

ACCELERATED CHEMICAL SHIFT ENCODED WATER-FAT IMAGING

by

Samir D. Sharma

**A Dissertation Presented to the
FACULTY OF THE USC GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA
In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(ELECTRICAL ENGINEERING)**

May 2012

Copyright 2012

Samir D. Sharma

Acknowledgements

I wish to thank the members of the Magnetic Resonance Engineering Lab (MREL) for all of their guidance and support. In particular, I must thank my adviser, Dr. Krishna S. Nayak, for all that he has done to bring me to USC and for all of the opportunities that he has made possible for me. I must also personally thank Dr. Houchun Harry Hu for introducing me to the field of water-fat MRI when I was a junior graduate student who was struggling to find a research project. His thoughtful and unwavering mentorship has helped me in ways for which I will be forever indebted to him. I look forward to a continued professional and personal relationship with both Krishna and Harry.

I thank my family for all of their continued love and support throughout my 24 or so years of school. To my Mom and Dad, to Sabina and Andy and Ayden, and to Shalu and Arvind, thank you for always being there for me.

And to Megan, for all that we've shared and experienced together and for all of your love and support, I thank you. I look forward to sharing with you what the future holds for us.

Table of Contents

Acknowledgements	ii
List of Tables	iv
List of Figures	v
Abstract	xiv
Chapter 1: Introduction	1
Chapter 2: MRI of Water and Fat	4
2.1: Basic MRI Physics	4
2.2: Water-Fat Imaging	18
2.3: Accelerated Imaging Techniques and Their Applications to Chemical Shift Encoded Water-Fat Imaging	34
Chapter 3: Multiscale B-splines for B ₀ Field Map Estimation	47
3.1: Introduction	47
3.2: Methods	48
3.3: Results	61
3.4: Discussion	68
Chapter 4: Accelerated Water-Fat Separation Using Parallel Imaging and Compressed Sensing	70
4.1: Introduction	70
4.2: Methods	73
4.3: Results	86
4.4: Discussion	92
Chapter 5: Accelerated T ₂ *-Compensated Quantitative Liver Fat Imaging Using Multiscale B-splines, Parallel Imaging, and Compressed Sensing	94
5.1: Introduction	94
5.2: Methods	96
5.3: Results	108
5.4: Discussion	113
Chapter 6: Future Work	116
Bibliography	119

List of Tables

- Table 3.1: The data acquisition parameters for each of the datasets are presented in this table. The first three datasets were collected on a GE 3T MRI system at HCC-II. The heart dataset, courtesy of Peter Kellman PhD, was acquired on a 1.5T Siemens MRI system. 60
- Table 5.1: This table presents the data acquisition parameters that were used to collect the fully sampled phantom and liver datasets. For the liver datasets, the field of view (FOV) was set based on the size of the subject. This affected the echo spacing (ΔTE) and the scan time. The slice thickness (Δz) was kept the same for all subjects. 104
- Table 5.2: This table presents the reconstruction parameters for the ARC/T2*-IDEAL and PI-CS-T2*-MSBS. N_{acs} denotes number of autocalibration lines. Note that the outer reduction factor is only applicable to the ARC/T2*-IDEAL reconstruction. 106

List of Figures

- Figure 2.1: (a) The magnetization (M) in the presence of the B_0 field. (b) The rotation of the magnetization as a result of the applied B_1 field. M^- denotes the magnetization just before the application of the B_1 field, and M^+ denotes the magnetization just after the application of the B_1 field. In this case, the amplitude and duration of the B_1 field caused a 90° rotation of the magnetization. 7
- Figure 2.2: Three receiver coils that can be used to measure the MR signal. (Left) This is a birdcage coil that can be used for many anatomies such as the head, ankle, and knee. (Middle) This is a torso coil that is typically used to acquire data from the torso. (Right) This is a knee coil that is typically used to acquire data from the knee. (Source: gehealthcare.com). 10
- Figure 2.3: An example timing diagram and associated k-space trajectory for a 2D Fourier Transform (2DFT) pulse sequence. The radiofrequency (RF) pulse represents the B_1 field that is applied to excite the spins. G_z denotes the z-axis gradient that is applied simultaneously with RF to choose the imaging slice. The G_y and G_x gradients define the trajectory through k-space as seen by the corresponding time points (a-d) in the right image. The area of the G_y gradient is modified between subsequent excitations to acquire a different k-space line along the k_y axis. DAQ stands for data acquisition. 13
- Figure 2.4: The k-space data (left) that were acquired using a 2DFT pulse sequence. The log of the magnitude is shown for better visibility. The image (right) that was reconstructed using an inverse 2D Fourier transform (IFT) of the acquired data. This particular image is a coronal slice of the brain. Both the k-space and image data are complex valued although only the magnitude is shown. 14
- Figure 2.5: Liver image without (left) and with (right) fat suppression. The fat suppressed image clearly shows the metastatic deposit in the left lobe of the liver (arrows). Images from [36]. 19
- Figure 2.6: Liver fat fraction images from two subjects. The fat fraction that was measured using a water-fat imaging technique is shown in each image. For reference, a 5.56% threshold has been cited to discern normal from abnormal liver fat [63]. Images from [27]. 19

- Figure 2.7: Chemical structures of water (left) and a glyceride (right), which is a main component of a fat molecule. Notice that the glyceride has many types of chemical bonds that contain hydrogen. Image courtesy Houchun H. Hu, PhD. 20
- Figure 2.8: Spectrum of a sample containing both water and fat. The water has one peak while the fat has multiple peaks although the CH_2 peak dominates. The horizontal axis is in parts per million (ppm), which is independent of magnetic field strength. At 3T, the frequency shift of the CH_2 peak relative to the water peak is about -420 Hz. Image courtesy Houchun H. Hu, PhD. 21
- Figure 2.9: The schematic (left) explains the STIR method. After an inversion pulse, the fat recovers more quickly than the water due to the shorter T_1 time of the former. At the time (200 ms here) of the zero-crossing of fat, imaging will occur to effectively null the fat component. The image (right) shows an example image using the STIR method. The fat signal was correctly suppressed. The arrows point to areas that may cause problems when using non T_1 -based techniques. Images from [3]. 22
- Figure 2.10: A schematic of the chemically shift selective (CHESS) method applied to a 2DFT pulse sequence. The 90° excitation pulse is spectrally-selective so that only the fat spins are excited and tipped into the transverse plane. The G_z gradient is applied to dephase the excited spins, resulting in zero net magnetization. A standard 2DFT imaging sequence immediately follows. Please refer to Figure 2.3 for details of the 2DFT sequence. 24
- Figure 2.11: The B_0 field map and water image estimates in the brachial plexus using voxel-independent IDEAL (IDEAL-VI) and IDEAL with region-growing (IDEAL-RG). The arrows in the IDEAL-VI images highlight incorrect B_0 field map estimates that cause water-fat swaps. The IDEAL-RG estimates do not exhibit these errors. The brachial plexus is typically a difficult area for uniform water-fat separation due to high off-resonance. Images from [71]. 28
- Figure 2.12: The least-squares cost as a function of B_0 field map for a voxel taken from an ankle dataset. Without additional knowledge, the field map estimate would be initialized to zero (solid white circle). Any gradient descent based algorithm would converge to a local minimum solution (dashed white circle) that would cause swapping of the water and fat signals. The true B_0 field map value is denoted by the solid black circle. 30

Figure 2.13: A head coil with eight receiver elements (top left) surrounds the object of interest (bottom left). The receiver elements simultaneously acquire functions of the k-space data. The image from a receiver element is weighted by a distinct spatial function (right) that is related to the element's spatial position relative to the object. The eight-channel head coil can, in theory, simultaneously acquire eight linearly-independent measurements that can be used to accelerate a data acquisition by up to a factor of eight. 35

Figure 2.14: Accelerating the data acquisition can be done by uniformly undersampling k-space. This results in aliasing in the image domain (top). In this example, k-space was undersampled by a factor of two, so each pixel in the aliased image is a weighted sum of two pixels from the underlying image. If one estimates the coil sensitivities (denoted by C_{xy}), a set of linear equations can be constructed that relates the underlying image pixels to the aliased measurements (see equation). The aliasing can be done by solving the equation, which results in the SENSE reconstructed image (bottom). 37

Figure 2.15: The GRAPPA reconstruction synthesizes non-acquired k-space points (red) by a linear combination of neighboring acquired k-space samples across all coils (gray). The weights that are used in the linear combination are estimated from a calibration region within the acquired data itself. After the non-acquired k-space samples from each individual coil have been synthesized, the individual coil images are calculated via the Fourier transform. The resulting images can then be combined via root-sum-of-squares. Image from [42]. 38

Figure 2.16: The SPIRiT reconstruction imposes the constraint that each k-space point, acquired (gray) or not (red), be a linear combination of neighboring k-space samples across all coils. Like GRAPPA, the weights that are used in the linear combination are estimated from a calibration region within the acquired data itself. After the reconstruction, the individual coil images are calculated via the Fourier transform. The resulting images are then combined via root-sum-of-squares. Image from [42]. 39

- Figure 2.17: The compressibility of an axial brain image using a discrete cosine transform (DCT), a Daubechies-4 wavelet transform, and a finite difference transform. Selected percentages of the highest-magnitude coefficients are used to reconstruct the image (middle three columns). The wavelet transform provides high compressibility of the image. Images from [39]. 41
- Figure 2.18: Radial and spiral are two common non-Cartesian trajectories that are used to reduce the data acquisition time. Each segment that ends with an arrowhead represents the k-space that is covered after one excitation pulse. The use of non-Cartesian trajectories may require advanced image reconstruction techniques that compensate for off-resonance effects. Images from [7]. 44
- Figure 2.19: This Venn diagram depicts the different acceleration techniques that have been applied to chemical shift encoded water-fat imaging. Only the first author of each work is shown. PI – parallel imaging, CS – compressed sensing, NC – non-Cartesian, PK – partial k-space. In this dissertation, both parallel imaging and compressed sensing are used to accelerate chemical shift encoded water-fat imaging. 46
- Figure 3.1: A 1D example to illustrate the vectors that span the different subspaces in which the B0 field map estimate is updated. (a) One vector that spans the level 0 subspace. The B0 field map estimate will be a constant value for all pixels. (b) The three vector that span the level 1 subspace. The current field map estimate, which originated from the level 0 subspace, will be updated in this subspace. (c) The five vectors that span the level 2 subspace. As in the previous case, the current B0 field map estimate will be updated in this subspace. 50
- Figure 3.2: Single cubic B-spline functions from coarser scale (top) to finer scale (bottom) and their corresponding cubic B-spline sets. The B-spline set is created by spatially shifting the single cubic B-spline in both spatial directions by multiples of the knot spacing. The cubic B-splines are nonnegative everywhere and the set sums to one at all spatial positions. The field map update term at the m^{th} scale is restricted to be in the space that is spanned by the m^{th} cubic B-spline set. 54

Figure 3.3: The B0 field map, water, and fat estimates using IDEAL-VI, IDEAL-RG and the proposed multiscale B-spline (MSBS) method. Notice the errors in the IDEAL-VI B0 field map estimate (white arrows). These errors propagate into the water and fat images producing swaps. That is, the estimated fat image contains signal from the water component and vice versa. The IDEAL-RG estimates serve as the reference. The proposed multiscale B-spline method produces estimates that are very similar to IDEAL-RG. It is important to remember that the benefit of the MSBS approach as compared to IDEAL-RG will become apparent in later chapters. 62

Figure 3.4: The evolution of the B0 field map estimate using the proposed MSBS method. The early iterations using coarse scale B-splines are needed to reduce the possibility of converging to a local minimum solution, while of course providing a rough approximation of the B0 field map. The later iterations provide refinement to the field map estimate. It is unlikely that convergence to a local minima solution will originate at these finer scales. 63

Figure 3.5: The B0 field map, water, and fat image estimates of the liver using IDEAL-RG and the proposed MSBS method. The estimates using the proposed technique appear similar to the reference IDEAL-RG implementation. In the next chapter, image estimates will be shown using the same dataset that has been retrospectively undersampled to simulate an accelerated acquisition. It will be seen that the estimates between the two cases are quite similar. 64

Figure 3.6: The B0 field map, water, and fat image estimates of the brachial plexus using IDEAL-RG and the proposed MSBS method. The brachial plexus is typically a challenging area for water-fat separation due to the high degree of off-resonance in this anatomy. The IDEAL-RG estimate exhibits a water-fat swap due to an incorrect B0 field map estimate (white arrows). Notice that this error occurs near a low signal region. The lack of a connecting signal region may have caused this swap. This is one commonly cited shortcoming of the IDEAL-RG implementation. The proposed MSBS accurately estimates the B0 field map and properly separated water and fat. 65

- Figure 3.7:** The B0 field map, water, and fat image estimates of a short axis view of the heart using IDEAL-VI, IDEAL-RG and the proposed MSBS method. The IDEAL-VI estimates exhibit regions of water-fat swaps due to errors in the B0 field map (black arrows). Water-fat swaps are also seen in the IDEAL-RG estimates. This is likely because the algorithm attempted to region grow through areas of low SNR. The proposed MSBS method accurately estimated the B0 field map, which led to proper water-fat separation. (Dataset courtesy of Peter Kellman PhD, National Institutes of Health) 67
- Figure 4.1:** Uniformly undersampled (US) and Poisson disk (PD) k-space sampling patterns and their corresponding magnitude of the point-spread function (PSF). Both schemes sample the central 24x24 region for kernel calibration and result in 7.7x acceleration. Additionally, both schemes avoid large gaps of non-sampling points in k-space, which is desirable for parallel imaging. However, the structure in the US PSF leads to the coherence aliasing artifacts that should be avoided for a compressed sensing reconstruction. The PD PSF leads to the incoherent, noise-like artifacts that are desired for a compressed sensing reconstruction. 76
- Figure 4.2:** The top row in each set of images shows the fully sampled coil images and the bottom row shows the corresponding coil sensitivities (magnitude) that were derived from the SPIRiT k-space kernel using only the central 16 phase encode lines. Using the coil sensitivities implicitly imposes the SPIRiT constraint, which avoids the need for an explicit calibration consistency expression in the reconstruction. 81
- Figure 4.3:** After initialization, the coil sensitivity maps are estimated and then held fixed. Subsequently, the water and fat images and the B0 field map are estimated via an alternating minimization. The water and fat images are estimated using parallel imaging and compressed sensing while the B0 field map is updated using multiscale cubic B-splines. 83
- Figure 4.4:** The B0 field map, water, and fat estimates of the liver that was acquired using an eight-channel torso coil. The ARC/IDEAL-RG estimates exhibit unresolved aliasing artifacts (arrows) and noise amplification (arrowhead). These artifacts were anticipated since the outer acceleration factor of 4 is greater than the number of coils along the axis of undersampling. In contrast, the estimates using the proposed method exhibit only slight incoherent artifacts as a result of the Poisson disk sampling and ℓ_1 -regularization. 88

Figure 4.5: The B0 field map, water, and fat estimates of the brachial plexus acquired using an eight-channel neurovascular coil. Field map estimate errors using IDEAL-VI (white ellipses) cause water-fat swaps. The MSBS field map estimation approach correctly estimates the field map to avoid the swaps. The ARC/MSBS estimates exhibit noise artifacts (arrowheads), especially in the brain. These artifacts were expected since the 3x by 3x outer acceleration factor is greater than the number of receiver elements. The estimates using the proposed PI-CS-MSBS approach exhibit a relatively reduced level of artifacts. Slight loss of subtle features is seen (white arrows). 89

Figure 4.6: The ARC/IDEAL-RG estimates exhibit noticeable artifacts as a result of unresolved aliasing (white arrows). This was expected since the outer acceleration factor (3x by 3x) was greater than the number of coils. The estimates using the proposed PI-CS-MSBS method correspond well with the reference estimates. The scan time of the prospectively undersampled acquisition was only 01:12 as compared to 10:11 of the fully-sampled acquisition. 90

Figure 4.7: The B0 field map, water, and fat estimates of the knee acquired using an eight-channel knee coil. The arrowheads and arrows in the ARC/IDEAL-RG estimates highlight, respectively, regions of noise amplification and image artifacts. These artifacts were anticipated since the outer acceleration factor of 3x by 3x is greater than the number of receiver elements. In contrast, the estimates using the proposed method exhibit incoherent artifacts as a result of the Poisson disk sampling and l1-regularization. System issues, such as eddy currents, that are associated with the k-space sampling order do not affect the quality of the water-fat separation as evidenced by the retrospective versus prospective results. 91

Figure 5.1: The least squares cost is a nonconvex function of both B0 field map and R2* parameters. However, the cost appears to be a convex function of both B0 field map and R2* in the neighborhood of the true B0 field map value. This suggests that signal estimation can be done in a pseudo-sequential manner in which the B0 field map is first approximated, followed by a joint refinement and estimation of the B0 field map and R2* values, respectively. 98

- Figure 5.2: An example to show the compressibility of R2* images using the Daubechies-4 wavelet transform. Each image is shown using only a selected percentage of the highest magnitude transform coefficients. Calculation of the mean R2* within a region of interest shows that nearly no to little degradation of the mean R2* value occurs with higher compression. 100
- Figure 5.3: The flowchart summarizes the steps for signal estimation. The coil sensitivity maps are estimated first and are then held fixed. Subsequently, the water, fat, B0 field map and R2* map are estimated via an alternating minimization. The R2* map is estimated only after an approximation of the B0 field map has been made. This has the benefit of avoiding local minima estimates of R2*. The following abbreviations are used in the flowchart: B0 field map (ψ), water/fat (w/f), B-spline set at m^{th} level (B_m). 102
- Figure 5.4: Liver region of interest (ROI) (white square) in which the fat fraction estimate was measured. The ROI was 3x3 pixels, which covers a region of approximately 7.5x7.5mm². The ROI is shown on the water image for better visibility. The reference fat fraction value was calculated by taking the average of the T2*-IDEAL fat fraction estimates in the ROI. 107
- Figure 5.5: The different concentrations of distilled water and Intralipid (IL) (left image). Each of the three labeled phantoms also included 1 mg of MnCl₂ to simulate the T2* shortening that would be caused by iron overload. The fat fraction estimates differ by approximately 5% without/with R2* compensation (top center image). The fat fraction estimates using the proposed method (bottom center image) closely agree with the reference estimates. The estimated R2* values are quite high due to the presence of MnCl₂ (right images). 109
- Figure 5.6: An example of the fat fraction and R2* estimates using T2*-IDEAL, 3.1x ARC/T2*-IDEAL, and 3.1x PI-CS-T2*-MSBS. The reported fat fraction was calculated as the average of the fat fraction values in the ROI. Notice the unresolved aliasing artifacts in the ARC/T2*-IDEAL fat estimate (white arrowhead). 111

Figure 5.7: The average fat fraction estimate in the ROI is plotted against the reference fat fraction value. The identity line is shown with dashes. Linear regression yielded slope, intercept, and R^2 values of 0.78 ± 0.22 , -0.0004 ± 0.024 , and 0.34, respectively, for ARC/T2*-IDEAL and 1.08 ± 0.04 , -0.007 ± 0.004 , and 0.97, respectively, for the proposed PI-CS-T2*-MSBS approach. 112

Figure 5.8: The mean square error (MSE), with respect to the reference fat fraction, of the fat fraction estimates in the ROI. Each triplet represents the MSE for T2*-IDEAL (left), 3.1x ARC/T2*-IDEAL (center), and 3.1x PI-CS-T2*-MSBS (right). Each MSE value is decomposed into its variance (blue) and squared bias (red) components. In 27 of 28 cases, the MSE of the ARC/T2*-IDEAL estimate was the highest. In general, most of the error was attributed to the variance component. 113

Abstract

Water-fat imaging techniques are commonly used in the research and clinical setting to separate and/or quantitate the two species. Within this class of techniques, chemical shift encoded methods have recently gained significant interest because of their insensitivity to magnetic field inhomogeneity.

The benefits of chemical shift encoding come at a cost of a long scan time because data must be acquired at multiple echo times. This cost increases the likelihood of image artifacts that result from patient motion during the scan while also reducing the overall level of patient comfort. The long scan time becomes especially problematic when imaging non-stationary regions of the body such as the abdomen. In these cases, a compromise with other parameters must be made. Most often, a sacrifice of the spatial resolution and/or volume coverage is required to avoid motion-related image artifacts.

The objective of this dissertation is to present a new method for chemical shift encoded water-fat separation and quantitation from an accelerated data acquisition. In this context, accelerating the acquisition implies that a fewer number of measurements are acquired as compared to the requirements from the Nyquist Sampling Theorem. The resulting image recovery is an ill-posed inverse problem that requires the incorporation of additional knowledge to find the desired solution. A combined parallel imaging and compressed sensing approach for accelerated water-fat separation and quantitation is introduced. To enable this work, I developed a novel method for B_0 field map estimation as an

alternative to traditional region-growing and region-merging schemes. Uniform water-fat separation is demonstrated in a variety of anatomies from both retrospectively and prospectively undersampled acquisitions. In addition, accurate quantitation is presented in water-fat phantoms as well as in liver datasets acquired from seven subjects.

Chapter 1: Introduction

Magnetic resonance imaging (MRI) is a powerful and noninvasive imaging modality that is routinely used for diagnostic studies in the clinical setting. Unlike x-ray computerized tomography (CT), MRI does not use ionizing radiation nor does it present any other known health risks. The versatility of MRI allows the user to control parameters that will influence certain aspects of the resultant image. This allows MRI to be applied to a wide variety of applications such as functional brain imaging, diffusion imaging, and water-fat imaging.

Water-fat imaging is a term that is used to describe any MRI technique that specifically aims to separate and/or quantitate the distribution of water and fat within the imaged object. These techniques are commonly used in the clinical setting to suppress the bright fat signal so that underlying pathology can be more clearly visualized. Likewise, the techniques can be applied to suppress the water signal in those applications seeking information about fat within the body. Recently, advanced water-fat imaging techniques have been developed that increase the robustness of water-fat separation and that provide accurate quantitation of water and fat.

The benefits of water-fat imaging and MRI in general come at a cost of a long scan time. While one CT scan will take seconds to a few tens of seconds, an MRI scan will typically take a few minutes and some may even take up to a few tens of minutes. As will be explained in subsequent chapters, the scan time when doing water-fat imaging can be even longer than standard MRI scans.

These long scan times increase the likelihood of image artifacts that result from patient motion during the scan while also reducing the overall level of patient comfort. In addition, the long scan time becomes especially problematic when imaging non-stationary regions of the body such as the abdomen. In these cases, a compromise with other parameters must be made. Most often, a sacrifice of the spatial resolution and/or volume coverage is required to avoid motion-related image artifacts.

The objective of this dissertation is to present a new method for water-fat separation and quantitation from an accelerated data acquisition. In this context, accelerating the acquisition implies that a fewer number of measurements are acquired as compared to conventional methods. The resulting image recovery is an ill-posed inverse problem that requires the incorporation of additional knowledge to find the desired solution.

The next chapter provides the basic background on MRI physics, water-fat imaging, and the current applications of accelerated imaging techniques to water-fat imaging. This chapter is meant to provide the reader with the necessary background for understanding the motivation and basic idea of the subsequent chapters.

Chapters 3-5 each present a different aspect of the proposed method. Each chapter builds upon the previous one by extending some component in the signal model and reconstruction routine. Chapter 3 develops the signal model and introduces the use of multiscale B-splines for B₀ field map estimation.

Chapter 4 extends this framework to allow for accelerated water-fat separation. Lastly, chapter 5 extends the method for water-fat separation by incorporating the necessary modifications for accurate quantitation. This chapter can be viewed as the culmination of the advances presented in the three chapters.

Finally, chapter 6 concludes this dissertation by discussing ongoing work and ideas for future work.

Chapter 2: MRI of Water and Fat

This chapter begins with an introduction to basic MRI physics. An emphasis is placed on those concepts that are most relevant to the topic of this dissertation. The next section focuses on the traditional methods for the imaging of water and fat using MRI. The molecular differences between water and fat are explained, and the prominent MR imaging methods that exploit these differences are presented. Chemical shift encoded techniques are discussed in greatest depth because the work presented in this dissertation lies within this category. The section concludes with the current limitations of chemical shift encoded techniques. Finally, the topic of accelerated MR imaging methods is presented.

