model in our simulations [67]. In this model the CA

concentration in the aorta over time is given by

$$\begin{aligned}
C_p^{Aorta}(t) = & 7.5527 \exp \left[\frac{-(t - 0.171)^2}{0.00605} \right] \\
& + 1.0003 \exp \left[\frac{-(t - 0.364)^2}{0.035912} \right] \\
& + 1.064 \exp(-0.083t) / [1 + \exp(-37.772(t - 0.482))], \quad (6.1)
\end{aligned}$$

where t is provided in minutes, and the calculated concentration, C_p^{Aorta} , is given in mMol.

As the contrast agent travels through the blood stream, it disperses somewhat. The contrast agent concentration in the blood plasma near the tissue of interest, $C_p(t)$, is modeled as a convolution of the aortic CA concentration with a dispersive transport function, $h(t)$:

$$C_p(t) = C_p^{Aorta}(t) * h(t), \quad (6.2)$$

where $*$ is the convolution operator. Yang et al. define the transport function, $h(t)$, as

$$h(t) = \beta^{-\alpha} (t - t_0)^{\alpha-1} \exp[-(t - t_0)/\beta] \Gamma(\alpha), \quad (6.3)$$

where Γ is the Gamma function, t_0 is the bolus arrival time, and they (and we) use $\alpha = 4$ and $\beta = 0.03$ min.

We have now established a mathematical representation for the CA concentration in the blood plasma *near* the tissue of interest. The following section explains the relationship between the plasma concentration, $C_p(t)$, and the CA concentration in the tissue of interest, $C_t(t)$.

6.1.1 Definition of Kinetic Parameters

The flow of contrast agent between the blood plasma and tissue depends on both blood flow and vascular permeability, and varies depending on tissue type. In fact, it is this inter-tissue variability that we are trying to determine in a DCE-MRI study. The two-compartment model in Figure 6.1 is commonly used to represent the flow of contrast agent between the blood plasma and the tissue of interest. According to this model, the tissue concentration, $C_t(t)$, is related to the plasma concentration, $C_p(t)$, via

$$C_t(t) = K^{\text{trans}} C_p(t) * e^{-k_{\text{ep}} t}, \quad (6.4)$$

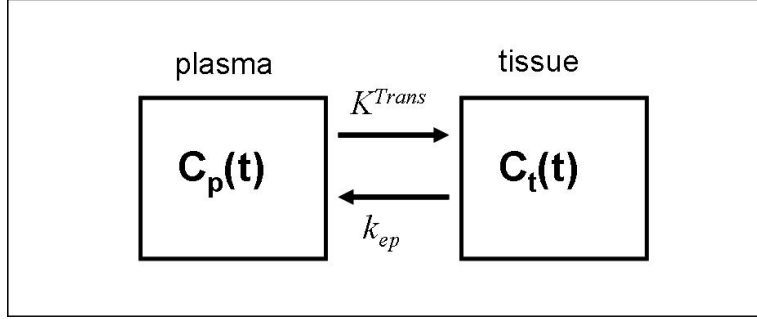


Figure 6.1: The two-compartment model relates the contrast agent concentration in the blood plasma, $C_p(t)$, to the tissue concentration, $C_t(t)$. The physiology of the tissue affects the flow of contrast agent between the plasma and the tissue, and is modeled by two kinetic parameters. The transfer constant K^{trans} represents the flow of contrast agent from the plasma to the tissue, and the rate constant k_{ep} , represents the flow of contrast agent from the tissue back to the plasma.

where K^{trans} represents the flow rate of CA from the plasma to the tissue and is known as the volume transfer constant, and k_{ep} represents the flow rate from tissue to plasma and is known as the flux rate constant. A third kinetic parameter is also regularly referenced in DCE-MRI: v_e , which is the volume of extravascular extracellular space, per unit volume of tissue, and is related to K^{trans} and k_{ep} by $k_{ep} = K^{trans}/v_e$ [14].

Typical concentration curves for $C_p^{Aorta}(t)$, $C_p(t)$, and $C_t(t)$ are shown in Fig. 6.2, for a slowly enhancing lesion. The transport function, $h(t)$, is shown in Fig. 6.3, and a plot of the plasma-tissue transfer function given by the two-compartment model, $h_t(t) = K^{trans} \exp(-k_{ep}t)$, is shown in Figure 6.4.

Contrast agents like Gd-DTPA act as T1 shortening agents, thereby affecting the MR signal received. Each contrast agent has a characteristic relaxivity that affects the relaxation rates of a tissue (T_1 and T_2). Relaxivity is denoted by the variable r and is defined as the increase in relaxation rate per unit increase in contrast agent concentration; thus it has units $\text{mMol}^{-1} \text{sec}^{-1}$, where concentration is measured by molarity (Mol), which is the number of moles of solute in 1 liter of solution, and 1 mole is 6.02×10^{23} particles. Longitudinal relaxivity is denoted as r_1 and relates the contrast concentration to T_1 according to:

$$\frac{1}{T_1(t)} = \frac{1}{T_{10}} + r_1 C_t(t), \quad (6.5)$$

where T_1 is the (new) longitudinal relaxation rate of the tissue, T_{10} is the inherent (pre-injection) longitudinal relaxation rate of the tissue, and r_1 is the relaxivity induced




Figure 6.2: Modeled contrast agent concentration in the aorta ($C_p^{Aorta}(t)$), blood plasma ($C_p(t)$), and tissue of interest ($C_t(t)$).




Figure 6.3: Transport function, $h_p(t)$, that models the dispersion in concentration of contrast agent as it travels through the blood stream from the injection site to the blood plasma near the tissue of interest.




Figure 6.4: Tissue transfer function, $h_t(t)$. This transfer function characterizes the flow of contrast agent between the blood plasma and the tissue of interest. This transfer function is based on the two-compartment model and incorporates the kinetic parameters K^{trans} and k_{ep} .

by the contrast agent

[13]. In addition to varying by contrast agent, the relaxivity also depends on the main field strength, B_0 . For our simulations we used $r_1 = 4.5 \text{ mMol}^{-1} \text{ sec}^{-1}$, which is based on the findings of Sasaki for studies of Gd at 3T [68]. Average values for T_{10} of breast tissue at 3T were found by Rakow-Penner to be 366.78 msec for fat tissue and 1444.83 msec for glandular tissue [69]. We used the T_{10} for glandular tissue in our simulations.

The relationship between T_1 and the MR signal, $\hat{\rho}$, is given by

$$\hat{\rho}(\theta, T_E) = \rho_0 \sin \theta \frac{(1 - E_1)}{(1 - E_1 \cos \theta)} e^{-T_E/T_2^*},$$

where T_1 is embedded in the variable E_1 :

$$E_1 \equiv e^{-T_R/T_1} \tag{6.6}$$

[9, p. 455]. Here ρ_0 is the voxel spin density, and θ is the flip angle. For our studies, we assume uniform voxel spin density and use a flip angle of $\theta = 10$ degrees and

$T_R = 4.6$ msec. Both the flip angle and TR were chosen to match those used in clinical MR breast studies at our institution.

We have now established all of the equations we need to relate tissue specific kinetic parameters to the enhancement we see in an MR image. This enables us to create realistic simulations of DCE-MRI based on chosen true kinetic parameter values, and in Chapter VII we do just that. The following section discusses the process of deriving kinetic parameter estimates from a reconstructed dynamic image sequence.

6.2 Estimating Kinetic Parameters from an Image Sequence

Ultimately we want to estimate these kinetic parameters from a reconstructed image sequence, so they can be of diagnostic value. This basically involves working backwards through the series of equations relating k_{ep} and K^{trans} to signal intensity, which are given in the previous section.

6.2.1 Computing Contrast Agent Concentration

Assuming a static, pre-enhancement scan is performed, and explicitly including the parameters' dependence on spatial location \vec{r} , we can write the magnitude of the transverse magnetization for this scan as

$$\begin{aligned} m_0(\vec{r}, \theta) &= \rho_0(\vec{r}) \sin \theta \exp(-T_E/T_2^*(\vec{r})) \frac{1 - E_{10}(\vec{r})}{1 - E_{10}(\vec{r}) \cos \theta} \\ &= \rho(\vec{r}, \theta) \frac{1 - E_{10}(\vec{r})}{1 - E_{10}(\vec{r}) \cos \theta}, \end{aligned} \quad (6.7)$$

where

$$\rho(\vec{r}, \theta) \equiv \rho_0(\vec{r}) \sin \theta \exp(-T_E/T_2^*(\vec{r})),$$

and

$$E_{10}(\vec{r}) \equiv e^{-T_R/T_{10}(\vec{r})}.$$

Because we know the tip angle θ and TR used in the acquisition, and assuming $T_{10}(\vec{r})$ is known, we can easily compute $\rho(\vec{r}, \theta)$ from the pre-enhancement image, $m_0(\vec{r}, \theta)$. Assumption of a known $T_{10}(\vec{r})$ is reasonable since in practice one could acquire a T1 map prior to the dynamic scan.

During the dynamic portion of the scan, the image is of course dependent on time,

so we rewrite Eq. (6.7) to explicitly include this dependence:

$$m(\vec{r}, t, \theta) = \rho(\vec{r}, \theta) \frac{1 - E_1(\vec{r}, t)}{1 - E_1(\vec{r}, t) \cos \theta}. \quad (6.8)$$

Solving for $E_1(\vec{r}, t)$ we get

$$E_1(\vec{r}, t) = \frac{m(\vec{r}, t, \theta)/\rho(\vec{r}, \theta) - 1}{m(\vec{r}, t, \theta) \cos \theta / \rho(\vec{r}, \theta) - 1}, \quad (6.9)$$

where $\rho(\vec{r}, \theta)$ is known from the pre-enhancement scan.

Next, we compute the estimated tissue concentration of the contrast agent by combining and rearranging Eqs. (6.6) and (6.5):

$$C_t(\vec{r}, t) = -\frac{1}{r_1 T_R} \ln(E_1(\vec{r}, t) e^{T_R/T_{10}(\vec{r})}). \quad (6.10)$$

Theoretically, we can now compute the CA tissue concentration for any location within the object, \vec{r} at any time t . This would require a continuous representation of the object in both space and time, which of course we don't have. In practice our reconstructed image sequences are discrete in both space and time, so Eq. (6.10) actually enables us to estimate the tissue concentration only at a discrete set of locations and a discrete set of time points. Namely, we can estimate the tissue concentration at any pixel location in any time frame.

6.2.2 Concentration to Kinetic Parameter Estimates

Once we have estimated $C_t(t)$, we then need to transform that estimate into estimates of the clinically relevant kinetic parameters, K^{trans} , k_{ep} , and v_e . The relationship between the kinetic parameters of a tissue and CA concentration within that tissue was established by the two-compartment model discussed in Section 6.1.1 (see Figure 6.1 and Eq. (6.4)). To determine K^{trans} and k_{ep} from our estimate of $C_t(t)$, we employ the method of Variable Projection, which is discussed in Section 6.3.

6.3 Variable Projection

The final step in determining the kinetic parameters K^{trans} , k_{ep} , and v_e , based on a reconstructed image sequence uses the known (or estimated) contrast agent concentrations, $C_p(t)$ and $C_t(t)$, and the two-compartment model pictured in Figure 6.1 and represented by the convolution relationship in Eq. (6.4). Note that for simula-

tions, we know $C_p(t)$ based on Eqs. (6.1), (6.3) and (6.2). For images reconstructed from real data, $C_p(t)$ can be determined from pixels in the dynamic image sequence corresponding to arterial blood vessels. To solve Eq. (6.4) for K^{trans} and k_{ep} , we use the variable projection (VP) method. The current section discusses the method of variable projection for the general case, and Section 6.3.1 covers our implementation of VP to estimate K^{trans} and k_{ep} .

Variable Projection is a well-established method of solving separable nonlinear least squares problems. In these problems, the measurement model is a linear combination of non-linear functions. Golub and Pereyra provide a nice review in [70], summarizing “the method of variable projections eliminates the linear variables obtaining a somewhat more complicated function that involves only the non-linear parameters.” One solves this more complicated function, obtaining the optimal estimates of the nonlinear parameters, and then uses those estimates to calculate the linear parameters.

The general model for a separable nonlinear least squares problem can be written as:

$$y_i = \sum_{j=1}^n x_j a_j(\boldsymbol{\theta}; t_i) + \varepsilon_i, \quad (6.11)$$

where y_i are the noisy observations, which are modeled as the linear combination of n basis functions, $a_j(\cdot)$, weighted by linear coefficients, x_j , and ε_i is noise. The basis functions depend on parameter vector $\boldsymbol{\theta}$ as well as t_i , which are independent variables associated with each observation (e.g., samples in time). In this model, the form of the nonlinear functions $a_j(\cdot)$ is known, and the goal is to estimate the n unknown linear parameters, x_j , as well as the unknown nonlinear parameters, $\boldsymbol{\theta}$, from the observations, y_i .

Given m observations, we can stack the measurements y_i into a vector, $\mathbf{y} = [y_1, \dots, y_m]^T$, and write the model in matrix-vector form as:

$$\mathbf{y} = \mathbf{A}(\boldsymbol{\theta})\mathbf{x} + \boldsymbol{\varepsilon}, \quad (6.12)$$

where $\mathbf{A}(\boldsymbol{\theta})$ is a matrix with $[\mathbf{A}(\boldsymbol{\theta})]_{ij} = a_j(\boldsymbol{\theta}, t_i)$, \mathbf{x} is the vector of linear parameters, $[x_1, \dots, x_n]^T$, and $\boldsymbol{\varepsilon} = [\varepsilon_1, \dots, \varepsilon_m]^T$ is the noise vector.

