Accelerated Computation of Regularized Estimates in Magnetic Resonance Imaging

by

Michael J. Allison

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in the University of Michigan
2014

Doctoral Committee:

Professor Jeffrey A. Fessler, Chair
Professor Anna C. Gilbert
Professor Alfred O. Hero III
Assistant Research Scientist Jon-Fredrik Nielsen
Professor Douglas C. Noll

ACKNOWLEDGEMENTS

This work would not have been possible without the guidance of many people. First and foremost, I would like to thank my supervisor, Prof. Jeffrey Fessler, for his insightful advice over the course of my graduate career. I am grateful to Dr. Sathish Ramani for mentoring me during my early years and helping me determine my research focus. Furthermore, I would like to thank my doctoral committee for their suggestions, which greatly improved this thesis.

I would also like to thank the people and agencies who provided material support. Specifically, the Natural Sciences and Engineering Research Council of Canada (NSERC), the National Institutes of Health (NIH), and the Department of Electrical Engineering and Computer Science for providing crucial financial support. Furthermore, Prof. Thomas Chenevert and Dr. Alexey Samsonov for providing data sets that allowed me to pursue a wide range of MRI topics.

Finally, I would like to acknowledge all those who provided emotional support and encouragement. First, the EE:Systems staff who helped me transition through the many stages of a graduate career. Becky Turanski for always providing a kind ear and an outside perspective. Everyone in the image processing and fMRI labs for their camaraderie and interesting discussions. Takanori Watanabe and the rest of the pot-luck crew for being great friends throughout my time in Michigan — I wish you all the best. It would be remiss of me to not thank my parents and sister for

keeping me focused on the road ahead. Finally, and most importantly, I am eternally grateful to my wife, Annie, for her unwavering support and for helping me determine my path when it was unclear. I promise to return the favor if you choose to do a Ph.D.

TABLE OF CONTENTS

CHAPTER

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

Magnetic resonance imaging (MRI) is a non-invasive medical imaging modality that uses electromagnetic fields to image both physical and metabolic properties of the body. A fundamental aspect of MRI is the generation and measurement of spatially varying parameters such as magnetic field strength and receive coil sensitivity. We often assume that these parameters are uniform; however, there are many situations in which they are not. In some cases, we exploit the spatial variation to acquire additional information (e.g., parallel imaging techniques [1]). In other cases, the spatial variation can lead to significant reconstruction artifacts (e.g., main magnetic field inhomogeneity [2]). However, in all of these situations, accurate estimates of the spatial variation can be used to improve image quality [1, 2].

Regularized estimation methods provide robust and accurate estimates of spatially varying parameters [2–4]. They are most often statistically driven and contain additional terms that introduce a priori information about the estimate. However, these methods require the minimization of cost functions that contain large matrices making direct solutions intractable. Instead, iterative minimization methods are used; however, many of the common iterative methods converge slowly on these problems and this detracts from the appeal of regularized estimation. In this thesis, we investigate the regularized estimation of several spatially varying parameters with a focus on developing faster minimization algorithms.

MRI acquisition times can be reduced using parallel imaging techniques in which data is simultaneously acquired from multiple receive coils [1]. Some of these techniques require estimates of the spatially varying sensitivity of each coil to reconstruct the final image [5–7]. These estimates must be computed from calibration data obtained at the time of acquisition as they are influenced by the patient and the surrounding environment [8]. Regularized coil sensitivity estimation methods impose a smoothness constraint to generate high quality estimates even in cases of low signal-to-noise ratio (SNR) [3,9] and patient motion [9]. However, the standard conjugate gradient (CG) minimization technique used for this regularized estimation problem can take several minutes per coil [10]. As there can be dozens of coils in a single acquisition [11], the resulting computational costs associated with regularized estimation can be significant. In Chapter III, we further investigate regularized coil sensitivity estimation and propose a faster minimization strategy.

Ideally, the main magnetic field used in MRI would be spatially uniform over the entire field-of-view. However, this is not possible even with careful shimming [12, p. 837] and the remaining inhomogeneity can cause artifacts in image reconstructions [2]. These artifacts can be avoided with accurate estimates of the spatial inhomogeneity, which can be computed from multiple scans acquired at different echo times [2,13]. Regularized field inhomogeneity estimators impose a smoothness constraint to generate high quality estimates but they require minimizing a nonconvex cost function [2,4,14]. An existing minimization strategy using separable quadratic surrogate (SQS) functions finds a desirable local minimum for this problem but at a significant computational cost [2]. This cost is further compounded in some algorithms where regularized field estimation is used as a sub-step (e.g., [15] and our work in Chapter V). In Chapter IV, we propose two faster minimization strategies for regularized main magnetic field inhomogeneity estimation.

Magnetic field inhomogeneity is also a nuisance parameter in water-fat imaging where the goal is to obtain separate water and fat images. One approach to water-fat imaging is to jointly estimate the water image, the fat image, and the field inhomogeneity from several scans acquired at different echo times [16]. However, these additional scans increase the acquisition times of such methods beyond those of the basic water or fat imaging methods [13]. To counter this, compressed sensing based water-fat (CS-WF) imaging techniques, which can reconstruct high quality images from significantly less data than traditional methods, have been proposed [17–20]. The disadvantage of these techniques is that their cost functions contain nonlinear terms and non-differentiable functions. The existing CS-WF minimization strategies use linearization and corner rounding approximations to create CG based methods that take tens of minutes to converge [17–20]. In Chapter V, we propose a new initialization strategy and minimization algorithm that reduces the computation time of CS-WF image reconstruction.

The remainder of this thesis is organized as follows. Chapter II presents an overview of the MRI fundamentals and the optimization techniques used in our work. Chapter III investigates the coil sensitivity estimation problem and is based on [9] and its supplemental material. Chapter IV explores magnetic field inhomogeneity estimation and is an extension of [21]. Chapter V presents our minimization technique for the compressed sensing based water-fat image reconstruction problem. Chapter VI summarizes our contributions and proposes topics for future investigation.

1.1 Specific Contributions

The specific contributions in this thesis pertain to the previously outlined parameter estimation problems.

For receive coil sensitivity estimation, we extend the work on a regularized estimator proposed in [3]. We first introduce an alternating direction method of multipliers (ADMM) based algorithm [22] that minimizes the quadratic cost function in half the time required by a CG method with a circulant preconditioner. In doing so, we propose a variable splitting strategy that reformulates a shift-variant finite differencing matrix in a manner that allows for exact ADMM update steps. Furthermore, we demonstrate the benefits of using a modified ADMM approach, based on updating the Lagrange multipliers between variable updates [23], for this problem. We also address several unanswered topics related to regularized coil sensitivity estimation by demonstrating its improved performance compared to existing low resolution approaches on sensitivity encoded (SENSE) reconstructions [5] as well as the artifacts generated by using shift-invariant finite differencing matrices in the regularizer and the SENSE reconstruction artifacts that can result from using tight estimation masks.

For main magnetic field inhomogeneity estimation, we present a general cost function that combines the cases of regularized multiple echo time field map estimation [2] and regularized water-fat imaging field map estimation [4,14]. We introduce two new minimization algorithms that use quadratic surrogate functions that reduce the estimation time to a thirtieth of the existing SQS method [2]. The first of these methods adapts Huber's algorithm for quadratic surrogates [24] by exploiting the sparsity of the Hessian matrix of the surrogate function using sparse Cholesky factorization [25]. The second method accelerates a nonlinear CG algorithm using a monotonic step size line search algorithm and a preconditioning matrix both based on quadratic surrogate functions. We also explore using edge preserving regularization for field inhomogeneity estimation at air-tissue interfaces.

For CS-WF image reconstruction, we present a novel alternating minimization strategy that combines aspects of our two previous topics to obtain estimates in under a twelfth the time of the existing corner rounding CG algorithm [17]. To update the water and fat images, which involves minimizing a cost function with an ℓ_1-norm, we adapt an augmented Lagrangian based method proposed in [26]. To update the field inhomogeneity estimate, which involves a smoothness promoting ℓ_2-norm, we use a CG algorithm with a sparse Hessian matrix preconditioner like in Chapter IV. We also present an initialization strategy, based on our previous field inhomogeneity estimator, that can efficiently compute a field map estimate with a similar level of regularization as the CS-WF estimator. This high quality initialization shifts the bottleneck in the CS-WF reconstruction from the nonlinear system model to the CS based water and fat image updates.

Bibliography

[1] D. J. Larkman and R. G. Nunes, “Parallel magnetic resonance imaging,” *Phys. Med. Biol.*, vol. 52, no. 7, pp. R15–R55, Apr. 2007.

[2] A. K. Funai, J. A. Fessler, D. T. B. Yeo, V. T. Olafsson, and D. C. Noll, “Regularized field map estimation in MRI,” *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1484–94, Oct. 2008.

[3] S. L. Keeling and R. Bammer, “A variational approach to magnetic resonance coil sensitivity estimation,” *Applied Mathematics and Computation*, vol. 158, no. 2, pp. 53–82, Nov. 2004.

[4] 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,” *Mag. Res. Med.*, vol. 59, no. 3, pp. 571–80, Mar. 2008.

[5] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, “SENSE: sensitivity encoding for fast MRI,” *Mag. Res. Med.*, vol. 42, no. 5, pp. 952–62, Nov. 1999.

[6] D. K. Sodickson and W. J. Manning, “Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays,” *Mag. Res. Med.*, vol. 38, no. 4, pp. 591–603, Oct. 1997.

[7] J. Tsao, P. Boesiger, and K. P. Pruessmann, “k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations,” *Mag. Res. Med.*, vol. 50, no. 5, pp. 1031–42, Nov. 2003.

[8] S. O. Schönberg, O. Dietrich, M. F. Reiser, and A. L. Baert, *Parallel Imaging in Clinical MR Applications*, Medical Radiology Diagnostic Imaging Series. Springer, 2007.

[9] M. J. Allison, S. Ramani, and J. A. Fessler, “Accelerated regularized estimation of MR coil sensitivities using augmented Lagrangian methods,” *IEEE Trans. Med. Imag.*, vol. 32, no. 3, pp. 556–64, Mar. 2013.

[10] M. J. Allison, S. Ramani, and J. A. Fessler, “Regularized MR coil sensitivity estimation using augmented Lagrangian methods,” in *Proc. IEEE Intl. Symp. Biomed. Imag.*, 2012, pp. 394–7.

[11] L. L. Wald, “Parallel Imaging Update: How Many Elements Do We Need?,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2006, p. 202.

[12] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: Physical principles and sequence design*, Wiley, New York, 1999.

[13] T. A. Bley, O. Wieben, C. J. François, J. H. Brittain, and S. B. Reeder, "Fat and water magnetic resonance imaging," *J. Mag. Res. Im.*, vol. 31, no. 1, pp. 4–18, Jan. 2010.

[14] W. Huh, J. A. Fessler, and A. A. Samsonov, "Water-fat decomposition with regularized field map," in *Proc. Intl. Soc. Mag. Res. Med.*, 2008, p. 1382.

[15] W. A. Grissom, V. Rieke, A. B. Holbrook, Y. Medan, M. Lustig, J. Santos, M. V. McConnell, and K. B. Pauly, "Hybrid referenceless and multibaseline subtraction MR thermometry for monitoring thermal therapies in moving organs," *Med. Phys.*, vol. 37, no. 9, pp. 5014–26, 2010.

[16] H. Eggers and P. Börnert, "Chemical shift encoding-based waterfat separation methods," *J. Mag. Res. Im.*, To appear in 2014.

[17] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig, "Compressed sensing for chemical shift-based water-fat separation," *Mag. Res. Med.*, vol. 64, no. 6, pp. 1749–59, Dec. 2010.

[18] S. D. Sharma, H. H. Hu, and K. S. Nayak, "Accelerated water-fat imaging using restricted subspace field map estimation and compressed sensing," *Mag. Res. Med.*, vol. 67, no. 3, pp. 650–9, Mar. 2012.

[19] S. D. Sharma, H. H. Hu, and K. S. Nayak, "Chemical shift encoded water-fat separation using parallel imaging and compressed sensing," *Mag. Res. Med.*, vol. 69, no. 2, pp. 456–66, Feb. 2013.

[20] C. N. Wiens, C. M. McCurdy, J. D. Willig-Onwuachi, and C. A. McKenzie, "R2*-corrected waterfat imaging using compressed sensing and parallel imaging," *Mag. Res. Med.*, vol. 71, no. 2, pp. 608–616, 2014.

[21] M. J. Allison and J. A. Fessler, "Accelerated computation of regularized field map estimates," in *Proc. Intl. Soc. Mag. Res. Med.*, 2012, p. 0413.

[22] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1-3, pp. 293–318, Apr. 1992.

[23] R. Glowinski and P. L. Tallec, *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*, Soc. Indust. Appl. Math., 1989.

[24] P. J. Huber, *Robust statistics*, Wiley, New York, 1981.

[25] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, "Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate," *ACM Trans. Math. Software*, vol. 35, no. 3, pp. 22:1–22:14, Oct. 2008.

[26] S. Ramani and J. A. Fessler, "Parallel MR image reconstruction using augmented Lagrangian methods," *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 694–706, Mar. 2011.

CHAPTER II

Background

This section presents the pertinent background for the remainder of this thesis. It begins with a derivation of the discrete MR imaging model starting from the classical description of the underlying MR physics (the organization of which is based on [1–3]). It then presents several advanced MR imaging techniques that are the bases of later chapters. It concludes with an overview of several image reconstruction concepts that are used throughout this work.

2.1 MRI Background

Magnetic resonance imaging is a non-invasive medical imaging modality that has gained widespread acceptance due to its flexibility and its lack of ionizing radiation [4]. It uses three types of electromagnetic fields to image both body tissue and its metabolic properties: the main magnetic field ($\mathbf{B}_0$), the radiofrequency (RF) field ($\mathbf{B}_1$), and field gradients ($\boldsymbol{G}$). It is the interaction of these fields with certain types of atoms in the body that gives rise to MRI [1].

2.1.1 Magnetic Moments and Electromagnetic Fields

Spins

In MRI, a spin is an atom that has an odd number of protons or neutrons and thus possesses a property called nuclear spin angular momentum. Hydrogen (^{1}H) is the most studied spin in the body as it is the most common (present in H_2O) and produces the largest signal [1]. Spins can be thought of in a classical sense as electrically charged, spinning gyroscopes that generate a magnetic dipole moment co-linear to their angular momentum [4, 5]. Under normal circumstances, the spins in the body do not exhibit a net magnetic moment as they are randomly oriented [1]. However, they can be manipulating into generating MR signals using electromagnetic fields [2].

Main magnetic field ($\mathbf{B}_0$)

Spins align either parallel or anti-parallel to an external magnetic field, $\mathbf{B}_0$ [4]. The direction of $\mathbf{B}_0$ is typically referred to as the longitudinal or z-direction [1]. The parallel orientation is the lower energy state and thus, there will be marginally more spins in this direction. This imbalance results in a net magnetic moment in the z-direction with magnitude proportional to the number of spins per unit volume, M_0 [2].

If a spin is tipped from its alignment, it will precess around the z-direction at the Larmor frequency, ω. The value of this frequency is dictated by the Larmor equation

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

where γ is the gyromagnetic ratio (a fixed value specific to each atom) and B is the magnetic field strength. Hydrogen has $\gamma/2\pi = 42.58$ MHz/T, which yields a Larmor frequency of 63 MHz for a typical field strength of 1.5 T [1, 2].

Radiofrequency field ($\mathbf{B}_1$)

Spins in the presence of $\mathbf{B}_0$ are resonant in the sense that they can be excited out of equilibrium. This is done by applying a radiofrequency field ($\mathbf{B}_1$), rotating at the Larmor frequency, in the plane perpendicular to the z-direction (referred to as the transverse or xy-plane) [1]. In the classical sense, the effect of the $\mathbf{B}_1$ field can be viewed as applying a torque to the spinning gyroscopes and tipping them towards the xy-plane [4]. The degree of tipping, referred to as the tip angle, is determined by the duration and strength of $\mathbf{B}_1$ (usually on the order of microteslas and milliseconds) [1]. For a single excitation, a 90° tip angle is often used as it causes the net magnetization to lie entirely in the xy-plane which results in a strong signal [1,3].

Field gradients ($\boldsymbol{G}$)

Localizing components of a signal generated within a uniform magnetic field (e.g., $\mathbf{B}_0$) is difficult as all of the spins of a specific type are excited by, and precess at, the same frequency ($\omega_0 = \gamma B_0$). To counter this, linear-gradient magnetic fields ($\boldsymbol{G}$) are added to $\mathbf{B}_0$. The direction of these field gradients is the same as $\mathbf{B}_0$, but their magnitudes vary linearly in either the x, y, or z-direction [1]. Thus, the spatial location of a spin, (x, y, z), is encoded in its frequency:

$$\omega(x, y, z) = \gamma(B_0 + G_x x + G_y y + G_z z), \tag{2.2}$$

where G_x, G_y, and G_z are the magnetization gradients in the x, y, and z directions, respectively [2,6].

Relaxation and the Bloch equation

After $\mathbf{B}_1$ has ceased, the spins precess back to equilibrium in a process called relaxation [2]. There are two dominant forms of relaxation: longitudinal and transverse.

Longitudinal relaxation is the return to equilibrium of the z-direction magnetization and is described by

$$M_z(t) = M_0(1 - e^{-t/T_1}), \tag{2.3}$$

where $M_z(t)$ is the magnitude of the magnetization in the z-direction at time t, M_0 is the steady state magnetization (which is parallel to $\mathbf{B}_0$)[1], and T_1 is the spin-lattice time constant.[2] T_1 accounts for the exchange of energy between the spins and their atomic surroundings and is thus tissue dependent [1,4]. Transverse relaxation is the decay of the magnetization in the xy-plane and is described by

$$M_{xy}(t) = M_0 e^{-t/T_2}, \tag{2.4}$$

where M_{xy} is the magnitude of the magnetization in the xy-plane and T_2 is the spin-spin time constant which accounts for the inter-spin interactions [1,5]. Typical values of T_1 are 100-2000 ms, while typical values of T_2 are 10-300 ms [1].

Macroscopically, the net magnetization vector of the excited spins, $\mathbf{M}$, can be thought of as precessing back to equilibrium at the Larmor frequency. The Bloch equation describes this time-varying behavior:

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma \mathbf{B} - \frac{M_x\,\mathbf{i} + M_y\,\mathbf{j}}{T_2} - \frac{(M_z - M_0)\,\mathbf{k}}{T_1}, \tag{2.5}$$

where M_x and M_y are the magnitudes of the magnetization in the x and y directions, $\mathbf{i}, \mathbf{j}, \mathbf{k}$ are the unit-length vectors corresponding to the x, y, z directions, and $\mathbf{B}$ is the combination of the previously discussed magnetic fields [1]. This rotating vector induces an electromotive force in a receive coil placed next to the patient [1]. It is

---

[1] To better differentiate between the magnetic field strength (B) and magnetization (M), it helps to consider the units. The units of B are Tesla or $\mathrm{N \cdot A^{-1} \cdot m^{-1}}$ whereas the units of M are $\mathrm{A \cdot m^{-1}}$.

[2] This equation assumes a 90° excitation.

this induced force that makes up the measured signal in MRI. Thus, an MR signal equation can be determined by solving the Bloch equation [2].[3]

2.1.2 Acquiring MR Imaging Data

A typical MRI acquisition involves two stages [1]. In the first stage, excitation, $\mathbf{B}_1$ is used to excite a volume of spins and the field gradients $\boldsymbol{G}$ are used to encode spatial location. In the second stage, reception, the precessing transverse magnetization is measured using a receive coil. The reception stage is typically short in duration so as to only measure the signal when it is excited, and not when it has returned to equilibrium. Thus, this process of excitation and reception is repeated many times during the acquisition of a single image to capture all of the required information [1].

Excitation

Excitation is the process of tipping the spins from equilibrium using $\mathbf{B}_1$. There are two types of excitation: non-selective and selective [1]. Non-selective excitation excites the entire volume of spins by applying $\mathbf{B}_1$ in the presence of only $\mathbf{B}_0$. The receive coil then detects signal from the entire volume and 3-D imaging techniques are required to distinguish signal locations [1]. Selective excitation excites only a specific region of the volume by adding field gradients to $\mathbf{B}_0$ during the excitation. Typically, the selected region is a thin slice which allows for 2-D imaging techniques. For example, a slice perpendicular to the z-direction can be excited by applying G_z during excitation. By doing so, the Larmor frequency of each spin depends on its z-location ($\omega(z) = \gamma(B_0 + G_z z)$). As $\mathbf{B}_1$ only excites spins with the same frequency, the location and shape of the excited slice can be controlled using the frequency and shape of the RF pulse [1].

---

[3] There is no known closed-form solution for the general Bloch equation with all fields present; however, solutions exist under certain reasonable assumptions [2].

Reception

In the common case of the RF field being inactive ($\mathbf{B}_1 = 0$) during reception, the Bloch equation (2.5) can be solved to yield an equation describing the behaviour of the local transverse magnetization, $M(\vec{r}, t)$, after excitation:

$$M(\vec{r}, t) = M(\vec{r}, 0)e^{-t/T_2(\vec{r})}e^{-i\omega_0 t}\text{exp}\left(-i\gamma \int_0^t \mathbf{G}(\tau) \cdot \vec{r} \ \mathrm{d}\tau\right) \tag{2.6}$$

where $M(\vec{r}, 0)$ is the initial transverse magnetization, $\vec{r} = [x \ \ y \ \ z]$, and $\mathbf{G} \cdot \vec{r} = G_x x + G_y y + G_z z$ [1].

The receive coil detects changes in the flux within the xy-plane [1], so assuming it has uniform sensitivity, the detected signal is proportional to

$$s_r(t) = \int_{vol} M(\vec{r}, t) \ \mathrm{dV}. \tag{2.7}$$

Substituting (2.6) into this equation yields,

$$s_r(t) = \iiint M(\vec{r}, 0)e^{-t/T_2(\vec{r})}e^{-i\omega_0 t}\text{exp}\left(-i\gamma \int_0^t \mathbf{G}(\tau) \cdot \vec{r} \ \mathrm{d}\tau\right) \mathrm{dx \ dy \ dz}. \tag{2.8}$$

Assuming that a perfectly rectangular slice parallel to the z-direction is excited, that G_z is off during receiving, and that the acquisition time is short enough that the T_2 decay term is virtually constant over the readout, (2.8) is demodulated at frequency ω_0 to obtain an equation for the 2-D imaging signal:

$$s_{xy}(t) = \iint m(x, y)\text{exp}\left(-i\gamma \int_0^t G_x(\tau)x + G_y(\tau)y \ \mathrm{d}\tau\right) \mathrm{dx \ dy}, \tag{2.9}$$

where

$$m(x, y) \approx \int M(\vec{r}, 0)e^{-T_\mathrm{E}/T_2(\vec{r})} \ \mathrm{dz} \tag{2.10}$$

is the integral of the magnetization over the slice which is actually a function of both the magnetic properties of the object being scanned (e.g., spin density, T_1, and T_2)

as well as the scan parameters such as the echo time, T_{E}, which is approximately the time when the readout occurs [1, 5].[4]

Equation (2.9) can be written in terms of spatial frequency (referred to as k-space in the MR literature) to obtain the essential 2-D imaging signal equation:

$$\begin{aligned} s(t) &= \iint m(x,y) e^{-i2\pi(k_x(t)x + k_y(t)y)} \,\mathrm{dx}\,\mathrm{dy}, \\ &= \mathcal{F}_{2D}\{m(x,y)\}\Big|_{\substack{k_x(t)\\ k_y(t)}} \end{aligned} \tag{2.11}$$

where

$$k_x(t) = \frac{\gamma}{2\pi}\int_0^t G_x(\tau)\,\mathrm{d}\tau \quad \text{and} \quad k_y(t) = \frac{\gamma}{2\pi}\int_0^t G_y(\tau)\,\mathrm{d}\tau \tag{2.12}$$

are the k-space (spatial frequency) sample locations with units cycles/cm [1, 5].

The magnitude and duration of the gradients are designed so that the signal equation adequately samples k-space. The manner in which k-space is sampled is referred to as the sampling trajectory. The most common trajectories acquire an evenly spaced grid of k-space locations, referred to as Cartesian sampling [6]. Such an acquisition is beneficial as it allows for the reconstruction of the image using inverse fast-Fourier transform (FFT) techniques.

Noise

The noise in the MR signal is predominantly from random spin fluctuations within the patient and thermal noise resulting from resistance in the receive coil. As such, it is modeled as complex, additive white noise with a Gaussian probability distribution [7]. In the case of fully sampled Cartesian imaging, the noise in the reconstructed image is also white Gaussian as the discrete Fourier transform (DFT) is unitary.

---

[4] 3-D MRI is a simple extension of the 2-D signal equation presented in this section. The difference is that the entire volume is excited and G_z is applied in the same manner as G_x and G_y (usually after excitation).

Complex MR images are often viewed as magnitude images, in this situation the noise in the object voxels has a Rician distribution [7]. In the case of undersampled imaging, the noise may no longer have a white Gaussian distribution as the associated transforms are typically non-unitary.

2.1.3 Discrete Image Model

To create a discrete image, the signal equation (2.11) is sampled to collect a series of complex valued MR measurements, $\{\mathrm{y}_1, \ldots, \mathrm{y}_{n_d}\}$. These measurements are modeled as

$$\mathrm{y}_i = s(t_i) + \epsilon_i, \qquad i = 1, \ldots, n_d, \tag{2.13}$$

where t_i is the time index of the ith measurement and ϵ is additive noise [3, 8]. To facilitate computer processing, the object $m(\vec{r})$ is also paramaterized using a linear combination of n_p evenly spaced spatial basis functions, b:

$$m(\vec{r}) = \sum_{j=1}^{n_p} \mathrm{x}_j b(\vec{r} - \vec{\mathrm{x}}_j), \tag{2.14}$$

where $\vec{\mathrm{x}}_j$ is the spatial location corresponding to coefficient x_j [3]. The signal equation (2.11) then becomes

$$\begin{aligned} s(t_i) &= \int \left[\sum_{j=1}^{n_p} \mathrm{x}_j b(\vec{r} - \vec{\mathrm{x}}_j) \right] e^{-i2\pi \vec{k}(t_i)\cdot\vec{r}} \, \mathrm{d}\vec{r} \\ &= \sum_{j=1}^{n_p} a_{ij} \mathrm{x}_j, \end{aligned} \tag{2.15}$$

where

$$a_{ij} = \int b(\vec{r} - \vec{\mathrm{x}}_j) \; e^{-i2\pi \vec{k}(t_i)\cdot\vec{r}} \, \mathrm{d}\vec{r}, \tag{2.16}$$

and $\vec{k}(t_i)$ is the k-space sample location at time t_i [3]. Equation (2.16) can be further simplified through a change of variables, yielding

$$\begin{aligned} a_{ij} &= \int b(\vec{q})\ e^{-i2\pi\vec{k}(t_i)\cdot(\vec{q}+\vec{\mathrm{x}}_j)}\ \mathrm{d}\vec{q} \\ &= e^{-i2\pi\vec{k}(t_i)\cdot\vec{\mathrm{x}}_j} \int b(\vec{q})\ e^{-i2\pi\vec{k}(t_i)\cdot\vec{q}}\ \mathrm{d}\vec{q} \\ &= e^{-i2\pi\vec{k}(t_i)\cdot\vec{\mathrm{x}}_j} B\{\vec{k}(t_i)\}, \end{aligned} \tag{2.17}$$

where $B\{\vec{k}\}$ is the Fourier transform of the basis function evaluated at $\vec{k}$ [3]. For simplicity, we consider impulse basis functions in our subsequent analysis. Thus, $B\{\vec{k}\} = 1$ for all $\vec{k}$ and the signal equation in (2.15) becomes

$$s(t_i) = \sum_{j=1}^{n_p} \mathrm{x}_j e^{-i2\pi\vec{k}(t_i)\cdot\vec{\mathrm{x}}_j}, \tag{2.18}$$

which, in the case of evenly spaced basis functions (and assuming Cartesian k-space sampling), is the DFT of $\{\mathrm{x}_1, \ldots, \mathrm{x}_{n_p}\}$ evaluated at $\vec{k}(t_i)$.

From (2.18), the measurement model can be written in the following matrix-vector form:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}, \tag{2.19}$$

where

$$\mathbf{y} = \begin{bmatrix} \mathrm{y}_1 \\ \vdots \\ \mathrm{y}_{n_d} \end{bmatrix}, \qquad \boldsymbol{\epsilon} = \begin{bmatrix} \epsilon_1 \\ \vdots \\ \epsilon_{n_d} \end{bmatrix}, \qquad \mathbf{x} = \begin{bmatrix} \mathrm{x}_1 \\ \vdots \\ \mathrm{x}_{n_p} \end{bmatrix} \tag{2.20}$$

and $[\mathbf{A}]_{(i,j)} = e^{-i2\pi\vec{k}(t_i)\cdot\vec{\mathrm{x}}_j}$ which is the element of $\mathbf{A}$ in the ith row and jth column [3,8]. Thus, the image reconstruction problem becomes one of estimating the n_p long parameter vector, $\mathbf{x}$, from the n_d measurements, $\mathbf{y}$.

In many applications, the k-space samples are evenly spaced and the acquisition is designed so that $n_p = n_d$ with $\mathbf{A}$ being a Fourier encoding matrix. Then, the measurement noise is often ignored and the magnetization image, $\hat{\mathbf{x}}$, is reconstructed using an inverse fast-Fourier transform. If $n_p = n_d$ but noise is significant, the magnetization image may be reconstructed using regularized estimation techniques such as penalized-likelihood methods. If $n_d \ll n_p$, more advanced regularized estimation techniques such as compressed sensing or parallel imaging techniques may be used (see Section 2.4).

2.2 Main Magnetic Field Inhomogeneity ($\Delta \mathbf{B}_0$)

Up to this point, the main magnetic field ($\mathbf{B}_0$) has been assumed to be spatially uniform over the entire field-of-view. However, there are several common sources of field inhomogeneity. First, typical main magnetic field coil designs result in non-uniform magnetic fields over large field-of-views and although shimming can reduce the degree of this inhomogeneity, it fails to eliminate it [4, p. 837]. Second, differences in the bulk magnetic susceptibility of tissues in the patient can result in large distortions in the main magnetic field [4, p. 762]. These distortions are particularly prevalent at air-tissue interfaces such as the lungs and at the surfaces of metallic implants.

To determine the effects of this inhomogeneity on the reconstructed image, we return to (2.2). Adding the field inhomogeneity at spatial location (x, y, z), $\Delta B_0(x, y, z)$, yields

$$\omega(x, y, z) = \gamma(B_0 + \Delta B_0(x, y, z) + G_x x + G_y y + G_z z). \tag{2.21}$$

Using this new term in (2.9) and ignoring inhomogeneity in the z-direction[5] gives

$$s_{xy}(t) = \iint m(x,y) \exp\left(-i\gamma \Delta B_0(x,y) t\right) \exp\left(-i\gamma \int_0^t G_x(\tau)x + G_y(\tau)y \, d\tau\right) \text{dx dy}, \tag{2.22}$$

resulting in the following modified signal equation

$$\begin{aligned} s(t) &= \iint m(x,y) e^{-i\gamma \Delta B_0(x,y)t} e^{-i2\pi(k_x(t)x + k_y(t)y)} \text{dx dy}, \\ &= \mathcal{F}_{2D}\{m(x,y) e^{-i\gamma \Delta B_0(x,y)t}\}\Big|_{\substack{k_x(t) \\ k_y(t)}}. \end{aligned} \tag{2.23}$$

Thus, the acquired signal is equal to the Fourier transform of the product of the magnetization image and a complex exponential containing the term $\gamma \Delta B_0(x,y)t$, which has units of radians.

The additional complex exponential term resulting from magnetic field inhomogeneity can cause several types of artifacts and distortions including scaling, stretching, and the addition of phase in the reconstructed image (see [4, Chap. 20]). These effects are particularly noticeable in scans with long readout times (as the value in the complex exponential term increases linearly with time) such as echo-planar imaging (EPI) and spiral acquisitions [9]. In these situations, it can be advantageous to estimate the field inhomogeneity which can then be used to perform field-corrected MR image reconstruction (e.g. [8–10]). Chapter IV addresses this estimation problem.

One possible source of confusion on this subject stems from the terminology used in field inhomogeneity estimation. In many of these works, the authors imply that the term 'field map' refers to the field inhomogeneity at each voxel in the field-of-view. However, they are actually estimating the difference in the rotational frequency at

---

[5] Z-direction inhomogeneity in 2-D imaging can alter the shape of the selected slice and the perceived location of tissue within that slice; however, this is an acceptable approximation for thin slices.

each voxel:

$$\Delta\omega(x, y) = \gamma \Delta B_0(x, y). \tag{2.24}$$

For most problems, the true magnetic field inhomogeneity can be determined from this term [11]; however, more advanced methods are required in the presence of chemical shift (see Section 2.3 and Chapter V for additional details).

2.3 Chemical Shift and Water-Fat Imaging

Depending upon the molecule in which they are present, hydrogen protons may experience slightly different magnetic environments [4, p. 7]. This can cause different tissues within $\mathbf{B}_0$ to be under slightly different effective magnetic fields, and thus, have local Larmor frequencies shifted from the expected ω_0. The degree of shift is described by the 'chemical shift' δ of the molecule, which is defined as

$$\delta = \frac{-\Delta\omega}{\omega_0} \tag{2.25}$$

and usually expressed in units of parts-per-million (ppm) [4, p. 200]. In cases of known ω_0 (i.e., a specific magnetic field strength), it is common to state the resulting frequency shift $\Delta\omega$ rather than the chemical shift.

To analyze the effects of chemical shift on the reconstructed image, we modify (2.2) by adding the frequency shift at spatial location (x, y), $-\Delta\omega(x, y)$:

$$\omega(x, y, z) = \gamma(B_0 + G_x x + G_y y + G_z z) - \Delta\omega(x, y). \tag{2.26}$$

Following the same steps as Section 2.2, we arrive at the following signal equation:

$$\begin{aligned} s(t) &= \iint m(x, y) e^{i\Delta\omega(x,y)t} e^{-i2\pi(k_x(t)x + k_y(t)y)} \mathrm{dx\ dy}, \\ &= \mathcal{F}_{2D}\{m(x, y) e^{i\Delta\omega(x,y)t}\}\Big|_{\substack{k_x(t) \\ k_y(t)}}. \end{aligned} \tag{2.27}$$

The resulting acquired signal is equal to the Fourier transform of the product of the magnetization image and a complex exponential which, like in the case of field inhomogeneity, can result in artifacts and distortions in the reconstructed images.

One particularly challenging situation where chemical shift artifacts can occur is when imaging objects with both water and fatty tissue. The electronic configuration of triglyceride molecules in the fatty tissue shields the protons resulting in a chemical shift δ of approximately 3.5 ppm [12].[6] Thus, the frequency shift of the fatty tissue $\Delta\omega_{fat}$ is approximately 220 Hz for a 1.5 T scan. The resulting signal equation (ignoring field inhomogeneity and other sources of chemical shift) can be modeled as [12]

$$s(t) = \mathcal{F}_{2D}\{w(x,y) + f(x,y)e^{i\Delta\omega_{fat}t}\}\Big|_{\substack{k_x(t)\\ k_y(t)}} \tag{2.28}$$

where $w(x,y)$ and $f(x,y)$ are the water and fat components of the magnetization at (x,y) respectively [14]. If not properly accounted for, the chemical shift will cause the fat component of the image to spatially shift relative to the water component [4]. Chapter V addresses water-fat imaging techniques that account for the chemical shift.

2.4 Accelerated MR Imaging

Section 2.1 describes how traditional MRI provides a means for signal localization by encoding object contrast in the spatial-frequency domain using field gradients ($\boldsymbol{G}$) [15]. The limitation of this approach is that only one k-space position can be sampled at a time. There are also physical and physiological limits on how fast k-space can be sampled [15]. Thus, to accelerate the acquisition of MR images, we

---

[6] The true chemical shift of fatty tissue is a spectrum with several distinct peaks [13]; however, most methods focus on the main peak.

must reduce the number of k-space samples. However, reducing the number of k-space samples for a given image violates the Nyquist criterion, resulting in aliasing artifacts in conventional reconstructions [16]. In this thesis, we use two accelerated MR imaging techniques capable of reconstructing images from undersampled k-space data: parallel imaging and compressed sensing.

2.4.1 Parallel Imaging

The previous derivation of the signal equation assumed that the receive coil had uniform sensitivity. However, for surface coils, the contribution of the object magnetization to the measured signal varies markedly with spatial location [15]. Knowledge of this varying sensitivity can be used to derive additional information about the spatial distribution of the magnetization. Furthermore, coil sensitivity is a receiver specific property. Thus, several distinct samples of the object can be simultaneously acquired by using multiple receive coils in parallel. This additional information presents the possibility of accurate reconstruction from highly undersampled k-space data and forms the basis of parallel MRI acceleration methods such as Sensitivity Encoding (SENSE) [15] and Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [17]. We focus on SENSE imaging as it is the underlying method of Chapter III.

SENSE imaging

In SENSE imaging, the signal equation (2.11) for each receive coil becomes

$$\begin{aligned} s_l(t) &= \iint m(x,y)c_l(x,y)e^{-i2\pi(k_x(t)x+k_y(t)y)}\mathrm{dx\ dy} \\ &= \mathcal{F}_{2D}\{m(x,y)c_l(x,y)\}\Big|_{\substack{k_x(t)\\k_y(t)}}\ , \end{aligned} \tag{2.29}$$

where $c_l(x, y)$ is the sensitivity of the lth coil at spatial location (x, y). Sampling and discretizing this equation as in Section 2.1.3 yields

$$\mathbf{y} = \mathbf{B}\mathbf{x} + \boldsymbol{\epsilon}, \qquad \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_L \end{bmatrix}, \qquad \mathbf{B} = \begin{bmatrix} \mathbf{A}\mathbf{C}_1 \\ \vdots \\ \mathbf{A}\mathbf{C}_L \end{bmatrix}, \tag{2.30}$$

where $\mathbf{y}_l \in \mathbb{C}^{n_d}$ is the measurement vector from the lth coil, $\mathbf{x} \in \mathbb{C}^{n_p}$ is the desired magnetization image, $\mathbf{C}_l = \operatorname{diag}\{c_{lj}\}$ with c_{lj} denoting the jth coefficient of the the lth parametrized sensitivity profile, $\mathbf{A}$ is the Fourier encoding matrix, $\boldsymbol{\epsilon} \in \mathbb{C}^{n_p}$ accounts for the measurement noise, and L denotes the number of coils.

Since the sensitivity profiles differ, the system of equations in (2.30) is overdetermined in the usual case where $Ln_d > n_p$. Thus, given the sensitivity profiles $\mathbf{C}_l$, the underlying image $\mathbf{x}$ can be reconstructed using either a least-squares or penalized least-squares method. Although the problem remains overdetermined for reduction factors of up to L, traditional parallel imaging methods (e.g., [15]) cannot achieve such extreme acceleration rates in practice due to increased noise levels and coil coupling. However, recent advances in sparse reconstruction, such as compressed sensing [16, 18, 19], have allowed for parallel imaging with such levels of acceleration [20, 21].

The major limitation of this method is that precise sensitivity profiles are necessary to obtain accurate reconstructions of $\mathbf{x}$ [22]. These profiles cannot be precomputed as they depend on varying environmental factors such as coupling with dielectric materials in the body [23]. Instead, they must be estimated from data collected at the time of each scan. Chapter III addresses this estimation problem.

2.4.2 Compressed Sensing

Compressed sensing (CS) provides a framework for accurately reconstructing a signal from far fewer samples than dictated by the Nyquist criterion [24]. CS theory states that a signal can be reconstructed with reasonable accuracy from a small number of samples if three conditions hold [16]. First, the desired signal or its transform must have a sparse representation. Second, any undersampling artifacts in the sparsifying domain resulting from a linear reconstruction should be incoherent. Third, the reconstruction must be performed using a nonlinear method which enforces both data consistency and sparsity [16, 25].

MRI satisfies these three requirements and is thus a good candidate for the application of CS [16]. First, nearly all MR images are sparse in some transform domain (often the wavelet transform domain or the domain of spatial finite differences) [25]. Second, a high degree of incoherence is guaranteed if the k-space samples are selected at random [19]. This condition is not practical in 2-D Cartesian MRI as sampling trajectories must follow relatively smooth paths due to physiological and physical constraints. However, by using randomly selected phase encode lines in 3-D Cartesian MRI, such sampling can be achieved for 2-D cross sections. Third, several nonlinear reconstruction methods have been proposed for MR imaging including the following optimization problem:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} ||\Psi\mathbf{x}||_1 \ \text{ s.t. } \ ||\mathcal{F}_u\mathbf{x} - \mathbf{y}||_2 < \lambda, \tag{2.31}$$

where Ψ is the sparsifying transform, $\mathcal{F}_u$ is the corresponding undersampled Fourier operator, and λ is a parameter that controls the degree of data consistency [16]. We explore several minimization strategies for cost functions with this form in Chapter V.

Optimization Transfer Method [3]

```
Repeat until stop criterion is satisfied:
    1. Select a surrogate function $\Phi^{(n)}$ that satisfies (2.32).
    2. Find $\mathbf{x}^{(n+1)}$ s.t. $\Phi^{(n)}(\mathbf{x}^{(n+1)}) \leq \Phi^{(n)}(\mathbf{x}^{(n)})$.
    3. $n = n + 1$.
```

Figure 2.1: Overview of the optimization transfer method.

2.5 Optimization Transfer Methods

Many of the methods in this thesis indirectly estimate parameters through the minimization of specially designed cost functions. Optimization transfer methods provide powerful tools for minimizing the complicated cost functions that are often encountered in medical imaging. They allow for more specialized solutions to problems than those typically obtained from the standard steepest descent methods [3, 26]. As summarized in Fig. 2.1, the premise behind optimization transfer is to replace the original cost function Ψ with a simpler surrogate function, $\Phi^{(n)}$. As long as this surrogate function satisfies the following surrogate conditions, the minimizer of this function will also reduce the original cost function [3]:

$$\begin{aligned} \Phi^{(n)}(\mathbf{x}^{(n)}) &= \Psi(\mathbf{x}^{(n)}) \\ \Phi^{(n)}(\mathbf{x}) &\geq \Psi(\mathbf{x}) \quad \forall \mathbf{x} \end{aligned} \tag{2.32}$$

In fact, the surrogate function only needs to be reduced and not fully minimized. By iteratively constructing and minimizing a surrogate function at each estimate, we are guaranteed to monotonically decrease the original cost function [3].

2.6 Augmented Lagrangian Methods

The augmented Lagrangian (AL) method is a minimization technique that adds Lagrange multiplier estimates to the function being minimized, thus reducing the

```
Augmented Lagrangian Method
Initialize: $\mathbf{x}^{(0)}$ and $\boldsymbol{\gamma}^{(0)}$.
Select $\mu_0$ and set $k = 0$.
Repeat until stop criterion is achieved:
    1. Find $\mathbf{x}^{(k+1)}$ that approximately minimizes $\mathcal{L}(\cdot, \boldsymbol{\gamma}^{(k)}, \mu_k)$,
    2. $\boldsymbol{\gamma}^{(k+1)} = \boldsymbol{\gamma}^{(k)} + \mu_k \left(\mathbf{A}\mathbf{x}^{(k+1)} - \mathbf{c}\right)$,
    3. Choose $\mu_{k+1} \geq \mu_k$,
    4. $k = k + 1$.
```

Figure 2.2: Overview of the general AL method [27, Ch. 17].

ill-conditioning that is commonly encountered by other penalty methods such as quadratic penalties [27, Ch. 17]. Although there are more general forms of the AL method, we consider the following equality-constrained optimization problem commonly encountered in medical imaging:

$$\arg\min_{\mathbf{x}} f(\mathbf{x}) \text{ s.t. } \mathbf{A}\mathbf{x} = \mathbf{c}, \tag{2.33}$$

where $\mathbf{x} \in \mathbb{R}^N$, f is a real, closed, proper convex function, $\mathbf{A} \in \mathbb{R}^{M \times N}$, and $\mathbf{c} = \mathbb{R}^M$ [27, 28]. The AL method introduces a vector of Lagrange multipliers $\boldsymbol{\gamma} \in \mathbb{R}^M$ to (2.33), resulting in the following AL function:

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\gamma}, \mu) = f(\mathbf{x}) + \langle \mathbf{A}\mathbf{x} - \mathbf{c}, \boldsymbol{\gamma} \rangle + \frac{\mu}{2} \|\mathbf{A}\mathbf{x} - \mathbf{c}\|_2^2, \tag{2.34}$$

where $\mu > 0$ is an AL parameter [27,28]. The resulting AL method finds a minimizer of the original problem (2.33) by updating the estimates for $\mathbf{x}$ and $\boldsymbol{\gamma}$ as outlined in Fig. 2.2. One important distinction between AL and simpler penalty methods is that, for typical problems and assuming a sufficiently large μ_0, the value of μ does not need to be increased with iteration to ensure convergence [29].

In this thesis, we will often use a scaled AL algorithm similar to the ones presented in [21, 28] as it is a more natural form for complex valued variables. This reformulation is found by completing the square in (2.34) and ignoring the constant

```
Scaled Augmented Lagrangian Method
Initialize: $\mathbf{x}^{(0)}$ and $\boldsymbol{\eta}^{(0)}$.
Select $\mu_0$ and set $k = 0$.
Repeat until stop criterion is achieved:
    1. Find $\mathbf{x}^{(k+1)}$ that approximately minimizes $\mathcal{L}(\cdot, \boldsymbol{\eta}^{(k)}, \mu_k)$,
    2. $\boldsymbol{\eta}^{(k+1)} = \boldsymbol{\eta}^{(k)} - \left(\mathbf{A}\mathbf{x}^{(k+1)} - \mathbf{c}\right)$,
    3. Choose $\mu_{k+1} \geq \mu_k$,
    4. $k = k + 1$.
```

Figure 2.3: Overview of a scaled AL method based on [21,28].

terms:

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\eta}, \mu) = f(\mathbf{x}) + \frac{\mu}{2}\|\mathbf{A}\mathbf{x} - \mathbf{c} - \boldsymbol{\eta}\|_2^2, \tag{2.35}$$

where $\boldsymbol{\eta} \triangleq -\frac{1}{\mu}\boldsymbol{\gamma}$. This results in the modified AL method outlined in Fig. 2.3.