2.1: Basic MRI Physics

2.1.1: Nuclear Spins

Atoms with an odd number of protons and/or an odd number of neutrons possess a nonzero nuclear magnetic moment [50]. The magnetic moment of an atom is a vector quantity that determines the torque that a magnetic field will exert on it. This quantity can be expressed as

$$\mu = \gamma \hbar I \tag{2.1}$$

where γ is the gyromagnetic ratio (a known constant that is unique to a nuclear species), \hbar is the reduced Planck's constant, and I is the spin operator of the atom. In the human body, hydrogen (^1H) is the most abundant atom that exhibits a nonzero magnetic moment; sodium (^{23}Na), phosphorous (^{31}P), and potassium (^{39}K) also exist but in much smaller quantities.

2.1.2: Magnetic Fields and Their Effects on Nuclear Spins

In MRI, three types of magnetic fields are used to manipulate the nuclear spins. These three fields are commonly referred to as the B0 field, B1 field, and gradient (G) field. Each of these fields performs a unique task. The following subsections describe each of the magnetic fields and their effects on nuclear spins.

B0 Field

The B0 field is an external static magnetic field that is used to polarize the object to be imaged. In the absence of an external magnetic field, the nuclear spins within the object are randomly oriented, which results in zero net macroscopic magnetic moment, or magnetization. The presence of an external magnetic field has two important effects.

First, the spins assume either a parallel or antiparallel orientation relative to the direction of the magnetic field. By convention, the longitudinal axis (z-axis) is chosen to be along the direction of the B0 field. Interestingly, the difference in

the number of spins that assume a parallel versus an antiparallel orientation is only about 7 parts per million [50], but this difference is enough to generate a discernible signal in MRI.

The second effect of an external magnetic field is the exhibition of resonance by the spins at the so-called Larmor frequency. The Larmor frequency is dependent on the nuclear species and the strength of the magnetic field, and can be written as

$$f = \frac{\gamma}{2\pi} B \quad (2.2)$$

where, for example, $\gamma/2\pi = 42.58$ MHz/Tesla (T) for ^1H and B is the strength of the magnetic field. Typical B_0 field strengths for clinical MRI are 1.5T and 3T. Resonance of the nuclear spins at the Larmor frequency is important as the discussion now moves to the B_1 field.

B1 Field

The B_1 field is a dynamic magnetic field that is used to rotate the magnetization away from the longitudinal axis in a process called excitation. This field is applied in the transverse plane, which is perpendicular to the longitudinal axis, at a frequency that is equal to the Larmor frequency of the nuclear species of interest. The amplitude and duration of the applied field dictate the degree of rotation from the longitudinal axis. Figure 2.1 depicts the rotation of the

magnetization from the longitudinal axis to the transverse plane as a result of the B1 field. After applying the B1 field, the magnetization rotates about the longitudinal axis at the Larmor frequency in a process known as precession. It is at this point that the MR signal is ready to be measured. First, however, the gradient fields must be applied to spatially encode the signal.




Figure 2.1: (a) The magnetization (M) in the presence of the B0 field. (b) The rotation of the magnetization as a result of the applied B1 field. M⁰ denotes the magnetization just before the application of the B1 field, and M⁺ denotes the magnetization just after the application of the B1 field. In this case, the amplitude and duration of the B1 field caused a 90° rotation of the magnetization.

Gradient (G) Fields

Just after excitation using the B1 field, the magnetization is rotated away from the longitudinal axis and each nuclear spin is precessing about that axis at the Larmor frequency. The signal that would be recorded is a superposition of all precessing spins. However, since all of the spins are rotating at the same frequency, it is generally not possible to spatially localize the signals.

Gradient fields are used to spatially localize the MR signal. For example, a gradient field that is applied in the x-direction will alter the magnetic field along that axis. Equation 2.2 indicates that the spins will then precess at a frequency that depends on their position along the x-axis. Thus, the spins can be localized along the x-axis. By applying gradients in the other two axes, the spins can be localized in the three spatial dimensions. Typically, linear gradients are used so that the magnetic field is a linear function of position. One benefit of linear gradients will become apparent as spatial encoding is discussed in detail in later sections.

2.1.3: Relaxation

After excitation of the magnetization by the B1 field, the spins are rotated from the longitudinal axis and begin precession in the transverse plane at a frequency that is linearly proportional to the applied magnetic field. The spins do not remain in the transverse plane indefinitely; they will eventually return to alignment along the longitudinal axis in a process known as relaxation. There are two types of relaxation in MRI: 1) longitudinal relaxation and 2) transverse relaxation.

Longitudinal relaxation describes the process by which the magnetization returns to equilibrium along the longitudinal axis. This relaxation, better thought of as signal recovery, is characterized by a time constant known as T1. In the case of a 90° excitation, the magnetization has been rotated entirely into the

transverse plane, as in Figure 2.1b. The recovery of the longitudinal magnetization can be written as

$$M_z(t) = M_0(1 - e^{-t/T1}) \quad (2.3)$$

where $M_z(t)$ is the longitudinal magnetization at time t and M_0 is the longitudinal magnetization before excitation. The time constant $T1$ is dependent on both the tissue type (e.g. gray matter, white matter, liver) and the magnetic field strength.

Transverse relaxation describes the process by which the magnetization decays in the transverse plane. The decay is caused by a broadening of the precessional frequencies, which results from nuclear spin-spin interactions. The broadening leads to dephasing, which then results in signal decay. Transverse relaxation is characterized by a time constant known as $T2$. Again in the case of a 90° excitation, the decay of the magnetization in the transverse- (xy -) plane can be written as

$$M_{xy}(t) = M_0 e^{-t/T2} \quad (2.4)$$

where $M_{xy}(t)$ is the transverse magnetization at time t .

A closely related time constant is $T2^*$. Whereas $T2$ characterizes only the signal decay due to spin-spin interactions, the $T2^*$ time constant reflects decay due to spin-spin interactions as well as B_0 field inhomogeneity. It is important to

note that the signal decay due to $T2^*$ is only approximated as exponential. $T2$ is dependent on the tissue type, while $T2^*$ is dependent on the tissue type and the magnetic field strength. The estimation of $T2^*$ will play a crucial role in accurately quantitating the relative amounts of water and fat. This is especially true in the liver because iron often coexists with fat. Because $T2^*$ also considers dephasing due to B_0 field inhomogeneity, $T2^* \leq T2$. It is also the case that $T2 \leq T1$.

2.1.4: The Signal Equation and Its Fourier Interpretation

The MR signal is measured using a receiver coil, examples of which are shown in Figure 2.2. The receiver coil is designed to measure magnetic flux changes in the transverse plane that are caused by the precessing spins [50].




Figure 2.2: Three receiver coils that can be used to measure the MR signal. (Left) This is a birdcage coil that can be used for many anatomies such as the head, ankle, and knee. (Middle) This is a torso coil that is typically used to acquire data from the torso. (Right) This is a knee coil that is typically used to acquire data from the knee. (Source: gehealthcare.com).

In a 2D imaging case, the received signal at time t , $s_r(t)$, can be written as

$$s_r(t) = \int_x \int_y m(x, y, t) dx dy \quad (2.5)$$

where $m(x, y, t)$ is a complex-valued quantity that denotes the magnetization at spatiotemporal location (x, y, t) . Taking into account the time-varying gradient field used for spatial encoding and ignoring T_2/T_2^* decay for simplicity, the baseband signal $s(t)$ is

$$s(t) = \int_x \int_y m(x, y) \exp \left[-i\gamma \left(\int_0^t G_x(\tau) d\tau \right) x \right] \exp \left[-i\gamma \left(\int_0^t G_y(\tau) d\tau \right) y \right] dx dy \quad (2.6)$$

where G_x and G_y denote, respectively, the applied gradient fields along the x and y axes. The expression for $s(t)$ can be rewritten as

$$s(t) = \int_x \int_y m(x, y) e^{-i2\pi[k_x(t)x + k_y(t)y]} dx dy \quad (2.7)$$

where

$$k_j(t) = \frac{\gamma}{2\pi} \int_0^t G_j(\tau) d\tau, \quad j = \{x, y\} \quad (2.8)$$

In the form presented in Eq. 2.7, one may realize that the signal $s(t)$ is the Fourier transform of the magnetization $m(x,y)$ evaluated at $[k_x(t),k_y(t)]$ in Fourier space.

That is,

$$s(t) = F[m(x,y)]|_{k_x(t),k_y(t)} \quad (2.9)$$

where F denotes the Fourier transform operator.

The domain of $s(t)$ is commonly referred to in MRI as k-space, where k_x and k_y represent spatial-frequencies within that space. The Fourier interpretation of the signal equation greatly simplifies the understanding of data acquisition and image reconstruction, as will be seen in the next section.

2.1.5: Pulse Sequence and Imaging Principles

A pulse sequence can be thought of as a set of time-dependent instructions that is used to acquire data from an MRI scanner. Figure 2.3 shows a timing diagram and associated trajectory through k-space for a 2D Fourier Transform (2DFT) pulse sequence. By altering the area of the G_y gradient, a different k-space 'line' (from (b) to (c) in the k-space trajectory image) can be acquired. Thus, k-space can be filled in a raster-like fashion by changing the area of the G_y gradient between subsequent excitations. Typically, the area of the G_y gradient is changed such that there is equal spacing of the k-space 'lines' along the k_y -axis.

The 2DFT pulse sequence is one of the most basic and most commonly used sequences in MRI. While more complex sequences can be implemented by altering the shapes of the G_x and G_y gradients, an understanding of the 2DFT sequence suffices for the subject of this dissertation.

Upon acquiring the data, they must be processed into a meaningful form in a procedure known as image reconstruction. One of the beauties of the 2DFT pulse sequence is that the acquired samples lie on a Cartesian grid so that the image can be reconstructed by an inverse 2D Fourier transform. Figure 2.4 shows an example of the k-space data acquired using a 2DFT pulse sequence and the reconstructed image.

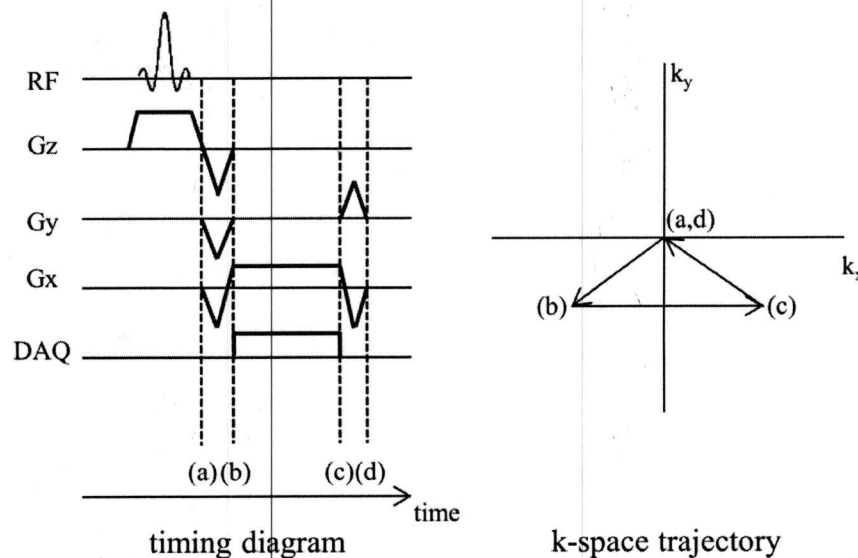


Figure 2.3: An example timing diagram and associated k-space trajectory for a 2D Fourier Transform (2DFT) pulse sequence. The radiofrequency (RF) pulse represents the B1 field that is applied to excite the spins. G_z denotes the z-axis gradient that is applied simultaneously with RF to choose the imaging slice. The G_y and G_x gradients define the trajectory through k-space as seen by the corresponding time points (a-d) in the right image. The area of the G_y gradient is modified between subsequent excitations to acquire a different k-space line along the k_y axis. DAQ stands for data acquisition.

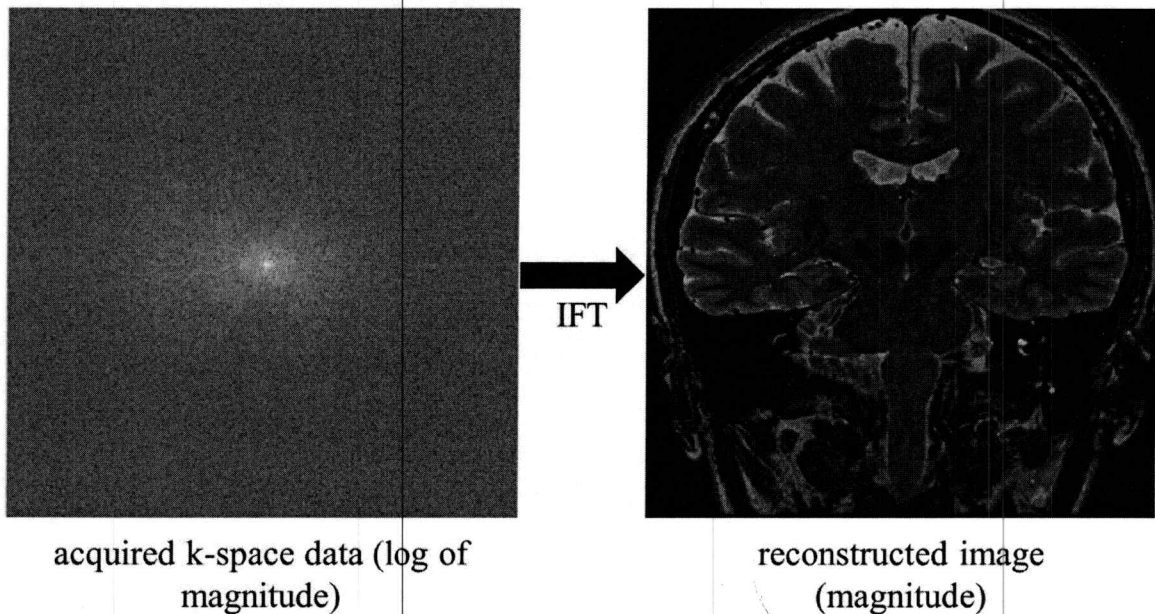


Figure 2.4: The k-space data (left) that were acquired using a 2DFT pulse sequence. The log of the magnitude is shown for better visibility. The image (right) that was reconstructed using an inverse 2D Fourier transform (IFT) of the acquired data. This particular image is a coronal slice of the brain. Both the k-space and image data are complex valued although only the magnitude is shown.

The Fourier interpretation of the signal equation is very helpful when defining the sampling requirements in a 2DFT sequence. The choices of sample spacing and sampling extent will, respectively, dictate the field of view and the resolution of the reconstructed image.

Since sampling occurs in k-space (i.e. the Fourier domain), aliasing can occur in the image domain. To avoid aliasing, the field of view is generally set to be greater than or equal to the extent of the object. Setting the field of view, in turn, determines the k-space sample spacing. Following is the relationship between sample spacing and field of view.

$$\Delta k_j = 1/FOV_j, j = \{x, y\} \quad (2.10)$$

In Eq. 2.10, Δk_j is the sampling spacing along the k_j -axis and FOV_j is the desired field of view along the j -axis.

The spatial resolution is determined by the sampling extent in k-space. The sampling extent can be modeled as a *rect* function in k-space, which corresponds to a *sinc* function in the image domain that introduces blurring. The spatial resolution can be written as

$$\delta_j = 1/2k_{j_max}, j = \{x, y\} \quad (2.11)$$

where k_{j_max} is the maximum sampled spatial frequency (in cm^{-1}) along the k_j -axis.

2.1.6: Off-resonance

Off-resonance is a term that describes a deviation from the expected resonant frequency ($f_0 = \frac{\gamma}{2\pi} B_0$) of a nuclear spin. In this subsection, the sources of off-resonance, their effects on the MR signal, and techniques to address off-resonance will be presented.

Sources of Off-resonance

The sources of off-resonance are main field inhomogeneity, susceptibility-induced variations, and chemical shift.

- 1) Main field inhomogeneity refers to a nonuniformity of the external static magnetic field. The nonuniformity may be caused by any imperfection in the materials used in or the process of building the magnet. Typically, the off-resonance attributed to main field inhomogeneity is relatively small due to the current state of magnet technology.
- 2) Susceptibility-induced variations describe off-resonance that is caused by differences in bulk magnetic susceptibility within the sample. For example, air is paramagnetic, which means that it is attracted to an externally applied magnetic field, whereas water (commonly found in tissue) is diamagnetic, which means that it creates a magnetic field in opposition to an externally applied magnetic field. Susceptibility-induced variations are most severely seen at air-tissue boundaries such as in the sinuses and lungs.
- 3) Chemical shift refers to the off-resonance that is caused by the electron shielding of a nucleus. The shielding slightly alters the magnetic field that is experienced by the nucleus, which shifts the resonant frequency as seen by the relationship in Eq. 2.2. For example, the hydrogen in a fat molecule experiences a different electron shielding than the hydrogen in a water molecule. This difference manifests as a shift in resonant frequency

between hydrogen in fat and hydrogen in water. The topic of chemical shift as it pertains to fat and water will be discussed more thoroughly in later sections.

Effects of Off-resonance

Off-resonance has two main effects on the MR signal. First, a spatially-varying off-resonance implies that neighboring spins will be precessing at different frequencies. This difference can increase the degree of dephasing, which would cause a faster decay of the apparent signal in the transverse plane. The faster rate of decay is captured by a decrease in the value of $T2^*$. Second, off-resonance will create phase errors in the measured MR signal. For a 2DFT pulse sequence, this error will manifest as a shift along the x-axis in the image domain. This is typically a relatively benign image artifact since the shift is often less than one pixel. However, for more complicated k-space trajectories, phase errors can manifest as a less tolerable blurring in the image domain.

Addressing Off-resonance

Two general techniques are used to address off-resonance. One that is specific to main field inhomogeneity is called shimming. Shim coils are found in most modern MRI scanners. Their primary and often only task is to fine-tune the main field to be more uniform. The second technique will be generally referred to as post-processing correction. In post-processing correction, the off-resonance is

estimated from the acquired data so that image artifacts can be subsequently removed. Note that these techniques cannot synthesize data that has been lost due to $T2^*$ decay; rather, they are used more often to correct shifts or blurring in the image. Post-processing correction techniques will be revisited and discussed more thoroughly when the topic of chemical shift encoding is presented.

2.2: Water-Fat Imaging

2.2.1: Clinical Motivation

Water-fat imaging techniques have an important role in clinical MRI. They are often used to suppress fat signals, which appear bright in MR images, so that underlying pathology can be clearly visualized. Fat suppression methods have been used for clearer assessment of meniscal tears [46, 73], for tumor detection [18], for spine imaging [35] and for detection and aging of vertebral fractures [68]. Figure 2.5 shows a liver image without and with fat suppression. The fat-suppressed image clearly shows the metastatic deposit in the left lobe of the liver.

More recently, water-fat imaging has been used to quantify the relative amounts of water and fat within an anatomy of interest [56]. Many studies have used these quantitative imaging techniques for liver and/or pancreatic fat quantification in the diagnosis and management of nonalcoholic fatty liver disease and as a potential biomarker for metabolic diseases such as obesity and

diabetes [27, 49, 28]. Figure 2.6 shows the relative amount of liver fat, termed the fat fraction, in two different subjects.



Figure 2.5: Liver image without (left) and with (right) fat suppression. The fat suppressed image clearly shows the metastatic deposit in the left lobe of the liver (arrows). Images from [36].




Figure 2.6: Liver fat fraction images from two subjects. The fat fraction that was measured using a water-fat imaging technique is shown in each image. For reference, a 5.56% threshold has been cited to discern normal from abnormal liver fat [63]. Images from [27].

2.2.2: MR Differences between Water and Fat

The T1 of fat and that of water are different. Recall that T1 refers to the time that it takes the magnetization to return to equilibrium along the longitudinal axis after

an excitation. Fat has a shorter T1 than water, which is why fat often appears brighter in MR images.

Another MR difference between water and fat is referred to as chemical shift. As shown in Figure 2.7, the appearance of hydrogen in water is different from that in lipid molecules (fat). Due to the different chemical structures, the hydrogen in each of the two molecules experiences a different electron shielding [50]. In the presence of a magnetic field, this difference manifests as a shift in resonant frequency. Figure 2.8 shows a spectrum of a sample containing both water and fat. Water has one peak while fat has multiple peaks in the spectrum. The multiple fat peaks are a result of the different chemical bonds in a glyceride that contain hydrogen.




Figure 2.7: Chemical structures of water (left) and a glyceride (right), which is a main component of a fat molecule. Notice that the glyceride has many types of chemical bonds that contain hydrogen. Image courtesy Houchun H. Hu, PhD.




Figure 2.8: Spectrum of a sample containing both water and fat. The water has one peak while the fat has multiple peaks although the CH₂ peak dominates. The horizontal axis is in parts per million (ppm), which is independent of magnetic field strength. At 3T, the frequency shift of the CH₂ peak relative to the water peak is about -420 Hz. Image courtesy Houchun H. Hu, PhD.

2.2.3: T1-based Method

The short TI inversion recovery (STIR) method is a commonly used T1-based technique for fat suppression [8]. Figure 2.9 is a schematic that displays the concept of the STIR method. An 180° excitation pulse, known as an inversion

pulse, is applied to rotate the magnetization from the +z-axis to the -z-axis. Since the T1 of fat is shorter than that of water, it will recover at a faster rate. The STIR method will image when fat is at the zero-crossing, which nulls the fat component in the resulting image.


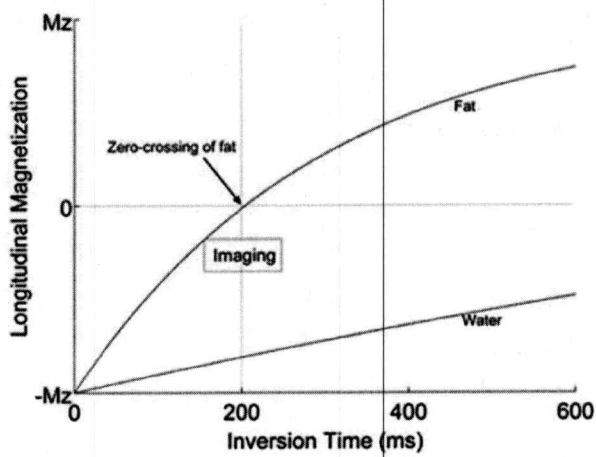


Figure 2.9: The schematic (left) explains the STIR method. After an inversion pulse, the fat recovers more quickly than the water due to the shorter T1 time of the former. At the time (200 ms here) of the zero-crossing of fat, imaging will occur to effectively null the fat component. The image (right) shows an example image using the STIR method. The fat signal was correctly suppressed. The arrows point to areas that may cause problems when using non T1-based techniques. Images from [3].

The main advantage of the STIR method is that it provides reliable fat suppression. It is also insensitive to B0 inhomogeneity. This method is available on most MRI systems and is widely used in the clinical setting.

A disadvantage of the STIR method is its inherent T1 weighting [3]. Since imaging occurs at the time of the fat zero-crossing, other species that have similar T1 times will also, inadvertently, be nulled. As an example, a commonly used contrast agent known as gadolinium reduces the T1 of its associated

species. This could lead to inadvertent nulling of a contrast-enhancing lesion. Another disadvantage is the degradation of signal-to-noise ratio of the water signal since imaging does not occur when the water spins are entirely in the transverse plane. It is estimated that the signal-to-noise ratio is decreased by up to 50% [3].

2.2.4: Chemically Selective Fat Suppression

Chemically shift selective (CHESS) [23] is a commonly used fat suppression technique that, unlike STIR, is based on the chemical shift between water and fat. Although these methods are commonly called *fat* suppression, the same concept can be applied for water suppression. Figure 2.10 shows a schematic of the CHESS method applied to a 2DFT pulse sequence. CHESS applies what is known as a spectrally-selective excitation pulse that excites only the fat spins by 90° . The excited spins are subsequently dephased to null the net magnetization. A standard imaging sequence, such as a 2DFT sequence, immediately follows. The image on the right in Figure 2.5 is an example of fat suppression.

Fat suppression techniques are extremely versatile since they can be appended to most existing pulse sequences. Some variant of the CHESS method is available on most MRI systems and is widely used in the clinical setting.

A primary limitation of fat suppression techniques is their sensitivity to B_0 inhomogeneity. B_0 inhomogeneity will cause an unknown shift in the resonant

frequencies of water and fat that will lead to incomplete fat suppression or worse, inadvertent suppression of the water signal. This severely restricts the use of fat suppression techniques in areas of high B_0 inhomogeneity such as the head and neck and the extremities.