The residual vector is

$$\mathbf{r}(\mathbf{x}, \boldsymbol{\theta}) = \mathbf{y} - \mathbf{A}(\boldsymbol{\theta})\mathbf{x}, \quad (6.13)$$

and we want to determine the linear parameters \mathbf{x} and nonlinear parameters $\boldsymbol{\theta}$ by

finding the arguments that minimize the norm of the residual:

$$\begin{aligned} (\mathbf{x}, \boldsymbol{\theta}) &= \frac{1}{2} \operatorname{argmin}_{\mathbf{x}, \boldsymbol{\theta}} \|\mathbf{r}(\mathbf{x}, \boldsymbol{\theta})\|_2^2 \\ &= \frac{1}{2} \operatorname{argmin}_{\mathbf{x}, \boldsymbol{\theta}} \|\mathbf{y} - \mathbf{A}(\boldsymbol{\theta})\mathbf{x}\|_2^2. \end{aligned} \quad (6.14)$$

Minimizing the residual as represented in Eq. (6.14) can be difficult however, because there are multiple unknown parameters ($\boldsymbol{\theta}$ and \mathbf{x}). Variable projection “simplifies” the expression by eliminating the linear parameters in Eq. (6.14), thereby reducing the expression to one that includes only the nonlinear parameters. Because the form of the basis functions $a_j(\cdot)$ is known, note that if we knew the nonlinear parameters $\boldsymbol{\theta}$, then we would also know $\mathbf{A}(\boldsymbol{\theta})$, and we could easily solve for the linear parameters \mathbf{x} in terms of the pseudo-inverse of $\mathbf{A}(\boldsymbol{\theta})$:

$$\mathbf{x} = \mathbf{A}(\boldsymbol{\theta})^+ \mathbf{y}. \quad (6.15)$$

Taking this expression for the linear parameters and plugging into the expression for the residual in Eq. (6.14), we can rewrite the minimization problem as

$$\boldsymbol{\theta} = \operatorname{argmin}_{\boldsymbol{\theta}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}(\boldsymbol{\theta})\mathbf{A}(\boldsymbol{\theta})^+ \mathbf{y}\|_2^2. \quad (6.16)$$

Although this functional expression is somewhat more complicated than the form in Eq. (6.14), it now involves fewer parameters (only the nonlinear ones), and as Golub and Pereyra discuss, this reduction in the dimension of the parameter space ultimately makes the problem more tractable [70].

The method of variable projection is basically composed of first solving this modified minimization problem to find the optimal estimates of the nonlinear parameters and then using those estimates to find the optimal estimates of the linear parameters, based on the relationship established in Eq. (6.15). To see the origin of the naming of the variable projection method, we can rewrite Eq. (6.16) as

$$\boldsymbol{\theta} = \operatorname{argmin}_{\boldsymbol{\theta}} \frac{1}{2} \|(\mathbf{I} - \mathbf{A}(\boldsymbol{\theta})\mathbf{A}(\boldsymbol{\theta})^+) \mathbf{y}\|_2^2. \quad (6.17)$$

In this form, the matrix in parenthesis, which operates on \mathbf{y} , is the *projector* of the orthogonal complement of the column space of $\mathbf{A}(\boldsymbol{\theta})$ and the expression inside the norm is called the variable projection of \mathbf{y} .

There are a variety of methods for solving Eq. (6.16) (or equivalently Eq. (6.17)),

some of which are better suited for certain types of problems and many of which are discussed in [70]. As we will see in Section 6.3.1, when variable projection is applied to the problem of estimating kinetic parameters from tissue CA concentrations, the form of Eq. (6.16) simplifies greatly and no fancy minimization techniques are required.

6.3.1 Using VP to Estimate Kinetic Parameters

Variable Projection can be used in the process of estimating kinetic parameters from a dynamic MR image sequence. Section 6.2 discusses the process of estimating the tissue concentration of contrast agent from a such a sequence, and the current section will discuss the subsequent step of estimating the kinetic parameters K^{trans} and k_{ep} from tissue concentrations, thereby completing the discussion on estimating kinetic parameters from a dynamic image sequence.

Our current goal is to estimate the kinetic parameters K^{trans} and k_{ep} from computed estimates of the tissue concentration, $C_t(t)$. The tissue concentration depends on a variety of physiological factors, which we represent using kinetic parameters, and also depends on the plasma concentration of the contrast agent. The mathematical relationship between plasma and tissue concentrations is given in Eq. (6.4) and rewritten here for convenience:

$$C_t(t) = K^{\text{trans}} C_p(t) * e^{(-k_{\text{ep}}t)}.$$

Recall that for simulations, $C_p(t)$ is known, and for real images, the value of $C_p(t)$ can be determined from image pixels corresponding to arterial blood vessels.

Looking at this convolution relationship, we see that the tissue concentration is a function of a single linear parameter, K^{trans} , and a single nonlinear parameter, k_{ep} . That is, Eq. (6.4) is simply an instantiation of Eq. (6.11) where $n = 1$, and

$$\begin{aligned} y_i &= C_t(t_i), \\ \mathbf{x} &= K^{\text{trans}}, \\ \boldsymbol{\theta} &= k_{\text{ep}}, \\ a(\boldsymbol{\theta}; t_i) &= C_p(t_i) * \exp(-k_{\text{ep}}t_i). \end{aligned}$$

Again, we can stack the m measurements into a vector, $\underline{C}_t = [C_t(t_1), \dots, C_t(t_m)]^T$, and rewrite the model in vector form as:

$$\underline{C}_t = \mathbf{a}(k_{\text{ep}}) K^{\text{trans}}, \tag{6.18}$$

where $\mathbf{a}(k_{\text{ep}}) = [a(k_{\text{ep}}, t_1), \dots, a(k_{\text{ep}}, t_m)]^T$. Note that because we have only a single basis function, the matrix \mathbf{A} from Eq. (6.12) has been reduced in dimension and is now simply a vector, $\mathbf{a}(\cdot)$.

Continuing to follow the VP formulation, we know that *if* we knew the nonlinear parameter, k_{ep} , we could quickly find an estimate for the linear parameter, \hat{K}^{trans} , by using Eq. (6.15). Note that because of the reduced dimensionality of our problem, we need not compute a pseudoinverse, but rather can rewrite Eq. (6.15) as:

$$\hat{K}^{\text{trans}} = \frac{\mathbf{a}(k_{\text{ep}})^T}{\mathbf{a}(k_{\text{ep}})^T \mathbf{a}(k_{\text{ep}})} \underline{C}_t. \quad (6.19)$$

Finally, plugging our parameters into Eq. (6.16), we can write:

$$\hat{k}_{\text{ep}} = \underset{k_{\text{ep}}}{\text{argmin}} \frac{1}{2} \left\| \underline{C}_t - \mathbf{a}(k_{\text{ep}}) \frac{\mathbf{a}(k_{\text{ep}})^T}{\mathbf{a}(k_{\text{ep}})^T \mathbf{a}(k_{\text{ep}})} \underline{C}_t \right\|_2^2. \quad (6.20)$$

Some simplification yields

$$\begin{aligned} \hat{k}_{\text{ep}} &= \underset{k_{\text{ep}}}{\text{argmin}} \left\| \underline{C}_t \right\|_2^2 - \frac{(\mathbf{a}(k_{\text{ep}})^T \underline{C}_t)^2}{\mathbf{a}(k_{\text{ep}})^T \mathbf{a}(k_{\text{ep}})} \\ &= \underset{k_{\text{ep}}}{\text{argmax}} \frac{(\mathbf{a}(k_{\text{ep}})^T \underline{C}_t)^2}{\mathbf{a}(k_{\text{ep}})^T \mathbf{a}(k_{\text{ep}})} \\ &= \underset{k_{\text{ep}}}{\text{argmax}} \left| \left(\frac{\mathbf{a}(k_{\text{ep}})}{\|\mathbf{a}(k_{\text{ep}})\|} \right)^T \underline{C}_t \right|^2, \end{aligned} \quad (6.21)$$

the last line of which indicates how we find our estimate \hat{k}_{ep} based on the contrast agent concentration in a tissue of interest. An estimate of k_{ep} can be found by searching in sufficiently fine increments over the range of physiologically feasible rate constants (k_{ep} 's), and choosing as the estimate the k_{ep} that maximizes the inner product with the measurements, \underline{C}_t , as indicated in Eq. (6.21). Finally, K^{trans} is calculated from this estimate of k_{ep} , according to the relationship established in Eq. (6.19), and v_e is simply the ratio $K^{\text{trans}}/k_{\text{ep}}$.

This completes our discussion on estimating the kinetic parameters, K^{trans} , k_{ep} , and v_e from a dynamic image sequence. Chapter VII presents simulation results comparing kinetic parameter estimates computed from conventional and TRUIR reconstructions.

CHAPTER VII

Simulation Study and Results

We performed a simulation study to examine the effects of phase encode acquisition trajectory and regularization parameter selection on TRUIR reconstructed image sequences of a dynamic digital breast phantom. This chapter provides the details of our experiment and presents our findings.

7.1 Experiment Setup

We simulated DCE-MR breast imaging with a dynamic, bilateral 2D digital phantom, representing a slice through the breast in the sagittal plane. In our simulations, this is the y-z plane at a single x coordinate. The phase encode directions are ky and kz, and we ignore the readout direction, kx.

7.1.1 Modeling the Dynamic Object

We modeled our dynamic object after breast tissue containing multiple lesions. We included 6 circular lesions of various sizes and kinetic properties. Three of the simulated lesions are quite large, with a radius of 10 pixels (pixel size $\approx 1 \text{ mm}^2$), and three are much smaller with a radius of 2 pixels. Within each size category, one lesion has slowly enhancing kinetics, one exhibits moderate enhancement, and one enhances rapidly, in order to model a range of tumor types. Figure 7.1 shows the true dynamic object at 1 minute, when the lesions are beginning to show enhancement. At this point in time, the two slowly enhancing lesions are just starting to become visible, the two moderately enhancing lesions show more contrast, while the two rapidly enhancing lesions are near their peak enhancement.

To compute the true dynamics of each lesion, we begin by assigning true kinetic parameter values to each lesion. Choosing appropriate values for K^{trans} , k_{ep} , and v_e

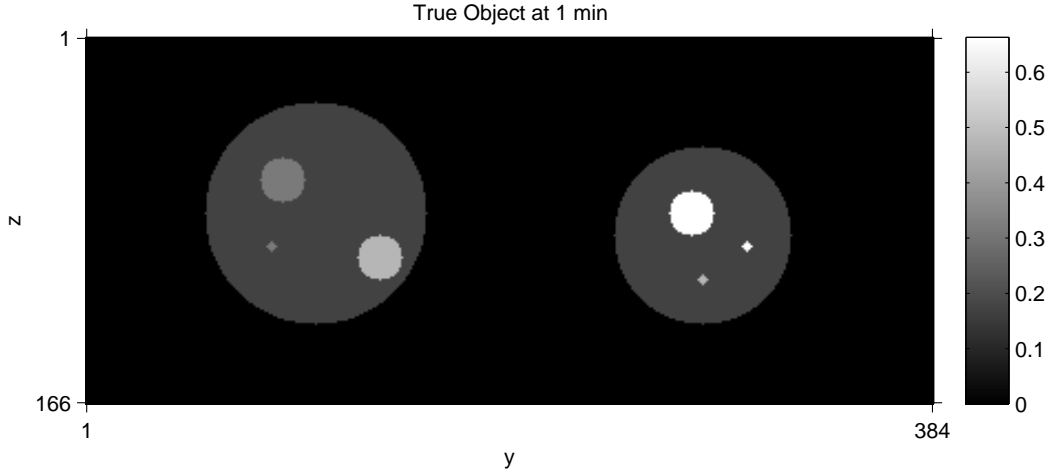


Figure 7.1: True dynamic object used in our simulations, shown at a single point in time (1 minute).

proved somewhat challenging, as reported kinetic parameter values for breast lesions vary quite a bit throughout the literature [21–23, 25, 26, 71]. For example, a 2008 study by Veltman et al. of 102 breast lesions found an average K^{trans} of 1.2 min^{-1} for benign lesions and 2.4 min^{-1} for malignant lesions, with corresponding average k_{ep} values of 3 min^{-1} for benign lesions and 3.8 min^{-1} for malignant ones [21]. However in 2004, Eliat et al.’s study of 100 patients found a mean K^{trans} of 0.14 min^{-1} and 0.35 min^{-1} , and a mean k_{ep} of 0.27 min^{-1} and 0.9 min^{-1} for benign and malignant lesions, respectively [22]. And, a 2005 study of 68 patients undergoing neoadjuvant chemotherapy found K^{trans} values as high as 10.6 min^{-1} in lesion hot spots [23]. We speculate that limitations in time/space resolution may have hampered consistent quantification of these parameters.

Our chosen true kinetic parameter values fall within the (wide) range of those reported in the literature, and correspond with values used by a previous PhD student to simulate enhancement of breast lesions [50]. The slowly enhancing lesion has $K^{\text{trans}} = 0.2 \text{ min}^{-1}$, $k_{\text{ep}} = 1.3 \text{ min}^{-1}$, and $v_e = 0.15$. The moderately enhancing lesion has $K^{\text{trans}} = 0.6 \text{ min}^{-1}$, $k_{\text{ep}} = 2.0 \text{ min}^{-1}$, and $v_e = 0.3$, while the rapidly enhancing lesion has $K^{\text{trans}} = 3.0 \text{ min}^{-1}$, $k_{\text{ep}} = 6.0 \text{ min}^{-1}$, and $v_e = 0.5$. The “healthy” tissue in our simulations does not exhibit any enhancement. Image size is based on the size of the real sensitivity maps (384 x 166 pixels), and the entire simulated object was chosen to fit within areas of good coil sensitivity.

We used Yang’s population averaged AIF to simulate the arterial concentration of injected Gd-DTPA with bolus arrival time, t_0 , of 0.5 min [67]. The series of relation-




Figure 7.2: True enhancement curves for the slow, moderate, and rapidly enhancing lesions in our simulated object.

ships detailed in Chapter VI, Section 6.1, determine the true object enhancement at each TR ($TR = 4.6$ msec), based on the true kinetic parameters. Figure 7.2 shows the true enhancement curves for the slow, moderate, and rapidly enhancing lesions. For simplicity, we disregarded any blood volume component when generating our enhancement curves. Tables 7.1 and 7.2 contain a summary of the variable values used in our simulations.

We use the object's true enhancement at each TR to generate the k-space data, one PE location at a time. As in an actual MR scan of a dynamic object, each PE sample in our simulations represents the dynamic object at a different point in time. As discussed in Chapters I and IV, in order to reconstruct a single image frame, it is necessary to group PE samples from a range in time, and we assign a single time point to this image frame (e.g., the midpoint of the group of PE sample times). This indicates a model mismatch between acquisition and reconstruction, and our simulations accurately represent that mismatch.

Variable	Value	Units	Long name
T_{10}	1444.83	msec	inherent T1 for glandular breast tissue
r_1	4.5	mMol ⁻¹ sec ⁻¹	relaxivity induced by Gd at 3T
TR	4.6	msec	repetition time for acquisition
θ	10	degrees	flip angle
t_0	0.5	min	bolus arrival time

Table 7.1: Values used in simulations.

Enhancement Type	$\mathbf{K}^{\text{trans}}$ min ⁻¹	\mathbf{k}_{ep} min ⁻¹	\mathbf{v}_e
Slow	0.2	1.3	0.15
Moderate	0.6	2.0	0.3
Rapid	3.0	6.0	0.5

Table 7.2: Kinetic Parameter values for simulated lesions.

