

Efficient Model-Based Reconstruction for Dynamic MRI

by

Claire Yilin Lin

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied and Interdisciplinary Mathematics)
in the University of Michigan
2021

Doctoral Committee:

Professor Jeffrey A. Fessler, Co-Chair
Professor Anna C. Gilbert, Co-Chair
Professor Selim Esedoğlu
Professor Douglas C. Noll

Claire Yilin Lin

yilinlin@umich.edu

ORCID iD: [0000-0002-0088-7065](https://orcid.org/0000-0002-0088-7065)

© Yilin Lin 2021

Acknowledgements

It has been a fulfilling experience to study and apply computational modeling and numerical optimization to the important field of medical imaging. For this I am immensely grateful to my advisor Jeff Fessler, for his abundant technical insight and guidance, his contagious enthusiasm and curiosity in research, and his patient and caring personality.

I thank Doug Noll for his thoughtful suggestions and generous help in my understanding of MRI acquisition and analysis. I thank Anna Gilbert and Selim Esedoğlu for their invaluable feedback and inspiring characters. I also thank Jon-Fredrik Nielsen and Scott Peltier for their gracious assistance in data collection.

My graduate work was supported in part by NIH under Grants R01 EB023618, U01 EB026977, and R21 AG061839, and in part by NSF under Grant IIS 1838179.

I had a fruitful internship at the FDA, thanks to the wonderful mentorship of Rongping Zeng. I have also greatly enjoyed my internship at KLA, owing to the pleasant collaboration with John Luo.

Thanks to the companionship of my lovely labmates, who have created a fun and enjoyable working environment: Amos Cao, Anish Lahiri, Cameron Blocker, Caroline Crockett, Melissa Haskell, Michelle Karker, Naveen Murthy, and Shouchang Guo.

Thanks to my brothers and sisters in Christ, who have shown me the love of Christ: David Hong, Deanna Montgomery, Emily Mueller, Hannah Kim, Jenna Moon, Joe Iafrate, Josh Kammeraad, Josiah Simeth, Matt Schnizlein, Ryan Han, Ryan Hayes, and Xianglong Wang; to David and Dorothy Durling who have welcomed me as part of their family, to Rob and Fietka Narske who have been examples as Christ's faithful servants, and to Rick Mattson who has many times reminded me of His all sufficiency.

I am deeply grateful to my mom and dad, for their selfless love and support in my life adventures, and their vision and courage in the pursuit of my education.

I am eternally grateful to Jesus Christ the Son of God, who was delivered over to death for my sins, and was raised to life for my justification. In Him are my joy, my peace, and my ultimate hope.

Contents

Acknowledgements	ii
List of Figures	vi
List of Tables	viii
Abstract	ix
Chapter	
1 Introduction	1
2 Background	3
2.1 Magnetic Resonance Imaging	3
2.1.1 Signal equation	3
2.1.2 Dynamic MRI	4
2.2 Model-based reconstruction for dynamic MRI	4
2.2.1 System operator	5
2.2.2 Model assumptions	9
2.2.3 Optimization methods	9
2.2.4 Evaluation and analysis	10
3 Efficient Algorithms for the Low-rank Plus Sparse Model	12
3.1 Introduction	12
3.2 Problem and related methods	14
3.2.1 Conventional proximal gradient scheme	15
3.2.2 Conventional variable splitting scheme	17
3.3 Accelerated algorithms	17
3.3.1 Proximal gradient scheme	18
3.3.2 Variable splitting scheme	20

3.4	Results	23
3.4.1	Cardiac perfusion dataset	24
3.4.2	Cardiac cine dataset	24
3.4.3	PINCAT phantom dataset	27
3.5	Discussion	28
3.5.1	Alternative variable splitting	28
3.5.2	ISTA implementation in comparison with [148]	29
3.6	Conclusion	30
3.7	Supplementary material	31
3.7.1	y - t images	31
3.7.2	Long-run behaviors	31
3.7.3	Non-Cartesian experiment	32
4	Efficient Regularized Field Map Estimation	35
4.1	Introduction	35
4.2	Problem formulation	37
4.3	Efficient algorithm	39
4.3.1	Initialization	39
4.3.2	Preconditioning matrices	40
4.3.3	Monotonic step size line search	43
4.4	Results	44
4.4.1	Brain simulation	45
4.4.2	Phantom dataset	47
4.4.3	Cardiac water-fat simulation	48
4.4.4	Ankle water-fat dataset	50
4.5	Conclusion	52
5	A Spatiotemporal Model for Task-based Functional MRI	53
5.1	Introduction	53
5.2	Problem and related models	55
5.2.1	L+S: low-rank plus sparse decomposition	55
5.2.2	L+UV: low-rank plus task-based decomposition	55
5.3	Proposed model	56
5.3.1	B+UV: background plus spatiotemporal decomposition	57
5.3.2	Optimization	58
5.4	Results	63
5.4.1	Reconstruction: simulated tasks	64

5.4.2	Reconstruction: visual and motor tasks	67
5.5	Conclusion	70
6	A Model-based Reconstruction Framework	71
6.1	Intravoxel field inhomogeneity correction	71
6.1.1	Implementation	72
6.1.2	Simulation	73
6.2	Head motion compensation	74
6.2.1	Implementation	75
6.2.2	Simulation	76
6.3	Summary: a general framework	77
7	Future Work	78
7.1	Physiological noise correction	78
7.2	Other models	78
7.3	Model-driven data acquisition	79
	Bibliography	80

List of Figures

3.4.1 Convergence of five algorithms for cardiac perfusion and cine datasets.	25
3.4.2 X_∞ , reconstructed and residual images on cardiac perfusion dataset. .	26
3.4.3 X_∞ , reconstructed and residual images on cardiac cine dataset.	26
3.4.4 Convergence of five algorithms on PINCAT phantom dataset.	27
3.4.5 X_{true} , X_∞ and residual images on PINCAT phantom dataset.	28
3.7.1 y - t images reconstructed from fully and undersampled perfusion dataset.	31
3.7.2 y - t ground truth and image reconstructed from undersampled PINCAT dataset.	31
3.7.3 Long-run cost convergence of five algorithms on three datasets.	32
3.7.4 Long-run NRMSD convergence of five algorithms on three datasets. .	32
3.7.5 NRMSD and cost convergence of three proximal algorithms on abdom- inal DCE dataset.	33
3.7.6 Reconstructed and residual images by ISTA, FISTA, and POGM with similar runtime on abdominal DCE dataset.	34
4.3.1 Factorization and error images on a toy problem.	42
4.3.2 Change of sparsity of \mathbf{L}_t with respect to the scaling factor of H_{\max} in its tolerance.	43
4.4.1 Selected slices of coil-combined simulation image, initial field map (in Hz), regularized field map estimate $\hat{\boldsymbol{\omega}}$, ground truth field map $\boldsymbol{\omega}_{\text{true}}$, and error $ \hat{\boldsymbol{\omega}} - \boldsymbol{\omega}_{\text{true}} $	46
4.4.2 RMSE convergence of seven algorithms on simulation dataset.	47
4.4.3 Selected slices of coil-combined phantom image, initial field map (in Hz), and regularized field map estimate.	48
4.4.4 RMSD convergence of four algorithms on phantom data.	48
4.4.5 RMSE convergence of seven algorithms on water-fat simulation. . . .	49
4.4.6 Simulated image for the 1 st echo, initial field map $\tilde{\boldsymbol{\omega}}^0$ (in Hz) by voxel- wise estimation, and initial fieldmap $\boldsymbol{\omega}^0$ by PWLS.	50

4.4.7 Ground truth and estimated field map, water, and fat images on simulated data.	50
4.4.8 Coil-combined water-fat image for the 1 st echo, initial field map $\tilde{\omega}^0$ (in Hz) by voxel-wise estimation, initial fieldmap ω^0 by PWLS, regularized field map estimate, estimated water image, and estimated fat image. .	51
4.4.9 RMSD convergence of four algorithms on ankle dataset.	52
5.4.1 $3 \times k_z$ - t undersampling.	64
5.4.2 Two simulated tasks as rows of V	64
5.4.3 Timeseries of three activated voxels and one non-activated voxel on simulated data.	66
5.4.4 ROC curves with (AUC values) for simulated task experiment.	67
5.4.5 Activation maps of one slice on simulated dataset.	67
5.4.6 Two tasks s rows of V on visual and motor task data.	68
5.4.7 Timeseries of three activated and one non-activated voxels on visual and motor task data.	69
5.4.8 Activation maps on visual and motor task data.	70
6.1.1 B_0 (in Hz) and gradient maps (in Hz/meter) in simulation.	74
6.1.2 Reconstructed and error images in simulation.	74
6.2.1 SPM motion-corrected and conjugate phase reconstructed images without and with motion correction.	76
6.2.2 Timeseries of a voxel in SPM motion-corrected image, and conjugate phase reconstructed image without and with motion correction.	76
6.3.1 A general model-based reconstruction framework.	77

List of Tables

3.1	Algorithm: proximal gradient for L+S	19
3.2	Algorithm: variable splitting for L+S	22
4.1	Number of nonzeros, memory usage, and NRMSE of factorizations on a toy problem.	42
4.2	Algorithm: preconditioned NCG-MLS for field map estimation	45
4.3	Time for each algorithm to reach an RMSD below 0.5 Hz.	47
5.1	Algorithm: NCG-LS for B+UV	63

Abstract

Dynamic magnetic resonance imaging (MRI) has important clinical and neuroscience applications (e.g., cardiac disease diagnosis, neurological behavior studies). It captures an object in motion by acquiring data across time, then reconstructing a sequence of images from them. This dissertation considers efficient dynamic MRI reconstruction using handcrafted models, to achieve fast imaging with high spatial and temporal resolution. Our *modeling framework* considers data acquisition process, image properties, and artifact correction. The reconstruction model expressed as a large-scale inverse problem requires *optimization algorithms* to solve, and we consider efficient implementations that make use of underlying problem structures.

In the context of dynamic MRI reconstruction, we investigate efficient updates in two frameworks of algorithms for solving a nonsmooth composite convex optimization problem for the low-rank plus sparse (L+S) model. In the proximal gradient framework, current algorithms for the L+S model involve the classical iterative soft thresholding algorithm (ISTA); we consider two accelerated alternatives, one based on the fast iterative shrinkage-thresholding algorithm (FISTA), and the other with the recent proximal optimized gradient method (POGM). In the augmented Lagrangian (AL) framework, we propose an efficient variable splitting scheme based on the form of the data acquisition operator, leading to simpler computation than the conjugate gradient (CG) approach required by existing AL methods. Numerical results suggest faster convergence of our efficient implementations in both frameworks, with POGM providing the fastest convergence overall and the practical benefit of being free of algorithm tuning parameters.

In the context of magnetic field inhomogeneity correction, we present an efficient algorithm for a regularized field inhomogeneity estimation problem. Most existing minimization techniques are computationally or memory intensive for 3D datasets, and are designed for single-coil MRI. We consider 3D MRI with optional consideration of coil sensitivity and a generalized expression that addresses both multi-echo field map estimation and water-fat imaging. Our efficient algorithm uses a preconditioned

nonlinear conjugate gradient method based on an incomplete Cholesky factorization of the Hessian of the cost function, along with a monotonic line search. Numerical experiments show the computational advantage of the proposed algorithm over state-of-the-art methods with similar memory requirements.

In the context of task-based functional MRI (fMRI) reconstruction, we introduce a space-time model that represents an fMRI timeseries as a sum of task-correlated signal and non-task background. Our model consists of a spatiotemporal decomposition based on assumptions of the activation waveform shape, with spatial and temporal smoothness regularization on the magnitude and phase of the timeseries. Compared with two contemporary task fMRI decomposition models, our proposed model yields better timeseries and activation maps on simulated and human subject fMRI datasets with multiple tasks.

The above examples are part of a larger framework for model-based dynamic MRI reconstruction. This dissertation concludes by presenting a general framework with flexibility on model assumptions and artifact compensation options (e.g., field inhomogeneity, head motion), and proposing future work ideas on both the framework and its connection to data acquisition.

Chapter 1

Introduction

Dynamic magnetic resonance imaging (MRI) is a non-invasive medical imaging technique that has important clinical applications. For example, cardiac MRI is used in the diagnosis and treatment of cardiovascular disease, and brain functional MRI (fMRI) is used in neurological and psychological studies. Dynamic MRI captures an object in motion by acquiring a sequence of data across time, then reconstructing a sequence of images from them. Goals in advancing dynamic MRI include fast and high quality imaging. In data acquisition, parallel imaging with multiple receiver coils [160], as well as different undersampling schemes [162,182], have been proposed to accelerate sampling while preserving image quality. Given undersampled multi-coil data, this dissertation focuses on the efficient reconstruction of a MR image sequence with high spatial and temporal resolution. Chapter 2 provides further background for dynamic MRI reconstruction.

Numerous reconstruction models have been proposed for MRI and extended to dynamic MRI [56,164], introducing inverse problems that require different optimization methods [57]. Since the additional time dimension in dynamic imaging scales up the problem size, it is essential to design fast algorithms and take advantage of the structure of our problem when possible. In particular, the low-rank plus sparse (L+S) model decomposes an image sequence into a temporally correlated background and a dynamic foreground, which can be formulated as a nonsmooth composite convex optimization problem. Chapter 3 investigates new efficient algorithms in the proximal gradient and variable splitting frameworks, with faster convergence than previous L+S reconstruction algorithms in each framework.

One subset of dynamic MRI is brain fMRI, which images neural activity in the brain based on blood-oxygenation-level dependent (BOLD) contrast [144]. fMRI uses relatively late echo times to ensure good BOLD contrast, and relatively long readout times to speed imaging to track functional changes [61]. These characteristics lead

to sensitivity of fMRI to magnetic field inhomogeneity, causing image distortion and artifacts if uncorrected, especially near air/tissue interfaces. One way to compensate for field inhomogeneity effect is to estimate a field map and account for it during reconstruction. Chapter 4 presents a regularized field map estimation problem, with a generalized formulation that considers multiple coils with coil sensitivity and addresses both the multi-echo field map estimation and the water-fat imaging settings. Our novel algorithm uses a preconditioner based on an incomplete Cholesky factorization, achieving faster convergence than state-of-the-art methods with similar memory requirements.

fMRI can be broadly categorized into resting state fMRI and task-based fMRI. Resting state fMRI explores functional connectivity across the brain, while task-based fMRI identifies brain regions that are functionally involved in a specific task performance. Conventional fMRI studies postprocess the reconstructed image sequences to correct for noise and artifacts, then analyze functional connectivity or task activation of the processed timeseries. Model-based reconstruction with undersampled data, on the other hand, attempts to use model assumptions to recover the original image sequence and extract neural signal of interest during reconstruction. Chapter 5 introduces a new model for task-based fMRI using *a priori* knowledge of the task waveform, with spatial and temporal smoothness regularization on the magnitude and phase of the timeseries. Compared with contemporary task fMRI decomposition models, our method shows better timeseries and activation maps on fMRI reconstruction with multiple tasks.

fMRI analysis favors high spatiotemporal resolution images, yet a challenge in fMRI reconstruction is its weak signal of interest. Sources of confounding factors include scanner drift, physiological noise (e.g., cardiac and respiratory rhythms), subject head motion during scan, and magnetic field inhomogeneity. Failure to separate noise and artifacts from the neural signal of interest can cause false correlation or anti-correlation [132,158]. Chapter 6 formulates a model-based reconstruction framework for 3D multi-coil dynamic MRI with additional compensation for intravoxel field inhomogeneity and undesired head motion, leading to a discussion of future work in Chapter 7.

Chapter 2

Background

2.1 Magnetic Resonance Imaging

Magnetic Resonance Imaging uses electromagnetic fields to excite magnetic spins in the body, causing hydrogen protons (spins) to precess around the magnetic field direction. After a radiofrequency (RF) pulse tips the spins into the transverse (xy) plane, relaxation occurs as the spins precess back to equilibrium. For spatial localization, locally varying gradients are applied, and the resulting location-dependent frequency of precession provides information of magnetization at each position $\vec{r} = (x, y, z)$. Our received time-dependent signal data $d(t)$ is a summation of the magnetization over the locations.

2.1.1 Signal equation

We can represent the magnetization at location \vec{r} by its image magnitude $x(\vec{r})$ and phase $\phi(\vec{r}, t) = 2\pi\mathbf{k}(t) \cdot \vec{r}$, where the k -space trajectory $\mathbf{k}(t)$ is the spatial frequency sample locations. Ideally, the demodulated MRI signal data seen by a single ideal receiver coil is

$$d(t) = \int x(\vec{r})e^{-i2\pi\mathbf{k}(t)\cdot\vec{r}}d\vec{r}. \quad (2.1)$$

This is a Fourier transform of the image value $x(\vec{r})$, and the k -space is the Fourier domain of image space.

In reality, samples are collected by a receiver coil at discrete time points t_i . For (static) MRI we can approximate the signal by $d_i = d(t_i)$ at k -space location $k_i = \mathbf{k}(t_i)$ and discretize the continuous $x(\vec{r})$ into x_j :

$$d_i \approx \sum_{j=1}^{N_v} x_j e^{-i2\pi k_i \cdot \vec{r}_j},$$

where N_v is the number of voxels. This can be compactly represented as

$$\mathbf{d} = \mathbf{E}\mathbf{x} + \boldsymbol{\epsilon}, \quad (2.2)$$

where $\mathbf{d} \in \mathbb{C}^{N_k}$ is a collection of N_k recorded signal samples, $\mathbf{E} : \mathbb{C}^{N_v} \rightarrow \mathbb{C}^{N_k}$ is an encoding operator that can be represented by an $N_k \times N_v$ matrix with elements $E_{ij} = e^{-i2\pi k_i \cdot \vec{r}_j}$, $\boldsymbol{\epsilon}$ is complex Gaussian noise, and $\mathbf{x} \in \mathbb{C}^{N_v}$ is an unknown image that we aim to reconstruct.

2.1.2 Dynamic MRI

In dynamic MRI with N_t time frames, the unknown image in (2.2) becomes a space-time image sequence matrix $X \in \mathbb{C}^{N_v \times N_t}$, and with multiple coils, the encoding operator generalizes to $\mathbf{E} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k N_c}$, where N_k is now the total number of k -space samples over the entire timeseries, and N_c is the number of signal receiver coils. Since the object is changing while data is collected, dynamic MRI is inherently undersampled. Our goal in reconstruction is to obtain high spatiotemporal resolution images from undersampled k -space data.

As a subset of dynamic MRI, brain function MRI has been developed and explored since the 1990's to study functional activity and cognitive behaviors of the brain. In particular, task-based fMRI localizes brain activations in response to short stimuli or tasks, and its analysis typically includes statistical tools such as the general linear model (GLM) [64,118]. On the other hand, resting state fMRI investigates the functional connectivity, where one uses seed-based analysis [22] or independent component analysis (ICA) [17] to find correlation across brain regions. These image analyses favor high spatiotemporal resolution images, and typically require image processing steps that correct for physiological noise (e.g., cardiac and respiratory rhythms), scanner drift, head motion, and image distortion [112]. Failure to separate noise and artifacts from the neural signal of interest can cause false correlation or anti-correlation [132,158].

2.2 Model-based reconstruction for dynamic MRI

For dynamic MRI, we aim to reconstruct a MR image sequence X from undersampled k -space data \mathbf{d} in (2.2) by solving an optimization problem

$$\operatorname{argmin}_X \frac{1}{2} \|\mathbf{E}X - \mathbf{d}\|_2^2 + \lambda R(X), \quad (2.3)$$

where $R(\cdot)$ is a regularizer with parameter λ that depends on model assumptions.

This dissertation focuses on parallel imaging with undersampling, hence \mathbf{E} is no longer a single Fourier transform. \mathbf{E} can also address field inhomogeneity and motion compensation, and Section 2.2.1 discusses its general formulation. The regularizer $R(\cdot)$ contains model assumptions and is critical for the reconstructed image quality, and Section 2.2.2 reviews some state-of-the-art models for dynamic MRI reconstruction. Section 2.2.3 presents some optimization methods to solve (2.3) once \mathbf{E} and $R(\cdot)$ are formulated, and Section 2.2.4 discusses common evaluation metrics and analysis techniques on reconstructed images.

2.2.1 System operator

The system operator \mathbf{E} in (2.3) can account for *a priori* knowledge of the imaging system and artifact correction. This section covers its formulation with k -space sampling, coil sensitivity, head motion, and field inhomogeneity compensation.

k-space sampling

In (single-coil, single-frame) MRI, the system $\mathbf{E} : \mathbb{C}^{N_v} \rightarrow \mathbb{C}^{N_k}$ has elements

$$E_{ij} = e^{-i2\pi k_i \cdot \vec{r}_j}.$$

With uniform sampling in k -space, it can be efficiently computed by a fast Fourier transform (FFT) operator $\mathbf{F} : \mathbb{C}^{N_v} \rightarrow \mathbb{C}^{N_v}$. In practice, however, undersampling schemes are often used to speed up acquisition. For example, undersampled Echo-Planar Imaging (EPI) combined with parallel imaging [46] accelerates fMRI acquisition while preserving image quality. In this case, the spatial Fourier operator is followed by an undersampling mask $\mathbf{\Omega} : \mathbb{C}^{N_v} \rightarrow \mathbb{C}^{N_k}$

$$\mathbf{E} = \mathbf{\Omega F}.$$

In addition to uniform undersampling, nonuniform trajectories are extensively used in fMRI sampling [36,69,152]. The nonuniform FFT (NUFFT) [55] efficiently approximates the nonuniform Fourier transform using an oversampled discrete Fourier transform (DFT) of scaled image values, then interpolates the DFT values onto desired frequency locations. The (single-coil, single-frame) MRI system matrix with NUFFT $Q \in \mathbb{C}^{N_k \times N_v}$ is

$$Q = \Gamma F \Lambda,$$

where $\Lambda \in \mathbb{C}^{N_v \times N_v}$ is a diagonal scaling matrix that precompensates for imperfections in the interpolation by Γ , $F \in \mathbb{C}^{K \times N_v}$ is the oversampled (K point) DFT that can be computed by FFT, and $\Gamma \in \mathbb{C}^{N_k \times K}$ is a frequency-domain interpolator that can be tabulated to reduce operations and memory storage.

In dynamic MRI, Γ and Λ are defined based on the sampling trajectory of each time frame, and can be reused if the same trajectory is repeated across frames. For nonuniform sampling, the (single-coil) system operator $\mathbf{E} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k}$ can be written as

$$\mathbf{E} = \mathbf{\Omega}\mathbf{Q},$$

where the dimensions of $\mathbf{\Omega}$ and \mathbf{Q} depend on the sampling trajectory. One practical trajectory is a stack of spirals with spiral NUFFT $Q_{xy} \in \mathbb{C}^{N_{kxy} \times N_x N_y}$ and k_z - t undersampling mask $\mathbf{\Omega} : \mathbb{C}^{N_{kxy} N_z N_t} \rightarrow \mathbb{C}^{N_k}$. In this case, the sampling operator $\mathbf{Q} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_{kxy} N_z N_t}$ is

$$\mathbf{Q} = I_{N_t} \otimes (I_{N_z} \otimes Q_{xy}),$$

where \otimes is the Kronecker product operation, and I_N is the identity matrix with dimension $N \times N$.

Coil sensitivity

Parallel MRI uses multiple signal receiver coils to accelerate data acquisition while maintaining image quality. Since its development, multiple strategies have been proposed that consider coil sensitivity in image reconstruction [72,160,181]. In particular, Sensitivity Encoding (SENSE) explicitly uses coil sensitivity maps in reconstruction. The coil sensitivity matrix $S \in \mathbb{C}^{N_v N_c \times N_v}$ can be represented by a stack of N_c coil maps from S_1 to S_{N_c} , where the c th coil map $S_c \in \mathbb{C}^{N_v \times N_v}$ is diagonal and can be estimated from individual coil images $\mathbf{x}_c = S_c \mathbf{x}$ [190]. In multi-coil dynamic MRI, the elements in $\mathbf{E} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k N_c}$ now become

$$E_{ctij} = e^{-i2\pi k_{ti} \cdot \vec{r}_{tj}} S_{cj},$$

and together with sampling $\mathbf{\Omega}\mathbf{Q}$, the system operator can be written as

$$\mathbf{E} = (I_{N_c} \otimes \mathbf{\Omega}\mathbf{Q})(I_{N_t} \otimes S),$$

where the coil sensitivity maps S are static across N_t time frames.

Head motion compensation

Subject motion during MRI scans causes image artifacts and can degrade MRI analysis results. One example is the 3D head movement during a brain fMRI scan, captured by a rigid motion with three rotation and three translation parameters. Numerous methods have been proposed for head motion correction, including k -space compensation [13,26,155,193], image space registration [91], and joint image reconstruction with motion estimation [42,78].

Given image space rigid transformation $\mathbf{T} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_v \times N_t}$ for all N_t time frames, the multi-coil dynamic system $\mathbf{E} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k N_c}$ now has elements

$$E_{ctij} = e^{-i2\pi k_{ti} \cdot \vec{r}_{tj}} S_{cj} \mathbf{T}_t,$$

and the operator becomes

$$\mathbf{E} = (I_{N_c} \otimes \Omega \mathbf{Q})(I_{N_t} \otimes S) \mathbf{T}.$$

In practice, we can implement \mathbf{T} by interpolation using a rotation with three shears and a translation [42,52]. The rotation and translation parameters can be estimated using an initial reconstructed image sequence. Section 6.2 provides implementation detail and a simulation example for reconstruction with motion correction.

Field inhomogeneity

For MR scans with long readout times, there are off-resonance effects caused by magnetic field (B_0) inhomogeneity. This is due to main field imperfections and magnetic susceptibility variations in tissues, which can cause signal loss and image distortions if uncorrected. For fMRI in brain regions near air/tissue interface, signal degradation can be severe [143]. Numerous compensation strategies have been proposed based on both data acquisition [47,141,208] and image reconstruction [76,96,139,184]. The field map can be estimated using MR scans at multiple echo times (usually 2 or 3) [10,65,147,185,201], and built into the system operator using a low-rank approximation [59,61,108].