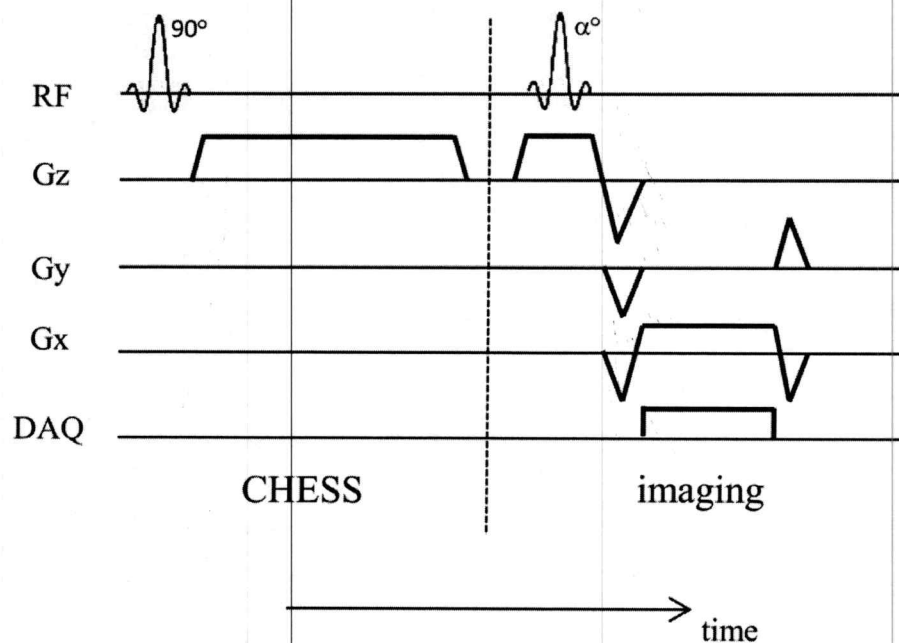


Figure 2.10: A schematic of the chemically shift selective (CHESS) method applied to a 2DFT pulse sequence. The 90° excitation pulse is spectrally-selective so that only the fat spins are excited and tipped into the transverse plane. The G_z gradient is applied to dephase the excited spins, resulting in zero net magnetization. A standard 2DFT imaging sequence immediately follows. Please refer to Figure 2.3 for details of the 2DFT sequence.

2.2.5: Chemical Shift Encoding

Like fat suppression, chemical shift encoding techniques are based on the chemical shift between water and fat. However unlike fat suppression, these techniques have evolved such that they are now insensitive to B_0 inhomogeneity,

which makes them especially useful for imaging areas of the body in which high off-resonance is often encountered [44]. In addition, chemical shift encoding techniques that quantify relative amounts of water and fat have recently been developed [55]. The remainder of this subsection discusses the historical development, the commercial implementations, the current state, and the current limitations of chemical shift encoding.

Historical Development

Chemical shift encoding for water-fat separation was first proposed by Dixon in 1984 [13]. Now commonly referred to as the 'two-point Dixon', this method measures the MR signal at two time points (called echo times), the first when water and fat are in-phase (angle is 0°) and the second when the two species are opposed-phase (angle is 180°). The resulting system of linear equations can be inverted to estimate the water and fat components.

The approach taken by Dixon fundamentally differed from that of fat suppression. Rather than separating the two species during data acquisition, the two-point Dixon method measured functions of both species and then separated them in a post-processing step. This gave the additional benefit of having both 'water-only' and 'fat-only' images rather than just one of them.

A major shortcoming of the two-point Dixon method was its disregard of B_0 inhomogeneity. The frequency shift of both water and fat that is caused by the inhomogeneity means that the two species will not necessarily be in-phase or

opposed-phase when measured. The end result of this sensitivity to B0 inhomogeneity, as with fat suppression, is inaccurate separation of water and fat.

In 1991, Glover et al. proposed an extension to three-points [20]. Signal measurements were made at relative phase shifts of 180° , 0° , and -180° . In the presence of zero B0 inhomogeneity, the third measurement is redundant. However when there exists B0 field inhomogeneity, the third measurement will be different from the other two and all measurements can be used to estimate the B0 inhomogeneity as well as the water and fat signals. Further work by Xiang et al. [69] generalized the data acquisition of the three-point method by allowing for an arbitrary, but common, phase shift between measurements. The water and fat signals were recovered by incorporating the presumed knowledge that the B0 field is smoothly varying and that living tissue is more likely to contain only water than only fat. In an interesting return to the original two-point method, Ma [43] proposed a phase-correction scheme that also assumed a smoothly-varying B0 field. Water and fat estimation was done in order of the confidence of the estimate. This allowed more confident regions to influence the final estimate of less confident regions. Reeder et al. [59, 58] introduced a general reconstruction approach that could be used for data acquisitions with arbitrary phase shifts between measurements. Known as Iterative Decomposition of Water and Fat with Echo Asymmetry and Least-squares Estimation (IDEAL), this method estimates the water and fat signals and the B0 field map independently for each voxel in an iterative manner. For each pixel, the B0 field estimate is first

initialized to zero. The water and fat signals are estimated, with a fixed field map estimate, via linear least-squares estimation. The estimate of the field map is then updated, also via linear least-squares. The process of water and fat estimation and field map update is repeated until predetermined stopping criteria are met.

IDEAL has become the most well-known and widely used chemical shift encoding technique. This is, in part, due to additional works that have shown the benefit of or increased the robustness of IDEAL. Pineda et al. used Cramer-Rao bound analysis to optimize the choice of phase shift between measurements [52]. They then showed via Monte Carlo simulation that the IDEAL reconstruction algorithm achieved this lower bound for the optimal choice of phase shift and for all water:fat ratios. At a similar time, Yu et al. introduced a region-growing scheme for improving the accuracy of the B₀ field inhomogeneity, or field map, estimate [71]. In areas of high B₀ inhomogeneity, the voxel-independent IDEAL (IDEAL-VI) estimation approach tends to converge to local minima field map estimates that can lead to swaps of the water and fat signals. The region-growing method proposed by Yu aimed to avoid local minima estimates by using the presumed knowledge that the true field map is smoothly varying. An initial seed voxel is chosen if it is believed that the field map value can be confidently estimated. The region-growing then moves in a square-spiral trajectory through the image. For each voxel, the initial guess of the B₀ field map is determined by a 2D linear extrapolation of the already-calculated neighboring

field map estimates. Figure 2.11 shows the benefit of incorporating region-growing into the IDEAL estimation.




Figure 2.11: The B0 field map and water image estimates in the brachial plexus using voxel-independent IDEAL (IDEAL-VI) and IDEAL with region-growing (IDEAL-RG). The arrows in the IDEAL-VI images highlight incorrect B0 field map estimates that cause water-fat swaps. The IDEAL-RG estimates do not exhibit these errors. The brachial plexus is typically a difficult area for uniform water-fat separation due to high off-resonance. Images from [71].

Commercial Implementations

IDEAL with region-growing (IDEAL-RG) has been adopted by General Electric (GE) Healthcare. In recent years, GE Healthcare has begun offering IDEAL-RG first as an “investigational-only” and now as a commercial product on many of its MRI systems.

Philips Healthcare has also recently begun offering their version of a chemical shift encoding sequence. Similar to IDEAL-RG, mDixon is available as a commercial product on the Philips MRI systems.

Current State of Chemical Shift Encoding

Further development of chemical shift encoding techniques continues primarily in two areas. The first is in the development of techniques with increased robustness to B_0 inhomogeneity while the other is in the extension of current methods to allow for quantitation of the relative amounts of water and fat.

At the heart of chemical shift encoding is estimation of the B_0 field map. Determining the field map is a challenging task as the least-squares cost function is non-convex and many global minima may potentially exist. Figure 2.12 shows the least-squares cost as a function of B_0 field map value for a voxel taken from an ankle dataset. Without additional knowledge, the field map estimate, when initialized to zero, would converge to a local minimum solution using any gradient descent based algorithm. To guide the estimate to the true value, the additional knowledge that the true field map is spatially smoothly varying is typically incorporated into the reconstruction.




Figure 2.12: The least-squares cost as a function of B0 field map for a voxel taken from an ankle dataset. Without additional knowledge, the field map estimate would be initialized to zero (solid white circle). Any gradient descent based algorithm would converge to a local minimum solution (dashed white circle) that would cause swapping of the water and fat signals. The true B0 field map value is denoted by the solid black circle.

Many methods that focus on the estimate of the B0 field map have been proposed, each essentially differing in how the presumed knowledge of smoothness is used. In addition to the previously mentioned region-growing approach from Yu et al., Berglund et al. have recently developed a multi-seed region-growing scheme [2]. Eggers et al. [17] and Berglund et al. [1] have separately proposed two-point methods with flexible choice of echo times. In their work, Lu and Hargreaves [37] first determine for each pixel all of the B0 field

map values that yield a local minimum in the least-squares cost function. They then use a field map growth scheme to assign estimates to each pixel while ensuring field map smoothness. Jacob and Sutton [30] take a conceptually similar route although the details quickly diverge. In their work, the possible field map values for each pixel are determined via a modification of the harmonic retrieval framework. The ambiguity is resolved by solving a cost function that promotes smoothness of the field map. A different perspective has emerged from the approach proposed by Tsao and Jiang [65]. Rather than considering each individual pixel, this work first treats each of images as one large pixel and estimates one B0 field map value for all pixels. Subsequently, the images are decomposed into smaller regions and the B0 field map value is updated for each region. This approach can be continued until the region size becomes one pixel. Lastly, recent work from Hernando et al. aims to estimate the B0 field map on a finely discretized grid [25]. Employing the VARPRO formulation [24] for dimensionality reduction and a regularization term that penalizes non-smooth B0 field maps, a graph cut-based algorithm is used to estimate the field map.

The recent surge of new chemical shift encoding techniques indicates the popularity of this research area as well as the general consensus that methods for improving the insensitivity to B0 inhomogeneity are still needed. While each of these works certainly shows promising results, no one method has yet stood out from all of the rest.

The other area of recent research interest has been in extending chemical shift encoding techniques to allow for quantitation of the relative amounts of water and fat. Water-fat quantitation has the potential of replacing tissue biopsy as the gold standard for the diagnosis and grading of fatty infiltration in the body. Separate works [9, 27, 28, 72] have demonstrated that quantitative chemical shift encoding techniques can accurately and/or precisely measure liver fat content.

Accurate quantitation requires the compensation of many confounding factors, the most notable of which are the spectral complexity of fat, T1 bias, and T2* decay. In the aforementioned quantitative works, these confounding factors are accounted for in similar manners. The spectral complexity of fat, as was seen in Figure 2.8, can be measured using spectroscopy. The fat spectrum is then assumed to be known and is modeled directly into the reconstruction. The T1 bias is minimized by using a small flip angle (usually between 5-10°) in the excitation. This ensures that the magnetization, comprised of both water and fat, has recovered fully to the longitudinal axis before the subsequent excitation. Note that a larger flip angle may result in only a fully recovered fat component (due to the shorter T1 of fat than water), which would result in T1-weighting that is manifested by a bright fat signal in the image. Lastly, T2* is compensated for by directly estimating its value concurrently with the water, fat, and B0 field map estimates. It has been estimated that approximately 31% of people with liver fat infiltration also exhibit iron overload [19]. The presence of iron increases the local B0 field inhomogeneity that, in turn, increases the rate of signal decay as

captured by the $T2^*$ parameter. Thus, $T2^*$ must be accounted for to have accurate estimates of the relative amounts of water and fat. A common $T2^*$ value is estimated for both water and fat since this has been shown to reduce the estimation error as compared to individual $T2^*$ values for each species [26].

Chemical shift encoding techniques for water-fat quantitation are not yet available as commercial products from any of the MRI vendors. However, "investigational-only" versions have been distributed to select locations for their testing. Given the recent aforementioned works, it is believed that a commercial product will soon be available from at least one of the major vendors.

Current Limitations

The insensitivity to B_0 inhomogeneity and the ability to quantify water and fat of chemical shift encoded techniques requires that measurements be collected at multiple time points, which increases the scan time as compared to a single time point acquisition. Data are acquired at typically three time points for water-fat separation and six time points for water-fat quantitation [72]. The increased scan time often requires that a compromise be made between the spatial resolution of the images, the volumetric coverage of the images, and the tolerance to image artifacts that are caused by subject motion. For example, a high-resolution 3D water-fat MRI scan of the knee can take in the tens of minutes [48]. Any subject motion during this relatively long scan would cause artifacts that may degrade the diagnostic quality of the images. In abdominal imaging, the total scan time,

which directly affects the spatial resolution and the volume coverage, is often limited by the subject's capacity to hold his/her breath. Respiratory-related motion artifacts can severely degrade image quality so a compromise is usually made by limiting the spatial resolution and/or volume coverage.

To address this compromise, techniques for reducing the total scan time have been proposed. The next section discusses accelerated imaging techniques for chemical shift encoding. A general explanation of the underlying theory and the application to chemical shift encoding is presented for each of the acceleration techniques.

2.3: Accelerated Imaging Techniques and Their Applications to Chemical Shift Encoded Water-Fat Imaging

2.3.1: Parallel Imaging

Parallel imaging implies the use of multiple receiver channels that simultaneously acquire functions of the k-space data. Each channel has a distinct complex-valued receive pattern that is related to its location relative to the object of interest [60]. That is, each receiver channel 'sees' the object from a unique perspective. Figure 2.13 shows a schematic of eight receiver channels and the corresponding images from each channel. Using M receiver channels, it is theoretically possible to simultaneously collect M linearly-independent measurements, which can be used to accelerate the data acquisition by up to a factor of M as compared to a single receiver coil.




Figure 2.13: A head coil with eight receiver elements (top left) surrounds the object of interest (bottom left). The receiver elements simultaneously acquire functions of the k-space data. The image from a receiver element is weighted by a distinct spatial function (right) that is related to the element's spatial position relative to the object. The eight-channel head coil can, in theory, simultaneously acquire eight linearly-independent measurements that can be used to accelerate a data acquisition by up to a factor of eight.

The two most commonly used parallel imaging techniques are called SENSE [54] and GRAPPA [21]. These two techniques differ in the manner by which undersampled multichannel k-space data are processed to recover the object of interest.

Sensitivity encoding (SENSE) relies on the knowledge of the complex-valued receive pattern for each channel, which is commonly referred to as the coil sensitivity. The coil sensitivities are estimated either from a separate calibration scan or from a region within the acquired data itself. The sensitivity information is then used to undo the aliasing that is caused by undersampling k-space. Figure 2.14 explains how SENSE is used to recover an image that has been undersampled by a factor of two using two receiver elements. Because of

the manner in which the aliasing is undone, SENSE is typically thought of as an image domain based method.

Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) is the other commonly used parallel imaging technique. In contrast to SENSE, the GRAPPA reconstruction is performed in the k-space domain by assuming that each non-acquired k-space sample can be synthesized by a linear combination of neighboring acquired samples. The weights in the linear combination are estimated from a calibration region within the acquired region itself. Figure 2.15 illustrates the GRAPPA reconstruction.

Recently, Lustig et al. have proposed a GRAPPA-like technique, called Iterative Self-consistent Parallel Imaging Reconstruction from Arbitrary k-Space (SPIRiT) [42] that instead imposes the constraint that each k-space point, whether acquired or not, be a linear combination of its both acquired and non-acquired neighbors. The constraint leads to a set of linear equations that can be solved to recover the k-space for each coil. Figure 2.16 illustrates the SPIRiT reconstruction.

Though not yet as widely used as SENSE and GRAPPA, SPIRiT is presented because it will be referred to in later chapters. Note that both GRAPPA and SPIRiT are typically thought of as k-space based methods due to the manner in which the missing data are recovered.

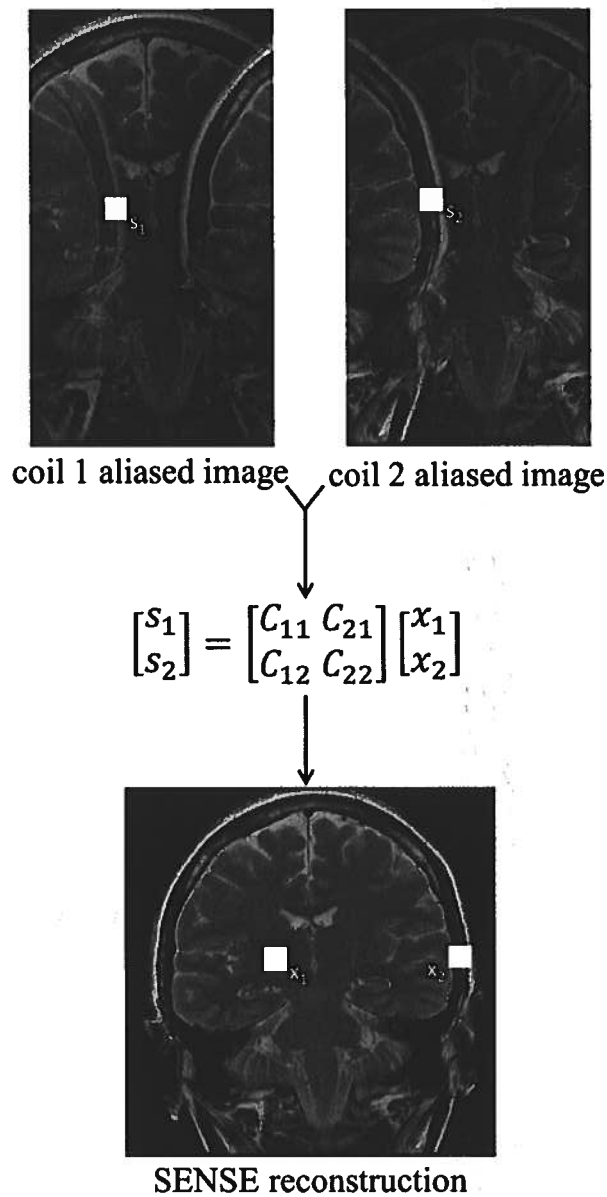


Figure 2.14: Accelerating the data acquisition can be done by uniformly undersampling k-space. This results in aliasing in the image domain (top). In this example, k-space was undersampled by a factor of two, so each pixel in the aliased image is a weighted sum of two pixels from the underlying image. If one estimates the coil sensitivities (denoted by C_{xy}), a set of linear equations can be constructed that relates the underlying image pixels to the aliased measurements (see equation). The aliasing can be done by solving the equation, which results in the SENSE reconstructed image (bottom).




Figure 2.15: The GRAPPA reconstruction synthesizes non-acquired k-space points (red) by a linear combination of neighboring acquired k-space samples across all coils (gray). The weights that are used in the linear combination are estimated from a calibration region within the acquired data itself. After the non-acquired k-space samples from each individual coil have been synthesized, the individual coil images are calculated via the Fourier transform. The resulting images can then be combined via root-sum-of-squares. Image from [42].




Figure 2.16: The SPIRiT reconstruction imposes the constraint that each k-space point, acquired (gray) or not (red), be a linear combination of neighboring k-space samples across all coils. Like GRAPPA, the weights that are used in the linear combination are estimated from a calibration region within the acquired data itself. After the reconstruction, the individual coil images are calculated via the Fourier transform. The resulting images are then combined via root-sum-of-squares. Image from [42].

Many works have applied parallel imaging to chemical shift encoding in a modular format [27, 45, 70]. First, images at each echo time are recovered from the undersampled k-space data using either SENSE or GRAPPA. The synthesized fully-sampled echo time images are then input into any water-fat imaging routine, such as IDEAL-RG. While straightforward to implement, it is not well understood how any magnitude or phase errors in the parallel imaging reconstruction would manifest in the subsequent water-fat step. Phase errors, especially, could lead to inaccurate separation and/or quantitation of the two species.

2.3.2: Compressed Sensing

The theory underlying compressed sensing (CS) suggests that MR images can be recovered with a high degree of fidelity if the images are known to be compressible. Put another way, if the image can be accurately represented using far fewer transform coefficients than in its native (i.e. image) domain, the number of measurements needed to recover the image should also be far fewer than is conventionally thought. In separate seminal works, Candes et al. [11] and Donoho [15] showed that a signal $x \in \mathcal{R}^N$ could be perfectly recovered from $K \ll N$ measurements under two conditions: sparsity and incoherence.

Sparsity refers to the ability to represent a signal with relatively few nonzero coefficients. In other words, if the signal x can be fully represented with only $S \ll N$ nonzero values, then x is said to be a sparse signal. Note that the signal need not be sparse in its native domain. For example, an MR image may be sparse in the wavelet domain, in which case the wavelet basis is referred to as the sparsifying system.

Incoherence is a measure of the 'spread' of the measurement vectors in the sparsifying system [10]. In essence, this measure seeks to quantify the dissimilarity between the measurement system Φ and the sparsifying system Θ . Incoherence is written as

$$\mu(\Phi, \Theta) = \max_{k,j} |\langle \phi_k, \theta_j \rangle| \quad (2.12)$$

With overwhelming probability, a signal with S nonzero entries can be perfectly recovered with a number of measurements

$$K \geq C\mu^2 S \log(N) \quad (2.13)$$

for some positive constant C .

MR images are typically not sparse but may be highly compressible. In this case, the claim made in Eq. 2.13 does not hold but a result from Candes et al. [11] provides an upper bound on the ℓ_2 -norm of the reconstruction error. Figure 2.17 shows the compressibility of an axial brain MR image using three different transforms. It is seen that relatively few wavelet coefficients are needed for a visually near-perfect image reconstruction.




Figure 2.17: The compressibility of an axial brain image using a discrete cosine transform (DCT), a Daubechies-4 wavelet transform, and a finite difference transform. Selected percentages of the highest-magnitude coefficients are used to reconstruct the image (middle three columns). The wavelet transform provides high compressibility of the image. Images from [39].

Lustig et al. first introduced the application of compressed sensing into MRI [39]. The authors proposed a variable-density sampling scheme in which the central k-space region is sampled with higher probability than the periphery. This was proposed to satisfy incoherence requirements for compressed sensing while also incorporating the information that the energy of MR images is concentrated in the low spatial frequencies. From an undersampled k-space acquisition, the image was estimated via Equation 2.14

$$\min_{\mathbf{m}} \|\mathbf{k} - \mathbf{F}_u \mathbf{m}\|_2^2 + \lambda \|\Theta \mathbf{m}\|_1 \quad (2.14)$$

where \mathbf{k} represents the undersampled k-space measurements, \mathbf{F}_u denotes the undersampled Fourier transform, \mathbf{m} is the desired image, Θ is the sparsifying transform and λ is the regularization parameter that represents the tradeoff between the data fidelity term and the ℓ_1 -regularization term. The expression in Eq. 2.14 is a convex optimization problem that can be solved in a variety of ways [5].

Compressed sensing has been applied to chemical shift encoded water-fat separation by exploiting the presumed compressibility of the water and fat images and the B0 field map using the finite difference transform [14]. This represents a departure from the conventional approach for water-fat separation in which estimation was done on a per-voxel basis. In this work, the authors have proposed to directly embed the compressed sensing reconstruction into the

water-fat separation. The potential for error propagation, as was discussed in the previous parallel imaging subsection, is not a concern in this reconstruction.

2.3.3: Additional Acceleration Techniques

Two additional acceleration techniques are discussed here. They can be used to accelerate water-fat MRI, but they are not directly related to this dissertation work.

Partial k-Space Techniques

Partial k-space techniques exploit the conjugate symmetry property of the Fourier transform [53]. The underlying assumption used by these methods is that the true image is real-valued and thus its k-space representation should exhibit conjugate symmetry. In theory, one would need to acquire only one-half of k-space and enforce conjugate symmetry to synthesize the other half. In practice, slightly more than one-half of k-space is acquired so that a low resolution estimate of phase can be demodulated before k-space synthesis.

Reeder et al. [57] proposed a homodyne reconstruction technique for water-fat separation. The central region of k-space is used to obtain low resolution estimates of the B0 field map as well as the complex-valued water and fat images. The estimated phases from the B0 field map and the water and fat images are used to demodulate the acquired signal. A conventional water-fat separation method, such as IDEAL [58], can then be used to estimate the true

water and fat images from all of the acquired data. The need for a central k-space region to estimate the phase necessarily limits this technique to acceleration factors of less than two.