A special case of the AL algorithm is the alternating direction method of multipliers (ADMM) where the optimization problem is of the form

$$\arg\min_{\mathbf{x},\mathbf{z}} \; f(\mathbf{x}) + g(\mathbf{z}) \text{ s.t. } \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} = \mathbf{c}, \tag{2.36}$$

with g denoting a real closed, proper, convex function, $\mathbf{z} \in \mathbb{R}^L$, and $\mathbf{B} \in \mathbb{R}^{M \times L}$ [28,30]. The resulting scaled AL function is

$$\mathcal{L}_M(\mathbf{x}, \mathbf{z}, \boldsymbol{\eta}, \mu) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\mu}{2}\|\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c} - \boldsymbol{\eta}\|_2^2. \tag{2.37}$$

The ADMM approach uses one pass of the Gauss-Seidel method to approximately find the joint minimizer of (2.37) [28]. The resulting algorithm, Fig. 2.4, is guaranteed to converge as long as the minimization errors in the variable update steps are absolutely summable [30].[7]

There are numerous convergence properties for the AL and ADMM algorithms proven for the case of real valued data [28,30]. However, the same cannot be said for

---

[7] Assuming the AL function has a saddle point [28].

```
Alternating Direction Method of Multipliers
Initialize: $\mathbf{x}^{(0)}$, $\mathbf{z}^{(0)}$, and $\boldsymbol{\gamma}^{(0)}$.
Select $\mu_0$ and set $k = 0$.
Repeat until stop criterion is achieved:
    1. Find $\mathbf{x}^{(k+1)}$ that approximately minimizes $\mathcal{L}_M(\cdot, \mathbf{z}^{(k)}, \boldsymbol{\eta}^{(k)}, \mu_k)$,
    2. Find $\mathbf{z}^{(k+1)}$ that approximately minimizes $\mathcal{L}_M(\mathbf{x}^{(k+1)}, \cdot, \boldsymbol{\eta}^{(k)}, \mu_k)$,
    3. $\boldsymbol{\eta}^{(k+1)} = \boldsymbol{\eta}^{(k)} - \big(\mathbf{A}\mathbf{x}^{(k+1)} + \mathbf{B}\mathbf{z}^{(k+1)} - \mathbf{c}\big)$,
    4. Choose $\mu_{k+1} \geq \mu_k$,
    5. $k = k + 1$.
```

Figure 2.4: Overview of the ADMM [30].

the complex valued data typically encountered in MRI. Still, these algorithms behave similarly when applied to complex valued experimental data and in some cases we have extended the existing proofs to this setting (e.g., Chapter III).[8]

2.7 Sparse Cholesky Factorization

Many of the regularized estimation methods in this thesis involve inverting a Hessian matrix of the following form:

$$\mathbf{H} = \mathbf{D} + \lambda \mathbf{C}^{\mathrm{T}} \mathbf{C} \tag{2.38}$$

where $\mathbf{D} \in \mathbb{C}^N$ is a diagonal matrix, N is the number of elements being estimated, λ is a regularization parameter, and $\mathbf{C} \in \mathbb{R}^{M \times N}$ is a finite differencing matrix. However, the size of $\mathbf{H}$ for typical data sets renders brute force inversion techniques computationally infeasible. Whereas an efficient solution might exist for the block tri-diagonal structure of $\mathbf{H}$ produced by first-order finite differences [31], there are no such algorithms for the block penta-diagonal with penta-diagonal block structures associated with the more common second-order finite differences (see Fig. 2.5). Furthermore, the addition of the diagonal matrix limits the applicability of circulant

---

[8] If using complex valued data, the AL function is better described by $\mathcal{L}(\mathbf{x}, \boldsymbol{\gamma}, \mu) = f(\mathbf{x}) + \text{real}\{\langle \mathbf{A}\mathbf{x} - \mathbf{c}, \boldsymbol{\gamma} \rangle\} + \frac{\mu}{2} \|\mathbf{A}\mathbf{x} - \mathbf{c}\|_2^2$.

inversion techniques (see Section 3.1). There are, however, several efficient methods for solving a linear system of equations based on $\mathbf{H}$, which is sufficient to solve these regularized estimation problems. One such method that we use in this thesis is sparse Cholesky factorization [32].

Since $\mathbf{H}$ is positive definite,[9] Cholesky factorization can solve $\mathbf{H}^{-1}\mathbf{x}$ for arbitrary $\mathbf{x}$. This is done by first using Cholesky factorization to compute a lower triangular matrix $\mathbf{L}$ such that $\mathbf{H} = \mathbf{L}\mathbf{L}^{\mathrm{T}}$ and then using $\mathbf{L}$ to solve the linear system of equations through forward and backward substitution [33]. For a dense matrix, this approach has a cost of $\mathcal{O}(N^3/3)$, which is approximately two times more efficient than similar techniques such as LU factorization [34]; however, for sparse matrices, standard Cholesky factorization may be inefficient as it can result in significant infilling. In these situations, sparse Cholesky factorization techniques are used where a permutation matrix $\mathbf{P}$ is included in the decomposition (i.e., $\mathbf{H} = \mathbf{P}\mathbf{L}\mathbf{L}^{\mathrm{T}}\mathbf{P}^{\mathrm{T}}$) to reduce the degree of infilling within $\mathbf{L}$. The resulting algorithm has a much lower cost than traditional Cholesky factorization of sparse matrices, with the exact degree of savings depending on $\mathbf{P}$. The experiments in this thesis use the CHOLMOD [32] implementation of sparse Cholesky factorization which uses approximate minimum degree ordering to determine an appropriate permutation matrix.

We evaluated the sparse Cholesky operation on Hessian matrices for varying image sizes where $\mathbf{D} = \mathbf{I}$ and $\mathbf{C}$ is a 2-D, second-order finite differencing matrix. Fig. 2.5 shows the non-zero elements of one such Hessian matrix as well as the non-zero elements in the $\mathbf{L}$ matrix generated by CHOLMOD. Clearly, there is only modest infilling in the lower triangular matrix. Table 2.1 presents the time and memory required to compute $\mathbf{H}^{-1}\mathbf{x}$ for $\mathbf{x} = \mathbf{1}$ (with double precision) using the CHOLMOD

---

[9] This requires at least one element of $\mathbf{D}$ to be positive, which is usually true for any non-trivial imaging problem.

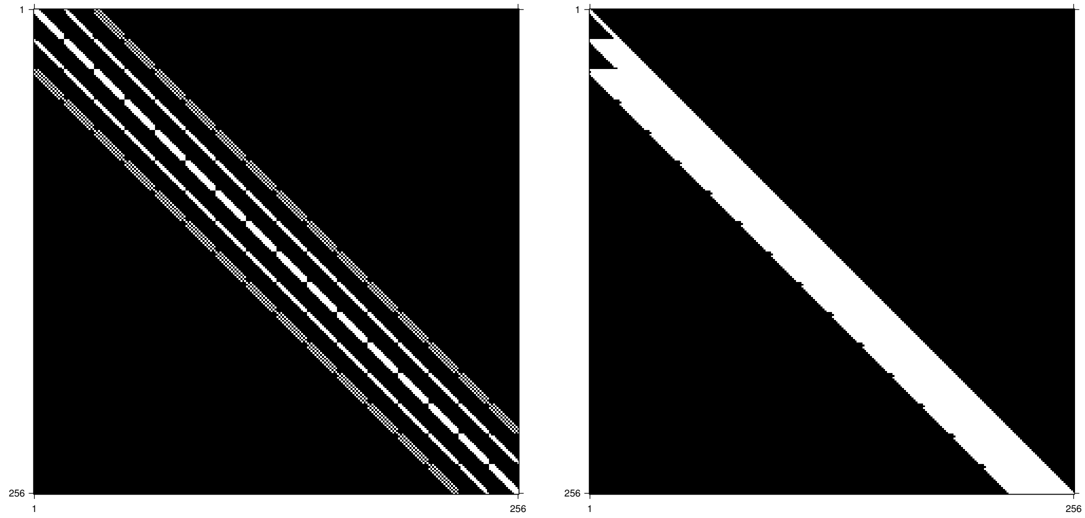

Figure 2.5: Example of the non-zero elements of a Hessian matrix with $\mathbf{D} = \mathbf{I}$ and a full 2-D second-order $\mathbf{C}$ (left) and the corresponding $\mathbf{L}$ matrix created by CHOLMOD (right).

Table 2.1: CHOLDMOD $\mathbf{H}^{-1}\mathbf{x}$ Computation Times and Memory Usage

Image Size (pixels)	**Time Solving (s)**	**Peak Memory (MB)**
16×16	0.001	0.2
32×32	0.006	1.0
64×64	0.035	5.6
128×128	0.224	28
256×256	1.237	143
512×512	7.003	715
1024×1024	274.548	3193
2048×2048	OUT OF MEMORY	

implementation included in MATLAB on a workstation with a 2.66 GHz, quad-core Intel Xeon CPU and 8 GB of memory. These experimental results demonstrate that the CHOLMOD method is efficient in both computation and memory requirements for small 2-D problems; however, it becomes computationally infeasible for larger problems. We therefore also investigate other methods for solving problems with such Hessian matrices in this thesis.

2.8 Spatial Resolution Analysis of Regularized Estimators

Many of the estimation methods found in this thesis use regularization to incorporate a priori information. Typically, these methods require the selection of a regularization parameter value that controls the influence of the regularization term on the estimate. Analyzing the spatial resolution properties of such estimators can assist with the selection of this parameter [3, 9, 35].

The spatial resolution of an estimator is described using its local impulse response, which is the effect that a perturbation of a single pixel of the true, noiseless object $\mathbf{x} \in \mathbb{C}^N$ has on the estimate $\hat{\mathbf{x}}$ [3, 35]:

$$\bar{\mathbf{l}}^{(j)}(\mathbf{y};\mathbf{x}) \triangleq \lim_{\epsilon \to 0} \frac{\hat{\mathbf{x}}(\mathbf{y} + [\bar{\mathbf{y}}(\mathbf{x} + \epsilon\, \mathbf{e}_j) - \bar{\mathbf{y}}(\mathbf{x})]) - \hat{\mathbf{x}}(\mathbf{y})}{\epsilon} = \nabla\hat{\mathbf{x}}(\mathbf{y})\nabla\bar{\mathbf{y}}(\mathbf{x})\mathbf{e}_j, \tag{2.39}$$

where $\mathbf{y} \in \mathbb{C}^M$ is the measurement vector, $\bar{\mathbf{y}}$ is the measurement vector from an ideal system (i.e., no noise), and $\mathbf{e}_j \in \mathbb{R}^N$ is the standard basis vector corresponding to the pixel being investigated (in this case, the jth).

For implied estimators based on the general cost function

$$\hat{\mathbf{x}}(\mathbf{y}) = \arg\min_{\mathbf{x}} \Psi(\mathbf{x},\mathbf{y}), \tag{2.40}$$

the estimator gradient is

$$\nabla\hat{\mathbf{x}}(\mathbf{y}) = \left[\nabla^{[2,0]}\Psi(\hat{\mathbf{x}}(\mathbf{y}),\mathbf{y})\right]^{-1}\left[-\nabla^{[1,1]}\Psi(\hat{\mathbf{x}}(\mathbf{y}),\mathbf{y})\right] \tag{2.41}$$

where $\nabla^{[2,0]}\Psi(\mathbf{x},\mathbf{y}) \in \mathbb{C}^{N\times N}$ with $\left[\nabla^{[2,0]}\Psi(\mathbf{x},\mathbf{y})\right]_{(j,k)} = \frac{\partial^2}{\partial x_j \partial x_k}\Psi(\mathbf{x},\mathbf{y})$ and $\nabla^{[1,1]}\Psi(\mathbf{x},\mathbf{y}) \in \mathbb{C}^{N\times M}$ with $\left[\nabla^{[1,1]}\Psi(\mathbf{x},\mathbf{y})\right]_{(j,k)} = \frac{\partial^2}{\partial x_j \partial y_k}\Psi(\mathbf{x},\mathbf{y})$ [3, 35]. This assumes that $\Psi(\mathbf{x},\mathbf{y})$ has a unique minimizer for each $\mathbf{y}$, that the required derivatives exist, and that the Hessian matrix $\nabla^{[2,0]}\Psi(\mathbf{x},\mathbf{y})$ is invertible [3]. The resulting local impulse response is given by

$$\bar{\mathbf{l}}^{(j)}(\mathbf{y};\mathbf{x}) = \left[\nabla^{[2,0]}\Psi(\hat{\mathbf{x}}(\mathbf{y}),\mathbf{y})\right]^{-1}\left[-\nabla^{[1,1]}\Psi(\hat{\mathbf{x}}(\mathbf{y}),\mathbf{y})\right]\nabla\bar{\mathbf{y}}(\mathbf{x})\mathbf{e}_j. \tag{2.42}$$

Most of the regularized estimation methods found in this thesis are penalized-likelihood estimators which have the form

$$\Phi(\mathbf{x}, \mathbf{y}) = \mathrm{L}(\mathbf{x}, \mathbf{y}) + \mathrm{R}(\mathbf{x}), \tag{2.43}$$

where L is the negative log-likelihood and R is the regularizer [3]. For these methods, the impulse response of (2.39) is approximately equal to

$$\bar{\mathbf{l}}^{(j)}(\mathbf{y}; \mathbf{x}) \approx \left[\nabla^{[2,0]}\mathrm{L}\left(\hat{\mathbf{x}}(\bar{\mathbf{y}}), \bar{\mathbf{y}}(\mathbf{x})\right) + \nabla^2\mathrm{R}\left(\hat{\mathbf{x}}(\bar{\mathbf{y}})\right)\right]^{-1} \nabla^{[2,0]}\mathrm{L}\left(\hat{\mathbf{x}}(\bar{\mathbf{y}}), \bar{\mathbf{y}}(\mathbf{x})\right)\mathbf{e}_j, \tag{2.44}$$

where ∇^2R is the Hessian of R [3].[10]

These resolution analysis techniques are used to relate the field map estimators in Chapters IV and V resulting in an effective water-fat imaging initialization method as outlined in Chapter V.

---

[10] Assuming the usual case that the negative log-likelihood with noiseless data $\bar{\mathbf{y}}(\mathbf{x})$ is minimized by the truth $\mathbf{x}$.

Bibliography

[1] D. G. Nishimura, "Principles of magnetic resonance imaging," 1996, Unpublished textbook.

[2] A. K. Funai, *Regularized estimation of main and RF field inhomogeneity and longitudinal relaxation rate in magnetic resonance imaging*, Ph.D. thesis, Univ. of Michigan, Ann Arbor, MI, 48109-2122, Ann Arbor, MI, 2011.

[3] J. A. Fessler, "Image reconstruction: Algorithms and analysis," Book in preparation, 2006.

[4] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: Physical principles and sequence design*, Wiley, New York, 1999.

[5] J. A. Fessler, "BME 516: Lecture Notes (PDF)," Retrieved from http://web.eecs.umich.edu/∼fessler/course/516/l/, 2014.

[6] D. C. Noll, "BME 516: Lecture Notes (PDF)," Retrieved from http://web.eecs.umich.edu/∼dnoll/BME516, 2014.

[7] J. Sijbers, A. Den Dekker, J. Van Audekerke, M. Verhoye, and D. Van Dyck, "Estimation of the noise in magnitude MR images," *Magnetic resonance imaging*, vol. 16, no. 1, pp. 87–90, 1998.

[8] B. P. Sutton, D. C. Noll, and J. A. Fessler, "Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities," *IEEE Trans. Med. Imag.*, vol. 22, no. 2, pp. 178–88, Feb. 2003.

[9] A. K. Funai, J. A. Fessler, D. T. B. Yeo, V. T. Olafsson, and D. C. Noll, "Regularized field map estimation in MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1484–94, Oct. 2008.

[10] K. Sekihara, S. Matsui, and H. Kohno, "NMR imaging for magnets with large nonuniformities," *IEEE Trans. Med. Imag.*, vol. 4, no. 4, pp. 193–9, Dec. 1985.

[11] P. Jezzard and R. S. Balaban, "Correction for geometric distortion in echo planar images from B_0 field variations," *Mag. Res. Med.*, vol. 34, no. 1, pp. 65–73, July 1995.

[12] T. A. Bley, O. Wieben, C. J. François, J. H. Brittain, and S. B. Reeder, "Fat and water magnetic resonance imaging," *J. Mag. Res. Im.*, vol. 31, no. 1, pp. 4–18, Jan. 2010.

[13] 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," *Mag. Res. Med.*, vol. 60, no. 5, pp. 1122–34, Nov. 2008.

[14] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig, "Compressed sensing for chemical shift-based water-fat separation," *Mag. Res. Med.*, vol. 64, no. 6, pp. 1749–59, Dec. 2010.

[15] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Mag. Res. Med.*, vol. 42, no. 5, pp. 952–62, Nov. 1999.

[16] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Mag. Res. Med.*, vol. 58, no. 6, pp. 1182–95, Dec. 2007.

[17] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Mag. Res. Med.*, vol. 47, no. 6, pp. 1202–10, June 2002.

[18] D. L. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[19] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Info. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[20] D. Liang, B. Liu, J. Wang, and L. Ying, "Accelerating SENSE using compressed sensing," *Mag. Res. Med.*, vol. 62, no. 6, pp. 1574–84, Dec. 2009.

[21] S. Ramani and J. A. Fessler, "Parallel MR image reconstruction using augmented Lagrangian methods," *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 694–706, Mar. 2011.

[22] R. Bammer, M. Auer, S. L. Keeling, M. Augustin, L. A. Stables, R. W. Prokesch, R. Stollberger, M. E. Moseley, and F. Fazekas, "Diffusion tensor imaging using single-shot SENSE-EPI," *Mag. Res. Med.*, vol. 48, no. 1, pp. 128–36, July 2002.

[23] S. O. Schoenberg, O. Dietrich, and M. F. Reiser, *Parallel Imaging in Clinical MR Applications*, Springer, Berlin, 2007.

[24] R. G. Baraniuk, "Compressive sensing," *IEEE Sig. Proc. Mag.*, vol. 24, no. 4, pp. 118–21, 2007.

[25] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.

[26] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *American Statistician*, vol. 58, no. 1, pp. 30–7, Feb. 2004.

[27] J. Nocedal and S. Wright, *Numerical Optimization*, Springer, 2 edition, 2006.

[28] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[29] D. P. Bertsekas, “Multiplier methods: A survey,” *Automatica*, vol. 12, no. 2, pp. 133–45, Mar. 1976.

[30] J. Eckstein and D. P. Bertsekas, “On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators,” *Mathematical Programming*, vol. 55, no. 1-3, pp. 293–318, Apr. 1992.

[31] D. E. Petersen, H. H. B. Sørensen, P. C. Hansen, S. Skelboe, and K. Stokbro, “Block tridiagonal matrix inversion and fast transmission calculations,” *J. Comput. Phys.*, vol. 227, no. 6, pp. 3174–3190, Mar. 2008.

[32] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, “Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate,” *ACM Trans. Math. Softw.*, vol. 35, no. 3, pp. 22:1–22:14, Oct. 2008.

[33] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical recipes in C*, Cambridge Univ. Press, New York, 2 edition, 1992.

[34] L. N. Trefethen and D. Bau, *Numerical linear algebra*, Soc. Indust. Appl. Math., Phildelphia, 1997.

[35] J. A. Fessler and W. L. Rogers, “Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs,” *IEEE Trans. Im. Proc.*, vol. 5, no. 9, pp. 1346–58, Sept. 1996.

CHAPTER III

Receive Coil Sensitivity Estimation

In this chapter, we develop and analyze algorithms for regularized sensitivity profile estimation. In Section 3.1, we present a regularized sensitivity profile estimator as well as an augmented Lagrangian based algorithm that can significantly accelerate the estimation process.[1] Section 3.2 presents additional analysis of the regularized sensitivity profile estimator as well as alternative AL based minimization algorithms.[2]

3.1 Accelerated Computation of Regularized Sensitivity Profile Estimates

3.1.1 Introduction

Accurate radio-frequency coil sensitivity profiles are required in many parallel imaging applications (e.g., sensitivity encoding (SENSE) [2], simultaneous acquisition of spatial harmonics (SMASH) [3], and *k*-*t* SENSE [4]). Due to coil deformation during patient setup and dielectric coupling, these profiles must be determined at the time of acquisition [5]. One common approach is to perform a calibration scan prior to the parallel imaging acquisition in which images from a large body coil and

---

[1] This section is based on the published paper [1].

[2] This section is based on the supplementary material that accompanies [1].

multiple surface coils are acquired and reconstructed. Since the body coil has near uniform sensitivity, its image can be used in conjunction with a surface coil image to estimate the surface coil sensitivity profile.[3]

The most straightforward method to estimate the coil sensitivity is to compute the ratio of the surface coil image voxel values (z_i) to the body coil image voxel values (y_i), z_i/y_i. However, ratio estimates can be corrupted by measurement noise, particularly in low signal regions. Furthermore, such estimates can have sharp discontinuities at object edges, contrary to the smooth nature of true coil sensitivity profiles [7]. It is also desirable to have reasonable sensitivity estimates in any low signal regions surrounding the object to avoid reconstruction artifacts that could arise due to patient motion [8]. The ratio estimator, however, does not extrapolate; thus, improved estimation methods can be beneficial.

One approach to generate smooth sensitivity estimates is to measure only the center of k-space [8]. Although simple, this approach does not accurately estimate sensitivities near object edges and can introduce Gibbs ringing artifacts. Filtering procedures have also been proposed including polynomial fitting [2, 9–11], wavelet denoising [12], and using thin-plate splines [13]. These methods do not completely eliminate the Gibbs ringing, while selecting a particular basis function is complicated by the varying size of low signal regions within the images [7,14]. Furthermore, many of these methods disregard the non-stationary variance of the noise in the ratio estimates. In contrast, regularized estimation methods [7, 15, 16] provide smooth sensitivity estimates and are capable of extrapolation without explicit basis function selection or filtering. These methods, however, can be computationally expensive for large problems [7] and this cost is compounded by the large number of coils

---

[3] This work differs from correction methods such as [6] in that we are estimating the receive coil sensitivity profile for use in subsequent parallel imaging reconstruction methods.

in some arrays [17]. Although sensitivity estimation can be performed off-line, the computational costs of regularized methods can increase the overall compute times of parallel imaging.

In this chapter, we take a regularized approach and pose sensitivity estimation as the minimization of a quadratic cost function like in [7]. The large matrices in the cost function prevent one from computing a simple, non-iterative solution. Instead, iterative methods must be used for large data sets; however, traditional methods like conjugate gradient (CG) converge slowly for this problem [7, 18]. Augmented Lagrangian (AL) based minimization techniques [19], and the related Bregman iterations method [20], have been used to accelerate convergence in imaging problems such as denoising [21] and reconstruction [21–29]. Those papers primarily focus on problems that contain non-differentiable regularization terms such as those based on the ℓ_1-norm. However, the underlying theory applies to a wide variety of optimization problems, including the quadratic problem considered here. We therefore propose a new AL based method for estimating sensitivity profiles. To derive this method, we introduce a reformulation of the finite differencing matrix and a subsequent variable splitting that lead to an algorithm with exact alternating minimization steps. This algorithm is equivalent to an alternating direction method of multipliers (ADMM) [30] formulation, which provides a guarantee of convergence. We also explore a variation of this algorithm that updates the Lagrange multipliers between alternating minimization steps. Such variations have been found to improve the convergence rates of other AL based algorithms [31].

Section 3.1.2 presents the derivation of our ADMM algorithm and its intermediate updating variant. Section 3.1.3 compares the convergence speeds of these algorithms with those of CG based methods by performing experiments on both simulated and real data. Section 3.1.4 discusses the results of these experiments and additional properties of the algorithms. Section 3.1.5 concludes by discussing other problems that have quadratic cost functions where our methods may provide an improvement over the traditional techniques.

3.1.2 Materials and Methods

This section introduces our proposed methods for MR coil sensitivity estimation. We begin by posing the estimator as an optimization problem. We then outline the general approach used to solve this problem and present our specific algorithm, with variations, in detail.

Cost function formulation

Regularized methods for MR coil sensitivity estimation are both robust to noise and effective at extrapolating the estimate in regions of low signal [7, 16]. These methods avoid computing the quotient (z_i/y_i) by finding the minimizer of a cost function containing a data-fidelity term and a regularization term that promotes smoothness in the estimate. Similar to [7, 15], we estimate the sensitivity profile by minimizing a weighted sum of quadratic terms:

$$\hat{\mathbf{s}} \triangleq \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{z} - \mathbf{D}\mathbf{s}\|_{\mathbf{W}}^2 + \frac{\lambda}{2}\|\mathbf{R}\mathbf{s}\|_2^2, \tag{3.1}$$

where $\mathbf{s} = [s_1, \ldots, s_N]^{\mathrm{T}}$ with $s_i \in \mathbb{C}$ denoting the desired coil sensitivity at the ith voxel and N denoting the number of voxels,[4] $\mathbf{z} = [z_1, \ldots, z_N]^{\mathrm{T}}$ with $z_i \in \mathbb{C}$ denoting the surface coil image value at the ith voxel, $\mathbf{D} = \operatorname{diag}\{y_i\}$ is a diagonal matrix containing the body coil image voxel values ($y_i \in \mathbb{C}$), $\mathbf{R} \in \mathbb{R}^{M \times N}$ is a finite differencing matrix for the case of non-periodic boundary conditions with M sets of finite

---

[4] This chapter uses s to refer to coil sensitivities. This is the same variable as c from Section 2.4.1 but is named differently to match the corresponding publication.

differences, and $\lambda > 0$ is a regularization coefficient. Additionally, $\mathbf{W} = \operatorname{diag}\{w_i\}$ is a diagonal weighting matrix (with $w_i \in [0, 1]$) that allows us to ensure that the estimate is based primarily on voxels that provide meaningful sensitivity information. Note that a finite differencing matrix with non-periodic boundary conditions is necessary as periodic boundary conditions introduce errors at the edges of the image that can propagate and corrupt the estimate near the object voxels.[5]

Equation (3.1) has a quadratic cost function and therefore has the closed-form solution $\hat{\mathbf{s}} = [\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \lambda\mathbf{R}^{\mathrm{H}}\mathbf{R}]^{-1}\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z}$ where $\mathbf{X}^{\mathrm{H}}$ denotes the Hermitian transpose of $\mathbf{X}$; however, computing this solution is impractical due to the size and complexity of $\mathbf{R}$. Memory constraints further restrict the use of other direct methods, such as Cholesky factorization, for large problems like 3-D data sets. Furthermore, standard iterative solution methods, such as CG, exhibit slow converge for this problem even when using carefully selected preconditioners [32]. To address this, we propose an augmented Lagrangian method to minimize the cost function, the development of which consists of three stages [22]. First, we use variable splitting [27, 33] to convert the unconstrained optimization problem into an equivalent constrained problem, thereby decoupling the effects of the matrices in (3.1). Second, we introduce vector Lagrange multipliers and express the constrained problem in an AL framework. Third, we solve the resulting AL problem using an alternating minimization scheme.

ADMM–Circ: ADMM sensitivity estimation algorithm with circulant substeps

Directly applying variable splitting to (3.1) results in an AL algorithm requiring an approximate solution for one of the alternating minimization steps due to the complexity of the finite differencing matrix $\mathbf{R}$ [18]. Section 3.2.2 presents one such

---

[5] See Section 3.2.6 for additional details.

$$\underbrace{\begin{bmatrix} 0 & 0 & 0 & 0 & \dots & 0 \\ -1 & 2 & -1 & 0 & & 0 \\ 0 & -1 & 2 & -1 & & 0 \\ \vdots & & & & \ddots & \\ 0 & 0 & 0 & 0 & & 0 \end{bmatrix}}_{\mathbf{R}} = \underbrace{\begin{bmatrix} 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & 0 & & 0 \\ 0 & 0 & 1 & 0 & & 0 \\ \vdots & & & & \ddots & \\ 0 & 0 & 0 & 0 & & 0 \end{bmatrix}}_{\mathbf{B}} \underbrace{\begin{bmatrix} 2 & -1 & 0 & 0 & \dots & -1 \\ -1 & 2 & -1 & 0 & & 0 \\ 0 & -1 & 2 & -1 & & 0 \\ \vdots & & & & \ddots & \\ -1 & 0 & 0 & 0 & & 2 \end{bmatrix}}_{\mathbf{C}}$$

Figure 3.1: The matrices **R**, **B**, and **C** for the case of 1D second-order finite differences. The top and bottom rows of **C** compute the difference between the first and last pixels, hence the need for the mask **B**.

algorithm, ADMM–CG. We can avoid this complication if we focus on traditional finite differencing matrices (those with spatially invariant stencils). For such regularizers, we can express the finite differencing matrix as $\mathbf{R} = \mathbf{B}\mathbf{C}$ where $\mathbf{C} \in \mathbb{R}^{M \times N}$ is a typical finite differencing matrix for the case of periodic boundary conditions, containing additional non-zero rows that penalize the differences between voxels on opposing boundaries of the image, and $\mathbf{B} \in \{0, 1\}^{M \times M}$ is a diagonal matrix that contains a mask to eliminate the effects of the added rows. The additional non-zero rows in **C** ensure that $\mathbf{C}^{\mathrm{H}}\mathbf{C}$ is block circulant with circulant blocks unlike **R**. Fig. 3.1 illustrates these matrices for the case of 1D second-order finite differences. We then write the estimation problem in (3.1) as

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{z} - \mathbf{D}\mathbf{s}\|_{\mathbf{W}}^2 + \frac{\lambda}{2}\|\mathbf{B}\mathbf{C}\mathbf{s}\|_2^2. \tag{3.2}$$

We introduce two splitting variables, $\mathbf{u}_0 \in \mathbb{C}^M$ and $\mathbf{u}_1 \in \mathbb{C}^N$, to this new formulation to decouple the matrices **D**, **B**, and **C**. The resulting equivalent constrained optimization problem is

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s},\mathbf{u}_0,\mathbf{u}_1} \frac{1}{2}\|\mathbf{z} - \mathbf{D}\mathbf{u}_1\|_{\mathbf{W}}^2 + \frac{\lambda}{2}\|\mathbf{B}\mathbf{u}_0\|_2^2 \ \text{ s.t. } \mathbf{u}_1 = \mathbf{s} \text{ and } \mathbf{u}_0 = \mathbf{C}\mathbf{s}. \tag{3.3}$$

Solving this constrained optimization problem is exactly equivalent to solving the unconstrained problem (3.1).

We express (3.3) in the more concise notation:

$$\arg\min_{\mathbf{s},\mathbf{u}} \frac{1}{2}\|\mathbf{h} - \mathbf{A}\mathbf{u}\|_2^2 \text{ s.t. } \mathbf{u} = \mathbf{T}\mathbf{s}, \tag{3.4}$$

where

$$\mathbf{u} \triangleq \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_0 \end{bmatrix}, \mathbf{T} \triangleq \begin{bmatrix} \mathbf{I} \\ \mathbf{C} \end{bmatrix}, \mathbf{h} \triangleq \begin{bmatrix} \mathbf{W}^{\frac{1}{2}}\mathbf{z} \\ \mathbf{0} \end{bmatrix}, \mathbf{A} \triangleq \begin{bmatrix} \mathbf{W}^{\frac{1}{2}}\mathbf{D} & \mathbf{0} \\ \mathbf{0} & \sqrt{\lambda}\mathbf{B} \end{bmatrix},$$

and $\mathbf{W}^{\frac{1}{2}} \triangleq \operatorname{diag}\{\sqrt{w_i}\}$.

We then introduce two vectors of Lagrange multipliers, $\boldsymbol{\eta}_0 \in \mathbb{C}^M$ and $\boldsymbol{\eta}_1 \in \mathbb{C}^N$, and express (3.4) as an AL problem. We use the general AL formulation outlined in [22] that incorporates the Lagrange multiplier into the quadratic penalty term. This formulation is a natural extension of the traditional AL to the case of complex values and it simplifies the derivation of the subsequent alternating minimization steps. The resulting AL function-based minimization problem is

$$\arg\min_{\mathbf{s},\mathbf{u}} \frac{1}{2}\|\mathbf{h} - \mathbf{A}\mathbf{u}\|_2^2 + \frac{1}{2}\|\mathbf{u} - \mathbf{T}\mathbf{s} - \boldsymbol{\eta}\|_{\mathbf{V}}^2, \tag{3.5}$$

where

$$\boldsymbol{\eta} \triangleq \begin{bmatrix} \boldsymbol{\eta}_1 \\ \boldsymbol{\eta}_0 \end{bmatrix}, \mathbf{V} \triangleq \begin{bmatrix} \nu_1\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \nu_0\mathbf{I} \end{bmatrix},$$

and $\nu_0, \nu_1 > 0$ are AL penalty parameters that influence the convergence rate of the algorithm but do not affect the final estimate [28].

Traditional AL methods would require jointly minimizing (3.5) over the vectors $\mathbf{s}$ and $\mathbf{u}$; however, such an approach is computationally expensive for typical image sizes. Instead, we use a block Gauss–Seidel type alternating minimization strategy that has been effective in solving other AL problems [21, 30] in which we alternate between minimizing (3.5) independently with respect to $\mathbf{s}$ and $\mathbf{u}$ as follows:

$$\mathbf{s}^{(j+1)} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{u}^{(j)} - \mathbf{T}\mathbf{s} - \boldsymbol{\eta}^{(j)}\|_{\mathbf{V}}^2, \tag{3.6}$$

$$\mathbf{u}^{(j+1)} = \arg\min_{\mathbf{u}} \frac{1}{2}\|\mathbf{h} - \mathbf{A}\mathbf{u}\|_2^2 + \frac{1}{2}\|\mathbf{u} - \mathbf{T}\mathbf{s}^{(j+1)} - \boldsymbol{\eta}^{(j)}\|_{\mathbf{V}}^2. \tag{3.7}$$

Update (3.7) has a simple closed-form solution:

$$\mathbf{u}^{(j+1)} = \left(\mathbf{A}^{\mathrm{H}}\mathbf{A} + \mathbf{V}\right)^{-1}\left[\mathbf{A}^{\mathrm{H}}\mathbf{h} + \mathbf{V}(\mathbf{T}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}^{(j)})\right]. \tag{3.8}$$

In fact, the block diagonal structures of $\mathbf{A}$ and $\mathbf{V}$ decouple the update of $\mathbf{u}$ into two parallel updates in terms of $\mathbf{u}_1$ and $\mathbf{u}_0$:

$$\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}\left[\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z} + \nu_1(\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_1^{(j)})\right], \tag{3.9}$$

$$\mathbf{u}_0^{(j+1)} = \mathbf{B}_2^{-1}(\mathbf{C}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_0^{(j)}), \tag{3.10}$$

where $\mathbf{B}_2 \triangleq \frac{\lambda}{\nu_0}\mathbf{B}^{\mathrm{H}}\mathbf{B} + \mathbf{I}$ and $\mathbf{D}_2 \triangleq \mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \nu_1\mathbf{I}$ are both diagonal matrices that are trivial to invert. The closed-form update for $\mathbf{s}$ may at first appear more complicated to compute:

$$\begin{aligned}\mathbf{s}^{(j+1)} &= \left(\mathbf{T}^{\mathrm{H}}\mathbf{V}\mathbf{T}\right)^{-1}\mathbf{T}^{\mathrm{H}}\mathbf{V}\left(\mathbf{u}^{(j)} - \boldsymbol{\eta}^{(j)}\right)\\ &= \left(\nu_1\mathbf{I} + \nu_0\mathbf{C}^{\mathrm{H}}\mathbf{C}\right)^{-1}\left[\nu_0\mathbf{C}^{\mathrm{H}}(\mathbf{u}_0^{(j)} - \boldsymbol{\eta}_0^{(j)}) + \nu_1(\mathbf{u}_1^{(j)} - \boldsymbol{\eta}_1^{(j)})\right].\end{aligned} \tag{3.11}$$

However, since $\mathbf{C}^{\mathrm{H}}\mathbf{C}$ is block circulant with circulant blocks, $\mathbf{C}^{\mathrm{H}}\mathbf{C} = \mathbf{Q}^{\mathrm{H}}\boldsymbol{\Phi}\mathbf{Q}$ where $\mathbf{Q}$ is a (multidimensional) discrete Fourier transform (DFT) matrix and $\boldsymbol{\Phi}$ is a diagonal matrix containing the spectrum of the convolution kernel of $\mathbf{C}^{\mathrm{H}}\mathbf{C}$. Substituting this decomposition into (3.11) yields:

$$\mathbf{s}^{(j+1)} = \mathbf{Q}^{\mathrm{H}}\boldsymbol{\Phi}_2^{-1}\mathbf{Q}\left[\nu_0\mathbf{C}^{\mathrm{H}}(\mathbf{u}_0^{(j)} - \boldsymbol{\eta}_0^{(j)}) + \nu_1(\mathbf{u}_1^{(j)} - \boldsymbol{\eta}_1^{(j)})\right], \tag{3.12}$$

where $\boldsymbol{\Phi}_2 \triangleq \nu_1\mathbf{I} + \nu_0\boldsymbol{\Phi}$. This formulation is simpler to compute since $\boldsymbol{\Phi}_2$ is a diagonal matrix and we implement $\mathbf{Q}$ efficiently using fast Fourier transforms (FFTs).

```
ADMM–Circ
Initialize: $\mathbf{u}_1^{(0)} = \mathbf{s}^{(0)}$, $\mathbf{u}_0^{(0)} = \mathbf{C}\mathbf{s}^{(0)}$, $\boldsymbol{\eta}_0^{(0)} = \mathbf{0}$,
$\boldsymbol{\eta}_1^{(0)} = \mathbf{0}$ and $j = 0$.
Set $\mathbf{D}_2^{-1} = \left[\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \nu_1\mathbf{I}\right]^{-1}$ and $\mathbf{z}_2 = \mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z}$.
Set $\mathbf{B}_2^{-1} = \left[\frac{\lambda}{\nu_0}\mathbf{B}^{\mathrm{H}}\mathbf{B} + \mathbf{I}\right]^{-1}$.
Set $\boldsymbol{\Phi}_2^{-1} = [\nu_1\mathbf{I} + \nu_0\boldsymbol{\Phi}]^{-1}$.
Repeat until stop criterion is achieved:
    1. $\mathbf{s}^{(j+1)}$ from (3.12) exactly,
    2. $\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}\left[\mathbf{z}_2 + \nu_1(\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_1^{(j)})\right]$,
       $\mathbf{u}_0^{(j+1)} = \mathbf{B}_2^{-1}(\mathbf{C}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_0^{(j)})$,
    3. $\boldsymbol{\eta}_1^{(j+1)} = \boldsymbol{\eta}_1^{(j)} - (\mathbf{u}_1^{(j+1)} - \mathbf{s}^{(j+1)})$,
       $\boldsymbol{\eta}_0^{(j+1)} = \boldsymbol{\eta}_0^{(j)} - (\mathbf{u}_0^{(j+1)} - \mathbf{C}\mathbf{s}^{(j+1)})$,
    4. $j = j + 1$.
```

Figure 3.2: Overview of the ADMM–Circ algorithm. Note that $\mathbf{C}\mathbf{s}^{(j+1)}$ only needs to be computed once per iteration.

Fig. 3.2 summarizes the resulting sensitivity profile estimation algorithm, ADMM–Circ. Each stage of the proposed algorithm consists of an exact, non-iterative update. Furthermore, it can be shown that the steps in this formulation are identical to those of an ADMM algorithm applied to the real valued case where we treat the complex valued terms as a stack of their real and imaginary components. As discussed in Section 3.1.4, this equivalence allows us to conclude that the ADMM–Circ algorithm converges to the solution of (3.1). In contrast, the parallel imaging reconstruction algorithm in [22] is an AL method that lacks a convergence proof due to the type of splitting used.

Alternating minimization with intermediate updating

Updating the Lagrange multipliers $\boldsymbol{\eta}$ between each alternating minimization step has been shown to increase the convergence rates of several AL based algorithms [31]. We also explore this variation in our proposed algorithm by updating the relevant

```
ADMM–Circ–IU
Initialize: $\mathbf{u}_1^{(0)} = \mathbf{s}^{(0)}$, $\mathbf{u}_0^{(0)} = \mathbf{C}\mathbf{s}^{(0)}$, $\boldsymbol{\eta}_0^{(0)} = \mathbf{0}$, $\boldsymbol{\eta}_1^{(0)} = \mathbf{0}$ and $j = 0$.
Set $\mathbf{D}_2^{-1} = \left[\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \nu_1\mathbf{I}\right]^{-1}$ and $\mathbf{z}_2 = \mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z}$.
Set $\mathbf{B}_2^{-1} = \left[\frac{\lambda}{\nu_0}\mathbf{B}^{\mathrm{H}}\mathbf{B} + \mathbf{I}\right]^{-1}$.
Set $\boldsymbol{\Phi}_2^{-1} = [\nu_1\mathbf{I} + \nu_0\boldsymbol{\Phi}]^{-1}$.
Repeat until stop criterion is achieved:
    1. $\mathbf{s}^{(j+1)}$ from (3.12) exactly,
    2. $\boldsymbol{\eta}_1^{(j+1/2)} = \boldsymbol{\eta}_1^{(j)} - (\mathbf{u}_1^{(j)} - \mathbf{s}^{(j+1)})$,
       $\boldsymbol{\eta}_0^{(j+1/2)} = \boldsymbol{\eta}_0^{(j)} - (\mathbf{u}_0^{(j)} - \mathbf{C}\mathbf{s}^{(j+1)})$,
    3. $\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}\left[\mathbf{z}_2 + \nu_1(\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_1^{(j+1/2)})\right]$,
       $\mathbf{u}_0^{(j+1)} = \mathbf{B}_2^{-1}(\mathbf{C}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_0^{(j+1/2)})$,
    4. $\boldsymbol{\eta}_1^{(j+1)} = \boldsymbol{\eta}_1^{(j+1/2)} - (\mathbf{u}_1^{(j+1)} - \mathbf{s}^{(j+1)})$,
       $\boldsymbol{\eta}_0^{(j+1)} = \boldsymbol{\eta}_0^{(j+1/2)} - (\mathbf{u}_0^{(j+1)} - \mathbf{C}\mathbf{s}^{(j+1)})$,
    5. $j = j + 1$.
```

Figure 3.3: The ADMM–Circ algorithm with intermediate Lagrange multiplier updating (ADMM–Circ–IU). Note that $\mathbf{C}\mathbf{s}^{(j+1)}$ only needs to be computed once per iteration.

Lagrange multipliers after each alternating minimization step, Fig. 3.3. The resulting algorithm, ADMM–Circ–IU, requires no additional variables and the added updates (Step 2) are computationally inexpensive. Section 3.1.4 describes the convergence properties of such adaptations.

Parameter selection

Regularized methods require the selection of a regularization parameter, λ in (3.1), which controls the smoothness of the sensitivity profile. We discuss how λ is selected for typical problems in Section 3.1.3.

In addition, our proposed AL methods require that we specify values for the AL penalty parameters ν_0 and ν_1. Following [22], we determine the parameter values using the condition numbers of the matrices requiring inversion in the alternating minimization steps. For both the ADMM–Circ method and its variation, we consider the matrices $\mathbf{B}_2$, $\mathbf{\Phi}_2$, and $\mathbf{D}_2$. We normalize the coil images before performing the estimate; thus, the condition number of $\mathbf{D}_2$, $\kappa(\mathbf{D}_2)$, is $(1+\nu_1)/(d_{\min}^2+\nu_1)$ where $d_{\min}^2$ is the smallest diagonal element of $\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D}$. Furthermore, effective weighting matrices should have near zero values to remove the effects of noise in the low signal regions of the body coil image. Thus, $d_{\min}^2 \approx 0$ and $\kappa(\mathbf{D}_2)$ does not typically depend on the data. We therefore set our parameters by considering the condition numbers of the other two matrices, $\kappa(\mathbf{B}_2)$ and $\kappa(\mathbf{\Phi}_2)$. Through extensive numerical simulation, we found that setting ν_0 such that $\kappa(\mathbf{B}_2) \in [225, 400]$ and then ν_1 such that $\kappa(\mathbf{\Phi}_2) \in [200, 1000]$ provided good convergence rates for a variety of data sets.

3.1.3 Results

We evaluated our proposed sensitivity estimation methods using two very different data sets. The first experiment used simulated brain data whereas the second used real breast phantom data. Previous publications investigated the accuracy of similar regularized estimators [7]; however, there have been few comparisons with other methods concerning their effects on SENSE reconstruction quality. We therefore included an illustration of the improved SENSE reconstruction quality obtained from using regularized sensitivity estimates over standard techniques in Section 3.2. The focus of this section is on accelerating these algorithms and thus, in this section, we compare the convergence speeds of our AL algorithms with those of conventional CG and PCG with the following circulant preconditioner (PCG–Circ):

$$\mathbf{P}_{\mathrm{C}} = \mathbf{Q}^{\mathrm{H}} \left(\mathbf{I} + \lambda \mathbf{\Omega}\right) \mathbf{Q}, \tag{3.13}$$

where $\mathbf{\Omega}$ is a diagonal matrix containing the spectrum of the convolution kernel of $\mathbf{R}^{\mathrm{H}}\mathbf{R}$ [34].

We initialized each algorithm with a sensitivity profile comprising the standard ratio estimate over the object voxels and the mean magnitude and phase of these values over the non-object voxels. All of the algorithms were implemented in MATLAB (The MathWorks, Natick, MA, USA) and the experiments were run on a PC with a 2.66 GHz, quad-core Intel Xeon CPU.

We compared the convergence properties of the algorithms using the normalized ℓ_2-distance between the current estimate, $\mathbf{s}^{(j)}$, and the minimizer of (3.1), $\hat{\mathbf{s}}$:

$$\mathcal{D}(\mathbf{s}^{(j)}) = \frac{\|\mathbf{s}^{(j)} - \hat{\mathbf{s}}\|_2}{\|\hat{\mathbf{s}}\|_2}. \tag{3.14}$$

We focused on 2-D estimation problems so that we could use Cholesky factorization to determine a non-iterative "exact" solution to (3.1). Using this non-iterative solution for $\hat{\mathbf{s}}$ avoids favoring a specific iterative algorithm.

Cost function setup

In defining the estimation problem (3.1), we chose a second-order finite differencing matrix for $\mathbf{R}$ as it resulted in more accurate sensitivity estimates than both first-order and fourth-order finite differences (results not shown). We used a binary mask, created by thresholding the body coil image, for the weighting matrix $\mathbf{W}$. This ensured that the majority of voxels in the object support were included in the estimate, while limiting the number of noisy, non-object voxels.

We selected the AL penalty parameters ν_0 and ν_1 for both experiments using the same set of condition numbers. In particular, we selected ν_0 and ν_1 such that $\kappa(\mathbf{B}_2) = 255$ and $\kappa(\mathbf{\Phi}_2) = 650$ for both the ADMM–Circ and ADMM–Circ–IU algorithms.

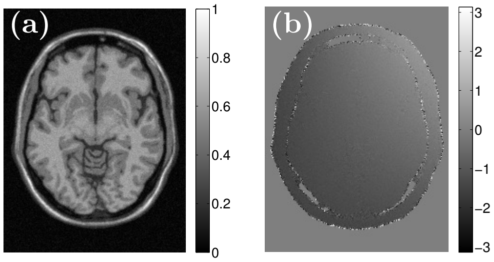

Figure 3.4: The (a) magnitude and (b) phase (masked) of the body coil image for the simulated brain data.

Simulated brain data

Our first experiment used a 256×192 pixel, T1 weighted, transverse plane brain image from the BrainWeb database [35] (1 mm isotroptic in-plane resolution, slice thickness = 1 mm, no noise). To create a more realistic MR image, we added a slowly varying phase component to the brain image. We then added complex random Gaussian noise to create a body coil image, $\mathbf{y}$, with a signal-to-noise ratio[6] (SNR) of 10. Fig. 3.4 presents the magnitude and phase of the resulting body coil image.

We simulated sensitivity profiles for four circular coils placed just outside the field-of-view (FOV) using an analytic method [36]. These sensitivities were then combined with our complex brain image and complex random Gaussian noise to create four surface coil images, $\mathbf{z}$, with SNRs of approximately 10. Fig. 3.5 presents the true sensitivities and their corresponding surface coil images.

We estimated the coil sensitivities using our proposed AL methods and the two CG methods. We set $\lambda = 2^5$ as this value produced accurate estimates (compared to the truth) over both the high intensity voxels and their surrounding regions. We ran 20 000 iterations of each method to ensure convergence. All of the algorithms converged to a normalized ℓ_2-distance of less than -200 dB from, and appeared nearly

---

[6] $\text{SNR} = \mu_o/\sigma_b$ where μ_o is the mean of the magnitudes of the non-zero object pixels in the true image and σ_b is the standard deviation of the background pixels in the noisy image.