Given a field map $\omega(\vec{r})$ at location \vec{r} , the signal equation (2.1) becomes

$$d(t) = \int x(\vec{r}) e^{-i\omega(\vec{r})t} e^{-i2\pi \mathbf{k}(t) \cdot \vec{r}} d\vec{r}. \quad (2.4)$$

With piecewise constant discretization of $\omega(\vec{r})$, the multi-coil dynamic system $\mathbf{E} :$

$\mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k N_c}$ now has elements

$$E_{ctij} = e^{-i\omega_j t_i} e^{-i2\pi k_{ti} \cdot \vec{r}_{tj}} s_{cj}.$$

Since it is impractical to carry out this large-scale matrix multiplication, a basis expansion [59] can approximate the exponential term

$$\sum_{l=1}^L B_{il} C_{lj} \approx e^{-i\omega_j t_i}$$

with estimated terms B_{il}, C_{lj} . While this piecewise constant discretization of $\omega(\vec{r})$ assumes a constant field map across a voxel, it can be extended to a piecewise linear discretization that accounts for intravoxel B_0 inhomogeneity [61,108]. With bases $B_l = \text{diag}(B_{il}) \in \mathbb{C}^{N_k \times N_k}$ and $C_l = \text{diag}(C_{lj}) \in \mathbb{C}^{N_v N_t \times N_v N_t}$, the system operator now becomes

$$\mathbf{E} = \left(I_{N_c} \otimes \left(\sum_{l=1}^L B_l(\Omega \mathbf{Q}) C_l \right) \right) (I_{N_t} \otimes S).$$

In functional MRI when the field map is assumed static after motion correction, one can collect multi-echo images before the dynamic scan, and reuse its low-rank approximation across frames. Dynamic field map estimation has also been explored in the joint reconstruction setting [84,128,145,185,192,205].

Chapter 4 discusses an efficient algorithm for regularized field map estimation, and Section 6.1 provides implementation details and a simulation example for reconstruction with intravoxel field inhomogeneity compensation.

A generalized expression for \mathbf{E}

We summarize this section 2.2.1 on the dynamic MRI system operator $\mathbf{E} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k N_c}$ by its generalized form, considering head motion $\mathbf{T} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_v \times N_t}$, coil sensitivity maps $S \in \mathbb{C}^{N_v N_c \times N_v}$, low-rank approximation factors of field inhomogeneity $B_l \in \mathbb{C}^{N_k \times N_k}$, $C_l \in \mathbb{C}^{N_v N_t \times N_v N_t}$, and undersampled NUFFT operator $\Omega \mathbf{Q} : \mathbb{C}^{N_v N_t} \rightarrow \mathbb{C}^{N_k}$:

$$\mathbf{E} = \left(I_{N_c} \otimes \left(\sum_{l=1}^L B_l(\Omega \mathbf{Q}) C_l \right) \right) (I_{N_t} \otimes S) \mathbf{T}.$$

In analytical reconstruction methods such as conjugate phase (CP) [139], one can apply an adjoint operator \mathbf{E}^* to data \mathbf{d} to get the images X . One can also use such

analytical reconstruction results as initial estimates for model-based reconstruction.

2.2.2 Model assumptions

In the image reconstruction problem (2.3), formulations of X and $R(X)$ capture the underlying model assumptions. Here we briefly review some state-of-the-art models for dynamic MRI reconstruction. For example, a low-rank constraint or a nuclear norm regularizer has been used to model the approximately low-rank k - t image [36,39]. A spatial-temporal factorization also achieves low-rankness [37,109,134]. An l_1 norm regularizer encourages sparsity in a transformed domain [54,135]. To combine low-rankness and sparsity, the low-rank plus sparse decomposition [149,154,180,199] represents a dynamic image as a summation of low-rank static background and Fourier-sparse dynamic foreground, while the low-rank and sparse model encourages the image sequence to be both low-rank and sparse [119]. Other regularizers such as total variation (TV) and temporal finite difference are also used to encourage spatial and temporal regularity [9,31,34,136].

In addition to predefined operators for regularization, data-driven approaches adaptively learn the operators. The dictionary learning model learns a synthesis dictionary and its sparse codes [12,120,163], while the transform learning model learns an analytical transform and its sparse codes [79,200]. Recently, model-based deep learning is also explored for dynamic MRI, where a convolutional neural network (CNN) denoising prior is included in the model-based iterative scheme [23,172].

This dissertation considers handcrafted models that capture desired properties of the reconstructed image sequence. Chapter 3, for example, focuses on a decomposition model that represents a dynamic image as a sum of a relatively static background and a dynamic foreground. Chapter 5 proposes a task fMRI model that separates task activation from non-task background signal.

2.2.3 Optimization methods

Given formulations of the system operator \mathbf{E} and regularizer $R(X)$, different optimization methods [57] have been proposed to solve the image reconstruction problem (2.3).

One common dynamic MRI reconstruction problem formulation uses a composite convex function consisting of a sum of a smooth and a nonsmooth function. This includes, for example, cost functions with nuclear norm or l_1 norm regularization. Proximal gradient methods (PGM), whose iterates are based on the proximal operator [14,41,151], specialize in solving composite optimization problems. Classical PGM,

also known as iterative soft thresholding algorithm (ISTA) [45], has a slow $O(1/k)$ convergence bound. The fast iterative soft thresholding algorithm (FISTA) [15,16] accelerates it by achieving a $O(1/k^2)$ convergence, and the proximal optimized gradient method (POGM) [104,186] further improves the worst-case convergence bound in FISTA by about two.

While PGM provides convenient updates for composite cost functions in a synthesis model, its update formulation may not be straightforward for cost functions in an analysis model, due to the analysis operator. In this case, variable splitting methods reformulate the cost as an equivalent constrained minimization problem with an auxiliary variable, and solve it with algorithms using augmented Lagrangian (AL) and alternating direction method of multipliers (ADMM) [8,25]. To achieve faster convergence of variable splitting methods, [202,203] explore adaptive selection of AL parameters. In MRI reconstruction, one can further exploit structure of the system operator \mathbf{E} for efficient update [161].

Quadratically regularized cost functions arise in dynamic MRI reconstruction, for example, when one considers finite difference regularizers, or as a step in AL-based methods. The conjugate gradient (CG) algorithm is well-suited for such cost functions, and one can explore faster convergence with CG using preconditioners, such as a circulant preconditioner [32] when the Hessian matrix is approximately Toeplitz.

Nonconvex regularized cost functions arise, for example, when one considers separate magnitude and phase regularization [60,209,211]. In this case, an optimization approach is to minimize a sequence of surrogate functions using optimization transfer [110]. Alternatively, one can use nonlinear CG with line search [43].

This dissertation includes efficient algorithms for different optimization problems. Chapter 3 discusses fast proximal gradient methods and efficient AL algorithms in the context of the low-rank plus sparse model for dynamic parallel MRI reconstruction. Chapter 4 provides a computationally and memory efficient preconditioned nonlinear CG scheme for the regularized field map estimation problem. Chapter 5 efficiently solves a reconstruction problem using CG and alternating minimization.

2.2.4 Evaluation and analysis

After reconstruction, image quality is evaluated using different metrics based on one’s focus and intended use of the reconstructed images. Given a reference “ground truth” image, the mean square error (MSE) and the structural similarity index (SSIM) [197] are useful for assessing difference of structural information and structural distortion from a desired image. The signal-to-noise ratio (SNR) and peak

SNR (PSNR) are also widely used to measure the reconstruction quality from under-sampled k -space data.

In addition to evaluating an image in its spatial dimension, the reconstruction quality of dynamic MRI also depends on its temporal dimension. In particular, an active research area of fMRI analysis aims to find relationships between brain regions using their timeseries. Conventional fMRI analysis processes the reconstructed images by smoothing and artifact correction, then evaluates the timeseries across voxels in the brain based on specific goals in fMRI. In task-based fMRI, a main goal is detecting activated brain regions. Given a predefined task waveform, one can compute the correlation map of the image sequence to the task, and define a task activation map based on correlation values. As activation maps can be sensitive to the choice of correlation threshold and experimental setup, robust methods have been proposed, such as false discovery rate (FDR)-controlling procedure [66], test-retest scheme [67, 140], and permutation tests [7,137]. In resting state fMRI, one goal is to analyze functional connectivity by identifying resting state networks (RSN) [44]. To this end, seed voxel / region-based regression [22] and independent component analysis (ICA)-based methods [17] have been used extensively.

In iterative reconstruction, we also evaluate the performance of an algorithm by monitoring its convergence behavior and speed. For example, we can plot the cost or the normalized root mean square difference (NRMSD) of each iterate to a converged image. In cases with a non-strictly convex or even nonconvex cost, a “converged image” is not unique, and one can take, for example, the average of converged images by different algorithms.

Throughout this dissertation, we measure image quality by the error to the ground truth or the difference to a reference. We also use metrics specific to our goals, such as timeseries and activation maps for fMRI analysis, and convergence plots in algorithm comparison.

Chapter 3

Efficient Algorithms for the Low-rank Plus Sparse Model

Among MRI reconstruction models, the low-rank plus sparse (L+S) model considers the spatiotemporal feature of dynamic images, and decomposes an image sequence into a temporally correlated background and a dynamic foreground. This assumption enables the reconstruction of undersampled dynamic parallel MRI data. Solving for the low-rank and the sparse components involves nonsmooth composite convex optimization, and algorithms for this problem can be categorized into proximal gradient methods and variable splitting methods. This chapter investigates new efficient algorithms for both schemes. While current proximal gradient techniques for the L+S model involve the classical iterative soft thresholding algorithm (ISTA), we consider two accelerated alternatives, one based on the fast iterative shrinkage-thresholding algorithm (FISTA), and the other with the recent proximal optimized gradient method (POGM). In the augmented Lagrangian (AL) framework, we propose an efficient variable splitting scheme based on the form of the data acquisition operator, leading to simpler computation than the conjugate gradient (CG) approach required by existing AL methods. Numerical results suggest faster convergence of the efficient implementations for both frameworks, with POGM providing the fastest convergence overall and the practical benefit of being free of algorithm tuning parameters.

3.1 Introduction

The application of compressed sensing (CS) to Magnetic Resonance Imaging (MRI) has been extensively explored to accelerate the data acquisition process [75,

This chapter presents our published work in the conference and journal papers [114,115].

124]. In particular, since dynamic MRI data is inherently undersampled, it is useful to use a CS-MRI model for image reconstruction. CS has also been combined with parallel MRI techniques such as SENSitivity Encoding (SENSE) [160], aiming to collect more data with multiple receiver coils, thereby possibly improving the spatiotemporal resolution trade-off of the reconstructed images. This combination is especially useful in dynamic MRI, where reconstruction of high spatial and temporal resolution is desired [94]. Compared with models that use coil-by-coil auto-calibration [72,125], the SENSE framework uses explicit knowledge of the sensitivity of the receiver coils.

In addition to image sparsity based on CS, the low-rank models of the space-time matrix have also been explored for dynamic MRI, based on assumptions of the similarities between temporal profiles [119,148,169,188]. In particular, a low-rank plus sparse (L+S) matrix decomposition assumes incoherence between a low-rank component L and a sparse component S , with L modeling the temporally correlated background, and S the dynamic information that lies on top of the background. The corresponding reconstruction problem can be formulated as a convex optimization problem, where the nuclear norm and l_1 norm are used to respectively promote low-rankness and sparsity regularization on L and S . The L+S formulation has various applications, such as motion estimation in dynamic contrast-enhanced MRI, and automated background suppression for angiography [148,188].

One technique for solving such optimization problems involves the class of proximal gradient methods (PGM), whose iterates are based on the proximal operator [14,41,151]. In particular, [148] solves the L+S decomposition with the iterative soft thresholding algorithm (ISTA). Although accelerated variants of ISTA have been applied to various non-parallel and parallel MRI models with sparsity regularization [19,28,85,129], to our knowledge, fast PGM has yet to be explored for the L+S reconstruction problem.

Variable splitting is another category of optimization schemes that has been used extensively for various MRI reconstruction models, with formulation in the augmented Lagrangian (AL) framework. In single-coil dynamic MRI, variable splitting has provided efficient alternating update schemes for L+S models [119,188]. Using the splitting of variables to decouple a cost function into simpler sub-problems, one can also apply accelerated schemes, such as the fast iterative shrinkage-thresholding algorithm (FISTA) [16], to the sub-problems, for more efficient computation [97,204]. For the L+S model, [188] proposed a splitting scheme for which the AL function leads to sub-problems with quadratic updates. That approach requires inverting a matrix of the form $(\mathbf{E}^*\mathbf{E} + \delta I)^{-1}$, where \mathbf{E} is an encoding operator (see Section 2.2.1), \mathbf{E}^*

is its adjoint, and δ denotes a penalty parameter. With non-Cartesian sampling, or with multiple coils in the case of parallel MRI, a computationally demanding iterative approach like the conjugate gradient (CG) method is required for the updates. Efficient formulations of this update have been investigated, including a singular value decomposition (SVD) of $\mathbf{E}^*\mathbf{E}$ [20], and pre-multiplication of a Fourier operator [169]. However, these implementations are still computationally intensive with multi-coil data.

For undersampled dynamic parallel MRI, this chapter presents efficient algorithms for the L+S image reconstruction problem. In particular, in the PGM category, we investigate two accelerated alternatives to ISTA, one based on FISTA, and the other the recent proximal optimized gradient method (POGM) [186]. For methods involving variable splitting, we adapt a splitting scheme that uses the matrix structures associated with the undersampling pattern, the Fourier encoding and the sensitivity maps, leading to faster MRI reconstruction [35,111,161]. In this case, we take advantage of the L+S model structure, leading to efficient updates with only two AL variables. This chapter is an extension of our conference work [114] that briefly investigated the two accelerated algorithms in the PGM category. Compared with this earlier work, here we discuss the algorithms in more detail, investigate another accelerated algorithm in the variable splitting scheme, and include an extension to non-Cartesian MRI in the Supplement (Section 3.7).

The rest of this chapter is organized as follows. Section 3.2 formulates the corresponding convex optimization problem, and reviews some related methods for solving it. Section 3.3 presents the efficient implementations for the two classes of algorithms. Section 3.4 reports experimental results, followed by discussion and conclusion in Sections 3.5 and 3.6 respectively.

3.2 Problem and related methods

In the L+S framework for dynamic MRI, the goal is to estimate an unknown image, modeled as a superposition of a low-rank component L and a sparse component S . In parallel MRI, we are provided with undersampled k -space data $\mathbf{d} \in \mathbb{C}^{N_s N_c}$, where N_s is the total number of samples received from each receiver coil (across all frames), and N_c is the number of coils. N_x and N_y denote the image dimensions of each image frame, and N_t is the number of time frames. The L+S formulation [148] uses the

following regularized convex optimization scheme:

$$\operatorname{argmin}_{L,S} \frac{1}{2} \|\mathbf{E}(L + S) - \mathbf{d}\|_2^2 + \lambda_L \|L\|_* + \lambda_S \|\mathbf{T}S\|_1, \quad (3.1)$$

where $L, S \in \mathbb{C}^{N_x N_y \times N_t}$ are the desired dynamic image components, $\mathbf{E} : \mathbb{C}^{N_x N_y \times N_t} \rightarrow \mathbb{C}^{N_s N_c}$ is the data acquisition operator that considers the coil sensitivities and the Fourier transform with undersampling, and $\mathbf{T} : \mathbb{C}^{N_x N_y \times N_t} \rightarrow \mathbb{C}^{N_x N_y N_t}$ a known sparsifying transform operator based on *a priori* assumptions of the domain of image sparsity. This work considers the (unitary) temporal Fourier transform operator, with $\mathbf{T}S = (T \otimes I_{N_x N_y}) \mathbf{vec}(S)$, where T is the $N_t \times N_t$ unitary temporal discrete Fourier transform matrix. This sparsifying transform has been extensively used to promote sparsity in dynamic MRI reconstruction [94,119,148,188]. Our accelerated methods adapt readily to other unitary operators and to 3D dynamic MRI problems. Here the data consistency is captured by the vector l_2 -norm term, the low-rankness of L by the matrix nuclear norm, and the sparsity of the transformed S by the vector l_1 norm. The contributions between these three terms are balanced by the regularization parameters λ_L and λ_S .

Methods for solving the optimization problem (3.1) fall into two classes: those based on the proximal gradient methods, and those using AL with variable splitting. Below, we review these two methods and existing implementations for the L+S reconstruction model.

3.2.1 Conventional proximal gradient scheme

To implement the classical PGM on the L+S optimization problem, we combine the two unknowns by forming a single “stacked” variable $\mathcal{X} = \begin{bmatrix} L \\ S \end{bmatrix}$. With this change, (3.1) can be equivalently expressed as

$$\min_{\mathcal{X}} g(\mathcal{X}) + h_1(\mathcal{X}) + h_2(\mathcal{X}), \quad \text{where}$$

$$g(\mathcal{X}) = \frac{1}{2} \|\mathbf{E} \mathbf{E} \mathcal{X} - \mathbf{d}\|_2^2, \quad (3.2)$$

$$h_1(\mathcal{X}) = \lambda_L \|[I \ 0]\mathcal{X}\|_*, \quad \text{and} \quad h_2(\mathcal{X}) = \lambda_S \|[0 \ \mathbf{T}]\mathcal{X}\|_1.$$

Here I and 0 denote respectively the identity and the zero matrices, of size $N_x N_y \times N_x N_y$. To verify the convergence assumptions of PGM, we note that $g(\mathcal{X})$ is a smooth,

convex, and continuously differentiable function, whose gradient is Lipschitz continuous with constant $l(\nabla g)$; $h_1(\mathcal{X})$, $h_2(\mathcal{X})$ are continuous, convex and nonsmooth functions. For arbitrary variables \mathcal{Y} , \mathcal{Z} , the k th iterate of the PGM is then given by the proximal operator:

$$\mathcal{Y}_k = \mathbf{prox}_h(\mathcal{Y}_{k-1} - t\nabla g(\mathcal{Y}_{k-1})) , \text{ where}$$

$$\mathbf{prox}_h(\mathcal{Z}) = \underset{\mathcal{Y}}{\operatorname{argmin}} h(\mathcal{Y}) + \frac{1}{2}\|\mathcal{Y} - \mathcal{Z}\|_2^2 .$$

Here t is a chosen step size, whose dependence on the Lipschitz constant $l(\nabla g)$ guarantees convergence of the algorithm.

The proximal maps for the nuclear norm in h_1 and the vector l_1 norm in h_2 have closed-form expressions. In particular, \mathbf{prox}_{h_2} is given by the soft thresholding operator

$$\mathbf{\Lambda}_\lambda(\mathcal{Y}) = \operatorname{sign}(\mathcal{Y}) \odot (|\mathcal{Y}| - \lambda)_+ ,$$

where \odot denotes element-wise multiplication. \mathbf{prox}_{h_1} is the singular value thresholding operator

$$\mathbf{SVT}_\lambda(\mathcal{Y}) = U\mathbf{\Lambda}_\lambda(\Sigma)V^* ,$$

where $U\Sigma V^*$ is a singular value decomposition of \mathcal{Y} . Since h_1 and h_2 are functions of L and S respectively, the k -th iterate can be written separately for L_k and S_k :

$$L_k = \mathbf{SVT}_{\lambda_L}(L_{k-1} - td(\mathcal{X}_{k-1})) , \text{ and}$$

$$S_k = \mathbf{T}^* \left(\mathbf{\Lambda}_{\lambda_S} \left[\mathbf{T}(S_{k-1} - td(\mathcal{X}_{k-1})) \right] \right) , \text{ where}$$

$$d(\mathcal{X}) = [I \ 0]\nabla g(\mathcal{X}) = \mathbf{E}^*([\mathbf{E} \ \mathbf{E}]\mathcal{X} - \mathbf{d}) = [0 \ I]\nabla g(\mathcal{X}) .$$

Here \mathbf{T}^* denotes the adjoint operator of \mathbf{T} , defined by the corresponding inverse Fourier transform. The expression of S_k uses the fact that \mathbf{T} represents the unitary temporal Fourier transform, and that the l_2 norm is unitary invariant.

Since the gradient $d(\mathcal{X}_{k-1})$ is the same in both expressions of L_k and S_k , only one gradient evaluation is necessary in each iteration. The methods in [148] use this fact, exploiting computational efficiency by jointly estimating L and S . However, Section 3.3 shows that fast PGM provides much further acceleration.

3.2.2 Conventional variable splitting scheme

It has been empirically observed that in some problem settings, AL-based methods can achieve higher accuracy than PGM in fewer iterations [30]. This has motivated another technique of solving (3.1), using variable splitting.

A splitting scheme is introduced in [188] to solve the L+S decomposition problem. In particular, (3.1) is reformulated with two constraints:

$$\begin{aligned} \operatorname{argmin}_{L,S} \min_{U,W} \frac{1}{2} \|\mathbf{E}(L+S) - \mathbf{d}\|_2^2 + \lambda_L \|U\|_* + \lambda_S \|W\|_1 \\ \text{subject to } \begin{cases} U = L \\ W = \mathbf{T}S. \end{cases} \end{aligned} \quad (3.3)$$

With this formulation, the associated modified AL function is

$$\begin{aligned} \frac{1}{2} \|\mathbf{E}(L+S) - \mathbf{d}\|_2^2 + \lambda_L \|U\|_* + \lambda_S \|W\|_1 \\ + \frac{\delta_1}{2} \|L - U + V_1\|_2^2 + \frac{\delta_2}{2} \|\mathbf{T}S - W + V_2\|_2^2, \end{aligned}$$

where V_1, V_2 are Lagrange multiplier arrays, and δ_1, δ_2 are two corresponding AL penalty parameters that affect the convergence rate, but not the final estimates.

This problem can be solved by iterative updates of the four unknowns, followed by updates of the Lagrange multipliers. In particular, each update of L and S is quadratic, requiring computation of $(\mathbf{E}^*\mathbf{E} + \delta_i I)^{-1}$ for $i = 1, 2$. With single-coil Cartesian data, as considered in [188], $\mathbf{E}^*\mathbf{E}$ is circulant and one can use FFT operations for efficient computation. However, in parallel MRI, the operator \mathbf{E} contains additional information of coil sensitivities, so $\mathbf{E}^*\mathbf{E}$ is not circulant, and the updates of the quadratic terms would require an iterative method like the CG approach.

Based on this observation, Section 3.3 presents a new AL algorithm that simplifies the computation by considering a more efficient variable splitting scheme for the L+S model.

3.3 Accelerated algorithms

This section presents three efficient algorithms for the minimization problem (3.1). Two of them are in the class of PGM, with additional momentum terms in the updates that help achieve faster convergence rates. The third is an AL method that uses a

different variable splitting scheme than (3.3), exploiting the structure of the data acquisition operator \mathbf{E} , improving computation efficiency for parallel MRI.

3.3.1 Proximal gradient scheme

The ISTA update for L and S , as given in [148], is based on classical PGM, for which the sequence of function values converges to the optimal function value at a rate of $O(1/k)$ [16]. We assume from now on that the operator \mathbf{E} is normalized such that the spectral norm $\|\mathbf{E}\|_2 = 1$ for fully sampled data. Then the Lipschitz constant of $g(\cdot)$ in (3.2) satisfies

$$l(\nabla g) = \|[\mathbf{E} \ \mathbf{E}]\|_2^2 = 2\|\mathbf{E}\|_2^2 \leq 2, \quad (3.4)$$

so ISTA converges for any step size t with $0 < t < \frac{2}{2\|\mathbf{E}\|_2^2} = 1$.

We now introduce two accelerated methods for (3.1) that have $O(1/k^2)$ convergence rates; their convergence analyses build on the work of Nesterov’s fast gradient methods [133]. We use the same algorithm framework for these two accelerated L+S variants of ISTA, formulated as **Algorithm 1**. Computing the gradient is the most expensive step in each iteration; because both the L and S updates involve the same gradient expression, we jointly update them by first computing

$$\mathcal{X}_k = \begin{bmatrix} L_k \\ S_k \end{bmatrix},$$

then evaluating the gradient in a data consistency term, denoted as M_k below.

FISTA update

Built upon the convergence analysis in [133], FISTA achieves the same rate of convergence of $O(1/k^2)$ [16]. In addition to the unknown \mathcal{X}_k in each iteration, the FISTA update involves a secondary sequence $\tilde{\mathcal{X}}_k = \begin{bmatrix} \tilde{L}_k \\ \tilde{S}_k \end{bmatrix}$ computed by adding a “momentum” term to the original sequence. This addition preserves the computational simplicity of ISTA, as the main computational effort of gradient evaluation remains unchanged from ISTA. Given the L+S framework of **Algorithm 1**, the additional FISTA initialization and updates are:

Algorithm 1: Proximal Gradient L+S

Inputs:

- d**: undersampled k-t data
- E**: data acquisition operator
- T**: temporal Fourier transform
- λ_L : singular value threshold
- λ_S : sparsity threshold

Initialization:

- $M_0 = L_0 = \mathbf{E}^* \mathbf{d}$, $S_0 = 0$
- additional initialization (*I*) for FISTA or POGM

for $k = 1, 2, \dots, N$ **do**

- update \mathcal{X}_k by FISTA or POGM scheme (\mathcal{X}_k)
- update M_k by FISTA or POGM scheme (M_k)

end for
output: \mathcal{X}_N

$$\begin{aligned}
 (I) \quad & \tilde{\mathcal{X}}_0 = \mathcal{X}_0, \quad \theta_0 = 1 \\
 (\mathcal{X}_k) \quad & L_k = \mathbf{S} \mathbf{V} \mathbf{T}_{\lambda_L} (M_{k-1} - \tilde{S}_{k-1}) \\
 & S_k = \mathbf{T}^* \left(\mathbf{\Lambda}_{\lambda_S} \left[\mathbf{T} (M_{k-1} - \tilde{L}_{k-1}) \right] \right) \\
 & \theta_k = \frac{1 + \sqrt{1 + 4\theta_{k-1}^2}}{2} \\
 & \tilde{\mathcal{X}}_k = \mathcal{X}_k + \frac{\theta_{k-1} - 1}{\tilde{\theta}_k} (\mathcal{X}_k - \mathcal{X}_{k-1}) \\
 (M_k) \quad & M_k = \tilde{L}_k + \tilde{S}_k - t \mathbf{E}^* (\mathbf{E} (\tilde{L}_k + \tilde{S}_k) - \mathbf{d}) .
 \end{aligned}$$

Based on the Lipschitz constant in (3.4), convergence for FISTA is guaranteed when the step size satisfies $0 < t \leq \frac{1}{2\|\mathbf{E}\|_2^2} = 0.5$.