Non-Cartesian Acquisitions

The 2DFT pulse sequence presented earlier is said to be a Cartesian acquisition because it samples k-space on a Cartesian grid. Non-Cartesian k-space trajectories exist and are used in MRI generally to reduce the data acquisition time. Figure 2.18 shows the k-space coverage of a radial and a spiral trajectory, which are two commonly used non-Cartesian acquisitions.

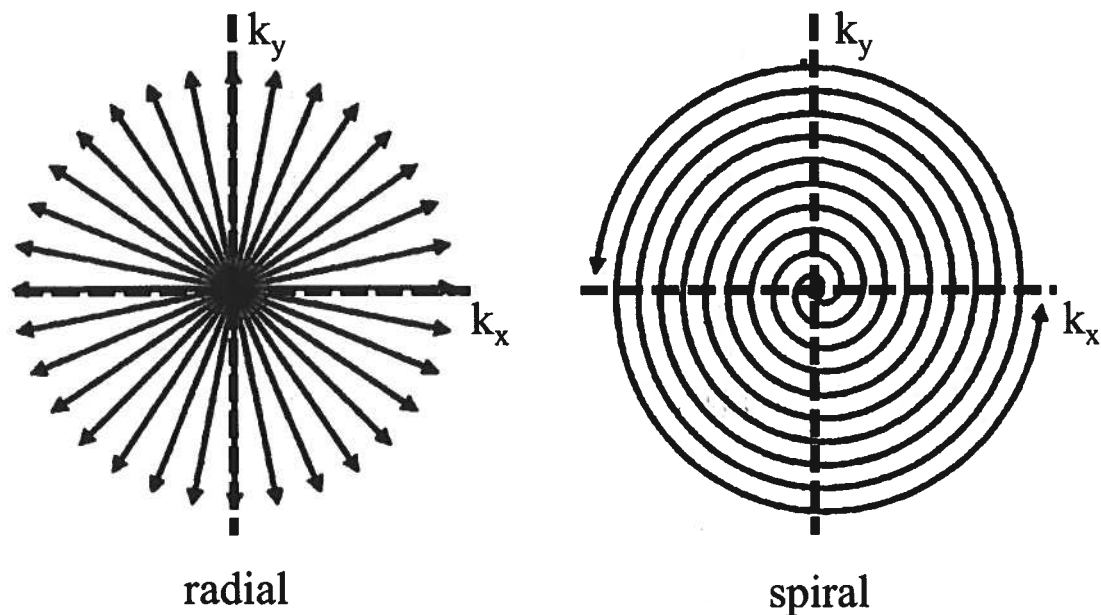


Figure 2.18: Radial and spiral are two common non-Cartesian trajectories that are used to reduce the data acquisition time. Each segment that ends with an arrowhead represents the k-space that is covered after one excitation pulse. The use of non-Cartesian trajectories may require advanced image reconstruction techniques that compensate for off-resonance effects. Images from [7].

Many works have been proposed that use non-Cartesian acquisitions for chemical shift encoded water-fat separation [4, 7, 67]. Unlike in the Cartesian setting, off-resonance during a non-Cartesian acquisition will typically cause image blurring. By accounting for the unique time that each k-space sample was collected, the off-resonance effects on the decomposed water and fat images can be removed.

2.3.4: Summary

Figure 2.19 summarizes the different acceleration techniques that have been applied to chemical shift encoded water-fat imaging. Each of the works exploits only one acceleration technique. In addition, those works that use parallel imaging do so in a sequential fashion, which introduces the possibility of error propagation into the water-fat imaging step.

This dissertation introduces a new approach for chemical shift encoded water-fat imaging that exploits both parallel imaging and compressed sensing. The two techniques for acceleration are integrated directly into the signal model and reconstruction method so that estimation occurs directly from the undersampled k-space data. In addition, a novel method for B0 field map estimation is presented. The method is general such that it can also be directly integrated into the parallel imaging and compressed sensing reconstruction framework.

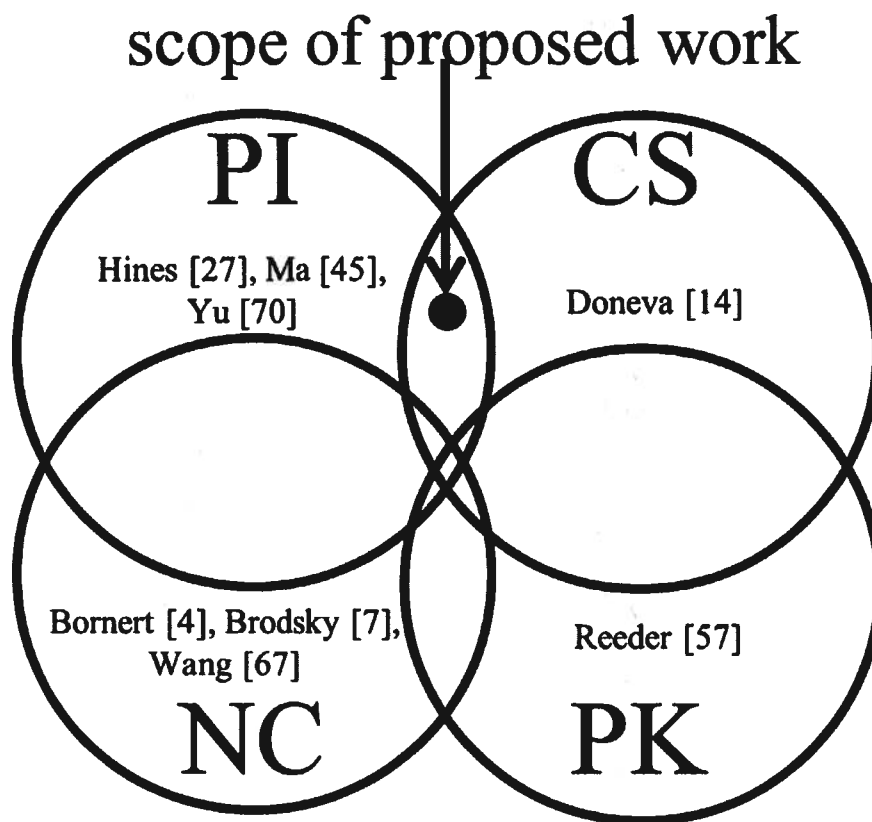


Figure 2.19: This Venn diagram depicts the different acceleration techniques that have been applied to chemical shift encoded water-fat imaging. Only the first author of each work is shown. PI – parallel imaging, CS – compressed sensing, NC – non-Cartesian, PK – partial k-space. In this dissertation, both parallel imaging and compressed sensing are used to accelerate chemical shift encoded water-fat imaging.

Chapter 3: Multiscale B-splines for B0 Field Map Estimation

3.1: Introduction

Recall that the B0 field map represents the off-resonance that is due to main field inhomogeneity and susceptibility-induced variations. This off-resonance causes a frequency shift in the water-fat spectrum so that the species may not be precessing at the desired frequency. An accurate estimate of the field map, or at least compensation of its phase effects, is necessary for accurate qualitative and quantitative measures of water and fat. An interesting aspect of water-fat imaging is that while water and fat are the primary signals of interest, the real technical challenge lies in estimating the B0 field map. This is because the least-squares cost as a function of field map value is non-convex with many local, and potentially many global, minima. Additional knowledge about the behavior of the B0 field map is often required to reconcile this ambiguity.

In the previous chapter, many works were presented that proposed methods to handle the ambiguity associated with estimating the B0 field map [2, 7, 25, 30, 37, 43, 65, 71]. While each of these methods certainly showed their robustness in particular applications, they also assumed that k-space is fully-sampled and therefore that one has access to the fully-sampled echo time images. This assumption is violated in an undersampled k-space acquisition, which prohibits the direct translation of these methods to this new setting.

This chapter presents a novel multiscale B-spline method for estimating the B0 field map. The method is general in the sense that it can be directly applied to arbitrarily undersampled Cartesian k-space acquisitions. A fully-sampled k-space or even a fully-sampled central region is not a prerequisite for this approach as it is with most other B0 field map estimation methods. The use of multiscale B-splines provides an accurate and compact representation of the B0 field map that avoids local minima estimates that would cause a water-fat swap.

The remainder of this chapter discusses the details of the proposed method. A general discussion, with no assumption of fully-sampled or undersampled k-space, will be presented. Results will be shown using a variety of fully-sampled datasets because the method for estimating the water and fat images from undersampled data has not yet been presented (this will be done in Chapter 4). However, it will be important to remember that the presented method can be directly applied to an undersampled k-space acquisition once the technique for estimating the water and fat images has been established.

3.2: Methods

3.2.1: Multiscale B-splines

In the proposed approach, the B0 field map is gradually updated using B-spline functions of successively finer scales. The coarse scale updates of the field map are used primarily to avoid local minima estimates that would cause a water-fat

swap while, of course, providing a rough approximation of the true B0 field map. The fine scale updates are then used to provide a more accurate estimate of the true B0 field map.

Tsao and Jiang proposed a multiresolution approach for B0 field map estimation using fully-sampled echo time images [65]. In this work, they first treated each of the echo time images as one large voxel for which one constant field map value was estimated. Each image was then decomposed into smaller regions and the B0 field map estimate was updated for each region. The decomposition and B0 field map update continued until the region size was one voxel. Though not presented in this form, this approach can be viewed as a gradual update of the B0 field map estimate using overlapping rectangular functions.

Inspired by their work, the proposed method seeks to restrict the subspace dimension in which the B0 field map estimate is updated. The subspace dimension is gradually increased to allow for a more accurate B0 field map estimate, but the dimension of the space always remains less than the number of field map voxels. This provides the benefit of estimating the B0 field map from a reduced number of measurements. Further, the vectors that span the subspace promote a spatially smoothly varying B0 field map.

The approach is first explained conceptually in 1D using a B0 field map of 256 pixels. Initially, strict restrictions are placed on the subspace in which the field map is updated. Figure 3.1a shows the one vector that spans the coarsest

scale (call it level 0) subspace. The single constant-valued vector restricts the B0 field map estimate to assuming only one value that is equal for all pixels. The field map is updated in this subspace. Subsequently, the subspace is enlarged such that it spans a space of greater dimension. Figure 3.1b shows the three vectors that span this level 1 subspace. The current B0 field map estimate, which originated from the level 0 subspace, is updated in the level 1 subspace. The same procedure is repeated using the vectors shown in Figure 3.1c, which span the level 2 subspace. In general, the current B0 field map estimate, which comes from the level $m-1$ subspace, is updated using the level m subspace vectors.




Figure 3.1: A 1D example to illustrate the vectors that span the different subspaces in which the B0 field map estimate is updated. (a) One vector that spans the level 0 subspace. The B0 field map estimate will be a constant value for all pixels. (b) The three vectors that span the level 1 subspace. The current field map estimate, which originated from the level 0 subspace, will be updated in this subspace. (c) The five vectors that span the level 2 subspace. As in the previous case, the current B0 field map estimate will be updated in this subspace.

The 1D example extends in a quite straightforward manner to higher dimensions. For example in 2D, a subspace element is created via an outer

product of one vector from each dimension within the same level. Referring to the previous example, the level 0 subspace would be represented by a constant-valued 2D vector while the level 1 subspace would be spanned by 9 ($=3^2$) 2D vectors.

The functions shown in Figures 3.1b/c are linear B-splines. A common argument, although personally believed to be a bit unfounded, is that the piecewise linear nature of the linear B-splines is insufficient for capturing the smoothly varying nature of the true B0 field map. It is important to realize that as the subspace vectors go to finer and finer scales, a locally linear approximation of the B0 field map may be valid. This appears to be the case as will be seen in later results. In addition, Irarrazabal et al. have shown that piecewise linear functions can provide accurate representations of off-resonance [29]. Nevertheless, the generality of the presented approach allows for the use of arbitrary functions that many will agree can accurately represent the true B0 field map.

Higher-order B-splines can be created from linear B-splines using the Cox-de Boor recursion formula. These higher-order B-spline functions may better approximate the true B0 field map. For example, Skare et al. have reported that cubic B-splines provide an accurate and compact representation of the B0 field inhomogeneity that arises from metallic implants [61]. The authors cite the local support of B-splines as a primary advantage since it provides a compact representation of the field inhomogeneity even at higher spatial

frequencies. It appears that this finding may be translated to the current problem at hand of estimating the B0 field map in water-fat imaging.

A 1D cubic B-spline is defined as

$$b(t) = \begin{cases} \frac{2}{3} - \left(1 - \frac{|t|}{2}\right) t^2 & 0 \leq |t| \leq 1 \\ \frac{(2-|t|)^3}{6} & 1 < |t| \leq 2 \\ 0 & \text{else} \end{cases} \quad (3.1)$$

This function is nonzero only on the interval $(-2, 2)$. To create the base 1D cubic B-spline, $b(t)$ is uniformly sampled on the interval $(-2, 2)$ such that the number of sampled points equals the desired B-spline size s (e.g. 128 pixels). A knot spacing, h_t , is defined as

$$h_t = \text{round}\left(\frac{s-1}{4}\right) \quad (3.2)$$

To create a base 2D cubic B-spline of size $M \times N$, two cubic B-splines, one with length M and the other with length N , and their associated knot spacing are first calculated. The base 2D cubic B-spline is then created via an outer product of the two 1D cubic B-spline functions. The 2D cubic B-spline set is created by spatially shifting the base 2D cubic B-spline by all combinations of the multiples (both positive and negative) of the knot spacing in both dimensions. For a $I \times J$ (e.g. 256 x 192 pixels) field map, the base 2D cubic B-spline is shifted in both

dimensions until it is entirely zero in the $I \times J$ image. Shifting by the knot spacing ensures that the sum of the B-spline set at any spatial position is equal to one. The creation of a base cubic B-spline and associated set in higher dimensions can be done using a straightforward extension of the 2D example presented here. Notice that the cubic B-spline is a separable function, which may potentially be exploited for a more computationally efficient implementation. Figure 3.2 shows example 2D cubic B-splines and their corresponding B-spline sets from coarser to finer scale. Just as in the previous discussion with linear B-splines, the same procedure for updating the B_0 field map at successively finer scales can be implemented using the cubic B-spline sets.

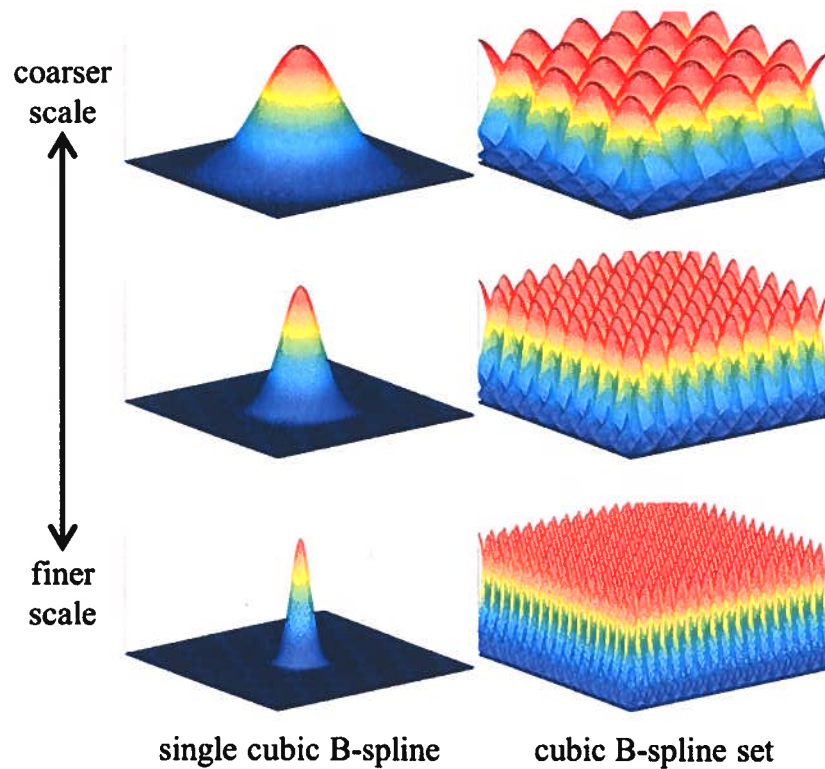


Figure 3.2: Single cubic B-spline functions from coarser scale (top) to finer scale (bottom) and their corresponding cubic B-spline sets. The B-spline set is created by spatially shifting the single cubic B-spline in both spatial directions by multiples of the knot spacing. The cubic B-splines are nonnegative everywhere and the set sums to one at all spatial positions. The field map update term at the m^{th} scale is restricted to be in the space that is spanned by the m^{th} cubic B-spline set.

Having presented the conceptual approach, the discussion turns to the mathematical details. The following subsections discuss the signal model and estimation routine. An understanding of both subsections will be important as future chapters will build upon these concepts.

3.2.2: Signal Model

Equation 3.3 presents the signal model that relates the unknown water and fat images (ρ) and the unknown B0 field map (represented in Ψ) directly to the acquired k-space measurements (\mathbf{k}) at all echo times, in the presence of additive white Gaussian noise.

$$\mathbf{k} = \mathbf{F}\Psi\mathbf{A}\rho + \mathbf{N}(\mathbf{0}, \sigma^2) \quad (3.3)$$

The vector ρ is a concatenation of the water image and the fat image (i.e. $\rho = [\rho_w^1, \rho_f^1, \rho_w^2, \rho_f^2, \dots, \rho_w^P, \rho_f^P]$) where the subscript denotes the species and the superscript denotes the pixel index. The matrix \mathbf{A} operates pixel-wise and a six-peak fat spectrum with known relative amplitudes and frequency shifts is assumed [72]. This fat spectrum was found via spectroscopy. The block diagonal matrix Ψ contains $\exp(j2\pi\psi^p t_n)$ on the p^{th} diagonal of the n^{th} block, where ψ^p is the field map value (in Hertz) at the p^{th} pixel and t_n is the time (in seconds) of the n^{th} echo. Note that the product $\Psi\mathbf{A}\rho$ represents the echo time images from all coils. The matrix \mathbf{F} represents the Fourier transform and a mask operator that selects only the k-space points that were acquired. This transform is applied individually to each echo time image to yield the corresponding k-space measurements.

The signal model presented in Equation 3.3 represents a departure from the traditional single-pixel model [59]. In a sense, the image model is a

generalization of the single-pixel model because the latter assumes that k-space has been fully-sampled so that an inverse Fourier transform can be applied to yield the echo times images. The image model does not make this assumption. It may seem that having to model the images rather than individual pixels is a disadvantage, but in fact, it will be shown to be the opposite.

3.2.3: Signal Estimation

The water and fat images and the B0 field map are decomposed via an alternating minimization similar to the method proposed by Reeder et al. [59]. In each step, the water and fat images are estimated using the current estimate of the B0 field map and then the B0 field map estimate is updated with the current water and fat images fixed. For now, it is assumed that a method exists for estimating the water and fat images. When k-space is fully-sampled, the water and fat images can be estimated using linear least-squares, which yields a closed-form solution. The issue of an undersampled k-space will be discussed in Chapter 4.

To derive the expression for the B0 field map update ($\Delta\psi$), the signal model in Equation 3.3 is rewritten to separate the current estimates from the (unknown) error in each estimate.

$$\mathbf{k} = \mathbf{F}(\hat{\Psi} \cdot \Delta\Psi)\mathbf{A}(\hat{\rho} + \Delta\rho) + \mathbf{N}(\mathbf{0}, \sigma^2) \quad (3.4)$$

The term $\hat{\rho}$ denotes the current estimate of the water and fat images while $\Delta\rho$ represents the error in the estimate. The same convention is used for the B0 field map-dependent term, Ψ . Taking the first order Taylor approximation for the terms in $\Delta\Psi$ results in

$$\mathbf{k} = \mathbf{F}(\hat{\Psi} \cdot \Delta\mathbf{T})\mathbf{A}(\hat{\rho} + \Delta\rho) + \mathbf{N}(\mathbf{0}, \sigma^2) \quad (3.5)$$

where $\Delta\mathbf{T}$ is a block diagonal matrix that contains $1 + j2\pi\Delta\psi^p t_n$ on the p^{th} diagonal of the n^{th} block, where $\Delta\psi^p$ is the field map update term at the p^{th} pixel and t_n is the time of the n^{th} echo. At this point, Equation 3.5 is modified by grouping the known terms on the left hand side, the unknown terms on the right hand side, and discarding the $\Delta\mathbf{S} \cdot \Delta\rho$ term, where $\Delta\mathbf{S}$ contains $j2\pi\Delta\psi^p t_n$ on the p^{th} diagonal of the n^{th} block.

$$\mathbf{k} - \mathbf{F}\hat{\Psi}\mathbf{A}\hat{\rho} = \mathbf{F}\hat{\Psi}\mathbf{A}\Delta\rho + \mathbf{F}\hat{\Psi}\Delta\mathbf{S}\mathbf{A}\hat{\rho} + \mathbf{N}(\mathbf{0}, \sigma^2) \quad (3.6)$$

Notice that the left hand side of Equation 3.6 is the residual error between the measurements and the current estimate while the right hand side is a function of the errors in the water, fat, and B0 field map estimates as well as noise. Denoting the left hand side as \mathbf{r} , the right hand side as $\mathbf{x}(\Delta\psi, \Delta\rho)$, the expression for the B0 field map update in the level m subspace can be written as

$$\min_{\Delta\psi, \Delta\rho} \|\mathbf{r} - \mathbf{x}(\Delta\psi, \Delta\rho)\|_2^2 \text{ s. t. } \Delta\psi \in \text{span}\{\mathbf{B}_m\} \quad (3.7)$$

where \mathbf{B}_m denotes the subspace vectors at the m^{th} level. In words, the convex minimization problem in Equation 3.7 seeks to find the B0 field map update term that minimizes the least-squares cost while also being restricted to reside in the space that is spanned by the m^{th} B-spline set. This expression resembles that proposed by Reeder et al. [59], however there are significant differences. First, the entire B0 field map is simultaneously updated rather than on a per-pixel basis. Relatedly, because the whole field map is simultaneously updated, the *a priori* knowledge of smooth variations can be exploited. The convex minimization problem is solved using conjugate gradients (code courtesy of Michael Lustig, PhD, University of California Berkeley).

The following algorithm summarizes the alternating minimization approach for estimating the B0 field map and the water and fat images. Again, at this point, it is assumed that a method exists for estimating the water and fat images given a current estimate of the B0 field map (Step 2).

1. Initialize the B0 field map (ψ) to zero and select the $m = 0$ level subspace vectors
2. Estimate the water and fat images with current B0 field map estimate
3. Update the B0 field map estimate ($\Delta\psi$) using subspace vectors at the current (m^{th}) level. This will update Ψ .
4. IF $\max |\Delta\psi| < \varepsilon$ (e.g. 1Hz)
 - Go to Step 5
 ELSE
 - Go to Step 2
5. IF current level is equal to predefined end level
 - Done
 ELSE
 - Update subspace to be one finer scale ($m = m + 1$)
 - Go to step 2

3.2.4: In Vivo Experiments

The proposed multiscale B-spline approach was tested using a variety of datasets from different anatomical regions of the body. All datasets were collected with volunteer consent. A majority of the datasets were acquired using a GE Signa Excite HDxt 3T MRI system located at the USC Healthcare Consultation II (HCC-II). The data were collected using an investigational version of the GE IDEAL 3D spoiled gradient echo sequence. The raw k-space

data were saved and were securely transferred to the MREL network. Image reconstruction was done offline using Matlab (The Mathworks, Inc., Natick, MA). Datasets were also provided courtesy of Peter Kellman PhD, National Institutes of Health. These datasets were processed in the same manner as those collected at HCC-II. Table 3.1 presents the applicable data acquisition parameters for each of the datasets.

Anatomy	Matrix Size	TE1 (ms)	Δ TE (ms)
Ankle	256x256	2.184	0.794
Liver	256x256	2.184	0.796
Brachial Plexus	192x192	2.184	0.796
Heart	256x192	1.55	2.27

Table 3.1: The data acquisition parameters for each of the datasets are presented in this table. The first three datasets were collected on a GE 3T MRI system at HCC-II. The heart dataset, courtesy of Peter Kellman PhD, was acquired on a 1.5T Siemens MRI system.

The B-spline support was scaled by a factor of three-fourths between successive scales. Using the ankle dataset as an example, the support of each B-spline in the level 1 subspace was 256x256 pixels. The support of each B-spline in the level 2 subspace was 192x192 pixels, 144x144 pixels in level 3, and so on. The support size in the final level was empirically chosen to be 16 pixels in each dimension.