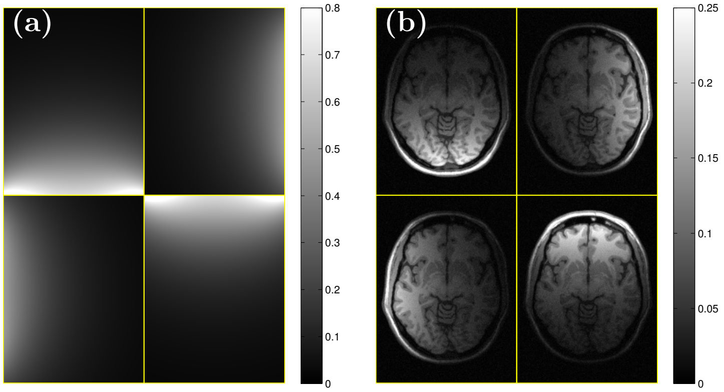

Figure 3.5: The magnitudes of the (a) simulated sensitivity profiles and the (b) simulated surface coil images for the brain data.

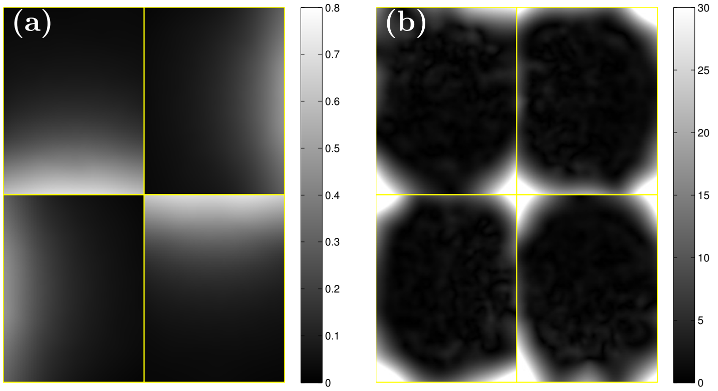

Figure 3.6: The magnitudes of the (a) estimated sensitivity profiles and (b) their percentage difference to the true sensitivities for the simulated brain data.

identical to, the Cholesky based solution $\hat{\mathbf{s}}$. Fig. 3.6 presents the estimated coil sensitivities as well as their percentage difference to the truth. The convergence rates of the algorithms were similar for all four coils so we present the results for one representative coil. Fig. 3.7 plots $\mathcal{D}(\mathbf{s}^{(j)})$ with respect to both iteration and time for the bottom left coil in Fig. 3.5. ADMM–Circ–IU was the fastest algorithm, converging to within $\mathcal{D}(\mathbf{s}^{(j)}) = 0.1\,\%$ in approximately 85 seconds. PCG–Circ was faster than ADMM–Circ with convergence times of nearly 130 and 165 seconds, respectively. Conventional CG took by far the longest time at 535 seconds.

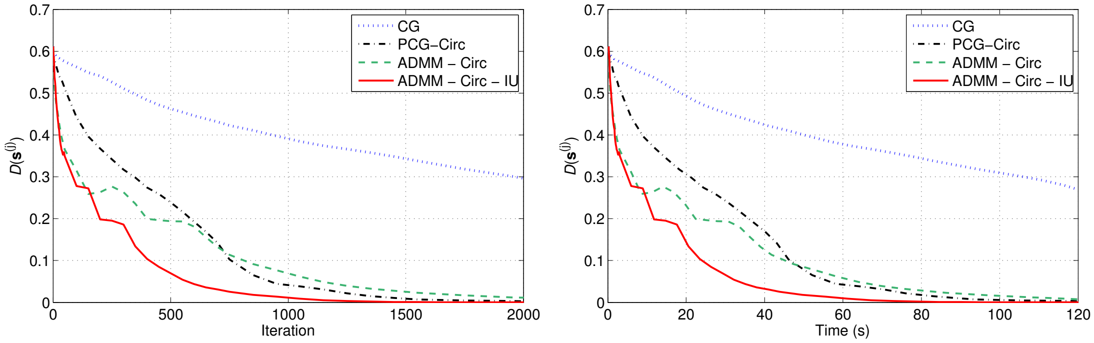

Figure 3.7: Plots of the normalized ℓ_2-distance between $\mathbf{s}^{(j)}$ and $\hat{\mathbf{s}}$, $\mathcal{D}(\mathbf{s}^{(j)})$, with respect to iteration (left) and time (right) for the bottom left brain data surface coil image in Fig. 3.5.

Figure 3.8: The magnitude of the breast phantom body coil image.

Breast phantom data

Our second experiment used a breast phantom consisting of two containers plastered with vegetable shortening and filled with “Super Stuff” bolus material (Radiation Products Design Inc., Albertville, MN, USA). Calibration data was acquired using four surface coils and one body coil on a Philips 3T scanner ($T_R = 4.6$ ms, $T_E = 1.7$ ms, matrix = 384×96). We reconstructed four surface coil images and one body coil image, each 384×96 pixels, using an inverse FFT. Figs. 3.8 and 3.9(a) show the magnitudes of the body coil image and surface coil images, respectively. This data set presents several challenges for sensitivity estimation due to the placement of coils near the center of the FOV and because of large regions of low signal both within and outside the object.

To determine a suitable regularization parameter, λ, we first estimated the coil sensitivities using the CG method for several values of λ. We then performed two-fold accelerated SENSE reconstructions [2] using each set of estimated sensitivities

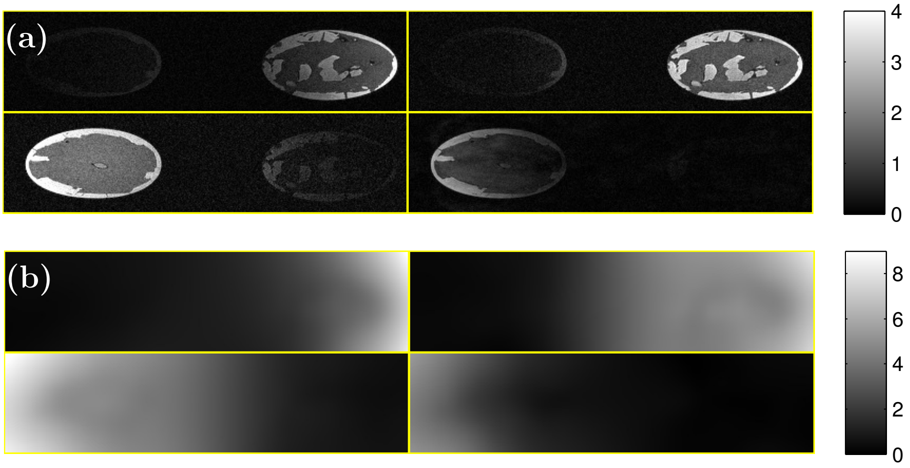

Figure 3.9: The magnitudes of the (a) breast phantom surface coil images and the (b) corresponding estimated sensitivity profiles.

and compared the resulting images to the body coil image (not shown). We selected $\lambda = 2^7$ as its corresponding reconstructed image had minimal artifacts and matched closely to the body coil image.

We estimated the coil sensitivities using our proposed AL methods and the two CG methods. We ran 20 000 iterations of each algorithm to ensure that convergence was achieved. Again, the resulting estimates all converged to a normalized ℓ_2-distance of less than -200 dB from the Cholesky based solution $\hat{\mathbf{s}}$. Fig. 3.9(b) presents the estimated coil sensitivities. The convergence rates of the algorithms were similar for all four coils so we present the results for one representative coil. Fig. 3.10 plots $\mathcal{D}(\mathbf{s}^{(j)})$ with respect to both iteration and time for the bottom left coil in Fig. 3.9. ADMM–Circ–IU was again the fastest algorithm, converging to within $\mathcal{D}(\mathbf{s}^{(j)}) = 0.1\,\%$ in approximately 50 seconds. Unlike in the brain data experiment, ADMM–Circ had a similar convergence rate to PCG–Circ with both algorithms requiring approximately 100 seconds. Conventional CG again took the longest time at 445 seconds.

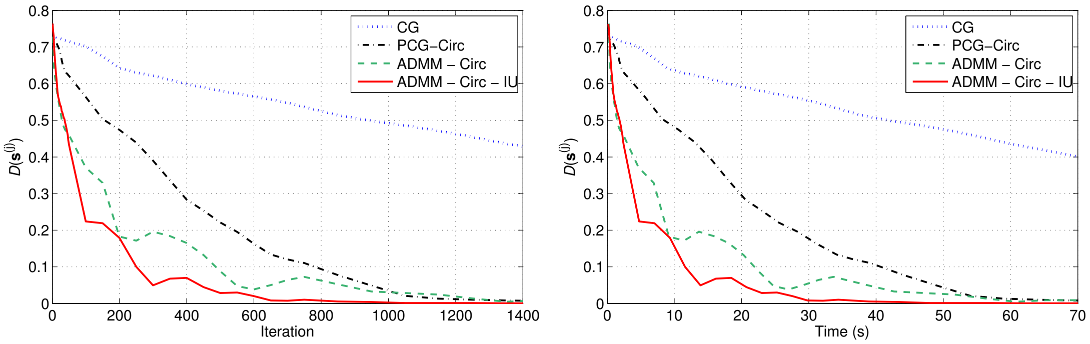

Figure 3.10: Plots of the normalized ℓ_2-distance between $\mathbf{s}^{(j)}$ and $\hat{\mathbf{s}}$, $\mathcal{D}(\mathbf{s}^{(j)})$, with respect to iteration (left) and time (right) for the bottom left breast data surface coil image in Fig. 3.9.

3.1.4 Discussion

The sensitivity estimates generated by minimizing the cost function in (3.1) are smooth like true coil sensitivity profiles. As further discussed in Section 3.2.4, the sensitivity estimates of the brain data are highly accurate over the object and surrounding pixels. The largest errors are at the extreme corners of the image where there is no information about the true sensitivities. The flexibility of the regularized estimation method is highlighted in the breast phantom experiment by its ability to simultaneously estimate the sensitivity within both breasts and smoothly extrapolate over the regions in-between. This is particularly evident for the coils that have near uniform sensitivity over a single breast (the top right and bottom left coils in Fig. 3.9). As illustrated in Section 3.2.4, SENSE reconstructions performed with these sensitivity profiles were artifact free unlike those created using low-pass filter techniques.

ADMM–Circ–IU was the fastest method in all experiments requiring as little as half the time of PCG–Circ and a ninth the time of conventional CG. ADMM–Circ, although much faster than the CG based methods over the first few iterations, had similar convergence times to PCG–Circ in our breast experiment and was slower in our simulated brain experiment. Thus, using intermediate updating significantly accelerated our ADMM algorithm. The CG algorithm remained the slowest method in all experiments. Interestingly, the relative convergence rate of the PCG–Circ algorithm depended on the experimental data. This behavior is partly a result of the varying accuracy of the preconditioner used in the PCG algorithm. Specifically, the circulant preconditioner used an identity matrix in place of the weighted body coil image voxel intensities (i.e., $\mathbf{I} + \lambda \mathbf{R}^{\mathrm{H}}\mathbf{R}$ for $\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \lambda \mathbf{R}^{\mathrm{H}}\mathbf{R}$). This approximation works best for images that have few low signal voxels as is apparent from the decreased performance of the PCG–Circ algorithm on the breast data compared to the simulated brain data which has a higher percentage of voxels with significant signal. In contrast, our proposed ADMM algorithms do not require such approximations and their convergence speeds are therefore more robust to differences in the data.

Table 3.1 presents the approximate number of complex multiplication and addition operations required by an iteration of each algorithm. For typical finite differencing matrices, ADMM–Circ–IU, ADMM–Circ, and PCG–Circ require a similar number of operations, whereas traditional CG requires fewer operations per iteration. The effect of these varying costs per iteration is highlighted by contrasting the convergence rates of each algorithm in terms of iteration and time as seen in Figs. 3.7 and 3.10. As with time, ADMM–Circ–IU needed approximately half as many iterations as PCG–Circ and ADMM–Circ. CG required significantly more iterations to converge than the other algorithms, offsetting any savings in cost per iteration.

The convergence curves for our ADMM methods exhibited non-monotonic behavior with respect to $\mathcal{D}(\mathbf{s}^{(j)})$. We found that the degree of non-monotonicity was influenced by the choice of AL penalty parameters, ν_0 and ν_1. In fact, the parameter settings that provided the fastest convergence rates typically resulted in

Table 3.1: Approximate Number of Complex Arithmetic Operations Per Iteration for the Case of Second-Order Finite Differences

Estimator	**Number of Operations**
ADMM–Circ–IU	$19N + 13M$[a]$+ 2 \cdot \mathrm{O}_{\mathrm{FFT}}$[b]
ADMM–Circ	$17N + 11M + 2 \cdot \mathrm{O}_{\mathrm{FFT}}$
PCG–Circ	$23N + 11M + 2 \cdot \mathrm{O}_{\mathrm{FFT}}$
CG	$22N + 11M$

[a] $M \approx 4N$ for 2-D problems.

[b] $\mathrm{O}_{\mathrm{FFT}}$ denotes the cost of the FFT operations ($O(N \log(N))$).

non-monotonicity in the $\mathcal{D}(\mathbf{s}^{(j)})$ plots.

All of our proposed algorithms converged to the solution of (3.1) in every experiment. As discussed after (3.12), our ADMM–Circ algorithm is equivalent to an ADMM algorithm with exact update steps. We can therefore conclude that this algorithm converges to the solution of (3.1) as per [30, Th. 8]. Our intermediate updating variant, ADMM–Circ–IU, does not have the exact formulation outlined in the hypothesis of [30, Th. 8]. However, a guarantee of convergence exists for similar ADMM variants with symmetric Lagrange multiplier updating [31]. We are currently investigating an extension of this proof to ADMM–Circ–IU.

The convergence rates of our proposed algorithms were robust to the particular choice of condition numbers used to determine the AL penalty parameters ν_0 and ν_1. In fact, we used the same condition numbers for our two very different experiments. Furthermore, our fastest algorithm, ADMM–Circ–IU, remained faster than PCG–Circ for $\kappa(\mathbf{B}_2)$ values nearly two times larger and smaller than the optimal value and for $\kappa(\mathbf{\Phi}_2)$ values three times larger or smaller than optimal. We also explored varying the λ value in our experiments and found that this set of condition numbers consistently worked well. The choice of the best condition numbers does not depend on the surface coil image. Therefore, if one wanted to fine-tune the convergence rate of the algorithms, a single coil of a multi-coil array would suffice.

It is a common practice in medical imaging to restrict estimates and reconstructions to within masked regions to improve both their computation time and quality over the region. If this is done for simple problems like our simulated brain data, which requires minimal interpolation within the object support, then the PCG–Circ algorithm estimating within a masked region will converge faster than our ADMM–Circ methods estimating over the entire FOV. However, this is not the case for more complicated problems like our breast phantom data. In particular, we found that our ADMM–Circ–IU algorithm, estimating over the full FOV, converged to $\mathcal{D}(\mathbf{s}^{(j)}) = 0.1\,\%$ at the same speed or faster than a PCG–Circ algorithm estimating within a masked region consisting of a convex hull[7] surrounding the object support, Fig. 3.11. Furthermore, the quality of the unmasked ADMM–Circ–IU estimates was similar to that of the masked PCG estimates over the masked region. This is partially because the weighting matrix $\mathbf{W}$ minimizes the impact of noisy voxels outside of the object support. A major disadvantage of masking is that the lack of an estimate outside the mask can lead to significant SENSE reconstruction artifacts if the object moves into this region during acquisition [8]. Thus, the mask would have to be carefully selected with this in mind. We therefore followed existing work [7] and focused on algorithms without support masks.

In addition to the algorithms presented in this section and Section 3.2, we also explored AL algorithms that incorporated simpler variable splittings. For instance, we introduced the single splitting variable $\mathbf{u}_0 = \mathbf{R}\mathbf{s}$ to (3.1) and similarly, $\mathbf{u}_0 = \mathbf{C}\mathbf{s}$ to (3.2). The AL formulations used to minimize the resulting cost functions had only two update equations. However, one of these equations required an approximate it-

---

[7] See Section 3.2.5 for an illustration of why a convex hull is required.

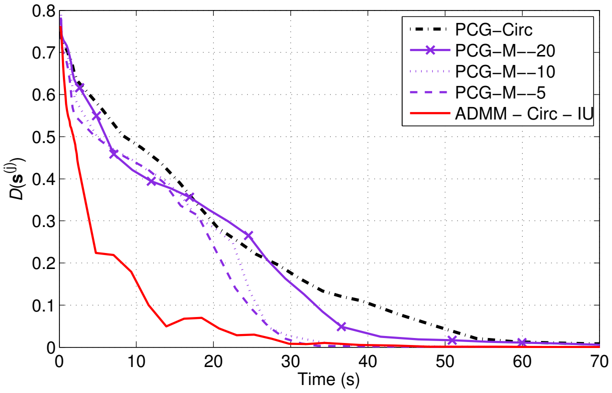

Figure 3.11: Plots of $\mathcal{D}(\mathbf{s}^{(j)})$ with respect to time for ADMM–Circ–IU and PCG–Circ without masks, as well as PCG–Circ using masks with various degrees of dilation (5, 10, 20 pixels), applied to the bottom left breast data surface coil image in Fig. 3.9. For each case, the $\hat{\mathbf{s}}$ used in $\mathcal{D}(\mathbf{s}^{(j)})$ is the regularized solution for the appropriate mask.

erative solution and the resulting AL algorithms were highly sensitive to inaccuracies in the approximation. In fact, when using PCG for the approximate update step, the optimal number of inner PCG iterations was so large that the overall algorithms were slower than regular CG. Curiously, this is the type of splitting that is used in the popular split Bregman approaches [21], although there it is used in cases where $\mathbf{R}^{\mathrm{H}}\mathbf{R}$ is circulant.

If the body coil data $\mathbf{y}$ is not available, one could use the square-root of the sum-of-squares of the surface coil images in its place [10, 37, 38]. Our algorithms would remain the same and only the elements of $\mathbf{D}$ would change. However, it may be more desirable in this situation to perform joint estimation of the final image and the sensitivity profiles (e.g., [10]). Such algorithms are more complicated to compute than (3.1) and might also benefit from an ADMM reformulation.

3.1.5 Conclusions

We developed a new iterative method, ADMM–Circ, using variable splitting and AL strategies that accelerates the regularized estimation of MR coil sensitivities. By separating the finite differencing matrix for the case of non-periodic boundary conditions into a finite differencing matrix for the case of periodic boundary conditions and a diagonal masking matrix, we were able to find a variable splitting strategy that resulted in an algorithm with exact update steps. Additionally, we demonstrated that intermediate updating of the Lagrange multipliers significantly accelerated our proposed AL algorithm. Our fastest method, ADMM–Circ–IU, had convergence speeds up to twice those of the PCG method with a circulant preconditioner.

More generally, we illustrated how AL methods can be used to accelerate convergence for imaging problems with certain classes of quadratic cost functions. There are many areas in MR imaging where similar cost functions are used. For instance, B_0 and B_1^+ map estimation can be performed by minimizing cost functions with quadratic regularization terms over the image domain [39–41]. One specific application is proposed in Chapter V.

3.2 Additional Topics in Regularized Sensitivity Profile Estimation

3.2.1 Introduction

This section contains further analysis of the regularized sensitivity profile estimation method presented in Section 3.1. It begins by exploring several other variable splittings that lead to similar AL estimation algorithms. Specifically, Section 3.2.2 presents a second ADMM algorithm for sensitivity profile estimation, while Section 3.2.3 presents an alternate formulation that leads to an AL (but not ADMM) estimation algorithm with similar performance. We then demonstrate several aspects that must be considered when implementing the regularized estimator in a clinical setting. In particular, Section 3.2.4 illustrates the improved SENSE reconstruction quality resulting from using regularized sensitivity estimates compared to traditional ratio based estimates. Section 3.2.5 demonstrates why a dilated convex hull surrounding the object voxels should be used for an estimation mask. Section 3.2.6 illustrates the importance of using a finite differencing matrix with non-periodic boundary conditions in our estimator (3.1).

3.2.2 ADMM Estimation Algorithm with Conjugate Gradient Substeps

In this section we present and evaluate a second AL algorithm, ADMM–CG, which does not use the reformulation of the finite differencing matrix as discussed in Section 3.1.2. We begin with the derivation of the algorithm which uses the same techniques as ADMM–Circ. We then compare this new algorithm to the methods presented in Section 3.1 using the same data sets and briefly discuss its properties.

Method derivation

We begin our derivation by introducing two new variables, $\mathbf{u}_0 \in \mathbb{C}^M$ and $\mathbf{u}_1 \in \mathbb{C}^N$, to the initial cost function in (3.1). The purpose of these variables is to isolate the finite differencing matrix $\mathbf{R}$ from the diagonal matrix $\mathbf{D}$. The resulting constrained optimization problem is

$$\hat{\mathbf{s}} \triangleq \arg\min_{\mathbf{s},\mathbf{u}_0,\mathbf{u}_1} \frac{1}{2}\|\mathbf{z} - \mathbf{D}\mathbf{u}_1\|_{\mathbf{W}}^2 + \frac{\lambda}{2}\|\mathbf{u}_0\|_2^2 \;\text{ s.t. }\; \mathbf{u}_1 = \mathbf{s} \text{ and } \mathbf{u}_0 = \mathbf{R}\mathbf{s}. \tag{3.15}$$

Solving this constrained optimization problem is exactly equivalent to solving the unconstrained problem (3.1).

As in (3.4), we can express (3.15) in the more concise notation:

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s},\mathbf{u}} \frac{1}{2}\|\mathbf{h} - \mathbf{K}\mathbf{u}\|_2^2 \;\text{ s.t. }\; \mathbf{u} = \mathbf{G}\mathbf{s}, \tag{3.16}$$

where $\mathbf{u}$ and $\mathbf{h}$ were defined in (3.4),

$$\mathbf{G} \triangleq \begin{bmatrix} \mathbf{I} \\ \mathbf{R} \end{bmatrix}, \text{ and } \mathbf{K} \triangleq \begin{bmatrix} \mathbf{W}^{\frac{1}{2}}\mathbf{D} & \mathbf{0} \\ \mathbf{0} & \sqrt{\lambda}\mathbf{I} \end{bmatrix}.$$

We then tackle (3.16) using the previously described AL formalism and obtain the following AL function-based minimization problem:

$$\arg\min_{\mathbf{s},\mathbf{u}} \frac{1}{2}\|\mathbf{h}-\mathbf{K}\mathbf{u}\|_2^2 + \frac{1}{2}\|\mathbf{u}-\mathbf{G}\mathbf{s}-\boldsymbol{\eta}\|_{\mathbf{V}}^2, \tag{3.17}$$

where $\boldsymbol{\eta}$ and $\mathbf{V}$ were defined in (3.5).

Due to the complexity of jointly minimizing (3.17) over $\mathbf{s}$ and $\mathbf{u}$, we again consider an alternating minimization scheme. In particular, we sequentially solve the following set of equations:

$$\mathbf{s}^{(j+1)} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{u}^{(j)}-\mathbf{G}\mathbf{s}-\boldsymbol{\eta}^{(j)}\|_{\mathbf{V}}^2, \tag{3.18}$$

$$\mathbf{u}^{(j+1)} = \arg\min_{\mathbf{u}} \frac{1}{2}\|\mathbf{h}-\mathbf{K}\mathbf{u}\|_2^2 + \frac{1}{2}\|\mathbf{u}-\mathbf{G}\mathbf{s}^{(j+1)}-\boldsymbol{\eta}^{(j)}\|_{\mathbf{V}}^2. \tag{3.19}$$

As with ADMM–Circ, the update equation for $\mathbf{u}$, (3.19), has a simple closed-form solution which can be decoupled into two parallel updates in terms of $\mathbf{u}_1$ and $\mathbf{u}_0$ due to the block diagonal structures of $\mathbf{K}$ and $\mathbf{V}$:

$$\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}\left[\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z} + \nu_1(\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_1^{(j)})\right], \tag{3.20}$$

$$\mathbf{u}_0^{(j+1)} = \frac{\nu_0}{\nu_0+\lambda}\left(\mathbf{R}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_0^{(j)}\right). \tag{3.21}$$

where $\mathbf{D}_2 \triangleq \mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \nu_1\mathbf{I}$ is a diagonal matrix.

Equation (3.18) does not have an efficient closed-form solution due to the size and complexity of $\mathbf{R}$. Instead, we approximately solve (3.18) using several iterations of the preconditioned conjugate gradient (PCG) method with warm starting, the optimal number of which is determined empirically. We design the specific preconditioner, $\mathbf{P}$, by considering the closed-form solution of (3.18):

$$\mathbf{G}_2\mathbf{s}^{(j+1)} = \mathbf{G}^{\mathrm{H}}\mathbf{V}\left(\mathbf{u}^{(j)} - \boldsymbol{\eta}^{(j)}\right), \tag{3.22}$$

where $\mathbf{G}_2 = \mathbf{G}^{\mathrm{H}}\mathbf{V}\mathbf{G} = \nu_1\mathbf{I} + \nu_0\mathbf{R}^{\mathrm{H}}\mathbf{R}$. Our goal is to create an easily invertible $\mathbf{P}$ that preconditions $\mathbf{G}_2$ to obtain fast convergence for this subproblem. For typical

```
ADMM–CG
Initialize: $\mathbf{u}_1^{(0)} = \mathbf{s}^{(0)}$, $\mathbf{u}_0^{(0)} = \mathbf{R}\mathbf{s}^{(0)}$, $\boldsymbol{\eta}_0^{(0)} = \mathbf{0}$, $\boldsymbol{\eta}_1^{(0)} = \mathbf{0}$ and $j = 0$.
Set $\mathbf{D}_2^{-1} = \left[\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \nu_1\mathbf{I}\right]^{-1}$ and $\mathbf{z}_2 = \mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z}$.
Repeat until stop criterion is achieved:
    1. $\mathbf{s}^{(j+1)}$ from PCG solution of (3.18) using (3.23).
    2. $\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}\left[\mathbf{z}_2 + \nu_1(\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_1^{(j)})\right]$,
       $\mathbf{u}_0^{(j+1)} = \frac{\nu_0}{\nu_0+\lambda}\left(\mathbf{R}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_0^{(j)}\right)$,
    3. $\boldsymbol{\eta}_1^{(j+1)} = \boldsymbol{\eta}_1^{(j)} - (\mathbf{u}_1^{(j+1)} - \mathbf{s}^{(j+1)})$,
       $\boldsymbol{\eta}_0^{(j+1)} = \boldsymbol{\eta}_0^{(j)} - (\mathbf{u}_0^{(j+1)} - \mathbf{R}\mathbf{s}^{(j+1)})$,
    4. $j = j + 1$.
```

Figure 3.12: Overview of the ADMM–CG algorithm. Note that $\mathbf{R}\mathbf{s}^{(j+1)}$ only needs to be computed once per iteration.

finite differencing matrices with non-periodic boundaries, $\mathbf{R}^{\mathrm{H}}\mathbf{R}$ has a near block circulant with circulant blocks structure. We therefore approximate $\mathbf{R}^{\mathrm{H}}\mathbf{R}$ in our preconditioner as $\mathbf{Q}^{\mathrm{H}}\boldsymbol{\Omega}\mathbf{Q}$ where $\mathbf{Q}$ is a (multidimensional) discrete Fourier transform (DFT) matrix and $\boldsymbol{\Omega}$ is a diagonal matrix containing the spectrum of the convolution kernel of $\mathbf{R}^{\mathrm{H}}\mathbf{R}$ [34]. Our resulting preconditioner is

$$\mathbf{P} = \mathbf{Q}^{\mathrm{H}}\left(\nu_1\mathbf{I} + \nu_0\boldsymbol{\Omega}\right)\mathbf{Q}. \tag{3.23}$$

Fig. 3.12 summarizes the resulting estimation algorithm composed of these update steps and the corresponding Lagrange multiplier updates, ADMM–CG. Note that the minimization in Step 1 is inexact, requiring an iterative solution; however, the optimal number of iterations is typically small. Furthermore, it can be shown that this algorithm is equivalent to an ADMM algorithm with an approximate update step for which the errors at each outer iteration can be made absolutely summable by using enough PCG iterations. We can therefore conclude that this algorithm converges to the solution of (3.1) as per [30, Th. 8].

```
ADMM–CG–IU
Initialize: $\mathbf{u}_1^{(0)} = \mathbf{s}^{(0)}$, $\mathbf{u}_0^{(0)} = \mathbf{R}\mathbf{s}^{(0)}$, $\boldsymbol{\eta}_0^{(0)} = \mathbf{0}$, $\boldsymbol{\eta}_1^{(0)} = \mathbf{0}$ and $j = 0$.
Set $\mathbf{D}_2^{-1} = \left[\mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{D} + \nu_1\mathbf{I}\right]^{-1}$ and $\mathbf{z}_2 = \mathbf{D}^{\mathrm{H}}\mathbf{W}\mathbf{z}$.
Repeat until stop criterion is achieved:
    1. $\mathbf{s}^{(j+1)}$ from PCG solution of (3.18) using (3.23).
    2. $\boldsymbol{\eta}_1^{(j+1/2)} = \boldsymbol{\eta}_1^{(j)} - (\mathbf{u}_1^{(j)} - \mathbf{s}^{(j+1)})$,
       $\boldsymbol{\eta}_0^{(j+1/2)} = \boldsymbol{\eta}_0^{(j)} - (\mathbf{u}_0^{(j)} - \mathbf{R}\mathbf{s}^{(j+1)})$,
    3. $\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}\left[\mathbf{z}_2 + \nu_1(\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_1^{(j+1/2)})\right]$,
       $\mathbf{u}_0^{(j+1)} = \frac{\nu_0}{\nu_0+\lambda}\left(\mathbf{R}\mathbf{s}^{(j+1)} + \boldsymbol{\eta}_0^{(j+1/2)}\right)$,
    4. $\boldsymbol{\eta}_1^{(j+1)} = \boldsymbol{\eta}_1^{(j+1/2)} - (\mathbf{u}_1^{(j+1)} - \mathbf{s}^{(j+1)})$,
       $\boldsymbol{\eta}_0^{(j+1)} = \boldsymbol{\eta}_0^{(j+1/2)} - (\mathbf{u}_0^{(j+1)} - \mathbf{R}\mathbf{s}^{(j+1)})$,
    5. $j = j + 1$.
```

Figure 3.13: The ADMM–CG algorithm with intermediate Lagrange multiplier updating (ADMM–CG–IU). Note that $\mathbf{R}\mathbf{s}^{(j+1)}$ only needs to be computed once per iteration.

Alternating minimization with intermediate updating

We also explored updating the Lagrange multipliers between each alternating minimization step. The resulting variation, ADMM–CG–IU, is presented in Fig. 3.13. As with the ADMM–Circ–IU algorithm, this algorithm lacks a guarantee of convergence although such guarantees exist for similar intermediate updating algorithms [31].

Parameter selection

The parameter selection strategy for our ADMM–CG based algorithms is similar to the strategy for ADMM–Circ because of the analogous structures of the alternating minimization steps. The major difference is that the update of $\mathbf{u}_0$ in (3.21) does not require the inversion of a matrix but rather a scalar term. In fact, this scalar term has the same form as $\kappa(\mathbf{B}_2)$ in Section 3.1.2. Subsequently, we found that setting ν_0 such that the scalar $\frac{\nu_0+\lambda}{\nu_0} \in [200, 400]$ and then setting ν_1 such that $\kappa(\mathbf{G}_2) \in [200, 1000]$

provided reasonable convergence rates.

Results

To evaluate our proposed ADMM–CG based algorithms, we performed the same experiments as in Section 3.1.3. The cost function was setup as described in Section 3.1.3 and the same ratio based estimate was used to initialize the algorithms. We used a single PCG iteration for the approximate update to $\mathbf{s}$ in the ADMM–CG based algorithms as this provided the fastest convergence rates with respect to time. We selected the AL penalty parameters ν_0 and ν_1 for ADMM–CG such that $\frac{\nu_0+\lambda}{\nu_0} = 225$ and $\kappa(\mathbf{G}_2) = 600$. As further discussed in Section 3.2.2, the optimal condition numbers for ADMM–CG–IU depended on the data and are therefore mentioned in the appropriate subsections.

(1) Simulated brain data: We ran 20 000 iterations of the ADMM–CG based algorithms on the simulated brain data described in Section 3.1.3. For the ADMM–CG–IU algorithm, we selected ν_0 and ν_1 such that $\frac{\nu_0+\lambda}{\nu_0} = 375$ and $\kappa(\mathbf{G}_2) = 600$. Our proposed ADMM–CG and ADMM–CG–IU algorithms converged to a normalized ℓ_2-distance of less than -200 dB from the Cholesky based solution to (3.1) and appeared nearly identical to Fig. 3.6. The convergence rates of the algorithms were similar for all four coils and thus we present the results for the same coil that was presented in Section 3.1.3. Fig. 3.14 plots $\mathcal{D}(\mathbf{s}^{(j)})$ with respect to both iteration and time for the ADMM–CG based algorithms as well as the algorithms evaluated in Section 3.1.3. ADMM–CG–IU and ADMM–CG were both slower than PCG–Circ, but faster than conventional CG, reaching $\mathcal{D}(\mathbf{s}^{(j)}) = 0.1\,\%$ in approximately 145 and 185 seconds, respectively.

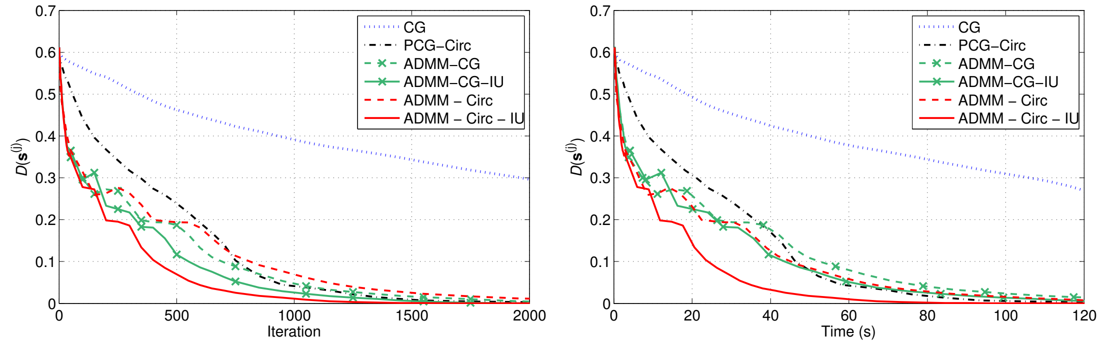

Figure 3.14: Plots of the normalized ℓ_2-distance between $\mathbf{s}^{(j)}$ and $\hat{\mathbf{s}}$, $\mathcal{D}(\mathbf{s}^{(j)})$, with respect to iteration (left) and time (right) for the bottom left brain data surface coil image in Fig. 3.5.

(2) Breast phantom data: We also ran 20 000 iterations of the ADMM–CG based algorithms on the breast phantom data described in Section 3.1.3. For the ADMM–CG–IU algorithm, we selected ν_0 and ν_1 such that $\frac{\nu_0+\lambda}{\nu_0} = 250$ and $\kappa(\mathbf{G}_2) = 600$. Again, both of our proposed algorithms converged to a normalized ℓ_2-distance of less than -200 dB from the Cholesky based solution to (3.1) and appeared nearly identical to Fig. 3.9(b). The convergence rates of the algorithms were similar for all four coils and thus we present the results for the same coil that was presented in Section 3.1.3. Fig. 3.15 plots $\mathcal{D}(\mathbf{s}^{(j)})$ with respect to both iteration and time for the ADMM–CG based algorithms as well as the algorithms evaluated in Section 3.1.3. ADMM–CG–IU was faster than both PCG–Circ and regular CG converging within $\mathcal{D}(\mathbf{s}^{(j)}) = 0.1\,\%$ in approximately 80 seconds. ADMM–CG was slower than its intermediate updating counterpart and PCG–Circ with a convergence time of nearly 120 seconds.

Discussion

The convergence rates with respect to iteration of the ADMM–CG based algorithms were close to their ADMM–Circ counterparts. However, the ADMM–CG based algorithms were much slower in time due to the added overhead of the PCG

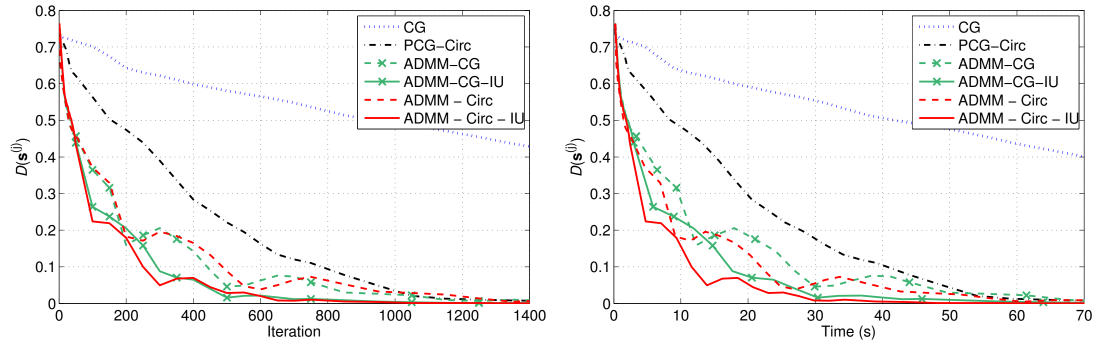

Figure 3.15: Plots of the normalized ℓ_2-distance between $\mathbf{s}^{(j)}$ and $\hat{\mathbf{s}}$, $\mathcal{D}(\mathbf{s}^{(j)})$, with respect to iteration (left) and time (right) for the bottom left breast data surface coil image in Fig. 3.9.

solution used to approximate Step 1. In fact, even when using only one iteration of PCG for this approximation, the per iteration costs of the ADMM–CG algorithms are much higher than those of the ADMM–Circ algorithms, Table 3.2.

The convergence curves of the ADMM–CG based algorithms exhibit a higher rate of non-monotonic behavior than the ADMM–Circ algorithms. This is partly caused by the approximate update in Step 1. If we run several more PCG sub-iterations in Step 1, the convergence curves with respect to iteration of the ADMM–CG algorithms appear similar to their ADMM–Circ counterparts (although much slower with respect to time). As with ADMM–Circ, the parameter settings that provided the fastest convergence rates typically resulted in non-monotonicity in the $\mathcal{D}(\mathbf{s}^{(j)})$ plots.

The proposed ADMM–CG–IU algorithm was faster than PCG–Circ in the breast phantom experiment, but slower in the simulated brain experiment. As discussed in Section 3.1.4, the relative speed of the PCG–Circ algorithm depends on the accuracy of the preconditioner in (3.13) and thus on the data. Contrarily, the preconditioning used for the approximation of Step 1 in the ADMM–CG algorithms does not depend on the data; thus, these algorithms are less sensitive to such differences. The ADMM–CG algorithm, although initially faster, converged slower than PCG–Circ in

Table 3.2: Approximate Number of Complex Arithmetic Operations Per Iteration for ADMM–CG Algorithms

Estimator	**Number of Operations**
ADMM–CG–IU [a]	$28N + 18M + 2 \cdot \mathrm{O}_{\mathrm{FFT}}$[b]
ADMM–CG [a]	$26N + 16M + 2 \cdot \mathrm{O}_{\mathrm{FFT}}$

[a] Step 1 uses a single PCG iteration.

[b] $\mathrm{O}_{\mathrm{FFT}}$ denotes the cost of the FFT operations ($O(N \log(N))$).

both experiments. Therefore, using intermediate updating also significantly accelerated the convergence rates of this ADMM algorithm. All of our proposed algorithms were significantly faster than traditional CG.

The convergence rates of our proposed ADMM–CG algorithms were robust to the choice of condition numbers used to determine the AL penalty parameters ν_0 and ν_1. We found that the convergence rates remained similar for condition numbers that differed from the optimal values by up to fifty percent. The chosen condition numbers also worked well for a wide range of regularization parameter values λ. However, we found that the optimal condition numbers for ADMM–CG–IU depended on the data unlike for ADMM–CG and the ADMM–Circ based algorithms. Still, like the ADMM–Circ algorithms, the choice of the optimal condition numbers does not depend on the surface coil image. Therefore, if one wanted to fine-tune the convergence rate of the algorithms, a single coil of a multi-coil array would suffice.

3.2.3 AL Estimation Method with Similar Variable Splitting

In formulating our proposed ADMM–Circ algorithm, we originally explored a different variable splitting strategy involving a double splitting within the regularization term of (3.1) [18, 32]:

$$\arg\min_{\mathbf{s},\mathbf{u}_0} \frac{1}{2}\|\mathbf{z} - \mathbf{D}\mathbf{s}\|_{\mathbf{W}}^2 + \frac{\lambda}{2}\|\mathbf{B}\mathbf{u}_0\|_2^2 \text{ s.t. } \mathbf{u}_0 = \mathbf{C}\mathbf{u}_1 \text{ and } \mathbf{u}_1 = \mathbf{s}. \tag{3.24}$$

This variable splitting led to update equations with nearly identical structures to those of ADMM–Circ. Furthermore, the resulting AL algorithm and its intermediate updating variation had similar convergence rates to their ADMM counterparts. However, analyzing the convergence properties of these algorithms was more complicated as they did not have ADMM structures. Thus, we focused on the ADMM formulations.

3.2.4 Effect on SENSE Reconstruction Quality

The advantages and accuracy of similar regularized sensitivity profile estimators have been discussed [7, 15]; however, there has been limited investigation into their effects on SENSE reconstruction quality. We therefore compare the quality of the SENSE reconstructions created with the coil sensitivities estimated using the regularized method in (3.1) to those estimated using the commonplace ratio and ratio of low resolution images methods.

Simulated brain data

Our first experiment was performed using the simulated brain data outlined in Section 3.1.3. We began by simulating a full resolution calibration scan using the same parameters as Figs. 3.4 and 3.5. Next, we estimated the coil sensitivities from the resulting body and surface coil images using our regularized method, the ratio of low resolution images method, and the ratio method.

We implemented the regularized method using our ADMM–Circ–IU algorithm with the same parameters as in Section 3.1.3. The ratio of low resolution images method was implemented by taking a set number of samples from the center of k-space of each coil, zero padding to get 256×192 element matrices (corresponding to a $256\,\text{mm} \times 192\,\text{mm}$ FOV), and reconstructing low resolution body coil and surface coil images using inverse DFTs. Smooth sensitivity estimates were then obtained by taking the ratio of these low resolution images. We present the results for two different amounts of sampling. The first uses the center 13×9 k-space samples resulting in sensitivity estimates that extend smoothly to the image edges. The second uses the center 51×38 samples which was found to provide the best SENSE reconstruction quality for this method. In both cases we applied a Hamming window to the selected k-space data to reduce any Gibbs ringing artifacts. The conventional ratio estimate ($\hat{s}_i = z_i/y_i$) was masked to remove the highly corrupt estimates of the non-object pixels using a binary mask created by thresholding the body coil image. The resulting sensitivity profile estimates for a single, representative coil are presented in Fig. 3.16.

As seen in Fig. 3.6, the regularized estimate is very close to the true sensitivity, differing only at the corners of the image. The minor discrepancies at the corners of the estimates are in part due to selecting a regularization parameter that emphasized accuracy over the object pixels and their immediate surrounding area as well as from the fact that there is no information about the true sensitivity in this region of the image. The conventional ratio estimate is much noisier over the object pixels and has no extrapolation. Both low resolution ratio estimates are smooth over the object support with varying degrees of extrapolation into the surrounding regions. However, the implicit smoothing of these methods creates inaccuracies in the estimates near object edges and in areas with predominantly low signal. Furthermore, any voxels significantly beyond the extrapolated regions exhibit large estimation errors. The typical errors that result from Gibbs ringing artifacts [7] have been reduced by the additional windowing.

Data from the four surface coils were then simulated for every other vertical line

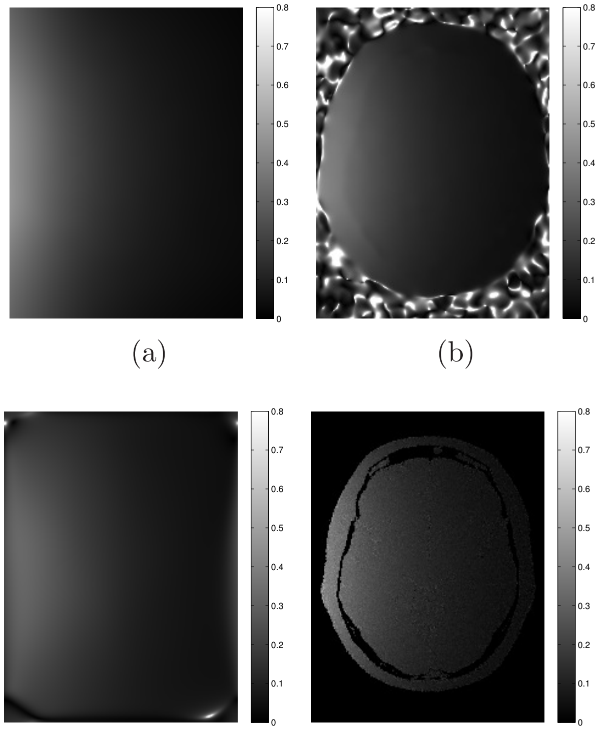

(a) (b)

(c) (d)

Figure 3.16: Example sensitivity profile estimates for the brain data using (a) the regularized method (b) the ratio of low resolution images method with the center 51×38 samples, (c) the ratio of low resolution images method with the center 13×9 samples, and (d) the conventional ratio method.

in k-space. SENSE reconstructions [2] were performed using this undersampled data set and the various sensitivity profile estimates. We restricted the reconstruction to a masked region found by dilating the threshholded body coil image by two pixels. These reconstructions and their differences to the truth are presented in Fig. 3.17.