POGM update

In the smooth unconstrained setting, the recent optimized gradient method (OGM) achieves a worst-case convergence bound twice as small as that of Nesterov’s fast gradient methods (FGM) [102,103] by optimizing the choice of the coefficients that determine the step size in a first-order algorithm by minimizing a relaxed worst-case performance bound of $f(X_N) - f(X_*)$, the cost function discrepancy at the N th iteration. This optimization problem is solved by semi-definite programming (SDP) in [50], and an analytical expression of the optimized step size is derived in [102], confirming the numerical observation that the worst-case performance is two times better than FGM’s bound. OGM was shown to have optimal complexity for large-scale smooth problems in [50].

OGM has been extended to the proximal case for nonsmooth composite problems, where a numerical worst-case performance twice better than FISTA is achieved by POGM [186]. In the L+S model, compared with the FISTA iterate, POGM introduces an additional sequence $\bar{\mathcal{X}}_k = \begin{bmatrix} \bar{L}_k \\ \bar{S}_k \end{bmatrix}$, whose update involves three momentum terms. The POGM formulation is guaranteed to converge when using the same step size as in FISTA, and it again achieves the same computational simplicity as ISTA. The initialization and updates with POGM for the L+S framework are:

$$\begin{aligned}
(I) \quad & \tilde{\mathcal{X}}_0 = \bar{\mathcal{X}}_0 = \mathcal{X}_0, \quad \theta_0 = \zeta_0 = 1 \\
(\mathcal{X}_k) \quad & \tilde{L}_k = M_{k-1} - S_{k-1} \\
& \tilde{S}_k = M_{k-1} - L_{k-1} \\
& \theta_k = \begin{cases} \frac{1 + \sqrt{1 + 4\theta_{k-1}^2}}{2}, & k < N \\ \frac{1 + \sqrt{1 + 8\theta_{k-1}^2}}{2}, & k = N \end{cases} \\
& \bar{\mathcal{X}}_k = \tilde{\mathcal{X}}_k + \frac{\theta_{k-1}-1}{\theta_k}(\tilde{\mathcal{X}}_k - \tilde{\mathcal{X}}_{k-1}) \\
& \quad + \frac{\theta_{k-1}}{\theta_k}(\tilde{\mathcal{X}}_k - \mathcal{X}_{k-1}) + \frac{\theta_{k-1}-1}{\zeta_{k-1}\theta_k}t(\bar{\mathcal{X}}_{k-1} - \mathcal{X}_{k-1}) \\
& \zeta_k = t\left(1 + \frac{\theta_{k-1}-1}{\theta_k} + \frac{\theta_{k-1}}{\theta_k}\right) \\
& L_k = \mathbf{SVT}_{\lambda_L}(\tilde{L}_k) \\
& S_k = \mathbf{T}^*\left(\mathbf{A}_{\lambda_S}[\mathbf{T}(\bar{S}_k)]\right) \\
(M_k) \quad & M_k = L_k + S_k - t\mathbf{E}^*(\mathbf{E}(L_k + S_k) - \mathbf{d}).
\end{aligned}$$

The empirical results in Section 3.4 show that POGM converges faster than ISTA and FISTA, yet requires essentially the same computation time per iteration (dominated by the M_k update needed in all methods).

3.3.2 Variable splitting scheme

We now consider variable splitting methods for the L+S reconstruction problem for parallel MRI. In this setting, the data acquisition operator is $\mathbf{E} = \mathbf{\Omega}Q\mathbf{C}$, where $\mathbf{\Omega} : \mathbb{C}^{N_x N_y N_t N_c} \rightarrow \mathbb{C}^{N_s N_c}$ contains the undersampling patterns for all frames, $Q \in \mathbb{C}^{N_x N_y N_t N_c \times N_x N_y N_t N_c}$ represents a Fourier encoding matrix, and $\mathbf{C} : \mathbb{C}^{N_x N_y \times N_t} \rightarrow \mathbb{C}^{N_x N_y N_t N_c}$ captures the sensitivity maps of the receiver coils [35,111,161]. While direct extension of the splitting scheme in [161] to the L+S model leads to at least four more variables in the AL function, here we make use of the L and S formulations, and introduce only two AL variables to capture the constrained cost function. For simpler formulation of the algorithm, we assume from now on that \mathbf{C} is normalized such that $\mathbf{C}^*\mathbf{C} = \mathbf{I}$ (identity). This normalization is valid since our model considers sparsity with temporal Fourier transform \mathbf{T} , and the spatial scaling does not affect

the rank of the low-rank component. After reconstruction, one can undo the image scaling if needed.

With this expression, we represent the following novel reformulation of (3.1) in the constrained form

$$\begin{aligned} & \operatorname{argmin}_{L,S} \min_{Z,X} \frac{1}{2} \|\boldsymbol{\Omega}Z - \mathbf{d}\|_2^2 + \lambda_L \|L\|_* + \lambda_S \|\mathbf{T}S\|_1 \\ & \text{subject to } \begin{cases} Z = Q\mathbf{C}X \\ X = L + S . \end{cases} \end{aligned} \quad (3.5)$$

Compared with (3.3), this splitting scheme also involves four variable updates, but leads to simpler updates, as shown next.

The modified AL function corresponding to (3.5) is

$$\frac{1}{2} \|\boldsymbol{\Omega}Z - \mathbf{d}\|_2^2 + \lambda_L \|L\|_* + \lambda_S \|\mathbf{T}S\|_1 + \frac{\delta_1}{2} \|Z - Q\mathbf{C}X + V_1\|_2^2 + \frac{\delta_2}{2} \|X - (L + S) + V_2\|_2^2 .$$

The L update involves the nuclear norm, and its proximal map is given by singular value thresholding:

$$\begin{aligned} & \operatorname{argmin}_L \lambda_L \|L\|_* + \frac{\delta_2}{2} \|X - (L + S) + V_2\|_2^2 \\ & = \mathbf{S}\mathbf{V}\mathbf{T}_{\lambda_L/\delta_2}(X - S + V_2) . \end{aligned} \quad (3.6)$$

The S update contains a vector l_1 -norm term, whose proximal operator is soft thresholding, where we use the fact that \mathbf{T} is a unitary operator, with the change of variables $\tilde{S} = \mathbf{T}S$:

$$\begin{aligned} & \operatorname{argmin}_S \lambda_S \|\mathbf{T}S\|_1 + \frac{\delta_2}{2} \|X - (L + S) + V_2\|_2^2 \\ & = \mathbf{T}^* \left(\operatorname{argmin}_{\tilde{S}} \lambda_S \|\tilde{S}\|_1 + \frac{\delta_2}{2} \|\mathbf{T}(X - L + V_2) - \tilde{S}\|_2^2 \right) \\ & = \mathbf{T}^* \boldsymbol{\Lambda}_{\lambda_S/\delta_2}(\mathbf{T}(X - L + V_2)) . \end{aligned} \quad (3.7)$$

The updates for Z and X involve quadratic terms:

$$\begin{aligned} & \operatorname{argmin}_Z \frac{1}{2} \|\boldsymbol{\Omega}Z - \mathbf{d}\|_2^2 + \frac{\delta_1}{2} \|Z - Q\mathbf{C}X + V_1\|_2^2 \\ & = (\boldsymbol{\Omega}^* \boldsymbol{\Omega} + \delta_1 \mathbf{I})^{-1} (\boldsymbol{\Omega}^* \mathbf{d} + \delta_1 (Q\mathbf{C}X - V_1)) , \end{aligned} \quad (3.8)$$

$$\begin{aligned}
& \underset{X}{\operatorname{argmin}} \frac{\delta_1}{2} \|Z - Q\mathbf{C}X + V_1\|_2^2 + \frac{\delta_2}{2} \|X - (L + S) + V_2\|_2^2 \\
&= (\mathbf{C}^*\mathbf{C} + \frac{\delta_2}{\delta_1}\mathbf{I})^{-1} (\mathbf{C}^*Q^*(Z + V_1) + \frac{\delta_2}{\delta_1}(L + S - V_2)) \\
&= \frac{\delta_1}{\delta_1 + \delta_2} (\mathbf{C}^*Q^*(Z + V_1) + \frac{\delta_2}{\delta_1}(L + S - V_2)) , \tag{3.9}
\end{aligned}$$

where we use the fact that Q is the unitary Fourier encoding matrix, and that $\mathbf{C}^*\mathbf{C} = \mathbf{I}$ by assumption.

Compared with the splitting scheme in (3.3), which involves the inverse $(\mathbf{E}^*\mathbf{E} + \delta_i\mathbf{I})^{-1}$, our proposed variable splitting scheme in (3.5) only involves computing $(\mathbf{\Omega}^*\mathbf{\Omega} + \delta_1\mathbf{I})^{-1}$. Representing the undersampling mask matrix as a Kronecker product $\mathbf{\Omega} = I_{N_c} \otimes \tilde{\mathbf{\Omega}}$, we note that $\tilde{\mathbf{\Omega}}^*\tilde{\mathbf{\Omega}}$ is diagonal, hence the inverse $(\mathbf{\Omega}^*\mathbf{\Omega} + \delta_1\mathbf{I})^{-1}$ is easy to compute.

Algorithm 2 summarizes the implementation of these updates, as well as updates for the updates of the Lagrange multipliers.

Algorithm 2: Variable Splitting L+S

Inputs:

- d:** undersampled multi-coil k-t data
- $\mathbf{\Omega}$:** undersampling mask
- Q :** Fourier encoding operator
- \mathbf{C} :** coil sensitivity maps
- \mathbf{T} :** temporal Fourier transform
- λ_L :** singular value threshold
- λ_S :** sparsity threshold
- δ_1, δ_2 :** AL penalty parameters

Initialization: $X_0 = L_0 = \mathbf{C}^*Q^*\mathbf{\Omega}^*\mathbf{d}$, $S_0 = V_{1,0} = V_{2,0} = 0$

for $k = 1, 2, \dots, N$ **do**

 compute Z_k by efficient inverse (3.8)

 compute X_k by efficient inverse (3.9)

 compute L_k by singular value thresholding (3.6)

 compute S_k by soft thresholding (3.7)

$V_{1,k} \leftarrow V_{1,k-1} + (Z_k - Q\mathbf{C}X_k)$

$V_{2,k} \leftarrow V_{2,k-1} + (X_k - (L_k + S_k))$

end for

output: L_N, S_N

3.4 Results

To compare the algorithms, we first performed experiments on two dynamic MRI datasets examined in [148]. Each dataset includes Cartesian undersampled multi-coil data d , the k -space undersampling mask Ω , and coil sensitivity maps \mathbf{C} . We compared the results of the three accelerated algorithms with ISTA [148] and the AL-based method that requires CG for parallel MRI [188]. We then tested our methods on the physiologically improved nonuniform cardiac torso (PINCAT) numerical phantom used in [119]. In this case, the data is undersampled with a pseudo-radial scheme, as in the original implementation [3]. To compare algorithms in the parallel MRI setting, we included simulated coil sensitivity maps based on [73], using the Michigan Image Reconstruction Toolbox (MIRT) [4]. In the Supplement, we also explore a non-Cartesian MRI dataset from [148], where we compare methods in the PGM scheme. All our experiments used MATLAB R2018a, with a 2.7-GHz dual-core Intel Core i5. The MATLAB code that reproduces the experiments with our efficient algorithms is at <https://github.com/JeffFessler/reproduce-l-s-dynamic-mri>.

For each of the three datasets, we kept the regularization parameters λ_L, λ_S consistent for all algorithms. For the *in vivo* data, we set them to align as closely as possible with those in the original code provided by [148]; Section 3.5 discusses further details of this procedure. Similarly, we set a stopping criterion for our ISTA implementation that provides analogous results to the reconstructed images in [148]. To ensure fast convergence, we used a step size t of 0.99 for ISTA, and 0.5 for FISTA and POGM, as provided by the convergence theory of those methods. In addition, FISTA and POGM used an adaptive restart scheme [150]; we explored both the function and the gradient restart schemes, and report the results with the function scheme due to its slightly faster convergence with both datasets. We tuned the penalty parameters δ_1, δ_2 for the AL-based methods by sweeping across a range of values and choosing the ones that achieve the fastest convergence among them. For both datasets, we applied 3 inner CG iterations for each outer iteration of the AL scheme (3.3), with warm-starting; i.e., each CG call starts with the estimate from the previous AL iteration. No such inner iterations are needed for the proposed AL approach (3.5). We examine convergence rate by computing the normalized root-mean-squared difference (NRMSD) of each iterate to a converged image, defined by $\|X_k - X_\infty\|_2 / \|X_\infty\|_2$, where $X_k = L_k + S_k$, and $\|\cdot\|_2$ denotes the vector l_2 norm. We obtained $X_\infty = L_\infty + S_\infty$ as a reference by averaging $X_\infty^{\text{AL-2}}$ and X_∞^{POGM} , as discussed below for each dataset, then computed the distance to the minimizer.

3.4.1 Cardiac perfusion dataset

Images for this dataset have size $N_x \times N_y = 128 \times 128$, with $N_t = 40$ temporal profiles and $N_c = 12$ coils. Data were retrospectively undersampled by a factor of 10, using the sampling pattern from [148], with fully sampled low spatial frequencies and low-density-sampled outer k -space. We used $\lambda_L = \lambda_S = 0.01$ as in [148], with scalings to match the original implementation, as discussed below in Section 3.5. For the variable splitting framework, the penalty parameters were empirically tuned to achieve fast convergence, with $\delta_1 = \delta_2 = 1/5$ for AL with CG, and $\delta_1 = 1/5, \delta_2 = 1/50$ for the efficient AL method. We ran the efficient AL and the POGM implementations for 24,000 seconds to obtain $X_\infty^{\text{AL-2}}$ and X_∞^{POGM} . In this case, the cost function values $f_\infty^{\text{AL-2}}$ and f_∞^{POGM} are within 10^{-16} relative difference from each other, and we averaged the results to obtain X_∞ and f_∞ . In this case, the NRMSD between $X_\infty^{\text{AL-2}}$ and X_∞^{POGM} is approximately $8.9 \times 10^{-13}\%$. As shown in Fig. 3.4.1, POGM converges the fastest overall. ISTA converges faster than the other three methods at the beginning, a phenomenon that could be due to the chosen step size. FISTA converges faster than the AL-based methods in this case, with the CG implementation being the slowest of all. The supplement contains additional figures showing the long-run behavior of the algorithms.

For this dataset, ISTA reached its stopping criterion from [148] at $k = 53$ iterations, after 30 seconds of elapsed compute time. Fig. 3.4.2 shows the magnitude of the reconstructed X_∞ , as well as images of all 5 algorithms at 30 seconds time elapsed.

To help visualize the reconstructed image in the spatial-temporal domain, the Supplement includes y - t images for a selected y -slice in the center, with comparison to the fully sampled case. The AL-2 and the POGM updates provide significantly faster convergence than the other methods in their corresponding algorithmic schemes.

3.4.2 Cardiac cine dataset

This dataset corresponds to images of size 256×256 , with 24 temporal frames and 12 coils, and a retrospective undersampling factor of 8. As in the cardiac perfusion case, we used $\lambda_L = 0.01, \lambda_S = 0.0025$ with additional scalings. For the AL-based methods, we used $\delta_1 = 1/10, \delta_2 = 1/20$ for AL-CG, and $\delta_1 = 1/10, \delta_2 = 1/100$ for AL-2. To obtain X_∞ and f_∞ , we ran the efficient AL and the POGM implementations for 30,000 seconds and averaged the results. The cost function values are within 10^{-16} relative difference from each other, and the NRMSD between $X_\infty^{\text{AL-2}}$ and X_∞^{POGM} is $1.4 \times 10^{-7}\%$.

Fig. 3.4.1 illustrates that AL-2 achieving faster convergence than the FISTA update, but slower than POGM. Fig. 3.4.3 shows results with the same run time cut-off of 48 seconds, with ISTA taking $k = 30$ iterations to reach the stopping criterion. We again observed superior rates of convergence of the efficient implementations, in both the proximal gradient and the variable splitting schemes.

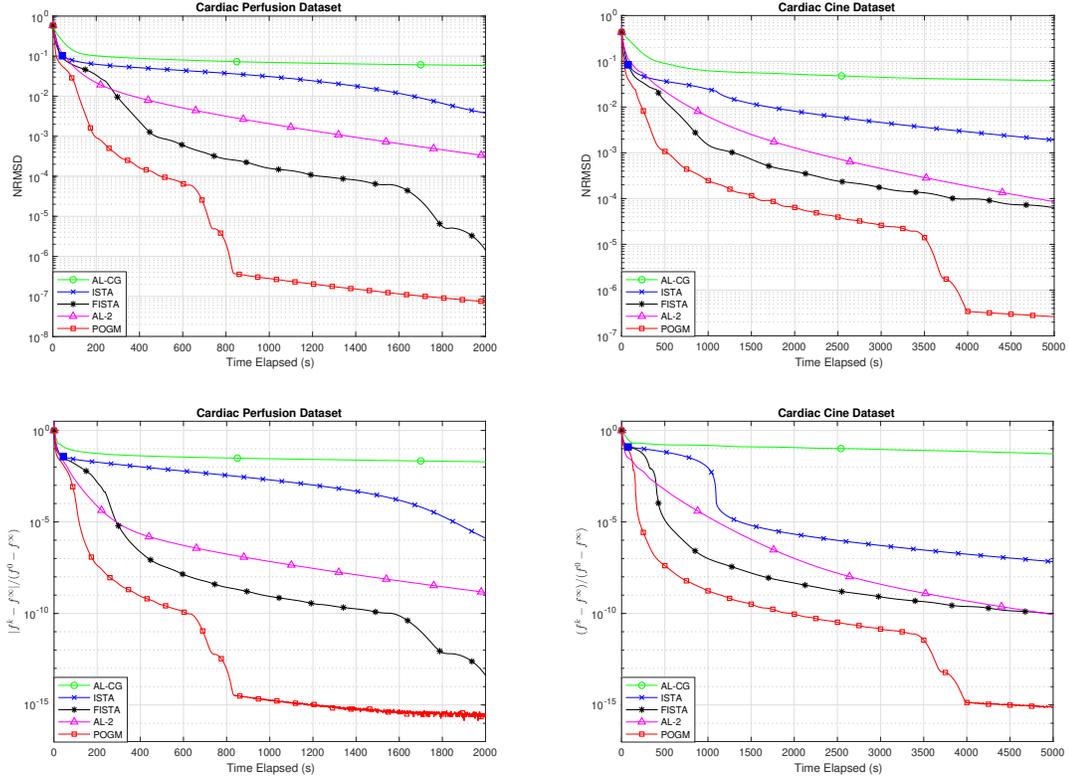


Figure 3.4.1: Convergence of five algorithms for cardiac perfusion and cine datasets, in terms of NRMSD to the minimizer (top row), as well as the cost function values (bottom row). AL-CG and AL-2 refer respectively to the AL methods with CG (3.3) and with **Algorithm 2** implementations. Every 100th iteration is marked by a dot, indicating their relative speeds. The blue square markers show when ISTA reaches the approximate stopping criteria corresponding to the implementation in [148].

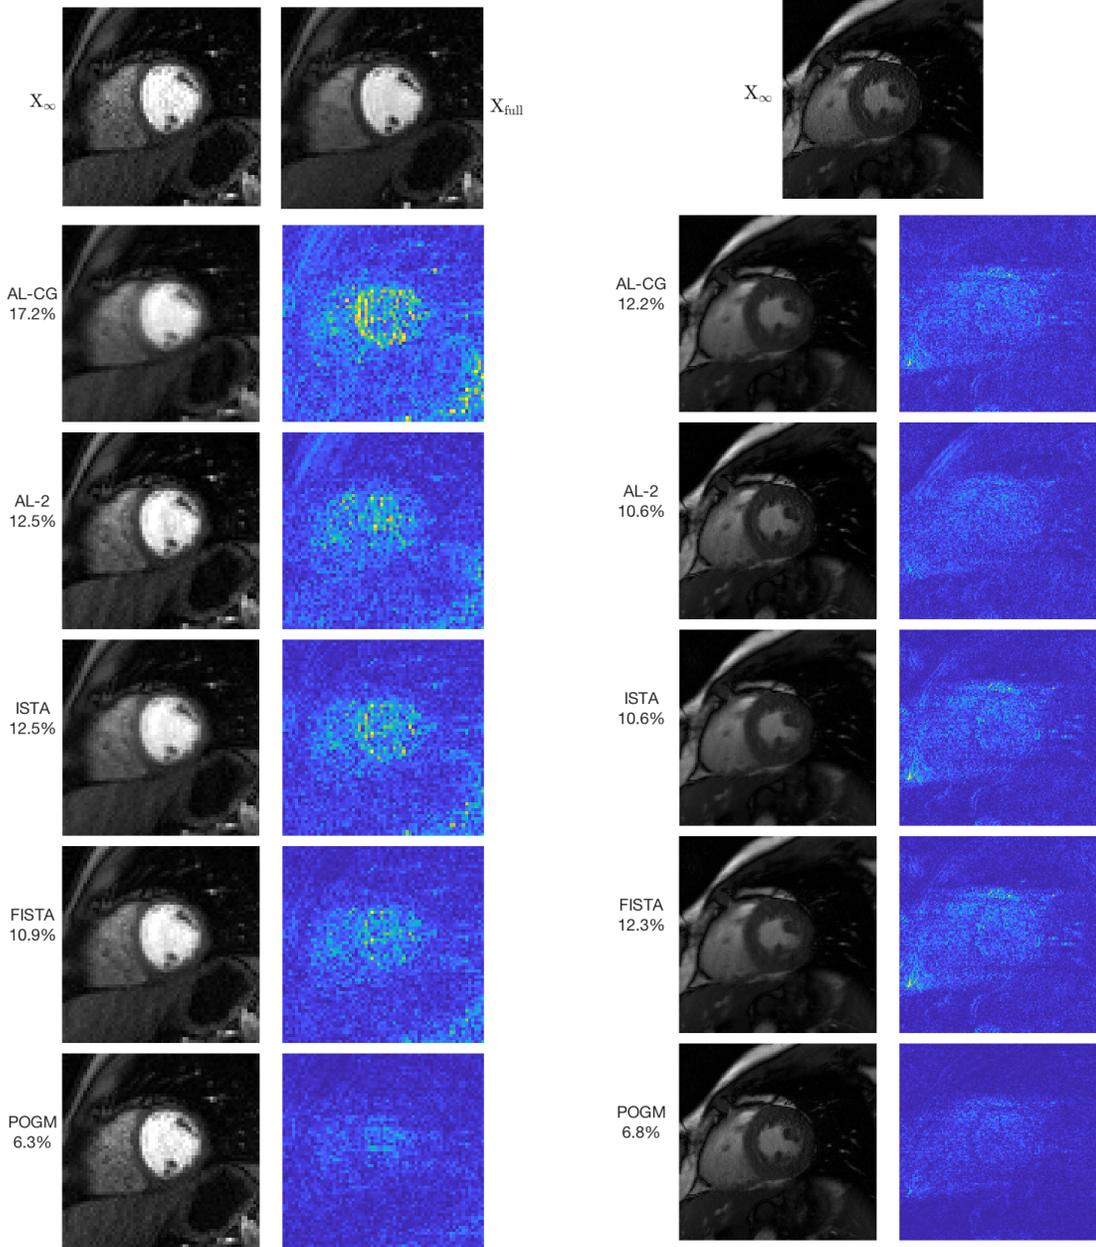


Figure 3.4.2: X_∞ , reconstructed and residual images on cardiac perfusion dataset. First row: X_∞ and X_{full} . Left column: reconstructed images on a scale of $[0,1]$ of one temporal frame, after the closest run time less than when ISTA reaches its stopping criterion (30 seconds). Right column: the residual images are plotted on a scale of $[0,0.2]$, with their corresponding NRMSD shown on the left of each row.

Figure 3.4.3: X_∞ , reconstructed and residual images on cardiac cine dataset. First row: X_∞ . Left column: reconstructed images on a scale of $[0,1]$ on one temporal frame, after the closest run time less than when ISTA reaches its stopping criterion (48 seconds). Right column: the residual images are plotted on a scale of $[0,0.2]$, with their corresponding NRMSD shown on the left of each row.

3.4.3 PINCAT phantom dataset

The ground truth phantom data provided by [119] has spatial dimension 128×128 with 50 temporal frames. To compare the algorithms in the multi-coil setting, we added simulated coil sensitivity maps of 32 coils (4 rings of 8 coils), with coil compression to reduce to $N_c = 8$ coils. Following the setup in [119], we applied a pseudo-radial undersampling mask Ω , i.e., a Cartesian trajectory that closely approximates a radial trajectory, with 24 spokes per frame, corresponding to an acceleration factor of $128/24 \approx 5.3$. We added zero mean Gaussian noise such that the signal to noise ratio is 46 dB. We tuned the regularization parameters λ_L to 0.0025 multiplied by the top singular value of L_0 , and λ_S to 0.05, each divided by a constant that captures the square root of the sum of squares of the coil sensitivity maps before the normalization $\mathbf{C}^* \mathbf{C} = \mathbf{I}$. The penalty parameters in the AL-based methods were empirically tuned to achieve fast convergence, with $\delta_1 = \delta_2 = 1/3$ for AL-CG, and $\delta_1 = 1/5, \delta_2 = 1/20$ for AL-2. We ran AL-2 and POGM for 30,000 seconds, and averaged the results to obtain X_∞ and f_∞ . The NRMSD between $X_\infty^{\text{AL-2}}$ and X_∞^{POGM} is approximately $4.4 \times 10^{-4}\%$, and the cost function values $f_\infty^{\text{AL-2}}$ and f_∞^{POGM} are within 10^{-12} relative difference from each other. The convergence behavior is similar to the *in vivo* case; Fig. 3.4.4 demonstrates again the superior convergence speed of AL-2 and POGM in the two schemes. For unknown reasons, AL-2 reaches a final cost that is about 10^{-10} higher than the proximal algorithms. This behavior is unimportant practically but still somewhat curious; it is unique to the PINCAT data. To visualize the results, Fig. 3.4.5 shows the ground truth X_{true} and the undersampled reconstruction $X_\infty = L_\infty + S_\infty$. See the Supplement for y - t images compared to the ground truth.