3.3: Results

Figure 3.3 shows B0 field map, water, and fat estimates using the voxel-independent IDEAL method (IDEAL-VI) [59], IDEAL with region-growing (IDEAL-RG) [58, 71], and the proposed multiscale B-spline (MSBS) method for an ankle dataset. The IDEAL-VI method suffers from water-fat swaps that are ultimately caused by incorrect estimates of the B0 field map. This example demonstrates the need for additional information in estimating the B0 field map, such as its smoothly varying nature. The proposed MSBS estimates closely match the reference IDEAL-RG estimates. It is important to remember that the benefit of the MSBS approach as compared to IDEAL-RG will become apparent in later chapters.

Figure 3.4 shows the B0 field map evolution using the MSBS method. In the early iterations, the B0 field map estimate appears to be a coarse approximation to the true value. At these stages, it is vital that the potential for converging to local minima solutions be minimized. The later iterations refine the B0 field map estimate. The possibility for converging to a local minimum, if it has not already occurred, is quite small.

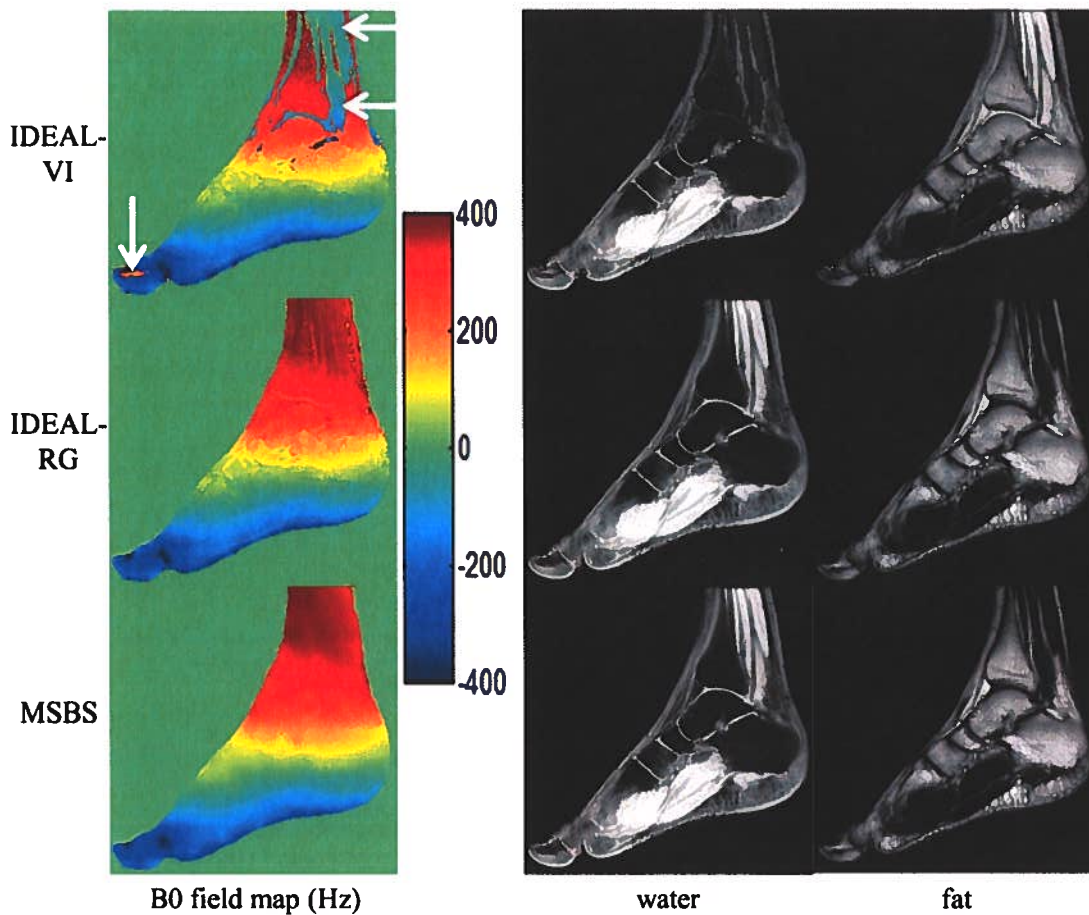


Figure 3.3: The B0 field map, water, and fat estimates using IDEAL-VI, IDEAL-RG and the proposed multiscale B-spline (MSBS) method. Notice the errors in the IDEAL-VI B0 field map estimate (white arrows). These errors propagate into the water and fat images producing swaps. That is, the estimated fat image contains signal from the water component and vice versa. The IDEAL-RG estimates serve as the reference. The proposed multiscale B-spline method produces estimates that are very similar to IDEAL-RG. It is important to remember that the benefit of the MSBS approach as compared to IDEAL-RG will become apparent in later chapters.




Figure 3.4: The evolution of the B0 field map estimate using the proposed MSBS method. The early iterations using coarse scale B-splines are needed to reduce the possibility of converging to a local minimum solution, while of course providing a rough approximation of the B0 field map. The later iterations provide refinement to the field map estimate. It is unlikely that convergence to a local minima solution will originate at these finer scales.

Figure 3.5 shows the B0 field map, water, and fat estimates for a liver dataset using IDEAL-RG and the proposed MSBS method. The IDEAL-VI

estimates are not shown because they were similar to the IDEAL-RG estimates. This example is shown primarily because image estimates from retrospectively undersampling this dataset will be shown in the next chapter. It will be shown that the B0 field map, as well as the water and fat image, estimates are quite similar in the non-accelerated and the accelerated cases.




Figure 3.5: The B0 field map, water, and fat image estimates of the liver using IDEAL-RG and the proposed MSBS method. The estimates using the proposed technique appear similar to the reference IDEAL-RG implementation. In the next chapter, image estimates will be shown using the same dataset that has been retrospectively undersampled to simulate an accelerated acquisition. It will be seen that the estimates between the two cases are quite similar.

Figure 3.6 shows the B0 field map, water, and fat estimates of the brachial plexus using IDEAL-RG and the proposed MSBS method. Errors in the IDEAL-RG B0 field map estimate are seen that subsequently cause water-fat swaps. This type of error when using IDEAL-RG is caused by the lack of connecting tissue regions and is a commonly cited shortcoming of this method. It is

interesting to see that the proposed MSBS method appears to have correctly estimated the B0 field map leading to proper water-fat separation.

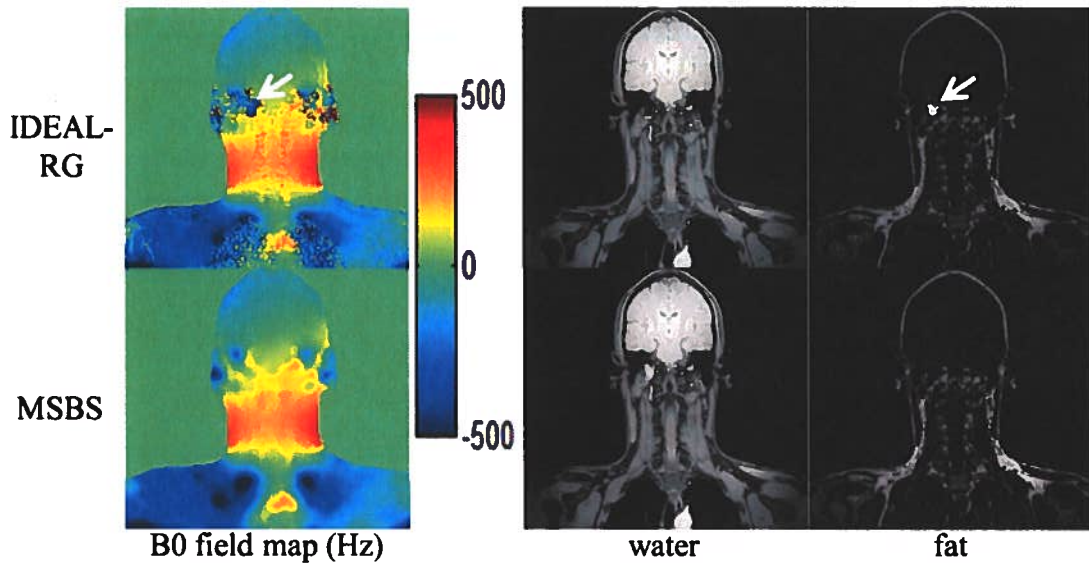


Figure 3.6: The B0 field map, water, and fat image estimates of the brachial plexus using IDEAL-RG and the proposed MSBS method. The brachial plexus is typically a challenging area for water-fat separation due to the high degree of off-resonance in this anatomy. The IDEAL-RG estimate exhibits a water-fat swap due to an incorrect B0 field map estimate (white arrows). Notice that this error occurs near a low signal region. The lack of a connecting signal region may have caused this swap. This is one commonly cited shortcoming of the IDEAL-RG implementation. The proposed MSBS accurately estimates the B0 field map and properly separated water and fat.

Figure 3.7 shows the B0 field map, water, and fat image estimates of a short axis view of the heart using IDEAL-VI, IDEAL-RG, and the proposed MSBS method. This dataset was collected on a Siemens 1.5T MRI system, which is different from the previous datasets in both manufacturer and main field strength. Both of the B0 field map estimates using IDEAL-VI and IDEAL-RG exhibit a region of inaccuracy that causes water-fat swaps. For IDEAL-RG, this error is

likely because the algorithm attempted to region grow through regions of low SNR. The estimates using the proposed MSBS method appear to yield proper water-fat separation.

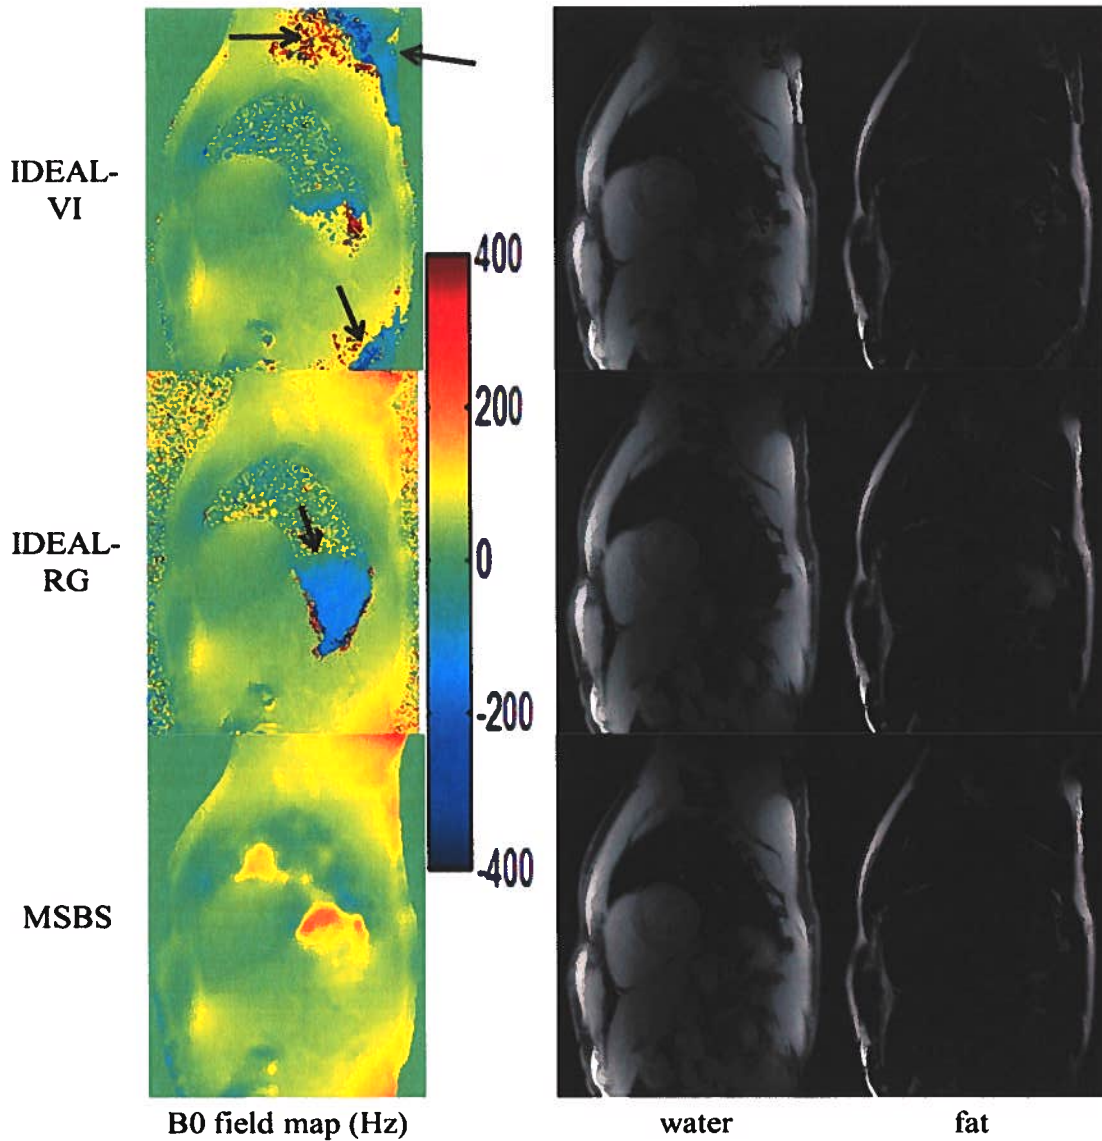


Figure 3.7: The B0 field map, water, and fat image estimates of a short axis view of the heart using IDEAL-VI, IDEAL-RG and the proposed MSBS method. The IDEAL-VI estimates exhibit regions of water-fat swaps due to errors in the B0 field map (black arrows). Water-fat swaps are also seen in the IDEAL-RG estimates. This is likely because the algorithm attempted to region grow through areas of low SNR. The proposed MSBS method accurately estimated the B0 field map, which led to proper water-fat separation. (Dataset courtesy of Peter Kellman PhD, National Institutes of Health)

3.4: Discussion

The proposed multiscale B-spline approach for B0 field map estimation avoids the local minima solutions that plague voxel-independent methods. This was evident in datasets from a variety of anatomies. In addition, it appears that the MSBS method was able to properly separate water and fat in cases where swaps were seen when implementing the commonly-used IDEAL-RG approach.

Generally, the B0 field map estimate using the proposed method appeared smoother than that from IDEAL-RG. This is likely a direct result of using restricted subspaces for the field map updates. Empirically, this degree of smoothing does not seem to affect the quality of the water-fat separation. By further increasing the dimension of the subspace, one may expect the field map estimates to approach those seen from IDEAL-RG.

Updating the field map in a restricted subspace significantly reduces the number of coefficients that must be estimated. This becomes extremely convenient if the B0 field map must be estimated from undersampled measurements. For example, the estimation of the B0 field map in the ankle dataset was finally computed with 529 ($= 23^2$) linear B-spline vectors. With an image size of 256 x 256 pixels, this represents a 'compression-ratio' of 0.8%. In 3D imaging at 256³ pixels, the number of coefficients to calculate would be 12,167 ($=23^3$), which is a ratio of 0.07%. By restricting the field map estimate to particular coefficients, the proposed method may achieve a greater acceleration than those approaches that impose ℓ_1 -regularization on the field map [14]. This is

because ℓ_1 -regularization methods must pay a penalty of additional measurements since the locations of significant coefficients are unknown prior to estimation. In contrast, the proposed method specifies which coefficients to estimate (i.e. those coefficients in the restricted subspace).

Chapter 4: Accelerated Water-Fat Separation Using Parallel Imaging and Compressed Sensing

In the previous chapter, a new technique was proposed for estimating the B0 field map. While no assumption was made about the data sampling as it pertained to estimation of the B0 field map, it was assumed that the water and fat images could be estimated via linear least-squares. This assumption is violated when k-space is undersampled. In this chapter, a framework is proposed for estimating the water and fat images from an undersampled k-space acquisition. After providing the motivation and necessary background, the details of the proposed method and results are presented.

4.1: Introduction

The insensitivity to off-resonance of chemical shift encoded techniques requires that data be acquired at multiple echo times, which increases the scan time as compared to a single echo acquisition. For example, a three-point acquisition for water-fat separation can lengthen the scan time by a factor of three to achieve the SNR-optimal echo spacing [52] while also balancing the spectral field-of-view. The increased scan time often requires that a compromise be made between the spatial resolution, the volume coverage, and the tolerance to motion artifacts. Often, the spatial resolution and volume coverage must be restricted to avoid any image artifacts that would be caused by subject motion.

compressed sensing [39] are not constrained to avoid large regions of non-sampled points.

To satisfy both requirements, a Poisson disk sampling pattern has been used in previous work [42]. Poisson disk sampling was introduced in the computer graphics community many years ago as a way to avoid the structured aliasing artifacts that occur when sampling at a regular interval under the Nyquist limit [12]. The artifacts that results from Poisson disk sampling appear much less objectionable visually than structured aliasing.

In this work, the Poisson disk sampling scheme is adopted for the joint parallel imaging and compressed sensing framework. Figure 4.1 shows two grids that are each of size 192x160 pixels. One of the grids is uniformly undersampled and the other is sampled with a Poisson disk scheme. Also shown in the figure are the corresponding PSF for each of the k-space sampling schemes. The central 24x24 region is fully sampled as it is used for kernel calibration (discussed later). The net acceleration of each sampling pattern is 7.7x. Similar to the uniform undersampling, the Poisson disk sampling does not contain large regions of non-sampled points. However, in contrast, the PSF of the Poisson disk pattern has much smaller sidelobe magnitudes, which implies a higher degree of incoherence. This, in turn, leads to incoherent artifacts that can be largely minimized using compressed sensing.




Figure 4.1: Uniformly undersampled (US) and Poisson disk (PD) k-space sampling patterns and their corresponding magnitude of the point-spread function (PSF). Both schemes sample the central 24×24 region for kernel calibration and result in 7.7x acceleration. Additionally, both schemes avoid large gaps of non-sampling points in k-space, which is desirable for parallel imaging. However, the structure in the US PSF leads to the coherence aliasing artifacts that should be avoided for a compressed sensing reconstruction. The PD PSF leads to the incoherent, noise-like artifacts that are desired for a compressed sensing reconstruction.

4.2.3: Signal Estimation

The water and fat images and the B0 field map are estimated from the undersampled k-space measurements using parallel imaging, compressed sensing, and multiscale B-splines. First, coil sensitivity maps are derived from the SPIRiT [42] k-space convolution kernel. Next, the coil-combined water and fat images and the B0 field map are estimated in an alternating manner. The water and fat images are estimated with a ℓ_1 -regularization on their respective transform coefficients and the B0 field map is estimated using the previously described multiscale B-splines. The following subsections describe in detail the application of parallel imaging and compressed sensing in the estimation routine.

Parallel Imaging

Traditionally, a distinction in parallel imaging has been made between image-domain methods such as SENSE [54] and k-space methods such as GRAPPA [21]. While differences between the two types of methods exist [22], they are often used interchangeably with little difference in reconstruction quality. Thus, it appears that the often made categorization may be invalid. Recently, Lai et al. [32] and Lustig et al. [41] have addressed the superficiality of the distinction between image-domain and k-space methods by showing that the coil sensitivity maps that are used in image-domain methods can be derived from the convolution kernels that are used in k-space methods. Thus, any 'k-space'

parallel imaging method can be equivalently implemented as an ‘image-domain’ method.

In this work, the coil sensitivity maps (\mathbf{C}) are derived from the estimated SPIRiT k-space kernel. This allows direct incorporation of the parallel imaging constraint into the data fidelity term in the reconstruction. The following mathematical derivation explains how this is done.

The calibration consistency expression proposed by Lustig et al. [42] is

$$\mathbf{k}_i = \sum_{j=1}^{N_c} \mathbf{g}_{ij} \otimes \mathbf{k}_j, \quad \forall i \quad (4.2)$$

where \mathbf{k}_i represents the fully-sampled k-space from the i^{th} coil, \mathbf{g}_{ij} is the convolution kernel, and \otimes denotes the circular convolution operator.

Taking the inverse Fourier transform of Equation 4.2 yields

$$\mathbf{I}_i = \sum_{j=1}^{N_c} \mathbf{G}_{ij} .* \mathbf{I}_j, \quad \forall i \quad (4.3)$$

where \mathbf{I}_i is the i^{th} coil image, \mathbf{G}_{ij} is the inverse Fourier transform of \mathbf{g}_{ij} , and $.*$ denotes voxel-wise multiplication.

Equation 4.3 can be rearranged to be written in terms of each voxel instead of each coil image as

$$\begin{bmatrix} I_1(x, y) \\ I_2(x, y) \\ \vdots \\ I_{N_c}(x, y) \end{bmatrix} = \begin{bmatrix} G_{11}(x, y) & G_{12}(x, y) & \dots & G_{1N_c}(x, y) \\ G_{21}(x, y) & G_{22}(x, y) & \dots & G_{2N_c}(x, y) \\ \vdots & \vdots & \ddots & \vdots \\ G_{N_c1}(x, y) & G_{N_c2}(x, y) & \dots & G_{N_cN_c}(x, y) \end{bmatrix} \begin{bmatrix} I_1(x, y) \\ I_2(x, y) \\ \vdots \\ I_{N_c}(x, y) \end{bmatrix} \quad \forall x, y \quad (4.4)$$

Equation 4.4 can be concisely written as

$$\mathbf{v}_{xy} = \mathbf{H}_{xy} \cdot \mathbf{v}_{xy} \quad \forall x, y \quad (4.5)$$

Equation 4.5 still represents the calibration consistency expression, but in a different form than in Equation 4.2. For a particular voxel (and dropping the subscripts in Equation 4.5 for brevity's sake), Equation 4.5 says that \mathbf{v} should be in the space spanned by the eigenvectors of \mathbf{H} that have an associated eigenvalue equal to one. With no FOV overlap, Lai et al. [32] have shown that \mathbf{H} will have one eigenvalue that is equal to one. In this case, \mathbf{v} is a scaled version of \mathbf{e}_1 , which is the eigenvector with eigenvalue equal to one and where the scalar $\alpha \in \mathbb{C}^1$, as seen in equation 4.6.

$$\mathbf{v} = \alpha \cdot \mathbf{e}_1 \quad (4.6)$$

The following equation can also be written for \mathbf{v} (still omitting subscripts).

$$\mathbf{v} = \begin{bmatrix} I_1 \\ I_2 \\ \cdot \\ \cdot \\ I_{N_c} \end{bmatrix} = m \cdot \begin{bmatrix} C_1 \\ C_2 \\ \cdot \\ \cdot \\ C_{N_c} \end{bmatrix} \quad (4.7)$$

where m is the magnetization and C_i denotes the i^{th} coil sensitivity. Equating terms in Equations 4.6 and 4.7, it is seen that \mathbf{e}_1 contains the coil sensitivities.

Figure 4.2 shows the coil images and derived coil sensitivities for an example liver dataset. The coil images were derived using the 16 central phase encodes lines for autocalibration and with a 7×7 SPIRiT k-space kernel.

Once the coil sensitivity maps have been derived, they remain fixed throughout the remainder of the reconstruction. An advantage of using the derived coil sensitivities is that the SPIRiT calibration consistency constraint is implicitly imposed whereas imposing this constraint in k-space would require an explicit calibration consistency expression in the reconstruction. The absence of the explicit consistency expression reduces the computational complexity of the iterative reconstruction.




Figure 4.2: The top row in each set of images shows the fully sampled coil images and the bottom row shows the corresponding coil sensitivities (magnitude) that were derived from the SPIRiT k-space kernel using only the central 16 phase encode lines. Using the coil sensitivities implicitly imposes the SPIRiT constraint, which avoids the need for an explicit calibration consistency expression in the reconstruction.

Compressed Sensing

The undersampled k-space measurements are modeled as a linear function of the coil-combined water and fat images. The linear map $F_u C \Psi A$ is, in general, unknown because the coil sensitivity maps and B0 field map are unknown. However, given an estimate of C and Ψ , the linear mapping becomes known and the water and fat images can be estimated using compressed sensing by

exploiting their presumed compressibility in a predetermined linear transform. Modeled in this way, the estimation of the water and fat images is a fairly standard CS-MRI problem.