The resulting normalized root-mean-square errors (NRMSE) between the SENSE reconstructions and the truth are presented in Table 3.3. The regularized method led to the most accurate SENSE reconstruction in terms of NRMSE as well as the one with the fewest structural artifacts (beyond the amplified noise inherent to SENSE reconstruction). The low resolution ratio method with the center 13×9 samples led to the least accurate SENSE reconstruction. It is clear from the artifacts in

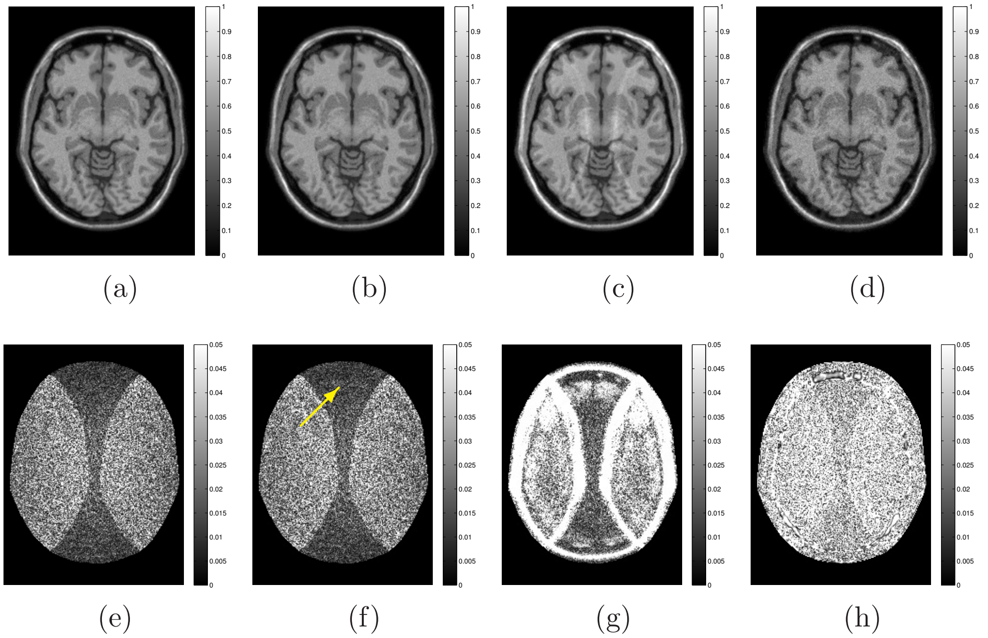

(a) (b) (c) (d)

(e) (f) (g) (h)

Figure 3.17: Resulting two-fold accelerated SENSE reconstructions for the brain data using (a) the regularized method (b) the ratio of low resolution images method with the center 51 × 38 samples, (c) the ratio of low resolution images method with the center 13 × 9 samples, and (d) the conventional ratio method sensitivity profile estimates. The corresponding differences to the truth are presented below in (e – h). The yellow arrow specifies an artifact in the SENSE reconstruction caused by inaccurate sensitivity estimates in a low signal region.

the difference image that oversmoothing led to large inaccuracies in the sensitivity profile estimates at the object edges. The low resolution ratio method with the center 51×38 samples led to the second most accurate reconstruction. Although the effects are less severe than for the 13×9 case, there are again structural artifacts in the reconstructions due to inaccurate sensitivity estimates at the object boundaries and low signal regions within the brain. The conventional ratio method also led to significant artifacts in the SENSE reconstruction. Specifically, the lack of smoothing in the sensitivity estimates led to high noise in the SENSE reconstruction, while the lack of extrapolation resulted in aliased object edges within the final reconstruction.

Table 3.3: NRMSEs Between the True Brain Images (Stationary and Shifted) and the SENSE Reconstructions

	Regularized	**Low Res. Ratio**		**Ratio**
		51×38	13×9	
Shift = 0	0.06	0.07	0.16	0.12
Shift = 2 pixels	0.06	0.07	0.16	0.16

Shifted simulated brain data

One possible complication when performing SENSE imaging is if the patient moves between the calibration and acquisition scans. In such cases, poorly extrapolated sensitivity profile estimates will introduce significant artifacts into the reconstruction [8]. To evaluate the different sensitivity estimators under such circumstances, we simulated a set of two times undersampled surface coil images in which the brain was moved two pixels to the right with respect to the coil sensitivities and the field-of-view. We then reconstructed the image using the previously estimated coil sensitivities over an equally shifted masked region. These reconstructions and their differences to the shifted truth are presented in Fig. 3.18.

The resulting NRMSEs between the SENSE reconstructions and the true shifted brain are presented in Table 3.3. The regularized method again led to the most accurate SENSE reconstruction with similar NRMSE and a lack of structural artifacts. The low resolution ratio method with the center 51×38 samples led to the second most accurate SENSE reconstruction; however, the inaccuracies in the sensitivity estimates at the object edges resulted in larger artifacts due to the shift, particularly at the far right side of the brain. The low resolution ratio method with the center 13×9 samples again resulted in the worst SENSE reconstruction. The shift of two pixels to the right emphasized the inaccuracies in the estimates near the object edges by introducing even larger artifacts (not visible with the current contrast windowing).

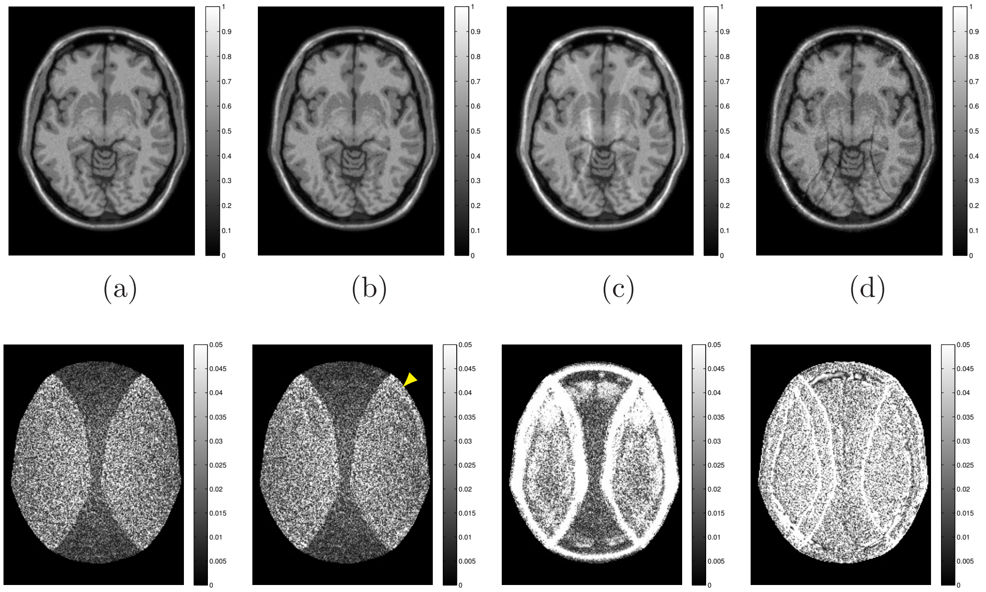

(a) (b) (c) (d)

(e) (f) (g) (h)

Figure 3.18: Resulting two-fold accelerated SENSE reconstructions of a brain shifted two pixels to the right with respect to Fig. 3.17 using the previous (a) regularized method (b) the ratio of low resolution images method with the center 51×38 samples, (c) the ratio of low resolution images method with the center 13×9 samples, and (d) the conventional ratio method sensitivity profile estimates. The corresponding differences to the shifted truth are presented below in (e – h). The yellow arrow indicates an area with increased artifacts due to inaccuracies in the sensitivity estimates at the object edges.

The SENSE reconstruction based on the conventional ratio method was significantly affected by the shift. In particular, the lack of any extrapolation in the estimated sensitivity profiles resulted in large artifacts within the object support.

High SNR simulated brain data

To better illustrate the typical inaccuracies produced by the ratio of low resolution images estimation method, we repeated the previous SENSE reconstruction experiments using simulated brain data with a higher SNR of 20. The specific body coil and four surface coil images are presented in Fig. 3.19. We performed sensitivity estimation using the regularized method and the ratio of low resolution images method with 51×38 samples. The resulting estimates for a representative coil are presented in Fig. 3.20.

The sensitivity profile estimates are similar to those for the case of lower SNR brain data found in Fig. 3.16. The regularized estimate is again very close to the true sensitivity differing only at the corners of the image. The ratio of low resolution images estimate is smooth over the object support and exhibits some extrapolation. However, there are noticeable inaccuracies in areas corresponding to regions of low signal within the brain.

We performed two-fold accelerated SENSE reconstructions with the higher SNR brain data and these sensitivity profile estimates. The results for both the case of no shift between calibration and scan, as well as a two pixel shift, are presented in Fig. 3.21. As with the case of low SNR brain data, the reconstructions created using the regularized sensitivity estimate have very low error and no major structural artifacts. Furthermore, the two pixel shift had little effect on the reconstruction quality indicating accurate extrapolation within the estimate. In contrast, the reconstruction created using the low resolution ratio estimates had several large structural artifacts (indicated with a yellow arrow) that were a result of the inaccurate sensitivity profile estimates in regions of low signal. The two pixel shift increased these artifacts indicating inaccurate extrapolation within the sensitivity estimates.

Breast phantom data

We also compared the sensitivity estimation methods using our breast phantom data from Section 3.1.3. In this case, we estimated the sensitivities of the breast phantom images presented in Fig. 3.9 using the same four methods as before: the regularized method with $\lambda = 2^7$, the ratio method, and the low resolution ratio method with both the center 77×19 and center 19×5 samples zero padded to

384×96 elements. The resulting estimates for a representative coil are presented in Fig. 3.22.

The regularized estimate is smooth over the entire field-of-view and closely matches the general trend in the ratio estimate. The low resolution ratio estimate with the center 77×19 samples is reasonably smooth over the object support with some extrapolation into the surrounding pixels. There are inaccuracies in the estimate near regions of low signal such as at the object edges and over the far right breast. The low resolution ratio estimate with the center 19×5 samples is smoother than the case of 77×19 samples and exhibits greater extrapolation. However, this estimate suffers from oversmoothing and is highly inaccurate at the object edges. Both of the low resolution ratio methods benefited from using a Hamming window to reduce the Gibbs ringing artifacts. The ratio estimate is very noisy over the object pixels and has no extrapolation.

To simulate the minor changes in the data that would occur between a calibration scan and an acquisition scan, we performed a SENSE reconstruction on a neighboring two-dimensional slice of our breast phantom data. The fully sampled body and surface coil images of this slice are presented in Fig. 3.23. First, we undersampled the surface coil images by selecting every other vertical line in k-space. As was done for the simulated brain data, we then performed SENSE reconstructions over a masked region using the previously estimated coil sensitivities. These reconstructions are presented in Fig. 3.24.

The SENSE reconstruction resulting from the regularized estimate has very high quality and few visible artifacts when compared to the body coil image in Fig. 3.23(a). The reconstruction resulting from the low resolution ratio estimate with the center 77×19 samples appears similar to that of the regularized estimate; however, the inner parts of the breasts are darker than in the body coil and surface coil images. This is largely a result of inaccurate sensitivity estimation in these low signal regions. In addition to the darkening artifact in the low signal regions of the image, the reconstruction resulting from the low resolution ratio estimate with the center 19×5 samples also has aliased edges of the breasts within the object support (indicated by a yellow arrow). These are a result of inaccurate sensitivity estimates at the object edges caused by oversmoothing. The reconstruction resulting from the conventional ratio estimate is very noisy and has several bright artifacts. This is due to inaccurate sensitivity estimation over the low signal pixels within the object support and a lack of extrapolation.

Discussion

From these experiments, we conclude that the regularized sensitivity estimation method outlined in (3.1), although more computationally expensive, provides improved sensitivity estimates for use in SENSE reconstructions compared to other commonly used non-parametric methods. Using a ratio of low resolution images provides reasonable estimates if the correct number of samples is selected. However, even after windowing to reduce the Gibbs ringing artifacts, these estimates are typically inaccurate at object edges and in areas of low signal. This results in artifacts in the SENSE reconstructions. The lack of smoothing and extrapolation in the conventional ratio method results in SENSE reconstructions that are very noisy and prone to large artifacts due to motion.

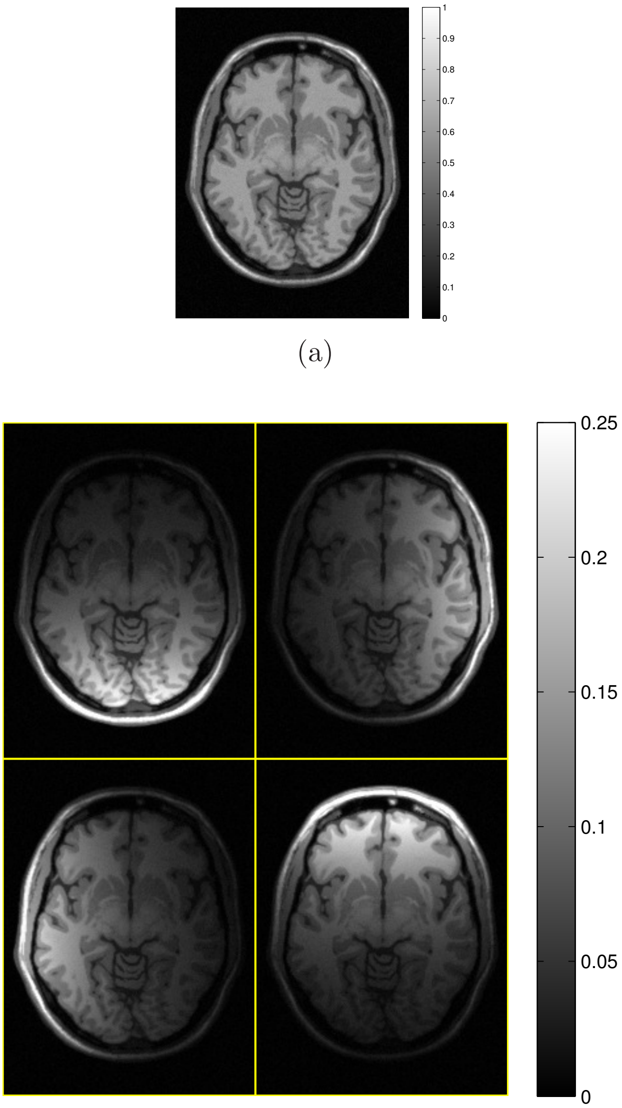

(a)

(b)

Figure 3.19: The magnitude of the fully sampled (a) body coil and (b) surface coil images for our high SNR simulated brain data.

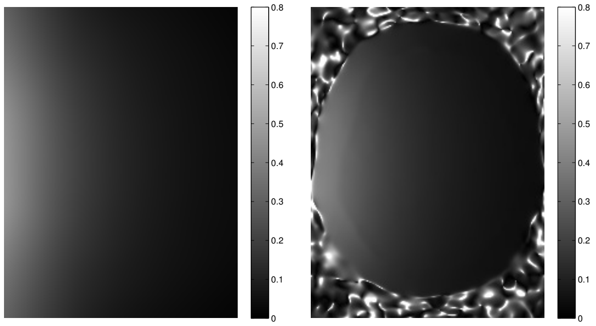

(a) (b)

Figure 3.20: Example sensitivity profile estimates found for the high SNR brain data using (a) the regularized method and (b) the ratio of low resolution images method with the center 51 × 38 samples.

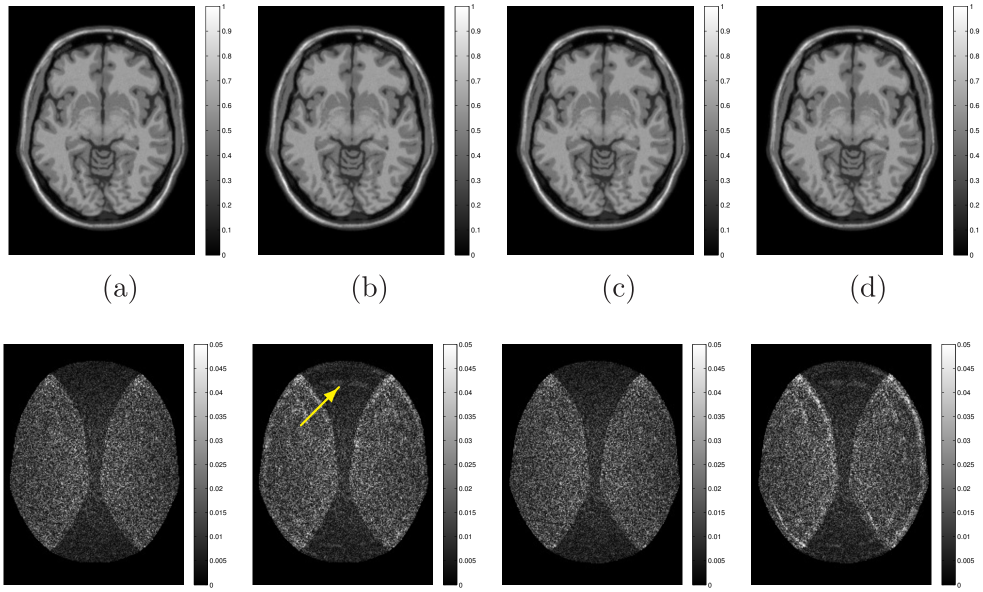

(e) (f) (g) (h)

Figure 3.21: Resulting two-fold accelerated SENSE reconstructions of the high SNR simulated brain data. For the case of no shift, (a) and (b) are the reconstructions corresponding to the regularized method and ratio of low resolution images method, respectively. (c) and (d) are the corresponding reconstructions for the case of a two pixel shift. The difference to the truth or shifted truth for each reconstruction is presented below in (e – h).

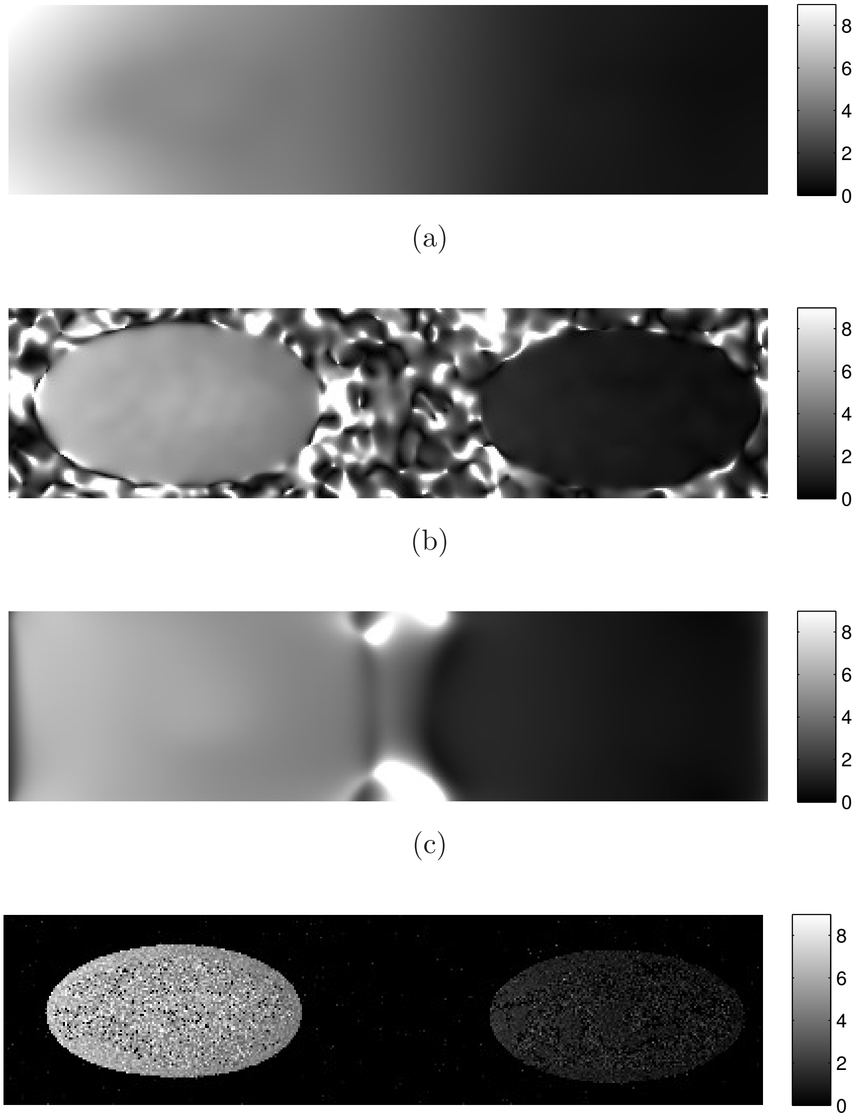

(a)

(b)

(c)

(d)

Figure 3.22: Example sensitivity profile estimates found for the breast phantom data using (a) the regularized method (b) the ratio of low resolution images method with the center 77×19 samples, and (c) the ratio of low resolution images method with the center 19×5 samples, and (d) the conventional ratio method.

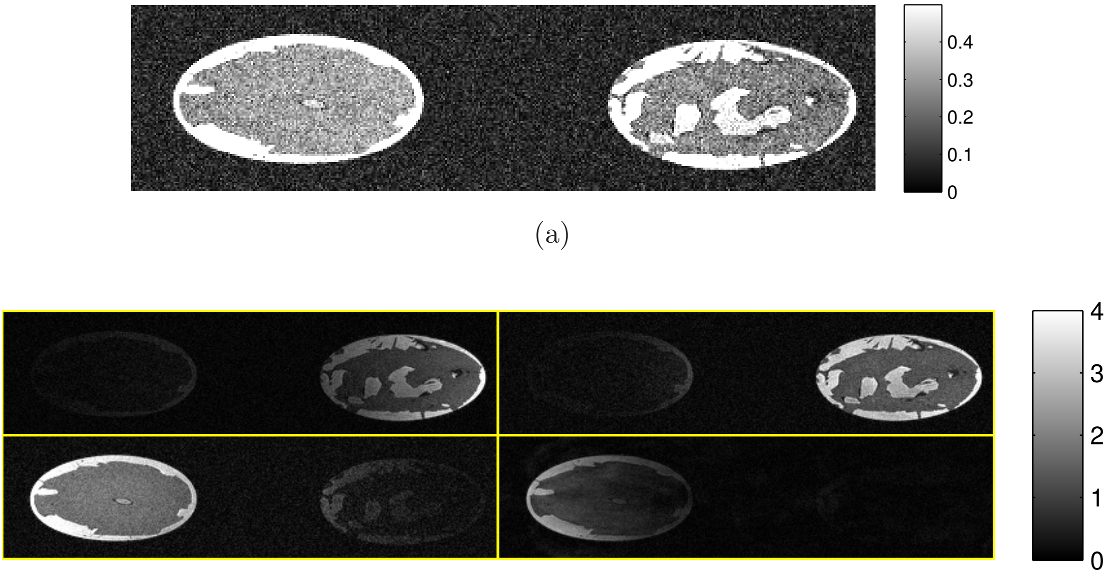

(a)

(b)

Figure 3.23: The magnitude of the fully sampled (a) body coil and (b) surface coil images for the neighboring two-dimensional slice of our breast phantom data.

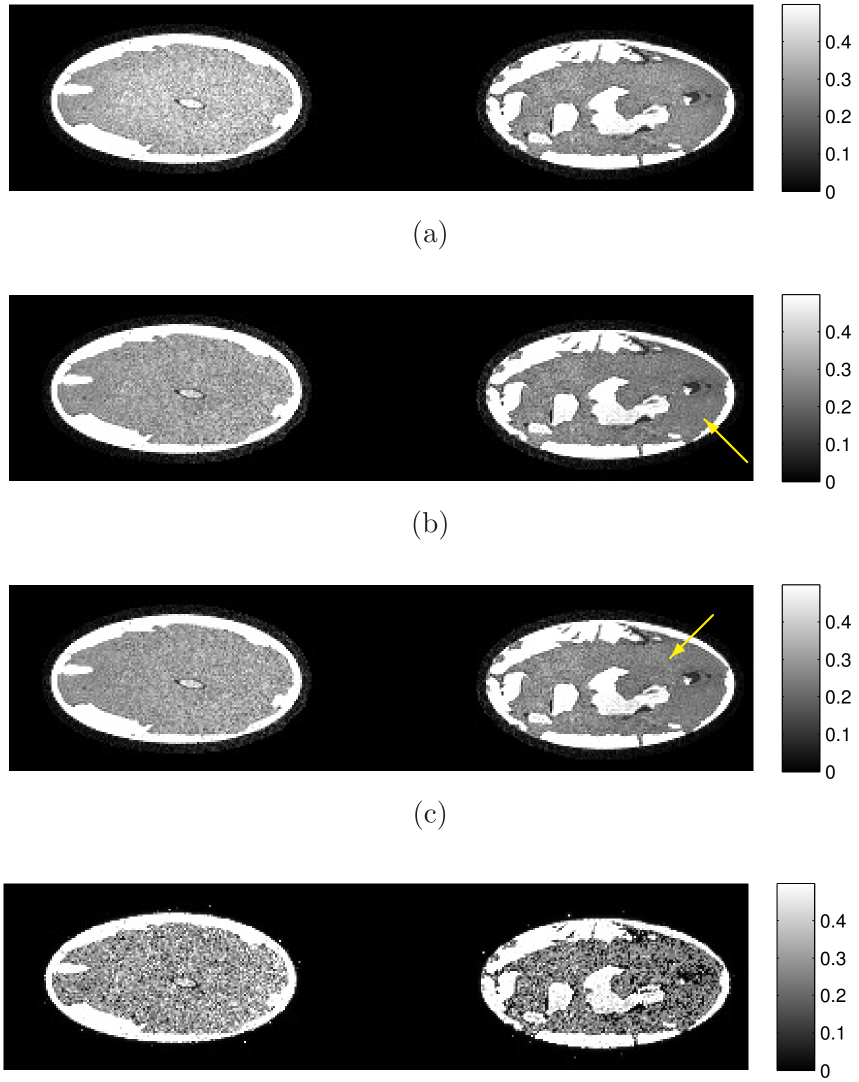

(a)

(b)

(c)

(d)

Figure 3.24: Resulting two-fold accelerated SENSE reconstructions of the neighboring slice of breast phantom data (Fig. 3.23) using the previous (a) regularized method (b) ratio of low resolution images method with the center 77 × 19 samples, (c) ratio of low resolution images method with the center 19×5 samples, and (d) conventional ratio method sensitivity profile estimates. The arrow in (b) points to a dark region in the reconstruction, while the arrow in (c) points to a reconstruction artifact caused by inaccurate sensitivity estimation at object edges.

3.2.5 Estimation Over a Convex Hull Mask

If estimating over a masked region to reduce computation time, the mask must be carefully selected to ensure accurate estimates. For data sets with spatially-contiguous support, such as the simulated brain in Section 3.1.3, this is relatively trivial; however, this is not the case for more complicated data sets such as our breast phantom data whose field-of-view (FOV) contains several spatially distinct objects. Due to the underlying physics, the typical coil sensitivity profile should smoothly vary across the entire FOV and generally decrease with distance from the coil. However, using a tight mask isolates the estimate over each object and this can result in large errors for objects that have low signal or only a few pixels. This can be avoided by using a mask consisting of a convex hull containing the spatially distinct objects.

To illustrate this phenomenon, we considered another slice of our breast phantom data, Fig. 3.25. This image has a small object to the left of the right breast (indicated by an arrow). We perform regularized estimation over a masked region consisting of spatially distinct objects as well as a masked region consisting of a convex hull of these points. Fig. 3.26 contains the two different masks and their corresponding sensitivity estimates. Fig. 3.27 presents line profiles of the absolute value of the sensitivity estimates taken horizontally through the center of the FOV for both estimates.

Comparing the two estimates, it is clear that they are similar for regions with relatively high SNR; however, they differ greatly over the small object next to the right breast. When using a tight mask, the estimated sensitivity in this region is very high in comparison to the nearby breast which does not match the underlying physics. This inaccuracy is a result of the estimate in this region being based on only a few low signal pixels. In contrast, the convex hull estimate is smooth over the

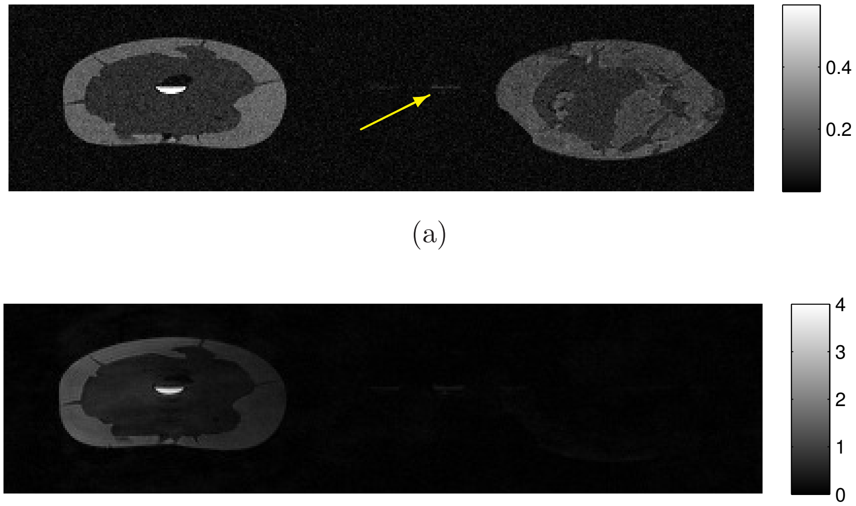

(a)

(b)

Figure 3.25: The magnitude of the (a) body coil and (b) surface coil image for an additional slice of our breast phantom data. The yellow arrow points to a small object within the FOV.

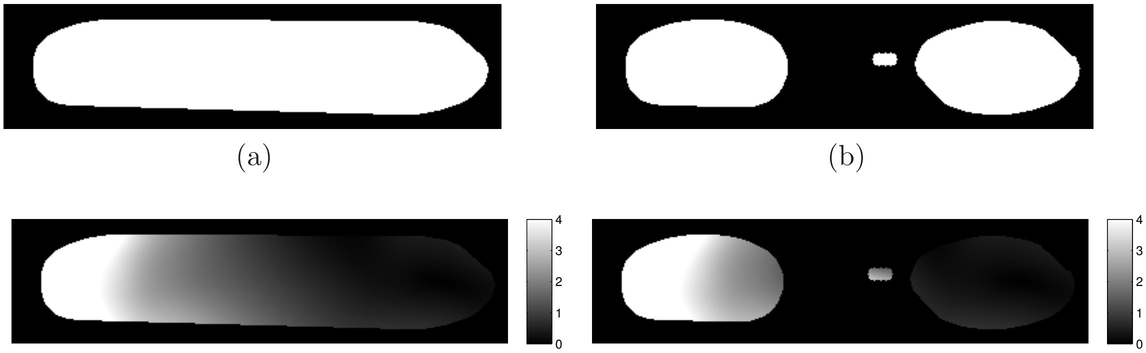

(a)

(b)

(c)

(d)

Figure 3.26: The masks for the cases of a (a) convex hull and (b) independent objects. The corresponding regularized estimates over the masked regions for the (c) convex hull and (d) independent objects.

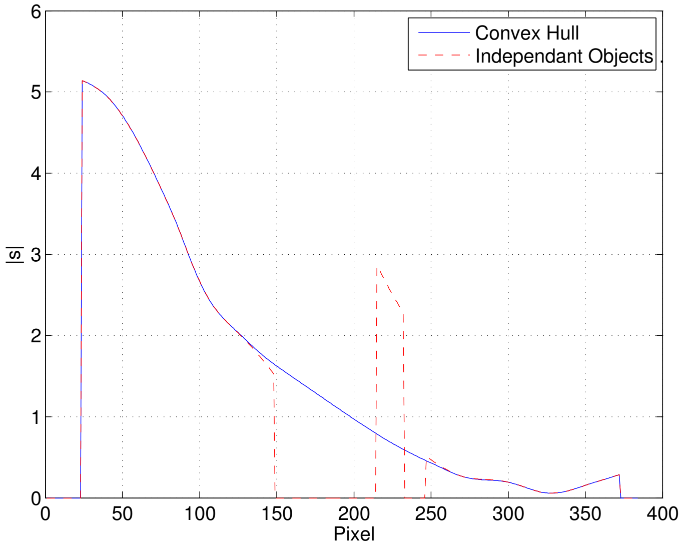

Figure 3.27: Horizontal line profiles taken through the center of the sensitivity estimates presented in Fig. 3.26.

entire masked region and the estimate over the small object is more realistic. This is because such a mask enforces smoothness both within and between all of the objects in the FOV. Thus, a convex hull should be used for the estimation mask to avoid inaccuracies in the final estimates.

3.2.6 Circulant Versus Non-Circulant Finite Differencing Matrices

In this section we demonstrate the importance of using a finite differencing matrix for the case of non-periodic boundary conditions (**R** or **BC**) rather than a finite differencing matrix for the case of periodic boundary conditions (**C**) in our cost function. Since $\mathbf{C}^{\mathrm{H}}\mathbf{C}$ is block circulant with circulant blocks, we will refer to the matrix for the case of periodic boundary conditions as the *circulant matrix*. In contrast, we will refer to the matrix for the case of non-periodic boundary conditions as the *non-circulant matrix*.

The receive coil is usually placed at or just beyond the boundary of the field-of-view. Since coil sensitivity is a physical phenomenon, its intensity will typically decrease with increased distance from the coil. However, if we use a circulant finite differencing matrix, we will be penalizing differences in the estimated sensitivities at opposing boundaries of the volume. Since there is often little information about the sensitivity at the edges of the volume, this penalization will result in a sensitivity estimate that dips near the coil and rises at the opposite side of the field-of-view. This is a clear mismatch with the underlying physics of the problem. Furthermore, because of the lack of meaningful information outside of the object voxels, this error will propagate to the estimate at the edges of the object. These errors within (and just outside) the object support can generate significant artifacts in SENSE reconstructions (see Section 3.2.4). Padding the image with zeros will not sufficiently remove this propagated estimation error. Thus, one must use a more realistic modeling assumption and select a non-circulant finite differencing matrix that avoids penalizing between opposite boundaries at the expense of increased complexity. To illustrate these claims, we recreated the estimates found in Section 3.1.3 using both the existing non-circulant finite differencing matrix ($\mathbf{R}$) and a circulant finite differencing matrix ($\mathbf{C}$).

Simulated brain data

We present the results for one coil of the simulated brain data. Fig. 3.28 presents the body coil image, true sensitivity, and resulting surface coil image used in this experiment. Fig. 3.29 presents the resulting estimates using both the non-circulant and circulant finite differencing matrices, as well as the percentage difference image for each estimate compared to the truth.

As stated before, the estimate using the circulant finite differencing matrix dips before the boundary near the coil and rises at the opposite edge of the field-of-view.

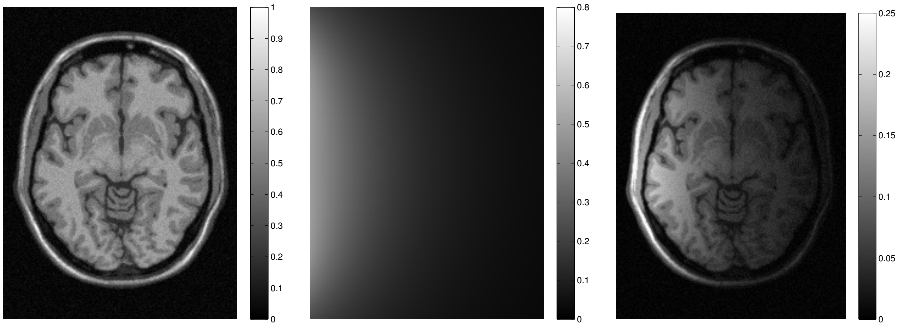

(a) (b) (c)

Figure 3.28: The (a) body coil, (b) true coil sensitivity, and (c) resulting surface coil magnitude images for the simulated brain data.

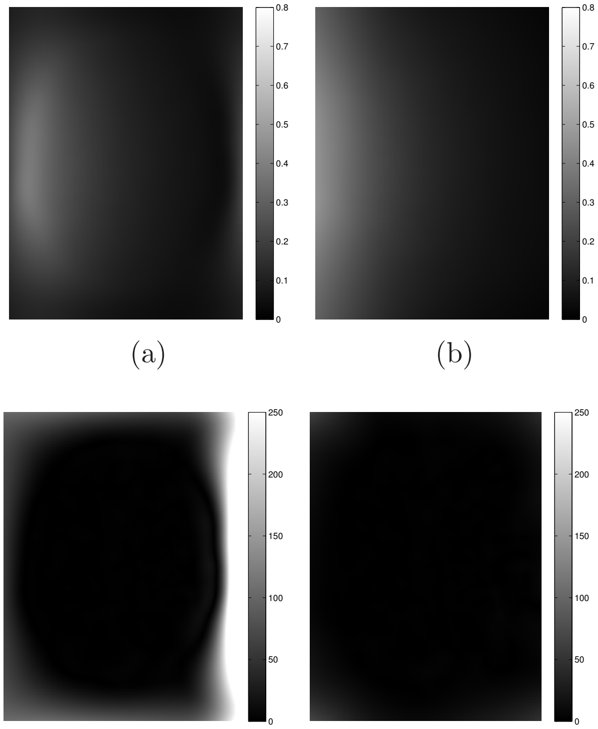

(a) (b)

(c) (d)

Figure 3.29: The resulting sensitivity estimates for the brain data using a (a) circulant matrix and a (b) non-circulant matrix. The percentage difference between the truth and the estimates from the (c) circulant matrix and the (d) non-circulant matrix are shown below.

This results in significant inaccuracies in the estimate at the image boundaries. In contrast, the estimate using the non-circulant matrix increases smoothly towards the image boundary closest to the coil. The overall estimation error is therefore much smaller and is confined to the outer corners of the image.

Fig. 3.30 presents the same estimates as Fig. 3.29, but masked in the spatial domain to highlight the error over the object support. In these images, we see that the error in the estimates from the non-circulant matrix has propagated to within the object support. This highly structured inaccuracy will cause large artifacts in SENSE reconstructions. In contrast, the error in the estimate from the circulant matrix is much lower over the entire object support and contains significantly less structure.

Padded simulated brain data

We also padded the brain data in Fig. 3.28 with zeros to get a 256×256 image (an addition of 32 pixels to both the left and right sides of the image). Fig. 3.31 presents the resulting estimates, masked to highlight the error over the object support. Similar inaccuracies to before are present in the circulant matrix estimate, while the estimate using the non-circulant matrix continues to have low error. Thus, the zero padding did not sufficiently mitigate the corruption of the estimate caused by using a circulant matrix.

Breast phantom data

We performed similar experiments on one coil of our breast phantom data found in Section 3.1.3. Fig. 3.32 presents the body coil and surface coil images used in this experiment.

Fig. 3.33 presents the resulting estimates using both the non-circulant and circu-

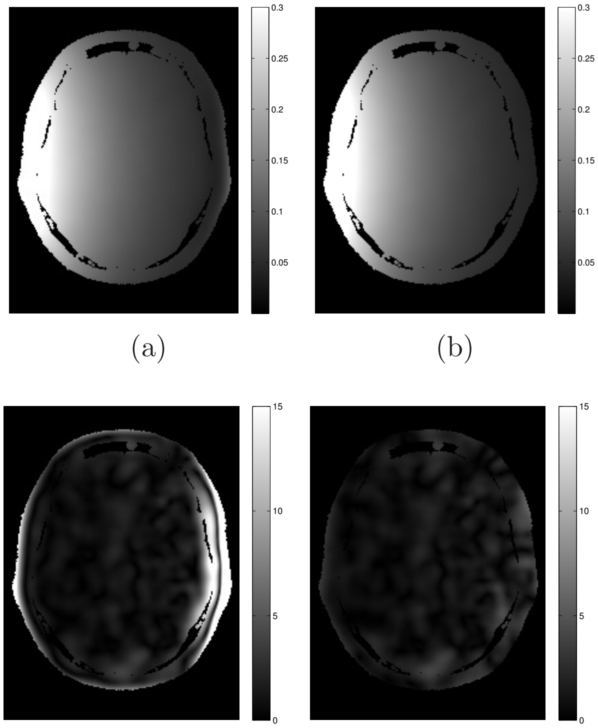

(a) (b)

(c) (d)

Figure 3.30: The same sensitivity estimates for the brain data as in Fig. 3.29 but masked to highlight the error over the object support. (a) and (c) are the resulting estimate and percentage difference to the truth, respectively, resulting from a circulant matrix. (b) and (d) are the same but resulting from a non-circulant matrix.

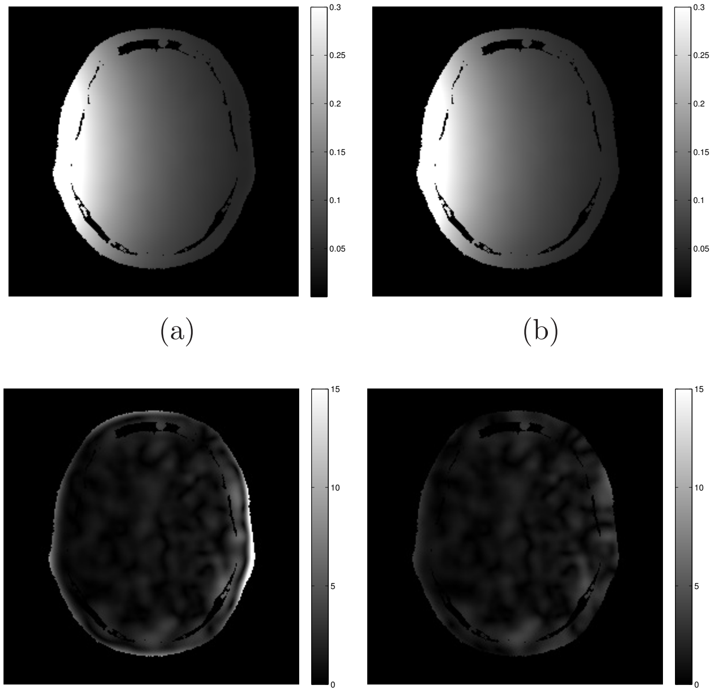

(a) (b)

(c) (d)

Figure 3.31: The masked sensitivity estimates for padded brain data generated using a (a) circulant matrix and a (b) non-circulant matrix. The masked percentage difference between the truth and the estimates from the (c) circulant matrix and the (d) non-circulant matrix are shown below.

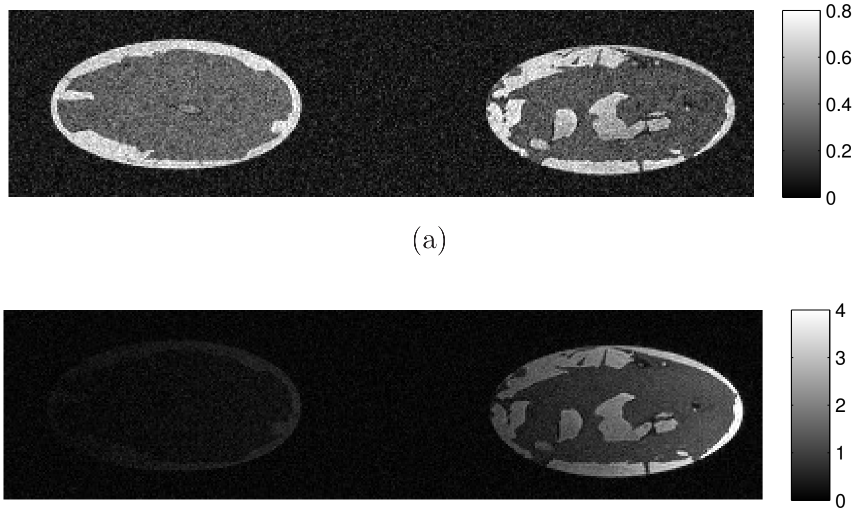

(a)

(b)

Figure 3.32: The (a) body coil and (c) surface coil magnitude images for the breast phantom data.

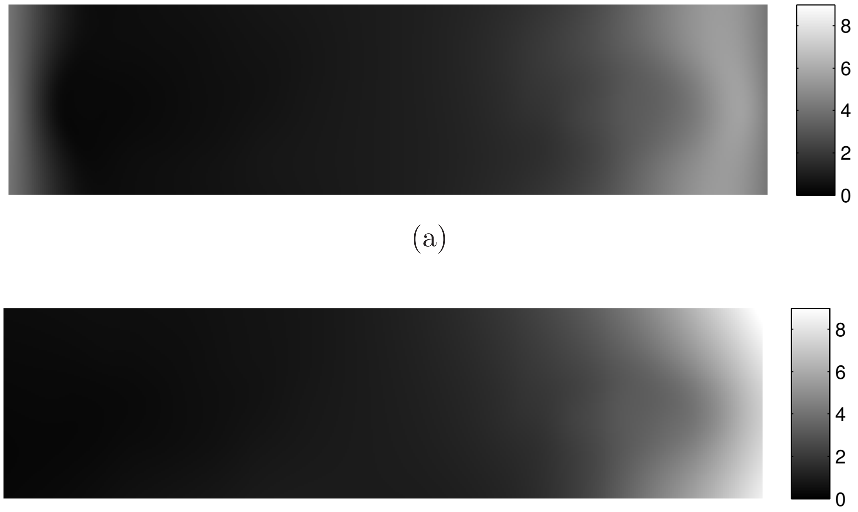

(a)

(b)

Figure 3.33: The sensitivity estimates for the breast phantom data generated using a (a) circulant matrix and a (b) non-circulant matrix.

lant finite differencing matrices. As with the brain data, there is an unrealistic dip in the estimate near the coil and a rise at the opposing boundary when using the circulant finite differencing matrix, Fig. 3.33(a). The estimate using the non-circulant matrix is more realistic, Fig. 3.33(b).

Fig. 3.34(a – b) presents the same estimates as Fig. 3.33, but masked in the spatial domain to highlight the error over the object support. Fig. 3.34(c) shows the difference between these two estimates. From these images, we see that the inaccuracies in the estimate at the boundaries of the image caused by the circulant finite differencing matrix propagated to within the object support. Thus, the need for a non-circulant finite differencing matrix is also evident for the case of real data.

Discussion

Using a finite differencing matrix for the case of periodic boundary conditions in our experiments caused substantial errors at the boundaries of the field-of-view and these propagated to within the object support. Furthermore, padding the images

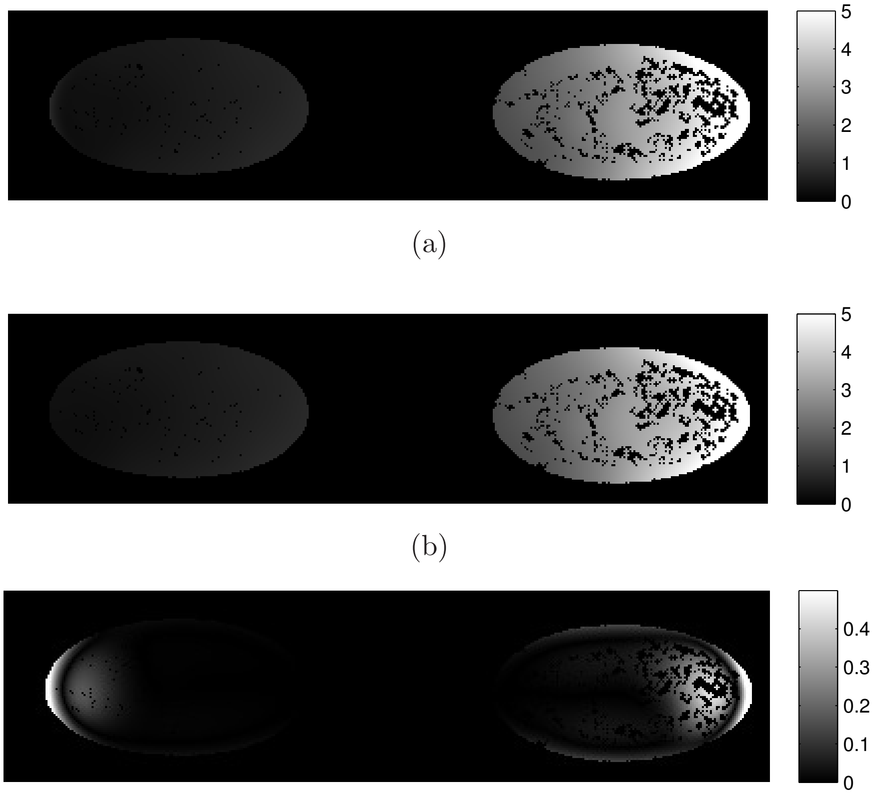

(a)

(b)

(c)

Figure 3.34: The same sensitivity estimates for the breast data as in Fig. 3.33 but masked to highlight the error over the object support. The difference between the estimates in (a) and (b) is presented in (c).

did not entirely mitigate the error. As seen in Section 3.2.4, these types of errors can cause significant artifacts in SENSE reconstructions. However, using a finite differencing matrix for the case of non-periodic boundary conditions avoided these errors. Thus, these experiments illustrate the need to use a non-circulant finite differencing matrix in the regularized estimator of (3.1).