7.1.2 Coil sensitivity

We used real sensitivity maps in our simulations. We computed these sensitivity maps from fully sampled (phantom) data collected using a 7-coil breast array on a Philips 3T scanner at the U of M hospital. We first reconstruct a (noisy) image for each coil using a standard IFFT of the fully sampled data, and then use a regularized estimation routine to compute the sensitivity map for each coil (mri_sensemap_denoise.m in [72]) [73].

A standard method to estimate coil sensitivities is to divide each fully sampled coil image by a fully sampled body coil image. Assuming the body coil has uniform spatial sensitivity, this image division will yield the sensitivity map for each coil. Sensitivity maps computed using this method, however, tend to be quite noisy and are also prone to gross inaccuracies in areas where there is very little signal in the body coil image, due to the classic divide by zero problem. We expect true sensitivity maps to vary smoothly in space, and the estimation routine that we use incorporates this expectation via a regularization term that encourages spatial smoothness.

The regularized estimation routine begins by estimating the coil sensitivities as the individual coil images divided by the body coil image, with an added feature that mitigates the divide by zero problem. Fully sampled data from the body coil was not available in this study, so we estimate the body coil image as the square root of the sum-of-squares of the 7 individual coil images, multiplied by the phase of the first coil image.

After the initial image division, pixel values in the resulting sensitivity map estimate are modified (or not) according to the corresponding pixel value in the body coil




Figure 7.3: Sensitivity maps used in simulations.

image. Sensitivity estimates at pixels corresponding to locations that have low signal in the body coil image are likely inaccurate, and therefore computed sensitivities at these locations are assigned a new value. Pixels that correspond to areas of decent signal in the body coil image are left unchanged. We use a minimum threshold for the acceptable body coil image value that is 5% of the maximum body coil image magnitude. Sensitivity values at locations that fall below this threshold are replaced with the median of the good sensitivity values. This is now the initial sensitivity estimate for the iterative regularized sensemap estimation routine. We used a smoothing parameter of $\beta = 2^6$ to compute our sensitivity maps, which are shown in Figure 7.3, for all 7 coils.

7.1.3 Acquired PE Trajectories

We generated three sets of k-space data, according to the three Phase Encode Trajectories discussed in Chapter IV. These three trajectories are showing in Figure 7.4 for reference. Each full trajectory samples the same 4766 PE locations, but the the order in which the samples are collected varies between trajectory. The Original Tra-




Figure 7.4: Three phase encode trajectories used in our simulations.

jectory is one that is currently used in breast MRI studies performed at the U of M hospital, and collects samples starting at low frequency locations and then moves to higher frequency locations as the acquisition continues. Reordered Trajectory 1 and Reordered Trajectory 2 both attempt to spread out the frequencies of sampled PE locations more evenly throughout the acquisition. Collection of 4766 PE locations per “fully-sampled” frame with our TR of 4.6 msec puts the nominal frame-rate for standard reconstruction at 21.9 seconds/frame. We generated a total of 12 full frames of data (for each PE trajectory), which corresponds to a total scan time of 4 mins and 23 seconds.

Note that, unlike the Keyhole method (and similar methods), our trajectories do not include a pre-injection reference frame that is fully sampled according to the Nyquist criteria. Nor do we acquire a fully-sampled reference frame at the end of the scan. Inclusion of initial and/or final reference frames can lead to desirable artifacts in the reconstructed image sequence if there is any motion between the reference frame(s) and the dynamic frames. Because our trajectories include no reference frames, we avoid any potential issues associated with motion relative to the reference frame(s).

7.1.4 Reconstruction Methods

We reconstructed dynamic image sequences from the simulated data using our proposed TRUIR method, as well as with a more traditional Homodyne+SENSE (HS) reconstruction for comparison [47]. We refer to the HS reconstruction method as such because it incorporates homodyning (to deal with the Partial Fourier acquisition) with SENSE (to deal with the 2x undersampling in k_y). A diagram of the HS reconstruction method is shown in Figure 7.5. Essentially the HS reconstruction is a homodyne reconstruction with SENSE unaliasing applied to the low resolution and full resolution images before the phase extraction and demodulation steps, respec-




Figure 7.5: Flow Chart of the Homodyne+SENSE reconstruction method. Steps corresponding to a traditional homodyne reconstructed are shown in blue, with the SENSE step shown in red.

tively.

For each PE trajectory, we reconstructed a 12-frame image sequence as well as a 24-frame image sequence using both TRUIR and HS reconstructions. The 12 and 24-frame image sequences are reconstructed from the exact same data set, i.e., there is no additional data available for the 24-frame reconstructions. Therefore the 24-frame image sequence represents the dynamic object over the same time range as the 12-frame image sequence. In doubling the number of dynamic frames, we are also doubling frame rate, which allows us to explore the flexibility of our TRUIR reconstruction method (as well as the HS method), and to evaluate the flexibility of each PE sample trajectory.

The TRUIR method naturally includes a flexible definition of frame, therefore reconstructing 12 vs 24 frames from a particular data set does not require any changes to the algorithm. To determine the number of data samples associated with each image frame, i.e., which samples comprise \mathbf{y}_m in the $\mathbf{y}_m - \mathbf{A}_m \mathbf{x}_m$ data-fit term of the cost function, TRUIR simply divides the total number of available samples by the desired number of reconstructed image frames. Therefore increasing the number of reconstructed frames from 12 to 24 simply means that the TRUIR method reduces the number of measurements associated with each frame by half, and relies on the regularization terms to counteract the otherwise expected reduction in image quality at each frame.

With HS reconstruction, it is quite a different story, as the HS method does not have this built-in flexibility. HS is a frame-by-frame reconstruction method, meaning that each frame is reconstructed independently from the other frames. If we halve the data available per frame by going to a 24-frame reconstruction, the standard HS method results in a dynamic image sequence with significant artifacts in each frame. Therefore, for 24-frame HS reconstructions we used a sliding window that divides the data into 24 overlapping groups, each containing a full frame’s worth of k-space samples, so that each HS image frame can be reconstructed from a “full” set of k-space samples.

Figure 7.6 illustrates how data is grouped into frames for both TRUIR and HS when reconstructing twice the number of frames as were acquired. The figure shows the case where a fully-sampled frame is comprised of 6 samples, and we have acquired 12 total samples, which is enough to reconstruct 2 fully-sampled (non-overlapping) image frames. If we want to reconstruct twice the number of frames, i.e., 4 image frames, from these 12 samples, TRUIR breaks the 12 samples into 4 even, non-overlapping groups, and thus associates 3 samples with each frame. The TRUIR grouping for the doubled number of image frames is shown in Figure 7.6 with the blue dashed lines. Despite the increase in frame rate, HS still needs a full frame’s worth of samples (here, 6 samples) to reconstruct each image frame. To reconstruct at the higher frame rate, HS uses a sliding window to divide the 12 total samples into 4 equally-spaced, overlapping groups, each containing 6 samples. Figure 7.6 shows the HS sliding window groupings with red dotted lines.

7.1.5 Assigning Frame Times

When we group multiple samples to reconstruct a single image frame, we are implicitly assuming that the grouped samples represent the same object. In particular, grouping samples for reconstruction in dynamic sequences assumes that the object is static over the time during which that group of samples is collected. Clearly this is not the case for dynamic objects; nonetheless it is necessary. The result is a reconstructed image frame that is an approximation of the dynamic object over the period of time during which the contributing k-space samples were collected. For our 12-frame reconstructions, that is a 22 second window.

In order to plot realistic enhancement curves, extract kinetic parameters, as well as to do any numerical comparison with the true object, we must assign the image frame to a single time point in that 22 second window. For most of our reconstructed image sequences, we use the mid-frame timepoint as the assigned frame time, with




Figure 7.6: Example of how TRUIR and HS group the available data when reconstructing at an increased frame rate. Here, 6 measurements comprise a fully-sampled frame, and we have collected two full frames of data (black lines). To reconstruct twice the number of image frames (four frames instead of two), TRUIR groups the data into four non-overlapping groups (blue dashed lines), while HS groups the data into four equally spaced, overlapping groups using a sliding window (red dotted lines).

the exception of reconstructions of samples from the Original Trajectory. Because the Original Trajectory samples all of the low frequencies first, and these samples contain most of the information about the object, we thought it would be more accurate to assign reconstructed frames based on the Original Trajectory to the beginning of each frame’s sample time window. The Original Trajectory’s low-frequency first sampling does not hold for a 24-frame sliding window, therefore 24-frame HS reconstructions use the mid-frame timepoint.

To summarize, we use the midframe timepoint for 12 and 24 frame TRUIR and HS reconstructions of samples from Reordered Trajectories 1 and 2, as well as for 24-frame HS reconstructions of the Original Trajectory. We use the beginning of the frame timepoint for 12 and 24 frame TRUIR reconstructions of the Original Trajectory, and for 12 frame HS reconstructions of the Original Trajectory. Note that the midframe times for 24-frame HS reconstructions differ from midframe times for 24-frame TRUIR reconstructions, due to the sliding window.

7.1.6 Regularization Parameters

As detailed in Chapter V, the spatial and temporal regularization parameters we use greatly affect the resulting TRUIR image sequences. Because no analytical method currently exists for choosing these regularization parameters for multi-coil imaging, we tested a range of values for both the spatial regularization parameter,

β , and the temporal regularization parameter α , and evaluated the resulting image sequences. We tested 3 values for β , with $\log_2 \beta = [-8, 0, 6]$, and 11 values for α , with $\log_2 \alpha = [0 : 2 : 20]$, which yields a total of 33 combinations of regularization parameters, and therefore a total of 66 dynamic TRUIR reconstructed image sequences of each data set (33 12-frame reconstructions + 33 24-frame reconstructions).

7.2 Results

We generated 3 sets of simulated data according to 3 different PE trajectories. Each data set was reconstructed using 12 and 24-frame HS, and 12 and 24-frame TRUIR, and 33 combinations of regularization parameters were tested for each TRUIR reconstruction. This means that we have a total of 204 reconstructed image sequences, each of which is comprised of 12 or 24 frames. In terms of temporal evaluation, we have six lesions of interest present in each of our 204 reconstructed image sequences, which yields 1,224 enhancement curves. In terms of spatial evaluation, with 12 or 24 frames per reconstructed sequence, we have a total of 7,344 different image frames we could evaluate. With 6 lesions per frame, that’s 44,064 regions of interest. Clearly we have too many results to present them all, so we instead present a representative subset, along with observations on overall trends for each reconstruction approach, and phase encode sampling scheme.

To establish a visual picture of our reconstructed image sequences, which will be helpful to have in mind when examining various results, Figure 7.7 shows a subset of frames from three of the 24-frame reconstructions. Frames 3-6 are shown, during which the lesions are exhibiting dynamic uptake of the contrast agent. Frames from the 24-frame HS reconstruction of samples from Reordered Trajectory 1 are shown on the left. Frames from 24-frame TRUIR reconstructions of the Original Trajectory are shown in the middle, and 24-frame TRUIR reconstructions of Reordered Trajectory 1 are shown on the right. The pictured Original Trajectory TRUIR was reconstructed with $\log_2 \alpha = 0$, which clearly does not provide enough temporal regularization. The TRUIR reconstruction of Reordered Trajectory 1 used $\log_2 \alpha = 6$. In evaluating the results of this study, we did not observe much dependence of the TRUIR reconstructions on the value of β . Therefore all presented TRUIR results use $\log_2 \beta = 6$, unless otherwise noted.




Figure 7.7: Frames from 3 reconstructed dynamic image sequences. Frames 3-6 are shown. Left: 24-frame HS reconstruction of samples from Reordered Trajectory 1. Middle: 24-frame TRUIR reconstruction of samples from the Original Trajectory, with $\log_2 \alpha = 0$. Right: 24-frame TRUIR reconstruction of samples from Reordered Trajectory 1, with $\log_2 \alpha = 6$.

7.2.1 Small Rapidly Enhancing Lesion

We first present some detailed results based on reconstructions of the small, rapidly enhancing lesion. We chose to present results for this lesion over the others because it presents the most challenging case for reconstruction: its small size requires high spatial resolution, and its fast kinetics require high temporal resolution.

Figure 7.8 shows enhancement curves for the small, rapid lesion from TRUIR reconstructions with a range of values of α . The figure includes 12 and 24-frame reconstructions of the Original Trajectory (top row), and 12 and 24-frame reconstructions of Reordered Trajectory 1 (bottom row). The true enhancement curve as well as the HS reconstructed enhancement curve also appear in each plot. The enhancement curves in Figure 7.8 provide a nice visualization of the effect of the temporal regularization parameter, α , on the temporal dynamics of the reconstructed image sequence.

For the 12-frame reconstructions, $\log_2 \alpha$ values less than 12 all result in reasonably good enhancement curves, while using larger values of α on the 12-frame reconstructions ($\log_2 \alpha = 16$ and 20) results in oversmoothing and loss of the temporal dynamics. So, for 12-frame reconstructions, the enhancement curves indicate an upper limit for

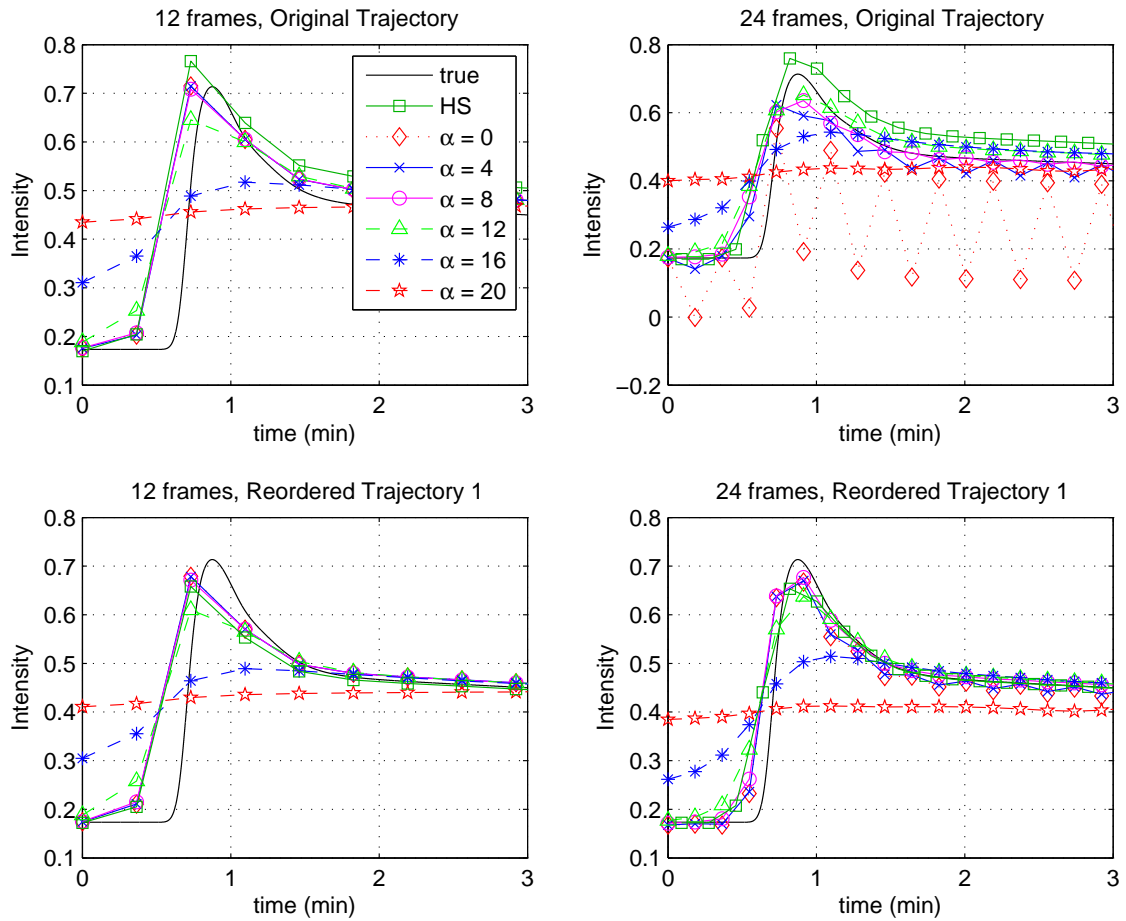


Figure 7.8: Effect of α on enhancement curves of the small, rapid lesion for different PE orderings.

our choice of α , but not a lower one (within the range of values we tested).