Specifically, given \mathbf{C} and Ψ , the water and fat images are estimated via Equation 4.8, where \mathbf{W} represents the Daubechies-4 wavelet transform that operates separately on the water and fat images. The experiments from Liang et al. [33] have shown that this transform provides high compressibility of many MRI images. The sparsity constraints are imposed on the coil-combined water and fat images rather than on the individual coil images because the latter potentially have less sparsity due to coil sensitivity-dependent magnitude and phase variations

$$\min_{\rho} \|\mathbf{k}_u - \mathbf{F}_u \mathbf{C} \Psi \mathbf{A} \rho\|_2^2 + \lambda \|\mathbf{W} \rho\|_1 \quad (4.8)$$

It is assumed that the noise is independently and identically distributed (i.e. $\Sigma = \sigma^2 \mathbf{I}$). The parameter λ balances data fidelity with transform sparsity. The expression in Equation 4.8 is a convex optimization problem that is solved using conjugate gradients.

Alternating Minimization

The water and fat images and the B0 field map are estimated via the algorithm that was presented at the end of §3.2.3. Now, however, the water and fat

images are estimated by solving the expression in Equation 4.8, rather than assuming that a linear least-squares estimate was sufficient. Figure 4.3 presents a flowchart that describes that alternating minimization procedure for estimating the water and fat images, and the B0 field map.

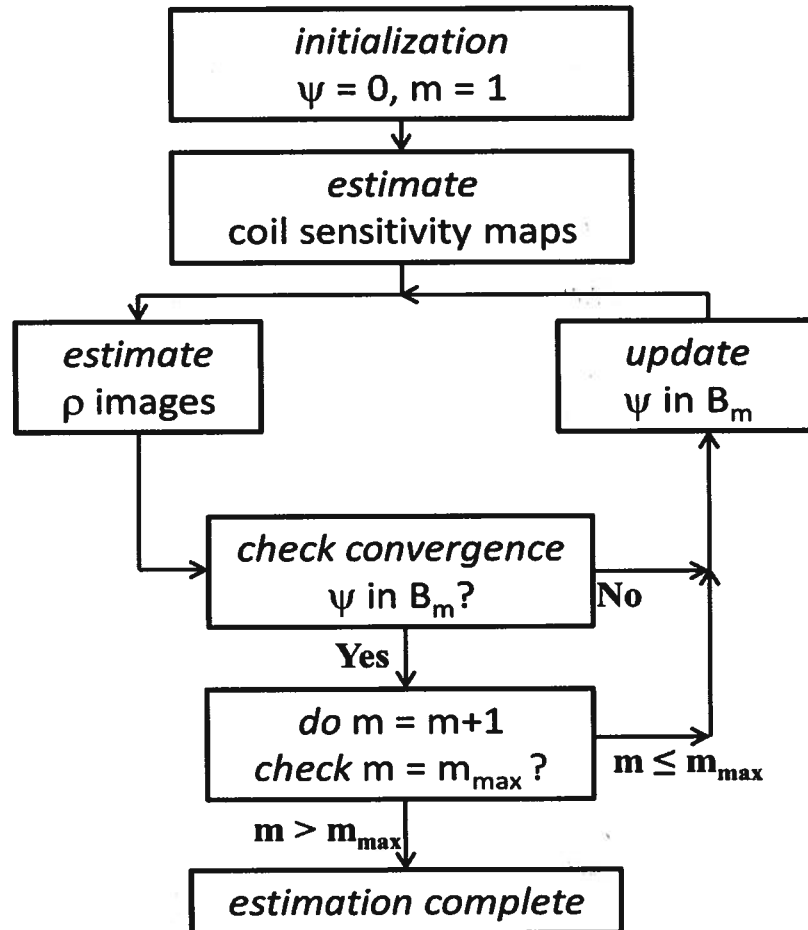


Figure 4.3: After initialization, the coil sensitivity maps are estimated and then held fixed. Subsequently, the water and fat images and the B0 field map are estimated via an alternating minimization. The water and fat images are estimated using parallel imaging and compressed sensing while the B0 field map is updated using multiscale cubic B-splines.

It is worth noting that the addition of parallel imaging and compressed sensing requires no conceptual change to the multiscale B-spline approach for estimating the B0 field map. Slight mathematical modifications are, of course, needed to account for the coil sensitivity information but the routine is unchanged.

4.2.4: In Vivo Experiments

Experiments were conducted with volunteer consent on a GE Signa EXCITE HDx 3T system (GE Healthcare, Waukesha, WI) at USC HCC-II. Fully sampled measurements were collected at three echo times using an investigational GE IDEAL 3D spoiled gradient echo sequence with one TE per TR and a sampling bandwidth of ± 125 kHz. The SNR-optimal echo spacing of 0.796ms required that the echo-train-length (ETL) was equal to one. This resulted in a three times increase in the scan time as compared to a single echo acquisition. The corrections proposed by Lu et al. [38] were not necessary since bipolar readouts were not used. In addition to the fully sampled datasets, two prospectively undersampled datasets, one of the thigh and the other of the knee, were also collected using a custom-built sequence.

To test the proposed approach, the fully sampled datasets were retrospectively undersampled using a Poisson disk sampling pattern. A liver dataset was undersampled along the phase encode axis, keeping the central 16 phase encode lines for kernel calibration. A brachial plexus, thigh, and knee

dataset were each undersampled along the phase encode and slice encode axes, keeping the central 24x24 region for kernel calibration. The coil sensitivity maps were then derived from the 7x7 SPIRiT k-space kernel. The coarsest scale B-spline had a 1D support size equal to the size of the corresponding image dimension (i.e. for an $M \times N$ image, the coarsest scale B-spline was of size $M \times N$). The support size of the B-spline functions was scaled by a factor of three-fourths in each dimension between successive scales. The finest scale B-spline had a 1D support size equal to 16. The proposed approach will hereafter be referred to as PI-CS-MSBS.

To compare the proposed approach, an existing parallel imaging and water-fat separation technique was implemented. To the author's knowledge, this is the most often used technique for accelerated water-fat separation in current MRI scanners. The fully sampled data were uniformly undersampled, keeping the same calibration region as used for the proposed approach. The retrospectively undersampled k-space data were then reconstructed using Autocalibrated Reconstruction for Cartesian Sampling (ARC) [6] with a 7x7 ARC k-space kernel. Subsequently, the data were coil combined [66] and then water-fat separation was done using IDEAL with region-growing. This reconstruction will hereafter be referred to as ARC/IDEAL-RG.

To serve as a reference, the fully sampled data were first coil combined and then passed to the IDEAL-RG reconstruction. This reconstruction will

hereafter be referred to as IDEAL-RG. Some of the datasets were also reconstructed using voxel-independent IDEAL (IDEAL-VI).

4.3: Results

Figure 4.4 shows the B0 field map, water, and fat estimates of the liver using 1x IDEAL-RG, 3.4x ARC/IDEAL-RG, and 3.4x PI-CS-MSBS. The 3.4x net acceleration used in ARC/IDEAL-RG was a result of 4x outer acceleration factor and 16 fully sampled central phase encode lines. The arrows in the ARC/IDEAL-RG estimates point to aliasing artifacts and the arrowheads highlight regions of noise amplification. This was the same liver dataset that was shown in the previous chapter. Notice that the B0 field map estimates are quite similar when using the proposed approach as 1x versus at 3.4x. This is not too surprising given the oversampling factor related to the B0 field map that was discussed in §3.4.

Figure 4.5 shows the B0 field map, water, and fat estimates of the brachial plexus using 1x IDEAL-VI, 1x MSBS, 6.4x ARC/MSBS, and 6.4x PI-CS-MSBS. In this example, the proposed multiscale B-spline approach was used in place of IDEAL-RG at 1x for B0 field map estimation since the latter exhibited water-fat swaps as was shown in the previous chapter. The 6.4x net acceleration used for ARC/MSBS was a result of 3x by 3x outer acceleration factor and the 24x24 fully sampled central region. The white ellipses in the IDEAL-VI B0 field map outline

areas of incorrect estimates that cause water-fat swaps. The white arrowheads point to regions of noise amplification in the ARC/MSBS estimates.

Figure 4.6 shows the B0 field map, water, and fat estimates of the thigh using 1x IDEAL-RG, 7.7x ARC/IDEAL-RG, and 8.4x PI-CS-MSBS (prospective downsampling). The 7.7x net acceleration was a result of 3x by 3x outer acceleration factor and the 24x24 fully sampled central region. The accelerated acquisition reduced the scan time from 10:11 to 01:12. The ARC/IDEAL-RG estimates exhibit noticeable artifacts as a result of unresolved aliasing (white arrows). This was expected since the outer acceleration factor was greater than the number of coils.

Figure 4.7 shows the B0 field map, water, and fat estimates of the knee using 1x IDEAL-RG, 7.7x ARC/IDEAL-RG (retrospective), 8.6x PI-CS-MSBS (retrospective), and 8.6x PI-CS-MSBS (prospective). The same Poisson disk pattern was used for both the retrospective and prospective undersampling. The arrows and arrowheads in the ARC/IDEAL-RG estimates highlight, respectively, artifacts and areas of noise amplification. Both the retrospective and prospective estimates using the proposed method are presented to demonstrate that system issues, such as eddy currents, associated with the k-space sampling order do not affect the quality of water-fat separation.




Figure 4.4: The B0 field map, water, and fat estimates of the liver that was acquired using an eight-channel torso coil. The ARC/IDEAL-RG estimates exhibit unresolved aliasing artifacts (arrows) and noise amplification (arrowhead). These artifacts were anticipated since the outer acceleration factor of 4 is greater than the number of coils along the axis of undersampling. In contrast, the estimates using the proposed method exhibit only slight incoherent artifacts as a result of the Poisson disk sampling and ℓ_1 -regularization.

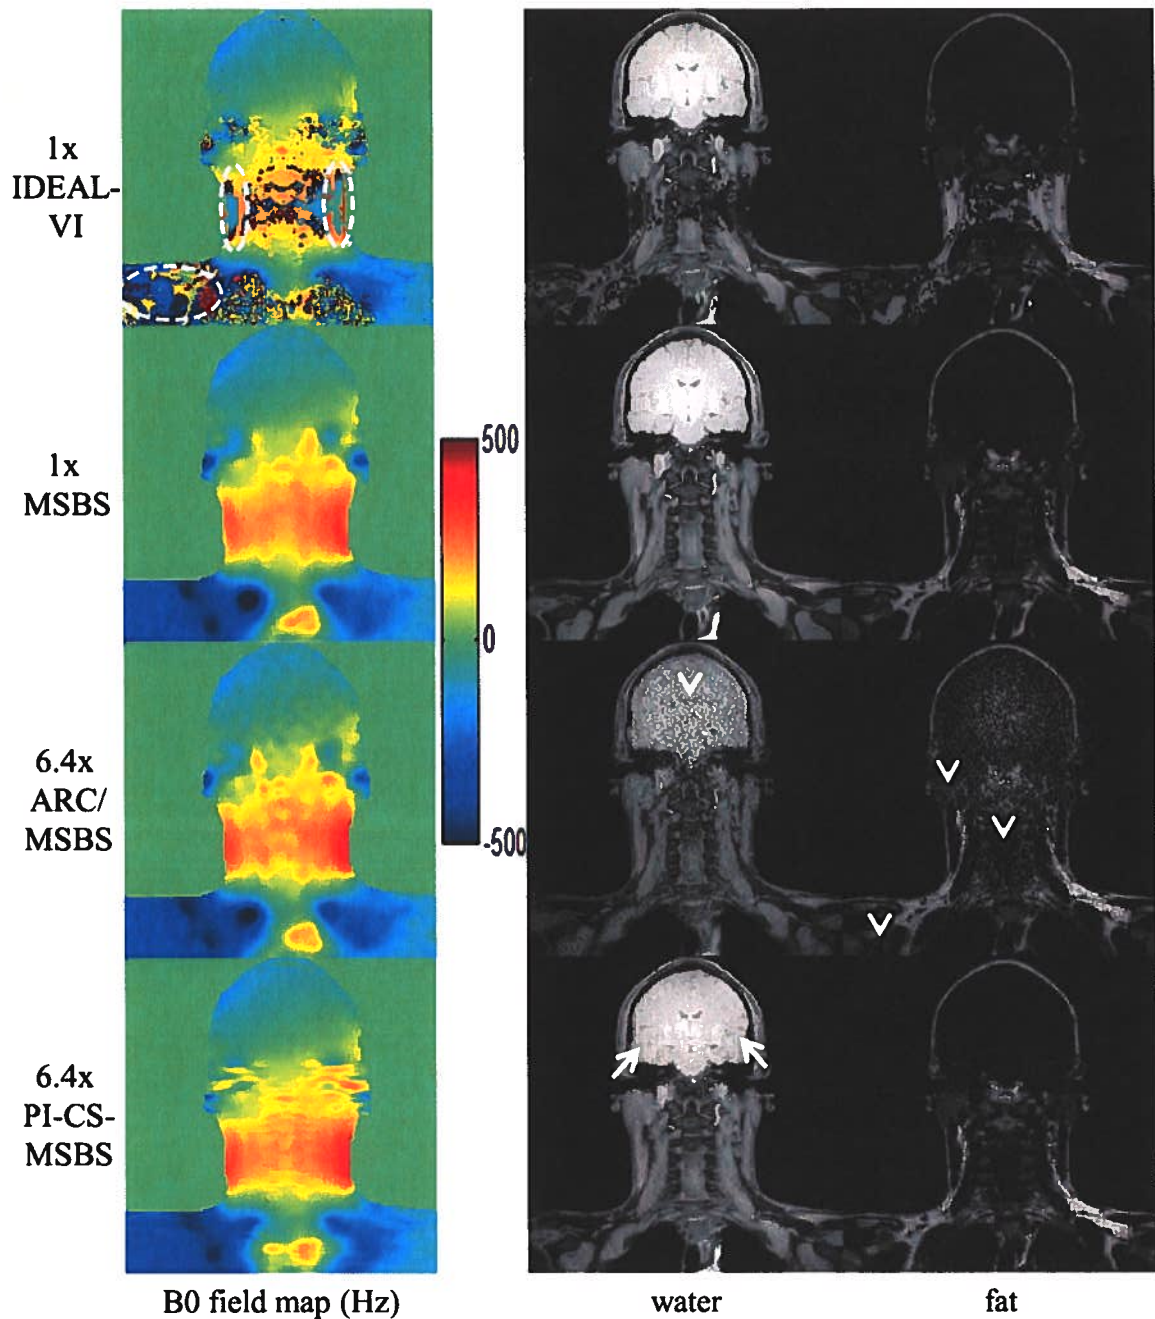


Figure 4.5: The B0 field map, water, and fat estimates of the brachial plexus acquired using an eight-channel neurovascular coil. Field map estimate errors using IDEAL-VI (white ellipses) cause water-fat swaps. The MSBS field map estimation approach correctly estimates the field map to avoid the swaps. The ARC/MSBS estimates exhibit noise artifacts (arrowheads), especially in the brain. These artifacts were expected since the 3x by 3x outer acceleration factor is greater than the number of receiver elements. The estimates using the proposed PI-CS-MSBS approach exhibit a relatively reduced level of artifacts. Slight loss of subtle features is seen (white arrows).

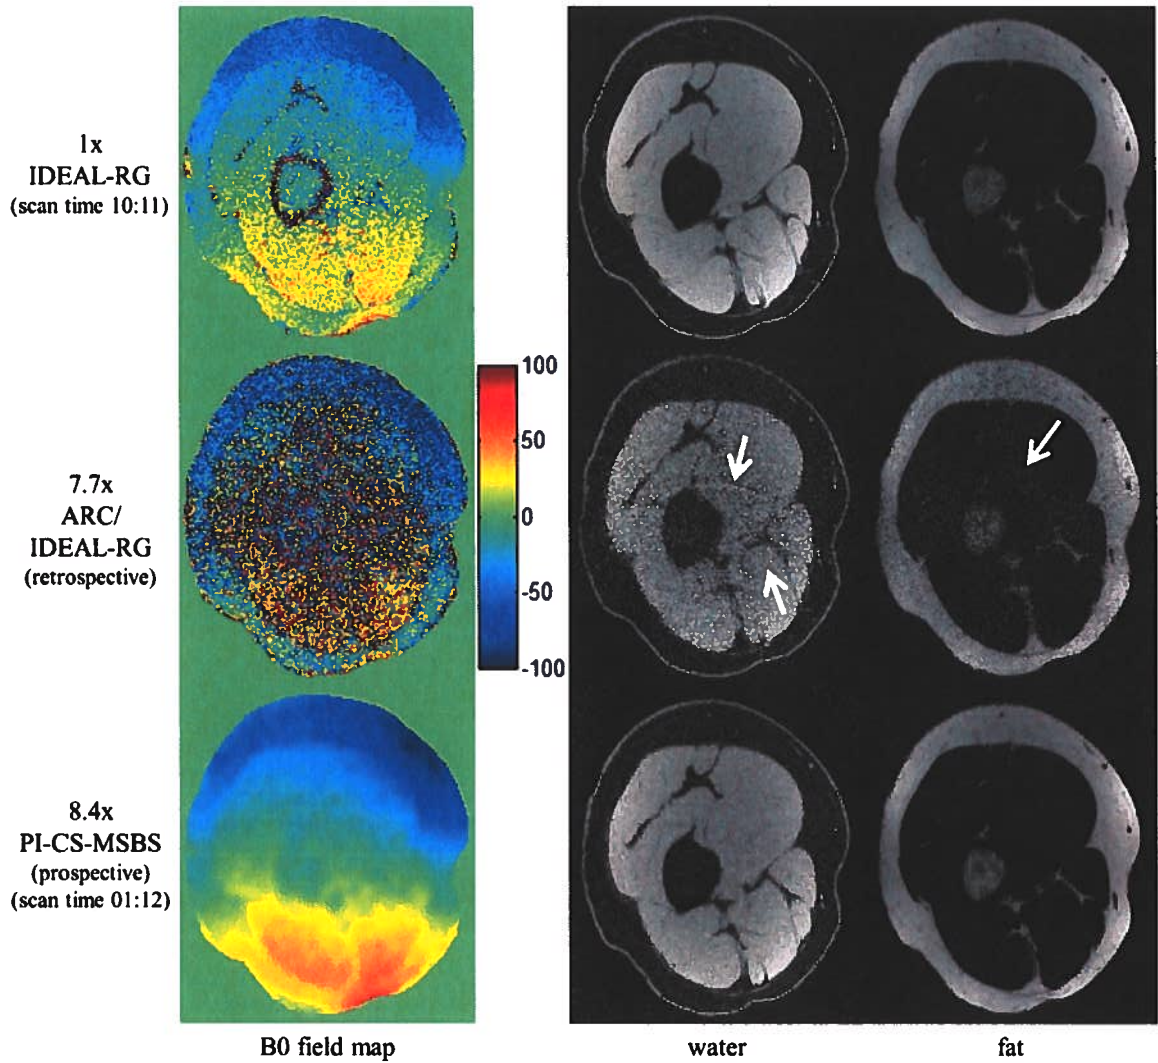


Figure 4.6: The ARC/IDEAL-RG estimates exhibit noticeable artifacts as a result of unresolved aliasing (white arrows). This was expected since the outer acceleration factor (3×3) was greater than the number of coils. The estimates using the proposed PI-CS-MSBS method correspond well with the reference estimates. The scan time of the prospectively undersampled acquisition was only 01:12 as compared to 10:11 of the fully-sampled acquisition.

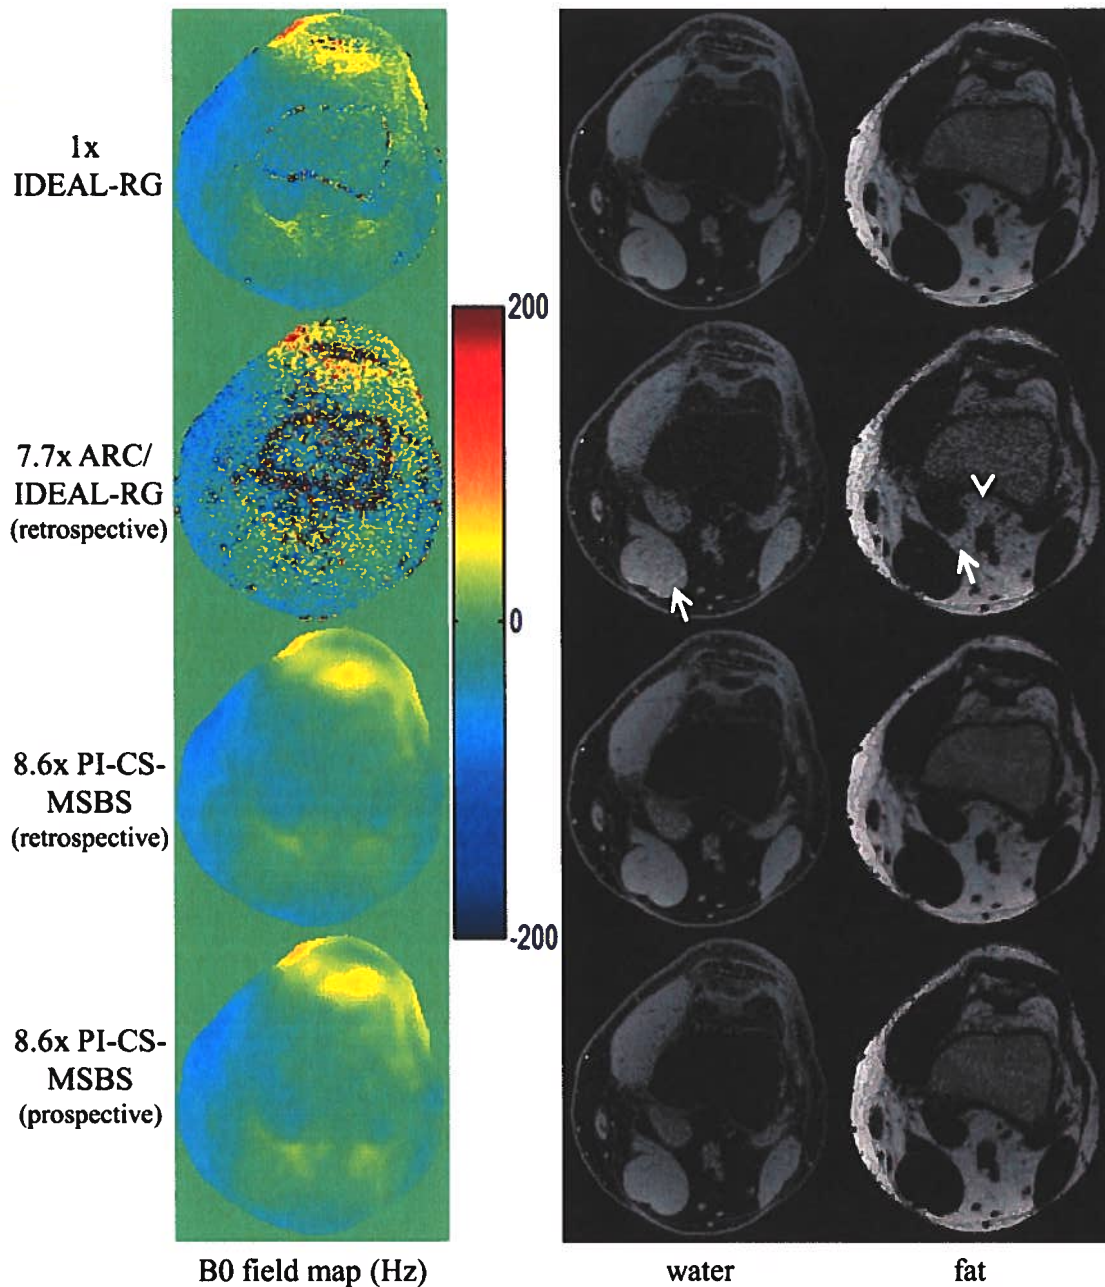


Figure 4.7: The B0 field map, water, and fat estimates of the knee acquired using an eight-channel knee coil. The arrowheads and arrows in the ARC/IDEAL-RG estimates highlight, respectively, regions of noise amplification and image artifacts. These artifacts were anticipated since the outer acceleration factor of 3x by 3x is greater than the number of receiver elements. In contrast, the estimates using the proposed method exhibit incoherent artifacts as a result of the Poisson disk sampling and l_1 -regularization. System issues, such as eddy currents, that are associated with the k-space sampling order do not affect the quality of the water-fat separation as evidenced by the retrospective versus prospective results.