3.2.7 Conclusions

In this section, we have presented two additional AL based minimization methods as well as analyzed several key aspects of the regularized sensitivity profile estimator. First, we derived an ADMM based minimization algorithm that does not use the finite differencing matrix reformulation from Section 3.1. Although this method was typically faster than the existing CG algorithms, it was significantly slower than our ADMM–Circ approach. We also presented an alternative variable splitting that leads to an AL algorithm with similar performance to ADMM-Circ but with no guarantee of convergence. We validated the improved accuracy of regularized sensitivity profile estimation methods over heuristic, low-resolution estimators by comparing SENSE reconstructions created using their estimates. We showed a type of SENSE reconstruction artifact that can occur when using tight sensitivity profile estimation masks. Finally, we demonstrated the need for a finite differencing matrix with non-periodic boundary conditions in our cost function, and thus, the importance of our proposed variable splitting.

Bibliography

[1] M. J. Allison, S. Ramani, and J. A. Fessler, “Accelerated regularized estimation of MR coil sensitivities using augmented Lagrangian methods,” *IEEE Trans. Med. Imag.*, vol. 32, no. 3, pp. 556–64, Mar. 2013.

[2] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, “SENSE: sensitivity encoding for fast MRI,” *Mag. Res. Med.*, vol. 42, no. 5, pp. 952–62, Nov. 1999.

[3] D. K. Sodickson and W. J. Manning, “Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays,” *Mag. Res. Med.*, vol. 38, no. 4, pp. 591–603, Oct. 1997.

[4] J. Tsao, P. Boesiger, and K. P. Pruessmann, “k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations,” *Mag. Res. Med.*, vol. 50, no. 5, pp. 1031–42, Nov. 2003.

[5] S. O. Schönberg, O. Dietrich, M. F. Reiser, and A. L. Baert, *Parallel Imaging in Clinical MR Applications*, Medical Radiology Diagnostic Imaging Series. Springer, 2007.

[6] C. R. Meyer, P. H. Bland, and J. Pipe, “Retrospective correction of intensity inhomogeneities in MRI,” *IEEE Trans. Med. Imag.*, vol. 14, no. 1, pp. 36–41, Mar. 1995.

[7] S. L. Keeling and R. Bammer, “A variational approach to magnetic resonance coil sensitivity estimation,” *Applied Mathematics and Computation*, vol. 158, no. 2, pp. 53–82, Nov. 2004.

[8] D. J. Larkman and R. G. Nunes, “Parallel magnetic resonance imaging,” *Phys. Med. Biol.*, vol. 52, no. 7, pp. R15–R55, Apr. 2007.

[9] R. Bammer, R. Stollberger, M. J. Augustin, T. Seifert, S. Strasser-Fuchs, P. Wach, H. P. Hartung, and F. Fazekas, “Parallel imaging strategies for high-speed magnetic resonance diffusion imaging,” in *Proc. SPIE 3978 Medical Imaging 2000: Physiology and Func. from Multidim. Im.*, 2000, pp. 12–22.

[10] L. Ying and J. Sheng, “Joint image reconstruction and sensitivity estimation in SENSE (JSENSE),” *Mag. Res. Med.*, vol. 57, no. 6, pp. 1196–1202, June 2007.

[11] P. Vemuri, E. G. Kholmovski, D. L. Parker, and B. E. Chapman, “Coil sensitivity estimation for optimal SNR reconstruction and intensity inhomogeneity correction in phased array MR imaging,” in *Information Processing in Medical Im.*, 2005, pp. 603–14.

[12] F. H. Lin, Y. J. Chen, J. W. Belliveau, and L. L. Wald, “A wavelet-based approximation of surface coil sensitivity profiles for correction of image intensity

inhomogeneity and parallel imaging reconstruction," *Hum. Brain Map.*, vol. 19, no. 2, pp. 96–111, Mar. 2003.

[13] M. A. G. Ballester, Y. Machida, Y. Kassai, Y. Hamamura, and H. Sugimoto, "Robust estimation of coil sensitivities for RF subencoded acquisition techniques," in *Proc. Intl. Soc. Mag. Res. Med.*, 2001, p. 799.

[14] J. Jin, F. Liu, E. Weber, Y. Li, and S. Crozier, "An electromagnetic reverse method of coil sensitivity mapping for parallel MRI - theoretical framework.," *J. Magn. Reson.*, vol. 207, no. 1, pp. 59–68, 2010.

[15] R. Bammer, M. Auer, S. L. Keeling, M. Augustin, L. A. Stables, R. W. Prokesch, R. Stollberger, M. E. Moseley, and F. Fazekas, "Diffusion tensor imaging using single-shot SENSE-EPI," *Mag. Res. Med.*, vol. 48, no. 1, pp. 128–36, July 2002.

[16] F. Huang, Y. Chen, G. R. Duensing, J. Akao, A. Rubin, and C. Saylor, "Application of partial differential equation-based inpainting on sensitivity maps," *Mag. Res. Med.*, vol. 53, no. 2, pp. 388–97, Feb. 2005.

[17] L. L. Wald, "Parallel Imaging Update: How Many Elements Do We Need?," in *Proc. Intl. Soc. Mag. Res. Med.*, 2006, p. 202.

[18] M. J. Allison and J. A. Fessler, "An augmented Lagrangian method for MR coil sensitivity estimation," in *Proc. Intl. Soc. Mag. Res. Med.*, 2011, p. 2881.

[19] D. P. Bertsekas, "Multiplier methods: A survey," *Automatica*, vol. 12, no. 2, pp. 133–45, Mar. 1976.

[20] L. M. Bregman, "The relaxation method for finding the common point of convex sets and its application to the solution of problems in convex programming," *USSR Comp Math and Math Phys*, vol. 7, pp. 200–17, 1967.

[21] T. Goldstein and S. Osher, "The split Bregman method for L1-regularized problems," *SIAM J. Imaging Sci.*, vol. 2, no. 2, pp. 323–43, 2009.

[22] S. Ramani and J. A. Fessler, "Parallel MR image reconstruction using augmented Lagrangian methods," *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 694–706, Mar. 2011.

[23] B. Liu, K. King, M. Steckner, J. Xie, J. Sheng, and L. Ying, "Regularized sensitivity encoding (SENSE) reconstruction using Bregman iterations," *Mag. Res. Med.*, vol. 61, no. 1, pp. 145–52, Jan. 2009.

[24] X. Ye, Y. Chen, and F. Huang, "Computational acceleration for MR image reconstruction in partially parallel imaging," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1055–63, May 2011.

[25] S. Ma, W. Yin, Y. Zhang, and A. Chakraborty, "An efficient algorithm for compressed MR imaging using total variation and wavelets," in *Proc. IEEE Conf. on Comp. Vision and Pattern Recognition*, 2008, pp. 1–8.

[26] M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems," *IEEE Trans. Im. Proc.*, vol. 20, no. 3, pp. 681–695, Mar. 2011.

[27] M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Fast image recovery using variable splitting and constrained optimization," *IEEE Trans. Im. Proc.*, vol. 19, no. 9, pp. 2345–56, Sept. 2010.

[28] S. Ramani and J. A. Fessler, "A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 31, no. 3, pp. 677–88, Mar. 2012.

[29] M. A. T. Figueiredo and J. M. Bioucas-Dias, "Restoration of Poissonian images using alternating direction optimization," *IEEE Trans. Im. Proc.*, vol. 19, no. 12, pp. 3133–45, Dec. 2010.

[30] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1-3, pp. 293–318, Apr. 1992.

[31] R. Glowinski and P. L. Tallec, *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*, Soc. Indust. Appl. Math., 1989.

[32] M. J. Allison, S. Ramani, and J. A. Fessler, "Regularized MR coil sensitivity estimation using augmented Lagrangian methods," in *Proc. IEEE Intl. Symp. Biomed. Imag.*, 2012, pp. 394–7.

[33] Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM J. Imaging Sci.*, vol. 1, no. 3, pp. 248–72, 2008.

[34] R. H. Chan, "Circulant preconditioners for Hermitian Toeplitz systems," *SIAM J. Matrix. Anal. Appl.*, vol. 10, no. 4, pp. 542–50, 1989.

[35] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans, "Design and construction of a realistic digital brain phantom," *IEEE Trans. Med. Imag.*, vol. 17, no. 3, pp. 463–8, June 1998.

[36] M. I. Grivich and D. P. Jackson, "The magnetic field of current-carrying polygons: An application of vector field rotations," *Amer. J. Phys.*, vol. 68, no. 5, pp. 469–74, May 2000.

[37] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, "The NMR phased array," *Mag. Res. Med.*, vol. 16, no. 2, pp. 192–225, Nov. 1990.

[38] C. A. McKenzie, E. N. Yeh, M. A. Ohliger, M. D. Price, and D. K. Sodickson, "Self-calibrating parallel imaging with automatic coil sensitivity extraction," *Mag. Res. Med.*, vol. 47, no. 3, pp. 529–38, Mar. 2002.

[39] A. K. Funai, J. A. Fessler, D. T. B. Yeo, V. T. Olafsson, and D. C. Noll, "Regularized field map estimation in MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1484–94, Oct. 2008.

[40] A. Funai, J. A. Fessler, W. Grissom, and D. C. Noll, "Regularized B1+ map estimation in MRI," in *Proc. IEEE Intl. Symp. Biomed. Imag.*, 2007, pp. 616–9.

[41] V. T. Olafsson, D. C. Noll, and J. A. Fessler, "Fast joint reconstruction of dynamic R_2^* and field maps in functional MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 9, pp. 1177–88, Sept. 2008.

CHAPTER IV

Main Magnetic Field Inhomogeneity Estimation

Estimating the main magnetic field inhomogeneity is a non-trivial but essential task for many magnetic resonance imaging techniques (see Section 2.2 for details). In this chapter, we explore regularized field map estimation for both multiple echo time field inhomogeneity estimation and water-fat imaging problems. Section 4.1 presents a generalized field map estimation cost function for both problems and proposes two novel minimization algorithms that reduce the computation time by over 30 times compared to the existing solutions.[1] Section 4.2 explores using edge preserving regularization to capture discontinuities in the magnetic field that can occur at tissue interfaces.

4.1 Accelerated Computation of Regularized Field Map Estimates

4.1.1 Introduction

Spatial inhomogeneity within the main magnetic field ($\mathbf{B}_0$) can degrade many magnetic resonance imaging (MRI) techniques. For instance, it can cause reconstruction artifacts particularly for scans with long readout times [2]. Field inhomogeneity is also a nuisance parameter in chemical shift based water-fat imaging

---

[1] This section is an extension of our earlier work on regularized field estimation [1].

techniques [3]. However, accurate estimates of the off-resonance frequency induced by the field inhomogeneity at each voxel (i.e., a field map) can mitigate both of these issues [2, 4–9].

Numerous methods have been proposed to estimate field maps. One approach is to acquire multiple scans at different echo times and then estimate the field inhomogeneity from the phase information in the resulting images [3]. Since field maps tend to be smooth within tissue [2], many estimates also enforce some form of smoothness requirement. There are several ways to do this including region growing techniques [10–15], filtering [16], curve fitting [17–19], multiresolution and subspace approaches [12, 19–21], and graph cut algorithms [22]. The drawbacks of those methods are that they use heuristic techniques or significant approximations in an attempt to correct for phase wrapping between the multiple acquisitions. As an alternative, regularized estimation methods such as [2, 23–25] estimate a smooth field map from multiple acquisition images while intrinsically accounting for phase wrapping. The disadvantage of these regularized methods is that they use nonconvex cost functions that require iterative minimization techniques.

An existing minimization technique for regularized field map estimation uses optimization transfer principles [26] to create a separable quadratic surrogate (SQS) [2, 23]. However, that method takes many iterations, and subsequently a long time, to reach a useful solution. This large computational cost impedes the adoption of these estimators. Other regularized field map estimation minimization techniques discretize the solution space [24, 25] and may require a second descent algorithm to produce sufficiently smooth estimates [25]. Similar cost functions appear in other medical imaging problems where the parameter of interest is contained within the phase of the cost function [15, 27–30]. Although certain minimization strategies work well for these specific problems (e.g., SQS methods [27,28], numerically implemented conjugate gradient method [29], and linearization techniques [9, 30]), they are not well suited for the regularized field map estimation problem.

This section presents two new methods that significantly decrease the computation time of the regularized field map estimators. We begin by extending the SQS method presented in [2] to a quadratic surrogate for the overall cost function. Our first method adapts Huber's algorithm for quadratic surrogates [31] by exploiting the structure of the Hessian matrix of the quadratic surrogate function. The second method modifies the nonlinear CG method by using the quadratic surrogate function to create an efficient monotonic line search as well as an effective preconditioning matrix. We then compare our methods with the existing SQS method [2] on both multiple echo time field map estimation and water-fat imaging data sets. We find that all of the methods converge to similar solutions for these data sets and that our fastest algorithms do so in less than a thirtieth of the time.

Section 4.1.2 presents a general form for the regularized field map estimation cost functions from [2,23] and derives our two novel minimization methods. Section 4.1.3 compares the performance of our proposed methods to the existing SQS method using both multiple echo time field map estimation and water-fat imaging data. Section 4.1.4 discusses the results and other key aspects of the algorithms, while Section 4.1.5 concludes.

4.1.2 Materials and Methods

This section proposes two quadratic surrogate based methods for faster computation of regularized field map estimates.

Generalized cost function formulation

The regularized field map estimation techniques for both multiple echo time field map estimation and water-fat imaging use similar image models [2, 23, 24]:

$$y_j^\ell = e^{i\omega_j t_\ell} x_j^\ell + \epsilon_j^\ell \tag{4.1}$$

where $y_j^\ell \in \mathbb{C}$ denotes the jth voxel in the reconstructed image of the ℓth scan, $\omega_j \in \mathbb{R}$ denotes the field map value at voxel j, t_ℓ denotes the echo time shift of the ℓth scan, and $\epsilon_j^\ell \in \mathbb{C}$ denotes the noise. x_j^ℓ has the following problem-dependent definitions:

$$x_j^\ell \triangleq \begin{cases} m_j & \text{for multiple echo time field map estimation} \\ w_j + f_j e^{i2\pi\Delta_f t_\ell} & \text{for water-fat imaging} \end{cases} \tag{4.2}$$

where $m_j, w_j, f_j \in \mathbb{C}$ denote the true magnetization, water component, and fat component at voxel j respectively, and Δ_f is the known frequency shift of fat relative to water (Hz).[2]

Assuming that we have zero-mean, white complex Gaussian noise ϵ_j^ℓ, the joint maximum-likelihood (ML) estimates are

$$\underset{\mathbf{w},\mathbf{x}}{\arg\min} \sum_{j=1}^{N}\sum_{l=1}^{L} \left| y_j^l - e^{i\omega_j t_\ell} x_j^\ell \right|^2 \tag{4.3}$$

where N and L are the number of voxels and scans respectively, $\mathbf{w}$ is a vector containing the field map variables, and $\mathbf{x}$ is a vector containing either the magnetization m_j or water-fat w_j, f_j variables as outlined in (4.2) [2, 23, 24].

Assuming the field map $\mathbf{w}$ is known, the ML estimates (4.3) yield analytic solutions for the unknown image variables found in x_j^ℓ. Substituting the solution back into (4.3) and simplifying yields the following negative log-likelihood for field map

---

[2] A multipeak fat model such as in [32] would require only minor changes to our algorithms.

estimation [2, 33]:

$$\Phi_F(\mathbf{w}) \triangleq \sum_{j=1}^{N} \sum_{m=1}^{L} \sum_{p=1}^{L} \varphi_{jmp}(\omega_j) \tag{4.4}$$

where

$$\varphi_{jmp}(\omega_j) = |g_{jmp}| \cdot [1 - \cos(\omega_j(t_m - t_p) + \angle g_{jmp})] \tag{4.5}$$

and $g_{jmp} = [\mathbf{G}_j]_{(m,p)}$ where (m, p) denotes the matrix indices and $\mathbf{G}_j = \text{diag}\{\underline{y}_j\}^{\text{H}} \mathbf{A}_j (\mathbf{A}_j^{\text{H}} \mathbf{A}_j)^{-1} \mathbf{A}_j^{\text{H}} \text{diag}\{\underline{y}_j\}$ with $\underline{y}_j = [y_j^1, \ldots, y_j^L]^{\text{T}}$. For multiple echo time field map estimation [2], we use the approximation $\mathbf{A}_j = [|y_j^1|, \ldots, |y_j^L|]^{\text{T}}$. For water-fat imaging [33], we use the following $L \times 2$ system matrix for each voxel:

$$\mathbf{A}_j = \begin{bmatrix} 1 & e^{i2\pi\Delta_f t_1} \\ 1 & e^{i2\pi\Delta_f t_2} \\ \vdots & \vdots \\ 1 & e^{i2\pi\Delta_f t_L} \end{bmatrix}. \tag{4.6}$$

Field maps tend to be smooth within body tissue [2], therefore a spatial regularizing term is added to the ML estimate to obtain a penalized-likelihood (PL) estimate [2, 33]. Furthermore, field map estimates are only needed for voxels where signal is present. We therefore incorporate an estimation mask yielding the final generalized PL cost function:

$$\Psi(\boldsymbol{\omega}_s) = \Phi(\boldsymbol{\omega}_s) + \frac{\beta}{2} \|\mathbf{C}\boldsymbol{\omega}_s\|_2^2, \tag{4.7}$$

where

$$\Phi(\boldsymbol{\omega}_s) \triangleq \sum_{j \in N_s} \sum_{m=1}^{L} \sum_{p=1}^{L} \varphi_{jmp}(\omega_j), \tag{4.8}$$

$\boldsymbol{\omega}_s \in \mathbb{R}^{N_s}$ is a vector containing the field map variables within the estimation mask,

N_s is the set of voxels within the mask, β is a regularization parameter, and $\mathbf{C}$ is a finite differencing matrix that accounts for the mask.[3]

Estimators based on minimizing this cost function provide accurate field map estimates [2, 23]; however, the cost function is difficult to minimize due to the nonlinear data-fit terms and the non-separability of the regularizer. Current minimization strategies use separable quadratic surrogate methods that can take many minutes to converge for 2-D images [1, 2, 23]. The remainder of this section proposes two new minimization strategies that reduce the computation time.

Quadratic surrogate function

Following work on similar cost functions [2, 23, 27], we use optimization transfer principles [26] to derive our minimization strategies. In particular, we use the same parabolic surrogate functions as in [2] to majorize the sinusoidal data-fit terms φ_{jmp}:

$$\varphi_{jmp}(\omega_j) \leq \varphi_{jmp}(\omega_j^{(n)}) + \dot{\varphi}_{jmp}(\omega_j^{(n)}) \left(\omega_j - \omega_j^{(n)}\right) + \frac{1}{2}\kappa_{jmp}\left(s_{jmp}(\omega_j^{(n)})\right)\left(\omega_j - \omega_j^{(n)}\right)^2 \tag{4.9}$$

where $\omega_j^{(n)}$ denotes the current field map estimate at voxel j and

$$\kappa_{jmp}(s) \triangleq |g_{jmp}|\,(t_m - t_p)^2\,\frac{\sin(s)}{s} \tag{4.10}$$

with

$$s_{jmp}(\omega) \triangleq (\omega \cdot (t_m - t_p) + \angle g_{jmp}) \bmod \pi. \tag{4.11}$$

are the optimal Huber's curvatures [2, 31].[4]

---

[3] $\mathbf{C}$ can be described in terms of a full FOV finite differencing matrix $\mathbf{C}_F$ as $\mathbf{C} = \mathbf{B} \cdot \mathbf{C}_F \cdot \mathbf{M}$ where $\mathbf{B}$ a diagonal matrix with the binary mask along the diagonal and $\mathbf{M} \in \{0,1\}^{N \times |N_s|}$ is a tall matrix that places the masked voxels into their correct positions within the full FOV.

[4] With Huber's curvatures, the functions satisfy the following conditions: the surrogate functions are differentiable, they are symmetric, and their curvatures are bounded and monotone non-increasing for $s > 0$. It can therefore be concluded that these surrogate functions have the optimal curvatures of all parabolic surrogate functions for (4.8) [2, 31].

We substitute these separable surrogate functions into the data-fit term (4.8) and, unlike in [2], vectorize the problem to obtain the following quadratic surrogate function for (4.7):

$$\Phi(\boldsymbol{\omega}_s^{(n)}) + \nabla\Phi(\boldsymbol{\omega}_s^{(n)})^{\mathrm{T}}(\boldsymbol{\omega}_s - \boldsymbol{\omega}_s^{(n)}) + \frac{1}{2}(\boldsymbol{\omega}_s - \boldsymbol{\omega}_s^{(n)})^{\mathrm{T}}\mathbf{D}^{(n)}(\boldsymbol{\omega}_s - \boldsymbol{\omega}_s^{(n)}) + \frac{1}{2}\beta\boldsymbol{\omega}_s^{\mathrm{T}}\mathbf{C}^{\mathrm{T}}\mathbf{C}\boldsymbol{\omega}_s \tag{4.12}$$

where ∇ denotes the gradient,

$$\mathbf{D}^{(n)} = \mathrm{diag}\{d_j^{(n)}\} \quad \text{with} \quad d_j^{(n)} = \sum_{m=1}^{L}\sum_{p=1}^{L} \kappa_{jmp}\left(s_{jmp}(\omega_j^{(n)})\right), \tag{4.13}$$

and $\boldsymbol{\omega}_s^{(n)}$ denotes a vector containing the current field map estimates.

The Hessian matrix of this surrogate cost function is $\mathbf{H}_n = \mathbf{D}^{(n)} + \beta\mathbf{C}^{\mathrm{T}}\mathbf{C}$ and it is positive definite as long as at least one value of $d_j^{(n)}$ is positive (which is true for any non-trivial problem). Thus, the quadratic surrogate function (4.12) has a unique, finite minimizer and we use it to derive our two novel minimization strategies.

Method 1: Huber's algorithm for quadratic surrogates

Although the Hessian matrix of (4.12) is large for typical image sizes ($\mathbb{R}^{|N_s|\times|N_s|}$ where $|N_s|$ is the number of voxels within the mask), it is positive definite and has a sparse banded structure. Sparse Cholesky factorization can therefore be used to efficiently solve a linear system of equations based on $\mathbf{H}_n$ (i.e., solving $\mathbf{H}_n^{-1}\mathbf{x}$ for arbitrary $\mathbf{x}$) [34]. Exploiting this efficiency, we use Huber's algorithm for quadratic surrogates [31] to obtain the following iterative algorithm for monotonically decreasing the original cost function (4.7):

$$\boldsymbol{\omega}_s^{(n+1)} = \boldsymbol{\omega}_s^{(n)} - \mathbf{H}_n^{-1}\nabla\Psi(\boldsymbol{\omega}_s^{(n)}). \tag{4.14}$$

Unlike traditional Newton's method which uses the Hessian of the original cost function, this minimization strategy guarantees monotonic convergence.

Method 2: Nonlinear conjugate gradient with a monotonic line search

We alter the nonlinear CG (NCG) method by considering two modifications, both based on our quadratic surrogates (4.12). First, we derive a monotonic step size line search algorithm using quadratic surrogates like in [35]. Second, we consider several quadratic surrogate based preconditioners that may change with iteration.

(1) Monotonic step size line search: The non-quadratic nature of the cost function (4.7) prevents direct computation of a step size. Instead, we must consider iterative line search methods. There are many existing line search methods capable of determining a 'sufficient' step size [36]. For example, a backtracking line search incorporating the Armijo rule [37] is a popular choice in medical imaging. The disadvantage of many of these methods is that they require multiple costly evaluations of the original cost function and they have parameter values that must be carefully selected given the nonconvex nature of our problem. Instead, we follow [35] and use a line search method based on Huber's algorithm for quadratic surrogates. This particular line search method is guaranteed to monotonically decrease the cost function.

To create the monotonic line search algorithm, we evaluate the original cost function (4.7) with respect to a scalar step size variable, α:

$$f(\alpha) = \Phi(\boldsymbol{\omega}_s + \alpha \mathbf{z}) + \frac{1}{2}\beta||\mathbf{C}(\boldsymbol{\omega}_s + \alpha \mathbf{z})||_2^2, \tag{4.15}$$

where $\mathbf{z} \in \mathbb{R}^{N_s}$ is the search direction.

We are again confronted with the fact that the data-fit term is nonconvex. Since this cost function has the same structure as (4.7), we follow the approach in Section 4.1.2 to create the following 1-D quadratic surrogate function:

$$\Phi(\boldsymbol{\omega}_s + \alpha^{(n)}\mathbf{z}) + \mathbf{z}^{\mathrm{T}}\nabla\Phi(\boldsymbol{\omega}_s + \alpha^{(n)}\mathbf{z})(\alpha - \alpha^{(n)}) + \frac{1}{2}d_\alpha^{(n)}(\alpha - \alpha^{(n)})^2 + \frac{1}{2}\beta||\mathbf{C}(\boldsymbol{\omega}_s + \alpha\mathbf{z})||_2^2, \tag{4.16}$$

where

$$d_\alpha^{(n)} = \sum_{j \in N_s}\sum_{m=1}^{L}\sum_{p=1}^{L}|z_j|^2\kappa_{jmp}\left(s_{jmp}(\omega_j + \alpha^{(n)}z_j)\right). \tag{4.17}$$

As in Section 4.1.2, we apply Huber's algorithm for quadratic surrogates to obtain the final monotonic line search algorithm:

$$\alpha^{(n+1)} = \alpha^{(n)} - \frac{\frac{\partial}{\partial\alpha}f(\alpha^{(n)})}{d_\alpha^{(n)} + \beta\mathbf{z}^{\mathrm{T}}\mathbf{C}^{\mathrm{T}}\mathbf{C}\mathbf{z}} \tag{4.18}$$

for which $\beta\mathbf{z}^{\mathrm{T}}\mathbf{C}^{\mathrm{T}}\mathbf{C}\mathbf{z}$ only needs to be computed once.

(2) Preconditioning matrices: We also explored several preconditioners to accelerate our CG based algorithm. Since the Hessian of the quadratic surrogate from Section 4.1.2 is positive definite and already computed during the monotonic line search, we use it to form our preconditioning matrices instead of the Hessian of the original cost function (4.7). Our first preconditioning matrix is the diagonal matrix, $\mathbf{P}_{\mathrm{D}} = \mathrm{diag}\{\mathbf{H}_n\}$, which can be efficiently implemented. Our second preconditioner is the full Hessian matrix of the quadratic surrogate $\mathbf{P}_{\mathrm{H}} = \mathbf{H}_n$ which is implemented using sparse Cholesky factorization like in Section 4.1.2.

4.1.3 Results

This section compares our two novel algorithms to the existing SQS minimization method using both multiple echo time field map estimation and water-fat imaging data sets.

For all of our experiments, we followed [2] and first normalized the data by dividing each image by $\mathbf{median}\{\sqrt{r_j}\}$ where $r_j \triangleq \sum_{m=1}^{L}\sum_{p=1}^{L}|g_{jmp}|(t_m - t_p)^2$ for the object voxels. This normalization accounts for some of the R_2^* effects in the multiple echo time field estimation data and simplifies selecting the regularization parameter β [2].

We performed field map estimation with our Huber's algorithm based method (4.14) (QS-Huber), our NCG with monotonic line search algorithm (4.18) (NCG-MLS), our NCG-MLS algorithm with the diagonal preconditioner (NCG-MLS-D), our NCG-MLS algorithm with the Hessian preconditioner (NCG-MLS-H), and with the existing SQS algorithm [2]. $\mathbf{C}$ was chosen to be a second-order finite differencing matrix as this provided high quality estimates for each experiment. We used one line search iteration for NCG-MLS and NCG-MLS-D and three for NCG-MLS-H.[5] These parameter values were determined empirically and were not further optimized for each experiment to demonstrate the robustness of the algorithms.

For each experiment, all of the algorithms were initialized with the same image. The computation of these images differed between the multiple echo time field map estimation and water-fat imaging data and is therefore outlined in the corresponding subsections. Since we have no information about the field map outside of the object voxels, we restricted the estimates to masked regions. The masks were found by taking a convex hull around those voxels with significant signal. We dilated the masks by two voxels to avoid reconstruction artifacts that would otherwise arise if the patient was to move between calibration and acquisition scans. All of the algorithms were implemented in MATLAB (The MathWorks, Natick, MA, USA) and the experiments were run on a PC with a 2.66 GHz, quad-core Intel Xeon CPU.

Although not guaranteed, we found that all of the methods converged to within machine precision of the same solution for each experiment when using our given

---

[5] We determined the conjugate gradient direction using the Polak-Ribiére method [38] as it typically converges faster on non-quadratic problems than the Fletcher-Reeves approach [39], [36].

initialization strategies. We therefore compared the convergence rates of each algorithm by computing the root-mean squared difference (RMSD) between the estimate at each iteration $\boldsymbol{\omega}_s^{(n)}$ and the mean of the final estimates from the two fastest methods (QS-Huber and NCG-MLS-H) $\hat{\boldsymbol{\omega}}_s$:

$$\text{RMSD}(\boldsymbol{\omega}_s^{(n)}) = \sqrt{\frac{\|\boldsymbol{\omega}_s^{(n)} - \hat{\boldsymbol{\omega}}_s\|_2^2}{|N_s|}}. \tag{4.19}$$

Using the mean of two estimates avoids favoring any one of the algorithms.

Multiple echo time field map estimation data

We initialized our multiple echo time field map estimation data using a tightly masked conventional estimate. This method takes the phase difference of the first two acquired images as follows [2]:

$$\hat{\omega}_j^{\text{conv}} = \angle\left(y_j^{1*} y_j^2\right)/t_1. \tag{4.20}$$

Although this conventional estimator does not explicitly account for phase wrapping and typically produces noisy estimates, it was sufficient for our experiments in that it has low computational complexity and its estimates resulted in all of the algorithms converging to within machine precision of the same solution.

We simulated a multiple echo time field map estimation data set using a 128×128 pixel brain image and field map acquired on a 3T GE scanner, Fig. 4.1. We combined the field map with the brain image to create three acquisition images with relative echo times $t_\ell = 0, 2, 10\,\text{ms}$ and $\text{R}_2^* = 20\,\text{s}^{-1}$. Complex Gaussian noise was also added to these images so that the SNR $\approx 20\,\text{dB}$. Fig. 4.2 presents the magnitude and phase of the simulated scan images.

We selected the regularization parameter $\beta = 2^{-3}$ as this was found to provide the most accurate estimates compared to the truth (details not shown). Fig. 4.3 shows

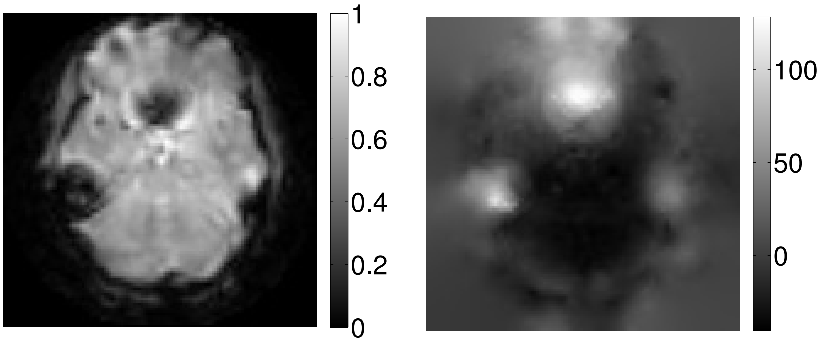

Figure 4.1: The magnitude of the brain image (left) and the “true” field map in Hz (right) used to create the simulated multiple echo time field map estimation data.

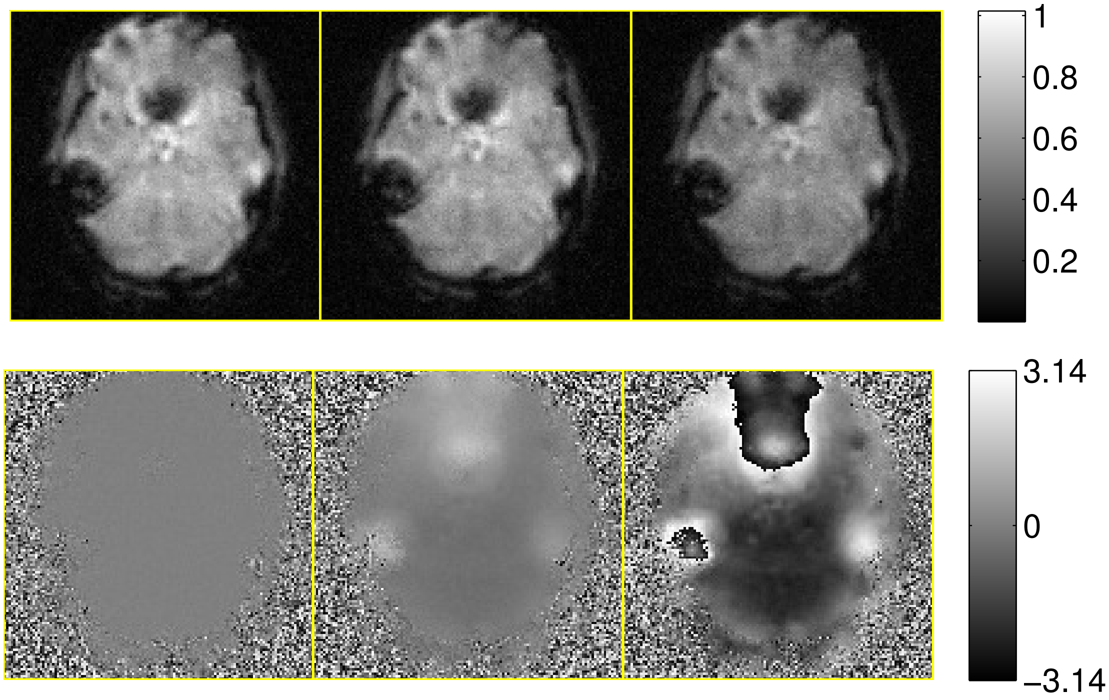

Figure 4.2: Simulated magnitude (top) and phase (bottom) images representing three acquisitions with relative echo times $t_\ell = 0, 2, 10\,\text{ms}$ (from left to right).

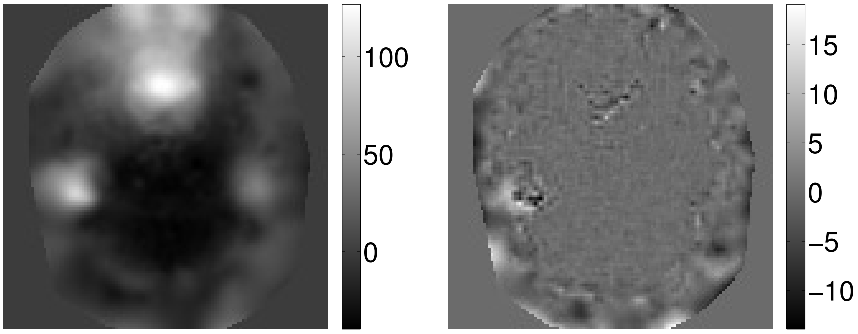

Figure 4.3: The final estimate (left) and its difference to the masked truth (right) in Hz for the multiple echo time field map estimation data set.

Table 4.1: Convergence Time and Iterations for the Multiple Echo Time Field Map Estimation Data Set

Method	**Time (s)**		**Iterations**	
	1 Hz	**1 Rad/s**	**1 Hz**	**1 Rad/s**
SQS [2]	16	48	3800	11400
NCG-MLS	1.1	1.8	98	180
NCG-MLS-D	0.9	1.5	83	150
NCG-MLS-H	0.5	0.9	2	4
QS-Huber	0.5	1.1	2	6

the regularized estimate and its difference to the masked truth for one representative method (QS-Huber). Fig. 4.4 plots the RMSDs in Hz versus time for all of the evaluated methods. Table 4.1 presents the time and number of iterations required to reach an RMSD of 1 Hz and 1 Rad/s for each algorithm.

Water-fat imaging data

We evaluated the algorithms on two water-fat imaging data sets. The first uses simulated images for which we know the truth. The second is an in-vivo knee data set.

The conventional estimator used to initialize the experiments in Section 4.1.3 is unsuitable for the case of water-fat imaging as it does not account for chemical shift due to fat. Instead, we used an initialization method proposed by [23, 24] in which

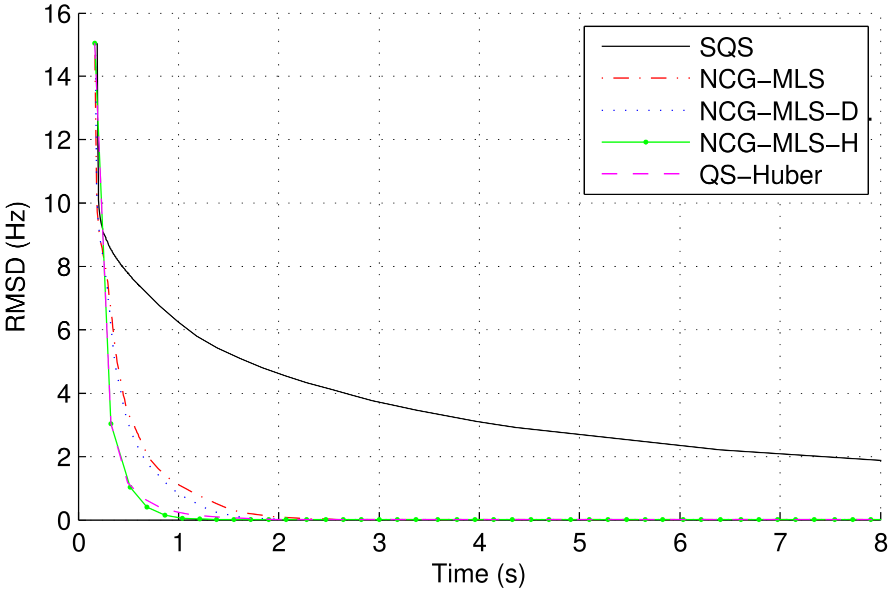

Figure 4.4: Plots of the RMSD in Hz versus computation time for all of the algorithms evaluated on the multiple echo time field map estimation data set.

we evaluated (4.8) over a discrete set of field map values spaced 2 Hz from $-|\Delta_f/2|$ to $|\Delta_f/2|$ [23]. Although this method generates noisy estimates, it accounts for the chemical shift due to fat, has low computational cost, and led to all of the algorithms converging to within machine precision of the same solution.

We also estimated the fat and water images using the proposed regularized field map estimates to further validate their accuracy. The water and fat estimation was performed voxel-wise using the following maximum-likelihood estimator [24, 33]:

$$\begin{bmatrix} w_j \\ f_j \end{bmatrix} = \left[\mathbf{A}_j^{\mathrm{H}} \mathbf{A}_j\right]^{-1} \mathbf{A}_j^{\mathrm{H}} \mathbf{W}(\omega_j)^{\mathrm{H}} \underline{y}_j, \tag{4.21}$$

where $\mathbf{W}(\omega_j) \triangleq \operatorname{diag}\{e^{i\omega_j t_1}, \ldots, e^{i\omega_j t_L}\}$.

(1) Simulated data: We simulated a water-fat imaging data set by combining the 256 × 256 pixel water and fat images presented in Fig. 4.5. We integrated a simulated field map (Fig. 4.5) with the resulting water-fat image to create three acquisition images with relative echo times $t_\ell = -0.4, 1.2, 2.8$ ms, $\Delta f = 220$ Hz, and

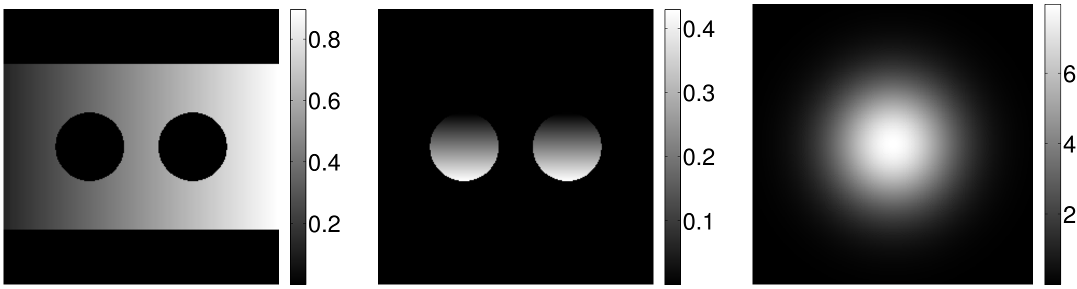

Figure 4.5: The magnitude of the water image (left), the fat image (center), and true field map in Hz (right) used to create the simulated water-fat data set.

Table 4.2: Convergence Time and Iterations for the Simulated Water-Fat Data Set

Method	**Time (s)**		**Iterations**	
	1 Hz	**1 Rad/s**	**1 Hz**	**1 Rad/s**
SQS [2]	280	2700	22600	218000
NCG-MLS	8.2	26	210	700
NCG-MLS-D	8.0	26	200	690
NCG-MLS-H	1.1	1.5	1	2
QS-Huber	1.0	1.5	1	2

$\mathrm{R}_2^* = 20\,\mathrm{s}^{-1}$. Complex noise was also added to these images so that the SNR $\approx 20\,\mathrm{dB}$. Fig. 4.6 shows the magnitude and phase of the final simulated images.

We selected the regularization parameter $\beta = 2^9$ as this was found to provide the most accurate estimates compared to the truth (details not shown). Fig. 4.7 shows the regularized estimate and its difference to the masked truth for one representative method (QS-Huber). Fig. 4.8 shows the resulting water and fat images generated using (4.21), while Fig. 4.9 plots the RMSDs in Hz versus time for all of the evaluated methods. Table 4.2 presents the time and number of iterations required to reach an RMSD of 1 Hz and 1 Rad/s for each algorithm.

(2) In-vivo knee data: We also evaluated the algorithms on a water-fat imaging knee data set acquired on a 1.5 T scanner using the IDEAL imaging acquisition protocol [15]. This data set consisted of three 256×256 pixel acquisition images with relative echo times $t_\ell = -0.4, 1.2, 2.8\,\mathrm{ms}$ and SNR $\approx 35\,\mathrm{dB}$. Fig. 4.10 presents

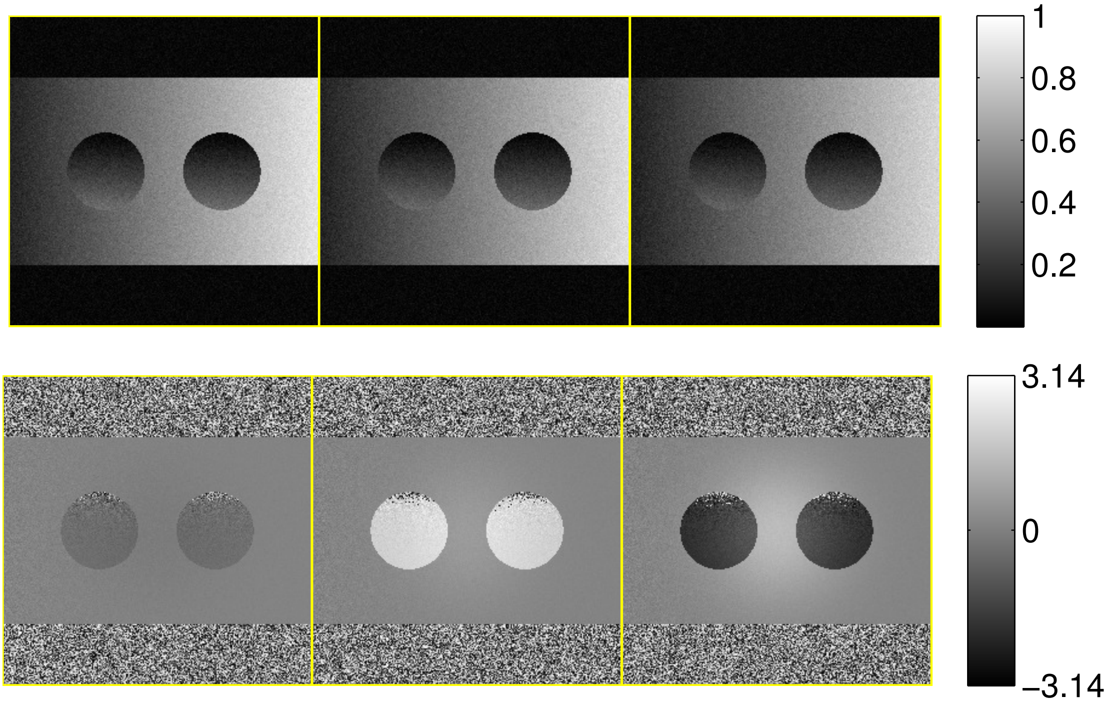

Figure 4.6: The magnitude (top) and phase (bottom) of the simulated water-fat scan images for relative echo times $t_\ell = -0.4, 1.2, 2.8\,\mathrm{ms}$ (from left to right).

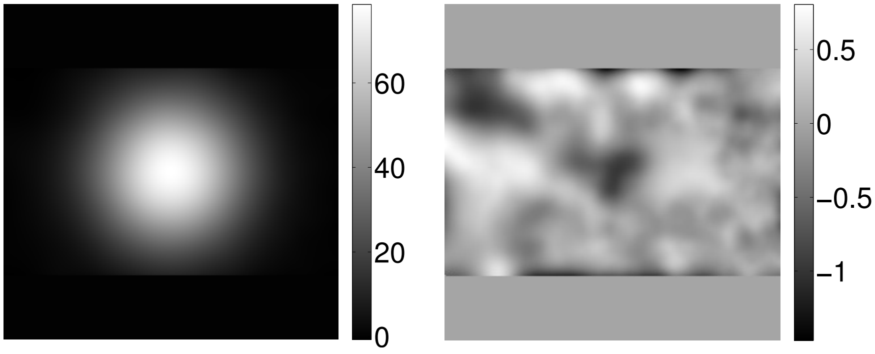

Figure 4.7: The final estimate (left) and its difference to the masked truth (right) in Hz for the simulated water-fat data set.

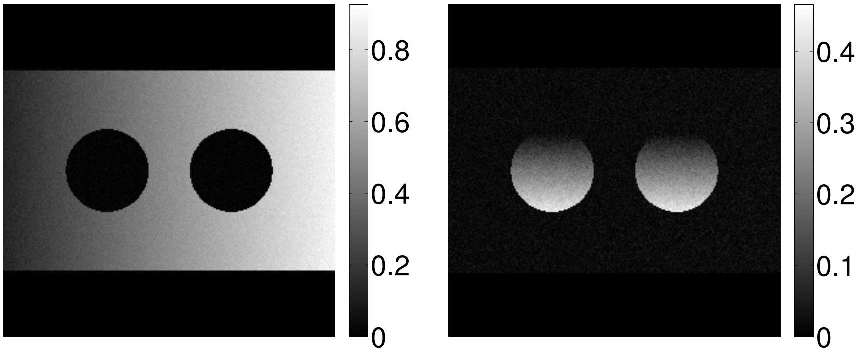

Figure 4.8: The resulting water image (left) and fat image (right) maximum-likelihood estimates computed using (4.21) and the field map in Fig. 4.7.