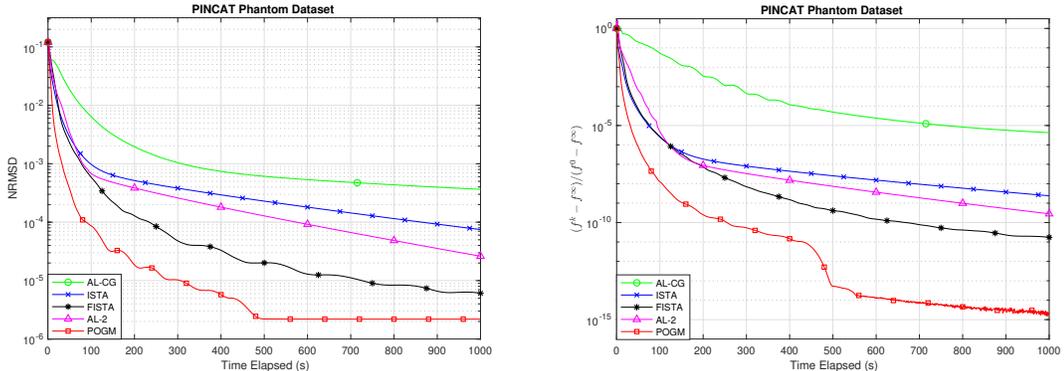


Figure 3.4.4: Convergence of five algorithms on PINCAT phantom dataset, in terms of NRMSD and cost to the minimizer. Every 100th iteration of each algorithm is marked by a dot, indicating their relative speeds.

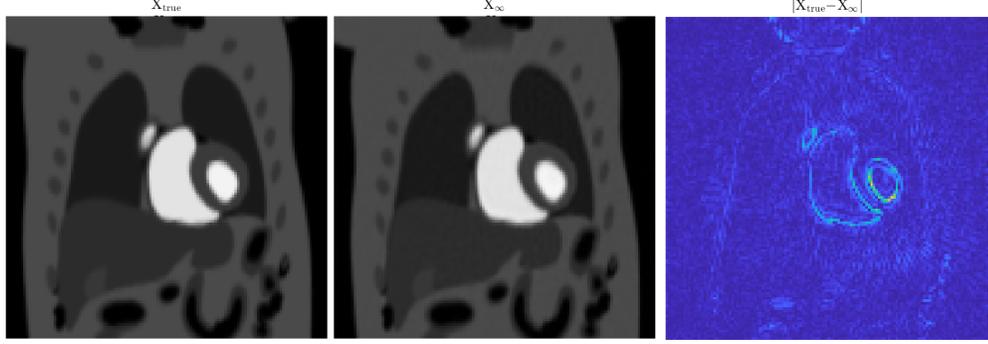


Figure 3.4.5: X_{true} , X_{∞} and residual images on PINCAT phantom dataset. From left to right: X_{true} of one temporal frame from the ground truth data, and its corresponding reconstruction X_{∞} as the average of $X_{\infty}^{\text{AL-2}}$ and X_{∞}^{POGM} , all on a scale of $[0,1]$. The residual image $|X_{\text{true}} - X_{\infty}|$ is on a scale of $[0,0.1]$.

3.5 Discussion

3.5.1 Alternative variable splitting

With undersampled multi-coil data, we have expressed the data acquisition operator as $\mathbf{E} = \mathbf{\Omega}Q\mathbf{C}$, and our proposed AL approach (3.5) splits the sampling mask $\mathbf{\Omega}$ from the Fourier encoding together with the coil sensitivity maps $Q\mathbf{C}$. An alternative is to split the Fourier encoding operator with undersampling, $\mathbf{\Omega}Q$, from the coil sensitivities \mathbf{C} , as proposed in [161]. Compared with (3.5), this splitting introduces a slight variation:

$$\min_{Z,X,L,S} \frac{1}{2} \|\mathbf{\Omega}QZ - \mathbf{d}\|_2^2 + \lambda_L \|L\|_* + \lambda_S \|\mathbf{T}S\|_1$$

$$\text{subject to } \begin{cases} Z = \mathbf{C}X, \\ X = L + S. \end{cases}$$

In this case, the update for Z involves the inverse of $Q^*\mathbf{\Omega}^*\mathbf{\Omega}Q + \delta_1 I$, which is block circulant and can be diagonalized by pre- and post-multiplication by FFT operators [111]. Hence with this splitting, efficient implementation of the inverse is again possible. We chose to implement the splitting (3.5) because of the simpler expressions (3.8), (3.9) for the Z and X updates, with diagonal matrix inverses due to the unitary property of Q . This leads to simpler updates with only inversion of diagonal matrices.

Although the proposed AL scheme shows empirical convergence, it does not have

analytical convergence guarantee as in generalized Alternating Direction Method of Multipliers (ADMM). To compare the variable splitting (3.5) to the ADMM scheme, we note that it is equivalent to a formulation in the monotropic programming framework [51]:

$$\min_U f(U)$$

subject to $AU = b$,

$$\text{where } U = \begin{bmatrix} Z \\ X \\ L \\ S \end{bmatrix}, A = \begin{bmatrix} I & -\mathbf{QC} & 0 & 0 \\ 0 & I & -I & -I \end{bmatrix}, \text{ and } b = 0.$$

Since the last two columns of A are linearly dependent, this splitting scheme does not satisfy the sufficient conditions for the convergence guarantee of ADMM [51]. To satisfy those conditions, one could introduce an alternative variable splitting that meets the convergence criteria of ADMM, but at the cost of more variables, and thus potentially slightly slower convergence [111]. We did not investigate that approach here since AL-2 empirically converged well, despite not satisfying the sufficient conditions in [51]. In addition, POGM converged faster, and is practically preferable because it does not require any AL-type tuning parameters.

3.5.2 ISTA implementation in comparison with [148]

Our ISTA implementation is based on the algorithm discussed mathematically in [148]. However, we did not directly use the MATLAB code provided for [148] at <http://cai2r.net/resources/software/ls-reconstruction-matlab-code>, because that code is slightly inconsistent with the math in [148].

The first inconsistency is the implementation of the operator \mathbf{E} . The MATLAB code for the Hermitian adjoint \mathbf{E}^* operation contains an additional division by the sum of squares of the coil sensitivity maps that causes inconsistency between the forward and the adjoint operations, preventing convergence to the minimizer of the stated cost function. We modified the code so that \mathbf{E}^* is the exact adjoint of \mathbf{E} , so that all of the algorithms, including ISTA, can converge to the same cost function. To obtain similar images as those in [148], we preprocess by dividing the given coil sensitivities \mathbf{C} by the square root of its sum of squares. Since this factor is close to being constant across the image, we absorb it into the regularization parameters λ_L and λ_S , to ensure a consistent setup with [148].

Another implementation difference involves the singular value threshold. In the cost function (3.1), the nuclear norm regularization parameter λ_L is a fixed constant, but in the provided code, λ_L changes across iterations, with a factor that depends on the leading singular value of L . This “moving target” cost function would make it impossible to compare the convergence rates of different algorithms. To ensure fair comparison of all the algorithms, while maintaining similar overall regularization as in [148], we fix λ_L by considering the leading singular value of L_∞ , produced by running the original implementation until convergence.

The provided implementation has stopping criteria based on the maximum number of iterations and the tolerance of the change in updates. With the above modifications, we stop our ISTA implementation when it reaches the same cost function value as at the stopped points, and compare the NRMSD at these points with other algorithms, as indicated by the blue square markers in Fig. 3.4.1.

3.6 Conclusion

This chapter presents efficient algorithms for the L+S reconstruction of dynamic parallel MRI. Within the proximal gradient category, in place of using ISTA to solve the optimization problem, we consider updates by FISTA and POGM. Both methods can be efficiently formulated within the L+S framework, preserving the computational simplicity of the original ISTA implementation. Experiments with two cardiac datasets in [148] and a phantom dataset in [119] verify their accelerated rates of convergence.

For AL-based approaches, we also proposed an efficient variable splitting scheme that considers the structure of the data acquisition operator. In particular, we split the variables based on the undersampling mask, the Fourier transform operator and the coil sensitivity maps. While the existing splitting scheme for the L+S model requires CG approach to solve for the quadratic updates [188], our proposed formulation leads to a diagonal matrix inverse that can be easily computed. Numerical experiments again confirms its superior convergence rate, compared with the existing implementation.

Although there is no strict convexity guarantee for the L+S optimization problem, our experimental results suggest high similarities between the reconstructed images by the AL and the PGM schemes, due to the observed low NRMSD of both $X_\infty^{\text{AL-2}}$ and X_∞^{POGM} . In the implementation perspective, however, AL-based methods in the L+S model requires the tuning of two additional penalty parameters, whereas POGM

has no extra tuning parameters. This practical benefit, combined with the empirical faster convergence of POGM seen in the examples, make POGM our recommended approach for solving L+S reconstruction problems for dynamic MRI.

3.7 Supplementary material

This supplement presents additional figures of space-time images and long-run convergence behaviors. We also illustrate the flexibility of our algorithm by a non-Cartesian experiment.

3.7.1 y - t images

To help visualize the reconstruction results for different time frames, we provide figures of the reconstructed images in the y - t domain, compared to their reference images, as shown in Figs. 3.7.1 and 3.7.2.

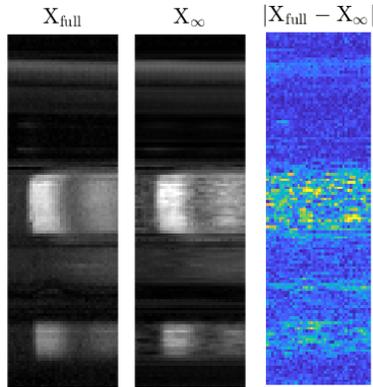


Figure 3.7.1: y - t images reconstructed from fully and undersampled perfusion dataset. The reconstructed images are on a scale of $[0,1]$, and the central slice for y is taken. The difference image is plotted on a scale of $[0,0.2]$.

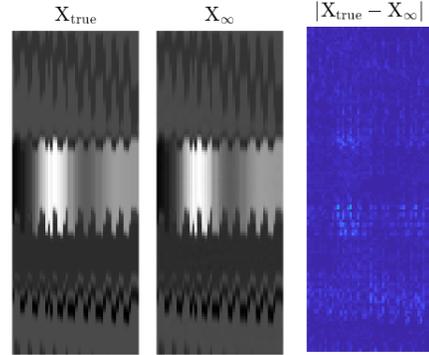


Figure 3.7.2: y - t ground truth and image reconstructed from undersampled PINCAT dataset. The images are on a scale of $[0,1]$, and the central slice for y is taken. The difference image is plotted on a scale of $[0,0.2]$.

3.7.2 Long-run behaviors

Figs. 3.7.3 and 3.7.4 show the results of running many more iterations of the algorithms investigated to illustrate the long-run behaviors of the methods, for the Cardiac Perfusion, Cardiac Cine and the PINCAT phantom datasets.

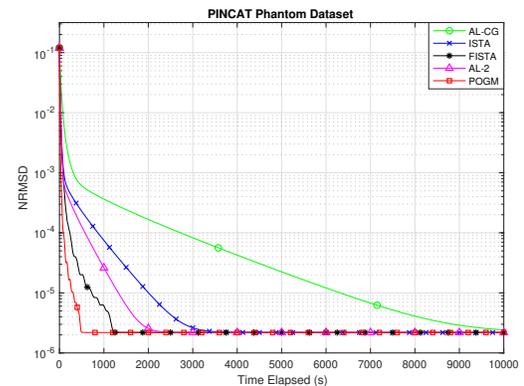
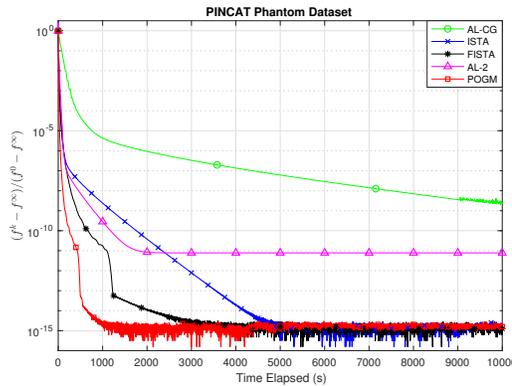
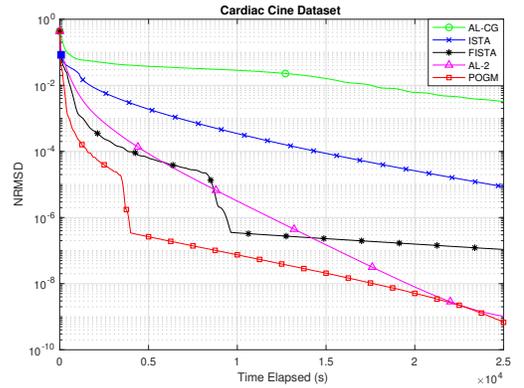
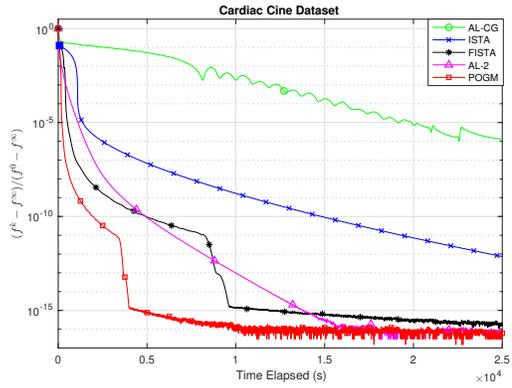
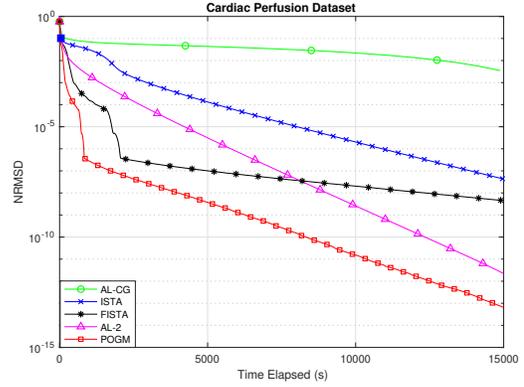
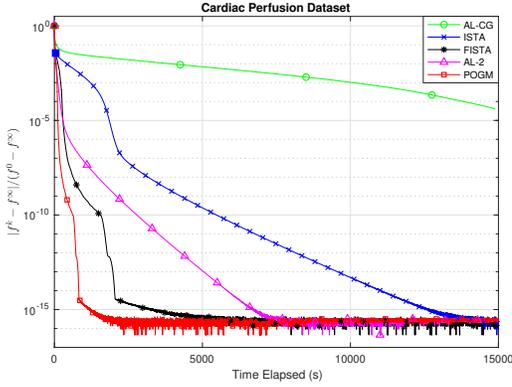


Figure 3.7.3: Long-run cost convergence of five algorithms on three datasets. Every 500th iteration is marked by a dot, indicating their relative speeds.

Figure 3.7.4: Long-run NRMSD convergence of five algorithms on three datasets. Every 500th iteration is marked by a dot, indicating their relative speeds.

3.7.3 Non-Cartesian experiment

We performed an additional experiment on the Abdominal dynamic contrast-enhanced (DCE) MRI dataset examined in [148]. This non-Cartesian dataset uses a golden-angle radial sampling pattern, and corresponds to images of size 384×384 , with 28 temporal frames and 12 receiver coils, having an acceleration factor of 12.

As in the setup in the MATLAB code provided by [148], we used $\lambda_L = 0.025$, $\lambda_S = 2e-5$, with adjusted scaling as described in Section 3.5.2. Here, the step size for PGM depends on the maximum eigenvalue of $\mathbf{E}^*\mathbf{E}$, and we estimated it using power iteration. Due to the observed faster convergence of POGM than the AL methods, and because an additional variable splitting would be needed for an AL approach for the non-Cartesian case, we focused on the proximal methods for this experiment. To obtain X_∞ and f_∞ , we ran the POGM implementations for $1e5$ seconds. Fig. 3.7.5 illustrates that POGM achieves the fastest convergence among the three methods. Fig. 3.7.6 shows reconstructed image results with the same run time cut-off of 153 seconds, with ISTA taking $k = 9$ iterations to reach the stopping criterion.

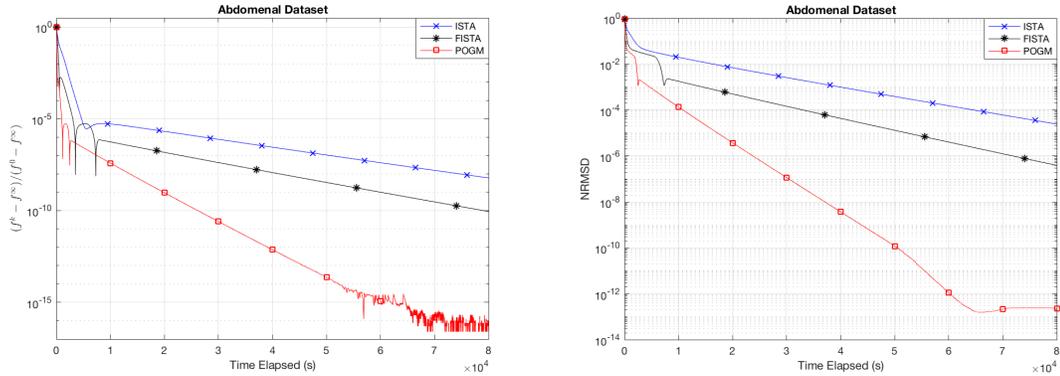


Figure 3.7.5: NRMSD and cost convergence of three proximal algorithms on abdominal DCE dataset. Every 500th iteration of each algorithm is marked by a dot, indicating their relative speeds.

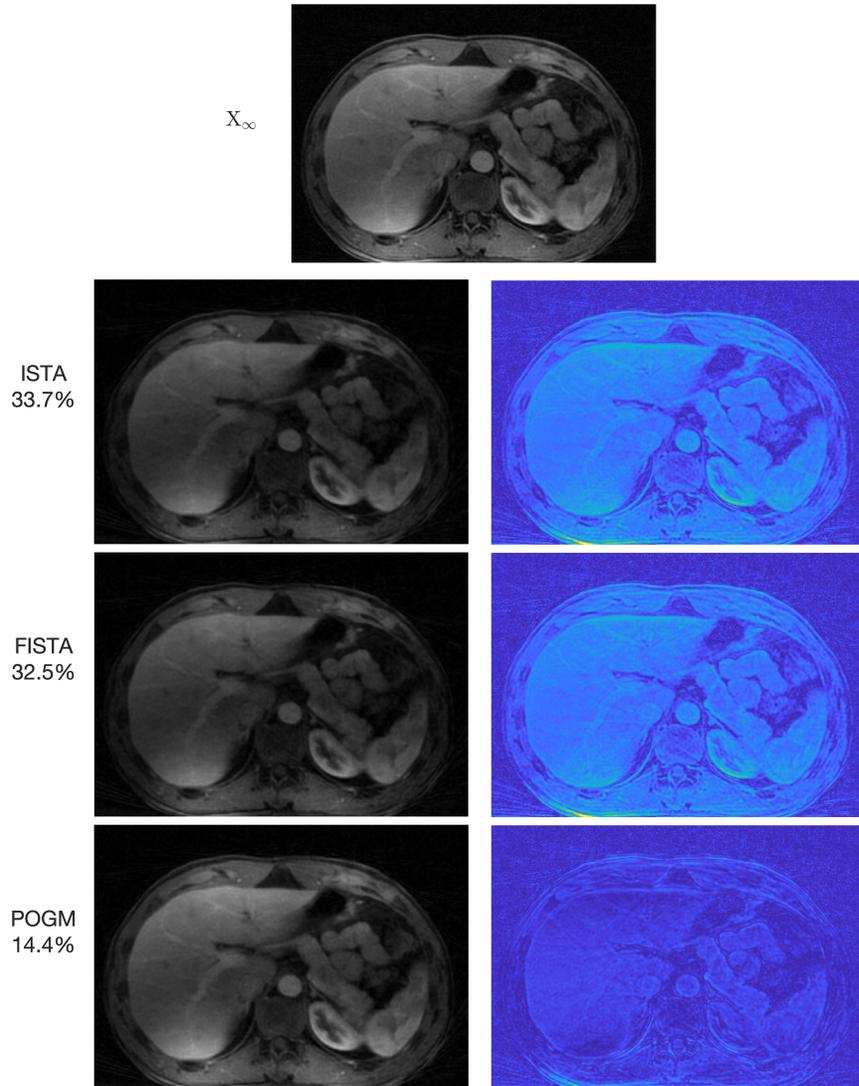


Figure 3.7.6: Reconstructed and residual images by ISTA, FISTA, and POGM with similar runtime on abdominal DCE dataset. First row: X_∞^{POGM} . Left column: reconstructed images on a scale of $[0, 5e-4]$ on one temporal frame, after the closest runtime less than when ISTA reaches its stopping criterion (153 seconds). Right column: the residual images are plotted on a scale of $[0, 2e-4]$, with their corresponding NRMSD shown on the left of each row.

Chapter 4

Efficient Regularized Field Map Estimation

Functional MRI with long readout times is sensitive to magnetic field inhomogeneity effects, which if uncorrected could lead to image distortion and artifacts, especially near air/tissue interfaces. Field inhomogeneity estimation is also important in chemical shift based water-fat imaging. Regularized field map estimation methods that account for phase wrapping and noise involve nonconvex cost functions that require iterative algorithms. Most existing minimization techniques were computationally or memory intensive for 3D datasets, and are designed for single-coil MRI. This chapter considers 3D MRI with optional consideration of coil sensitivity, and addresses the multi-echo field map estimation and water-fat imaging problem. Our efficient algorithm uses a preconditioned nonlinear conjugate gradient method based on an incomplete Cholesky factorization of the Hessian of the cost function, along with a monotonic line search. Numerical experiments show the computational advantage of the proposed algorithm over state-of-the-art methods with similar memory requirements.

4.1 Introduction

In magnetic resonance imaging (MRI), scans with long readout times require correction for magnetic field inhomogeneity during reconstruction to avoid artifacts [65, 92,139,165,174]. Field inhomogeneity is also a nuisance parameter in chemical shift based water-fat imaging techniques [24,68,82,82,88,168]. Field map estimation is thus crucial to field-corrected MR image reconstruction, and for fat and water image separation.

One field map estimation approach is to acquire MR scans at multiple echo times

This chapter presents our published work in the journal paper [116].

(usually 2 or 3), where a small echo time difference can help resolve any phase wrapping issues and a large echo time difference can help improve SNR. One can then estimate field inhomogeneity using images reconstructed from these scans [65]. Since field maps tend to be smooth within tissue, estimation methods with smoothness assumptions have been proposed for water-fat separation, including region growing techniques [49,90,123,167,206,210], filtering [201], curve fitting [89,173,178], multiresolution and subspace approaches [123,177,178,189], and graph cut algorithms [21]. To improve robustness of water and fat separation and reduce ambiguity of assignment, field map pre-estimation methods such as demodulation [48] and magnetization transfer [171] have been proposed as part of the water-fat imaging framework. Most of these methods, however, use various approximations to account for phase wrapping between different acquisitions. In contrast, regularized estimation methods [65,81,82] have been proposed to account for both phase wrapping and the smoothness of the field map from multiple acquisition images. Because the field map affects image phase, these approaches involve a nonconvex optimization problem that requires iterative methods.

To solve such optimization problems, [10,65,88] use a majorization-minimization (MM) approach by introducing a quadratic majorizer for their cost functions. The MM approach decreases the cost monotonically, but is computationally intensive, especially for large-scale datasets. Other regularized field map estimation minimization techniques quantize the solution space [81,82] and may require a second descent algorithm to produce sufficiently smooth estimates. An alternative minimization technique [147] uses nonlinear conjugate gradient (NCG) with a monotonic line search (MLS), and explored various preconditioners in the 3D single-coil case.

This chapter considers the regularized field map estimation problem in the 3D multi-coil MRI setting. In particular, we consider a generalized cost function in the multi-coil case for both multi-echo field map estimation and water-fat imaging. We minimize it by a NCG algorithm with an efficient MLS and an iteration-dependent preconditioner based on an incomplete Cholesky factorization [127] of the Hessian of the cost function. The incomplete Cholesky factorization has been applied to field inhomogeneity estimation using surface fitting [107], and recently to single-coil field map estimation with a similar cost function [147]. In addition to faster convergence, this preconditioner exploits the sparse structure of the Hessian, thus it is memory efficient and scales to 3D datasets. Compared to previous works [10,88,147], our new approach unifies the field map correction and the water-fat imaging problems, with a generalized expression that optionally considers multiple coils in MRI. Our efficient

algorithm on this problem shows significant computational and storage advantages compared with existing MM and NCG methods.

The rest of this chapter is organized as follows. Section 4.2 describes the optimization problem for the field map estimation problems for multi-coil MRI. Section 4.3 presents the NCG-MLS optimization scheme with the proposed preconditioner. Section 4.4 reports simulated and real experimental results, followed by conclusions in Section 4.5.

4.2 Problem formulation

We are given reconstructed images $\mathbf{y}_{cl} \in \mathbb{C}^{N_v}$ for the c th receiver coil of the l th scan, with $c = 1, \dots, N_c$, $l = 1, \dots, L$, where N_v denotes the total number of voxels in the image, N_c denotes the number of coils, and $L \geq 2$ denotes the number of echo times. We model the field inhomogeneity effect as

$$y_{clj} = e^{i\omega_j t_l} s_{cj} x_{lj} + \epsilon_{clj}, \quad (4.1)$$

where $j = 1, \dots, N_v$ is the voxel index, $\boldsymbol{\omega} \in \mathbb{R}^{N_v}$ is the unknown field map, $t_l \in \mathbb{R}$ is the echo time shift of the l th scan, $\mathbf{s}_c \in \mathbb{C}^{N_v}$ is the (known) coil sensitivity map for the c th coil, and $\epsilon_{cl} \in \mathbb{C}^{N_v}$ denotes the noise. For single-coil MRI, or when the coil images are combined as a preprocessing step, we have $N_c = 1$ and $\mathbf{s} = \mathbf{1}$ in (4.1).