4.4: Discussion

The proposed PI-CS-MSBS approach yielded B0 field map, water, and fat estimates of better quality than ARC/IDEAL-RG. The ARC/IDEAL-RG estimates exhibited severe artifacts, which were anticipated since the outer acceleration factor was greater than the number of coils along the dimension(s) of undersampling in all cases. In the proposed approach, the Poisson disk sampling caused incoherent aliasing artifacts that were reduced by the ℓ_1 -regularization in the reconstruction. Any remaining artifacts appeared much more benign than the structured artifacts in the ARC/IDEAL-RG estimates. The smaller number of slice encodes in the brachial plexus versus the thigh and knee datasets limited the degree of wavelet compressibility along that dimension, and may have contributed to the slight loss of subtle features in the brachial plexus estimates as compared to the knee estimates.

The joint parallel imaging and compressed sensing approach imposed two complementary constraints on the reconstruction; one was based on the distinct spatial sensitivities of the receiver elements, the other on the presumed compressibility of the underlying images. By using coil sensitivities that were derived from the SPIRiT k-space kernel, the SPIRiT constraint was imposed without requiring an explicit calibration consistency expression in the reconstruction. The use of coil sensitivities allowed for the reconstruction of coil-combined images, which reduced the computational complexity as compared to reconstructing one set of images (i.e. water, fat, B0 field map) for each coil. In

addition, reconstructing the coil-combined images freed the sparsifying transform from the responsibility of capturing the magnitude and phase variations of the coil sensitivities.

The current framework has some limitations. First, the topic of quantitation has not yet been addressed. Accurate water-fat quantitation requires compensation for confounding factors, the most prominent of which is $R2^*$ ($=1/T2^*$). Fortunately, the signal model that was proposed is easily amended to account for the $R2^*$ parameter as will be seen in the next chapter. Second, the phase accrual during the readout due to chemical shift off-resonance was ignored because its effects were minimal for the Cartesian trajectory and high sampling bandwidth that we used. If a non-Cartesian trajectory was used, one would have to account for the distinct time that each k-space point was sampled. Next, the regularization parameter λ was chosen empirically. This is acceptable for showing feasibility of the proposed method but an automated procedure would be required for wide acceptance. In addition, the reconstruction time would need to be shortened to permit online reconstruction. Lastly, the pulse sequence that we used for prospective undersampling was limited to sampling the same phase encode line for all echo times. This limitation restricted the sampling incoherence to the $ky-kz$ plane rather than the $ky-kz-TE$ volume. Blipping the phase encode between consecutive echoes would increase sampling incoherence, which may slightly improve results.

Chapter 5: Accelerated T2*-Compensated Quantitative Liver Fat Imaging Using Multiscale B-splines, Parallel Imaging, and Compressed Sensing

5.1: Introduction

Chemical shift encoded water-fat separation techniques provide qualitative information that may be used to infer that the amount of liver fat in a patient is 'higher' than 'normal'. However, this assessment lacks the objectivity that is required to make it a reliable source for the diagnosis, prognosis, and treatment of disease. The recent push for quantitative chemical shift encoded techniques represents an effort to replace, or at least supplement, the qualitative descriptions with objective measures. A primary motivation for doing so is to replace liver biopsy, the current clinical gold standard, with a noninvasive measure of liver fat [56]. Known as the fat fraction, this quantity reflects the relative amount of 'MR-visible' fat protons to the total number of 'MR-visible' water and fat protons. Fat fraction has the potential to be used as an imaging biomarker of nonalcoholic fatty liver disease, which can lead to fibrosis and cirrhosis [47].

An accurate and precise measure of fat fraction requires the compensation of confounding factors such as T2* decay, T1 bias, and noise bias. Unlike T1 bias and noise bias, T2* decay is typically compensated for by including it in the signal fitting. In theory, three echoes suffice for recovering the unknowns (e.g. water, fat, B0 field map, and T2*), but in practice it has been

found that at least six echoes are typically needed for accurate water-fat quantitation.

The need for at least six echoes lengthens the scan time as compared to a single echo acquisition. The increased scan time becomes especially problematic in liver imaging because the patient must maintain a breath-hold during the scan to avoid respiratory-induced image artifacts that would degrade the accuracy and confidence of the fat fraction measurement. The problem is often addressed by fixing the maximum scan time, which directly impacts other sequence/image parameters. Typically, a compromise must be made between the spatial resolution and the volume coverage that can be achieved within the allotted scan time.

In this chapter, a technique for accelerated water-fat quantitation is presented. The method can be used to significantly reduce or eliminate the degree of compromise between the spatial resolution and the volume coverage. Building upon the framework for accelerated water-fat separation, a slight modification to the signal model is presented to account for the $T2^*$ parameter. Likewise, the reconstruction routine is modified to include estimation of $T2^*$. Results are shown using datasets from seven volunteers who encompass a range of fat fractions observed clinically.

5.2: Methods

5.2.1: Signal Model

The effects of $T2^*$ can be modeled as an exponential decay of the transverse magnetization. For mathematical convenience, the rate of exponential decay is typically captured through the parameter $R2^*$, which is equal to $1/T2^*$. The incorporation of $R2^*$ into the signal model is shown in Equation 5.1.

$$\mathbf{k}_u = \mathbf{F}_u \mathbf{C} \Psi^R \mathbf{A} \boldsymbol{\rho} + \mathbf{N}(\mathbf{0}, \boldsymbol{\Sigma}) \quad (5.1)$$

Notice that the only difference between this expression and that in Equation 4.1 is that Ψ has been replaced with Ψ^R . The parameter Ψ^R contains $\exp(j2\pi(\psi^p + jR2^{*p}/2\pi)t_n)$ on the p^{th} diagonal of the n^{th} block, where ψ^p is the B0 field map value and $R2^{*p}$ is the $R2^*$ value (both in Hertz) at the p^{th} pixel and t_n is the time (in seconds) of the n^{th} echo. Put another way, $R2^*$ has been added as the imaginary component of the unknown term in the exponential. Since $j^2 = -1$, this formulation correctly reflects the modeled effects of $T2^*$.

The signal model assumes a single exponential decay (i.e. that water and fat share the same $R2^*$ value), which is not true in general since the $T2$ s of water and fat are different. This simplification is known to introduce bias into the fat fraction estimate, however Hernando et al. have found that the mean square error of the fat fraction estimate at clinically-relevant SNRs is smaller when using the single decay model than when using the two decay model (i.e. separate $R2^*$

values for water and fat) [26]. This finding is explained by a smaller variance of the fat fraction estimate when using the single decay model as compared to the two decay model, which more than makes up for the introduced bias.

5.2.2: Signal Estimation

Compensating for $R2^*$

The addition of $R2^*$ to the signal model requires a reevaluation of the signal estimation presented in §4.2.3. Since $R2^*$ is not a linear parameter in the signal model, the relatively straightforward approach that was used to estimate the water and fat images is not directly applicable here. Thus it is important to understand how the addition of $R2^*$ affects the least-squares cost function.

Figure 5.1 shows the least squares cost for one pixel from a fully-sampled liver dataset as a function of $B0$ field map and $R2^*$. This example shows that the cost is a nonconvex function of both parameters, which indicates that care should be taken to avoid local minima estimates. It is interesting to note that the shape of this exemplary cost function is much more dependent on the value of the $B0$ field map than on the value of $R2^*$. Further, in the neighborhood of the true $B0$ field map value, the cost appears to be a convex function of both the $B0$ field map and $R2^*$ parameters. This suggests that signal estimation can be done in a pseudo-sequential manner in which the $B0$ field map is first approximated using the method proposed earlier, followed by a joint refinement and estimation of the $B0$ field map and $R2^*$ values, respectively.




Figure 5.1: The least squares cost is a nonconvex function of both B0 field map and R2* parameters. However, the cost appears to be a convex function of both B0 field map and R2* in the neighborhood of the true B0 field map value. This suggests that signal estimation can be done in a pseudo-sequential manner in which the B0 field map is first approximated, followed by a joint refinement and estimation of the B0 field map and R2* values, respectively.

As with the other unknown parameters, a regularization criterion should be imposed on the R2* map because it is to be estimated from undersampled k-space data. In general, this criterion is not sufficient to avoid convergence to local minima solutions, however the proposed pseudo-sequential reconstruction permits the assumption of a convex least squares cost as a function of R2*. In this light, techniques such as ℓ_1 -regularization on the R2* estimate seem applicable if a sparsifying transform can be found.

Figure 5.2 shows an example R2* image that has been represented using different percentages of Daubechies-4 (db4) wavelet coefficients. Like water and

fat images, the db4 transform provides a highly compressible representation of the $R2^*$ map. This finding is somewhat expected since $R2^*$ is a tissue-dependent parameter that should reflect aspects of the underlying anatomy. In particular, the edges that define the boundary between two anatomical regions are apparent in the $R2^*$ image seen in Figure 5.2.

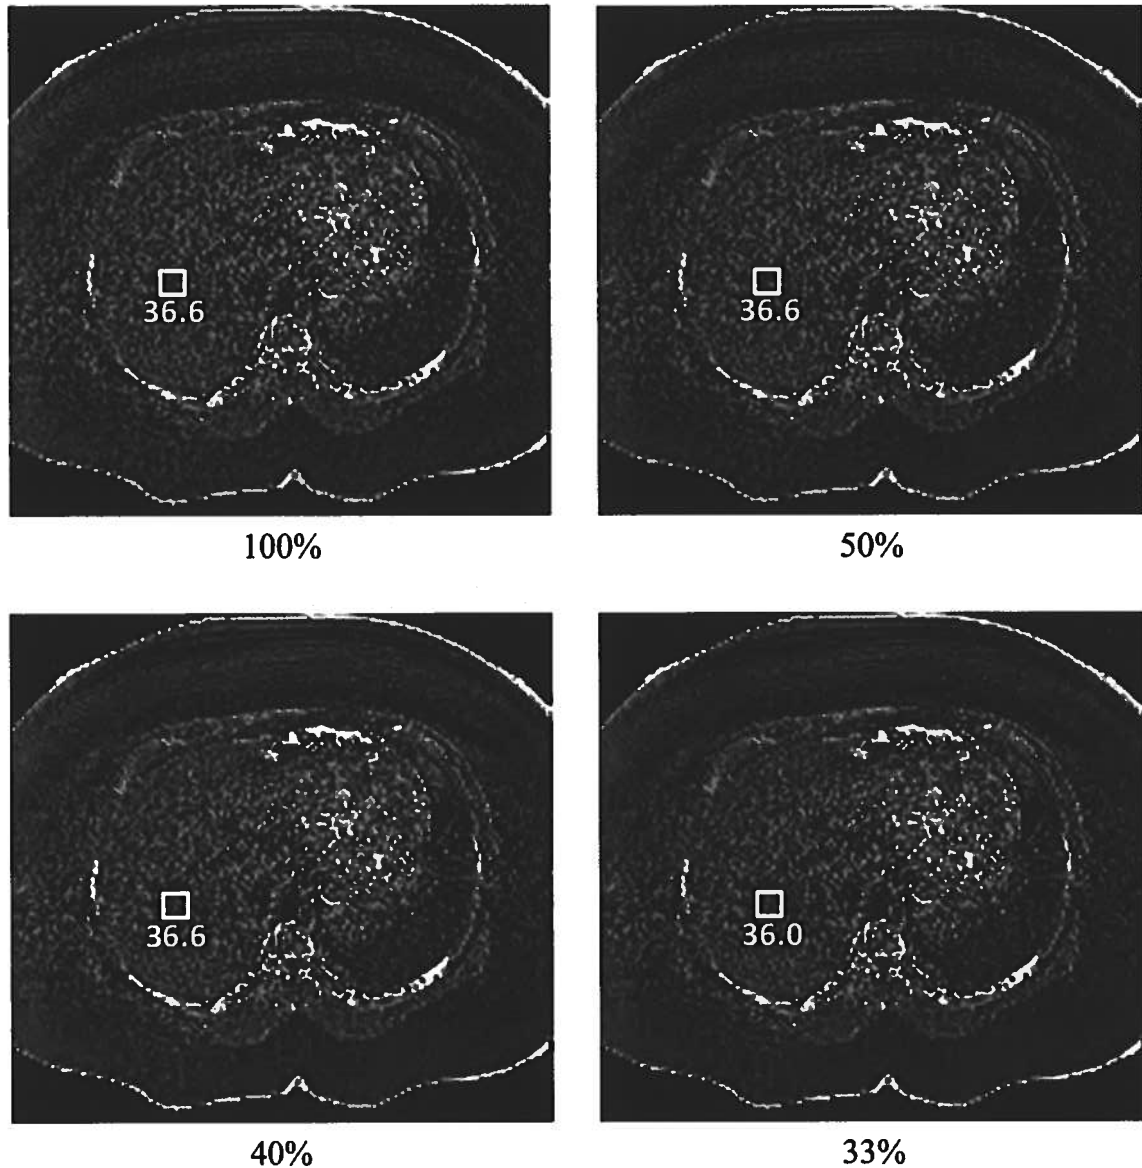


Figure 5.2: An example to show the compressibility of R2* images using the Daubechies-4 wavelet transform. Each image is shown using only a selected percentage of the highest magnitude transform coefficients. Calculation of the mean R2* within a region of interest shows that nearly no to little degradation of the mean R2* value occurs with higher compression.

A linear approximation of the exponential function is made using the first-order Taylor expansion. This is very similar to the approximation that was presented in Chapter 3 for the B0 field map estimation. With this approximation,

the B0 field map and R2* map are updated by the expression shown in Equation 5.2.

$$\begin{aligned} \min_{\Delta\psi, \Delta\rho, \Delta R2^*} \quad & \|r - x(\Delta\psi, \Delta\rho, \Delta R2^*)\|_2^2 + \lambda|\mathbf{W} \cdot \Delta\rho|_1 + \lambda|\mathbf{W} \cdot \Delta R2^*|_1 \\ \text{s. t. } \quad & \Delta\psi \in \text{span}\{\mathbf{B}_{\max}\} \end{aligned} \quad (5.2)$$

This expression is derived in a very similar manner as that in Equation 3.7 except for the addition of the R2* parameter. Further, the l1-regularization is imposed on both the R2* map update as well as the water and fat estimate errors ($\Delta\rho$). The estimate of $\Delta\rho$ is discarded once it has been estimated, but including it in the estimation has been shown to speed convergence. As in Equation 4.8, \mathbf{W} represents the Daubechies-4 wavelet transform. For simplicity, the same value of the regularization parameter (λ) is used for $\Delta\rho$ and $\Delta R2^*$, although further optimization may be possible by using independent values. The convex minimization problem is solved using conjugate gradients.

Alternating Minimization

Figure 5.3 presents a flowchart that summarizes the full signal estimation routine for water-fat separation and quantitation. An alternating minimization approach is taken to estimate the water/fat images and the B0 field/R2* maps. The R2* map is estimated only after an approximation of the B0 field map has been made, which has the benefit of avoiding local minima solutions.

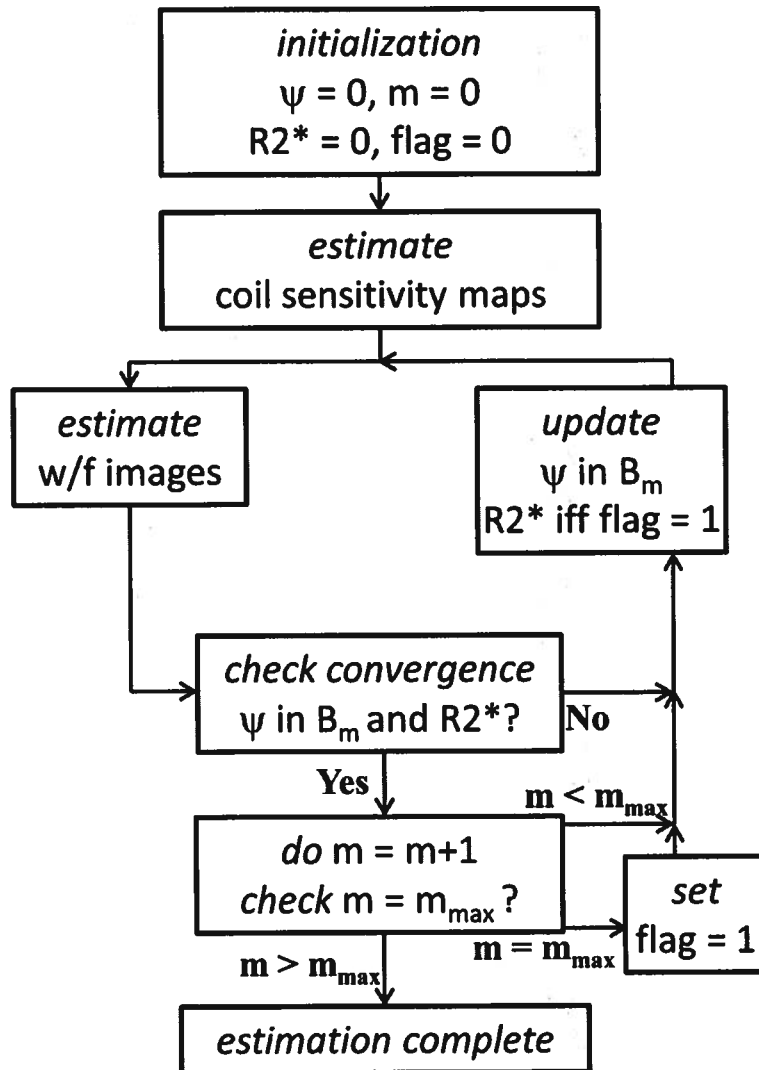


Figure 5.3: The flowchart summarizes the steps for signal estimation. The coil sensitivity maps are estimated first and are then held fixed. Subsequently, the water, fat, B0 field map and R2* map are estimated via an alternating minimization. The R2* map is estimated only after an approximation of the B0 field map has been made. This has the benefit of avoiding local minima estimates of R2*. The following abbreviations are used in the flowchart: B0 field map (ψ), water/fat (w/f), B-spline set at m^{th} level (B_m).

5.2.3: Experiments

All experiments were conducted on a GE Signa EXCITE HDx 3T system (GE Healthcare, Waukesha, WI) at USC HCC-II. Fully sampled measurements were collected at six echo times using an investigational GE IDEAL 3D spoiled gradient echo sequence. The echo train length was equal to two and the sampling bandwidth was ± 125 kHz. The multiecho acquisition with unipolar readouts resulted in a scan time that was three times longer than a single echo acquisition. The flip angle was set to 5° to minimize the effects of T1 bias [34].

Phantom

A collection of three water-fat phantoms was constructed to demonstrate the need for T2* compensation for accurate quantitation. The first phantom contained only distilled water, the second contained equal parts by mass of distilled water and 20% Intralipid (Baxter Healthcare, Deerfield, IL), and the third contained only 20% Intralipid. In addition, 1 mg of manganese(II) chloride was added to each phantom to replicate the T2* shortening that would be caused by the presence of iron in the liver. The phantoms were placed in a circular glass container along with other phantoms that were not a part of this experiment. The container was filled with tap water to decrease the likelihood of large susceptibility-induced B0 field map fluctuations. Table 5.1 shows the parameters of the data acquisition.

Liver

Fully sampled datasets were acquired with informed consent from seven subjects. Based on previous MRI scans, an effort was made to choose a subject cohort that would give a variety of liver fat fractions. The presence of iron in the liver was not an inclusion/exclusion criterion in this study. Each subject was required to maintain a breath-hold of 18-20s during the scan. Table 5.1 shows the data acquisition parameters, which varied slightly by subject.

Object	Coil	Matrix Size	FOV (cm)	Δz (mm)	TE₁ (ms)	ΔTE (ms)	Scan Time (mm:ss)
Phantom	eight-channel head	256x256x8	20	4.1	1.568	0.9	01:01
Liver	eight-channel torso	160x160x4	32 - 44	5	1.08	0.668 - 0.754	00:18 - 00:20

Table 5.1: This table presents the data acquisition parameters that were used to collect the fully sampled phantom and liver datasets. For the liver datasets, the field of view (FOV) was set based on the size of the subject. This affected the echo spacing (ΔTE) and the scan time. The slice thickness (Δz) was kept the same for all subjects.

Image Reconstruction

The phantom dataset was reconstructed using IDEAL-RG and IDEAL-RG with T2* compensation (T2*-IDEAL) [70] to demonstrate the effects of T2* compensation on fat quantitation. In addition, the phantom dataset was retrospectively undersampled by a factor of 3.4x along the phase encode axis and was then reconstructed using the proposed approach (PI-CS-T2*-MSBS).

Table 5.2 lists the applicable reconstruction parameters used by PI-CS-T2*-MSBS.

The liver datasets were reconstructed using three different methods. Each slice was processed independently from the other three slices in the acquired volume.

First, the proposed approach was tested by retrospectively undersampling each dataset by a factor of 3.1x along the phase encode axis using a Poisson disk sampling pattern. The regularization parameter λ was set to six for all reconstructions. Table 5.2 lists the applicable reconstruction parameters used in this reconstruction.

Second, a commonly-cited parallel imaging and water-fat quantitation method [27, 72] was implemented. This method has been shown to yield accurate and precise fat fraction measurements at a modest (~2x) acceleration factor [27]. The fully sampled data were undersampled by a factor of 3.1x and were then reconstructed using ARC [6]. Subsequently, the data were coil combined [66] and then water-fat quantitation was done using T2*-IDEAL. This reconstruction will hereafter be referred to as ARC/T2*-IDEAL. Table 5.2 lists the applicable reconstruction parameters used in this reconstruction.

Third, the fully sampled data were coil combined and then passed to T2*-IDEAL estimation routine. The estimates from this reconstruction served as the references.

In all cases, the fat fraction image estimate was calculated from the water and fat image estimates using the magnitude discrimination technique proposed by Liu et al. [34]. This technique has been shown to minimize the effects of noise bias especially when the true fat fraction is low.

Object	N_{acs}	Kernel Size	Outer Reduction	Net Acceleration
Phantom	16	7x7	4	3.1
Liver				3.4

Table 5.2: Reconstruction parameters for the ARC/T2*-IDEAL and PI-CS-T2*-MSBS. N_{acs} denotes number of autocalibration lines. Note that the outer reduction factor is only applicable to the ARC/T2*-IDEAL reconstruction.

Data Analysis

The middle slice of the phantom dataset was chosen to determine the effects of T2* compensation on the accuracy of the fat fraction estimate. In each of the three water-fat phantoms, the average fat fraction estimate using each of the reconstruction methods was calculated and compared.

For each of the 28 2D slices of the liver, a region of interest (ROI) was manually selected in the liver using the water estimate from the T2*-IDEAL reconstruction. Care was taken to ensure that the ROI avoided any vessels. Each ROI was 3x3 pixels, which covered an area of approximately 7.5x7.5mm². It was assumed that the true fat fraction was uniform over this region. Figure 5.4 shows a representative example of the ROI placement on the water image.



Figure 5.4: Liver region of interest (ROI) (white square) in which the fat fraction estimate was measured. The ROI was 3x3 pixels, which covers a region of approximately $7.5 \times 7.5 \text{mm}^2$. The ROI is shown on the water image for better visibility. The reference fat fraction value was calculated by taking the average of the T2*-IDEAL fat fraction estimates in the ROI.

The ROI was copied to the fat fraction image estimates from each of the three reconstruction techniques. The reference fat fraction value was calculated by taking the mean of the T2*-IDEAL fat fraction estimates in the ROI. The average fat fraction estimates in the ROI from the ARC/T2*-IDEAL and the proposed PI-CS-T2*-MSBS methods were also calculated and were compared to the reference value for each slice. The accuracy of each method was determined through linear regression. In addition, the mean square error (MSE) of the fat fraction estimates in the ROI was computed using the reference fat fraction value as ground truth. The error was decomposed into the variance and squared bias components.

5.3: Results

Phantoms

Figure 5.5 shows the estimated fat fractions and $R2^*$ estimates using $T2^*$ -IDEAL and the proposed PI-CS- $T2^*$ -MSBS method. In addition, the fat fraction estimates using IDEAL-RG (i.e. without $T2^*$ compensation) are also shown. The average fat fraction and $R2^*$ estimate for each of the three phantoms is also shown. A difference of about 5% in fat fraction is observed when $R2^*$ is not accounted for. The average fat fraction estimate with $R2^*$ compensation for the water phantom is close to the expected 0% and the estimate for the 20% Intralipid phantom is in close agreement with that reported by Bydder et al. [9]. Although the true fat fraction for the phantom that contains equal parts of distilled water and Intralipid is not known, it is expected to be close to 10%.