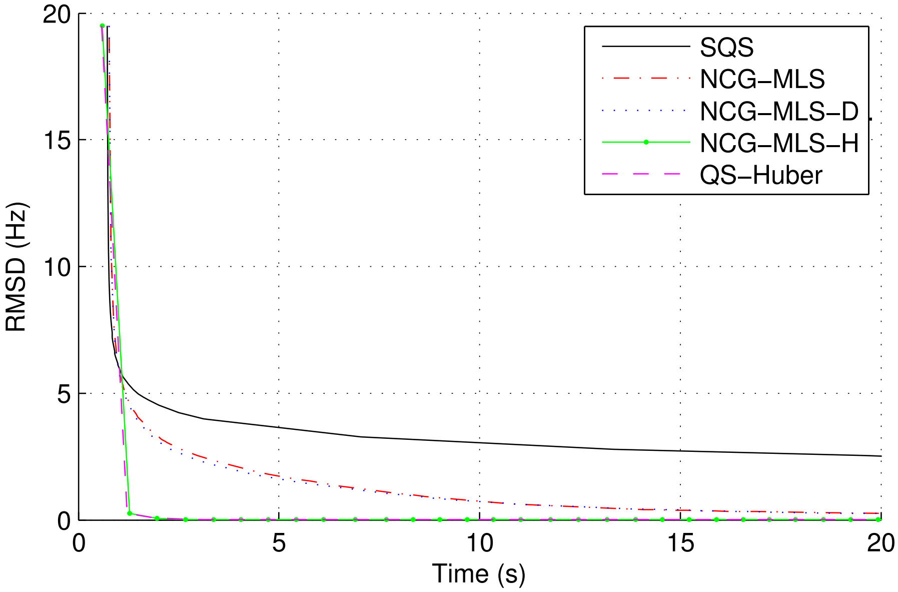

Figure 4.9: Plots of the RMSD in Hz versus time for all of the algorithms evaluated on the simulated water-fat data set.

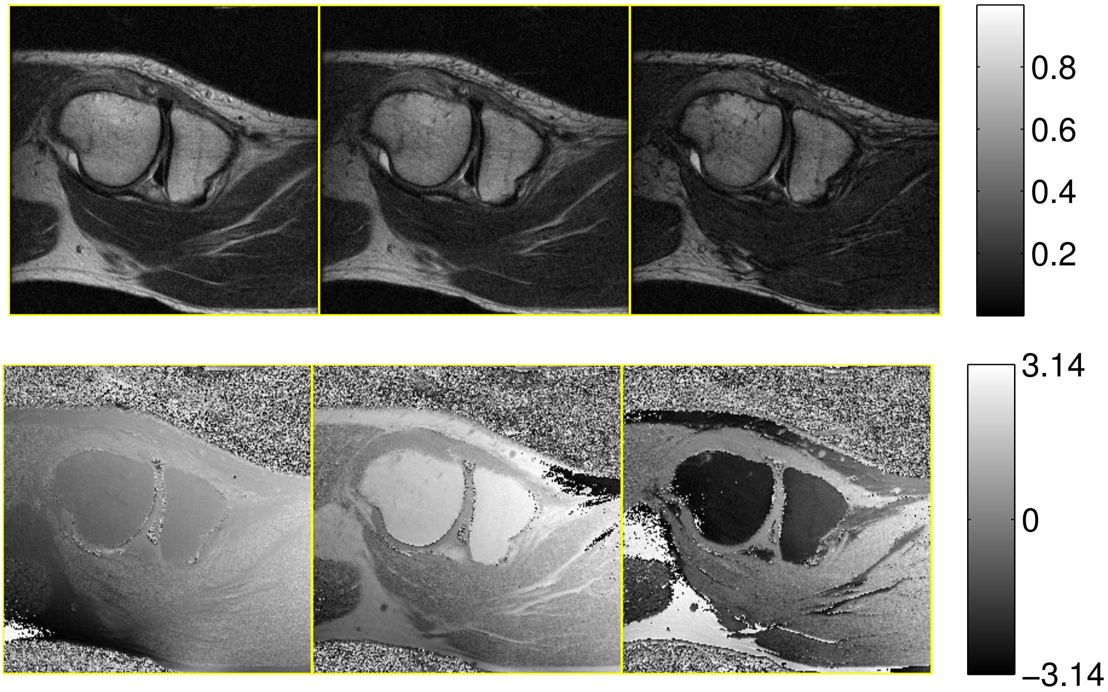

Figure 4.10: The magnitude (top) and phase (bottom) of the knee water-fat images with relative echo times $t_\ell = -0.4, 1.2, 2.8\,\mathrm{ms}$ (from left to right).

the magnitude and phase of the three scan images.

We selected the regularization parameter $\beta = 2^8$ as this provided high quality water and fat images. Fig. 4.11 shows the regularized estimate for one representative method (QS-Huber) and the resulting water and fat images generated using (4.21). Fig. 4.12 plots the RMSDs in Hz versus time for all of the evaluated methods. Table 4.3 presents the time and number of iterations required to reach an RMSD of 1 Hz and 1 Rad/s for each algorithm.

4.1.4 Discussion

Our novel regularized field map estimation methods were more than 30 times faster than the existing SQS method for multiple echo time field map estimation and more than 250 times faster for water-fat imaging data. Our fastest methods were those that used the exact Hessian at each iteration (i.e., QS-Huber and NCG-

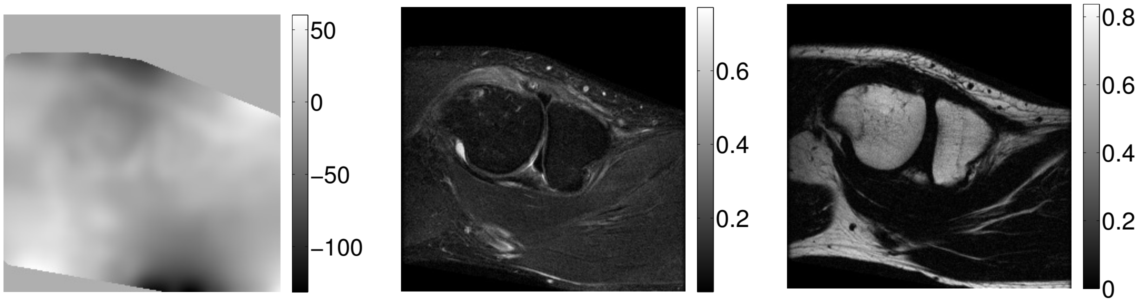

Figure 4.11: The regularized field map estimate in Hz (left), the resulting water image (center), and the resulting fat image (right) for the knee water-fat data set.

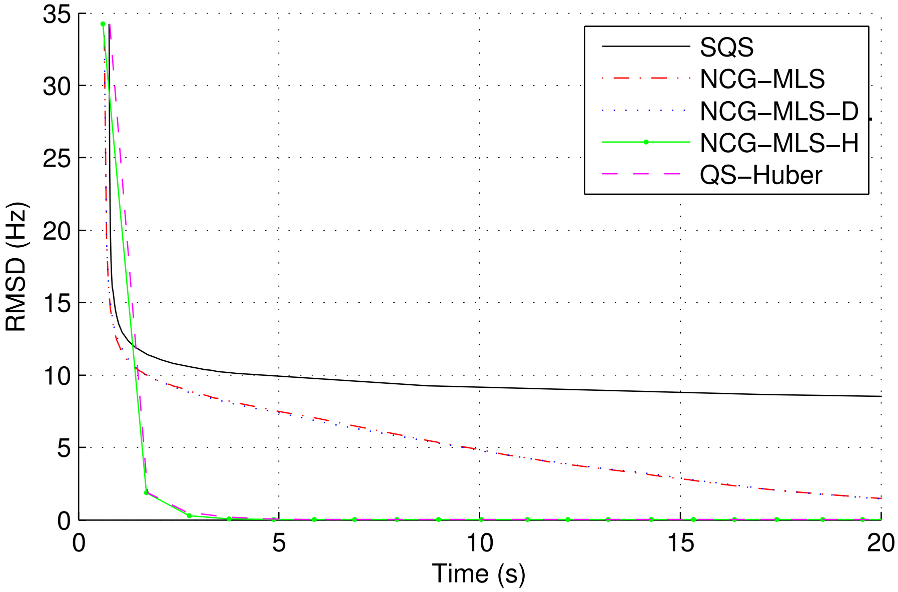

Figure 4.12: Plots of the RMSD in Hz versus time for all of the algorithms evaluated on the knee water-fat data set.

Table 4.3: Convergence Time and Iterations for the Knee Water-Fat Data Set

Method	**Time (s)**		**Iterations**	
	1 Hz	**1 Rad/s**	**1 Hz**	**1 Rad/s**
SQS [2]	2000	4400	126000	281000
NCG-MLS	23	39	400	700
NCG-MLS-D	24	42	420	740
NCG-MLS-H	2.0	3.1	2	3
QS-Huber	2.2	3.8	2	4

MLS-H). Our simpler NCG-MLS and NCG-MLS-D algorithms were slower than our fastest methods, but remained significantly faster than the SQS method. This is in part because the monotonic line search was able to find a near optimal step size in only a few iterations due to the similarity between the step size cost function (4.15) and our quadratic surrogate. The diagonal preconditioner did not drastically improved the NCG-MLS convergence rate and was actually a detriment for the in-vivo knee experiment. This preconditioner required additional line search iterations to realize its full per iteration acceleration; however, these savings were outweighed by the added computational cost.

Preconditioner computation could be reduced by computing the sparse Cholesky factorization only once and storing the resulting permutation matrix and the relatively sparse lower triangular matrix for subsequent iterations. The resulting preconditioning steps would require only permuted forward and backward substitutions making them faster, but with the potential drawback of using a non-updating preconditioner. We did not explore this variation as the current NCG-MLS-H algorithm already converges in only a few seconds.

Like our QS-Huber method, the existing SQS algorithm can be written as an application of Huber's algorithm [2],

$$\boldsymbol{\omega}_s^{(n+1)} = \boldsymbol{\omega}_s^{(n)} - \mathrm{diag}\left(\frac{1}{d_j^{(n)} + \beta c_j}\right) \nabla\Psi(\boldsymbol{\omega}_s^{(n)}). \tag{4.22}$$

where c_j are the elements of $|\mathbf{C}|^{\mathrm{T}}|\mathbf{C}|\mathbf{1}$ with $[|\mathbf{C}|]_{(m,n)} \triangleq |[\mathbf{C}]_{(m,n)}|$ and $\mathbf{1}$ denoting a ones vector [40]. Thus, the SQS method is equivalent to a vectorized quadratic surrogate function with a diagonal Hessian matrix. The proposed (non-diagonal) Hessian based methods converged quickly because their (non-separable) surrogates matched the original cost function closely. By comparison, the SQS algorithm required orders of magnitude more iterations to obtain the same solution because the diagonal Hessian approximation had much higher curvatures. Although each of these iterations had a lower cost than the (non-diagonal) Hessian based methods, the overall computation time remained much longer.

Although the Hessian based methods require only a few iterations to converge, the cost per iteration is significantly greater than the other methods and their lower overall computation time depends on having efficient algorithms for solving expressions with the inverted Hessian matrix. Sparse Cholesky factorization [34] is capable of quickly performing this task, but at a cost of greater memory usage. For instance, the sparse Cholesky method included in MATLAB [34] requires 30 MB of memory for a 128×128 pixel image but needs over 700 MB of memory for a typical 516×516 pixel image when using second-order finite differences (with double precision). Thus, for large problems such as 3-D data sets, the method may not be tractable even when using a tight reconstruction mask. In such situations, it may be advantageous to use an approximation to the full Hessian such as incomplete Cholesky factorization [41]. This change would likely require more iterations of the algorithms, but less cost per iteration. It may also be worth exploring partitioning the volume into smaller sections (e.g., 2-D slices), performing the estimation on each partition in parallel, and then recombining the estimates in a manner that ensures smoothness.

Since field maps are typically smooth over body tissue, low resolution acquisition images may be sufficient for multiple echo time field map estimation. This would reduce both the acquisition time and the computation times of the estimation algorithms. It would also greatly reduce the memory requirements of sparse Cholesky factorization approaches and may allow for the direct application of our fastest methods to 3-D data sets. However, because full resolution acquisitions are required for some water-fat imaging problems (including those in Chapter V), we have focused on traditional image sizes.

Due to the nonconvex nature of the cost function, the initialization of the algorithms is important for ensuring convergence to a desirable local minimum. We presented two different initialization strategies that had low computational costs yet sufficient accuracy for our experimental data sets. If these methods are insufficient, more computationally intensive initialization strategies may be considered. First, since field maps are typically smooth, a multiresolution estimate may provide a reasonable initialization as it would be based on images with higher SNR. One disadvantage of a multiresolution initialization is that it may be inaccurate at object edges or near water-fat tissue interfaces. Second, solving a simpler iterative problem where the sinusoidal data-fit terms in the cost function are replaced with quadratic functions (rather than using quadratic surrogates to solve the complex, sinusoidal cost function) has been found to provide good initial field maps for the water-fat field map estimation problem [23]. Third, using one of the discretized methods outlined in [24, 25] could provide a high quality initialization with no phase wrapping.

The water and fat images estimated from our field maps appear plausible; however, the maximum-likelihood water-fat estimation method (4.21) did produce some errors. First, the water-fat images were noisy due to the lack of regularization in the maximum-likelihood estimator (4.21). To reduce this noise, a joint water-fat, field map estimation method could be used (see Chapter V). Second, the water and fat components were not fully separated in the in-vivo knee data estimates. Given that the simulated data had nearly complete separation, the error in the in-vivo data is likely a result of using a single peak model for the chemical shift spectrum of fat. A more accurate multipeak fat model should increase the fat-water separation [14], while requiring trivial changes to our algorithms.

Although we have focused on field map estimation problems, our proposed methods may also provide significant speed improvements in other areas. For instance, regularized field map estimation is performed as a substep of other iterative algorithms [28]. In such algorithms, any savings provided by our methods would be magnified by the number of outer iterations. There are also problems that require minimizing similar cost functions which currently use SQS or numerical approximations [27, 29, 30]. It may be possible to derive similar fast minimization techniques for these.

4.1.5 Conclusions

We have presented two methods for minimizing the nonconvex cost function associated with regularized field map estimation. Both methods are based on quadratic surrogate functions that majorize the original cost function. The first method uses the quadratic surrogate and sparse Cholesky factorization in a Huber's algorithm based approach. The second method modifies the nonlinear conjugate gradient method by including a monotonic step size line search algorithm based on the quadratic surrogate. Our fastest algorithms were those that used the (sparse) Hessian of the quadratic surrogate function. These converged to the same estimate as the existing separable quadratic surrogate method at least 30 times faster in both multiple echo time field map estimation and water-fat imaging experiments.

4.2 Edge Preserving Field Map Estimation

4.2.1 Introduction

A primary assumption within [2] is that the magnetic field inhomogeneity is smooth over the object support. Under this assumption, it is reasonable to assume that quadratic regularization with a second-order finite differencing matrix will adequately capture the inhomogeneity. In fact, such regularization has been used to accurately estimate smooth field maps in both simulated and experimental settings [2, 14, 23]. However, while the magnetic field is typically smooth over uniform tissue, it can have large discontinuities at tissue interfaces (e.g., the surface of metallic implants and at the boundaries of the lungs and sinuses) [42–46]. In these regions, using quadratic regularization with second-order finite differences can result in errors in the field map estimate [14]. In this section, we evaluate the effects of other regularization strategies (particularly edge preserving regularization and lower-order finite differences) on field map estimation accuracy.

4.2.2 Field Map Estimation with Edge Preserving Regularization

We consider two different edge preserving regularization functions with low computational costs and simple quadratic surrogate functions. The first is the hyperbola function [47]

$$R_{\mathrm{H}}(x) \triangleq \delta^2 \left[\sqrt{1 + (x/\delta)^2} - 1 \right], \tag{4.23}$$

which has the following Huber's curvature [48]:

$$c_{\mathrm{H}}(x) = \frac{1}{\sqrt{1 + (x/\delta)^2}}. \tag{4.24}$$

The second is the Lange3 [49]

$$R_{\mathrm{L}}(x) \triangleq \delta^2 \left[\frac{|x|}{\delta} - \log\left(1 + \frac{|x|}{\delta}\right) \right], \tag{4.25}$$

which has the following Huber's curvature [48]:

$$c_{\mathrm{L}}(x) = \frac{1}{1 + |x/\delta|}. \tag{4.26}$$

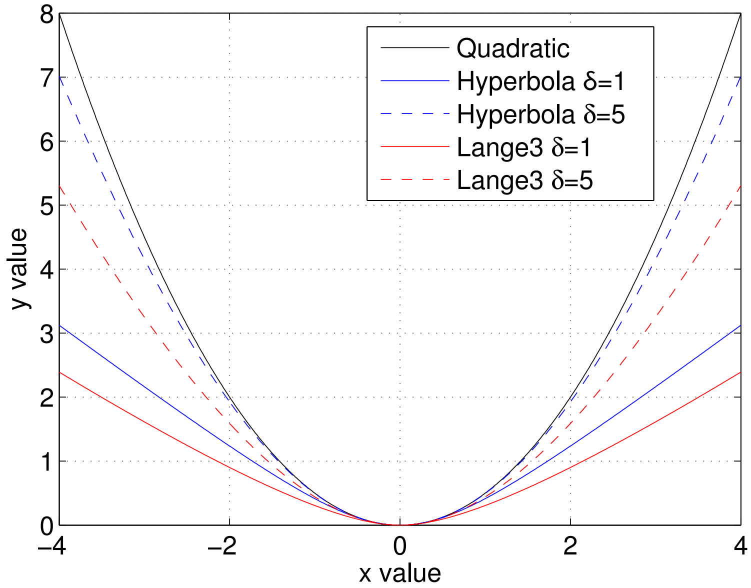

Figure 4.13: Hyperbola and Lange3 edge preserving functions for several values of δ compared to the standard quadratic function.

Fig. 4.13 compares these two functions to the quadratic function, for varying values of δ. For a given δ value, the Lange3 function is the most edge preserving out of the three as it is the "widest" function. Furthermore, a smaller δ parameter corresponds to greater edge preserving behavior in both edge preserving functions.

These functions replace the quadratic regularization term in (4.7). To solve this modified problem, we develop parabolic surrogate functions for the new regularization terms using the stated Huber's curvatures. Combining these parabolic functions with the quadratic surrogate function for the data-fit term, we get overall quadratic surrogate functions for the modified cost function (4.7) that we can solve using any of the methods in Section 4.1.

4.2.3 Simulation of Field Inhomogeneity at Tissue Interfaces

Simulations of the magnetic field inhomogeneity at tissue interfaces are needed to evaluate the different regularization methods. There are several common geometric shapes for which there exist analytic formulas for the magnetic field inhomogeneity. These include a sphere [43, 44], a double walled tube [45], a thin circular disk [44], and an infinite cylinder [44]; all of which are in a uniform medium. For our experiments, we simulated a sphere of air surrounded by water as this was the closest representation of a sinus cavity in the head. For magnetic susceptibilities $\ll 1$, the magnetic fields inside the sphere ($\mathbf{B}_i$) and in the surrounding medium ($\mathbf{B}_e$) are given by

$$\mathbf{B}_i \approx \mathbf{B}_0 \left(1 + \frac{\chi_e}{3}\right) \tag{4.27}$$

$$\mathbf{B}_e \approx \mathbf{B}_0 \left(1 + \frac{\chi_e}{3} + r^3 \frac{(\chi_e - \chi_i)(x^2 + y^2 - 2z^2)}{3(x^2 + y^2 + z^2)^{5/2}}\right) \tag{4.28}$$

where $\mathbf{B}_0$ is the static field oriented in the z-direction, r is the radius of the sphere, χ_i and χ_e are the magnetic susceptibilities of the sphere and the surrounding medium respectively (in SI units), and x, y, z are the Cartesian coordinates [43]. For our problem of air in water, $\chi_i = 0.36 \times 10^{-6}$ and $\chi_e = -9.05 \times 10^{-6}$ [44].[6] The added term $\chi_e/3$ is the Lorentz correction which can be accounted for during MR calibration and is therefore ignored in our experiments [44].

These equations describe the actual magnetic field strength values and not the frequency ω that is typically estimated using field map estimation (see Section 2.2). We use (2.24) with the gyromagnetic ratio of hydrogen to convert between the two terms.

4.2.4 Results

We evaluated the different regularization strategies using a data set that simulates the field map inhomogeneity resulting from an air cavity in the head (e.g., a sinus). To do this, we combined a brain magnetization volume with a simulated field map

---

[6] A large table of magnetic susceptibilities for different materials can be found in [44].

Figure 4.14: (a) Tranverse, (b) sagittal, and (c) coronal planes of the BrainWeb volume.

for a sphere of air in water. The simulated brain magnetization volume was the standard 181 × 217 × 181 voxel, noiseless, 1 mm isotropic resolution, T1 weighted data set from the BrainWeb database [50]. Fig. 4.14 shows the anatomical planes through this volume.

To evaluate the effects of an air cavity in the brain, we simulated the field map that would result from a sphere of air, 5 mm in radius, being placed in an infinite medium of water using (4.28). We assumed a field strength of 1.5 T and ignored other possible sources of field inhomogeneity. To mimic the effects of the finite resolution of a true scan, we computed the field map values at twice the resolution ($\boldsymbol{\omega}$), determined the resulting oversampled $e^{i\boldsymbol{\omega}t_\ell}$ image for each echo time difference, and then downsampled these with a $2 \times 2 \times 2$ box filter. The sphere of air was aligned with the upper portion of the sphenoid sinus in the BrainWeb data set and the field map was oriented as though the body was parallel to the main magnetic field. We selected the transverse slice at the center of the sphere for our experiments. Fig. 4.15 presents the resulting field map and a diagram highlighting the location of the sphere with respect to the true magnetization image.

We combined the $e^{i\boldsymbol{\omega}t_\ell}$ images with the brain image to create three acquisition images with relative echo times $t_\ell = 0, 1, 8\,\mathrm{ms}$ and $\mathrm{R}_2^* = 20\,\mathrm{s}^{-1}$. Complex Gaussian white noise was also added to these images so that the $\mathrm{SNR} \approx 20\,\mathrm{dB}$. Fig. 4.16

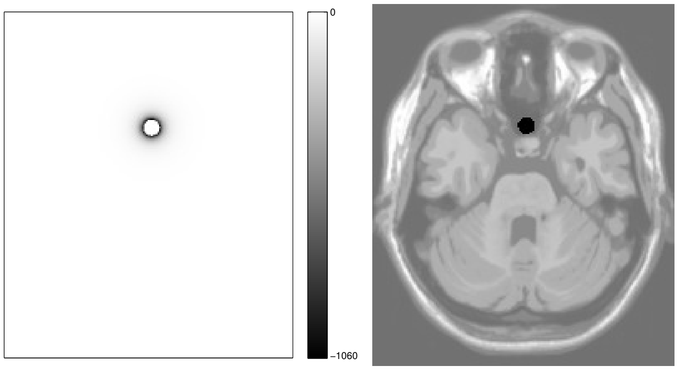

Figure 4.15: The simulated field map in Hz (left) and its location with respect to the magnitude image as identified with the black circle (right).

presents the magnitudes and phase of the final scan images.

We performed regularized field map estimates on this data set for each of the regularizers listed in Table 4.4. The estimation was performed using the Huber's algorithm based method outlined in Section 4.1.2 with the appropriate curvatures for each potential function (see Section 4.2.2). Furthermore, the estimate was restricted to a mask consisting of a dilated convex hull around the object pixels. We initialized all of the algorithms with the same conventional estimate and ran each algorithm until convergence. We computed the root-mean squared error (RMSE) between each estimate and the true field map[7] over the entire mask and over a region of interest (ROI) primarily containing the simulated sphere as highlighted in Fig. 4.17. A grid search was used to determine the parameter values for each potential function that resulted in the estimates with the lowest RMSE and no phase wrapping (phase wrapping was common with very small regularization values). The resulting estimates are presented in Fig. 4.18 and their corresponding RMSE values are found in Table 4.4.

---

[7] The "true" low-resolution field map was created by downsampling the oversampled field map with a $2 \times 2 \times 2$ box filter. This approximation was necessary because the field maps recovered from the undersampled $e^{i\boldsymbol{\omega} t_\ell}$ images using the conventional method contained phase wrapping, even for the smallest echo time differences. Although the effect of directly smoothing the field map is not the same as undersampling the $e^{i\boldsymbol{\omega} t_\ell}$ images, the resulting field map was nearly identical to the conventional estimate for those pixels without phase wrapping.

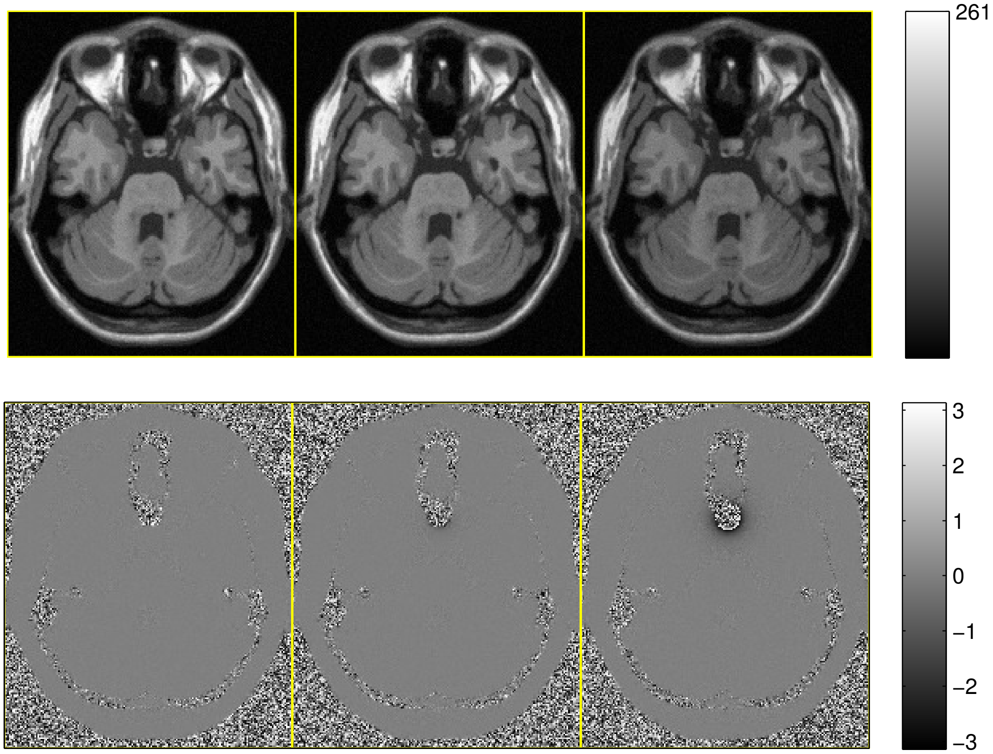

Figure 4.16: Magnitude (top) and phase (bottom) images of the three simulated acquisitions ($\mathbf{y}_l$) with $t_\ell = 0, 1, 8\,\mathrm{ms}$ from left to right.

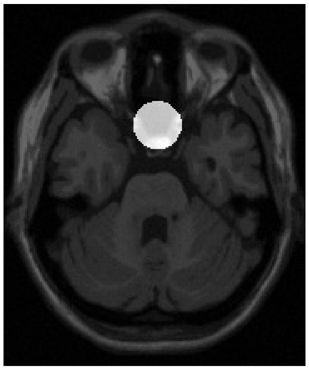

Figure 4.17: Image highlighting the region of interest used to compute RMSE values.

Table 4.4: RMSE and RMSE Over the ROI in Fig. 4.17 for Varying Regularizers

Function	**Order of C**	**Reg. Param**	δ	**RMSE (Hz)**	**RMSE - ROI (Hz)**
Quadratic	2	2^2		4.46	29.97
Quadratic	1	2^{-2}		4.76	31.89
Hyperbola	1	2^{-3}	2^6	4.60	30.64
Lange3	1	2^{-3}	2^6	4.56	30.32

We also present the line profiles of each estimate through the center of the sphere in Fig. 4.19.

4.2.5 Discussion

The estimate with the lowest RMSE came from the quadratic regularizer with second-order finite differences as it produced more accurate estimates of the uniform field away from the sphere. However, this method was highly inaccurate at the edges of the sphere, something the RMSE metric does not fully capture. The first-order finite differencing matrix estimators were able to better capture these sharp discontinuities, but at a cost of higher noise in the uniform regions, and thus, larger RMSEs. Among these regularizers, the edge preserving penalty functions provided a marginal improvement ($\approx 2\%$) in RMSE over the standard quadratic potential function. There was no significant difference between the estimates of the two edge preserving potential functions.

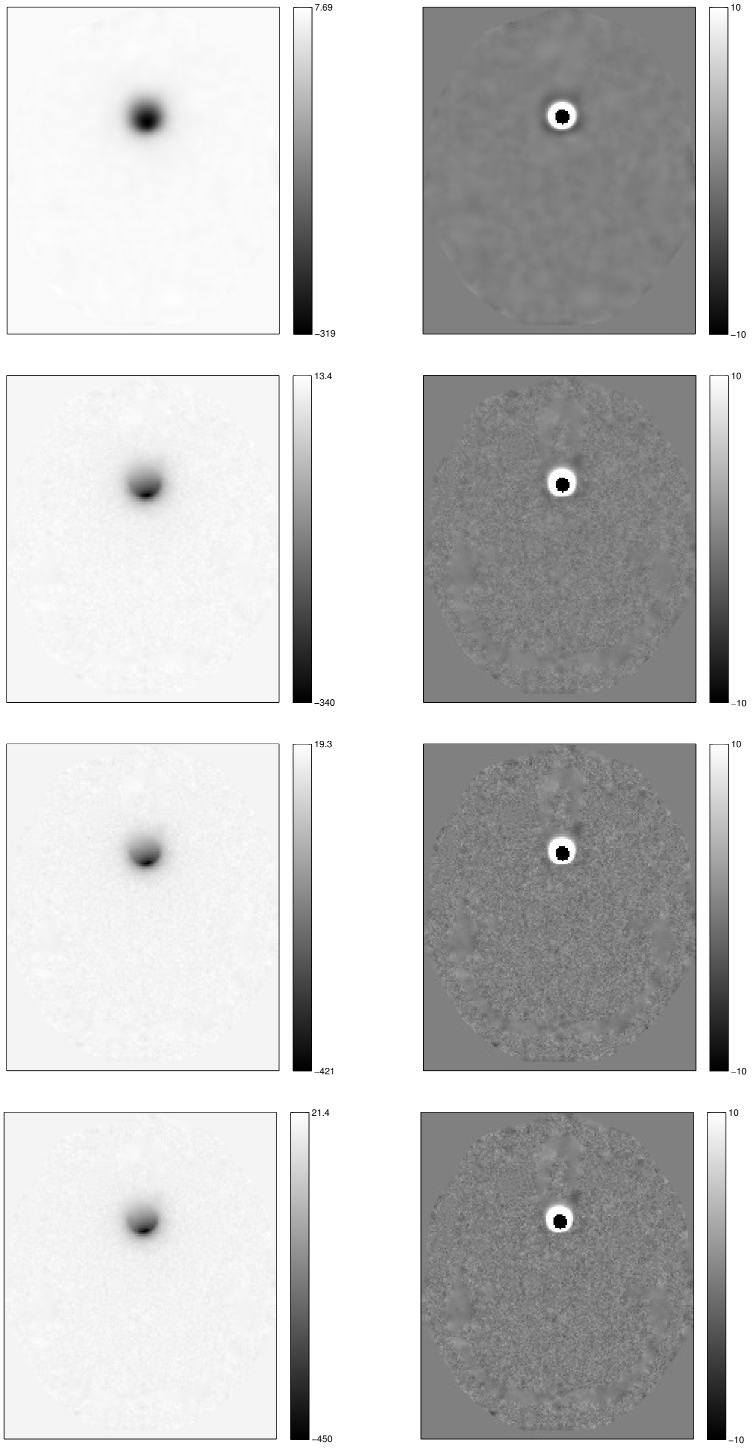

Figure 4.18: Magnitudes of the final estimates in Hz from each of the methods in Table 4.4 (left) and their difference in Hz to the truth (right). The order from top to bottom is quadratic with second-order differences, quadratic with first-order differences, hyperbola with first-order differences, and Lange3 with first-order differences.

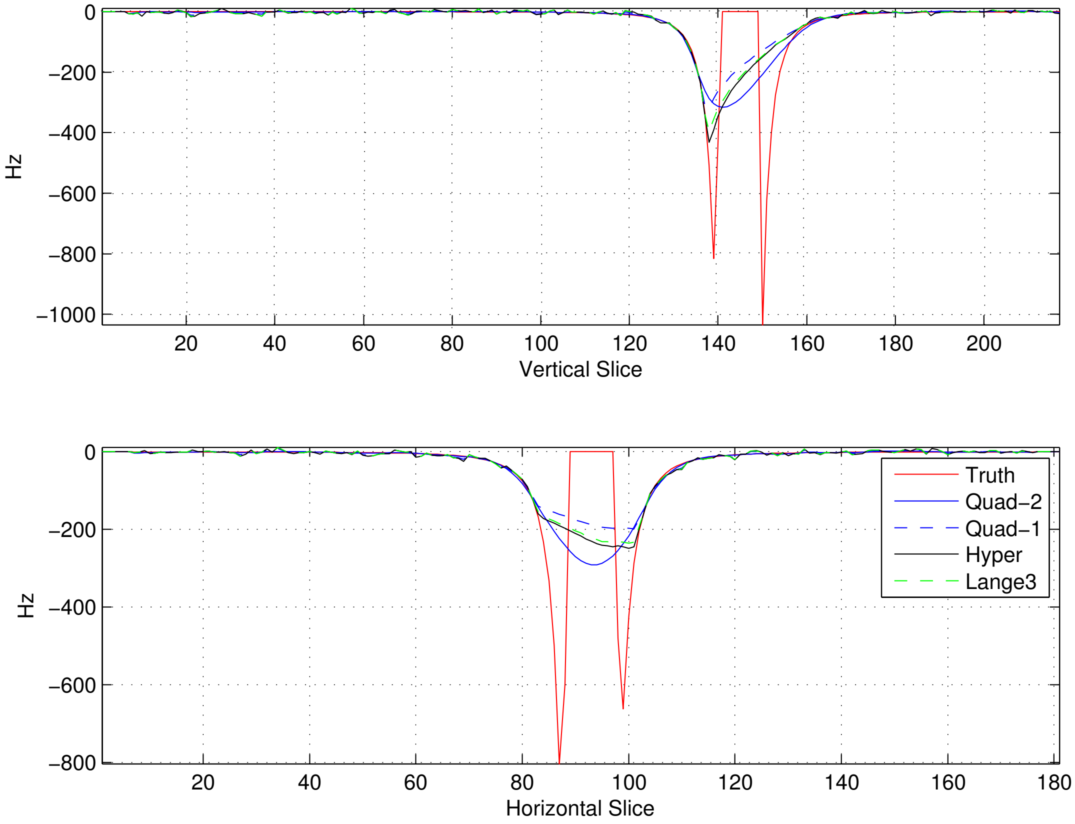

Figure 4.19: Comparison of vertical and horizontal line profiles of the estimates through the center of the sphere.

We can evaluate the ability of each regularizer to capture the sharp edges of the simulated field map by considering the cross sections of their corresponding estimates. The performance of the estimators with first-order finite differences were all similar, capturing part of the steep outer edge but missing the homogeneity in the center of the sphere. The inability of these methods to capture the field map within the sphere is not surprising given that there is limited signal in this region. What is more surprising is the similarity between the quadratic and edge preserving potential functions when using the same finite differencing matrix. Higher levels of edge preserving regularization could be used to capture more of the edges; however, the parameter values required to avoid phase wrapping in this data set resulted in potential functions with weak levels of edge preserving. The quadratic regularizer with second-order finite differences did not capture as much of the edge as those with first-order finite differences and improperly estimated a peak in the center of the sphere. This is again unsurprising given that second-order finite differences typically result in smoother estimates and interpolate with higher degrees of freedom over regions with low signal.

It is possible to preserve more of the edges in the field inhomogeneity with these regularizers, but such estimates are typically corrupted by noise. Fig. 4.20 shows one such estimate that was created using first-order finite differences and a quadratic regularizer with regularization parameter 2^{-8}. The estimate is very noisy and exhibits significant errors in the low signal regions. This is expected as the parameter value is small enough that the effect of the regularization term is minimal.

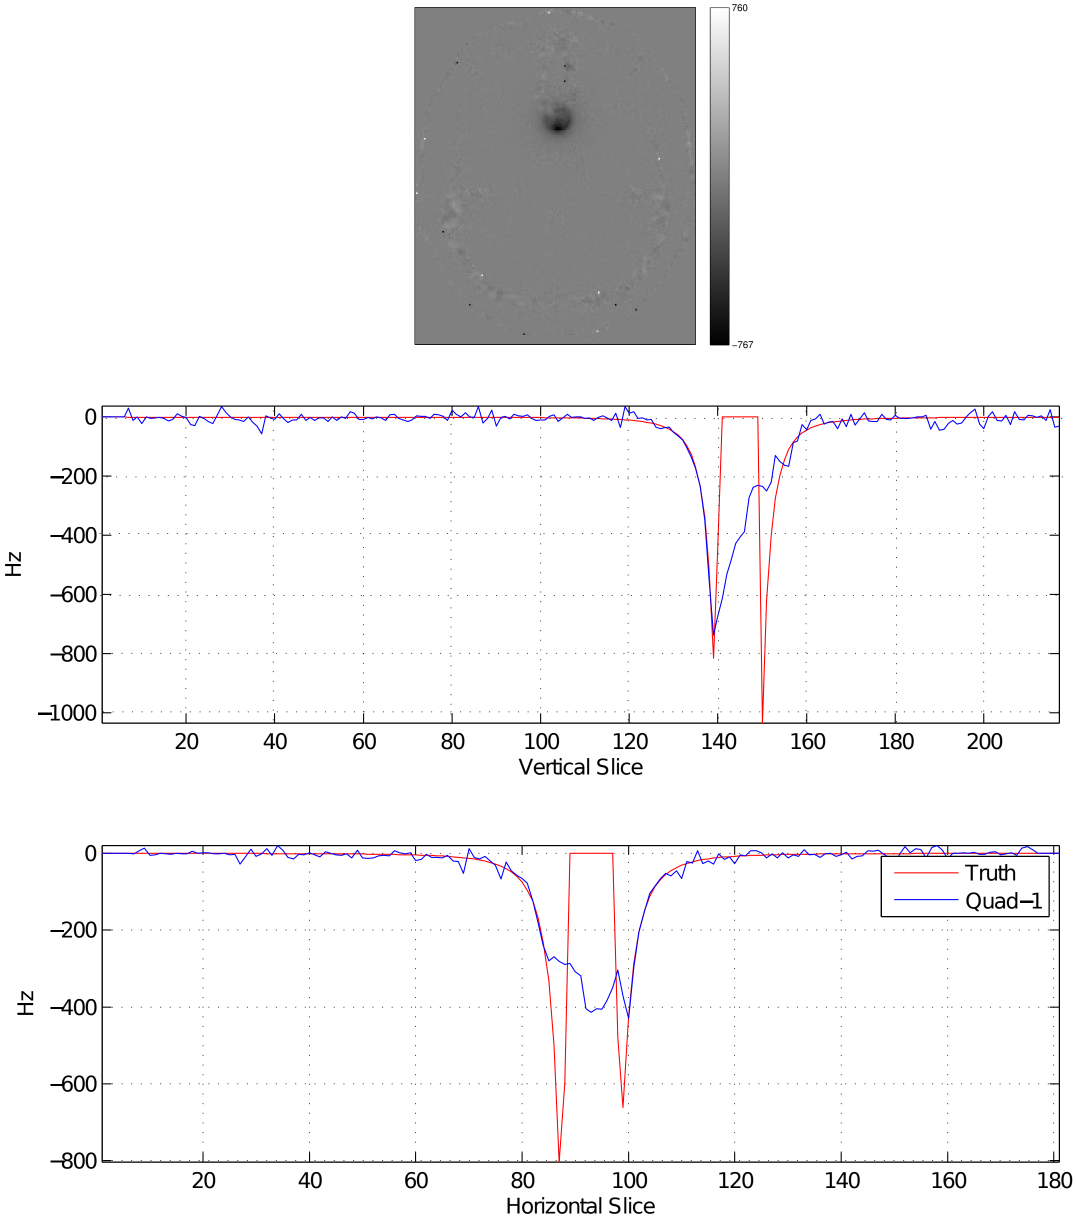

Figure 4.20: (Top) Magnitude of an estimate using quadratic regularization with first-order finite differences and $\beta = 2^{-8}$ (Hz). (Bottom) Line profiles through the center of the estimate.

4.2.6 Conclusions

In this section, we explored using both edge preserving regularization and different finite differencing orders in regularized field map estimation. Only slight modifications to our algorithms from Section 4.1 were required to incorporate these changes. Based on experiments using a simulated sphere of air in the sphenoid sinus, we found that the biggest effect on field map estimation quality was the order of the finite differencing matrix. In particular, for parameter values that avoided phase wrapping errors, first-order finite differences resulted in better edge preservation than second-order finite differences. The use of edge preserving regularization did improve estimation quality, as measured by RMSE, but only by 2%. The parameter values that fully captured the sharp discontinuities in the simulated field map resulted in noisy estimates with significant phase wrapping. This is partially a result of the large field inhomogeneity values produced by a simulated sphere of air in a 1.5 T magnetic field. If this experiment was repeated with smaller inhomogeity values, such as those typically seen in lower resolution in-vivo data, an acceptable compromise may be available. Furthermore, edge preserving regularization may have a greater impact on water-fat imaging, which often has a large number of tissue interfaces within the field-of-view. Such experiments are presented in Chapter VI as future work.

Bibliography

[1] M. J. Allison and J. A. Fessler, “Accelerated computation of regularized field map estimates,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2012, p. 0413.

[2] A. K. Funai, J. A. Fessler, D. T. B. Yeo, V. T. Olafsson, and D. C. Noll, “Regularized field map estimation in MRI,” *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1484–94, Oct. 2008.

[3] T. A. Bley, O. Wieben, C. J. François, J. H. Brittain, and S. B. Reeder, “Fat and water magnetic resonance imaging,” *J. Mag. Res. Im.*, vol. 31, no. 1, pp. 4–18, Jan. 2010.

[4] P. Jezzard and R. S. Balaban, “Correction for geometric distortion in echo planar images from B_0 field variations,” *Mag. Res. Med.*, vol. 34, no. 1, pp. 65–73, July 1995.

[5] K. Sekihara, S. Matsui, and H. Kohno, “NMR imaging for magnets with large nonuniformities,” *IEEE Trans. Med. Imag.*, vol. 4, no. 4, pp. 193–9, Dec. 1985.

[6] P. J. Reber, E. C. Wong, R. B. Buxton, and L. R. Frank, “Correction of off resonance-related distortion in echo-planar imaging using EPI-based field maps,” *Mag. Res. Med.*, vol. 39, no. 2, pp. 328–30, Feb. 1998.

[7] D. C. Noll, J. A. Fessler, and B. P. Sutton, “Conjugate phase MRI reconstruction with spatially variant sample density correction,” *IEEE Trans. Med. Imag.*, vol. 24, no. 3, pp. 325–36, Mar. 2005.

[8] G. H. Glover and E. Schneider, “Three-point Dixon technique for true water/fat decomposition with B_0 inhomogeneity correction,” *Mag. Res. Med.*, vol. 18, no. 2, pp. 371–83, Apr. 1991.

[9] 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,” *Mag. Res. Med.*, vol. 51, no. 1, pp. 35–45, Jan. 2004.

[10] K. Zhou, M. Zaitsev, and S. Bao, “Reliable two-dimensional phase unwrapping method using region growing and local linear estimation,” *Mag. Res. Med.*, vol. 62, no. 4, pp. 1085–1090, Oct. 2009.

[11] 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,” *Mag. Res. Med.*, vol. 54, no. 4, pp. 1032–9, Oct. 2005.

[12] W. Lu and B. A. Hargreaves, “Multiresolution field map estimation using golden section search for water-fat separation,” *Mag. Res. Med.*, vol. 60, no. 1, pp. 236–44, July 2008.

[13] M. Jacob and B. P. Sutton, "Algebraic decomposition of fat and water in MRI," *IEEE Trans. Med. Imag.*, vol. 28, no. 2, pp. 173–84, Feb. 2009.

[14] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig, "Compressed sensing for chemical shift-based water-fat separation," *Mag. Res. Med.*, vol. 64, no. 6, pp. 1749–59, Dec. 2010.

[15] 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," *Mag. Res. Med.*, vol. 54, no. 3, pp. 636–44, Sept. 2005.

[16] C. Windischberger, S. Robinson, A. Rauscher, M. Barth, and E. Moser, "Robust field map generation using a triple-echo acquisition," *J. Mag. Res. Im.*, vol. 20, no. 4, pp. 730–4, Oct. 2004.

[17] E. Schneider and G. Glover, "Rapid in vivo proton shimming," *Mag. Res. Med.*, vol. 18, no. 2, pp. 335–47, Apr. 1991.

[18] P. Irarrazabal, C. H. Meyer, D. G. Nishimura, and A. Macovski, "Inhomogeneity correction using an estimated linear field map," *Mag. Res. Med.*, vol. 35, no. 2, pp. 278–82, Feb. 1996.

[19] S. D. Sharma, H. H. Hu, and K. S. Nayak, "Chemical shift encoded water-fat separation using parallel imaging and compressed sensing," *Mag. Res. Med.*, vol. 69, no. 2, pp. 456–66, Feb. 2013.

[20] S. D. Sharma, H. H. Hu, and K. S. Nayak, "Accelerated water-fat imaging using restricted subspace field map estimation and compressed sensing," *Mag. Res. Med.*, vol. 67, no. 3, pp. 650–9, Mar. 2012.

[21] J. Tsao and Y. Jiang, "Hierarchical IDEAL: Fast, robust, and multiresolution separation of multiple chemical species from multiple echo times," *Mag. Res. Med.*, vol. 70, no. 1, pp. 155–9, 2013.

[22] J. M. Bioucas-Dias and G. Valado, "Phase unwrapping via graph cuts," *IEEE Trans. Im. Proc.*, vol. 16, no. 3, pp. 698–709, Mar. 2007.

[23] W. Huh, J. A. Fessler, and A. A. Samsonov, "Water-fat decomposition with regularized field map," in *Proc. Intl. Soc. Mag. Res. Med.*, 2008, p. 1382.

[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," *Mag. Res. Med.*, vol. 59, no. 3, pp. 571–80, Mar. 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," *Mag. Res. Med.*, vol. 63, no. 1, pp. 79–90, Jan. 2010.

[26] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *American Statistician*, vol. 58, no. 1, pp. 30–7, Feb. 2004.

[27] H. Wang and Y. Cao, "Spatially regularized T1 estimation from variable flip angles MRI," *Med. Phys.*, vol. 39, no. 7, pp. 4139–48, 2012.

[28] W. A. Grissom, V. Rieke, A. B. Holbrook, Y. Medan, M. Lustig, J. Santos, M. V. McConnell, and K. B. Pauly, "Hybrid referenceless and multibaseline subtraction MR thermometry for monitoring thermal therapies in moving organs," *Med. Phys.*, vol. 37, no. 9, pp. 5014–26, 2010.

[29] W. Tang, *Reconstruction of Parametric Image Maps in Single- and Multiple-Coil Functional Magnetic Resonance Imaging*, Ph.D. thesis, Auburn University, 2009.

[30] S. E. Maier, S. Vajapeyam, H. Mamata, C.-F. Westin, F. A. Jolesz, and R. V. Mulkern, "Biexponential diffusion tensor analysis of human brain diffusion data," *Mag. Res. Med.*, vol. 51, no. 2, pp. 321–30, Feb. 2004.

[31] P. J. Huber, *Robust statistics*, Wiley, New York, 1981.