For the 24-frame TRUIR reconstructions, each image frame is associated with undersampled data, and we rely on the temporal regularizer to provide a degree of connectivity between adjacent image frames. The enhancement curves for the 24-frame TRUIR reconstructions of the Original Trajectory (top right, Figure 7.8) highlight the importance of choosing an appropriate α in this case. As with the 12-frame reconstructions, using a too large α ($\log_2 \alpha = 16$ and 20) results in oversmoothing and loss of temporal dynamics. However, for 24-frame reconstructions of the Original Trajectory, we also clearly see the negative effect of choosing an α that is too small, i.e., undersmoothing in time. For $\log_2 \alpha = 0$, which was the smallest value in our test range, the enhancement curve looks terrible, with wild oscillations between one frame and the next. The disjoint distribution of sample locations for even and odd frames in the 24-frame TRUIR reconstruction of the Original Trajectory means we have to more heavily enforce the sharing of information between image frames, i.e., we need to use a sufficiently large value for α . So, for 24-frame reconstructions of the Original Trajectory, we see both an upper and lower bound on effective values for the temporal regularization parameter, with good values in the range $4 < \log_2 \alpha < 16$. Based on the results of further investigation of the temporal behavior of TRUIR reconstructed image sequences, some of which will be presented later, we concluded that $\log_2 \alpha = 6$ works well in most cases, and therefore many of the presented results use this value. Note that in the 24-frame reconstructions based on Reordered Trajectory 1, which has more uniformly distributed sample locations for even and odd frames, the enhancement curves shown in the bottom right plot of Figure 7.8 do not indicate a minimum required value for α .

Figure 7.9 shows reconstructed enhancement curves for the small, rapidly enhancing lesion from HS reconstructions of each of the PE sample trajectories, along with the true enhancement of the lesion. Figure 7.10 shows corresponding enhancement curves for TRUIR reconstructions (with $\alpha = 6$ and $\beta = 6$). Enhancement curves for all six HS and TRUIR reconstructed image sequences are shown: 12 and 24-frame reconstructions of the Original, Reordered 1, and Reordered 2 Trajectories. We see immediately that all of the TRUIR enhancement curves for the small lesion more closely match the truth than any of the HS enhancement curves. All of the HS reconstructed enhancement curves overshoot the true enhancement. The 24-frame HS enhancement curves (dashed lines) show slightly better representation of the true curve (black line) than their corresponding 12-frame reconstructions (solid lines). The same frame rate observation holds for the TRUIR enhancement curves: the 24-frame




Figure 7.9: Enhancement curves of HS reconstructions of the small, rapidly enhancing lesion. Reconstructions using the Original Trajectory, Reordered Trajectory 1, Reordered Trajectory 2, and 12 or 24 frames are shown.

reconstructed curves are closer to the true enhancement than the 12-frame curves.

The reconstructed enhancement curves for both HS and TRUIR reconstructions in Figures 7.9 and 7.10 all show bias, most clearly seen in the plateau region of the enhancement curves. This bias is a result of the overshoot of the image estimate at the center of the lesion (from which the enhancement curves are derived), which can be seen in the profiles through the lesion in Figure 7.16. We see that the bias is present in both HS and TRUIR reconstructions of the small rapidly enhancing lesion, but is more pronounced in HS reconstructions. The bias in the enhancement curves is consistent across the small lesions (slow, moderate and rapidly enhancing), but is significantly reduced or eliminated altogether in reconstructions of the large lesions (not shown).

For both HS and TRUIR reconstructions, the enhancement curve corresponding to 12-frame reconstruction of the Original Trajectory shows a much poorer fit than the others. Remember that we assign reconstructed image frames to different timepoints, depending on the acquisition and reconstruction schemes. For 12-frame reconstructions of the Original Trajectory, we assign each frame to the time at which




Figure 7.10: Enhancement curves of TRUIR reconstructions of the small, rapidly enhancing lesion, using $\log_2 \alpha = 6$ and $\log_2 \beta = 6$. Reconstructions using the Original Trajectory, Reordered Trajectory 1, Reordered Trajectory 2, and 12 or 24 frames are shown.

we begin to acquire that frame’s samples (since we sample all the low frequencies first). The poor fit of the 12-frame Original Trajectory enhancement curves shown in Figures 7.9 and 7.10 suggests that using the midframe timepoint may be a better choice. However, for larger lesions the enhancement fit is quite good using the assigned beginning frame times (not shown), and changing the assigned frame times would negatively affect the fit of those curves. Additionally, we use the same “low-frequency first” logic in assigning beginning frame times to the 24-frame TRUIR reconstructions of the Original Trajectory, and, as Figure 7.10 shows, this results in a good fit for the resulting enhancement curve. Therefore, we kept the originally assigned beginning frame times for the 12-frame reconstructions of the Original Trajectory. Exploring the effect of using different assigned frame times for this trajectory (e.g. mid-frame times, or quarter-frame times) may be a subject of future work.

We have established that for all tested trajectories and reconstructed frame rates, TRUIR reconstructions show superior temporal dynamics to the corresponding HS reconstructions. Next we look at spatial resolution properties for both reconstruction methods.

Figures 7.11 and 7.12 show a single frame of the true and reconstructed small, rapidly enhancing lesion, from HS and TRUIR reconstructions, respectively. Within each figure, the top plot shows the true lesion (A), 12-frame reconstructions of the Original Trajectory, Reordered Trajectory 1, and Reordered Trajectory 2 (B-D), and 24-frame reconstructions of the same trajectories (F-H). The bottom half of the figure shows the normalized absolute error of each reconstruction.

Within the HS reconstructions (Figure 7.11), the images reconstructed from the 3 different trajectories with 2 frame rates look largely similar. The plots of normalized error show that the 12 and 24-frame HS reconstructions of the Original Trajectory have larger structured error just outside the lesion, while the 12 and 24-frame reconstructions of both reordered trajectories have more error distributed throughout the background tissue. The TRUIR reconstructions in Figure 7.12 show similar behavior.

Close examination of the reconstructed lesions shows that TRUIR reconstructions have worse spatial resolution, particularly in the y-direction (which is the x-axis in the figures), than HS reconstructions. This can be better seen in the plots of the profiles through the reconstructed lesions for HS and TRUIR reconstructions, which appear in Figures 7.13 and 7.14, respectively. The profiles show that all of the reconstructions, both HS and TRUIR, show some spatial blur in both y and z. This is not surprising given the extreme undersampling of the elliptically-shuttered trajectories we use, compared to full $n_y \times n_z$ sampling (see Figure 4.1). The HS reconstructions generally




Figure 7.11: HS reconstructions of the small, rapidly enhancing lesion and the normalized absolute error of each reconstruction. The frame with peak intensity is shown, which was frame 3 in 12-frame reconstructions, and frame 5 in 24-frame reconstructions. A: true, B: HS 12 orig, C: HS 12 reord 1, D: HS 12 reord 2, E: HS 24 orig, F: HS 24 reord 1, G: HS 24 reord 2.

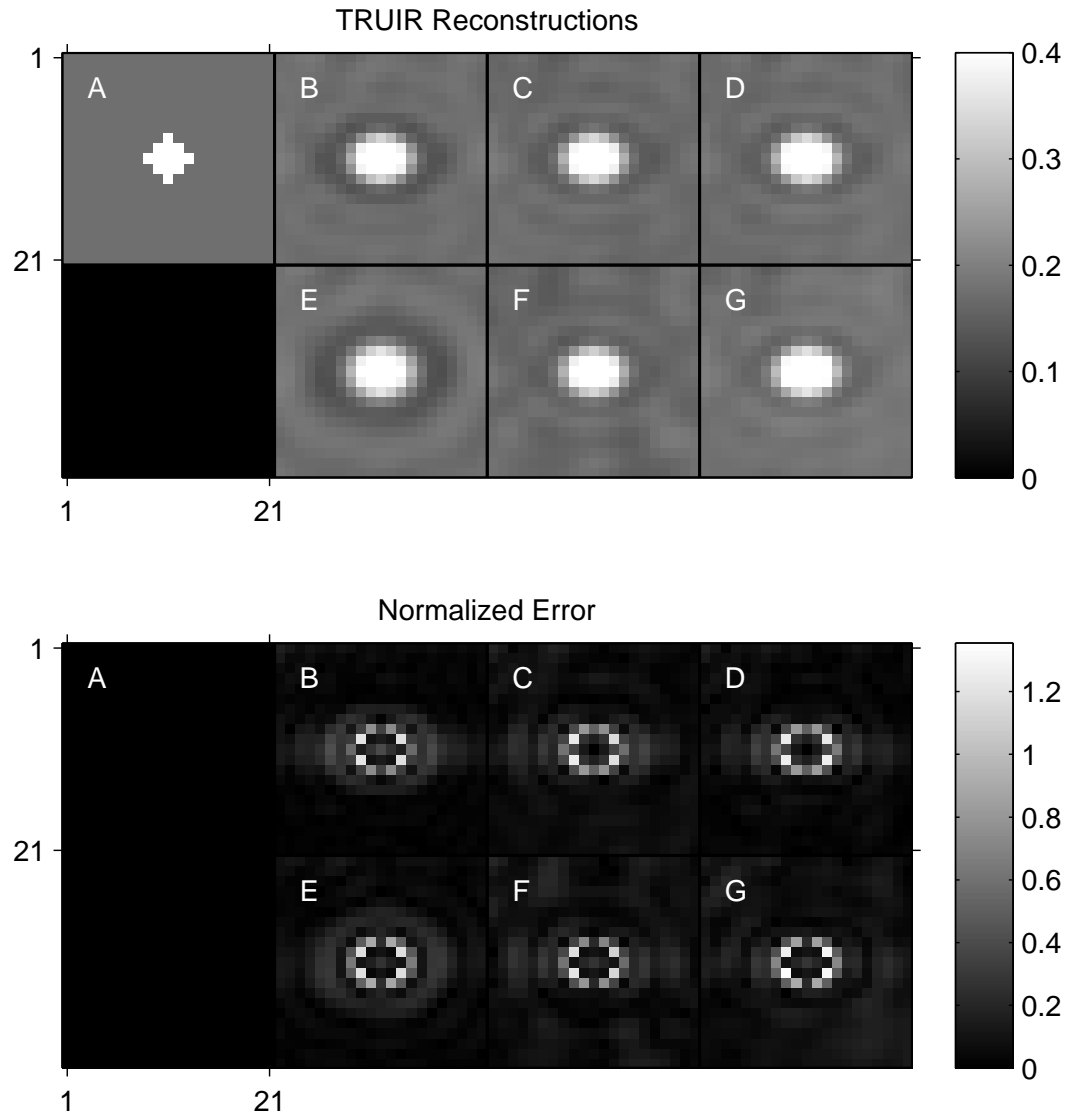


Figure 7.12: TRUIR reconstructions of the small, rapidly enhancing lesion and the normalized absolute error of each reconstruction. The frame with peak intensity is shown, which was frame 3 in 12-frame reconstructions, and frame 5 in 24-frame reconstructions. All shown TRUIR results use $\log_2 \alpha = 6$ and $\log_2 \beta = 6$ A: true, B: TRUIR 12 orig, C: TRUIR 12 reord 1, D: TRUIR 12 reord 2, E: TRUIR 24 orig, F: TRUIR 24 reord 1, G: TRUIR 24 reord 2.

overshoot the intensity at the center of the lesion, while the TRUIR reconstructions undershoot the lesion intensity, even at the lesion center in some cases.

Comparing the profiles in y from Figures 7.13 and 7.14, we see that the TRUIR profiles are wider than the HS profiles, which matches our previous observation of increased blur in the y -direction for TRUIR images in Figure 7.12. We hypothesized that the increased spatial resolution of HS reconstructions can be attributed to the homodyning part of the reconstruction process.

To test our hypothesis, we looked at images of the small, rapidly enhancing lesion from 12-frame reconstructions of the Original Trajectory, reconstructed using SENSE only, HS, TRUIR with very small regularization parameters ($\log_2 \alpha = 0$ and $\log_2 \beta = -8$), and TRUIR with $\log_2 \alpha = 6$ and $\log_2 \beta = 6$. We included a TRUIR reconstruction with the small regularization parameters to verify that the increased spatial blur seen in the TRUIR reconstructions of Figure 7.12 is not due to over-smoothing, i.e., too large regularization parameters.

The 3rd frame of the SENSE-only, HS, and TRUIR reconstructed image sequences appears in Figure 7.15, along with the normalized absolute error of each reconstruction. Profiles through the reconstructed lesions are shown in Figure 7.16. We see that the shape, normalized error and profile through the SENSE-only reconstructed lesion (B in Figure 7.15) closely match both of our TRUIR reconstructed lesions (D and E in Figure 7.15). Comparing the two TRUIR reconstructions to each other, we see essentially no difference. The TRUIR reconstruction in D uses the smallest tested values for the regularization parameters ($\log_2 \alpha = 0$ and $\log_2 \beta = -8$), while the TRUIR reconstruction in E uses moderate values for the regularization parameters ($\log_2 \alpha = 6$ and $\log_2 \beta = 6$). The fact that these two reconstructed image frames look almost identical confirms that the spatial blur we see in TRUIR is not a result of using overly large regularization parameters.