The fat fraction and $R2^*$ estimates using the proposed method closely agree with the corresponding reference estimates. In addition, both estimates appear qualitatively similar to the reference images. Any unresolved aliasing artifacts are largely not visible.




Figure 5.5: The different concentrations of distilled water and Intralipid (IL) (left image). Each of the three labeled phantoms also included 1 mg of MnCl₂ to simulate the T2* shortening that would be caused by iron overload. The fat fraction estimates differ by approximately 5% without/with R2* compensation (top center image). The fat fraction estimates using the proposed method (bottom center image) closely agree with the reference estimates. The estimated R2* values are quite high due to the presence of MnCl₂ (right images).

Liver Datasets

Figure 5.6 shows the fat fraction and R^2 estimates of one 2D slice using the three reconstruction methods. The reported fat fraction is the average of the fat fraction estimates within the ROI. Both qualitatively and quantitatively, the fat fraction estimate using the proposed method is quite similar to the reference estimate. In contrast, the ARC/T2*-IDEAL estimates suffer from unresolved aliasing, which is highlighted by the white arrowhead.

Figure 5.7 is a scatterplot of the average fat fraction estimate in the ROI versus the reference fat fraction value for each of the two accelerated techniques. Linear regression analysis yielded slope, intercept, and R^2 values of 0.78 ± 0.22 , -0.0004 ± 0.024 , and 0.34 , respectively for ARC/T2*-IDEAL and 1.08 ± 0.04 , -0.007 ± 0.004 , and 0.97 , respectively, for the proposed PI-CS-T2*-MSBS approach. The identity line is shown with dashes. At both low and high reference fat fractions, the estimated fat fractions from the ARC/T2*-IDEAL exhibit high variability from the identity line.

Figure 5.8 shows the MSE of the fat fraction estimates in the ROI using the reference fat fraction value as ground truth for each of the 28 slices. The error is decomposed into the sum of the variance and squared bias components. Note that the bias for the T2*-IDEAL estimator is zero in all cases because the reference fat fraction was calculated from the T2*-IDEAL fat fraction estimates in the ROI. Each triplet shows the MSE for T2*-IDEAL (left), 3.1x ARC/T2*-IDEAL (center), and 3.1x PI-CS-T2*-MSBS (right).




Figure 5.6: An example of the fat fraction and R2* estimates using T2*-IDEAL, 3.1x ARC/T2*-IDEAL, and 3.1x PI-CS-T2*-MSBS. The reported fat fraction was calculated as the average of the fat fraction values in the ROI. Notice the unresolved aliasing artifacts in the ARC/T2*-IDEAL fat estimate (white arrowhead).




Figure 5.7: The average fat fraction estimate in the ROI is plotted against the reference fat fraction value. The identity line is shown with dashes. Linear regression yielded slope, intercept, and R^2 values of 0.78 ± 0.22 , -0.0004 ± 0.024 , and 0.34 , respectively, for ARC/T2*-IDEAL and 1.08 ± 0.04 , -0.007 ± 0.004 , and 0.97 , respectively, for the proposed PI-CS-T2*-MSBS approach.




Figure 5.8: The mean square error (MSE), with respect to the reference fat fraction, of the fat fraction estimates in the ROI. Each triplet represents the MSE for T2*-IDEAL (left), 3.1x ARC/T2*-IDEAL (center), and 3.1x PI-CS-T2*-MSBS (right). Each MSE value is decomposed into its variance (blue) and squared bias (red) components. In 27 of 28 cases, the MSE of the ARC/T2*-IDEAL estimate was the highest. In general, most of the error was attributed to the variance component.

5.4: Discussion

The experiment using the water-fat phantoms demonstrated the need for T2* compensation to accurately estimate the fat fraction. An absolute difference of approximately 5% was seen in the fat fraction estimate with and without T2* compensation. The proposed method accurately estimated the fat fraction and R2* at 3.4x acceleration. In addition, no significant qualitative differences were seen in the estimates using the proposed method.

Linear regression analysis using the estimates of liver fat fraction showed that the proposed method was a more accurate estimator of fat fraction than was ARC/T2*-IDEAL. As seen in Figure 5.7, the estimated fat fractions using the proposed method are, in general, concentrated around the identity line. In contrast, those estimates using ARC/T2*-IDEAL show much more variability, both positive and negative, with regard to the reference fat fraction value. The estimated slope and R^2 values for the proposed method were much closer to the ideal value of one. The estimated intercept value for the ARC/T2*-IDEAL value was close to the ideal value of zero although the standard error was much higher.

In 27 of the 28 2D liver slices, the proposed approach yielded an estimate of fat fraction that had lower MSE than the corresponding estimate from ARC/T2*-IDEAL. By decomposing the MSE into its variance and squared bias components, it was found that the MSE was mostly attributable to the variance component for all three reconstruction methods. The especially high variance of the ARC/T2*-IDEAL estimator was observed as severe noise-like artifacts in the fat fraction estimate, although structured artifacts were occasionally visible (see Figure 5.6).

A potential concern with using multiscale B-splines to represent the B0 field map was that a slight error in the field map estimate may cause substantial quantitation errors. This has not been observed in any studies to date. That is,

any errors introduced by using multiscale B-splines for B0 field map estimation did not affect the accuracy of water-fat quantitation.

A shortcoming in this work was the absence of MR spectroscopy to serve as the reference for fat fraction. However, the studies from Hines et al. [27] and Meisamy et al. [49] have shown that the method used for reference (e.g. T2*-IDEAL) is an accurate and precise measure of fat fraction. Thus, results that are similar to the current finding would be expected if MR spectroscopy was used as the reference in the study.

Chapter 6: Future Work

In this chapter, opportunities are presented for advancing the work that was discussed in this dissertation. A technical opportunity specific to water-fat imaging is first proposed, followed by an extension to general off-resonance correction, and concluding with aspects for clinical translation.

Calibrationless Parallel Imaging and Water-Fat Separation

The method proposed in Chapter 4 for accelerated water-fat separation relied on first estimating the coil sensitivity maps. Once estimated, the coil maps were held fixed while the water and fat images and the B0 field map were estimated using an alternating minimization. The sequential nature of this approach offered the potential for errors in the coil sensitivity estimates to propagate into the subsequent water-fat separation step. An approach that implicitly exploits the redundant characteristics of the receiver system may allow for joint estimation of the water and fat images from each coil as well as the B0 field map, which would avoid the potential for error propagation.

Separate works from Lustig [40] and Trzasko et al. [64] have proposed calibrationless parallel imaging methods, which implicitly impose the expected smoothness of coil sensitivity maps. Neither of these methods requires a separate calibration stage nor a fully-sampled k-space region.

It appears these either of these methods could be directly incorporated with the multiscale B-spline approach (Chapter 3) to jointly estimate the B0 field

map, water, and fat images. To the best of the author's knowledge, this would be the first method for accelerated water-fat separation that does not require a fully-sampled center of k-space. It is expected that an extension to quantitation could be made with proper compensation for T2* effects.

General off-resonance correction

Off-resonance correction is required in many MRI applications, especially those that acquire k-space data using non-Cartesian trajectories. Unlike in 2DFT acquisitions, off-resonance in a spiral or radial acquisition causes blurring that severely degrades the image quality. Though off-resonance techniques for arbitrary trajectories have been proposed [51, 62], the need for a robust technique still remains.

It may be possible to extend the multiscale B-spline approach to the problem of general off-resonance correction. In doing so, a modification to the signal model would be required to account for the distinct time that each k-space point was sampled. The reconstruction would require a slight modification as well since the data points may not lie on a Cartesian grid.

Clinical Translation

An ultimate goal of the work presented in this dissertation is to translate the developed technique into the preclinical and clinical settings. Doing so would require many challenges, both technical and clinical, to be overcome.

The convex minimization problems presented in this dissertation were solved using conjugate gradients. While this algorithm achieves the best known undersampling-sparsity tradeoff [16], its computational complexity severely hinders the reconstruction time in large scale applications such as MRI. Recent work has led to reconstruction algorithms, such as message-passing algorithms [16], that dramatically decrease the reconstruction time while preserving the undersampling-sparsity tradeoff of convex minimization techniques.

The regularization parameter was empirically determined for all of the presented reconstructions. While this is acceptable for demonstrating feasibility of a new technique, wide clinical acceptance would require an automated procedure for choosing this parameter. Interestingly, very recent work from Khare et al. [31] has demonstrated a fast reconstruction algorithm that does not require the user to select the values of any free parameters. This completely automated and fast reconstruction algorithm was shown to perform at least as well as an empirically-tuned approach.

Further testing of the method would be required before it could be routinely applied on a patient population. A strong clinical collaboration is essential to ensure the success of this undertaking. This would require the recruitment of a large patient cohort for the study as well as a dedication from clinicians to qualitatively and quantitatively analyze each of the reconstructed images.

Bibliography

- [1] J. Berglund, H. Ahlstrom, L. Johansson, and J. Kullberg. Two-point Dixon method with flexible echo times. *Magnetic Resonance in Medicine*, 65(4):994–1004, 2011.
- [2] J. Berglund, L. Johansson, H. Ahlstrom, and J. Kullberg. Three-point Dixon method enables whole-body water and fat imaging of obese subjects. *Magnetic Resonance in Medicine*, 63(6):1659–1668, 2010.
- [3] T. A. Bley, O. Wieben, C. J. Francois, J. H. Brittain, and S. B. Reeder. Fat and water magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 31(1):4–18, 2010.
- [4] P. Boernert, P. Koken, and H. Eggers. Spiral water-fat imaging with integrated off-resonance correction on a clinical scanner. *Journal of Magnetic Resonance Imaging*, 32(5):1262–1267, 2010.
- [5] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [6] A. C. Brau, P. J. Beatty, S. Skare, and R. Bammer. Comparison of reconstruction accuracy and efficiency among autocalibrating data-driven parallel imaging methods. *Magnetic Resonance in Medicine*, 59(2):382–395, 2008.
- [7] E. K. Brodsky, J. H. Holmes, H. Yu, and S. B. Reeder. Generalized k-space decomposition with chemical shift correction for non-Cartesian water-fat imaging. *Magnetic Resonance in Medicine*, 59(5):1151–1164, 2008.
- [8] G. Bydder, R. Steiner, L. Blumgart, S. Khenia, and I. Young. MR imaging of the liver using short TI inversion recovery sequences. *Journal of Computer Assisted Tomography*, 9(6):1084–9, 1985.
- [9] M. Bydder, T. Yokoo, G. Hamilton, M. S. Middleton, A. D. Chavez, J. B. Schwimmer, J. E. Lavine, and C. B. Sirlin. Relaxation effects in the quantification of fat using gradient echo imaging. *Magnetic Resonance Imaging*, 26(3):347 – 359, 2008.
- [10] E. Candes and J. Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23:969, 2007.

- [11] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Transactions on Information Theory*, 52:5406 – 5425, 2006.
- [12] R. L. Cook. Stochastic sampling in computer graphics. *ACM Transactions on Graphics*, 5:51–72, January 1986.
- [13] W. T. Dixon. Simple proton spectroscopic imaging. *Radiology*, 153(1):189–194, 1984.
- [14] M. Doneva, P. Boernert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig. Compressed sensing for chemical shift-based water-fat separation. *Magnetic Resonance in Medicine*, 64(6):1749–1759, 2010.
- [15] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, 2006.
- [16] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [17] H. Eggers, B. Brendel, A. Duijndam, and G. Herigault. Dual-echo Dixon imaging with flexible choice of echo times. *Magnetic Resonance in Medicine*, 65(1):96–107, 2011.
- [18] S. Eustace, R. Tello, V. Decarvalho, J. Carey, E. Melhem, and E. K. Yucel. Whole body turbo STIR MRI in unknown primary tumor detection. *Journal of Magnetic Resonance Imaging*, 8(3):751–753, 1998.
- [19] D. George, S. Goldwurm, G. A. Macdonald, L. L. Cowley, N. I. Walker, P. J. Ward, E. C. Jazwinska, and L. W. Powell. Increased hepatic iron concentration in nonalcoholic steatohepatitis is associated with increased fibrosis. *Gastroenterology*, 114(2):311 – 318, 1998.
- [20] G. H. Glover and E. Schneider. Three-point Dixon technique for true water/fat decomposition with B₀ inhomogeneity correction. *Magnetic Resonance in Medicine*, 18(2):371–383, 1991.
- [21] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.

- [22] M. A. Griswold, S. Kannengiesser, R. M. Heidemann, J. Wang, and P. M. Jakob. Field-of-view limitations in parallel imaging. *Magnetic Resonance in Medicine*, 52(5):1118–1126, 2004.
- [23] A. Haase, J. Frahm, W. Hanicke, and D. Matthaei. 1H NMR chemical shift selective (CHESS) imaging. *Physics in Medicine and Biology*, 30:341, 1985.
- [24] D. Hernando, J. P. Haldar, B. P. Sutton, J. Ma, P. Kellman, and Z.-P. Liang. Joint estimation of water/fat images and field inhomogeneity map. *Magnetic Resonance in Medicine*, 59(3):571–580, 2008.
- [25] D. Hernando, P. Kellman, J. P. Haldar, and Z.-P. Liang. Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm. *Magnetic Resonance in Medicine*, 63(1):79–90, 2010.
- [26] D. Hernando, Z.-P. Liang, and P. Kellman. Chemical shift-based water/fat separation: a comparison of signal models. *Magnetic Resonance in Medicine*, 64(3):811–822, 2010.
- [27] C. D. Hines, A. Frydrychowicz, G. Hamilton, D. L. Tudorascu, K. K. Vigen, H. Yu, C. A. McKenzie, C. B. Sirlin, J. H. Brittain, and S. B. Reeder. T1 independent, T2* corrected chemical shift based fat-water separation with multi-peak fat spectral modeling is an accurate and precise measure of hepatic steatosis. *Journal of Magnetic Resonance Imaging*, 33(4):873–881, 2011.
- [28] H. H. Hu, H.-W. Kim, K. S. Nayak, and M. I. Goran. Comparison of fat-water MRI and single-voxel MRS in the assessment of hepatic and pancreatic fat fractions in humans. *Obesity*, 18(4):841–847, Apr. 2010.
- [29] P. Irarrazabal, C. H. Meyer, D. G. Nishimura, and A. Macovski. Inhomogeneity correction using an estimated linear field map. *Magnetic Resonance in Medicine*, 35(2):278–282, 1996.
- [30] M. Jacob and B. P. Sutton. Algebraic decomposition of fat and water in MRI. *IEEE Transactions on Medical Imaging*, 28:173 – 184, 2009.
- [31] K. Khare, C. J. Hardy, K. F. King, P. A. Turski, and L. Marinelli. Accelerated MR imaging using compressive sensing with no free parameters. *Magnetic Resonance in Medicine*, (in early view).

- [32] P. Lai, M. Lustig, S. Vasanawala, and A. C. Brau. ESPIRiT (efficient eigenvector-based L1SPIRiT) for compressed sensing parallel imaging - theoretical interpretation and improved robustness for overlapped FOV prescription. In *Proceedings of the 19th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, 2011.
- [33] D. Liang, B. Liu, J. Wang, and L. Ying. Accelerating SENSE using compressed sensing. *Magnetic Resonance in Medicine*, 62(6):1574–1584, 2009.
- [34] C.-Y. Liu, C. A. McKenzie, H. Yu, J. H. Brittain, and S. B. Reeder. Fat quantification with IDEAL gradient echo imaging: correction of bias from T1 and noise. *Magnetic Resonance in Medicine*, 58(2):354–364, 2007.
- [35] R. N. Low, M. J. Austin, and J. Ma. Fast spin-echo triple echo Dixon: Initial clinical experience with a novel pulse sequence for simultaneous fat-suppressed and nonfat-suppressed T2-weighted spine magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 33(2):390–400, 2011.
- [36] D. S. Lu, S. Saini, P. F. Hahn, M. Goldberg, M. J. Lee, R. Weissleder, B. Gerard, E. Halpern, and A. Cats. T2-weighted MR imaging of the upper part of the abdomen: should fat suppression be used routinely? *American Journal of Roentgenology*, 162(5):1095–1100, 1994.
- [37] W. Lu and B. A. Hargreaves. Multiresolution field map estimation using golden section search for water-fat separation. *Magnetic Resonance in Medicine*, 60(1):236–244, 2008.
- [38] W. Lu, H. Yu, A. Shimakawa, M. Alley, S. B. Reeder, and B. A. Hargreaves. Water-fat separation with bipolar multiecho sequences. *Magnetic Resonance in Medicine*, 60(1):198–209, 2008.
- [39] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- [40] M. Lustig, M. Elad, and J. Pauly. Calibrationless parallel imaging reconstruction by structured low-rank matrix completion. In *Proceedings of the International Society of Magnetic Resonance in Medicine*, 2010.
- [41] M. Lustig, P. Lai, M. Murphy, S. Vasanawala, M. Elad, J. Zhang, and P. J. An eigen-vector approach to autocalibrating parallel MRI, where SENSE meets GRAPPA. In *Proceedings of the 19th Annual Meeting of the International Society of Magnetic Resonance in Medicine.*, 2011.

- [42] M. Lustig and J. M. Pauly. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic Resonance in Medicine*, 64(2):457–471, 2010.
- [43] J. Ma. Breath-hold water and fat imaging using a dual-echo two-point Dixon technique with an efficient and robust phase-correction algorithm. *Magnetic Resonance in Medicine*, 52(2):415–419, 2004.
- [44] J. Ma. Dixon techniques for water and fat imaging. *Journal of Magnetic Resonance Imaging*, 28(3):543–558, 2008.
- [45] J. Ma, J. A. Bankson, and R. Stafford. Multipoint Dixon imaging using sensitivity encoding. In *Proceedings of the International Society of Magnetic Resonance in Medicine*, 2003.
- [46] T. H. Magee and G. W. Hinson. MRI of meniscal bucket-handle tears. *Skeletal Radiology*, 27:495–499, 1998.
- [47] C. A. Matteoni, Z. M. Younossi, T. Gramlich, N. Boparai, Y. C. Liu, and A. J. McCullough. Nonalcoholic fatty liver disease: A spectrum of clinical and pathological severity, June 1999.
- [48] C. J. McMahon, A. J. Madhuranthakam, J. S. Wu, C. M. Yablon, J. L. Wei, N. M. Rofsky, and M. G. Hochman. High-resolution proton density weighted three-dimensional fast spin echo (3D-FSE) of the knee with IDEAL at 1.5 Tesla: Comparison with 3D-FSE and 2D-FSE - initial experience. *Journal of Magnetic Resonance Imaging*, (in early view).
- [49] S. Meisamy, C. D. G. Hines, G. Hamilton, C. B. Sirlin, C. A. McKenzie, H. Yu, J. H. Brittain, and S. B. Reeder. Quantification of hepatic steatosis with T1-independent, T2*-corrected MR imaging with spectral modeling of fat: Blinded comparison with MR spectroscopy. *Radiology*, 258(3):767–775, 2011.
- [50] D. G. Nishimura. *Principles of Magnetic Resonance Imaging*. 1996.
- [51] D. Noll, C. Meyer, J. Pauly, D. Nishimura, and A. Macovski. A homogeneity correction method for magnetic resonance imaging with time-varying gradients. *IEEE Transactions on Medical Imaging*, 10:629–637, 1991.
- [52] A. R. Pineda, S. B. Reeder, Z. Wen, and N. J. Pelc. Cramer-Rao bounds for three-point decomposition of water and fat. *Magnetic Resonance in Medicine*, 54(3):625–635, 2005.

- [53] B. Porat. *A course in digital signal processing*. John Wiley & Sons, Inc., 1997.
- [54] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. Sense: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999.
- [55] S. B. Reeder, E. K. Bice, H. Yu, D. Hernando, and A. R. Pineda. On the performance of T2* correction methods for quantification of hepatic fat content. *Magnetic Resonance in Medicine*, 67(2):389–404, 2012.
- [56] S. B. Reeder, I. Cruite, G. Hamilton, and C. B. Sirlin. Quantitative assessment of liver fat with magnetic resonance imaging and spectroscopy. *Journal of Magnetic Resonance Imaging*, 34(4):729–749, 2011.
- [57] S. B. Reeder, B. A. Hargreaves, H. Yu, and J. H. Brittain. Homodyne reconstruction and IDEAL water-fat decomposition. *Magnetic Resonance in Medicine*, 54(3):586–593, 2005.
- [58] S. B. Reeder, A. R. Pineda, Z. Wen, A. Shimakawa, H. Yu, J. H. Brittain, G. E. Gold, C. H. Beaulieu, and N. J. Pelc. Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): Application with fast spin-echo imaging. *Magnetic Resonance in Medicine*, 54(3):636–644, 2005.
- [59] S. B. Reeder, Z. Wen, H. Yu, A. R. Pineda, G. E. Gold, M. Markl, and N. J. Pelc. Multicoil Dixon chemical species separation with an iterative least-squares estimation method. *Magnetic Resonance in Medicine*, 51(1):35–45, 2004.
- [60] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The NMR phased array. *Magnetic Resonance in Medicine*, 16(2):192–225, 1990.
- [61] S. Skare and J. Andersson. Correction of MR image distortions induced by metallic objects using a 3D cubic B-spline basis set: Application to stereotactic surgical planning. *Magnetic Resonance in Medicine*, 54(1):169–181, 2005.
- [62] B. Sutton, D. Noll, and J. Fessler. Fast, iterative reconstruction for MRI in the presence of field inhomogeneities. *IEEE Transactions on Medical Imaging*, 22:178 – 188, 2003.

- [63] L. S. Szczepaniak, P. Nurenberg, D. Leonard, J. D. Browning, J. S. Reingold, S. Grundy, H. H. Hobbs, and R. L. Dobbins. Magnetic resonance spectroscopy to measure hepatic triglyceride content: prevalence of hepatic steatosis in the general population. *American Journal of Physiology - Endocrinology and Metabolism*, 288(2):E462–E468, 2005.
- [64] J. D. Trzasko and A. Manduca. Calibrationless parallel MRI using ORACLE. In *Asilomar Conference on Signals, Systems, and Computers*, 2011.
- [65] J. Tsao and Y. Jiang. Hierarchical IDEAL - robust water-fat separation at high field by multiresolution field map estimation. In *Proceedings of the International Society of Magnetic Resonance in Medicine*, 2008.
- [66] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin. Adaptive reconstruction of phased array MR imagery. *Magnetic Resonance in Medicine*, 43(5):682–690, 2000.
- [67] K. Wang, H. Yu, J. H. Brittain, S. B. Reeder, and J. Du. k-space water-fat decomposition with T2* estimation and multifrequency fat spectrum modeling for ultrashort echo time imaging. *Journal of Magnetic Resonance Imaging*, 31(4):1027–1034, 2010.
- [68] S. N. Wiener, D. R. Neumann, and M. S. Rzeszutarski. Comparison of magnetic resonance imaging and radionuclide bone imaging of vertebral fractures. *Clinical Nuclear Medicine*, 14(9):666-670, 1989.
- [69] Q.-S. Xiang and L. An. Water-fat imaging with direct phase encoding. *Journal of Magnetic Resonance Imaging*, 7(6):1002–1015, 1997.
- [70] H. Yu, C. A. McKenzie, A. Shimakawa, A. T. Vu, A. C. Brau, P. J. Beatty, A. R. Pineda, J. H. Brittain, and S. B. Reeder. Multiecho reconstruction for simultaneous water-fat decomposition and T2* estimation. *Journal of Magnetic Resonance Imaging*, 26(4):1153–1161, 2007.
- [71] H. Yu, S. B. Reeder, A. Shimakawa, J. H. Brittain, and N. J. Pelc. Field map estimation with a region growing scheme for iterative 3-point water-fat decomposition. *Magnetic Resonance in Medicine*, 54(4):1032–1039, 2005.

- [72] H. Yu, A. Shimakawa, C. A. McKenzie, E. Brodsky, J. H. Brittain, and S. B. Reeder. Multiecho water-fat separation and simultaneous $R2^*$ estimation with multifrequency fat spectrum modeling. *Magnetic Resonance in Medicine*, 60(5):1122–1134, 2008.
- [73] M. Zanetti, C. W. A. Pfirrmann, M. R. Schmid, J. Romero, B. Seifert, and J. Hodler. Patients with suspected meniscal tears: Prevalence of abnormalities seen on MRI of 100 symptomatic and 100 contralateral asymptomatic knees. *American Journal of Roentgenology*, 181(3):635–641, 2003.