The unknown image $\mathbf{x}_l \in \mathbb{C}^{N_v}$ for the l th echo is problem-dependent, where

$$x_{lj} = \begin{cases} m_j & \text{in field map estimation,} \\ m_{w,j} + m_{f,j} \sum_{p=1}^P \alpha_p e^{i2\pi \Delta_{f,p} t_l} & \text{in water-fat imaging,} \end{cases}$$

where $\mathbf{m}, \mathbf{m}_w, \mathbf{m}_f \in \mathbb{C}^{N_v}$ are respectively the magnetization, water, and fat components, and $\Delta_{f,p} \in \mathbb{R}$ denotes the (known) frequency shifts of P fat peaks in the multipeak fat model [27,82,166] with relative amplitudes $\sum_{p=1}^P \alpha_p = 1$ that can be estimated and averaged over all fat pixels as a preprocessing step by existing methods [207]. The goal of the field map estimation problem is to estimate $\boldsymbol{\omega}$ and \mathbf{x} given \mathbf{y} and \mathbf{s} .

Assuming the noise ϵ is zero-mean, white complex Gaussian, the joint maximum-likelihood (ML) estimates of the field map $\boldsymbol{\omega}$ and image \mathbf{x} are the minimizers of the

negative log-likelihood as follows:

$$\begin{aligned} & \underset{\boldsymbol{\omega}, \mathbf{x}}{\operatorname{argmin}} \tilde{\Phi}(\boldsymbol{\omega}, \mathbf{x}), \text{ where} \\ \tilde{\Phi}(\boldsymbol{\omega}, \mathbf{x}) &= \sum_{j=1}^{N_v} \sum_{l=1}^L \sum_{c=1}^{N_c} |y_{clj} - e^{i\omega_j t_l} s_{cj} x_{lj}|^2. \end{aligned} \quad (4.2)$$

For a given field map $\boldsymbol{\omega}$, the ML estimate of \mathbf{x} has a closed-form expression [10,81] that one can substitute into (4.2) to give a cost function in terms of $\boldsymbol{\omega}$:

$$\Phi(\boldsymbol{\omega}) = \min_{\mathbf{x}} \tilde{\Phi}(\boldsymbol{\omega}, \mathbf{x}) = \sum_{j=1}^{N_v} \sum_{m,n=1}^L \sum_{c,d=1}^{N_c} \phi_{cdmnj}(\omega_j), \quad (4.3)$$

where

$$\begin{aligned} \phi_{cdmnj}(\omega_j) &:= |r_{cdmnj}| \left[1 - \cos(\angle r_{cdmnj} + \omega_j(t_m - t_n)) \right], \\ r_{cdmnj} &:= \frac{\Gamma_{mn}}{\sum_{c'=1}^{N_c} |s_{c'j}|^2} s_{cj} s_{dj}^* y_{cmj}^* y_{dnj}, \\ \boldsymbol{\Gamma} &:= \boldsymbol{\gamma}(\boldsymbol{\gamma}^* \boldsymbol{\gamma})^{-1} \boldsymbol{\gamma}^*, \end{aligned} \quad (4.4)$$

where \cdot^* denotes the complex conjugate, and $L \times L$ matrix $\boldsymbol{\Gamma}$ is defined in terms of

$$\boldsymbol{\gamma} = \begin{cases} \mathbf{1} & \text{in field map estimation,} \\ \left[\mathbf{1} \quad \sum_{p=1}^P \alpha_p e^{i2\pi \Delta_{f,p} \mathbf{t}} \right] & \text{in water-fat imaging,} \end{cases} \quad (4.5)$$

in which $\mathbf{1}$ denotes an all one vector of length L , and the exponential is applied element-wise. In the field map estimation case, this simplifies to $\Gamma_{mn} = 1/L \forall m, n$.

As B_0 field maps tend to be spatially smooth in MRI, we add a regularization term to (4.3) to form a penalized-likelihood (PL) cost function

$$\Psi(\boldsymbol{\omega}) = \Phi(\boldsymbol{\omega}) + \frac{\beta}{2} \|\mathbf{C}\boldsymbol{\omega}\|_2^2, \quad (4.6)$$

where \mathbf{C} is a first or second order finite difference operator with optional spatial weights as in [82]. Such regularization has been used in many other prior works [10, 65, 147].

4.3 Efficient algorithm

Several approaches have been proposed to solve the field map estimation problem in the single-coil setting, but are demanding in computation or memory. In particular, a quadratic majorizer with a diagonal Hessian [65] takes many iterations to converge even for 2D images, and a quadratic majorizer with an optimal curvature that inverts a $N_v \times N_v$ Hessian matrix [10] is memory-limited to small-scale data. In water-fat imaging, [82,176] process data in a single-coil manner using the graph cut method. Since graph cut requires discretization, [82] proposes to overcome this limitation by additionally running a descent algorithm such as in [88], which considers a quadratic majorizer with a diagonal Hessian that converges slowly.

Here, we optimize (4.6) using NCG with a monotonic line search [147], and consider a preconditioner with efficient computation and memory storage. Our field map estimation procedure is tabulated in the **Algorithm** below. For NCG, we choose the Polak-Ribiere update to compute a μ^i that satisfies the conjugacy condition [157].

After estimating the field map $\hat{\omega}$, we estimate the water and fat components for each voxel in water-fat imaging by applying the closed-form expression [81] using $\hat{\omega}$:

$$\begin{bmatrix} m_{w,j} \\ m_{f,j} \end{bmatrix} = \left((\gamma \cdot \text{diag}(e^{i\omega_j \mathbf{t}})) \otimes \mathbf{s}_j \right)^\dagger \mathbf{y}_j, \quad (4.7)$$

where \otimes denotes the Kronecker product, $(\cdot)^\dagger$ denotes the pseudo inverse, and $\mathbf{s}_j \in \mathbb{C}^{N_c}$ denotes the coil sensitivity map for the j th voxel.

Next we present our initialization, choice of preconditioner, and derive our iterative monotone line search algorithm in the multi-coil setting.

4.3.1 Initialization

For field map estimation, we initialize ω by a field map computed from the phase of the first two echoes of the coil combined images:

$$(\omega_j)^0 = \angle \left[\left(\sum_{c=1}^{N_c} s_{c1j}^* y_{c1j} \right)^* \left(\sum_{d=1}^{N_c} s_{d2j}^* y_{d2j} \right) \right] / (t_2 - t_1). \quad (4.8)$$

To initialize ω for water-fat imaging, we follow [88] and sweep through a range of 100 values from $-|\Delta_f/2|$ to $|\Delta_f/2|$ for each voxel, and choose the value with minimal cost (4.3), denoted as $\tilde{\omega}^0$. We then run a few CG iterations to minimize a penalized

weighted least squares (PWLS) problem

$$\boldsymbol{\omega}^0 = \underset{\boldsymbol{\omega}}{\operatorname{argmin}} \sum_{j=1}^{N_v} \rho_j (\omega_j - \tilde{\omega}_j^0)^2 + \frac{\beta}{2} \|\mathbf{C}\boldsymbol{\omega}\|_2^2, \quad (4.9)$$

where the spatial weights

$$\rho_j = \sum_{m,n=1}^L \sum_{c,d=1}^{N_c} |r_{cdmnj}|$$

are given by (4.4). We then use $\boldsymbol{\omega}^0$ as our initial estimate in the water-fat case.

To reduce ambiguity of water and fat assignment, one can also consider robust initialization schemes such as demodulation [48] or magnetization transfer [171].

4.3.2 Preconditioning matrices

To accelerate the NCG-based algorithm, given gradient \mathbf{g}^i of the cost at the i th NCG iteration, we explore a preconditioner \mathbf{P}^i with memory efficient implementation of $(\mathbf{P}^i)^{-1}\mathbf{g}^i$ using an incomplete Cholesky factorization [127]. In particular, the gradient $\mathbf{g} \in \mathbb{R}^{N_v}$ is given by

$$\mathbf{g} = \nabla\Psi(\boldsymbol{\omega}) = \nabla\Phi(\boldsymbol{\omega}) + \beta\mathbf{C}^\top\mathbf{C}\boldsymbol{\omega}, \quad (4.10)$$

where

$$\begin{aligned} (\nabla\Phi(\boldsymbol{\omega}))_j &= \sum_{m,n=1}^L \sum_{c,d=1}^{N_c} |r_{cdmnj}| (t_m - t_n)^2 \\ &\quad \cdot \sin(\angle r_{cdmnj} + \omega(t_m - t_n)). \end{aligned}$$

The Hessian of the cost (4.6) at the i th iteration is the sum of a diagonal matrix and an (approximately, due to the support mask) block Toeplitz with Toeplitz block (BTTB) matrix:

$$\mathbf{H}^i = \mathbf{D}^i + \beta\mathbf{C}^\top\mathbf{C} \in \mathbb{R}^{N_v \times N_v}, \quad (4.11)$$

where \mathbf{C} is the finite difference operation and $\mathbf{D}^i = \operatorname{diag}(d_j^i) \succeq 0$, where the Hessian of the negative log-likelihood has diagonal elements given by

$$d_j^i = \sum_{m,n=1}^L \sum_{c,d=1}^{N_c} \kappa_{cdmnj}(u_{cdmnj}(\omega_j^i)), \quad (4.12)$$

with

$$\begin{aligned} \kappa_{cdmnj}(u) &= |r_{cdmnj}|(t_m - t_n)^2 \frac{\sin(u)}{u}, \text{ and} \\ u_{cdmnj}(\omega) &= (\angle r_{cdmnj} + \omega(t_m - t_n)) \bmod \pi. \end{aligned} \quad (4.13)$$

In fact, the finite difference term can be seen as analogous to the Laplacian that arises in partial differential equations (PDE), where incomplete Cholesky factorization has also been used in its solvers [101].

Since the terms r_{cdmnj} and $t_m - t_n$ are shared across iterations, we precompute them at the initialization stage to efficiently calculate the gradient and Hessian at each iteration i . Note also that \mathbf{H}^i is positive definite as long as at least one value of d_j^i is positive (which is true for any nontrivial problem).

Although \mathbf{H}^i is sparse and banded, its inverse is approximately full, so directly computing the inverse would require far too much memory. To reduce memory, we propose to use a preconditioner that approximates the symmetric Hessian with a LU factorization of the form

$$\mathbf{P}^i = \mathbf{L}^i(\mathbf{L}^i)^\top \approx \mathbf{H}^i, \quad (4.14)$$

where $\mathbf{L}^i \in \mathbb{R}^{N_v \times N_v}$ is sparse lower triangular, enabling efficient computation (via back-substitution) of $(\mathbf{P}^i)^{-1} \mathbf{g}^i$ in the precondition step. Taking advantage of the sparsity and positive definiteness of our Hessian (4.11), preconditioning with an incomplete Cholesky factorization reduces both computation and memory. A popular form of the incomplete Cholesky factorization matches the matrix \mathbf{H} on its nonzero set, thus is at least as sparse as \mathbf{H} . Similar preconditioning with incomplete LU factorization has been used for simulating anisotropic diffusion in MRI [98]. In practice, for a better approximation one can control the sparsity of the factors by defining a tolerance on the magnitude of the elements of \mathbf{H} (below which entries in the factors are set to zero), with the trade-off between approximation accuracy and memory storage.

Fig. 4.3.1 illustrates the memory improvement by a toy problem of image size $20 \times 16 \times 8$, where we compute $\mathbf{H} = \mathbf{D} + \beta \mathbf{C}^\top \mathbf{C}$ and its inverse, with randomly chosen diagonal elements $d_j \in (0, 0.1)$ and $\beta = 0.1$. Fig. 4.3.1 considers the incomplete Cholesky factorization without tolerance, denoted \mathbf{L}_0 , and with a tolerance of $H_{\max} \times 10^{-3}$, denoted \mathbf{L}_t , where H_{\max} is the element in \mathbf{H} with maximum magnitude. Fig. 4.3.1 shows the sparse structure of \mathbf{H} , its nonsparse inverse \mathbf{H}^{-1} , and the Cholesky factorizations as well as their approximation errors. Table 4.1 shows the number of nonzero elements of each matrix, their memory storage, and their errors that affect the convergence rate, using the normalized root mean square error (NRMSE)

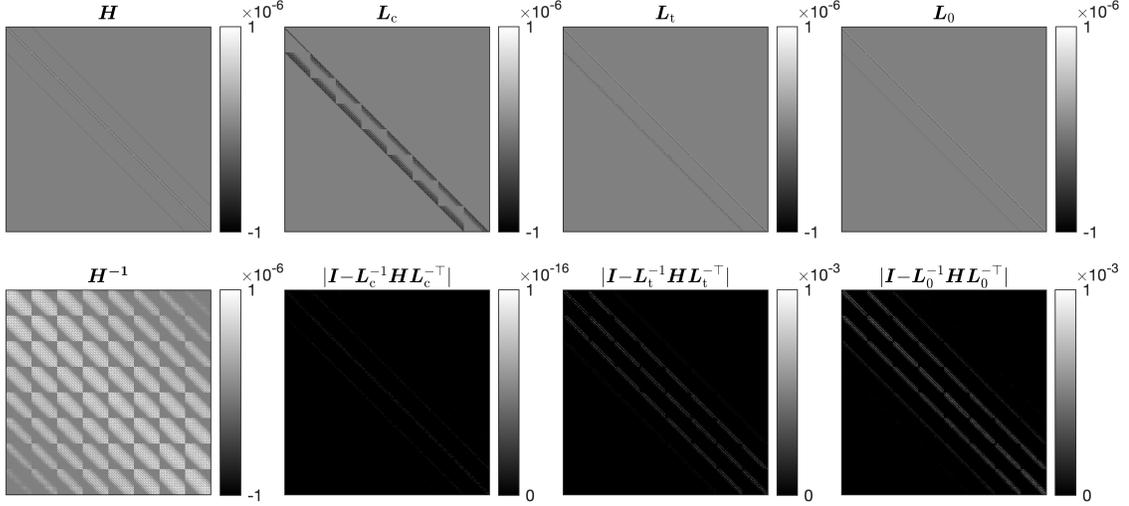


Figure 4.3.1: Factorization and error images on a toy problem of size $20 \times 16 \times 8$.

	\mathbf{H}	\mathbf{H}^{-1}	\mathbf{L}_c	\mathbf{L}_t	\mathbf{L}_0
Number of nonzeros ($\times 10^5$)	1.67	655	72.5	1.77	0.96
Storage (megabytes)	0.31	100.1	11.9	0.53	0.27
NRMSE			3e-16	4e-3	3e-2

Table 4.1: Number of nonzeros, memory usage, and NRMSE of factorizations on a toy problem of size $20 \times 16 \times 8$.

$\|\mathbf{I} - \mathbf{L}^{-1}\mathbf{H}\mathbf{L}^{-T}\|_{\text{F}}/\sqrt{N_v}$ for each factorization \mathbf{L} in our example. Fig. 4.3.2 illustrates how the sparsity of \mathbf{L}_t changes with respect to its tolerance by showing the percentage of nonzero elements in \mathbf{L}_t versus the scaling factor of H_{\max} in the tolerance.

For memory storage in this case, the number of nonzero elements in the incomplete Cholesky factor without tolerance \mathbf{L}_0 is more than 70 times less than that in the (complete) Cholesky factor \mathbf{L}_c , with more than 40 times memory saving. In general, we observe (by the banded structures) that the number of nonzero elements of \mathbf{L}_c is lower bounded by $(N_v - N_x N_y)N_x N_y$, while that of \mathbf{L}_0 is upper bounded by $4N_v$. This leads to the generalization that \mathbf{L}_0 is at least $(N_v - N_x N_y)/(4N_z)$ times more sparse than \mathbf{L}_c , which scales significantly with the problem size. The storage of the incomplete Cholesky factor with tolerance \mathbf{L}_t depends on the tolerance, and with the choice of tolerance here we observe a 40 times fewer nonzero values, saving memory by a factor of more than 20 compared with \mathbf{L}_c .

The trade-off with a sparser factorization, however, is a worse approximation error. This is reflected in the error matrices in Fig. 4.3.1 and the NRMSE in Table 4.1. While \mathbf{L}_0 has lower memory usage than \mathbf{L}_t , the inverse is a worse approximation to \mathbf{H}^{-1} .

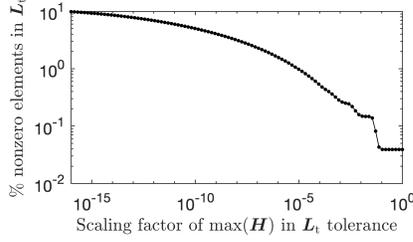


Figure 4.3.2: Change of sparsity of \mathbf{L}_t with respect to the scaling factor of H_{\max} in its tolerance.

In practice, nevertheless, both incomplete factorizations $\mathbf{L}\mathbf{L}^\top$ are positive definite, so as preconditioners they provide a descent direction in addition to storage advantage, whereas storing \mathbf{L}_c is infeasible for realistically sized 3D datasets.

4.3.3 Monotonic step size line search

With a search direction given by NCG, the choice of step size is important for convergence of the algorithm. To avoid multiple function evaluations required by backtracking line search algorithms [138], we implement a recursive line search algorithm using a quadratic majorizer with an optimal curvature, which guarantees monotone decrease of the cost function [58].

In the line search step, given a current field map estimate $\boldsymbol{\omega}^i$ and a search direction $\mathbf{z}^i \in \mathbb{R}^{N_v}$, we aim to find a step size that minimizes the cost (4.6):

$$\hat{\alpha} = \underset{\alpha}{\operatorname{argmin}} f(\alpha), \text{ where}$$

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

We iteratively minimize the nonconvex problem (4.15) using a quadratic majorizer based on Huber's method [87, p. 184] at the k th inner iteration (dropping outer iteration i for brevity):

$$\begin{aligned} q_k(\alpha) &= \Phi(\boldsymbol{\omega} + \alpha^{(k)} \mathbf{z}) \\ &+ \mathbf{z}^\top \nabla \Phi(\boldsymbol{\omega} + \alpha^{(k)} \mathbf{z})(\alpha - \alpha^{(k)}) \\ &+ \frac{1}{2} d^{(k)} (\alpha - \alpha^{(k)})^2 + \frac{\beta}{2} \|\mathbf{C}(\boldsymbol{\omega} + \alpha \mathbf{z})\|_2^2, \end{aligned}$$

where the optimal curvature is given by [10]

$$d^{(k)} = \sum_{j=1}^{N_v} |z_j|^2 d_j^{(k)}, \text{ where}$$

$$d_j^{(k)} = \sum_{m,n=1}^L \sum_{c,d=1}^{N_c} \kappa_{cdmnj} (u_{cdmnj}(\omega_j + \alpha^{(k)} z_j)), \quad (4.16)$$

with $\kappa_{cdmnj}(\cdot)$ and $u_{cdmnj}(\cdot)$ defined in (4.13).

Using one step of Newton’s method on the quadratic majorizer $q_k(\alpha)$ gives the step size update

$$\begin{aligned} \alpha^{(k+1)} &= \alpha^{(k)} - \frac{\frac{\partial}{\partial \alpha} q_k(\alpha^{(k)})}{\frac{\partial^2}{\partial \alpha^2} q_k(\alpha^{(k)})} \\ &= \alpha^{(k)} - \frac{\frac{\partial}{\partial \alpha} f(\alpha^{(k)})}{d^{(k)} + \beta \|\mathbf{Cz}\|_2^2}. \end{aligned} \quad (4.17)$$

We implement (4.17) efficiently by computing $\|\mathbf{Cz}\|_2^2$ only once per outer NCG iteration i . Since the majorizer satisfies $q_k(\alpha) \geq f(\alpha)$ for all step size α and inner line search iteration k , the update (4.17) guarantees monotonic decrease of the cost (4.15).

4.4 Results

We investigated our algorithm and its efficiency with two multi-echo field map estimation experiments and two water-fat imaging experiments. Due to the large data size, memory intensive methods with a direct solver using the full Hessian are excluded from our experiments. In particular, we compare our incomplete Cholesky preconditioner (NCG-MLS-IC) method versus a quadratic majorizer update with diagonal Hessian (QM) [65] and versus the NCG algorithm without any preconditioner (NCG-MLS) and with a diagonal preconditioner (NCG-MLS-D) [10]. In addition, we used the Poblano toolbox [5] to compare the convergence of the quasi-Newton (QN) and truncated Newton (TN) methods in our simulations.

For each dataset, we define a mask using the convex hull of all voxels that contribute to the signal (with coil-combined image magnitude thresholded below by $0.1y_{\max}$, where y_{\max} denotes the maximum image magnitude in the coil-combined image for the first echo time.), with a dilation of two voxels. We then computed ω within the mask, and tuned the regularization parameter β by sweeping across a range of val-

Algorithm: Preconditioned NCG-MLS

Inputs: $\mathbf{y}, \mathbf{s}, \mathbf{t}, \mathbf{C}, \beta$ **Initialization:** $\boldsymbol{\omega}^0$ by (4.8) or (4.9) $\mathbf{z}^0 = -\nabla\Psi(\boldsymbol{\omega}^0)$ $\alpha^{(0)} = 0$ precompute r_{cdmnj} by (4.4) and $t_m - t_n$ **for** $i = 0, 1, \dots, N - 1$ **do**compute gradient $\mathbf{g}^i = \nabla\Psi(\boldsymbol{\omega}^i)$ with (4.10)precondition $\mathbf{p}^i = -(\mathbf{P}^i)^{-1}\mathbf{g}^i$ with (4.14)compute μ^i with conjugacysearch direction $\mathbf{z}^{i+1} = \mathbf{p}^i + \mu^i\mathbf{z}^i \in \mathbb{R}^{N_v}$ **for** $k = 0, 1, \dots, N_i - 1$ **do**update step size $\alpha^{(k+1)}$ by (4.17)**end for**update $\boldsymbol{\omega}^{i+1} = \boldsymbol{\omega}^i + \alpha^{(N_i)}\mathbf{z}^{i+1}$ **end for****output:** $\boldsymbol{\omega}^N$

ues. All our experiments used MATLAB R2020a, with a 2.4-GHz dual-core Intel Core i7. The MATLAB code that reproduces the experiments with our efficient algorithm is available at <https://github.com/ClaireYLin/regularized-field-map-estimation>.

4.4.1 Brain simulation

We first simulated a 3D brain dataset with 40 64×64 slices, 4 simulated coils and 3 echo times $t_l = 0, 2, 10$ ms, with added complex Gaussian noise so that the SNR ≈ 20 dB. To generate multi-coil data, we simulated coil sensitivity maps with 4 coils based on [73] using the Michigan Image Reconstruction Toolbox (MIRT) [4]. We set $\beta = 2^{-4}$ with first order regularization to achieve visual resemblance to the ground truth field map. In light of the trade-off between storage and approximation error discussed in Section 4.3.2, we explored preconditioners using the incomplete Cholesky factorization both without tolerance (NCG-MLS-IC-0) and with a tolerance of $H_{\max}^i \times 10^{-3}$ for each iteration i (NCG-MLS-IC).

Fig. 4.4.1 shows four selected slices, their initial field map, and the regularized estimate by our algorithm. To examine the speed of convergence, we plot the root mean square error (RMSE) $\|\boldsymbol{\omega}^i - \boldsymbol{\omega}_{\text{true}}\|_2 / \sqrt{N_v}$ versus wall time in Fig. 4.4.2. The RMSE plots show that the QM and all the NCG-MLS methods converge to RMSE

≈ 5.6 Hz, though going through a slightly lower RMSE in the iterative process. Both the quasi-Newton and the truncated Newton methods converge to minimizers with higher RMSE, hence we omitted their comparison in the phantom experiment below. The plots show a significant computational gain of NCG-MLS preconditioned with the incomplete Cholesky factorization over all the other methods.

We also observe that using a nonzero tolerance in the incomplete Cholesky factorization gives a slightly faster convergence than not using one, hence we adopt that choice for the NCG-MLS-IC implementations in our next experiments.

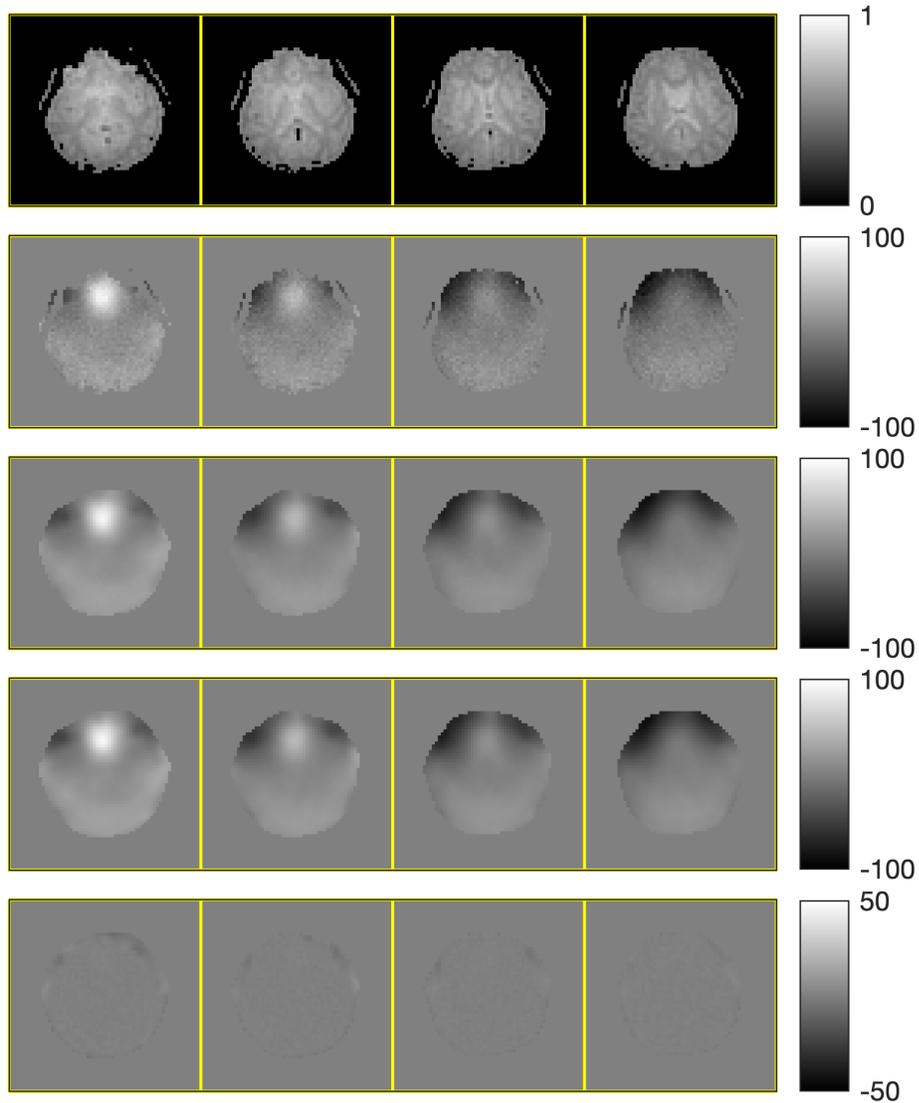


Figure 4.4.1: Selected slices of coil-combined simulation image, initial field map (in Hz), regularized field map estimate $\hat{\omega}$, ground truth field map ω_{true} , and error $|\hat{\omega} - \omega_{\text{true}}|$.