In comparison to the SENSE only and TRUIR reconstructions, the HS reconstructed lesion (C in Figure 7.15) shows less spatial blur in y and lower normalized error. This confirms our hypothesis that the superior spatial resolution of HS reconstructions is attributable to the homodyne part of the reconstruction algorithm (and not the SENSE part). The shape of the HS y -profile in Figure 7.16 also reflects the reduced spatial blur in y . The side lobes of the HS y -profile are closer in space to the main lobe than are the side lobes of the TRUIR y -profiles. In other words, we see tighter ripples in the y -direction for the HS reconstruction, which indicates that HS reconstructions contain more high spatial frequencies than our TRUIR reconstructions.

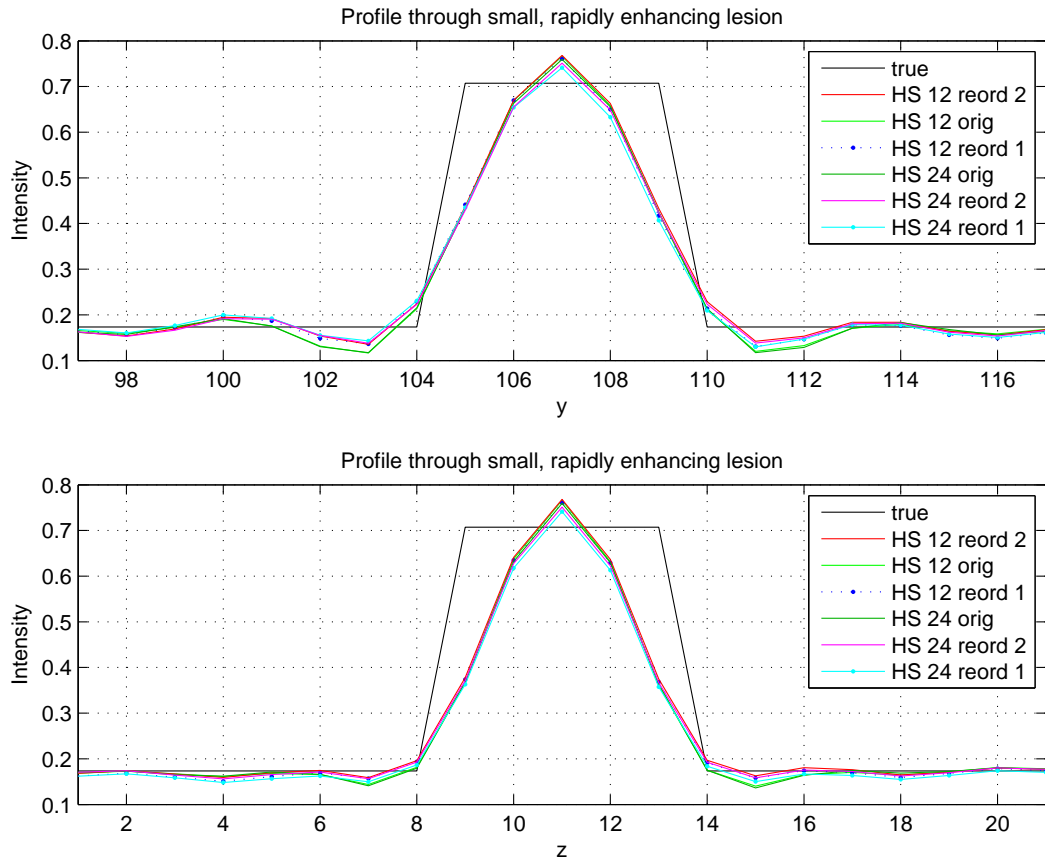


Figure 7.13: Profiles through HS reconstructions of the small, rapidly enhancing lesion shown in Fig. 7.11. Profile in the y-direction is shown in the top plot and profile in the z-direction is shown in the bottom one.

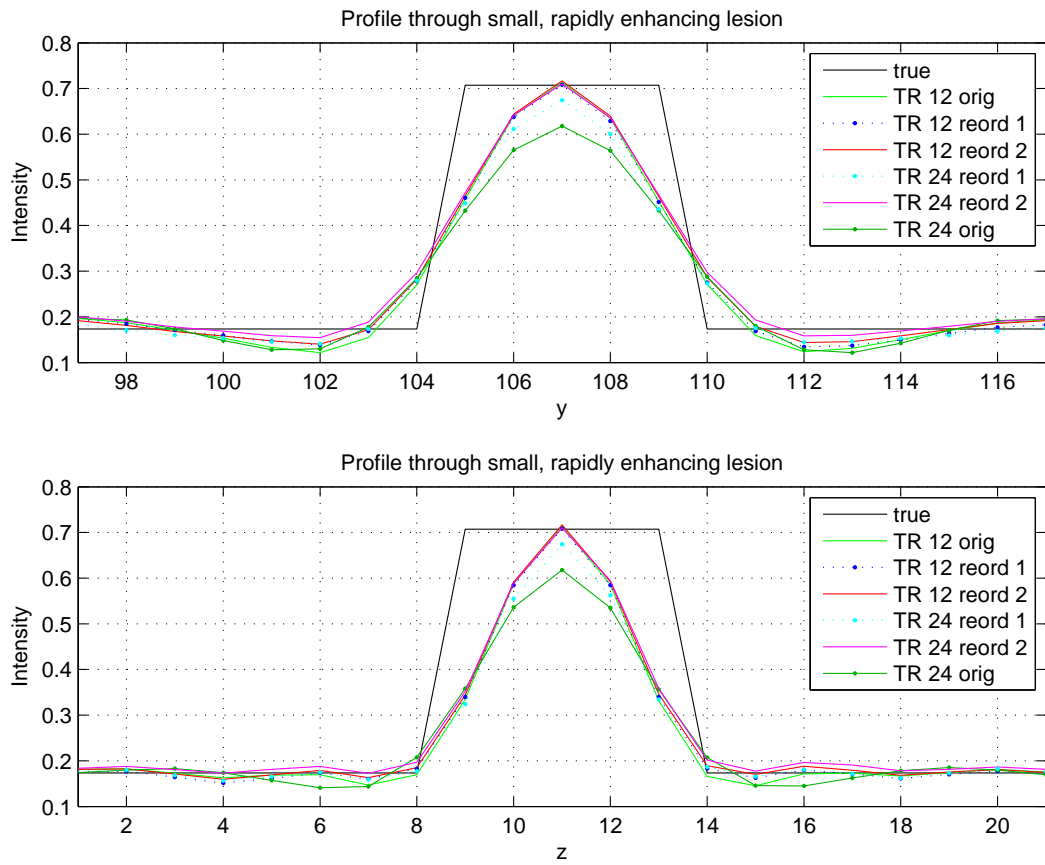


Figure 7.14: Profiles through TRUIR reconstructions of the small, rapidly enhancing lesion shown in Fig. 7.12. Profile in the y-direction is shown in the top plot and profile in the z-direction is shown in the bottom one.

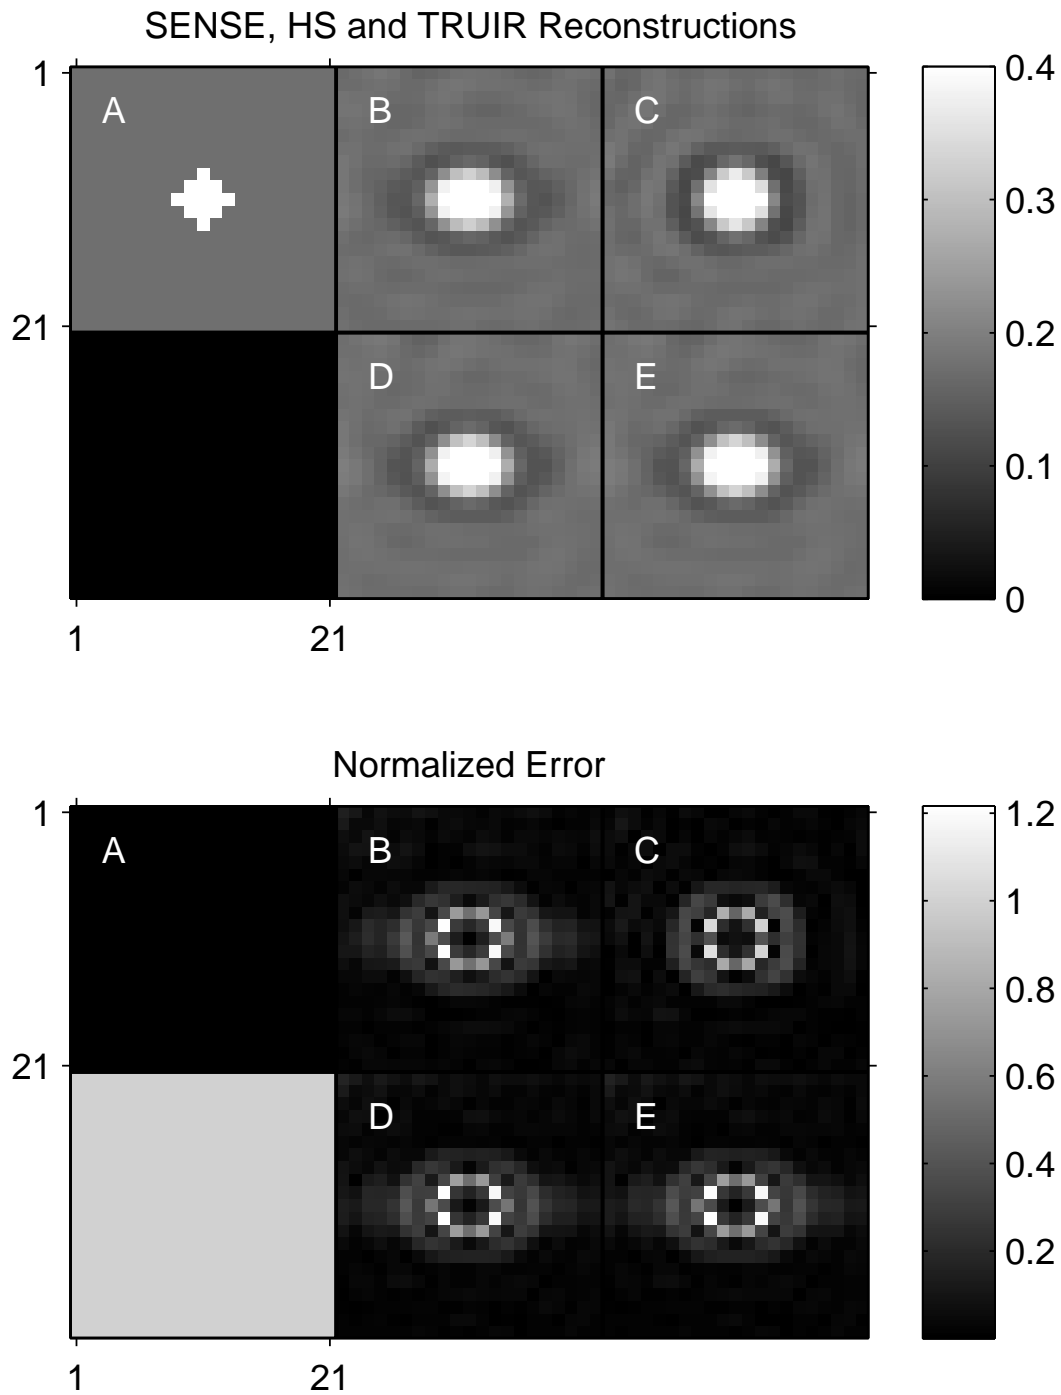


Figure 7.15: SENSE-only, HS, and TRUIR 12-frame reconstructions of the small, rapidly enhancing lesion. Reconstructions (top), and the normalized absolute error of each reconstruction (bottom). Frame 3 is shown. A: true, B: SENSE only, C: HS, D: TRUIR with $\log_2 \alpha = 0$, $\log_2 \beta = -8$, E: TRUIR with $\log_2 \alpha = 6$, $\log_2 \beta = 6$, H: .




Figure 7.16: Profiles through SENSE, HS, and TRUIR reconstructions of the small, rapidly enhancing lesion shown in Fig. 7.15. Profile in the y -direction is shown in the top plot and profile in the z -direction is shown in the bottom one.

Recall that all of our simulated trajectories are partial Fourier acquisitions, meaning the number of high frequency measurements is reduced by half in comparison to a traditional symmetric full Fourier acquisition. It is this reduction in high frequency data that results in reduced spatial resolution for our TRUIR reconstructions, in comparison to HS reconstructions. One of the first steps in homodyning (and HS reconstruction) is to run the partial Fourier data through a high pass filter that doubles the high frequencies, which sort of “makes up” for the other half of high frequency measurements that are missing in a partial Fourier acquisition. Our current TRUIR formulation includes nothing akin to this. Our data fidelity term fits the image sequence to the measured data as is - there is no initial doubling of high frequency components, despite the partial Fourier acquisition. Neither does our TRUIR formulation enforce a smooth phase assumption, which is also a key feature of homodyne reconstruction. Adjusting the TRUIR formulation to include some means of accounting for the reduced high frequency data in partial Fourier acquisitions will be an important area of future work. We believe that adding this adjustment will enable TRUIR reconstructions to approach or match the spatial resolution seen in HS reconstructions.

Figure 7.17 summarizes our results for the small, rapid lesion. The top plot shows the normalized enhancement error for each PE ordering and reconstructed frame rate, for both HS and TRUIR reconstructions. The values shown are the mean normalized absolute enhancement error, averaged over all 12 or 24 time frames of the specified reconstructed sequence. Again, the TRUIR results in this figure are for $\log_2 \alpha = 6$, $\log_2 \beta = 6$. This bar graph summarizes and quantifies the behavior of the HS and TRUIR enhancement curves from Figures 7.9 and 7.10. TRUIR reconstructions had significantly smaller enhancement error than HS reconstructions for all PE trajectories and reconstructed frame rates. For each trajectory, using TRUIR reconstruction reduces the normalized enhancement error about 5%, compared to HS reconstruction, sometimes more. For both HS and TRUIR reconstructions, we found that going to a 24-frame reconstruction reduced the enhancement error compared to the corresponding 12-frame reconstruction of the same PE trajectory. The improvement is slight for HS reconstructions, but quite pronounced for the 24-frame TRUIR reconstructions of both reordered trajectories.