[32] 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," *Mag. Res. Med.*, vol. 60, no. 5, pp. 1122–34, Nov. 2008.

[33] W. Huh, *Regularized statistical material decomposition in medical imaging*, Ph.D. thesis, University of Michigan, 2011.

[34] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, "Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate," *ACM Trans. Math. Software*, vol. 35, no. 3, pp. 22:1–22:14, Oct. 2008.

[35] J. A. Fessler and S. D. Booth, "Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction," *IEEE Trans. Im. Proc.*, vol. 8, no. 5, pp. 688–99, May 1999.

[36] C. T. Kelley, *Iterative methods for optimization*, Soc. Indust. Appl. Math., 1999.

[37] L. Armijo, "Minimization of functions having continuous derivatives," *Pacific J. Math*, vol. 16, no. 1, pp. 1–3, 1966.

[38] E. Polak and G. Ribiere, "Note sur la convergence de méthodes de directions conjuguées," *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique*, vol. 3, pp. 35–43, 1969.

[39] R. Fletcher and C. M. Reeves, "Function minimization by conjugate gradients," *Comput. J*, vol. 7, no. 2, pp. 149–54, 1964.

[40] A. R. De Pierro, “A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography,” *IEEE Trans. Med. Imag.*, vol. 14, no. 1, pp. 132–7, Mar. 1995.

[41] D. S. Kershaw, “The incomplete Cholesky-conjugate gradient method for the iterative solution of systems of linear equations,” *J. Comp. Phys.*, vol. 26, no. 1, pp. 43–65, Jan. 1978.

[42] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: Physical principles and sequence design*, Wiley, New York, 1999.

[43] D. A. Yoder, Y. Zhao, C. B. Paschal, and J. M. Fitzpatrick, “MRI simulator with object-specific field map calculations,” *Mag. Res. Im.*, vol. 22, no. 3, pp. 315–28, Apr. 2004.

[44] J. F. Schenck, “The role of magnetic susceptibility in magnetic resonance imaging: MRI magnetic compatibility of the first and second kinds,” *Med. Phys.*, vol. 23, no. 6, pp. 815–50, June 1996.

[45] A. J. Lucas, S. J. Gibbs, E. W. G. Jones, M. Peyron, J. A. Derbyshire, and L. D. Hall, “Diffusion imaging in the presence of static magnetic-field gradients,” *J. Mag. Res.*, vol. 104, no. :, pp. 273–82, 1993.

[46] P. Aksit, J. A. Derbyshire, and J. L. Prince, “Three-point method for fast and robust field mapping for EPI geometric distortion correction,” in *Proc. IEEE Intl. Symp. Biomed. Imag.*, 2007, pp. 141–4.

[47] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, “Two deterministic half-quadratic regularization algorithms for computed imaging,” in *Proc. IEEE Intl. Conf. on Image Processing*, 1994, vol. 2, pp. 168–71.

[48] J. A. Fessler, “Image reconstruction: Algorithms and analysis,” Book in preparation, 2006.

[49] K. Lange, “Convergence of EM image reconstruction algorithms with Gibbs smoothing,” *IEEE Trans. Med. Imag.*, vol. 9, no. 4, pp. 439–46, Dec. 1990, Corrections, T-MI, 10:2(288), June 1991.

[50] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans, “Design and construction of a realistic digital brain phantom,” *IEEE Trans. Med. Imag.*, vol. 17, no. 3, pp. 463–8, June 1998.

CHAPTER V

Water-Fat Image Reconstruction

In this chapter, we present novel initialization and minimization algorithms for the compressed sensing based water-fat image reconstruction problem. Our initialization strategy is an extension of our work in Chapter IV, while our minimization algorithm builds upon the augmented Lagrangian methods and preconditioning methods outlined in Chapters III and IV, respectively. The resulting overall minimization strategy reduces the time required to compute the water and fat images by over a factor of 12 compared to the existing Gauss-Newton based nonlinear CG algorithm in [1].

5.1 Introduction

Simultaneously imaging both water and fat presents a challenge for MRI as the protons in the fat molecules experience a chemical shift relative to those in water. If disregarded, this shift can cause artifacts in the reconstructed images (see Section 2.3). Furthermore, any main magnetic field inhomogeneity (as described in Chapter IV) obfuscates the distinction between fat and water components. To overcome these complications, several types of water-fat MR imaging techniques have been proposed.

There are four major classifications of water-fat imaging techniques (see [2, 3] for thorough reviews). The first is fat suppression in which the fat portion of the spectrum is excited and crushed prior to acquisition, thereby eliminating it from the image [4]. The second is water selective excitation, where the pulse sequence is designed to only excite the water molecules and not the fat [5,6]. The third is steady state free precession (SSFP) techniques that have carefully designed passbands in the phase response [2]. The fourth is chemical shift based methods (or Dixon methods) that encode chemical shift information into the signal phase by acquiring images at different echo times [1, 7–16]. Chemical shift based methods have increased in popularity because, unlike the other three water-fat imaging types, they intrinsically account for field map inhomogeneity and can provide separate water and fat images, which is beneficial in some clinical settings (e.g., fat quantification [17]) [3].

One of the major limitations of chemical shift based methods is that they require multiple acquisitions (ranging from 3 to over 10). This results in much longer acquisition times compared to the other water-fat imaging techniques. To address this limitation, recent work has focused on combining compressed sensing (CS) principles (see Section 2.4.2) with the existing chemical shift based methods to obtain accurate reconstructions from significantly less data, and subsequently, shorter acquisition times [1, 18–20]. The disadvantage of these CS based water-fat imaging methods is that they are computationally intensive [1, 18–20]. This chapter proposes a method to reduce the computation time of these estimators.

5.2 Compressed Sensing Based Water-Fat Imaging

We focus on the CS-WF method presented in [1] which uses the most direct combination of CS and water-fat imaging (other more recent CS based water-fat imaging methods [18–20] are extensions of this work). The system model used in [1] is

$$\mathbf{g}_l(\mathbf{x}) = \mathbf{F}_l\left[(\boldsymbol{\rho}_w + \boldsymbol{\rho}_f e^{2\pi i t_l \Delta_f}) \cdot e^{2\pi i t_l \boldsymbol{\phi}}\right], \qquad l = 1, \ldots, L \tag{5.1}$$

where L is the number of scans (typically three), $\mathbf{g}_l$ is the nonlinear measurement operator (representing the k-space samples for the lth acquisition), $\mathbf{F}_l \in \mathbb{C}^{K \times N}$ is an undersampled Fourier transform for the lth acquisition, $\boldsymbol{\rho}_w \in \mathbb{C}^N$ and $\boldsymbol{\rho}_f \in \mathbb{C}^N$ are the unknown water and fat images respectively, $t_l \in \mathbb{R}$ is the echo time shift of the lth acquisition, $\Delta_f \in \mathbb{R}$ is the (known) frequency shift between water and fat[1], $\boldsymbol{\phi} \in \mathbb{R}^N$ is the unknown field map in Hz, and $\mathbf{x} = [\boldsymbol{\rho}_w, \boldsymbol{\rho}_f, \boldsymbol{\phi}]$ is the vector of all unknowns.

The CS-WF method in [1] is derived by combining regularized chemical shift based water-fat imaging methods [11, 12, 16], which assume smoothness in the field map $\boldsymbol{\phi}$, with CS concepts, which assume the fat and water images $\boldsymbol{\rho}$ are sparse in a known domain. The resulting estimator uses the optimization problem

$$\underset{\boldsymbol{\rho},\boldsymbol{\phi}}{\arg\min}\ \|\mathbf{g}(\mathbf{x}) - \mathbf{y}\|_2^2 + \lambda_\rho \|\boldsymbol{\Psi}\boldsymbol{\rho}\|_1 + \lambda_\phi \|\boldsymbol{\Phi}\boldsymbol{\phi}\|_2^2 \tag{5.2}$$

where λ_ρ and λ_ϕ are regularization parameters, $\boldsymbol{\Phi} \in \mathbb{R}^{M_3 \times N}$ is a finite-differencing matrix,

$$\mathbf{g}(\mathbf{x}) = \begin{bmatrix} \mathbf{g}_1(\mathbf{x}) \\ \vdots \\ \mathbf{g}_L(\mathbf{x}) \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_L \end{bmatrix}, \quad \boldsymbol{\rho} = \begin{bmatrix} \boldsymbol{\rho}_w \\ \boldsymbol{\rho}_f \end{bmatrix}, \quad \boldsymbol{\Psi} = \begin{bmatrix} \boldsymbol{\Psi}_w & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Psi}_f \end{bmatrix} \tag{5.3}$$

with $\mathbf{y}_l$ denoting the measured k-space data from the lth acquisition and $\boldsymbol{\Psi}_w \in \mathbb{C}^{M_1 \times N}$, $\boldsymbol{\Psi}_f \in \mathbb{C}^{M_2 \times N}$ denoting sparsifying transforms.

---

[1] As in Chapter IV, extending the methods in this section to include a linear combination of several chemical shift values is trivial.

The difficulty with the CS-WF method is that it requires minimizing a non-quadratic cost function (5.2) containing the nonlinear $\mathbf{g}(\mathbf{x})$ term. The approach used in [1] is to first linearize the function around the current estimate (i.e., $\mathbf{g}(\mathbf{x}_n + d\mathbf{x}) \approx \mathbf{g}(\mathbf{x}_n) + d\mathbf{g}(\mathbf{x}_n)d\mathbf{x}$) to create the modified problem

$$(5.4) \quad \underset{d\boldsymbol{\rho}, d\boldsymbol{\phi}}{\arg\min} \; \|\mathbf{g}(\mathbf{x}_n) + d\mathbf{g}(\mathbf{x}_n)d\mathbf{x} - \mathbf{y}\|_2^2 + \lambda_\rho \|\boldsymbol{\Psi}(\boldsymbol{\rho}_n + d\boldsymbol{\rho})\|_1 + \lambda_\phi \|\boldsymbol{\Phi}(\boldsymbol{\phi}_n + d\boldsymbol{\phi})\|_2^2$$

where $d\mathbf{g}(\mathbf{x}_n)$ is the Jacobian of $\mathbf{g}(\mathbf{x})$ evaluated at the current estimate $\mathbf{x}_n \triangleq [\boldsymbol{\rho}_{w_n}, \boldsymbol{\rho}_{f_n}, \boldsymbol{\phi}_n]$. Then after solving (5.4) using a nonlinear CG (NCG) method (with a corner rounding approximation of the ℓ_1-norm as in [21, Appendix]), they update the estimate using $\mathbf{x}_{n+1} = \mathbf{x}_n + t d\mathbf{x}$ where t is found using backtracking line search. This is repeated until a stop criterion is achieved.

Although the regularized nature of this method makes it robust to low SNR, it is computationally intensive [1]. For example, a C-code implementation on a 2.4 GHz CPU required 9 min to compute an estimate of a 240 × 192 × 54 voxel data set [1]. Other CS based water-fat imaging methods use similar reconstruction algorithms. Wiens et al. proposed an R_2^* corrected version of the CS-WF algorithm and used the same NCG minimization strategy as above [20]. Sharma et al. proposed two CS based water-fat imaging methods that differ in their field map estimation: one fitting B-splines [18] and the other using a restricted subspace approach [19]. In both cases, they used an alternating minimization strategy where they first updated $\boldsymbol{\rho}$ using a corner rounding approximation and NCG method and then updated $\boldsymbol{\phi}$ by jointly minimizing a linearized cost function using a second corner rounding, NCG method. All of these methods report MATLAB reconstruction times in the tens of minutes. We propose a novel minimization strategy, the key difference being the use of variable splitting and an augmented Lagrangian based method, that greatly reduces the estimation time. We also propose a new initialization method that provides field map estimates with comparable regularization to the CS-WF method.

5.3 Novel Minimization Strategy

Linearizing the signal model in (5.2) is an effective strategy for handling the nonlinear field inhomogeneity term as only a few outer iterations are required to converge to a desirable local minimum when using a reasonably accurate initialization [1]. However, jointly minimizing the resulting linearized cost function like in [1,18–20] is challenging due to the presence of an ℓ_1-norm as well as the sizes and structures of the sparsifying and finite differencing matrices. In contrast, efficient methods for solving the linearized cost function (5.4) in terms of only $d\boldsymbol{\rho}$ or $d\boldsymbol{\phi}$ are easier to create. We therefore minimize (5.4) using an alternating minimization strategy in which we sequentially solve for one of $d\boldsymbol{\rho}$ or $d\boldsymbol{\phi}$ while keeping the other variable constant:

$$d\boldsymbol{\rho}^{(k+1)} = \arg\min_{d\boldsymbol{\rho}} \|\mathbf{g}(\mathbf{x}_n) + \mathbf{A}_{\rho_n} d\boldsymbol{\rho} + \mathbf{A}_{\phi_n} d\boldsymbol{\phi}^{(k)} - \mathbf{y}\|_2^2 + \lambda_\rho \|\boldsymbol{\Psi}(\boldsymbol{\rho}_n + d\boldsymbol{\rho})\|_1 \tag{5.5}$$

$$d\boldsymbol{\phi}^{(k+1)} = \arg\min_{d\boldsymbol{\phi}} \|\mathbf{g}(\mathbf{x}_n) + \mathbf{A}_{\rho_n} d\boldsymbol{\rho}^{(k+1)} + \mathbf{A}_{\phi_n} d\boldsymbol{\phi} - \mathbf{y}\|_2^2 + \lambda_\phi \|\boldsymbol{\Phi}(\boldsymbol{\phi}_n + d\boldsymbol{\phi})\|_2^2 \tag{5.6}$$

where $\mathbf{A}_{\rho_n}$ and $\mathbf{A}_{\phi_n}$ are the portions of the Jacobian matrix corresponding to $\boldsymbol{\rho}$ and $\boldsymbol{\phi}$ respectively (i.e., $d\mathbf{g}(\mathbf{x}_n) = [\mathbf{A}_{\rho_n}\ \mathbf{A}_{\phi_n}]$). Specifically,

$$\mathbf{A}_{\rho_n} = \mathbf{F}_{\mathrm{B}} \mathbf{T}_n \text{ where } \mathbf{T}_n = \begin{bmatrix} \mathbf{D}_1 & \mathbf{D}_1 e^{2\pi i t_1 \Delta_f} \\ \mathbf{D}_2 & \mathbf{D}_2 e^{2\pi i t_2 \Delta_f} \\ \vdots & \vdots \\ \mathbf{D}_L & \mathbf{D}_L e^{2\pi i t_L \Delta_f} \end{bmatrix} \tag{5.7}$$

with $\mathbf{D}_l = \text{diag}\{e^{2\pi i \boldsymbol{\phi}_n t_l}\}$ and $\mathbf{F}_\text{B} = \text{diag}\{\mathbf{F}_l\}$. While,

$$\mathbf{A}_{\phi_n} = \mathbf{F}_\text{B}\mathbf{C}_n \text{ where } \mathbf{C}_n = \begin{bmatrix} \hat{\mathbf{C}}_1 \\ \hat{\mathbf{C}}_2 \\ \vdots \\ \hat{\mathbf{C}}_L \end{bmatrix} \tag{5.8}$$

with $\hat{\mathbf{C}}_l = \text{diag}\{2\pi i t_l \mathbf{D}_l(\boldsymbol{\rho}_{w_n} + \boldsymbol{\rho}_{f_n} e^{2\pi i \Delta_f t_l})\}$.

Although $\mathbf{A}_{\phi_n}$ and $\mathbf{A}_{\rho_n}$ contain block matrices with diagonal blocks, the undersampled Fourier transform complicates minimizing both (5.5) and (5.6). There are, however, existing optimization methods that we can adapt to this problem. We present two such methods, one for (5.5) and one for (5.6).

5.3.1 Water and Fat Images Update ($d\boldsymbol{\rho}^{(k+1)}$)

The non-differentiable ℓ_1-norm in the cost function of (5.5) prevents the direct application of traditional descent-based algorithms. Although corner rounding techniques in which the ℓ_1-norm is approximated with a hyperbola make a descent-based approach feasible [21], they result in algorithms with slow convergence, particularly near local minima [22]. Other methods have been proposed for solving similar cost functions without the need for corner rounding (e.g., ISTA and its derivatives [23,24], split-Bregman iterations [25], and augmented Lagrangian (AL) based methods [22]). The AL approach in [22] was developed to minimize the cost function associated with parallel MR image reconstruction which has the same structure as (5.5). That method uses variable splitting to isolate the large matrices (i.e., $\mathbf{F}_\text{B}$, $\mathbf{T}_n$, Ψ) and an AL based minimization strategy to obtain faster convergence than its contemporaries. We therefore adopt the AL based method for this sub-problem.

Following [22], we introduce three splitting variables $\mathbf{u}_0 \in \mathbb{C}^{N \cdot L}$, $\mathbf{u}_1 \in \mathbb{C}^{M_1 \cdot M_2}$, and $\mathbf{u}_2 \in \mathbb{C}^{2N}$ to our cost function (5.5):

$$\underset{d\boldsymbol{\rho},\mathbf{u}_0,\mathbf{u}_1,\mathbf{u}_2}{\arg\min} \; \|\mathbf{F}_{\mathrm{B}}\mathbf{u}_0 - \mathbf{z}_n^{(k)}\|_2^2 + \lambda_\rho\|\mathbf{u}_1 + \mathbf{d}_n\|_1 \;\; s.t. \;\; \mathbf{u}_0 = \mathbf{T}_n d\boldsymbol{\rho}, \; \mathbf{u}_1 = \boldsymbol{\Psi}\mathbf{u}_2, \mathbf{u}_2 = d\boldsymbol{\rho} \tag{5.9}$$

where $\mathbf{z}_n^{(k)} = \mathbf{y} - \mathbf{g}(\mathbf{x}_n) - \mathbf{A}_{\phi_n} d\boldsymbol{\phi}^{(k)}$ and $\mathbf{d}_n = \boldsymbol{\Psi}\boldsymbol{\rho}_n$. Minimizing this new cost function is equivalent to solving the original problem in (5.5).

Introducing three vectors of scaled Lagrange multipliers $\boldsymbol{\eta}_0 \in \mathbb{C}^{N\cdot L}$, $\boldsymbol{\eta}_1 \in \mathbb{C}^{M_1\cdot M_2}$, and $\boldsymbol{\eta}_2 \in \mathbb{C}^{2N}$ leads to the following AL cost function:[2]

$$\begin{aligned} &\|\mathbf{F}_{\mathrm{B}}\mathbf{u}_0 - \mathbf{z}_n^{(k)}\|_2^2 + \lambda_\rho\|\mathbf{u}_1 + \mathbf{d}_n\|_1 + \\ &\qquad \frac{\mu}{2}\|\mathbf{u}_0 - \mathbf{T}_n d\boldsymbol{\rho} - \boldsymbol{\eta}_0\|_2^2 + \frac{\mu\nu_1}{2}\|\mathbf{u}_1 - \boldsymbol{\Psi}\mathbf{u}_2 - \boldsymbol{\eta}_1\|_2^2 + \frac{\mu\nu_2}{2}\|\mathbf{u}_2 - d\boldsymbol{\rho} - \boldsymbol{\eta}_2\|_2^2, \end{aligned} \tag{5.10}$$

where $\mu, \nu_1, \nu_2 > 0$ are AL penalty parameters that influence the convergence rate of the algorithm but do not affect the final estimate for convex problems [22].

We use an alternating minimization strategy outlined in [22] to approximately minimize the AL based cost function (5.10) jointly in terms of $d\boldsymbol{\rho}, \mathbf{u}_0, \mathbf{u}_1, \mathbf{u}_2$:

$$\mathbf{u}_0^{(j+1)} = \mathbf{H}_\mu^{-1}\left[\mathbf{F}_{\mathrm{B}}^{\mathrm{H}}\mathbf{z}_n^{(k)} + \frac{\mu}{2}(\mathbf{T}_n d\boldsymbol{\rho}^{(j)} + \boldsymbol{\eta}_0^{(j)})\right], \tag{5.11}$$

$$\mathbf{u}_1^{(j+1)} = \mathbf{shrink}\left\{\boldsymbol{\Psi}\mathbf{u}_2^{(j)} + \boldsymbol{\eta}_1^{(j)} + \mathbf{d}_n, \frac{\lambda}{\mu\nu_1}\right\} - \mathbf{d}_n, \tag{5.12}$$

$$\mathbf{u}_2^{(j+1)} = \mathbf{H}_{\nu_1\nu_2}^{-1}\left[\boldsymbol{\Psi}^{\mathrm{H}}(\mathbf{u}_1^{(j+1)} - \boldsymbol{\eta}_1^{(j)}) + \frac{\nu_2}{\nu_1}(d\boldsymbol{\rho}^{(j)} + \boldsymbol{\eta}_2^{(j)})\right], \tag{5.13}$$

$$d\boldsymbol{\rho}^{(j+1)} = \mathbf{H}_{\nu_2}^{-1}\left[\mathbf{T}_n^{\mathrm{H}}(\mathbf{u}_0^{(j+1)} - \boldsymbol{\eta}_0^{(j)}) + \nu_2(\mathbf{u}_2^{(j+1)} - \boldsymbol{\eta}_2^{(j)})\right] \tag{5.14}$$

where $\mathbf{shrink}\{\mathbf{v}, \alpha\} \triangleq \mathrm{sign}\{v_i\}\max\{|v_i| - \alpha, 0\} \; \forall \; v_i \in \mathbf{v}$ and[3]

$$\mathbf{H}_\mu = \mathbf{F}_{\mathrm{B}}^{\mathrm{H}}\mathbf{F}_{\mathrm{B}} + \frac{\mu}{2}\mathbf{I}, \tag{5.15}$$

$$\mathbf{H}_{\nu_1\nu_2} = \boldsymbol{\Psi}^{\mathrm{H}}\boldsymbol{\Psi} + \frac{\nu_2}{\nu_1}\mathbf{I}, \tag{5.16}$$

$$\mathbf{H}_{\nu_2} = \mathbf{T}_n^{\mathrm{H}}\mathbf{T}_n + \nu_2\mathbf{I}. \tag{5.17}$$

---

[2] Section 2.6 outlines the relationship between this scaled formulation and the traditional AL formulation.

[3] For complex valued data, $\mathrm{sign}\{v_i\} \triangleq v_i/|v_i| \; \forall \; v_i \neq 0$.

AL Algorithm for Computing $d\boldsymbol{\rho}^{(k+1)}$

```
Initialize variables and set $j = 0$.
Repeat until stop criterion is achieved:
    1. $\mathbf{u}_0^{(j+1)}$ using (5.11).
    2. $\mathbf{u}_1^{(j+1)}$ using (5.12).
    3. $\mathbf{u}_2^{(j+1)}$ using (5.13).
    4. $d\boldsymbol{\rho}^{(j+1)}$ using (5.14).
    5. $\boldsymbol{\eta}_0^{(j+1)} = \boldsymbol{\eta}_0^{(j)} - (\mathbf{u}_0^{(j+1)} - \mathbf{T}_n d\boldsymbol{\rho}^{(j+1)})$
    6. $\boldsymbol{\eta}_1^{(j+1)} = \boldsymbol{\eta}_1^{(j)} - (\mathbf{u}_1^{(j+1)} - \boldsymbol{\Psi}\mathbf{u}_2^{(j+1)})$
    7. $\boldsymbol{\eta}_2^{(j+1)} = \boldsymbol{\eta}_2^{(j)} - (\mathbf{u}_2^{(j+1)} - d\boldsymbol{\rho}^{(j+1)})$
    8. $j = j + 1$.
```

Figure 5.1: AL based minimization algorithm for solving (5.5) [22].

The matrix inversions in these update steps can be efficiently computed for typical CS-WF parameters. $\mathbf{H}_\mu$ can be inverted using fast Fourier transforms (FFTs). $\mathbf{H}_{\nu_2}$ is a 2×2 block matrix whose blocks are scalar multiples of the identity matrix; thus, $\mathbf{H}_{\nu_2}^{-1}$ has a trivial, analytic solution. Since the regularizer in (5.5) is edge preserving, shift-invariant sparsifying transforms can be used for $\boldsymbol{\Psi}_w$ and $\boldsymbol{\Psi}_f$ without the wrap-around artifacts seen in Section 3.2.6. Thus, $\mathbf{H}_{\nu_1\nu_2}$ can be inverted using FFTs (see Section 3.1.2). If shift-variant sparsifying transforms are used, (5.13) can be updated using a few iterations of a CG algorithm with a circulant preconditioner and warm starting [22]. Combining these update steps with the Lagrange multiplier updates yields our method for updating $d\boldsymbol{\rho}^{(k+1)}$ presented in Fig. 5.1.

5.3.2 Field Map Update ($d\boldsymbol{\phi}^{(k+1)}$)

As demonstrated in Section 3.2.6, the finite differencing matrix within the quadratic regularization term Φ must be shift-variant to avoid introducing wrap-around artifacts into the estimates. Furthermore, the Hessian matrix of the cost function in (5.6) is not sparse due to the undersampled Fourier transform in the data-fit term. Thus, non-iterative methods based on FFTs (i.e., circulant matrices) or sparse Cholesky factorization cannot be used to minimize (5.6). These approaches do, however, provide effective preconditioners for descent-based algorithms as demonstrated in Chapter IV. We therefore solve (5.6) using the CG method preconditioned with the sparse Hessian matrix of (5.6) for the case of fully sampled data:

$$\mathbf{P}_{\phi_n} = \mathbf{C}_n^{\mathrm{H}} \mathbf{C}_n + \lambda_\phi \mathbf{\Phi}^{\mathrm{H}} \mathbf{\Phi}, \tag{5.18}$$

which is implemented using sparse Choleksy factorization [26] like in Section 4.1.2. Although the accuracy of this preconditioner depends on the level of undersampling, it was highly effective for our 2-D image experiments over a wide range of sampling rates. As discussed in Section 4.1.2, the sparse Cholesky factorization method may become intractable for large data sets due to memory constraints. Section 5.5 proposes other strategies that are better suited for solving (5.6) in these situations.

5.3.3 Initialization Method

The accuracy of the CS-WF initialization is crucial for obtaining a desirable local minimum as well as reducing the overall computation time. The initialization strategy in [1] reconstructs each scan image independently using CS techniques [21] and then estimates the field map from these images using a region growing method. Although this approach avoids phase wrapping, the resulting field map initialization may have errors as the scan images used in its creation were reconstructed from much less data than (5.2).[4] Using this field map estimate $\hat{\boldsymbol{\phi}}$, the water and fat images are

---

[4] Changing the sampling pattern between scans ensures that the overall CS reconstruction in (5.2) has a higher effective rate of k-space sampling than any one of the individual scans.

estimated voxel-wise by solving the following least-squares problem:

$$\begin{bmatrix} \hat{\boldsymbol{\rho}}_w(j) \\ \hat{\boldsymbol{\rho}}_f(j) \end{bmatrix} = \underset{\boldsymbol{\rho}_w(j),\boldsymbol{\rho}_f(j)}{\arg\min} \left\| \mathbf{B}(j)\mathbf{M} \begin{bmatrix} \boldsymbol{\rho}_w(j) \\ \boldsymbol{\rho}_f(j) \end{bmatrix} - \mathbf{s}(j) \right\|_2^2 = \left[\mathbf{M}^{\mathrm{H}}\mathbf{M}\right]^{-1} \mathbf{M}^{\mathrm{H}}\mathbf{B}(j)^{\mathrm{H}}\mathbf{s}(j) \tag{5.19}$$

where j denotes the voxel index, $\mathbf{B}(j) = \mathrm{diag}\left\{\left[e^{2\pi i\hat{\boldsymbol{\phi}}(j)t_1}, e^{2\pi i\hat{\boldsymbol{\phi}}(j)t_2}, \ldots, e^{2\pi i\hat{\boldsymbol{\phi}}(j)t_L}\right]\right\}$,

$$\mathbf{M} = \begin{bmatrix} 1 & e^{2\pi i t_1 \Delta_f} \\ 1 & e^{2\pi i t_2 \Delta_f} \\ \vdots & \vdots \\ 1 & e^{2\pi i t_L \Delta_f} \end{bmatrix}, \text{ and } \mathbf{s}(j) = \begin{bmatrix} \mathbf{s}_1(j) \\ \vdots \\ \mathbf{s}_L(j) \end{bmatrix} \tag{5.20}$$

with $\mathbf{s}_l$ denoting the lth reconstructed image [1].

We follow the same initialization strategy as [1] except, instead of using a region growing method, we estimate a smooth field map using our fast regularized field map estimator from Chapter IV. This estimator reduces any inaccuracies resulting from the CS reconstruction artifacts and can yield an initialization that is close to the final field map estimated by (5.2). Furthermore, using this smooth field map estimate in (5.19) yields more accurate water and fat initialization images than using field maps containing reconstruction artifacts.

Ideally, the effective level of regularization used in the regularized field map estimator from Chapter IV matches that of the CS-WF method. However, using the same regularization parameter values in both algorithms does not lead to the same effective level of regularization due to differences in their cost functions. In the following section, we present a method to assist in selecting the parameter values for both algorithms such that they have similar effective levels of regularization.

CS-WF resolution analysis

We use resolution analysis (as outlined in Section 2.8) to assist in determining equivalent regularization parameter values for the regularized field map estimator in Chapter IV and the CS-WF method (5.2). Specifically, we analyze the local impulse response of the estimators [27].

The regularized field map estimation cost function (4.7) is a penalized-likelihood in which the data-fit term is minimized by the truth in the absence of noise. We therefore use the penalized-likelihood form of the impulse response (2.44) giving

$$\begin{aligned} \bar{\mathbf{l}}_R^{(j)}(\mathbf{y};\boldsymbol{\phi}) &\approx \left(\mathbf{R}_j + \beta \boldsymbol{\Phi}^{\mathrm{H}} \boldsymbol{\Phi}\right)^{-1} \mathbf{R}_j \mathbf{e}_j, \\ &= \left(\mathbf{I} + \beta \mathbf{R}_j^{-1} \boldsymbol{\Phi}^{\mathrm{H}} \boldsymbol{\Phi}\right)^{-1} \mathbf{e}_j, \\ &\approx \left(\mathbf{I} + \frac{\beta}{r_j} \boldsymbol{\Phi}^{\mathrm{H}} \boldsymbol{\Phi}\right)^{-1} \mathbf{e}_j \end{aligned} \tag{5.21}$$

where $\mathbf{R}_j = \operatorname{diag}\{r_j\}$ with $r_j = \sum_{m=1}^{L} \sum_{p=1}^{L} |g_{jmp}| \, (t_m - t_p)^2$ and g_{jmp} as defined in Chapter IV.

To determine the resolution properties of the CS-WF field map estimate, we assume we have reasonable estimates of the fat and water images $\boldsymbol{\rho}$. To address the nonlinear data-fit term, we follow [27] and first linearize the original cost function (5.2) around the current estimate $\boldsymbol{\phi}_n$ to get

$$\|\mathbf{z} - \mathbf{A}_{\phi_n} \boldsymbol{\phi}\|_2^2 + \lambda_\phi \|\boldsymbol{\Phi} \boldsymbol{\phi}\|_2^2, \tag{5.22}$$

where $\mathbf{z} = \mathbf{y} - \mathbf{g}([\boldsymbol{\rho}, \boldsymbol{\phi}_n]) + \mathbf{A}_{\phi_n} \boldsymbol{\phi}_n$. This linearized cost function has similar properties to a penalized-likelihood and we therefore use the penalized-likelihood form of the impulse response (2.44) which only requires the Hessian matrices of the data-fit and regularization terms evaluated at the truth. The Hessian matrix for the data-fit term evaluated at the truth is

$$2\mathbf{A}_{\phi_n}^{\mathrm{H}}\mathbf{A}_{\phi_n} = 2\sum_{l=1}^{L}(2\pi t_l)^2 \mathrm{diag}\left\{\left[\boldsymbol{\rho}_w + \boldsymbol{\rho}_f e^{2\pi i t_l \Delta_f}\right] \cdot e^{2\pi i t_l \boldsymbol{\phi}_n}\right\}^{\mathrm{H}} \cdot \mathbf{F}_l^{\mathrm{H}}\mathbf{F}_l \mathrm{diag}\left\{\left[\boldsymbol{\rho}_w + \boldsymbol{\rho}_f e^{2\pi i t_l \Delta_f}\right] \cdot e^{2\pi i t_l \boldsymbol{\phi}_n}\right\}, \tag{5.23}$$

while the Hessian matrix for the regularization term is

$$2\lambda_\phi \boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{\Phi}. \tag{5.24}$$

To simplify the analysis, we assume fully sampled Cartesian data (i.e., $\mathbf{F}_l^{\mathrm{H}}\mathbf{F}_l = \mathbf{F}^{\mathrm{H}}\mathbf{F} = \mathbf{I} \ \forall \ l$), and thus,

$$\begin{aligned} 2\mathbf{A}_{\phi_n}^{\mathrm{H}}\mathbf{A}_{\phi_n} = 2\mathbf{C}_n^{\mathrm{H}}\mathbf{C}_n &= 2\sum_{l=1}^{L}(2\pi t_l)^2 \mathrm{diag}\left\{\left|\boldsymbol{\rho}_w(j) + \boldsymbol{\rho}_f(j) e^{2\pi i t_l \Delta_f}\right|^2\right\}, \\ &\approx 2\sum_{l=1}^{L}(2\pi t_l)^2 \mathrm{diag}\left\{\left|s_l(j)\right|^2\right\}, \end{aligned} \tag{5.25}$$

where $s_l(j)$ is the jth voxel of the lth reconstructed image, which we assume to have similar magnitude to the true combined water-fat image (a reasonable assumption for high quality CS reconstructions). As this approximate Hessian of the data-fit term is diagonal, the local impulse response is

$$\begin{aligned} \bar{\mathbf{l}}_C^{(j)}(\mathbf{y}; \boldsymbol{\phi}) &\approx \left(2\mathbf{C}_n^{\mathrm{H}}\mathbf{C}_n + 2\lambda_\phi \boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{\Phi}\right)^{-1} 2\mathbf{C}_n^{\mathrm{H}}\mathbf{C}_n \mathbf{e}_j, \\ &= \left(\mathbf{I} + \lambda_\phi \left[\mathbf{C}_n^{\mathrm{H}}\mathbf{C}_n\right]^{-1} \boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{\Phi}\right)^{-1} \mathbf{e}_j, \\ &\approx \left(\mathbf{I} + \frac{\lambda_\phi}{a_j}\boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{\Phi}\right)^{-1} \mathbf{e}_j \end{aligned} \tag{5.26}$$

where $a_j = \sum_{l=1}^{L}(2\pi t_l)^2 \left|s_l(j)\right|^2$.

Since both impulse responses are shift variant, we follow [28] and consider the median values.[5] Assuming the typical case where both estimators use the same

---

[5] Ignoring the 2π associated with the Rad/s to Hz conversion, the two median values are close to one another for typical data sets.

finite differencing matrix, we can achieve similar degrees of effective regularization by selecting the regularization parameters β and λ_ϕ that balance

$$\frac{\beta}{\textbf{median}\left\{\sum_{m=1}^{L}\sum_{p=1}^{L}\left|g_{jmp}\right|(t_m - t_p)^2\right\}} = \frac{\lambda_\phi}{\textbf{median}\left\{\sum_{l=1}^{L}(2\pi t_l)^2\left|s_l(j)\right|^2\right\}}. \tag{5.27}$$

5.4 Results

This section compares our proposed alternating minimization based algorithm (GN-AM) to the existing linearized, corner rounding CG method (GN-CR) [1] using the same simulated and in-vivo water-fat imaging data sets as in Chapter IV.

We followed [1] and used shift-invariant first-order finite differencing matrices for the sparsifying transforms $\boldsymbol{\Psi}_f, \boldsymbol{\Psi}_w$ and a second-order finite differencing matrix with non-periodic boundary conditions for $\boldsymbol{\Phi}$ in all of our experiments. For all of the algorithms, we ran 10 outer iterations updating the estimates using $\mathbf{x}_{n+1} = \mathbf{x}_n + td\mathbf{x}$ where t was computed using a backtracking line search like in [1]. For each data set, we used two k-space sampling patterns. The first pattern provided an undersampling factor of 2.5 per scan by randomly undersampling approximately 40% of the k-space data using a uniform distribution, while including the center 33×33 k-space samples. The second pattern provided an undersampling factor of 5 per scan by randomly undersampling approximately 20% of the k-space data using a uniform distribution, while including the same center k-space samples. In both cases, the undersampling pattern differed between scans to increase the amount of information available for the joint reconstruction.[6]

For our GN-AM algorithm, we used only one alternating minimization iteration and two iterations of our PCG $d\boldsymbol{\phi}^{(k+1)}$ update algorithm. As in [22], we selected

---

[6] When using the typical three scans, the overall number of samples for the 2.5 and 5 times undersampling factors is 120% and 60% of one fully sampled scan, respectively.

the AL parameters μ, ν_1, ν_2 for our $d\boldsymbol{\rho}^{(k+1)}$ update algorithm using the condition numbers of the Hessian matrices in the variable update steps of Fig. 5.1. The actual parameter values differed between data sets and are therefore presented in the corresponding subsections. For the GN-CR method, we used the corner rounding parameter $h = 10^{-8}$, an Armijo backtracking line search ($c_1 = 10^{-6}, c_2 = 0.2$) [29], and determined the conjugate gradient direction using the Fletcher-Reeves method [30] with a direction reset every 100 iterations.[7] To provide a fair comparison, we used the following diagonal preconditioner for the GN-CR method to balance the weight of the water-fat image and field map estimation variables:

$$\mathbf{P}_n \triangleq \begin{bmatrix} (2L)\cdot\mathbf{I} + \frac{\lambda_\rho}{\sqrt{h}}\mathbf{D}_\Psi & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & (2L)\cdot\mathbf{I} + \frac{\lambda_\rho}{\sqrt{h}}\mathbf{D}_\Psi & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & 2\mathbf{C}_n^{\mathrm{H}}\mathbf{C}_n + 2\lambda_\phi\mathbf{D}_\Phi \end{bmatrix} \tag{5.28}$$

where $\mathbf{D}_\Psi$ and $\mathbf{D}_\Phi$ are the diagonals of $\boldsymbol{\Psi}_w^{\mathrm{H}}\boldsymbol{\Psi}_w$ and $\boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{\Phi}$ respectively.[8] We ran our AL based $d\boldsymbol{\rho}^{(k+1)}$ update algorithm and the existing corner rounded CG method until the normalized root-mean squared difference (NRMSD) between the two most recent $\boldsymbol{\rho}$ estimates was less than 0.1%, up to a maximum of 150 and 350 iterations for the GN-AM and GN-CR methods, respectively.[9] All of the algorithms were implemented in MATLAB (The MathWorks, Natick, MA, USA) and the experiments were run on a PC with a 2.66 GHz, quad-core Intel Xeon CPU.

We computed the initialization images for both algorithms using our method from Section 5.3.3. We used ISTA [23] and a level-4, orthonormal, Haar wavelet transform

---

[7] The Polak-Ribiére method [31] required significantly more resets, resulting in a slower overall convergence rate for our experiments (results not shown).

[8] We used an upper bound of the approximate gradient of the linearized cost function with respect to $\boldsymbol{\rho}$, $2\mathbf{A}_{\rho_n}^{\mathrm{H}}\mathbf{A}_{\rho_n} + \lambda_\rho\boldsymbol{\Psi}^{\mathrm{H}}\mathbf{W}_n^{-1}\boldsymbol{\Psi}$ where $\mathbf{W}_n = \mathrm{diag}\{w_i\}$ and $w_i = \sqrt{[\boldsymbol{\Psi}(\boldsymbol{\rho}_n + d\boldsymbol{\rho})]_i^*[\boldsymbol{\Psi}(\boldsymbol{\rho}_n + d\boldsymbol{\rho})]_i + h}$ [21], as it provides an appropriate level of scaling while requiring significantly less computation.

[9] This stopping criteria matched closely with the empirically determined optimal number of iterations for the first few outer iterations.

to compute the CS reconstruction of each scan image. We then estimated the field maps using the QS-Huber implementation from Chapter IV with a second-order finite differencing matrix. The regularization parameter for the QS-Huber algorithm was determined using (5.27) and the particular λ_ϕ specified for each experiment (details of which are in the corresponding subsections). We normalized the k-space data for each experiment by dividing by the maximum absolute value of the independent CS reconstructions.

Although not guaranteed, we found that all of the methods converged to similar estimates. Furthermore, due to the accuracy of our initialization method, the field map variables converged by less than 1 Hz RMSD over the object pixels in all of our experiments. We therefore compared the convergence rates of the algorithms using the NRMSD between the $\boldsymbol{\rho}$ estimate at each iteration and the final $\boldsymbol{\rho}$ estimate of each algorithm. Although the estimates were computed over the entire field-of-view, we restricted the calculation of the NRMSD to a masked region consisting of a convex hull surrounding those pixels containing signal.[10]

5.4.1 Simulated Water-Fat Imaging Data

We used the same simulated water-fat imaging data as in Chapter IV. In particular, we combined the 256 × 256 pixel water and fat images with the simulated field map (Fig. 5.2) to create three acquisition images with relative echo times $t_\ell = -0.4, 1.2, 2.8\,\text{ms}$, $\Delta f = 220\,\text{Hz}$, and $\text{R}_2^* = 20\,\text{s}^{-1}$. We randomly undersampled the data using the previously described sampling schemes to obtain two data sets with 2.5 and 5 times undersampling factors. Complex Gaussian noise was added to the k-space data so that the SNR $\approx 26\,\text{dB}$ and $23\,\text{dB}$ for the 2.5 and 5 times un-

---

[10] The relative NRMSD convergence rates of the algorithms were similar when calculated over the entire field-of-view (results not shown).

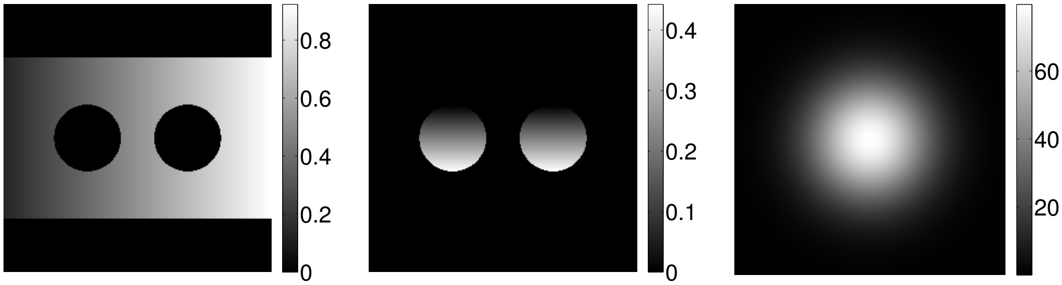

Figure 5.2: The magnitudes of the true water image (left), fat image (center), and field map in Hz (right) used to create our simulated data set.

dersampling factors, respectively. For the GN-AM method, we determined the AL parameters for the first iteration by setting the condition numbers of $\mathbf{H}_\mu, \mathbf{H}_{\nu_1\nu_2}, \mathbf{H}_{\nu_2}$ to approximately $2, 2, 1.06$, respectively. For subsequent iterations, the AL parameters were determined by setting the condition numbers to $5, 17,$ and 1.

For the 2.5 times undersampled data, we set the regularization parameters to $\lambda_\rho = 2^{-6}, \lambda_\phi = 2^{-6}$ as this was found to provide the most accurate estimates compared to the truth (details not shown). For the 5 times undersampled data, we used the same approach to select $\lambda_\rho = 2^{-5}, \lambda_\phi = 2^{-6}$. Figs. 5.3 and 5.4 show the magnitudes of the initialization images used by both algorithms for the cases of 2.5 and 5 times undersampling factors (masked to show detail over the object pixels). Figs. 5.5 and 5.6 show the resulting regularized estimates for both methods and their differences (masked to show detail). Figs. 5.7 and 5.8 are plots of the convergence of the $\boldsymbol{\rho}$ variables in terms of NRMSDs versus time for both methods, while Table 5.1 contains the times required by each algorithm to converge to 1% NRMSD over the mask for both 2.5 and 5 times undersampling factors.

5.4.2 In-vivo Knee Water-Fat Imaging Data

We used the same in-vivo knee water-fat imaging data as in Chapter IV. Specifically, we used a data set consisting of three 256×256 pixel acquisition images with

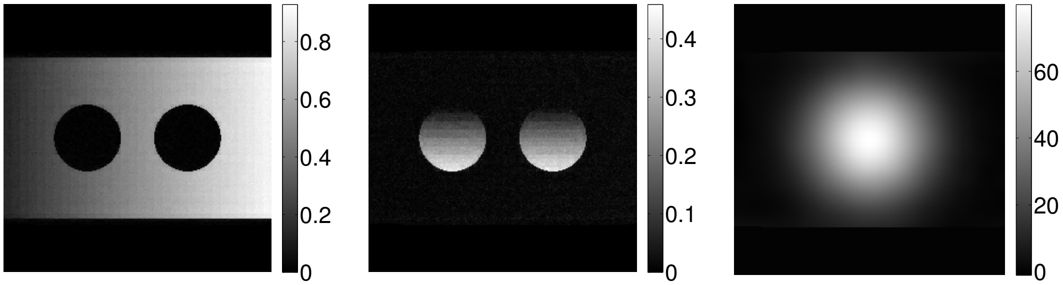

Figure 5.3: The (masked) magnitudes of the 2.5 times undersampled simulated data initialization for the water image (left), fat image (center), and field map in Hz (right) used by both algorithms.

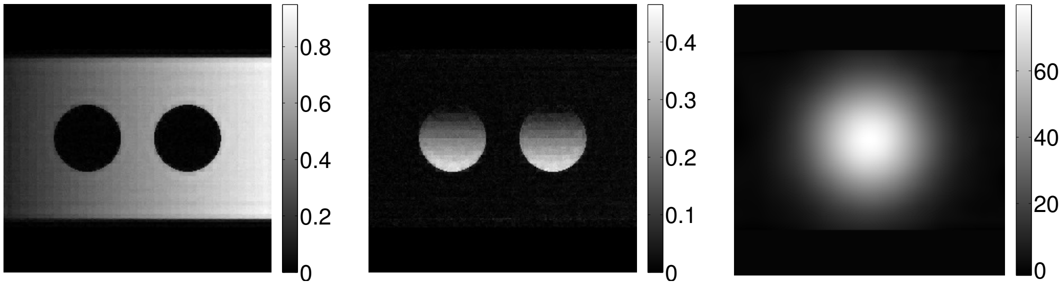

Figure 5.4: The (masked) magnitudes of the 5 times undersampled simulated data initialization for the water image (left), fat image (center), and field map in Hz (right) used by both algorithms.