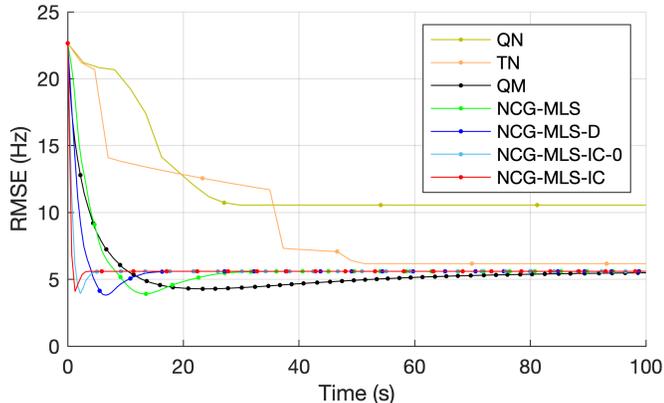


Figure 4.4.2: RMSE convergence of seven algorithms on simulation dataset. Every 10 iteration is marked by a dot.

4.4.2 Phantom dataset

Our second experiment uses a Function Biomedical Informatics Research Network (FBIRN) phantom [100] with two pieces of metal staple to induce field inhomogeneity, collected on a GE MR750 3T scanner with a 32-channel Nova Head Coil receiver. This dataset has size $74 \times 74 \times 10$ with 3 mm^3 isotropic voxel size, $\text{TR} = 10.5 \text{ ms}$, with 3 echo times $t_l = 0, 1, 2.3 \text{ ms}$. We computed coil sensitivity maps using ESPIRiT [190], and set $\beta = 2^{-3}$ with first-order finite difference regularization.

Fig. 4.4.3 shows four selected slices, their initial field map, and the regularized estimate by our algorithm. To compare convergence, we computed the root mean square difference (RMSD) $\|\omega^i - \omega^\infty\|_2 / \sqrt{N_v}$ to the converged ω^∞ of the QM method. The RMSD plots in Fig. 4.4.4 show that our algorithm converges much faster than the other three, reaching 0.33 Hz RMSD in 1 iteration, and 0.005 Hz RMSD in 2 iterations. Since this 3D dataset has a more realistic problem size than the simulated data, we quantify the convergence speedup by comparing the time it takes for each method to reach an RMSD below 0.5 Hz. Table 4.3 shows that our NCG-MLS algorithm with an incomplete Cholesky preconditioner provides a speedup of 15 times from NCG-MLS with a diagonal preconditioner, 18 times from that without a preconditioner, and 21 times from the quadratic majorizer implementation.

	QM	NCG-MLS	NCG-MLS-D	NCG-MLS-IC
Time (s)	96	81	69	4.5
vs. IC time	21×	18×	15×	

Table 4.3: Time for each algorithm to reach an RMSD below 0.5 Hz, and their relative proportions to the time taken by NCG-MLS-IC.

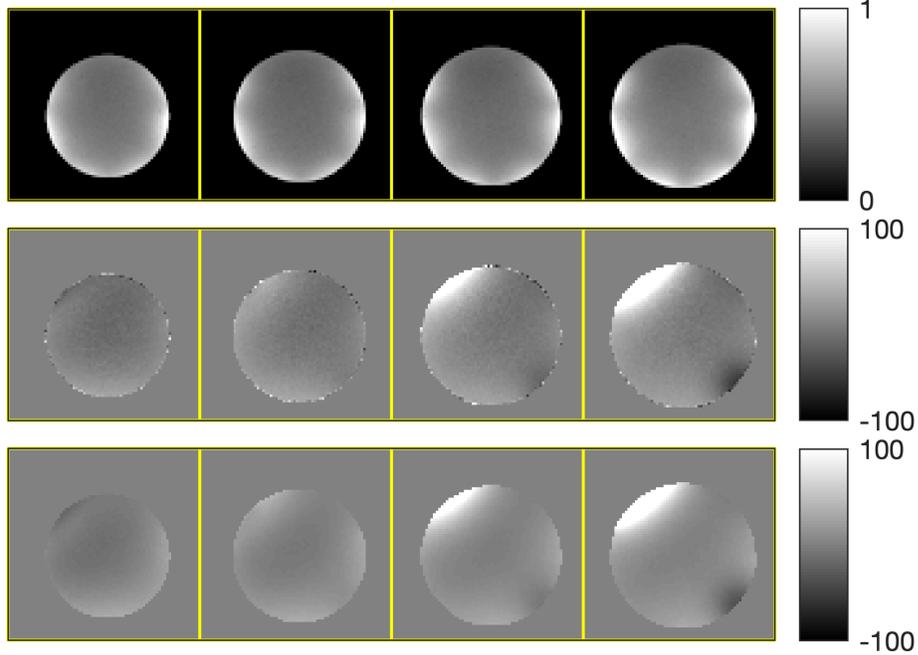


Figure 4.4.3: Selected slices of coil-combined phantom image, initial field map (in Hz), and regularized field map estimate.

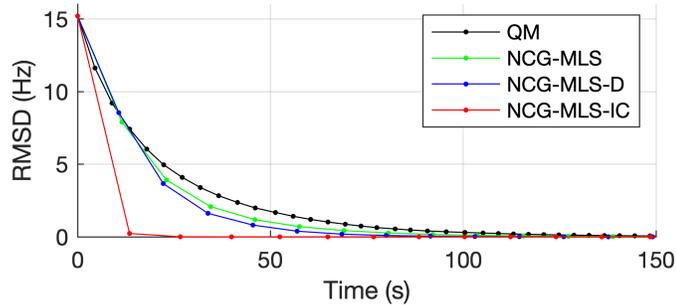


Figure 4.4.4: RMSD convergence of four algorithms on phantom data. Every iteration is marked by a dot.

4.4.3 Cardiac water-fat simulation

For water-fat imaging, we first performed a cardiac simulation based on one of the 8-echo datasets used in the ISMRM Fat-Water Toolbox [2]. Since implementations in the toolbox work only for 2D datasets, and coil combination such as [195] is often used in practice, we illustrate the flexibility of our algorithm in a 2D coil-combined case by simply setting the number of coils $N_c = 1$ and the coil sensitivity map $\mathbf{s} = \mathbf{1}$. We also consider the multippeak model in water-fat imaging.

This dataset has size 256×192 with 8 echo times from 1.5 to 17.4 ms (each 2.3 ms

apart). We generated ground truth field map and water and fat images using golden section search with multiresolution [123]. We used the same values $\{\alpha_p\}$ and $\{\Delta_{f,p}\}$ as in the toolbox implementations both for simulating images with 8 echo times using the model (4.1) and for estimation. For comparison, we also ran the graph cut (GC) method [82] using the same cost (4.6) with second-order finite differences as in [82], and $\beta = 2^{-7}$ as the regularization parameter.

Fig. 4.4.6 shows the first echo image, the initial field map $\tilde{\omega}^0$ by voxel-wise estimation, and the initial ω^0 after 10 CG iterations of PWLS minimization (4.9). Fig. 4.4.7 shows the ground truth field map, water and fat images, and the estimates and error images by the graph cut and by our algorithm. Compared with graph cut, our algorithm achieves slightly lower NRMSE on the water image (20.09% vs. 23.57%) and the fat image (20.93% vs. 23.43%), with lower final RMSE on the field map, shown in Fig. 4.4.5. To explore a combination suggested by [82], we ran 100 graph cut iterations followed by 100 optimal transfer iterations using a quadratic majorizer [88]. We used the implementation in the toolbox [2] which did not precompute r_{cdmnj} by (4.4). Fig. 4.4.5 shows the graph cut RMSE curve jumps up (to 615 Hz) on its first iteration, and converges to its own minimizer. The subsequent quadratic majorizer update lowers the RMSE further, which opens a promising future direction of combining graph cut with the faster NCG-MLS-IC with precomputation of common terms. Fig. 4.4.5 also shows the truncated Newton and quasi-Newton methods again converge to their minimizers with higher RMSE. We omitted all methods with higher RMSE in the real data experiment below.

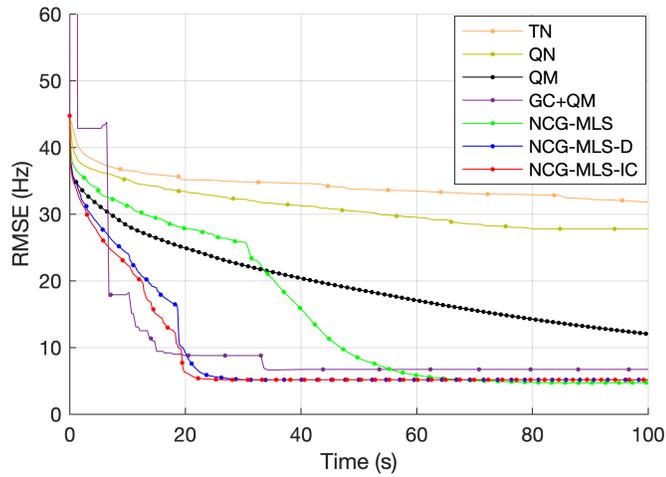


Figure 4.4.5: RMSE convergence of seven algorithms on water-fat simulation. Every 20 iterations is marked by a dot.

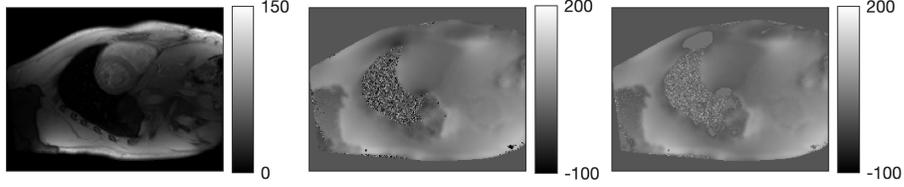


Figure 4.4.6: Simulated image for the 1st echo, initial field map $\tilde{\omega}^0$ (in Hz) by voxel-wise estimation, and initial fieldmap ω^0 by PWLS (4.9).

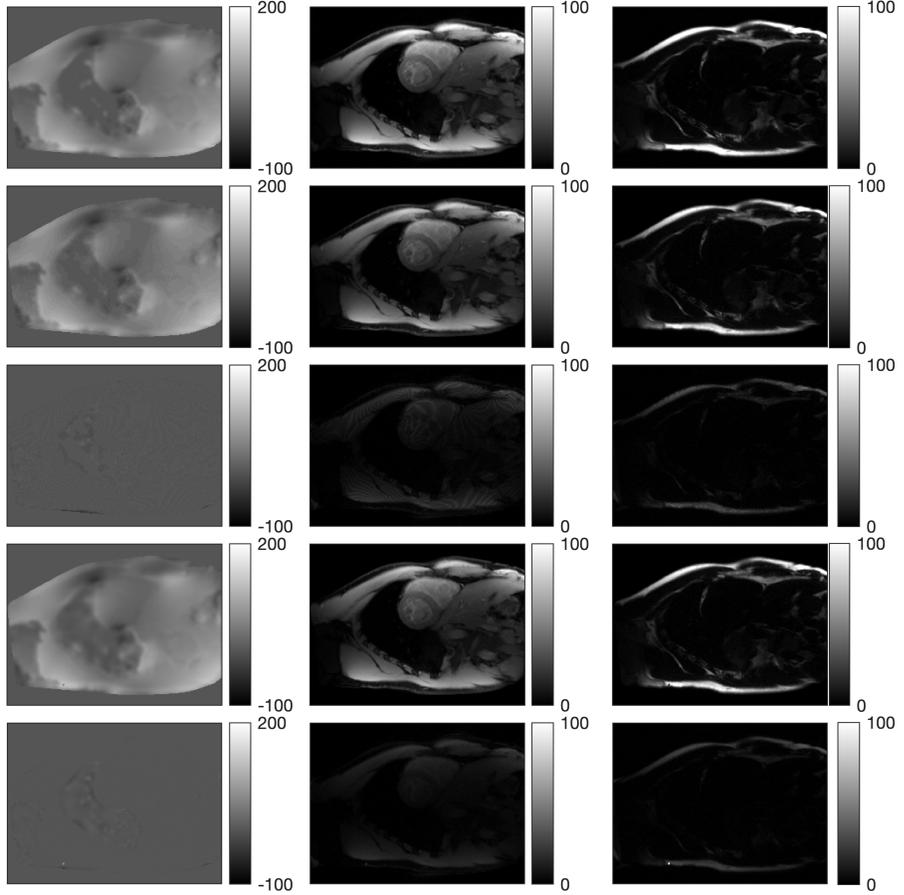


Figure 4.4.7: Ground truth and estimated field map, water, and fat images on simulated data. 1st row: ground truth field map, water, and fat images. 2nd and 3rd rows: graph cut estimates and their error images. 4th and 5th rows: NCG-MLS-IC estimates and their error images.

4.4.4 Ankle water-fat dataset

We further illustrate our algorithm in the 3D multi-coil setting using an ankle dataset from the ISMRM Fat-Water Separation Dataset [2]. This dataset has 4 256×256 slices, 8 coils and 3 echo times $t_l = 2.2, 3, 3.8$ ms, in a 3T scanner that corresponds to a single $\Delta_f \approx 440$ Hz. We chose $\beta = 2^{-10}$ with first-order finite

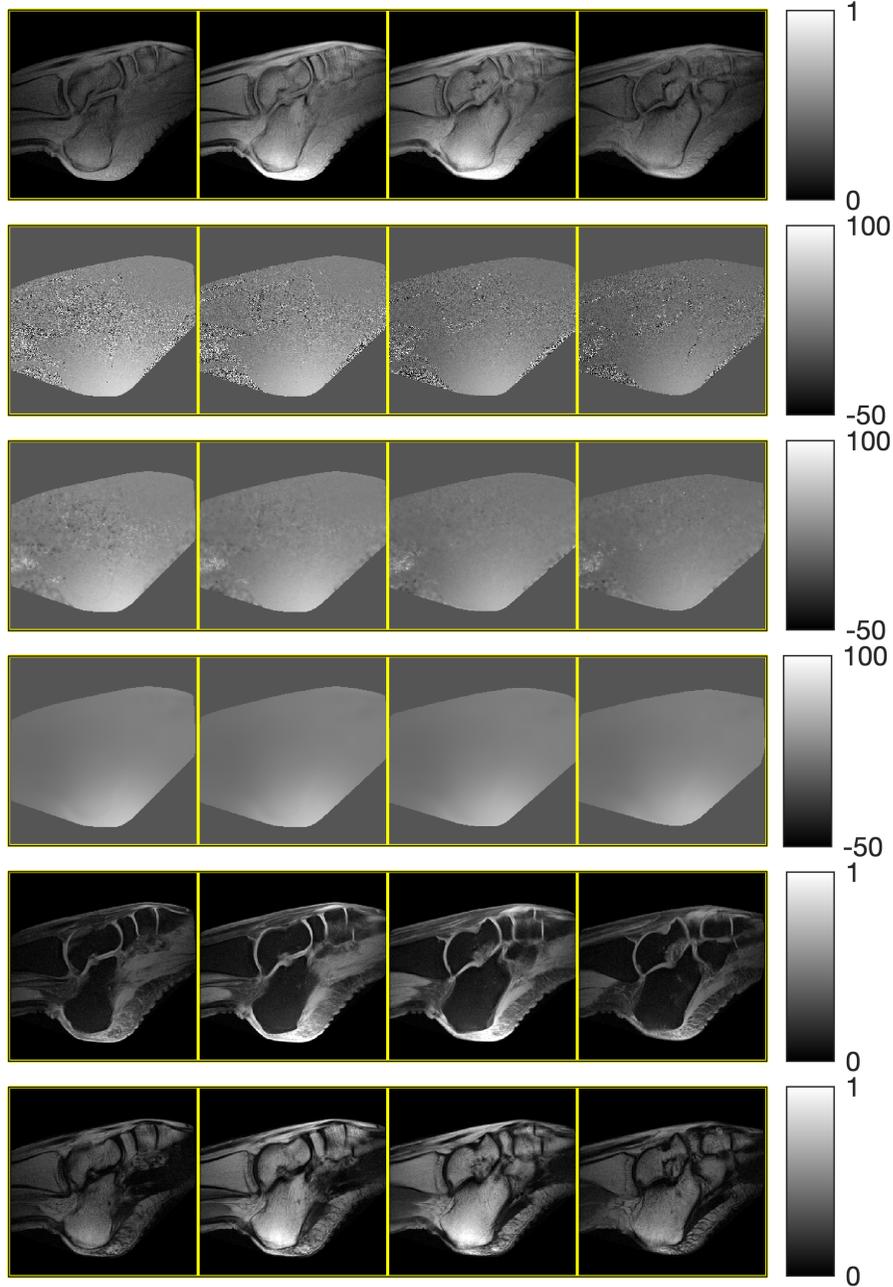


Figure 4.4.8: Image, initial and regularized field maps, and estimated water and fat images on ankle dataset. Top to bottom: coil-combined water-fat image for the 1st echo, initial field map $\tilde{\omega}^0$ (in Hz) by voxel-wise estimation, initial fieldmap ω^0 by PWLS (4.9), regularized field map estimate, estimated water image, and estimated fat image.

difference regularization to achieve visual separation of water and fat components.

Fig. 4.4.8 shows the first echo image, the initial field map $\tilde{\omega}^0$ by voxel-wise estimation, the initial ω^0 after 10 CG iterations of PWLS minimization (4.9), and the

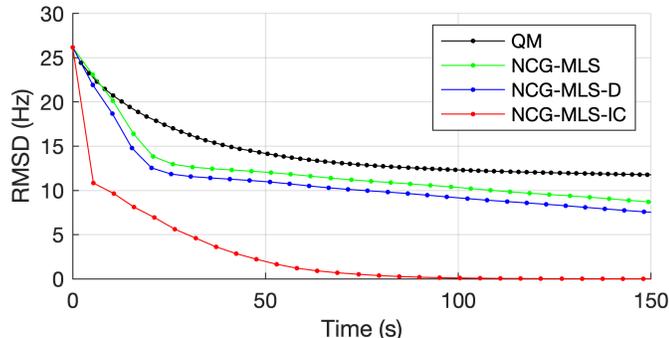


Figure 4.4.9: RMSD convergence of four algorithms on ankle dataset. Every iteration is marked by a dot.

regularized estimate by our algorithm. For completeness, Fig. 4.4.8 also shows the estimated water and fat images using (4.7), which achieve a visual separation of the two components. However, it is worth emphasizing that our main interest is in the speed of finding a minimizer of the problem (4.6). In this case, since QM converged to a different local minimum than the other three methods, we computed the RMSD to ω^∞ of the NCG-MLS method (without preconditioner). The RMSD plots in Fig. 4.4.9 show a significant computational gain of our algorithm over the other algorithms.

4.5 Conclusion

This chapter presents an efficient algorithm for both multi-echo field map estimation and water-fat imaging problem in the 3D multi-coil MRI setting. Given the nonconvex cost function, our algorithm uses the nonlinear conjugate gradient method with a preconditioner based on an incomplete Cholesky factorization, and a monotonic step size line search based on a quadratic majorizer with optimal curvatures. This is the first work to use the incomplete Cholesky factorization as a preconditioner for multi-coil field map estimation. Experiments with simulation and real data show that our method has faster convergence than existing memory-efficient methods.

This method is suitable for routine use; the main practical issue is choosing the regularization parameter in (4.6); the spatial resolution analysis in [65] is helpful as a guide for that in the single-coil case, and that work should be easily extended to the multi-coil case and the fat-water separation problem.

Chapter 5

A Spatiotemporal Model for Task-based Functional MRI

As a category of fMRI, the main goal of task-based fMRI is to identify brain regions that are functionally involved in performing specific tasks. To better identify task-activated brain regions in task-based functional magnetic resonance imaging (fMRI), various space-time models have been used to reconstruct image sequences from undersampled k -space data. These models represent an fMRI timeseries as a sum of two components, aiming to separate task-correlated signal from non-task background. This chapter proposes a model based on assumptions of the activation waveform shape, with spatial and temporal smoothness regularization on the magnitude and phase of the timeseries. We compare the proposed approach to two contemporary task fMRI decomposition models, and show that the proposed model yields better timeseries and activation maps on simulated and human subject fMRI datasets with multiple tasks.

5.1 Introduction

FMRI uses blood oxygenation level dependent (BOLD) contrast to measure dynamic neural activity [142]. As rapid whole brain imaging is desired to reduce noise and motion artifacts while acquiring more data, various methods have been proposed for fMRI acceleration [36,80,175,198]. In the case of undersampled acquisition, a significant challenge is to reconstruct functional MR images with both high spatial and temporal resolution. This has led to reconstruction models based on assumptions such as low-rankness, sparsity, and temporal smoothness [9,34,38,39,95,113,149,154,180].

This chapter extends our image recovery work in the conference paper [117] to image reconstruction from k -space data.

Due to the low strength of the signal of interest in fMRI, the reconstruction results can depend heavily on model assumptions, and it is desirable to develop a goal-driven model that considers the purpose of the fMRI experiment, such as identifying activated regions.

This chapter considers task-based fMRI, with the goal to detect task-activated regions of the brain. Conventional image analysis for conventionally reconstructed fMRI timeseries involves processing a reconstructed image sequence to correct for noise and artifacts, then comparing it with a predefined hemodynamic response function (HRF)-convolved task waveform using general linear modeling (GLM) [130]. Each voxel’s temporal correlation to the task waveform is a test statistic that indicates its activation by the task, and one can visualize the activated areas in the brain by, for example, plotting a thresholded correlation map for all the voxels.

Since MR images reconstructed from k -space are complex-valued, analysis approaches on complex images have also been proposed. For example, [11,170] investigate functional activation separately on the magnitude and phase images, whereas [29,106] combine magnitude and phase information of the complex independent components extracted from the complex image sequence. In the reconstruction setting, separate magnitude and phase regularization has been developed for static MRI, with quadratic spatial smoothness penalty on the phase [60,122,209,211].

With the aim of integrating task detection into the reconstruction process itself, we propose a reconstruction model based on *a priori* knowledge of the task waveform shape, with spatial and temporal smoothness assumptions on the magnitude and phase images. We compare it to two contemporary task fMRI decomposition models, and show that our model reconstructs better timeseries and activation maps on undersampled fMRI datasets with multiple simulated and real tasks. Compared to our conference work [117], this chapter considers the multi-coil image reconstruction problem with spatiotemporal undersampled k -space data, and uses an additional regularizer for the phase images. To the best of our knowledge, this is the first fMRI reconstruction model with separate magnitude and phase regularization.

The rest of this chapter is organized as follows. Section 5.2 formulates a general optimization framework for fMRI image reconstruction, and reviews two related task fMRI models using special cases of the formulation. Section 5.3 presents the proposed model and its optimization scheme. Section 5.4 reports simulated and human task experiment results, followed by conclusions in Section 5.5.

5.2 Problem and related models

To reconstruct task-based fMRI images from undersampled multi-coil k -space data, several space-time models have been proposed that all involve an optimization problem with a cost function of the form

$$\operatorname{argmin}_X \frac{1}{2} \|\mathbf{E}X - \mathbf{d}\|_2^2 + \lambda R(X), \quad (5.1)$$

where $X \in \mathbb{C}^{N_v \times N_t}$ is the desired image sequence with N_v voxels and N_t time frames, $\mathbf{E} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_k N_c}$ is an encoding operator with N_c signal receiver coils (see Section 2.2.1), $\mathbf{d} \in \mathbb{C}^{N_k N_c}$ is the undersampled multi-coil k -space data, and $R(\cdot)$ is a regularizer with parameter λ that depends on specific model assumptions.

We now review two existing task fMRI reconstruction models that are essentially special cases of framework (5.1), with underlying assumptions that inspire the proposed model in Section 5.3.

5.2.1 L+S: low-rank plus sparse decomposition

The low-rank plus sparse (L+S) approach [149,180] decomposes an fMRI image sequence into $X = L + S$, where the low-rank component L aims to model the non-task temporally correlated background, and the sparse component S aims to model the pseudo-periodic BOLD activity. With these assumptions, the L+S approach replaces (5.1) with the following optimization problem

$$\operatorname{argmin}_{L,S} \frac{1}{2} \|\mathbf{E}(L + S) - \mathbf{d}\|_2^2 + \lambda_L \|L\|_* + \lambda_S \|\mathbf{T}S\|_1, \quad (5.2)$$

where $\mathbf{T} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_v \times N_t}$ is the (unitary) temporal Fourier transform operator. The nuclear norm encourages L to be low-rank, and the l_1 norm encourages S to be periodic.

One way to solve (5.2) is to alternate between two proximal gradient updates, with singular value thresholding (SVT) for L and soft thresholding for S [149,180]. In our experiments, we used POGM to accelerate convergence (see Chapter 3 and [115]).