The bottom plot in Figure 7.17 shows the normalized error within the 21 x 21 pixel ROI surrounding the lesion, for each PE ordering, reconstructed frame rate, and reconstruction method. The error shown is from a single frame of the reconstructed image sequence (the frame that exhibited peak enhancement), and is the mean of the

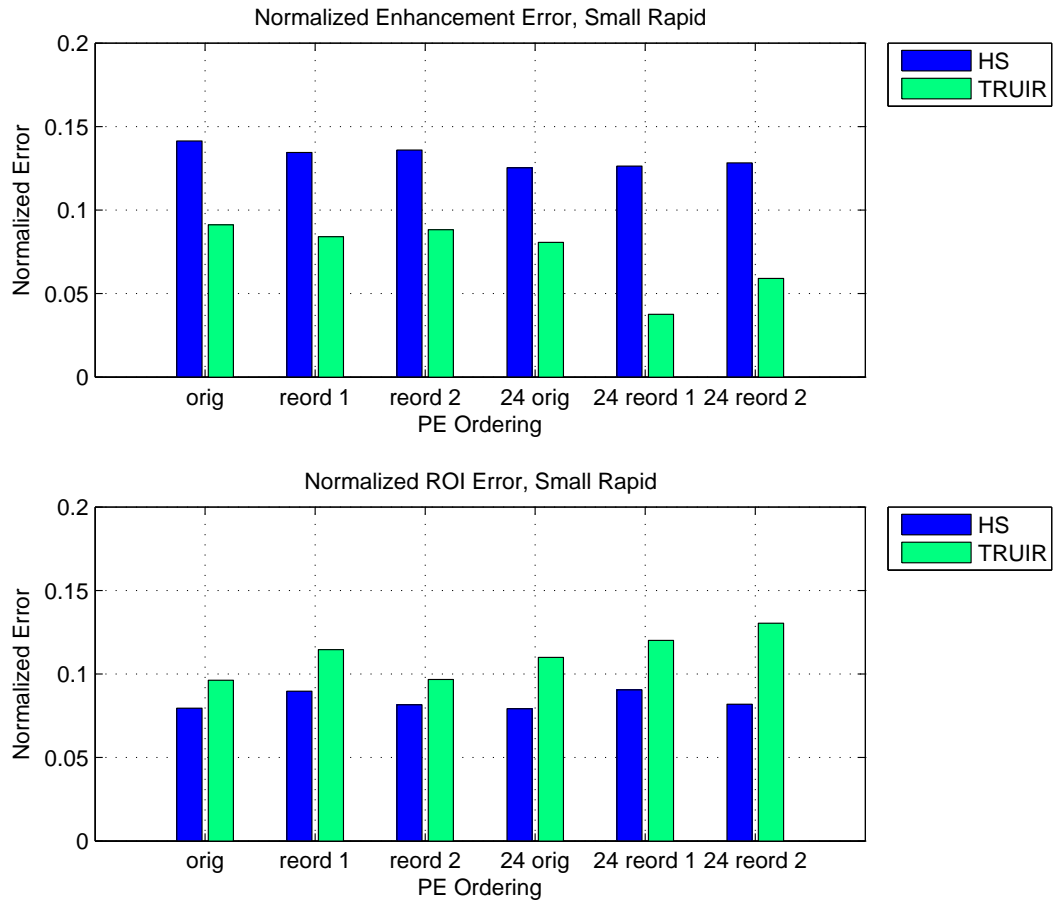


Figure 7.17: Top: Mean normalized absolute enhancement error for HS and TRUIR reconstructions (with $\log_2 \alpha = 6$, $\log_2 \beta = 6$), averaged over all 12 or 24 frames. Bottom: Mean normalized absolute error for an ROI surrounding the small, rapidly enhancing lesion.

normalized absolute error over all pixels within the ROI. We have already seen that the HS reconstructions have superior spatial resolution to our TRUIR reconstructions, and therefore it is no surprise that the normalized ROI error for HS reconstructions is lower than for TRUIR reconstructions, for all trajectories and frame rates. HS reconstructions reduce the normalized ROI error about 2.5%, compared to TRUIR reconstructions.

In summary, for the small, rapidly enhancing lesion, our TRUIR reconstructions offer reduced enhancement error (which suggests better temporal resolution), compared to HS reconstructions, but have increased ROI error (reduced spatial resolution). The increase in ROI error is smaller than the decrease in enhancement error, and, as previously discussed, we believe modification of the TRUIR formulation to better account for partial Fourier imaging would improve spatial resolution and reduce the ROI error for TRUIR reconstructions.

7.2.2 General Results

Now we look at more general results comparing HS and TRUIR for different PE trajectories and reconstructed frame rates. Figure 7.18 summarizes the temporal performance for HS reconstructions of each acquired PE trajectory. The figure shows the mean normalized enhancement error for each lesion within each reconstruction, as well as the mean error over all 6 lesions for each reconstruction (dark red). The mean normalized error for each lesion is the normalized absolute error at the lesion center, averaged over all frames. Not surprisingly, the 3 small lesions have significantly higher enhancement error than the 3 large lesions, for all PE orderings, with the small, rapidly enhancing lesion showing the worst enhancement error. For the large lesions, the 12-frame HS reconstruction of the Original Trajectory resulted in the lowest enhancement error, but the same reconstruction resulted for the highest enhancement error in the small lesions. Using either of the reordered trajectories reduced the mean enhancement error compared to the Original Trajectory, for both 12 and 24-frame HS reconstructions. The 24-frame HS reconstruction of Reordered Trajectory 1 had the lowest mean enhancement error over all HS reconstructions, although the mean error for all four reordered HS reconstructions is quite close.

Figure 7.19 shows the TRUIR enhancement error as a function of the temporal regularization parameter, α . All TRUIR reconstructions represented in this figure use $\log_2 \beta = 6$. The enhancement value shown for each reconstruction is the mean normalized absolute enhancement error, averaged over all 6 lesions and all frames. The enhancement for each lesion is measured at the center pixel of the lesion. The




Figure 7.18: Mean normalized absolute enhancement error for HS reconstructions. Normalized absolute error was calculated at the center of each lesion, and then averaged over all time frames to yield the values shown in the figure.




Figure 7.19: Enhancement error as a function of temporal regularization parameter, α . The x-axis shows the values of $\log_2 \alpha$ ('l2a').

mean enhancement error of a traditional 12-frame HS reconstruction of the Original Trajectory is included for reference, as is the enhancement error of the best HS reconstruction, which was a 24-frame reconstruction of Reordered Trajectory 1. Figure 7.19 shows that all TRUIR reconstructions beat even the best HS reconstruction, for a range of values of α . The 24-frame TRUIR reconstruction of the Original Trajectory is the most sensitive to the choice of temporal regularization parameter, giving huge errors for most values of α and only beating the best HS reconstruction for a single value of α ($\log_2 \alpha = 8$). TRUIR reconstructions based on the Reordered Trajectories exhibit more robustness to the choice of temporal regularization parameter, showing good performance for a large range of α values, as does the 12-frame TRUIR reconstruction of the Original Trajectory. All 12-frame TRUIR reconstructions show fairly constant enhancement error for $\log_2 \alpha = 0$ to 8, regardless of trajectory.

The 24-frame TRUIR reconstructions of the reordered trajectories show the overall best performance in terms of enhancement error, with the error minimized with a temporal regularization parameter in the range of $\log_2 \alpha = 6$ to 8. Somewhat surprisingly, we found that reconstructions of samples from Reordered Trajectory 1 had lower error than reconstructions of Reordered Trajectory 2, even though Reordered Trajectory 1 has more structure to the PE sample pattern, and the samples in Reordered Trajectory 2 look more uniformly distributed (see Figures 4.3, 4.5 and 4.8 in Chapter IV). The improvement in enhancement error for Reordered Trajectory 1 over Reordered Trajectory 2 is slight, but is consistent across reconstruction methods and frame rates. That is, Reordered Trajectory 1 had lower enhancement error than Reordered Trajectory 2 for 12 and 24-frame TRUIR reconstructions as well as for 12 and 24-frame HS reconstructions (see Figure 7.18). The lowest enhancement error over all methods was achieved with a 24-frame TRUIR reconstruction of Reordered Trajectory 1 using a temporal regularization parameter of $\log_2 \alpha = 6$.

Figure 7.20 summarizes the enhancement error in TRUIR reconstructions for each PE trajectory and each lesion. The presented results are for TRUIR reconstructions with $\log_2 \alpha = 6$ and $\log_2 \beta = 6$. As in Figure 7.18, the error shown is the mean normalized absolute enhancement error seen at the center of each lesion, averaged over all image frames. The mean enhancement error over all lesions is also shown for each PE ordering (dark red). As we saw in Figure 7.19, 24-frame TRUIR reconstruction of the Original Trajectory has the highest enhancement error, while 24-frame TRUIR reconstruction of Reordered Trajectory 1 has the lowest enhancement error. For the small lesions, 24-frame reconstruction of either Reordered Trajectory significantly reduces the enhancement error, compared to 12-frame reconstructions of any PE

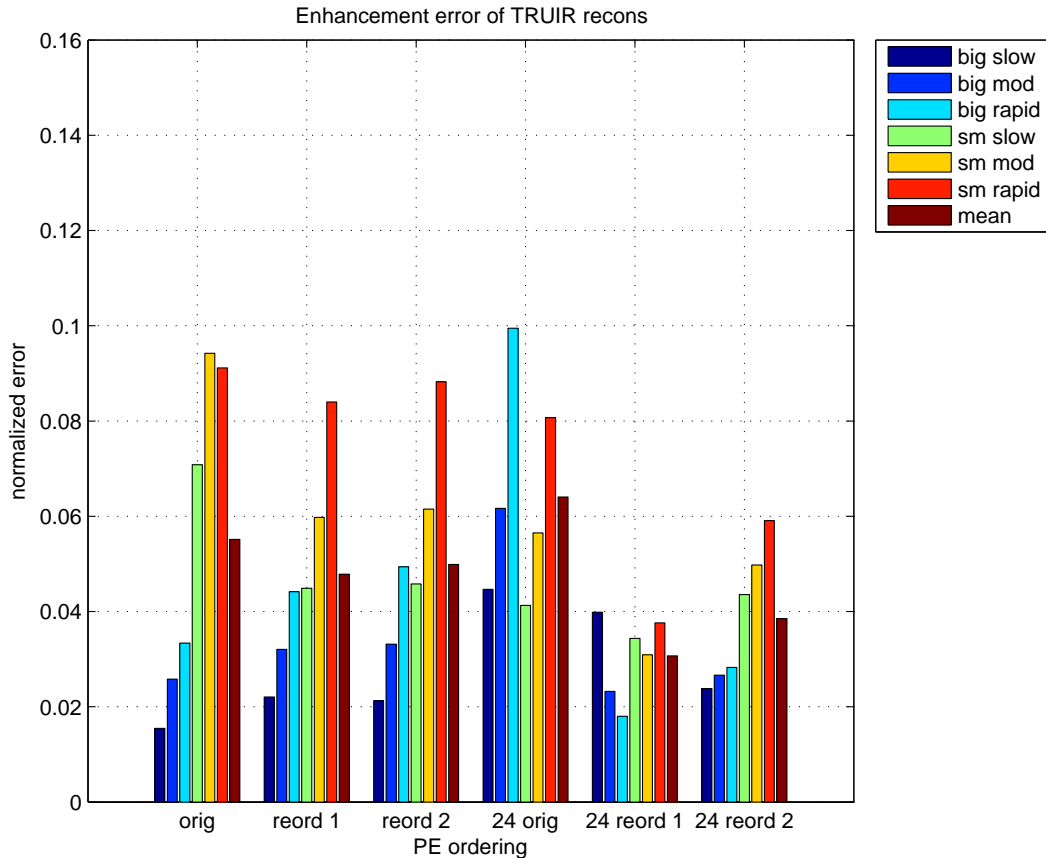


Figure 7.20: Mean normalized absolute enhancement error for TRUIR reconstructions with $\log_2 \alpha = 6$ and $\log_2 \beta = 6$. Normalized absolute error was calculated at the center of each lesion, and then averaged over all time frames to yield the values shown in the figure.

trajectory.

7.2.3 Kinetic Parameters

We computed the kinetic parameters K^{trans} , k_{ep} , and v_e for each lesion in each reconstructed image sequence, and compared these estimated values to the true kinetic parameter values. Figures 7.21 - 7.23 show the error in the three estimated kinetic parameters for TRUIR reconstructed image sequences as a function of α . All represented TRUIR reconstructions used $\log_2 \beta = 6$. The figures also include kinetic parameter estimates calculated from a traditional 12-frame HS reconstruction of the Original Trajectory for reference. The best HS kinetic parameter estimates are included as well. As with the enhancement error, the lowest error in kinetic parameter estimation over all HS reconstructions was achieved with a 24-frame HS reconstruction




Figure 7.21: K^{trans} estimation error as a function of α . Error shown for each reconstruction is the mean normalized absolute error of the K^{trans} estimate at the lesion centers, averaged over all lesions.

of Reordered Trajectory 1. In all three figures, represented results are from TRUIR reconstructions unless otherwise indicated in the legend. Similarly, 24-frame reconstructions are indicated in the legend, and all others are 12-frame reconstructions. The legend entry ‘reord’ refers to Reordered Trajectory 1.

Kinetic parameters were estimated based on the measured enhancement curve at the center of each lesion. The error shown in the figures is the mean normalized absolute error of the specified kinetic parameter, averaged over the 6 lesions in each reconstructed image sequence.

Figure 7.21 shows the K^{trans} estimation error, Figure 7.22 shows the k_{ep} estimation error, and Figure 7.23 shows the error in estimates of v_e . In Figure 7.21 and Figure 7.22, we see that all of the presented TRUIR reconstructions result in large improvements in K^{trans} and k_{ep} estimates, compared to the traditional 12-frame HS reconstruction of the Original Trajectory. Additionally, a variety of TRUIR reconstructions outperform even the best HS reconstruction in estimating each of the




Figure 7.22: k_{ep} estimation error as a function of α . Error shown for each reconstruction is the mean normalized absolute error of the k_{ep} estimate at the lesion centers, averaged over all lesions.




Figure 7.23: v_e estimation error as a function of α . Error shown for each reconstruction is the mean normalized absolute error of the v_e estimate at the lesion centers, averaged over all lesions.

three kinetic parameters. In Figure 7.21, 12 and 24-frame TRUIR reconstructions from both Reordered Trajectory 1 and Reordered Trajectory 2 have lower K^{trans} estimation error than the best HS reconstruction, for $\log_2 \alpha < 8$. Figure 7.22 shows that with $\log_2 \alpha < 6$, 24-frame TRUIR reconstructions of both Reordered Trajectories reduce the error in the estimate of k_{ep} , compared to the best HS reconstruction, although the difference is slight. In estimating v_e , Figure 7.23 shows that all TRUIR reconstructions (12 and 24 frames with all PE trajectories) with $\log_2 \alpha < 12$ provided better estimates than the best HS reconstruction, with the exception of the 24-frame TRUIR reconstruction of the Original Trajectory, which offers improved v_e estimates only for $\log_2 \alpha$ between 6 and 12.