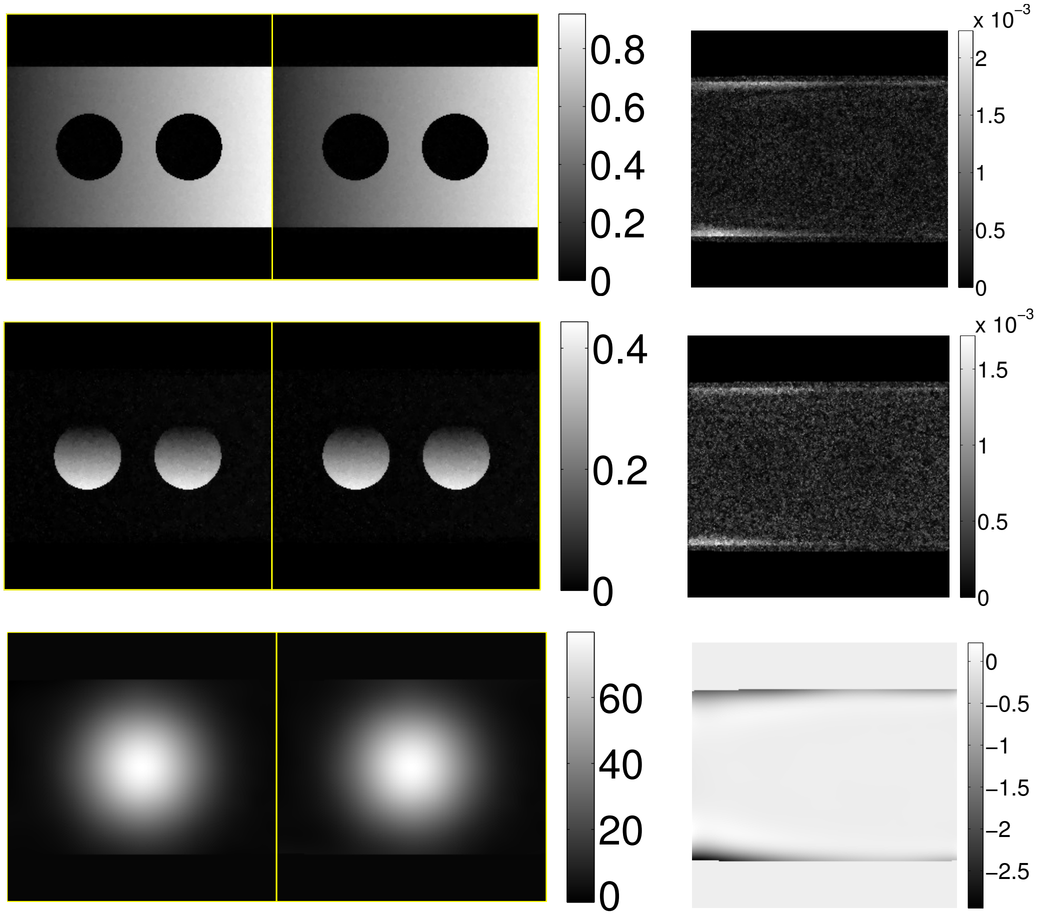

Figure 5.5: The (masked) magnitudes of the final estimate of our GN-AM algorithm (left), the final estimate of the GN-CR algorithm (center), and the difference between these two estimates (right) for the 2.5 times undersampled simulated data. The top row is the water image, the center row is the fat image, and the bottom row is the field map estimate in Hz.

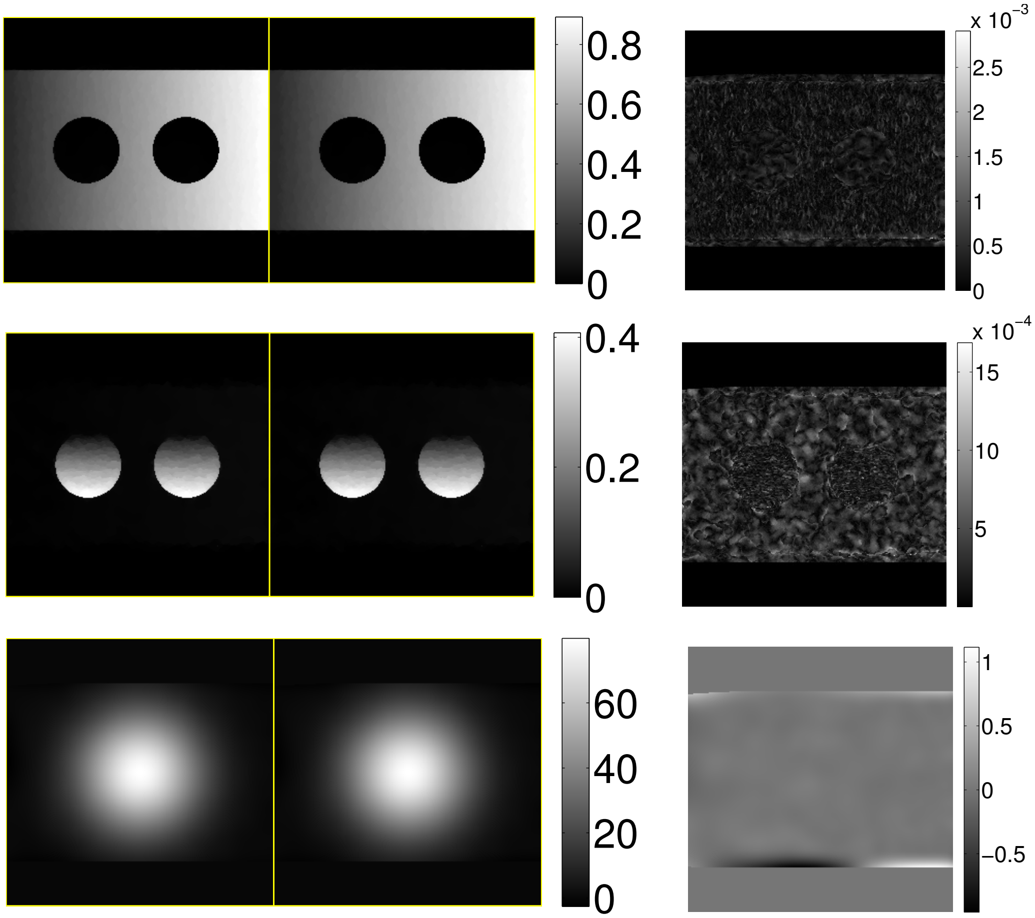

Figure 5.6: The (masked) magnitudes of the final estimate of our GN-AM algorithm (left), the final estimate of the GN-CR algorithm (center), and the difference between the two estimates (right) for the 5 times undersampled simulated data. The top row is the water image, the center row is the fat image, and the bottom row is the field map estimate in Hz.

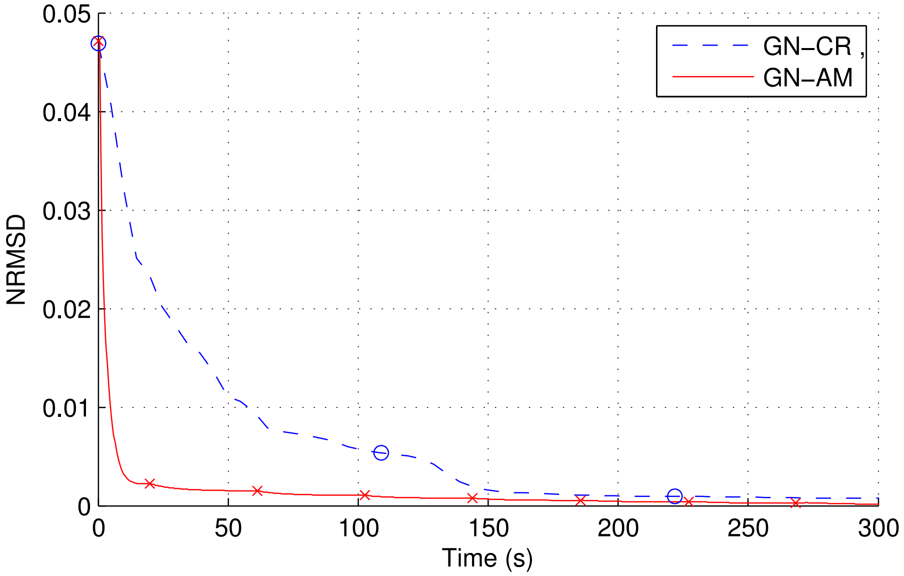

Figure 5.7: The NRMSDs versus time for the 2.5 times undersampled simulated water-fat images estimate $\boldsymbol{\rho} = [\boldsymbol{\rho}_w, \boldsymbol{\rho}_f]$ computed over a mask for both algorithms. The markers designate the outer iterations for each algorithm.

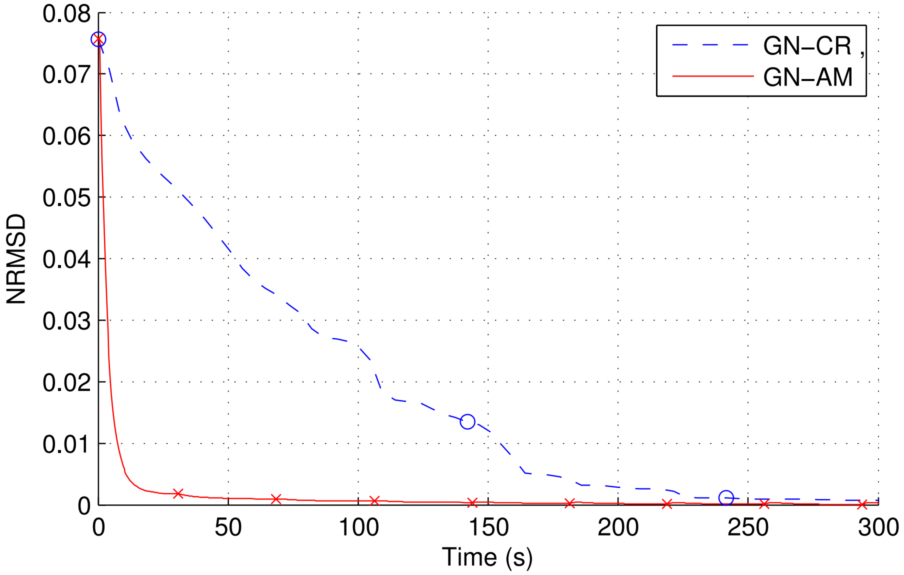

Figure 5.8: The NRMSDs versus time for the 5 times undersampled simulated water-fat images estimate $\boldsymbol{\rho} = [\boldsymbol{\rho}_w, \boldsymbol{\rho}_f]$ computed over a mask for both algorithms. The markers designate the outer iterations for each algorithm.

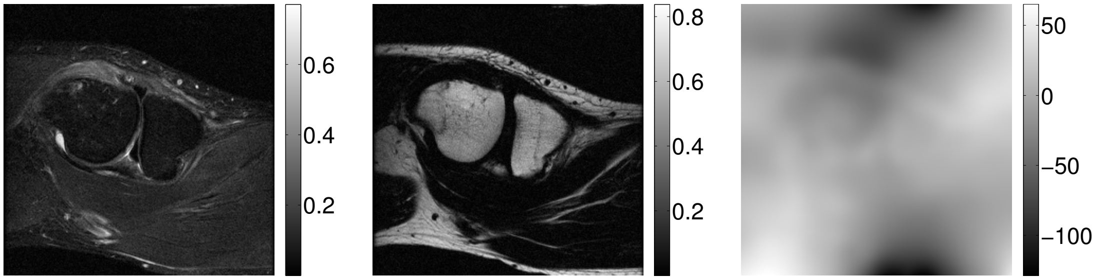

Figure 5.9: The magnitudes of the water image (left), fat image (center), and field map in Hz (right) estimated from the fully sampled in-vivo knee data set using our initialization method.

relative echo times $t_\ell = -0.4, 1.2, 2.8\,\text{ms}$ and $\text{SNR} \approx 35\,\text{dB}$ acquired on a 1.5 T scanner using the IDEAL imaging acquisition protocol [32]. Fig. 5.9 presents the water image, fat image, and field map estimated using our initialization strategy on the fully sampled data set. We randomly undersampled the data using the previously described sampling schemes to obtain data sets with 2.5 and 5 times undersampling factors and $\text{SNR} \approx 40\,\text{dB}$ and $37\,\text{dB}$, respectively. For the GN-AM method, we determined the AL parameters for the first iteration by setting the condition numbers of $\mathbf{H}_\mu, \mathbf{H}_{\nu_1\nu_2}, \mathbf{H}_{\nu_2}$ to approximately $17, 5, 1.06$, respectively. For subsequent iterations, the AL parameters were determined by setting the condition numbers to $5, 17,$ and 1.

For the 2.5 times undersampled data, we selected the regularization parameters $\lambda_\rho = 2^{-7.5}, \lambda_\phi = 2^{-7}$ as these provided high quality water and fat images. While for the 5 times undersampled data, we selected $\lambda_\rho = 2^{-7}, \lambda_\phi = 2^{-7}$. Figs. 5.10 and 5.11 show the magnitudes of the initialization images used by both algorithms for the cases of 2.5 and 5 times undersampling factors (masked to show detail over the object pixels). Figs. 5.12 and 5.13 show the resulting regularized estimates for both methods and their differences (masked to show detail). Figs. 5.14 and 5.15 are plots of the convergence of the $\boldsymbol{\rho}$ variables in terms of NRMSDs versus time for both methods, while Table 5.1 contains the times required by each algorithm to converge

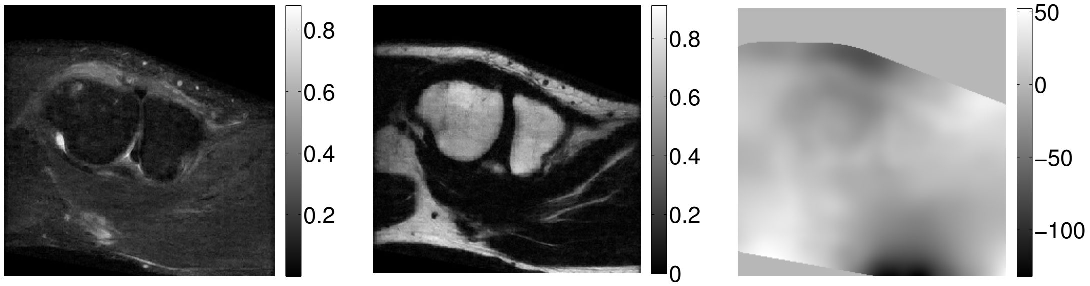

Figure 5.10: The (masked) magnitudes of the 2.5 times undersampled in-vivo knee data initialization for the water image (left), fat image (center), and field map in Hz (right) used by both algorithms.

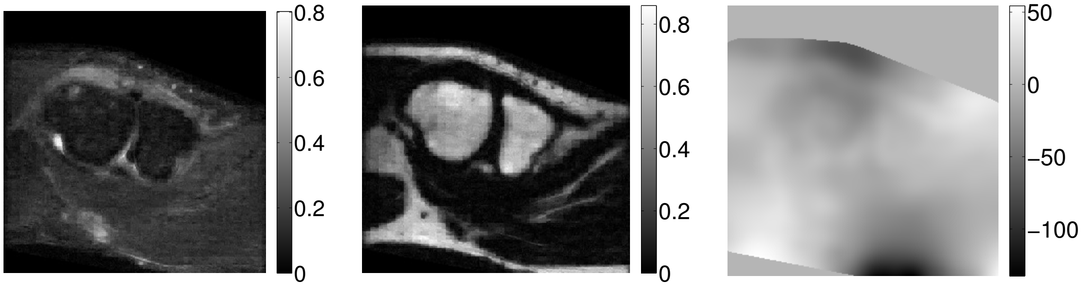

Figure 5.11: The (masked) magnitudes of the 5 times undersampled in-vivo knee data initialization for the water image (left), fat image (center), and field map in Hz (right) used by both algorithms.

to 1% NRMSD over the mask for both 2.5 and 5 times undersampling factors.

5.5 Discussion

Our novel minimization strategy converged at least 12 times faster than the existing NCG based method (GN-CR) for our 2.5 times undersampled data and 21 times faster for our 5 times undersampled data. The field map estimates did not change sig-

Table 5.1: Convergence Time to NRMSD = 1% Over Mask

Data Set	**Undersampling Factor**	**Time to NRMSD = 1%**		**Speed Increase**
		GN-AM	**GN-CR [1]**	
Simulated	2.5	4.8 s	58 s	12 ×
	5	7.3 s	160 s	21 ×
In-vivo Knee	2.5	11 s	270 s	24 ×
	5	12 s	260 s	21 ×

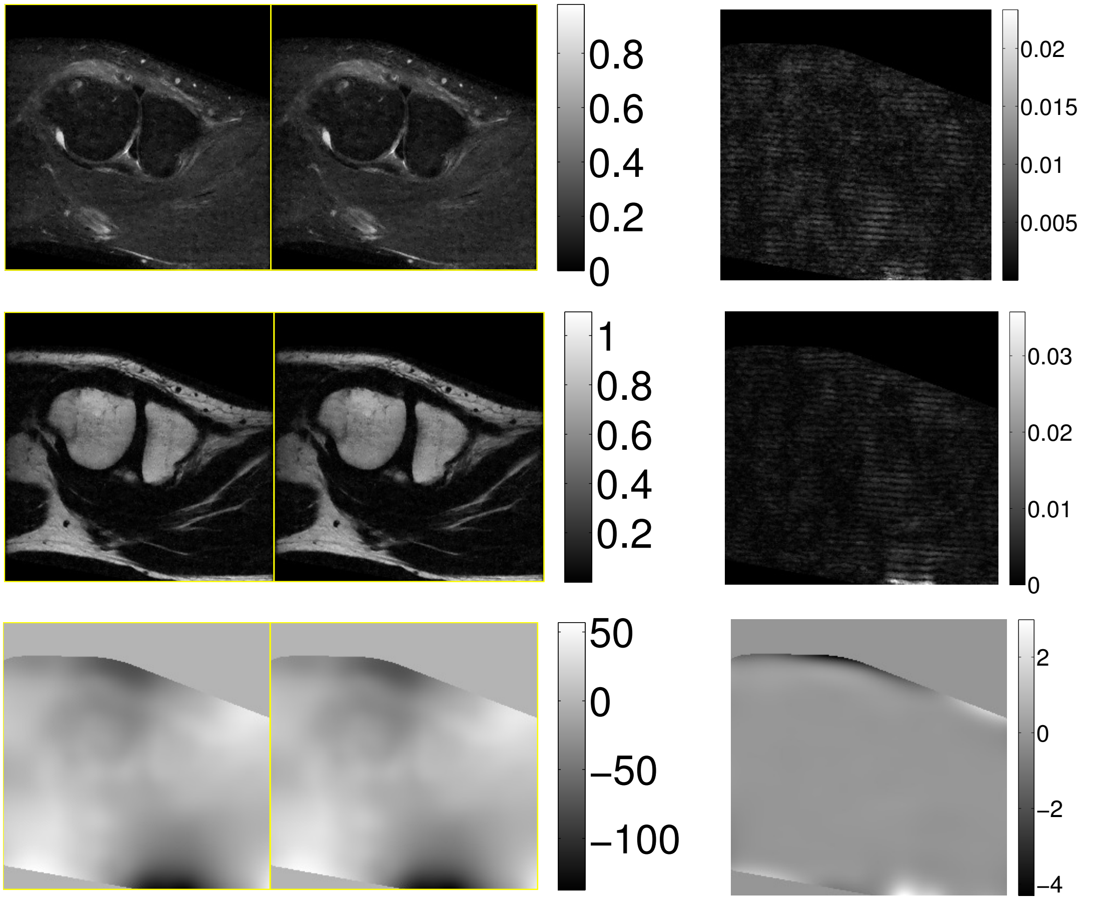

Figure 5.12: The (masked) magnitudes of the final estimate of our GN-AM algorithm (left), the final estimate of the GN-CR algorithm (center), and the difference between the two estimates (right) for the 2.5 times undersampled in-vivo knee data. The top row is the water image, the center row is the fat image, and the bottom row is the field map estimate in Hz.

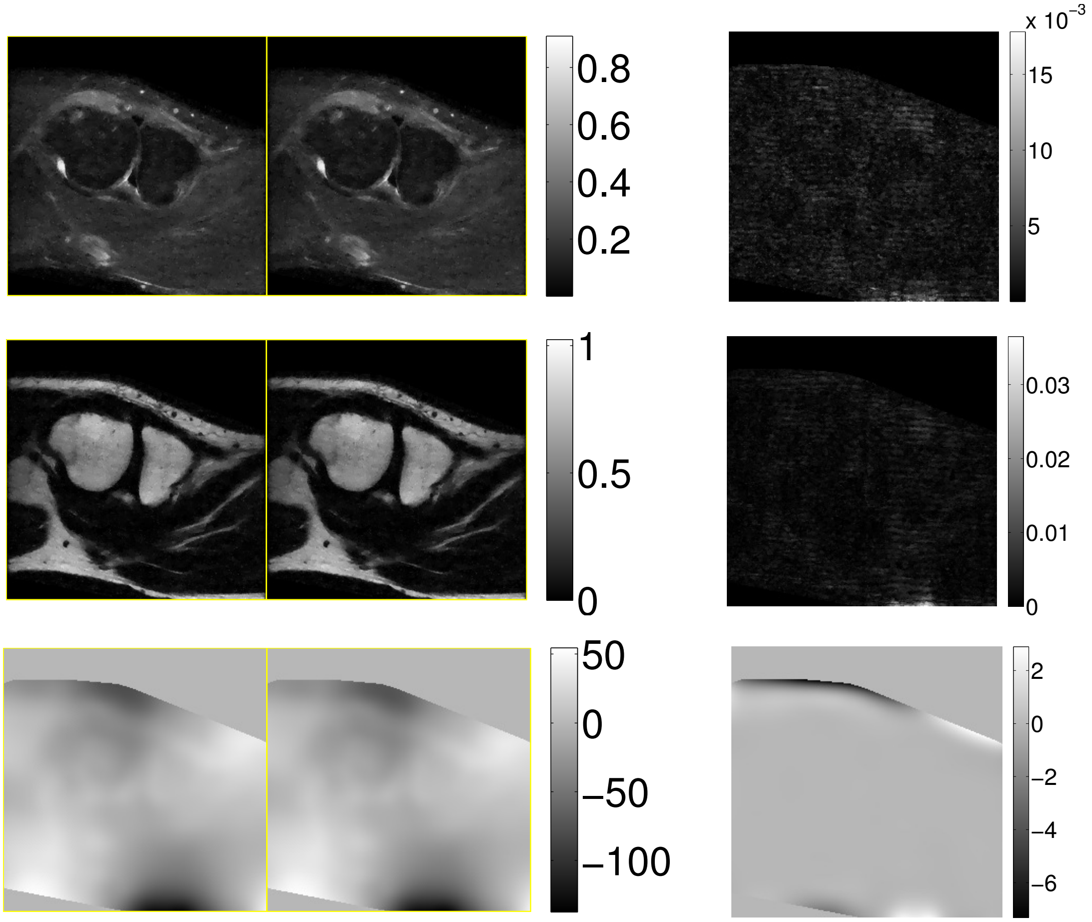

Figure 5.13: The (masked) magnitudes of the final estimate of our GN-AM algorithm (left), the final estimate of the GN-CR algorithm (center), and the difference between the two estimates (right) for the 5 times undersampled in-vivo knee data. The top row is the water image, the center row is the fat image, and the bottom row is the field map estimate in Hz.

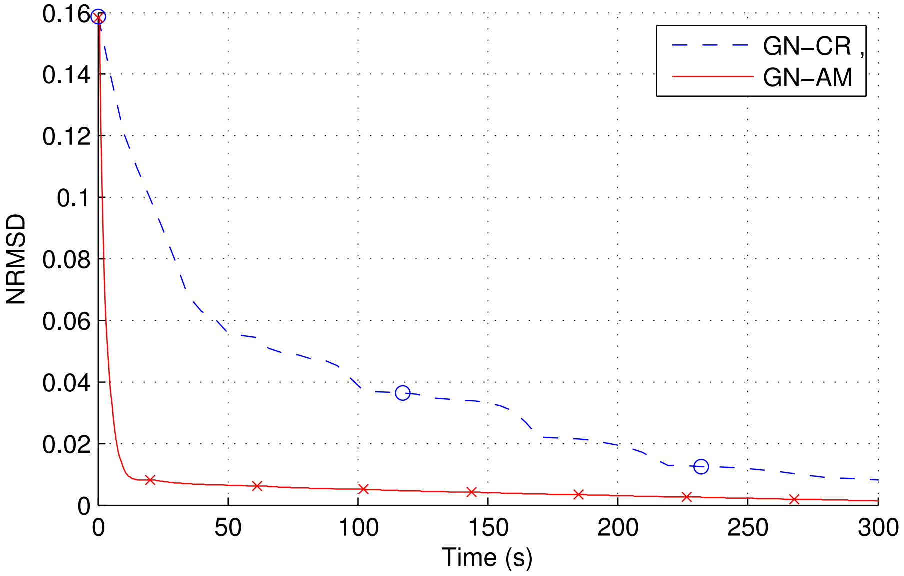

Figure 5.14: The NRMSDs versus time for the 2.5 times undersampled in-vivo knee water-fat images estimate $\boldsymbol{\rho} = [\boldsymbol{\rho}_w, \boldsymbol{\rho}_f]$ computed over a mask for both algorithms. The markers designate the outer iterations for each algorithm.

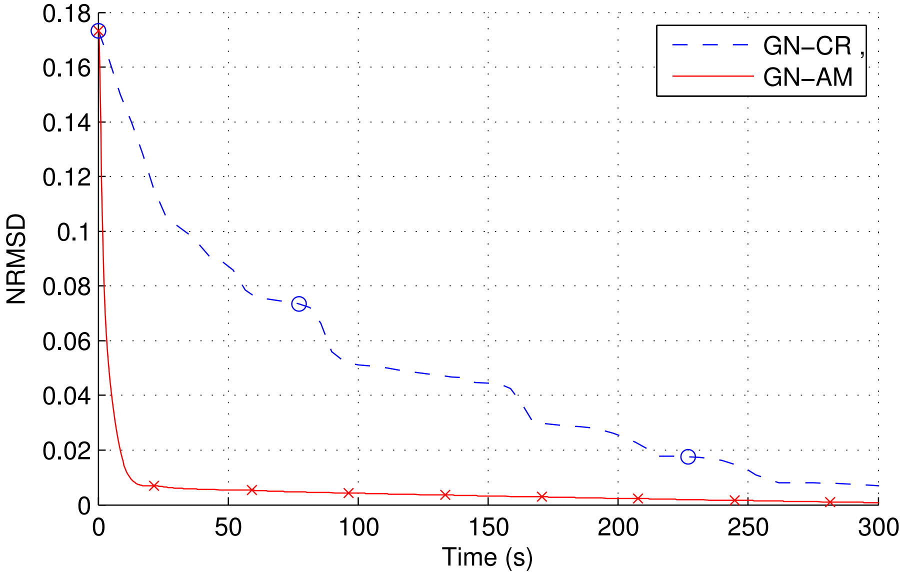

Figure 5.15: The NRMSDs versus time for the 5 times undersampled in-vivo knee water-fat images estimate $\boldsymbol{\rho} = [\boldsymbol{\rho}_w, \boldsymbol{\rho}_f]$ computed over a mask for both algorithms. The markers designate the outer iterations for each algorithm.

nificantly due to the high accuracy of our initialization method. As such, the CS-WF problem could be solved to a practical level of precision in a single outer iteration. Furthermore, our accurate field map initialization meant that our proposed GN-AM method was most efficient when using only one alternating minimization iteration in which each update was computed to a high degree of convergence. The overall convergence speeds of both algorithms were primarily related to the minimization of the water-fat images $\boldsymbol{\rho}$ which the AL based method was capable of performing much faster than the existing corner rounding NCG method. Had we used the existing region growing initialization method [1], both algorithms would likely have required additional outer iterations to converge. Still, our alternating minimization strategy would solve each of these outer iterations faster.

The convergence of both algorithms slows near the local minima in all of our experiments. This is primarily caused by the NRMSD based automated stopping criteria used by the inner minimization algorithms. As both algorithms approach a local minimum, the improvement obtained with each outer step decreases resulting in many additional outer steps being required to converge further. However, at this point only a few iterations of either algorithm are needed to minimize each of the linearized cost functions. Our current automated stopping criteria results in too many inner iterations being performed on these later linearized cost functions, leading to slower than necessary convergence. A better stopping criterion would decrease the maximum allowable inner iterations with each outer iteration; however, this more complicated method was not implemented since a practical solution was always found within the first few outer iterations.

As outlined in Chapter IV, memory constraints can make implementing the Hessian preconditioner of our $d\boldsymbol{\phi}^{(k+1)}$ update step using sparse Cholesky factorization intractable for 3-D data sets. As an alternative, smaller but less accurate preconditioners can be used such as an FFT implemented circulant preconditioner of the form $\mathbf{P}_C = |c|_{\max}\mathbf{I} + \lambda_\phi \Phi_C^{\mathrm{H}} \Phi_C$ where $|c|_{\max}$ is the maximum absolute value of $\mathbf{C}_n^{\mathrm{H}} \mathbf{C}_n$ and Φ_C is a shift-invariant approximation of Φ. Although the accuracy of this preconditioner also depends on the sampling rate, we found that the circulant approximation provided significant acceleration over standard CG even at 20% undersampling (results not shown). Other minimization strategies with less dependence on the sampling rate are also possible. The AL method from our $d\boldsymbol{\rho}^{(k+1)}$ update (5.5) can be used for this problem where the $\mathbf{u}_1$ update becomes trivial and the $\mathbf{u}_2$ update requires the minimization of a quadratic function with the Hessian matrix $\mathbf{H}_\phi = \Phi^{\mathrm{H}}\Phi + \frac{\nu_2}{\nu_1}\mathbf{I}$ (which contains shift-variant terms). Although this update step appears to have similar complexity to the original problem, the lack of an undersampled Fourier transform means that sparse Cholesky factorization can solve this problem directly and our previous preconditioners are more effective. Alternatively, the ADMM algorithm from Chapter III, which was found to be faster than PCG methods for one specific quadratic problem, could be adapted for our $d\boldsymbol{\phi}^{(k+1)}$ update step. However, this would require an additional variable splitting and a more complicated parameter selection process.

The optimal AL parameters μ, ν_1, ν_2 for our $d\boldsymbol{\rho}^{(k+1)}$ update change with each outer iteration due to the update variables $d\boldsymbol{\rho}, d\boldsymbol{\phi}$ converging to zero. In our experiments, using a fixed set of parameter values led to non-monotonic convergence in later $d\boldsymbol{\rho}^{(k+1)}$ updates. Although any negative effects from this non-monotonic convergence can be avoided by running a sufficient number of AL iterations, faster overall convergence is obtained by adjusting the parameter values between outer iterations. We empirically selected two sets of parameter values for each experiment, one for the first $d\boldsymbol{\rho}$ update and the other for subsequent steps. However, determining an automated parameter selection strategy remains an open problem.

The current AL parameter selection strategy outlined in [22], in which the values are determined using condition numbers, does not explicitly account for the undersampling rate. However, we found that this approach was relatively robust to changes in the regularization and sampling rates. To demonstrate this, we used the same condition numbers for both undersampling factors in each experiment. Although we could have further optimized our algorithm for each sampling rate, our GN-AM method remained consistently faster than the existing GN-CR method.

The convergence rate of the existing GN-CR algorithm is highly dependent on the parameters used in the backtracking line search. Too small of a common ratio and the line search requires many expensive cost function evaluations; while too large of a common ratio can result in faulty step sizes. For instance, using the recommended parameter settings from [21] resulted in significantly slower convergence of the GN-CR algorithm compared to our settings. Alternatively, a monotonic step size line search algorithm created using optimization transfer principles as in Chapter IV would avoid this source of complexity. Unlike in Chapter IV, such a monotonic line search strategy would likely provide only minor overall convergence rate improvements compared to the properly tuned backtracking line search algorithm used in our experiments due to the convex nature of the linearized problem.

Although the Gauss-Newton like linearization proposed in [1] is an effective method for handling the nonlinear nature of the field inhomogeneity term in the cost function (particularly with our accurate initialization strategy), it is complicated by having to select a step size for each update. The backtracking line search method provided sufficient step sizes for our experiments, but this is not guaranteed due to the nonconvex nature of the original cost function (5.2). There would be no need for step size selection if a surrogate function for the original cost function could be found like in Chapter IV. The undersampled Fourier transform prevents the direct application of the surrogates in Chapter IV and we must explore other functions such as those used in [33]. However, finding an appropriate surrogate function for CS-WF reconstruction remains an open problem.

The water and fat images reconstructed with the CS-WF method were close to the fully sampled data but did contain a few minor artifacts. For instance, the reconstructed images had some block-like artifacts (particularly the 5 times undersampled data). These were a result of using a first-order finite differencing matrix as the sparsifying transform and they could be reduced by adding a second sparsifying transform to the cost function (e.g., orthonormal wavelets) [21]. As with the field estimation problem shown in Chapter IV, the fat and water components were not fully separated in the in-vivo data. As previously discussed, the separation could be improved by using a multipeak fat model, which would require only minor changes to our algorithms [1]. These modifications to the original cost function (5.2) remain to be explored.

In our experiments, we found that the estimates from 5 times undersampled data had more error than those from 3 times undersampling. This was not surprising as the overall number of samples was far less in the 5 times undersampled case. However, it would be interesting to investigate the trade-off between the undersampling rate and the number of acquisitions for a fixed number of samples. For instance, the same number of samples would be required for three scans with 3 times undersampling as for five scans with 5 times undersampling. Our proposed fast algorithm for computing CS-WF reconstructions, facilitates exploring this question; however, such an investigation requires additional data sets with numerous scans that can be retroactively undersampled.

5.6 Conclusions

We have presented a novel minimization strategy for the compressed sensing based water-fat image reconstruction problem proposed by [1]. Our method uses the same linearization technique as the existing minimization method [1] but solves the linearized problem using an alternating minimization approach in which the water-fat images are updated using an AL method with variable splitting and the field map is updated using a CG algorithm with a sparse Hessian matrix preconditioner. We also introduced a new initialization strategy, based on the regularized field map estimator in Chapter IV, which provides accurate initializations and reduces the number of outer iterations required for convergence. Our novel minimization algorithm converged at least 12 times faster than the existing minimization strategy (using the same initialization) for both simulated and in-vivo experiments. This minimization strategy might also accelerate other CS based water-fat imaging problems that use both linearization and joint minimization techniques [18–20].

Bibliography

[1] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig, "Compressed sensing for chemical shift-based water-fat separation," *Mag. Res. Med.*, vol. 64, no. 6, pp. 1749–59, Dec. 2010.

[2] T. A. Bley, O. Wieben, C. J. François, J. H. Brittain, and S. B. Reeder, "Fat and water magnetic resonance imaging," *J. Mag. Res. Im.*, vol. 31, no. 1, pp. 4–18, Jan. 2010.

[3] H. Eggers and P. Börnert, "Chemical shift encoding-based waterfat separation methods," *J. Mag. Res. Im.*, To appear in 2014.

[4] A. Haase, J. Frahm, W. Hanicke, and W. Matthaei, "1 H NMR chemical shift selective (CHESS) imaging," *Phys. Med. Biol.*, vol. 30, no. 4, pp. 341, 1985.

[5] C. H. Meyer, J. M. Pauly, A. Macovski, and D. G. Nishimura, "Simultaneous spatial and spectral selective excitation," *Mag. Res. Med.*, vol. 15, no. 2, pp. 287–304, Aug. 1990.

[6] F. Schick, "Simultaneous highly selective MR water and fat imaging using a simple new type of spectral-spatial excitation," *Mag. Res. Med.*, vol. 40, no. 2, pp. 194–202, 1998.

[7] W. Dixon, "Simple proton spectroscopic imaging," *Radiology*, vol. 153, no. 1, pp. 189–94, Oct. 1984.

[8] G. H. Glover and E. Schneider, "Three-point Dixon technique for true water/fat decomposition with B_0 inhomogeneity correction," *Mag. Res. Med.*, vol. 18, no. 2, pp. 371–83, Apr. 1991.

[9] 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," *Mag. Res. Med.*, vol. 51, no. 1, pp. 35–45, Jan. 2004.

[10] Q.-S. Xiang and L. An, "Water-fat imaging with direct phase encoding," *J. Mag. Res. Im.*, vol. 7, no. 6, pp. 1002–15, Nov. 1997.

[11] 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," *Mag. Res. Med.*, vol. 59, no. 3, pp. 571–80, Mar. 2008.

[12] 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," *Mag. Res. Med.*, vol. 63, no. 1, pp. 79–90, Jan. 2010.

[13] 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," *J. Mag. Res. Im.*, vol. 31, no. 4, pp. 1027–1034, 2010.

[14] W. Lu and B. A. Hargreaves, “Multiresolution field map estimation using golden section search for water-fat separation,” *Mag. Res. Med.*, vol. 60, no. 1, pp. 236–44, July 2008.

[15] M. Jacob and B. P. Sutton, “Algebraic decomposition of fat and water in MRI,” *IEEE Trans. Med. Imag.*, vol. 28, no. 2, pp. 173–84, Feb. 2009.

[16] W. Huh, J. A. Fessler, and A. A. Samsonov, “Water-fat decomposition with regularized field map,” in *Proc. Intl. Soc. Mag. Res. Med.*, 2008, p. 1382.

[17] J. Machann, C. Thamer, B. Schnoedt, M. Haap, H. Haring, C. D. Claussen, M. Stumvoll, A. Fritsche, and F. Schick, “Standardized assessment of whole body adipose tissue topography by MRI,” *J. Mag. Res. Im.*, vol. 21, no. 4, pp. 455–462, 2005.

[18] S. D. Sharma, H. H. Hu, and K. S. Nayak, “Chemical shift encoded water-fat separation using parallel imaging and compressed sensing,” *Mag. Res. Med.*, vol. 69, no. 2, pp. 456–66, Feb. 2013.

[19] S. D. Sharma, H. H. Hu, and K. S. Nayak, “Accelerated water-fat imaging using restricted subspace field map estimation and compressed sensing,” *Mag. Res. Med.*, vol. 67, no. 3, pp. 650–9, Mar. 2012.

[20] C. N. Wiens, C. M. McCurdy, J. D. Willig-Onwuachi, and C. A. McKenzie, “R2*-corrected waterfat imaging using compressed sensing and parallel imaging,” *Mag. Res. Med.*, vol. 71, no. 2, pp. 608–616, 2014.

[21] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, “Compressed sensing MRI,” *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.

[22] S. Ramani and J. A. Fessler, “Parallel MR image reconstruction using augmented Lagrangian methods,” *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 694–706, Mar. 2011.

[23] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.

[24] A. Beck and M. Teboulle, “Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems,” *IEEE Trans. Im. Proc.*, vol. 18, no. 11, pp. 2419–34, Nov. 2009.

[25] T. Goldstein and S. Osher, “The split Bregman method for L1-regularized problems,” *SIAM J. Imaging Sci.*, vol. 2, no. 2, pp. 323–43, 2009.

[26] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, “Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate,” *ACM Trans. Math. Software*, vol. 35, no. 3, pp. 22:1–22:14, Oct. 2008.

[27] J. A. Fessler and W. L. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs," *IEEE Trans. Im. Proc.*, vol. 5, no. 9, pp. 1346–58, Sept. 1996.

[28] A. K. Funai, J. A. Fessler, D. T. B. Yeo, V. T. Olafsson, and D. C. Noll, "Regularized field map estimation in MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1484–94, Oct. 2008.

[29] L. Armijo, "Minimization of functions having continuous derivatives," *Pacific J. Math*, vol. 16, no. 1, pp. 1–3, 1966.

[30] R. Fletcher and C. M. Reeves, "Function minimization by conjugate gradients," *Comput. J*, vol. 7, no. 2, pp. 149–54, 1964.

[31] E. Polak and G. Ribiere, "Note sur la convergence de méthodes de directions conjuguées," *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique*, vol. 3, pp. 35–43, 1969.

[32] 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," *Mag. Res. Med.*, vol. 54, no. 3, pp. 636–44, Sept. 2005.

[33] C. Yip, W. A. Grissom, J. A. Fessler, and D. C. Noll, "Joint design of trajectory and RF pulses for parallel excitation," *Mag. Res. Med.*, vol. 58, no. 3, pp. 598–604, Sept. 2007.

CHAPTER VI

Conclusions and Future Work

In this thesis, we have explored the regularized estimation of several MR parameters. We have addressed aspects related to the estimators themselves as well as methods for efficiently minimizing their cost functions. There are, however, several topics that remain to be explored.

All of the regularized estimation methods in this work require the selection of at least one regularization parameter. We have provided several heuristic methods for selecting this parameter but no automated processes. There are several possible approaches for this task including Stein's unbiased risk estimate (SURE) based methods (e.g., [1]) and homotopy methods (e.g., [2]). However, all of these approaches face a similar problem in that those images with the lowest error, as quantified by some metric, may not be the best for clinical diagnoses. That being said, an advantage of our accelerated algorithms is that they can efficiently compute several estimates, with differing regularization parameters, allowing for faster parameter selection. This contributes to reducing the overall cost of regularized estimation.

In Chapter III, we proposed an ADMM based algorithm that minimizes the quadratic cost function associated with regularized coil sensitivity estimation in half the time required by a CG method with a circulant preconditioner. A key development in this algorithm was a novel variable splitting strategy that reformulates the shift-variant finite differencing matrix to allow for exact, non-iterative update steps. There are many areas of MR imaging where cost functions with similar structures are used. For example, B_1^+ map estimation can be performed by minimizing cost functions with quadratic regularization terms over the image domain [3]. This variable splitting technique can also be used to facilitate the application of AL methods to regularized estimators with reconstruction masks. The extension to these applications remains unexplored.

In Chapter IV, we developed two minimization methods for the nonconvex cost function associated with regularized main magnetic field map estimation. Our fastest methods used sparse Cholesky factorization to achieve estimation times that were at least 30 times less than the existing SQS method. There are other areas in MR imaging where similar cost functions are encountered [4–8]. Although our algorithms have already been adapted to some of these problems [8], they have yet to be applied to others. One limitation of the sparse Cholesky approach is that memory constraints render it intractable on 3-D data sets. We have proposed several alternatives to this problem including segmenting the data into smaller components and using incomplete sparse Cholesky factorization [9]; however, investigating these alternatives remains an open problem.

We also explored the effects of edge preserving regularization on magnetic field map estimation near tissue interfaces. Although our results provided valuable insight into the importance of the order of the finite differencing matrices, we have yet to evaluate this modification on real data. Extending this experiment to fat-water imaging is also a compelling research topic as regularized estimation near the air-tissue interface has been identified as a challenge in this imaging technique [10].

In Chapter V, we presented an alternating minimization strategy that accelerated the computation of the CS-WF image reconstruction problem by a factor of 12 over the existing NCG method. Our method used an AL based minimization strategy to update the water-fat images and a CG method with a sparse Cholesky preconditioner to update the field map estimate. We also proposed an initialization strategy based on the regularized field map estimator from Chapter IV. One complication of our minimization method is that it requires updating the AL parameters between outer iterations. Although the current heuristic method provides fast convergence, an automated parameter selection strategy would greatly simplify the algorithm. Furthermore, all of the current minimization strategies (including our proposed method) use a Gauss-Newton like linearization for which an appropriate step size must be selected with each outer iteration. Finding a surrogate function for the CS-WF image reconstruction cost function would mitigate the need for step size selection and simplify all of the algorithms (see [11]). Our alternating minimization strategy could also be applied to other CS based water-fat imaging methods [12–14].

There are several modifications of the original CS-WF cost function [10] that are of interest. As outlined in Chapter V, the fat and water image reconstructions could be improved by using multipeak fat models [10] and by adding additional sparsifying transforms to the cost function [15]. Furthermore, higher quality water and fat images may be possible by exploiting the fact that the object support in the fat image is nearly the compliment to that of the water image [10]. Such an extension would require modifications to both the cost function and our minimization algorithms (see [16,17]). One modification of particular interest is the introduction of mixed–norms, which allow for the incorporation of group sparsity [18]. It would also be interesting to evaluate our algorithms on additional experimental data. First, we could use data sets with numerous scans to explore the trade-off between the undersampling rate and the number of scans in the CS-WF algorithm. Second, it would be clinically relevant to evaluate the algorithms on real undersampled acquisition data, rather than the retrospectively undersampled data used in this work.

Bibliography

[1] S. Ramani, Z. Liu, J. Rosen, J.-F. Nielsen, and J. A. Fessler, "Regularization parameter selection for nonlinear iterative image restoration and MRI reconstruction using GCV and SURE-based methods," *IEEE Trans. Im. Proc.*, vol. 21, no. 8, pp. 3659–72, Aug. 2012.

[2] Y. Chen and A. O. Hero, "Recursive $\ell_{1,\infty}$ group lasso," *IEEE Trans. Sig. Proc.*, vol. 60, no. 8, pp. 3978–87, Aug. 2012.

[3] V. T. Olafsson, D. C. Noll, and J. A. Fessler, "Fast joint reconstruction of dynamic R_2^* and field maps in functional MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 9, pp. 1177–88, Sept. 2008.

[4] W. A. Grissom, V. Rieke, A. B. Holbrook, Y. Medan, M. Lustig, J. Santos, M. V. McConnell, and K. B. Pauly, "Hybrid referenceless and multibaseline subtraction MR thermometry for monitoring thermal therapies in moving organs," *Med. Phys.*, vol. 37, no. 9, pp. 5014–26, 2010.

[5] H. Wang and Y. Cao, "Spatially regularized T1 estimation from variable flip angles MRI," *Med. Phys.*, vol. 39, no. 7, pp. 4139–48, 2012.

[6] S. E. Maier, S. Vajapeyam, H. Mamata, C.-F. Westin, F. A. Jolesz, and R. V. Mulkern, "Biexponential diffusion tensor analysis of human brain diffusion data," *Mag. Res. Med.*, vol. 51, no. 2, pp. 321–30, Feb. 2004.

[7] W. Tang, *Reconstruction of Parametric Image Maps in Single- and Multiple-Coil Functional Magnetic Resonance Imaging*, Ph.D. thesis, Auburn University, 2009.

[8] H. Sun, W. A. Grissom, and J. A. Fessler, "Regularized estimation of Bloch-Siegert B1+ Maps in MRI," in *Proc. IEEE Intl. Conf. on Image Processing*, 2014, Submitted.

[9] D. S. Kershaw, "The incomplete Cholesky-conjugate gradient method for the iterative solution of systems of linear equations," *J. Comp. Phys.*, vol. 26, no. 1, pp. 43–65, Jan. 1978.

[10] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig, "Compressed sensing for chemical shift-based water-fat separation," *Mag. Res. Med.*, vol. 64, no. 6, pp. 1749–59, Dec. 2010.

[11] C. Yip, W. A. Grissom, J. A. Fessler, and D. C. Noll, "Joint design of trajectory and RF pulses for parallel excitation," *Mag. Res. Med.*, vol. 58, no. 3, pp. 598–604, Sept. 2007.

[12] S. D. Sharma, H. H. Hu, and K. S. Nayak, "Accelerated water-fat imaging using restricted subspace field map estimation and compressed sensing," *Mag. Res. Med.*, vol. 67, no. 3, pp. 650–9, Mar. 2012.

[13] S. D. Sharma, H. H. Hu, and K. S. Nayak, "Chemical shift encoded water-fat separation using parallel imaging and compressed sensing," *Mag. Res. Med.*, vol. 69, no. 2, pp. 456–66, Feb. 2013.

[14] C. N. Wiens, C. M. McCurdy, J. D. Willig-Onwuachi, and C. A. McKenzie, "R2*-corrected waterfat imaging using compressed sensing and parallel imaging," *Mag. Res. Med.*, vol. 71, no. 2, pp. 608–616, 2014.

[15] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.

[16] J. Valenzuela and J. A. Fessler, "Joint reconstruction of Stokes images from polarimetric measurements," *J. Opt. Soc. Am. A*, vol. 26, no. 4, pp. 962–8, Apr. 2009.

[17] X. He, J. A. Fessler, L. Cheng, and E. C. Frey, "Regularized image reconstruction algorithms for dual-isotope myocardial perfusion SPECT (MPS) imaging using a cross-tracer edge-preserving prior," *IEEE Trans. Med. Imag.*, vol. 30, no. 6, pp. 1169–83, June 2011.

[18] M. Kowalski, "Sparse regression using mixed norms," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 303 – 324, 2009.