5.2.2 L+UV: low-rank plus task-based decomposition

Another task fMRI model [37] also considers a low-rank component in hopes of capturing the background signal, inspired by a resting state model [39]. Instead of using temporal Fourier sparsity as a regularizer, this model considers a spatiotemporal

decomposition UV , where $U \in \mathbb{C}^{N_v \times r}$ is an estimated spatial map corresponding to a predefined temporal basis $V \in \mathbb{R}^{r \times N_t}$, where r is the number of task-associated basis vectors. This real-valued V can be a task waveform, or more generally a sinusoidal or block-like function, along with its temporal derivatives [37]. Compared with the L+S model, the temporal basis V constrains the shape of activation signal to lie in a given r -dimensional subspace, largely paralleling what is done with GLM approaches that are used for postprocessing. Together with the low-rank component L , the L+UV approach replaces (5.1) with the following optimization problem

$$\operatorname{argmin}_{L,U} \frac{1}{2} \|\mathbf{E}(L + UV) - \mathbf{d}\|_2^2 + \lambda_L \|L\|_*, \quad (5.3)$$

which has one fewer regularization parameter to tune than (5.2).

In our experiments, we solved (5.3) by alternating minimization as in [37], with POGM to accelerate the L update, and a closed-form linear least squares update for U . One can also apply an accelerated proximal gradient method to simultaneously update L and U , by a relatively simple extension of the methods in Chapter 3 and [115].

5.3 Proposed model

A desirable decomposition in task-based fMRI is to separate a timeseries into its non-task background signal and task-activated foreground dynamics. However, both models (5.2) and (5.3) contain overlapping assumptions in their components. For example, a periodic task can be both low rank and temporally Fourier sparse, hence might appear in either L or S in the L+S model. Similarly in L+UV, a block-like task that is supposed to be capture in UV (by design of V) can also appear in L due to its low-rankness.

Another desired property is inspired by fMRI analysis. As most methods use the magnitude images and decompose them for activation analysis, we hope to account for the magnitude operation in reconstruction. While one could analyze the magnitude of the reconstructed sum, or the real and imaginary parts of each component in the sum, there lacks an intuitive way to visualize the timeseries and activation of the reconstructed complex non-task and task components.

To this end, we propose a decomposition model that strives to better capture separate features of the non-task background and the task-activated foreground, using a sum of two real-valued components. This section discusses its formulation under framework (5.1) and an optimization scheme.

Notations: We use \odot as the Hadamard product and \otimes as the Kronecker product. $\text{vec}(\cdot)$ denotes vectorization, \cdot^* denotes complex conjugate, and $\Re(\cdot)$ takes the real part of a complex input. I_N is the identity matrix with dimension $N \times N$.

5.3.1 B+UV: background plus spatiotemporal decomposition

We propose to model the complex image sequence as $X = (B + UV) \odot e^{iP}$, where $B, UV, P \in \mathbb{R}^{N_v \times N_t}$ respectively capture the magnitude of background signal, magnitude of task signal, and phase. Under these assumptions, we consider a predefined temporal basis $V \in \mathbb{R}^{r \times N_t}$ with task waveforms, similar to the L+UV model [37]. We also use an edge-preserving spatial smoothness regularizer on the spatial coefficients in $U \in \mathbb{R}^{N_v \times r}$. Inspired by spatial and temporal similarity assumptions in resting state fMRI reconstruction models [9,34,38,113], we consider spatial and temporal smoothness penalty on the background BOLD magnitude signal B and the phase P . To accommodate phase wrapping, we apply spatial and temporal finite difference on the exponentiated phase e^{iP} as opposed to P itself [122,209,211]. Any temporally smooth terms that would normally be included in a GLM analysis of an fMRI timeseries, such as a linear drift term and/or discrete cosine transform (DCT) low frequency components, should appear in the background component B , whereas the task signals are represented in V and hence appear as coefficients in U . In cases where the hemodynamic response function is uncertain or subject to delays, one can also include temporal derivatives of the HRF-convolved task waveforms in V [37]. Our proposed cost function for task-based fMRI is

$$\begin{aligned} \underset{U}{\text{argmin}} \min_{B,P} \frac{1}{2} & \|\mathbf{E}((B + UV) \odot e^{iP}) - \mathbf{d}\|_2^2 \\ & + \frac{1}{2} \|\mathbf{D}_B B\|_2^2 + \frac{1}{2} \|\mathbf{D}_P e^{iP}\|_2^2 + \Psi(\mathbf{D}_U U), \end{aligned} \quad (5.4)$$

where

$$\mathbf{D}_B = \begin{bmatrix} \sqrt{\lambda_B^v} (D_t \otimes I_{N_v}) \\ \sqrt{\lambda_B^t} (I_{N_t} \otimes D_v) \end{bmatrix} : \mathbb{R}^{N_v \times N_t} \rightarrow \mathbb{R}^{2N_v N_t}$$

and

$$\mathbf{D}_P = \begin{bmatrix} \sqrt{\lambda_P^v} (I_{N_t} \otimes D_v) \\ \sqrt{\lambda_P^t} (D_t \otimes I_{N_v}) \end{bmatrix} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{2N_v N_t}$$

include first-order finite matrices $D_t \in \mathbb{R}^{N_t \times N_t}$ and $D_v \in \mathbb{R}^{N_v \times N_v}$ in temporal and spatial dimensions. For U we consider (columnwise) spatial finite difference

$$\mathbf{D}_U = (I_r \otimes D_v) : \mathbb{R}^{N_v \times r} \rightarrow \mathbb{R}^{N_v r},$$

with $\Psi(\cdot)$ as the sum of an edge-preserving potential function on its elements to account for spatial task-activated regions, weighted by a regularization parameter λ_U :

$$\Psi(\mathcal{U}) = \lambda_U \sum_i \psi(\mathcal{U}_i).$$

There are many typical edge-preserving potential functions, e.g., Cauchy [83], Huber [86], hyperbola [33], etc.. In this chapter we use the Fair potential function [53] as an illustration:

$$\psi_\delta(u) = \delta^2 (|u/\delta| - \log(1 + |u/\delta|)).$$

5.3.2 Optimization

To minimize the nonconvex cost (5.4), we use the nonlinear conjugate gradient (NCG) method with line search (LS) alternately on $\tilde{X} = [B \ U]$ and P . Our implementation for both variables is tabulated in the **Algorithm** below. For NCG, we use the Fletcher–Reeves update [63] to compute a μ^i that satisfies the conjugacy condition. Our algorithm decreases the cost (5.4) monotonically by design and converges to a local minimizer. Below we show the initialization scheme, and the NCG algorithm with line search for \tilde{X} and P updates.

Initialization

A good initialization is important for the nonconvex cost (5.4). In cases with spatiotemporal undersampling, we use data sharing among nearby frames to fill in missing k -space samples, for initialization only. This temporal interpolation technique is a generalization of the keyhole method [93,194]. For an initial guess of the overall image sequence X_0 , we apply a conjugate phase (CP) reconstruction [139] with sampling density compensation computed by an iterative method [156] to the data-shared k -space data. Due to nonconvexity, we provide the option to further estimate X_0 by running a few CG iterations of the (convex) B+UV without magnitude and phase separation [117], where the variables B, U become complex-valued:

$$\operatorname{argmin}_U \min_B \frac{1}{2} \|\mathbf{E}(B + UV) - \mathbf{d}\|_2^2 + \frac{1}{2} \|\mathbf{D}_B B\|_2^2 + \Psi(\mathbf{D}_U U). \quad (5.5)$$

To initialize U , we use the least squares solution

$$U_0 = M_0 V^\top (V V^\top)^{-1}$$

where M_0 is the magnitude of the demeaned X_0 . The background is initialized as

$$B_0 = |X_0| - U_0 V,$$

and the initial phase is

$$P_0 = \angle X_0.$$

The solution to the optimization problem (5.4) has ambiguity due to its nonuniqueness, as a sign flip of B, U and a π shift of P would result in the same cost. Using the above initialization scheme in our experiments, the reconstructed \hat{B} and \hat{U} were nonnegative and did not suffer from this effect; if needed, one could consider enforcing nonnegativity constraints on B or $B + UV$ to avoid potential ambiguity.

Update of image sequences $\tilde{X} = [B \ U]$

To simplify the B, U updates, we can express

$$\mathbf{E}_P B := \mathbf{E} \cdot \text{diag}(e^{iP}) B = \mathbf{E}(B \odot e^{iP}),$$

and

$$\begin{aligned} \mathbf{E}_{P,V} U &:= \mathbf{E} \cdot \text{diag}(e^{iP})(V^\top \otimes I_r) \text{vec}(U) \\ &= \mathbf{E}((UV) \odot e^{iP}), \end{aligned}$$

where we use the compatibility of vectorization with the Kronecker product

$$\text{vec}(UV) = (V^\top \otimes I_{N_v}) \text{vec}(U).$$

Then for a fixed P , the cost with respect to $\tilde{X} = [B \ U]$ can be written as

$$\Phi_P(\tilde{X}) = \frac{1}{2} \|\tilde{\mathbf{E}}_P \tilde{X} - \mathbf{d}\|_2^2 + \tilde{\Psi}(\tilde{\mathbf{D}} \tilde{X}),$$

where

$$\tilde{\mathbf{E}}_P = \begin{bmatrix} \mathbf{E}_P & \mathbf{E}_{P,V} \end{bmatrix},$$

$$\tilde{\mathbf{D}} = \begin{bmatrix} \mathbf{D}_B & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_U \end{bmatrix},$$

and $\tilde{\Psi}(\cdot)$ consists of the l_2 norm on the B term and $\Psi(\cdot)$ on the U term. Its gradient is

$$\begin{aligned} \nabla \Phi_P(\tilde{X}) &= \Re\{\tilde{\mathbf{E}}_P^*(\tilde{\mathbf{E}}_P \tilde{X} - \mathbf{d})\} \\ &+ \begin{bmatrix} \mathbf{D}_B^* \mathbf{D}_B B & \mathbf{D}_U^* \nabla \Psi(\mathbf{D}_U U) \end{bmatrix} \end{aligned} \quad (5.6)$$

where, for the Fair potential,

$$[\nabla \Psi_\delta(\mathcal{U})]_i = \frac{\mathcal{U}_i}{1 + |\mathcal{U}_i/\delta|}.$$

With a search direction given by NCG, the choice of step size is important for convergence of the algorithm. We implement a recursive line search algorithm using a quadratic majorizer with an optimal curvature, which guarantees monotone decrease of the cost function [58].

In the line search step, given a current estimate $\tilde{X}^i = [B^i \ U^i]$ and a search direction $Z^i = [Z_B^i \ Z_U^i] \in \mathbb{R}^{N_v \times (N_t + r)}$, we aim to find a step size

$$\underset{\alpha}{\operatorname{argmin}} f_P(\alpha), \text{ where}$$

$$\begin{aligned} f_P(\alpha) &= \frac{1}{2} \|\tilde{\mathbf{E}}_P(\tilde{X}^i + \alpha Z^i) - \mathbf{d}\|_2^2 + \tilde{\Psi}(\tilde{\mathbf{D}}(\tilde{X}^i + \alpha Z^i)) \\ &= \frac{1}{2} \|\mathbf{E}_P((B^i + \alpha Z_B^i) + (U^i + \alpha Z_U^i)V) - \mathbf{d}\|_2^2 \\ &\quad + \frac{1}{2} \|\mathbf{D}_B(B^i + \alpha Z_B^i)\|_2^2 + \Psi(\mathbf{D}_U(U^i + \alpha Z_U^i)). \end{aligned}$$

We iteratively minimize $f_P(\alpha)$ using a quadratic majorizer based on Huber's method [87, p. 184] at the k th inner iteration (dropping outer iteration i for brevity):

$$\begin{aligned}
q_k(\alpha) &= \frac{1}{2} \|\tilde{\mathbf{E}}_P(\tilde{X} + \alpha Z) - \mathbf{d}\|_2^2 + \frac{1}{2} \|\mathbf{D}_B(B + \alpha Z_B)\|_2^2 \\
&\quad + \Psi(\mathbf{D}_U(U + \alpha^{(k)} Z_U)) \\
&\quad + \text{vec}(Z_U)^\top \mathbf{D}_U^* \nabla \Psi(\mathbf{D}_U(U + \alpha^{(k)} Z_U)) (\alpha - \alpha^{(k)}) \\
&\quad + \frac{1}{2} d^{(k)} (\alpha - \alpha^{(k)})^2,
\end{aligned}$$

where the optimal curvature is given by

$$\begin{aligned}
d^{(k)} &= \text{vec}(Z_U)^\top \mathbf{D}_U^* \nabla^2 \Psi(\mathbf{D}_U(U + \alpha^{(k)} Z_U)) \mathbf{D}_U Z_U \\
&= \text{vec}(Z_U)^\top \mathbf{D}_U^* \text{diag}\left(\boldsymbol{\omega}_\Psi(\mathbf{D}_U(U + \alpha^{(k)} Z_U))\right) \mathbf{D}_U Z_U,
\end{aligned}$$

where

$$[\boldsymbol{\omega}_{\Psi_\delta}(\mathcal{U})]_i = \frac{1}{1 + |\mathcal{U}_i/\delta|}.$$

Using one step of Newton's method on the quadratic majorizer $q_k(\alpha)$ gives the step size update

$$\begin{aligned}
\alpha^{(k+1)} &= \alpha^{(k)} - \frac{\frac{\partial}{\partial \alpha} q_k(\alpha^{(k)})}{\frac{\partial^2}{\partial \alpha^2} q_k(\alpha^{(k)})} \\
&= \alpha^{(k)} - \frac{\frac{\partial}{\partial \alpha} f_P(\alpha^{(k)})}{d^{(k)} + \|\tilde{\mathbf{E}}_P Z\|_2^2 + \|\mathbf{D}_B Z_B\|_2^2}. \tag{5.7}
\end{aligned}$$

We implement (5.7) efficiently by computing $\|\tilde{\mathbf{E}}_P Z\|_2^2$ and $\|\mathbf{D}_B Z_B\|_2^2$ only once per outer NCG iteration i . Since the majorizer satisfies $q_k(\alpha) \geq f(\alpha)$ for all step size α and inner line search iteration k , the update (5.7) guarantees monotonic decrease of $f_P(\alpha)$.

Update of phase sequence P

For a fixed \tilde{X} , the cost with respect to the phase sequence P is

$$\Phi_X(P) = \frac{1}{2} \|\mathbf{E}_X e^{iP} - \mathbf{d}\|_2^2 + \frac{1}{2} \|\mathbf{D}_P e^{iP}\|_2^2,$$

where $\mathbf{E}_X = \mathbf{E} \cdot \text{diag}(B + UV)$. Its gradient is

$$\begin{aligned} \nabla \Phi_X(P) &= -\Re\{\text{diag}(ie^{-iP})\mathbf{E}_X^*(\mathbf{E}_X P - \mathbf{d})\} \\ &\quad - \Re\{ie^{-iP} \odot (\mathbf{D}_P^* \mathbf{D}_P e^{iP})\}. \end{aligned} \quad (5.8)$$

Similar to the line search iterates for \tilde{X} , we use the Newton's method to update the step size. In this case, due to nonconvexity of the exponential term, we use a backtracking strategy [138] to ensure monotone decrease of the cost. Given a current estimate P^i and a search direction $Z_P^i \in \mathbb{R}^{N_v \times N_t}$, we aim to find a step size

$$\underset{\alpha}{\text{argmin}} f_X(\alpha), \text{ where}$$

$$f_X(\alpha) = \Phi_X(P + \alpha Z_P^i).$$

The Newton step size update on $f_X(\alpha)$ is

$$\alpha^{(k+1)} = \alpha^{(k)} - \frac{\frac{\partial}{\partial \alpha} f_X(\alpha^{(k)})}{\frac{\partial^2}{\partial \alpha^2} f_X(\alpha^{(k)})}, \quad (5.9)$$

where

$$\frac{\partial}{\partial \alpha} f_X(\alpha^{(k)}) = \text{vec}(Z_P^i)^\top \nabla \Phi_X(P + \alpha Z_P^i),$$

and

$$\frac{\partial^2}{\partial \alpha^2} f_X(\alpha^{(k)}) = \text{vec}(Z_P^i)^\top \nabla^2 \Phi_X(P + \alpha Z_P^i) Z_P^i$$

where

$$\begin{aligned} \nabla^2 \Phi_X(P) &= -\Re\{\text{diag}(e^{-iP} \odot \mathbf{E}_X^*(\mathbf{E}_X P - \mathbf{d})) \\ &\quad - \text{diag}(e^{-iP})\mathbf{E}_X^* \mathbf{E}_X \text{diag}(e^{iP})\} \\ &\quad + \Re\{\text{diag}(e^{-iP})\mathbf{D}_P^* \mathbf{D}_P \text{diag}(e^{iP}) \\ &\quad - \text{diag}(e^{-iP} \odot (\mathbf{D}_P^* \mathbf{D}_P e^{iP}))\}. \end{aligned}$$

Normalization of V

Before our optimization procedure, we normalize the rows of V ; this leads to similar magnitudes across columns of U , and justifies our use of a single regularization parameter λ_U for the regularizer on U . This scaling also ensures the diagonal elements of $\tilde{\mathbf{E}}_P^* \tilde{\mathbf{E}}_P$ have similar magnitudes, so the same μ^i can be used for both B and U

Algorithm: NCG-LS for $\mathcal{Y} = \tilde{X}$ or $\mathcal{Y} = P$

Initialization: $\mathcal{Y}^0 =$ previous iterate (or CP reconstruction for 1st iteration) $Z^0 = -\nabla\Phi(\mathcal{Y}^0)$ by (5.6) or (5.8) $\alpha^{(0)} = 0$ **for** $i = 0, 1, \dots, N - 1$ **do** compute descent direction $\mathbf{g}^i = -\nabla\Phi(\mathcal{Y}^i)$ by (5.6) or (5.8) compute μ^i with conjugacy search direction $Z^{i+1} = \mathbf{g}^i + \mu^i Z^i$ **for** $k = 0, 1, \dots, N_i - 1$ **do** update step size $\alpha^{(k+1)}$ by (5.7) or (5.9) **end for** update $\mathcal{Y}^{i+1} = \mathcal{Y}^i + \alpha^{(N_i)} Z^{i+1}$ **end for****output:** \mathcal{Y}^N

without preconditioning.

5.4 Results

We validate our B+UV model using two fMRI experiments. The first experiment reconstructs images using a resting state fMRI dataset with simulated task waveforms imposed on specified regions of interest. The second experiment reconstructs images using an *in vivo* visual and motor fMRI dataset.

To compare how the signal changes within a voxel, we plot the timeseries of each component for randomly selected activated and non-activated voxels. To evaluate task detection accuracy, we use a blockwise permutation method [7] on the k -space data to compute an activation map for each algorithm. During each permutation, we first perform a circular shift at a random time frame of the k - t data sequence, then randomly permute it by blocks of size 50 s to account for fMRI autocorrelation. After reconstruction, we use the distribution of correlation between the task waveform and permuted timeseries to compute a p -value with Bonferroni correction for each voxel. We repeat this permutation process for $N_{\text{perm}} = 5000$ times, and define the final activation maps using p -value maps with a threshold at $10/N_{\text{perm}} = 0.002$.

Since the components in L+S and L+UV are complex-valued, for visualization purpose we show the magnitude of sums and the real part of individual components in the timeseries plots; we omit plots of the imaginary parts, as they show similar behaviors to the real parts. We compare their results with the (real-valued) B and

UV in our B+UV model.

5.4.1 Reconstruction: simulated tasks

Simulation setup

Our first experiment uses unprocessed resting state images from a Human Connectome Project (HCP) 1200 Subjects Dataset [1] with added simulated tasks. The advantage of using a human dataset is its resting state behavior with motion and physiological artifacts, making it a realistic background signal useful for validation of the non-task component in our model. The selected 3T resting state dataset has image size 60×72 and 5 z -slices, with a TR of 720 ms. To match the 200 s experiment for the human subject task (Section 5.4.2) with first 10 s of data discarded, we selected 265 time frames for reconstruction, and generated k -space data using a stack of spiral trajectory ($N_k = 3393$ for each slice) with retrospective k_z - t undersampling by a factor of 3 (Fig. 5.4.1).

We then simulated two tasks with different periodicities and generated their spiral k -space corresponding to three 3D trapezoidal regions of interest (ROI). Fig. 5.4.2 shows the task waveforms as rows of V ; one task is periodic with a block design, and the other is nonperiodic. To match the tasks in the real task experiment (Section 5.4.2), we simulated an interleaving left-right visual task using Task 1 and its negative waveform (Task 2) to generate two ROIs, and Task 3 to simulate a motor ROI. We added k -space task components to the synthetic resting state data to get our simulated undersampled spiral k -space data for reconstruction.

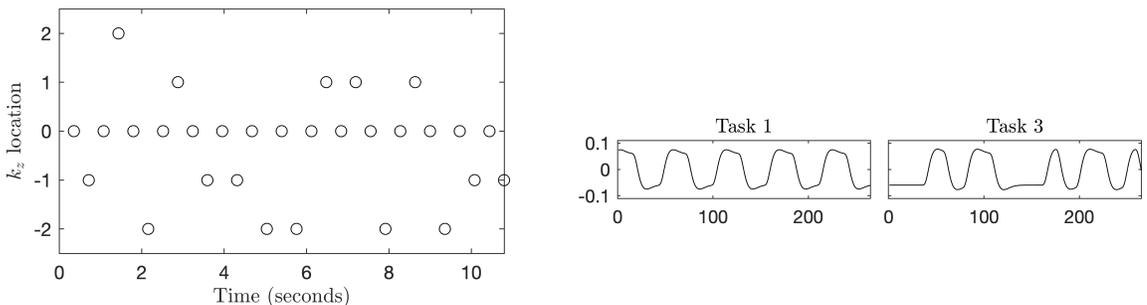


Figure 5.4.1: $3 \times k_z$ - t undersampling, with fully sampled k_z center and randomly sampled outer k_z space. **Figure 5.4.2:** Two simulated tasks as rows of V , across $N_t = 265$ time frames.

One limit of our simulation is the lack of functional information in the phase, as both the background and the task k -space data were generated by the Fourier

transform of real-valued images. Future work will consider more a realistic simulation; one idea is to add smooth phase images estimated from reconstructed images on human subjects.

Reconstruction result

To compare results reconstructed by three models, we ran each model by sweeping across and selecting regularization parameters that gave the lowest total false positive plus false negative thresholded activation.

To visualize individual components in the sum, Fig. 5.4.3 shows the timeseries of three task-activated voxels and one non-task voxel. In all three models, the magnitudes of the sum of two components are capable of recovering the shape and correlation values of the ground truth timeseries; however, some individual components do not capture background and task as they are designed to. For L+S, without knowledge of the task waveform shape, S tries to capture the task by its periodic behavior. The L+UV model, on the other hand, fails to separate task from background; in all three task voxels, UV captures a negated and shifted version of L . Although $|L+UV|$ recovers the shape of the ground truth, its two components did not represent background and task as the model assumes. In comparison, our B+UV consistently captures the background signal in B and the task in UV , as designed by the model assumptions.

To analyze task correlation of the task component in each model, we computed their receiver operating characteristic (ROC) curves (Fig. 5.4.4). The U maps in B+UV show consistent performance across three tasks; in comparison, S in L+S performs slightly worse for all tasks with lower area under curve (AUC) values, and U in L+UV performs worst with low true positive ratios.

Fig. 5.4.5 shows the activation maps for a center slice, based on the p -values of each model after permutation tests. For the L+S and L+UV models, we show the p -value maps that best match the ground truth, which in this case are the magnitude of the sum of components; this is consistent with our previous observation that their task components might not capture tasks as designed by the models. For our B+UV model, we show the p -value map of U that is designed to capture task correlations. Compared with U in our B+UV, we observe that $|L+S|$ shows more false negatives, while $|L+UV|$ shows slightly more false positives.

Although our B+UV performs better in AUC values and p -value maps, its ROC curves in Fig. 5.4.4 have the potential for lower false positive ratios. Future work will explore different regularization parameter tuning criteria; one idea is to use the

mean of “partial AUC” with a range of small false positive ratio values as a tuning criterion.