Looking at all three figures, we observe that reconstructions based on the Reordered Trajectories generally provide better kinetic parameter estimates than reconstructions of the Original Trajectory, and 24-frame reconstructions (solid lines) result in better kinetic parameter estimates than their associated 12-frame reconstructions (dashed lines). As with enhancement error, using a temporal regularization parameter of $\log_2 \alpha \leq 8$ generally provides the best estimates of the kinetic parameters. Although Figure 7.19 shows that 24-frame TRUIR reconstructions of the Reordered Trajectories have reduced enhancement error around $\log_2 \alpha$ of 6 to 8, we do not see the same effect in the kinetic parameter estimates. For $\log_2 \alpha < 8$, all 3 kinetic parameters have fairly constant error within each Reordered Trajectory’s TRUIR reconstruction.

Using our proposed TRUIR method, reordered trajectories and higher frame-rate reconstructions, we were able to reduce the estimation error for all three kinetic parameters, compared to the traditional 12-frame HS reconstruction of the Original Trajectory. We see the greatest improvement in estimates of K^{trans} , and the smallest improvement in estimates of v_e . The smallest error for all three kinetic parameters was achieved with 24-frame TRUIR reconstructions of the Reordered Trajectories.

7.2.4 Summary

The main results of our simulation study can be summarized as follows:

- Our proposed Reordered Trajectories 1 and 2 are more suitable for reconstructing dynamic image sequences with an increased frame rate, compared to the Original Trajectory currently in use.
- Choosing the right value for the temporal regularization parameter, α , is very important when reconstructing data from highly ordered, poorly distributed sample trajectories, such as the Original Trajectory, at higher frame rates.

- Our Reordered Trajectories are more robust to the choice of α than the Original Trajectory. Both Reordered Trajectories produced TRUIR image sequences with good temporal resolution and kinetic parameter estimates, over a wide range of temporal regularization parameter values.
- TRUIR reconstructions offer better temporal dynamics than even the best HS reconstruction. In particular, 24-frame TRUIR reconstructions of data acquired according to Reordered Trajectory 1 or Reordered Trajectory 2 showed the best overall temporal dynamics (for $\log_2 \alpha < 8$), compared to all other reconstructed image sequences, and also resulted in the best kinetic parameter estimates. We would like to conclude from this that TRUIR offers better temporal resolution than HS, but, as discussed in Chapter V, Section 5.3.3, temporal resolution is somewhat of an elusive concept in this setting, since measuring a true impulse response is a bit of a nonsensical affair. However, our results certainly show that TRUIR reconstructions have better temporal dynamics than HS, which suggests that TRUIR may offer improved temporal resolution.
- TRUIR reconstructions (with quadratic regularization in space) of the partial Fourier data in this study exhibited poorer spatial resolution properties than HS reconstructions. The HS reconstruction method is designed for reconstruction of partial Fourier data, and includes measures that address the asymmetry of the data, such as doubling of high frequency components and utilizing a smooth phase assumption. Our current TRUIR formulation does not include any such compensation measures for partial Fourier acquisitions. We believe that modification of the TRUIR formulation to better accommodate partial Fourier data will improve the spatial resolution of TRUIR reconstructed image sequences, and this is an important area of future work. Additionally, use of edge preserving regularization in space instead of quadratic regularization will also improve spatial resolution of TRUIR images.
- In this study, the spatial resolution of our TRUIR reconstructions was dominated by the undersampled k-space as opposed to the spatial regularization. Figure 7.15 shows TRUIR reconstructions of the small lesion using two different sets of regularization parameters. Subfigure (D) is a TRUIR reconstruction using $\log_2 \alpha = 0$ and $\log_2 \beta = -8$, which are the smallest values within the tested range for the temporal and spatial regularization parameters. Subfigure (E) is a TRUIR reconstruction using $\log_2 \alpha = 6$ and $\log_2 \beta = 6$, which generally resulted in good quality reconstructions. Because the spatial resolutions seen

in (D) and (E) are not visibly different, this tells us that the dominant factor contributing to TRUIR's poor spatial resolution is the undersampled k-space trajectory, rather than oversmoothing in space. Therefore, it would definitely be worth exploring the effect of using an extended PE trajectory that samples locations farther out in k-space. This should be an area of future work.

CHAPTER VIII

Conclusions and Future Work

8.1 Summary

This dissertation presents an object-domain model for image reconstruction in dynamic MRI, which we refer to as TRUIR: Temporal Regularization Use in Image Reconstruction. Our reconstruction model is formulated as penalized likelihood estimator that explicitly includes a temporal smoothness assumption in object space. TRUIR is therefore well-suited for MR imaging of dynamic objects that vary smoothly in time, such as the objects of interest in dynamic contrast-enhanced MRI. We extended our TRUIR model to incorporate parallel imaging, and also accelerated TRUIR reconstructions by utilizing Toeplitz matrices.

This work explores various aspects of the proposed TRUIR method including selection of spatial and temporal regularization parameters, flexibility for increased frame rate reconstruction, and the effect of different phase encode acquisition patterns on TRUIR reconstructions.

Chapter V explores TRUIR’s resolution properties through evaluation of a local impulse response, and presents an accelerated method we developed to compute the local impulse response. Our DCE-MRI simulation results Chapter VII also included analysis of spatial and temporal resolution properties, and resolution effects of regularization parameter choice.

We presented 2 new phase encode sampling trajectories that are based on re-ordering of sample locations from a current clinical sampling scheme. Our Reordered Trajectories are designed to distribute sampling of low and high frequency locations more uniformly in time. These trajectories offer significantly more flexibility for reconstructing at higher frame rates, and reconstructions of our Reordered Trajectories were also relatively robust to regularization parameter choice, especially when compared to the current clinical trajectory.

In evaluating our proposed TRUIR method, we focused on the application of DCE-MRI in the characterization and assessment of breast cancer. We designed a multi-coil simulation study that models dynamic contrast agent uptake in the breast with 6 representative lesions. We compared TRUIR reconstructions to a more traditional frame-by-frame Homodyne + SENSE reconstruction, for a variety of acquired trajectories and reconstructed frame rates. We found that TRUIR provides improved temporal resolution dynamics in the reconstructed dynamic image sequence, compared to the traditional frame-by-frame reconstruction, and that TRUIR achieved the best temporal dynamics when using samples from the Reordered Trajectories and reconstructing at an increased frame rate. TRUIR reconstructions showed improved lesion enhancement curves, and resulted in better estimates of the kinetic parameters K^{trans} , k_{ep} , and v_e . Our DCE-MRI simulation confirmed our earlier assessment that the Reordered Trajectories are more robust to reconstructing at higher frame rates than the current clinical trajectory, and are also more robust to the choice of the temporal regularization parameter.

Our simulation results also showed the negative result that TRUIR reconstructions with quadratic spatial regularization exhibited worse spatial resolution than the frame-by-frame Homodyne+SENSE method for the tested PE sampling patterns. We attribute this reduction in spatial resolution to the partial Fourier nature of the PE sampling schemes in this study. Unlike HS reconstruction, our current TRUIR formulation does not include any compensation measures for partial Fourier data. We believe that modification of the TRUIR formulation to better accommodate partial Fourier acquisitions will improve spatial resolution and is an important area of future work.

8.2 Future Work

There are plenty of directions for interesting future work with TRUIR:

- **Modification to accommodate Partial Fourier data:** TRUIR reconstructions of partial Fourier data (with quadratic regularization) exhibited poorer spatial resolution than HS reconstructions, because homodyne reconstruction includes measures to address the asymmetry of partial Fourier data, while our current TRUIR formulation does not. Modification of the TRUIR formulation to include some means of accounting for the reduced high frequency data in partial Fourier acquisitions should be investigated.

- **Alternate temporal basis function:** In parameterizing the dynamic object, the measurement model we employ currently uses a rect basis function in time. However, because we assume the object varies smoothly in time, rect functions are likely not the best temporal basis functions. Future work could explore using smoother temporal basis functions, e.g., b-splines, in our measurement model. This would change the structure of the system matrix \mathbf{B} . The matrix would no longer be block diagonal, but would also have off-diagonal components, which correspond to overlapping portions of the temporal basis functions. The presence of these off-diagonal components may also be desirable in terms of the mixing of k-space data between frames.
- **Edge preserving penalty:** Our current TRUIR formulation uses a quadratic penalty in space. Using instead an edge preserving penalty in space may increase spatial resolution, and is certainly an area worth exploring.
- **Extended PE trajectory:** The phase encode trajectories studied so far are severely limited in their kspace coverage compared to full Nyquist sampling (see Figure 4.1), which results in reduced spatial resolution in our reconstructed image sequences, even when we reconstruct using a full trajectory’s worth of samples for each image frame. We are interested in implementing one or more ‘extended’ trajectories that sample frequencies farther out in kspace. The extended trajectory could also re-visit low frequencies more often than the current PE trajectories. Collection of higher frequency sample locations provides the potential for increased spatial resolution in the reconstructed images, but may come at the cost of reduced temporal resolution or other image artifacts. Future work will include evaluating TRUIR reconstructions of data from an extended trajectory, and, in particular, assessing the ability of these reconstructions to provide both high spatial and high temporal resolution (with the right regularization parameters).
- **Regularization parameter selection:** Although TRUIR reconstructions of well-distributed data sets (e.g. Reordered Trajectories 1 and 2) have been shown here to be relatively robust to the choice of the temporal regularization parameter, we have not established an analytical method for choosing the spatial and temporal regularization parameters. An analytical method is certainly desirable, but may not be feasible.
- **Motion Correction:** Because our TRUIR approach is formulated in the object

domain, it easily lends itself to incorporation of motion correction, which is also in the object domain.

- **Phantom and human studies:** Future validation of TRUIR with a dynamic phantom study is necessary, as is eventual validation with a human studies.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] C. K. Kuhl, P. Mielcareck, S. Klaschik, C. Leutner, E. Wardelmann, Jürgen Gieseke, and H. H. Schild, “Dynamic breast MR imaging: are signal intensity time course data useful for differential diagnosis of enhancing lesions?,” *Radiology*, vol. 211, no. 1, pp. 101–10, Apr. 1999.
- [2] A. Jemal, R. Siegel, E. Ward, T. Murray, J. Xu, C. Smigal, and M. J. Thun, “Cancer statistics, 2006,” *CA Cancer J Clin*, vol. 56, no. 2, pp. 106–130, Mar. 2006.
- [3] F. Sardanelli, G. M. Giuseppetti, P. Panizza, M. Bazzocchi, A. Fausto, G. Simonetti, V. Lattanzio, and A. Del Maschio, “Sensitivity of mri versus mammography for detecting foci of multifocal, multicentric breast cancer in fatty and dense breasts using the whole-breast pathologic examination as a gold standard,” *Am. J. Roentgenol.*, vol. 183, pp. 1149–1157, Oct. 2004.
- [4] L. W. Nunes, M. D. Schnall, E. S. Siegelman, C. P. Langlotz, S. G. Orel, D. Sullivan, L. A. Muenz, C. A. Reynolds, and M. H. Torosian, “Diagnostic performance characteristics of architectural features revealed by high spatial-resolution MR imaging of the breast,” *Am. J. Roentgenol.*, vol. 169, no. 2, pp. 409–15, Aug. 1997.
- [5] E. Furman-Haran, D. Grobgeld, F. Kelcz, and H. Degani, “Critical role of spatial resolution in dynamic contrast-enhanced breast MRI,” *J. Mag. Res. Im.*, vol. 13, no. 6, pp. 862–7, June 2001.
- [6] L. Liberman, E. A. Morris, M. J-Y. Lee, J. B. Kaplan, L. R. LaTrenta, J. H. Menell, A. F. Abramson, S. M. Dashnaw, D. J. Ballon, and D. D. Dershaw, “Breast lesions detected on MR imaging: Features and positive predictive value,” *Am. J. Roentgenol.*, vol. 179, no. 1, pp. 171–8, July 2002.
- [7] U. Hoffmann, G. Brix, M. V. Knopp, T. Heß, and W. J. Lorenz, “Pharmacokinetic mapping of the breast: A new method for dynamic MR mammography,” *Mag. Res. Med.*, vol. 33, no. 4, pp. 506–14, Apr. 1995.
- [8] S. G. Orel, M. D, and M. D. Schnall, “MR imaging of the breast for the detection, diagnosis, and staging of breast cancer,” *Radiology*, vol. 220, no. 1, pp. 13–30, 2001.

- [9] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: Physical principles and sequence design*, Wiley, New York, 1999.
- [10] D. G. Nishimura, “Principles of magnetic resonance imaging,” 1996, Unpublished textbook.
- [11] P. C. Lauterbur, “Image formation by induced local interactions: examples employing nuclear magnetic resonance,” *Nature*, vol. 242, no. 5394, pp. 190–1, Mar. 1973.
- [12] P. Mansfield and P. K. Grannell, “NMR ‘diffraction’ in solids?,” *J. Phys. C: Solid State Phys.*, vol. 6, no. 22, pp. L422–6, Nov. 1973.
- [13] P. S. Tofts and A. G. Kermode, “Measurement of the blood-brain barrier permeability and leakage space using dynamic MR imaging. 1. Fundamental concepts,” *Mag. Res. Med.*, vol. 17, no. 2, pp. 357–67, Feb. 1991.
- [14] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T. Y. Lee, N. A. Mayr, G. J. Parker, R. E. Port, J. Taylor, and R. M. Weisskoff, “Estimating kinetic parameters from dynamic contrast-enhanced T(1)-weighted MRI of a diffusable tracer: standardized quantities and symbols,” *J. Mag. Res. Im.*, vol. 10, no. 3, pp. 223–32, Sept. 1999.
- [15] M. V. Knopp, E. Weiss, H. P. Sinn, J. Mattern, H. Junkermann, J. Radeleff, A. Magener, G. Brix, S. Delorme, I. Zuna, and G. . Kaick, “Pathophysiologic basis of contrast enhancement in breast tumors,” *J. Mag. Res. Im.*, vol. 10, no. 3, pp. 260–6, Sept. 1999.
- [16] S. Sinha and U. Sinha, “Recent advances in breast mri and mrs,” *NMR Biomed.*, vol. 22, no. 1, pp. 1099–1492, Jan. 2009.
- [17] Dustin Newell, Ke Nie, Jeon-Hor Chen, Chieh-Chih Hsu, Hon Yu, Orhan Nalcioğlu, and Min-Ying Su, “Selection of diagnostic features on breast mri to differentiate between malignant and benign lesions using computer-aided diagnosis: differences in lesions presenting as mass and non-mass-like enhancement,” *European Radiology*, vol. 20, no. 4, pp. 771–781, 2010.
- [18] M. N. Mariano, M. A.A.J. van den Bosch, B. L. Daniel, K. W. Nowels, R. L. Birdwell, K. J. Fong, P. S. Desmond, S. Plevritis, L. A. Stables, M. Zakhour, R. J. Herfkens, and D. M. Ikeda, ,” .
- [19] M. D. Schnall, J. Blume, David A. Bluemke, Gia A. DeAngelis, Nanette DeBruhl, Steven Harms, Sylvia H. Heywang-Kbrunner, Nola Hylton, Christiane K. Kuhl, Etta D. Pisano, Petrina Causer, Stuart J. Schnitt, David Thickman, Carol B. Stelling, Paul T. Weatherall, Constance Lehman, and Constantine A. Gatsonis, “Diagnostic architectural and dynamic features at breast mr imaging: Multicenter study,” *Radiology*, vol. 238, no. 1, pp. 42–53, jan 2006.

- [20] Riham H. El Khouli, Katarzyna J. Macura, Peter B. Barker, Mohamed R. Habba, Michael A. Jacobs, and David A. Bluemk, “Relationship of temporal resolution to diagnostic performance for dynamic contrast enhanced mri of the breast,” *J. Mag. Res. Im.*, vol. 30, no. 5, pp. 9991004,, nov 2009.
- [21] J. Veltman, M. Stoutjesdijk, R. Mann, H. J. Huisman, J. O. Barentsz, J. G. Blickman, and C. Boetes, “Contrast-enhanced magnetic resonance imaging of the breast: the value of pharmacokinetic parameters derived from fast dynamic imaging during initial enhancement in classifying lesions,” *Eur. Radiol.*, vol. 18, no. 6, pp. 1123–1133, Feb. 2008.
- [22] P-A. Eliat, V. Dedieu, C. Bertino, V. Boute, J. Lacroix, J-M. Constans, B. . Korvin, C. Vincent, C. Bailly, F. Joffre, J. . Certaines, and D. Vincensini, “Magnetic resonance imaging contrast-enhanced relaxometry of breast tumors: an MRI multicenter investigation concerning 100 patients,” *Mag. Res. Im.*, vol. 22, no. 4, pp. 475–81, May 2004.
- [23] M. D. Pickles, M. Lowry, D. J. Manton, P. Gibbs, and L. W. Turnbull, “Role of dynamic contrast enhanced mri in monitoring early response of locally advanced breast cancer to neoadjuvant chemotherapy,” *Breast Cancer Res Treat*, vol. 91, pp. 1–10, 2005.
- [24] R. Johansen, L. R. Jensen, J. Rydland, P. E. Goa, K. A. Kvistad, T. F. Bathen, D. E. Axelson, S. Lundgren, and I. S. Gribbestad, “Predicting survival and early clinical response to primary chemotherapy for patients with locally advanced breast cancer using dce-mri,” *J. Mag. Res. Im.*, vol. 29, no. 6, pp. 1300–1307, June 2009.
- [25] A. Thukral, D. M. Thomasson, C. K. Chow, R. Eulate, S. B. Wedam, S. N. Gupta, B. J. Wise, S. M. Steinberg, D. J. Liewehr, P. L. Choyke, and S. M. Swain, “Inammatory breast cancer: Dynamic contrast-enhanced mr in patients receiving bevacizumab- initial experience,” *Radiol.*, vol. 244, no. 3, pp. 727–735, Sept. 2007.
- [26] T. E. Yankeelov, M. Lepagea, A. Chakravarthy, E. E. Broome, K. J. Niermann, M. C. Kelley, I. Meszoely, I. A. Mayer, C. R. Herman, K. McManus, R. R. Price, and J. C. Gore, “Integration of quantitative dce-mri and adc mapping to monitor treatment response in human breast cancer: initial results,” *Mag. Res. Im.*, vol. 25, no. 1, pp. 1–13, Jan. 2007.
- [27] D. K. Sodickson and W. J. Manning, “Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays,” *Mag. Res. Med.*, vol. 38, no. 4, pp. 591–603, Oct. 1997.
- [28] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, “SENSE: sensitivity encoding for fast MRI,” *Mag. Res. Med.*, vol. 42, no. 5, pp. 952–62, Nov. 1999.

- [29] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger, “Advances in sensitivity encoding with arbitrary k-space trajectories,” *Mag. Res. Med.*, vol. 46, no. 4, pp. 638–51, Oct. 2001.
- [30] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, “Generalized autocalibrating partially parallel acquisitions (GRAPPA),” *Mag. Res. Med.*, vol. 47, no. 6, pp. 1202–10, June 2002.
- [31] M. Bydder, J. Perthen, and J. Du, “Optimization of sensitivity encoding with arbitrary k-space trajectories,” *Mag. Res. Im.*, vol. 25, no. 8, pp. 1123–9, Oct. 2007.
- [32] L. Ying and J. Sheng, “Joint image reconstruction and sensitivity estimation in SENSE (JSENSE),” *Mag. Res. Med.*, vol. 57, no. 6, pp. 1196–1202, June 2007.
- [33] J. J. Van Vaals, M. E. Brummer, W. T. Dixon, H. H. Tuithof, H. Engels, R. C. Nelson, B. M. Gerety, J. L. Chezmar, and J. A. Den Boer, “Keyhole method of accelerating imaging of contrast agent uptake,” *J. Mag. Res. Im.*, vol. 3, no. 4, pp. 671–5, July 1993.
- [34] R. A. Jones, O. Haraldseth, T. B. Müller, P. A. Rinck, and A. N. Øksendal, “K-space substitution: A novel dynamic imaging technique,” *Mag. Res. Med.*, vol. 29, no. 6, pp. 830–4, June 1993.
- [35] E. N. Yeh, C. A. McKenzie, M. A. Ohliger, and D. K. Sodickson, “3Parallel magnetic resonance imaging with adaptive radius in k-space (PARS): Constrained image reconstruction using k-space locality in radiofrequency coil encoded data,” *Mag. Res. Med.*, vol. 53, no. 6, pp. 1383–92, June 2005.
- [36] A. A. Samsonov, “On optimality of parallel MRI reconstruction in k-space,” *Mag. Res. Med.*, vol. 59, no. 1, pp. 156–64, Jan. 2008.
- [37] J. G. Pipe, “Motion correction with PROPELLER MRI: Application to head motion and free-breathing cardiac imaging,” *Mag. Res. Med.*, vol. 42, no. 5, pp. 963–9, Nov. 1999.
- [38] K. Arfanakis, A. A. Tamhane, J. G. Pipe, and M. A. Anastasio, “K-space undersampling in PROPELLER imaging,” *Mag. Res. Med.*, vol. 53, no. 3, pp. 675–83, Mar. 2005.
- [39] F. R. Korosec, R. Frayne, T. M. Grist, and C. A. Mistretta, “Time-resolved contrast-enhanced 3D MR angiography,” *Mag. Res. Med.*, vol. 36, no. 3, pp. 345–51, Sept. 1996.
- [40] K. K. Vigen, D. C. Peters, T. M. Grist, W. F. Block, and C. A. Mistretta, “Undersampled projection-reconstruction imaging for time-resolved contrast-enhanced imaging,” *Mag. Res. Med.*, vol. 43, no. 2, pp. 170–6, Feb. 2000.

- [41] C. R. Haider, H. H. Hu, N. G. Campeau, J. Huston, and S. J. Riederer, “3D high temporal and spatial resolution contrast-enhanced MR angiography of the whole brain,” *Mag. Res. Med.*, vol. 60, no. 3, pp. 749–60, Sept. 2008.
- [42] Z-P. Liang and P. C. Lauterbur, “An efficient method for dynamic magnetic resonance imaging,” *IEEE Trans. Med. Imag.*, vol. 13, no. 4, pp. 677–86, Dec. 1994.
- [43] J. M. Hanson, Z-P. Liang, E. C. Wiener, and P. C. Lauterbur, “Fast dynamic imaging using two reference images,” *Mag. Res. Med.*, vol. 36, no. 1, pp. 172–5, July 1996.
- [44] Z-P. Liang, B. Madore, G. H. Glover, and N. J. Pelc, “Fast algorithms for GS-model-based image reconstruction in data-sharing Fourier imaging,” *IEEE Trans. Med. Imag.*, vol. 22, no. 8, pp. 1026–30, Aug. 2003.
- [45] D. C. Noll, D. G. Nishimura, and A. Macovski, “Homodyne detection in magnetic resonance imaging,” *IEEE Trans. Med. Imag.*, vol. 10, no. 2, pp. 154–63, June 1991.
- [46] G. McGibney, M. R. Smith, S. T. Nichols, and A. Crawley, “Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI,” *Mag. Res. Med.*, vol. 30, no. 1, pp. 51–9, July 1993.
- [47] K. F. King and L. Angelos, “SENSE with partial Fourier homodyne reconstruction,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2000, p. 153.
- [48] M. Bydder and M. D. Robson, “Partial Fourier partially parallel imaging,” *Mag. Res. Med.*, vol. 53, no. 6, pp. 1393–401, June 2005.
- [49] S. Krishnan and T. L. Chenevert, “Spatio-temporal bandwidth-based acquisition for dynamic contrast-enhanced magnetic resonance imaging,” *J. Mag. Res. Im.*, vol. 20, no. 1, pp. 129–37, July 2004.
- [50] S. Krishnan, *K-space acquisition method for dynamic contrast-enhanced MRI: Application to breast tumors*, Ph.D. thesis, Univ. of Michigan, Ann Arbor, MI, 48109-2122, Ann Arbor, MI, 2004.
- [51] J. Tsao, P. Boesiger, and K. P. Pruessmann, “k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations,” *Mag. Res. Med.*, vol. 50, no. 5, pp. 1031–42, Nov. 2003.
- [52] J. Tsao, S. Kozerke, P. Boesiger, and K. P. Pruessmann, “Optimizing spatiotemporal sampling for k-t BLAST and k-t SENSE: Application to high-resolution real-time cardiac steady-state free precession,” *Mag. Res. Med.*, vol. 53, no. 6, pp. 1372–82, June 2005.

- [53] M. S. Hansen, . C. Baltés, J. Tsao, S. Kozerke, K. P. Pruessmann, and H. Eggers, “k-t BLAST reconstruction from non-Cartesian k-t space sampling,” *Mag. Res. Med.*, vol. 55, no. 1, pp. 85–91, Jan. 2006.
- [54] D. Xu, K. F. King, and Z-P. Liang, “Improving k-t SENSE by adaptive regularization,” *Mag. Res. Med.*, vol. 57, no. 5, pp. 918–30, May 2007.
- [55] K. T. Block, M. Uecker, and J. Frahm, “Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint,” *Mag. Res. Med.*, vol. 57, no. 6, pp. 1086–98, June 2007.
- [56] L. Chen, M. C. Schabel, and E. V. R. DiBella, “Reconstruction of dynamic contrast enhanced magnetic resonance imaging of the breast with temporal constraints,” *Mag. Res. Im.*, vol. 28, no. 5, pp. 637–45, June 2010.
- [57] B. P. Sutton, D. C. Noll, and J. A. Fessler, “Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities,” *IEEE Trans. Med. Imag.*, vol. 22, no. 2, pp. 178–88, Feb. 2003.
- [58] R. H. Chan and M. K. Ng, “Conjugate gradient methods for Toeplitz systems,” *SIAM Review*, vol. 38, no. 3, pp. 427–82, Sept. 1996.
- [59] J. A. Fessler, S. Lee, V. T. Olafsson, H. R. Shi, and D. C. Noll, “Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity,” *IEEE Trans. Sig. Proc.*, vol. 53, no. 9, pp. 3393–402, Sept. 2005.
- [60] K. A. Khalsa and J. A. Fessler, “Accelerated iterative reconstruction of temporally regularized dynamic MRI,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2007, p. 1922.
- [61] K. A. Khalsa and J. A. Fessler, “Resolution properties in regularized dynamic MRI reconstruction,” in *Proc. IEEE Intl. Symp. Biomed. Imag.*, 2007, pp. 456–9.
- [62] J. A. Fessler and W. L. Rogers, “Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs,” *IEEE Trans. Im. Proc.*, vol. 5, no. 9, pp. 1346–58, Sept. 1996.
- [63] E. Asma and R. M. Leahy, “Mean and covariance properties of dynamic PET reconstructions from list-mode data,” *IEEE Trans. Med. Imag.*, vol. 25, no. 1, pp. 42–54, Jan. 2006.
- [64] J. A. Fessler, *Image reconstruction: Algorithms and analysis*, ?, 2006, Book in preparation.
- [65] T. Yankeelov, J. Luci, M. Lepage, R. Li, L. Debusk, P. C. Lin, R. Price, and J. Gore, “Quantitative pharmacokinetic analysis of dce-mri data without an arterial input function: a reference region model,” *Mag. Res. Med.*, vol. 23, no. 4, pp. 519–529, May 2005.

- [66] G. J. Parker, A. Macdonald, S. Cheung, D. L. Buckley, A. Jackson, G. Jayson, and C. Roberts, “An experimentally-derived functional form for a population-averaged high temporal resolution arterial input function,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2005, p. 2100.
- [67] C. Yang, G. S. Karczmar, M. Medved, and W. M. Stadler, “Multiple reference tissue method for contrast agent arterial input function estimation,” *Mag. Res. Med.*, vol. 58, no. 6, pp. 1266–75, Dec. 2007.
- [68] M. Sasaki, E. Shibata, Y. Kanbara, and S. Ehara, “Enhancement effects and relaxivities of gadolinium-DTPA at 1.5 versus 3 Tesla: A phantom study,” *Magnetic Resonance in Medical Sciences*, vol. 4, no. 3, pp. 145–9, 2005.
- [69] R. Rakow-Penner, B. Daniel, H. Yu, A. Sawyer-Glover, and G. H. Glover, “Relaxation times of breast tissue at 1.5T and 3T measured using IDEAL,” *J. Mag. Res. Im.*, vol. 23, no. 1, pp. 87–91, Jan. 2006.
- [70] G. Golub and V. Pereyra, “Separable nonlinear least squares: the variable projection method and its applications,” *Inverse Prob.*, vol. 19, no. 2, pp. R1–26, Apr. 2003.
- [71] J. Li, Y. Yu, Y. Zhang, S. Bao, C. Wu, X. Wang, J. Li, X. Zhang, and J. Hu, “A clinically feasible method to estimate pharmacokinetic parameters in breast cancer,” *Med. Phys.*, vol. 36, no. 8, pp. 3786–3794, Aug. 2009.
- [72] J. A. Fessler, “Matlab tomography toolbox,” 2004, Available from <http://www.eecs.umich.edu/~fessler>.
- [73] Y. Kim, J. A. Fessler, and D. C. Noll, “Smoothing effect of sensitivity map on fMRI data using a novel regularized self-calibrated estimation method,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2008, p. 1267.