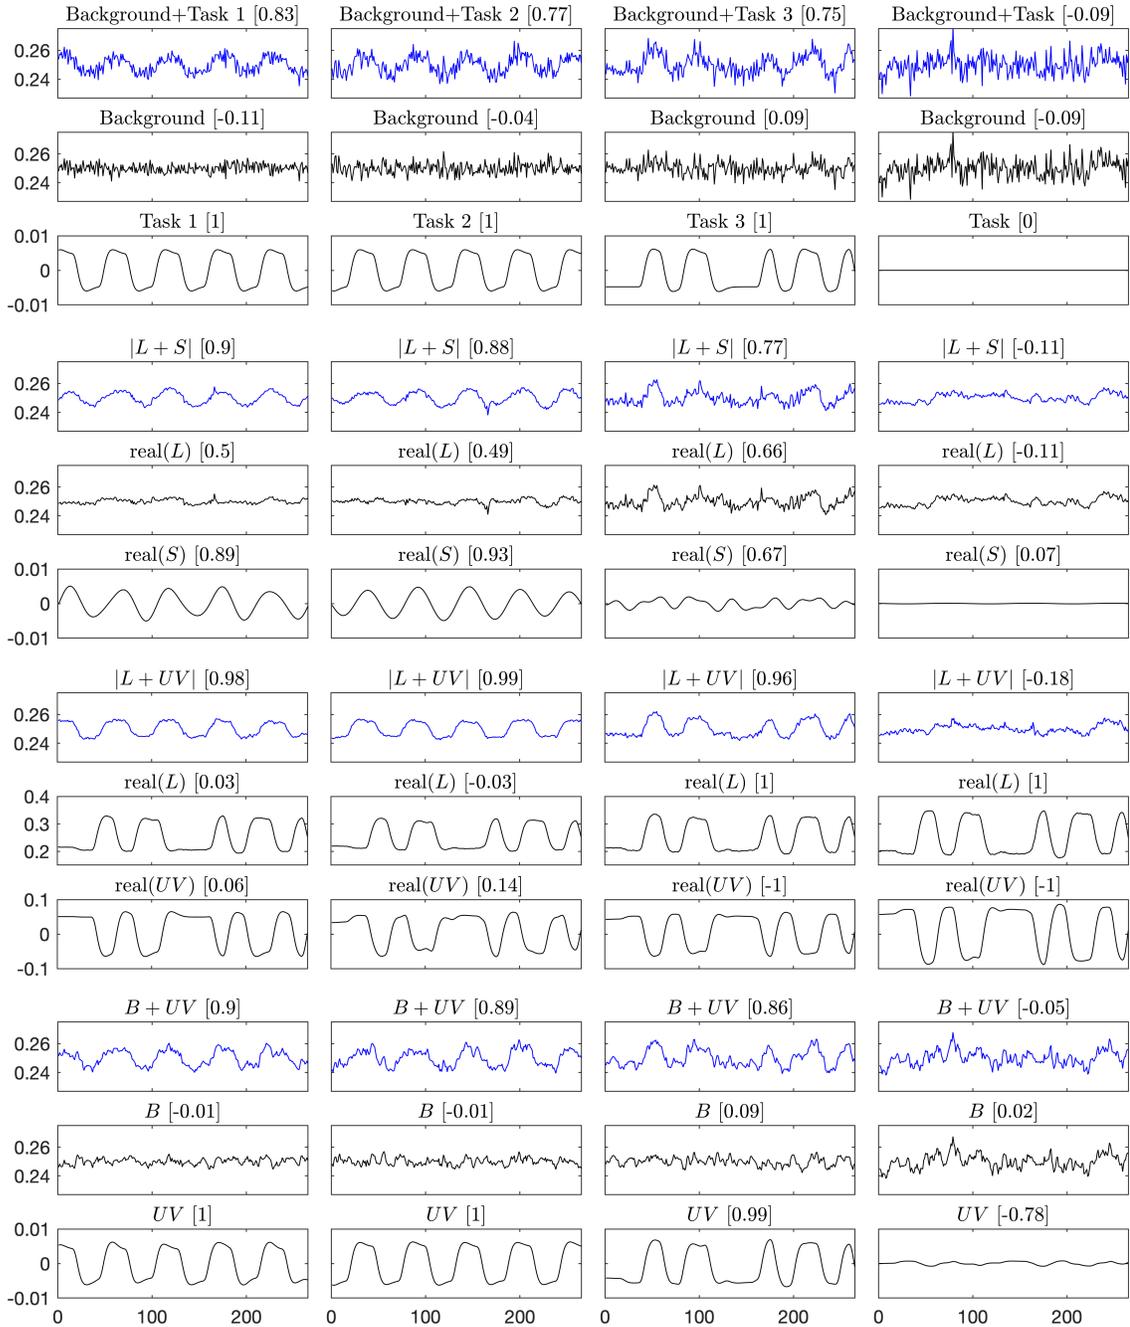


Figure 5.4.3: Timeseries of three activated voxels (corresponding to three simulated tasks) and one non-activated (rightmost column) voxels on simulated data. Top to bottom: ground truth, L+S, L+UV, and B+UV decompositions, with their [correlation] to corresponding activation waveforms (or highest absolute correlation in the non-task case), across $N_t = 265$ time frames.

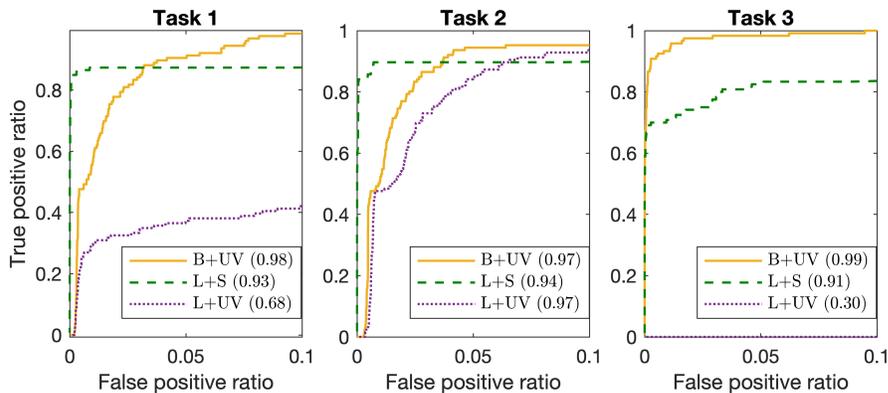


Figure 5.4.4: ROC curves with (AUC values) for simulated task experiment, on task component of each model: U in B+UV, $\text{real}(S)$ in L+S, and $\text{real}(U)$ in L+UV.

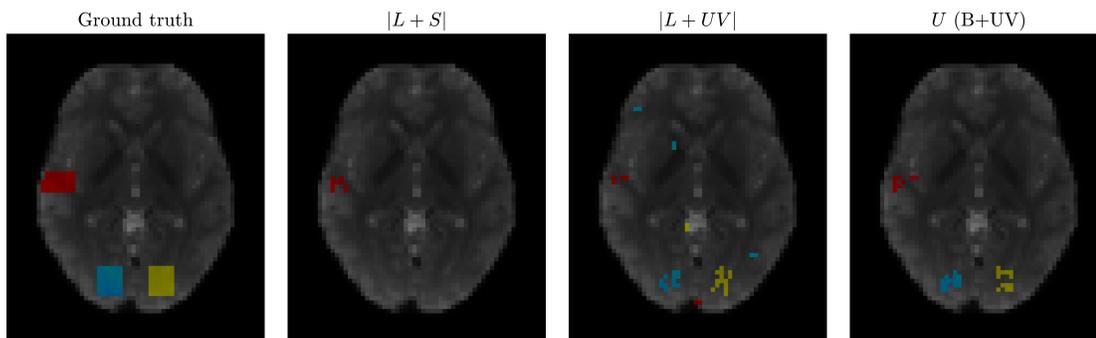


Figure 5.4.5: Activation maps of one slice on simulated dataset, contiguity-thresholded at 0.002 p -values after block permutation (cluster size = 2 voxels), and colored for three tasks.

5.4.2 Reconstruction: visual and motor tasks

Experiment setup

This human dataset uses a 2-shot spiral trajectory with volume $\text{TR} = 500$ ms (slice $\text{TR} = 250$ ms), $N_t = 380$ (with first 10 s of data discarded to justify the use of one task waveform for two visual tasks), and 5 image slices of size of 96×96 . Fig. 5.4.6 shows the rows in V with left visual and motor task waveforms. The visual task is an interleave of 20 seconds of left and 20 seconds of right visual checkerboards task repeated five times, while the motor task is nonperiodic; we assume the right visual task waveform is the negative of the left visual one. For each time frame, we retrospectively undersampled by randomly choosing k -space data from one of the two shots, with $N_k = 7432$ for each slice. We coil-compressed the 32-coil data into

$N_c = 4$ virtual coils, and used ESPIRiT [190] to compute sensitivity maps \mathbf{C} ; we also accounted for field inhomogeneity effects [59] in our system \mathbf{E} , with field maps estimated based on reconstructed spin-warp images using our regularized estimation method (Chapter 4). To compare model performance in task detection, we used the p -value map of the CP reconstruction on fully sampled k -space data as a reference.

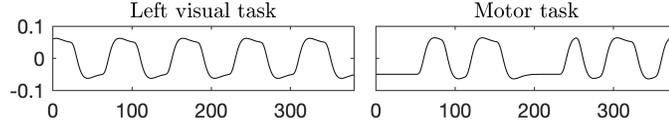


Figure 5.4.6: Two tasks as rows of V on visual and motor task data, across $N_t = 380$ time frames.

Reconstruction result

For this human dataset, we initialized our B+UV by 10 CG iterations of the (convex) B+UV without magnitude and phase separation (5.5).

Fig. 5.4.7 shows the timeseries of three task-activated voxels and one non-task voxel. Similar to the simulation results, all three models recovered the overall shape of the fully sampled timeseries in their magnitudes of sums. The components in L+S display different behaviors, with S trying to capture the task by its periodic assumption. The two components in L+UV fail to separate background and task information as designed by its model, though L contains most background oscillation. In comparison, our B+UV consistently captures background signal in B , and task signal in UV .

Fig. 5.4.8 shows the activation maps for a center slice, based on the p -values of each reconstruction model after block permutation tests. The L+S result fails to distinguish the three tasks, with false negatives for the visual tasks; L+UV and B+UV show comparable results, with slightly more false positives in L+UV.

One difference between this human dataset and our simulation in Section 5.4.1 is in the scale of their reconstructed images. Future work will address this difference by, for example, scaling the human k -space data to match our simulation. With more comparable image values, we can explore using the regularization parameters in our simulation to guide their selection in human experiments.

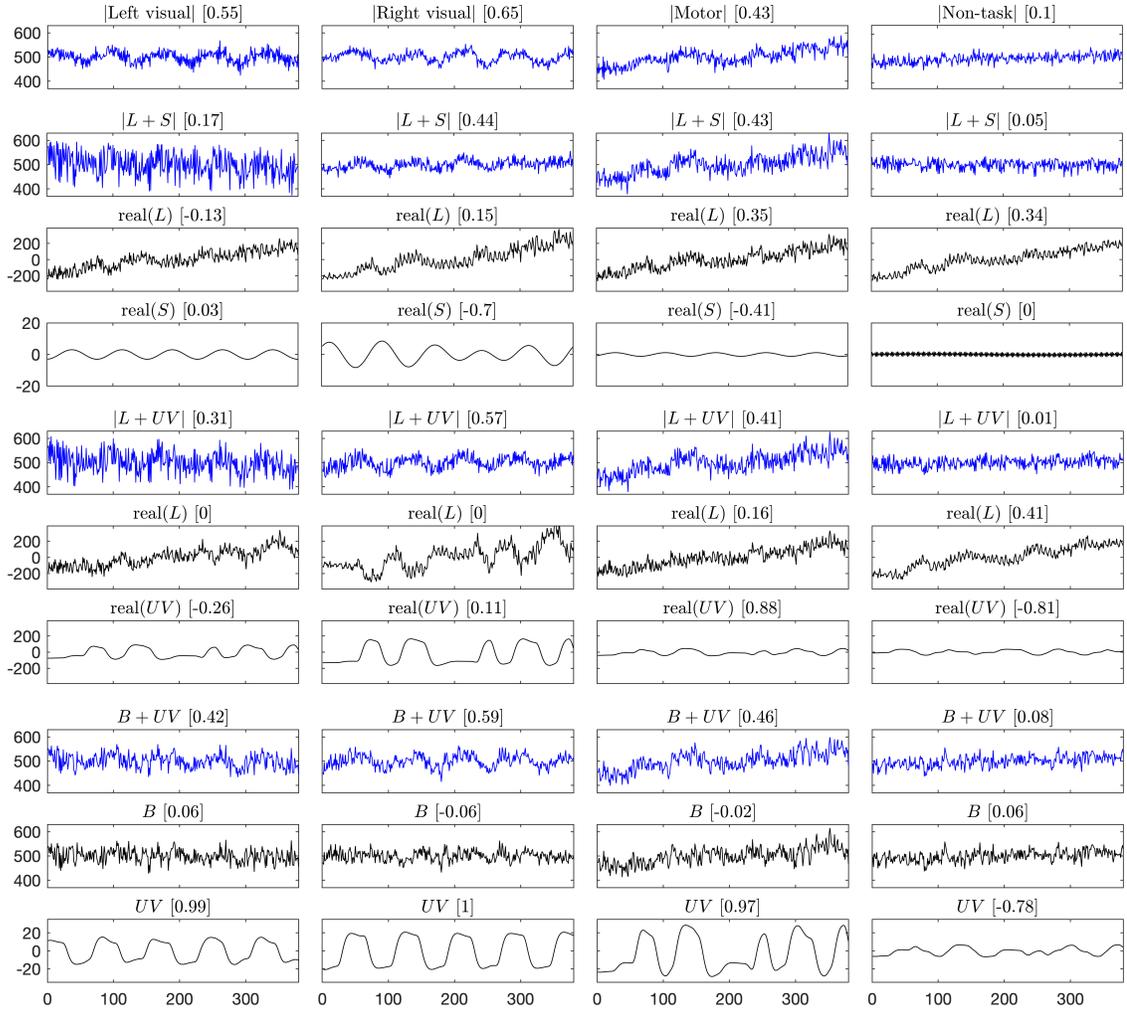


Figure 5.4.7: Timeseries of three activated and one non-activated (rightmost column) voxels on visual and motor task data. Top to bottom: fully sampled CP reconstruction, L+S, L+UV, and B+UV reconstructions, with their [correlation] to left visual, right visual, and motor activation waveforms (or highest absolute correlation in the non-task case), across $N_t = 380$ time frames.

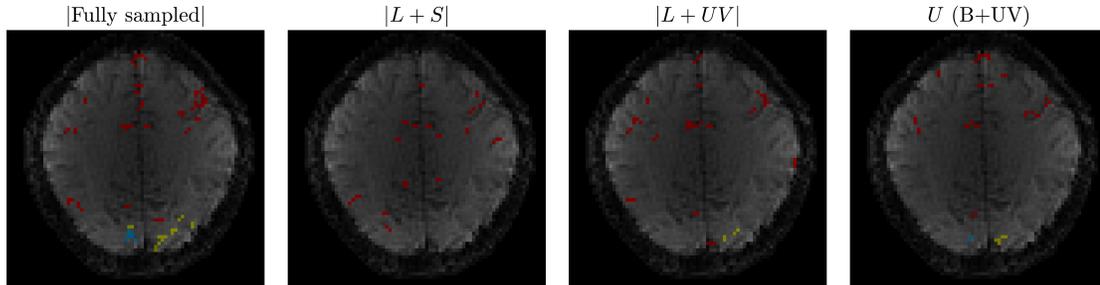


Figure 5.4.8: Activation maps on visual and motor task data, contiguity-thresholded at 0.002 p -values after block permutation, and colored for three tasks. The model-based reconstruction results were all from data undersampled by a factor of 2.

5.5 Conclusion

We present a task-based fMRI model for image sequence reconstruction from spatiotemporal undersampled multi-coil k -space data, with an aim to detect task-activated brain regions. Our proposed B+UV model assumes separation of task from non-task background signal, with spatial and temporal finite difference regularization on the background magnitude B and phase images, and task waveforms in the temporal basis V that drives UV to capture task activations. We use the nonlinear conjugate gradient method with line search to solve the optimization problem, and our algorithm is guaranteed to monotonically decrease the cost. Compared with existing reconstruction models for task-based fMRI, the proposed method yields better timeseries and activation maps on simulated and human subject task fMRI data.

One practical concern of our method is regularization parameter tuning for (5.4). One way to guide this selection is based on spatial and temporal local impulse responses; to this end, [62,74,146] can be extended to our setting. Other regularization scheme, such as second-order finite difference on B , can also be explored. To increase model robustness to subject task performance, the temporal basis V can be generalized to account for mismatch and deviation from the designed task waveforms. One might even consider joint estimation of spatial map U and temporal basis V .

The 3D multi-coil dynamic setting is a large-scale problem that is computationally and memory intensive; to this end, convergence speedup can be explored by exploiting problem structures and applying efficient preconditioners [105]. Finally, one may consider extending this model to general fMRI reconstruction (e.g., resting state fMRI) under spatial and temporal smoothness assumptions, with separate magnitude and phase regularization.

Chapter 6

A Model-based Reconstruction Framework

The preceding chapters have addressed some ingredients that are part of the optimization problem (2.3) for model-based dynamic MRI reconstruction, such as the regularizer $R(X)$ and optimization algorithms given specific cost formulations. We have also considered \mathbf{E} with nonuniform sampling and coil sensitivities. To further address artifacts, such as magnetic field inhomogeneity and motion effects, this chapter discusses implementation details of the system operator \mathbf{E} introduced in Section 2.2.1. Section 6.1 considers reconstruction accounting for intravoxel field inhomogeneity, and Section 6.2 discusses the implementation of head motion compensation. Finally, Section 6.3 presents a general framework for dynamic MRI reconstruction with all ingredients considered in this dissertation.

6.1 Intravoxel field inhomogeneity correction

FMRI with BOLD contrast typically requires relatively long readouts to track rapid functional changes, using acquisitions such as spiral and echo-planar imaging (EPI). This leads to its sensitivity to magnetic field inhomogeneity effects, causing geometric distortions in EPI and blurring in spiral imaging, especially near air/tissue interface. After a field map ω is estimated from MR scans having multiple echo times, as discussed in Chapter 4, we can account for its effect in the system operator \mathbf{E} during reconstruction. This section assumes a static field map; dynamic field map estimation has also been explored with joint image reconstruction [84,128,145,185,192,205]. Below we discuss the implementation of reconstruction with field map correction, and illustrate its improvement using a simulation.

6.1.1 Implementation

Signal equation

The signal equation with field inhomogeneity (2.4) at sampling time t_i is

$$d_i = \int x(\vec{r}) e^{i\omega(\vec{r})t_i} e^{-i2\pi k_i \cdot \vec{r}} d\vec{r}. \quad (6.1)$$

For a single time frame, when discretizing the image $x(\vec{r})$ and field map $\omega(\vec{r})$; different choices of their spatial basis functions lead to different formulations of \mathbf{E} in the matrix equation (2.2). A common basis function for a (3D) image is the rectangular function, where we assume each voxel has a constant image value:

$$x(\vec{r}) \approx \sum_{j=1}^{N_v} x_j \text{rect}_3((\vec{r} - \vec{r}_j) \oslash \Delta_r),$$

where Δ_r is a voxel resolution vector. Similarly, one might consider the same rectangular basis function for the field map [184]. However, this piecewise constant basis does not consider field inhomogeneity effects within a voxel, which can have nonzero gradients that lead to intravoxel spin dephase. To account for intravoxel susceptibility, [108,183] consider a piecewise linear basis for $\omega(\vec{r})$:

$$\omega(\vec{r}) \approx \sum_{j=1}^{N_v} (\omega_j + 2\pi \mathbf{g}_j \cdot (\vec{r} - \vec{r}_j)) \text{rect}_3((\vec{r} - \vec{r}_j) \oslash \Delta_r),$$

where \mathbf{g}_j is intravoxel field gradient at the j th voxel that can be computed using the difference of field map values at neighboring voxels. With this image-space discretization, our signal equation (6.1) becomes

$$d_i \approx \text{sinc}_3((k_i - t_i \mathbf{g}_j) \odot \Delta_r) \sum_{j=1}^{N_v} x_j e^{i\omega_j t_i} e^{-i2\pi k_i \cdot \vec{r}_j}. \quad (6.2)$$

Basis expansion approximation

Given the large problem size, it is computationally and memory impractical to form a system operator to compute (6.2) for all i, j . To this end, we adopt a basis

expansion approximation [59] for the intravoxel field inhomogeneity terms:

$$W_{ij} := \text{sinc}_3((k_i - t_i \mathbf{g}_j) \odot \Delta_r) e^{i\omega_j t_i} \approx \sum_{l=1}^L B_{il} C_{lj}. \quad (6.3)$$

Then the approximate signal equation (6.2) becomes

$$d_i \approx \sum_{l=1}^L B_{il} \left(\sum_{j=1}^{N_v} (x_j C_{lj}) e^{-i2\pi k_i \cdot \bar{r}_j} \right),$$

or in matrix form

$$\mathbf{d} \approx \sum_{l=1}^L \mathbf{B}_l \mathbf{Q} \mathbf{C}_l \mathbf{x},$$

where $\mathbf{B}_l = \text{diag}(B_{il})$, $\mathbf{C}_l = \text{diag}(C_{lj})$, and $\mathbf{Q} : \mathbb{C}^{N_v} \rightarrow \mathbb{C}^{N_k}$ is a Fourier transform operator.

To obtain the approximation bases $\{B_{il}\}, \{C_{lj}\}$, we consider an extension of the histogram principal components method in [59], where a 1D histogram is used to approximate the exponential function. In our case with intravoxel field gradient, the basis expansion (6.3) is a product of four functions, and a direct generalization to 4D histogram is computationally impractical. Hence we adopt a sketching scheme described in [108]. To estimate $\{B_{il}\}$, we first sample $N \ll N_v$ field map and gradient values randomly from all (masked) voxels, then perform an SVD on the matrix $W_N \in \mathbb{C}^{N_k \times N}$ in (6.3) formed at the sampled voxels. We speed up the computation by a truncated SVD using the eigenvalues and eigenvectors of $W_N^* W_N$ as in [199] to get the L dominant left singular vectors. To estimate $\{C_{lj}\}$, we use a sketching scheme on W in (6.3) by uniformly sampling K values from all N_k sampling time values $\{t_i\}$, then perform a least-squares fit using the estimated $\{B_{il}\}$.

In functional MRI, after applying motion correction on an image, we use a static field map and its approximations for all time frames during reconstruction. This is an approximation that disregards field drifts during the scan and changes in the field map due to changes in susceptibility effects as the head moves in the main field. Such changes are typically disregarded in fMRI and the approximation should be reasonable provided the motion is relatively small.

6.1.2 Simulation

This simulation uses images and field maps estimated in Section 4.4.2 to illustrate 3D image reconstruction with intravoxel field inhomogeneity correction. We synthe-

size k -space data for two image slices with a stack of spiral trajectory, using the signal equation (6.2). We then reconstruct images using a quadratic cost with spatial smoothing, with four system operator formulations: (1) \mathbf{E} without field inhomogeneity compensation, (2) $\mathbf{E}(\omega)$ with piecewise constant basis function for field map, (3) $\mathbf{E}(\omega, \mathbf{g}_z)$ with through-plane field gradient effects (disregarding in-plane gradients), and (4) $\mathbf{E}(\omega, \mathbf{g})$ with 3D intravoxel field inhomogeneity compensation.

Fig. 6.1.1 shows the field map ω and its field gradients in three dimensions. Fig. 6.1.2 shows the reconstructed images and their error maps, using the four system operators described above. Images reconstructed with compensation of intravoxel field inhomogeneity effects in all three dimensions give the lowest normalized root mean square error (NRMSE) to the ground truth.

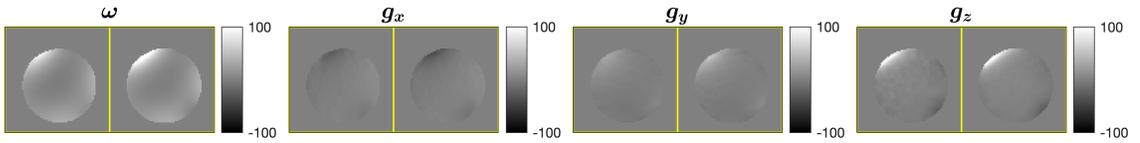


Figure 6.1.1: B_0 (in Hz) and gradient maps (in Hz/meter) in simulation.

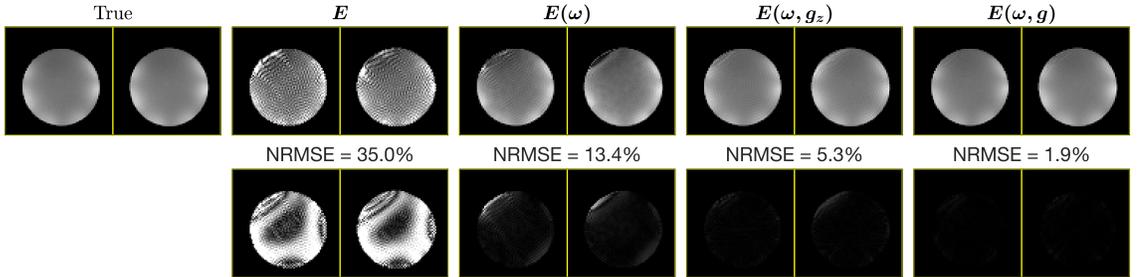


Figure 6.1.2: Reconstructed and error images in simulation, plotted on the same scale of $[0, 1]$.

6.2 Head motion compensation

Head motion is known to produce artifacts in MRI and spurious correlation in fMRI studies [155,159]; hence it is crucial to identify and correct for motion effects prior to analysis and diagnosis. One way to correct for motion in fMRI is to reconstruct images at each time frame, then register them to a reference image. In spatiotemporally undersampled fMRI, it is efficient to incorporate motion correction in the reconstruction process. Joint image reconstruction and motion estimation has been explored in both static and dynamic MRI [42,78,121], with a nonconvex cost

function. To avoid nonconvexity and focus on reconstruction of the image sequence, we can estimate the motion parameters using a crude initial reconstruction, and build them into a (fixed) system operator \mathbf{E} . Below we discuss the implementation of a system operator with motion compensation, and illustrate its improvement in reconstruction using a simulation.

6.2.1 Implementation

The 3D bulk movement of the head is a rigid motion that can be measured using registration and captured by six parameters (three for rotation and three for translation). To estimate these parameters for each time frame, we compute an initial reconstruction of the image sequence, such as a conjugate phase reconstruction [139]. We use image registration tools such as MATLAB’s Image Processing toolbox or the Statistical Parametric Mapping (SPM) software [6] to estimate 3D rigid motion of the image from each frame to the first frame.

To account for motion, we include the estimated rotation and translation parameters in our system operator \mathbf{E} during reconstruction. Motion correction can be implemented in k -space using their equivalent rotation and linear phase shift [13,26,155,193]. In multi-coil reconstruction, however, since the coils are static, coil sensitivity maps are in motion relative to the image sequence; thus, with k -space-based correction methods one needs image space registration of the coil maps, which scales up computational complexity by the number of coils. This motivates us to focus on image space registration of the image sequence during reconstruction. There are numerous techniques for image registration with different interpolation schemes [91]; here we adopt the Fourier interpolation method [52] in our reconstruction framework.

Fourier interpolation in k -space represents a 3D rotation matrix as a product of three 2D rotation matrices, each of which is approximated by the product of three shear matrices [191]. [42] adopts this scheme to image space transformation, with Fourier transform and inverse Fourier transform along each dimension before and after each shearing operation. After rotation by shearing, a translation can be directly applied to the image, or in k -space with linear phase shift as in [42]. In reconstruction, we apply this transformation $\mathbf{T} : \mathbb{C}^{N_v \times N_t} \rightarrow \mathbb{C}^{N_v \times N_t}$ for all N_t time frames to get a rigid transformed image \mathbf{TX} , before the Fourier transform and coil sensitivity operations.

6.2.2 Simulation

This simulation uses unprocessed resting state images from the NYU dataset [179] that contains natural (i.e., not cued) subject head motion during scan. We synthesize k -space data using a stack of spiral trajectory, keeping the original image size $64 \times 64 \times 39$ and resolution 3 mm, with 197 time frames. This k -space data has the motion effects present in the unprocessed image sequence. We then use a conjugate phase reconstruction using an \mathbf{E} with sampling density compensation computed by an iterative method [156] to obtain an image sequence with motion, and use MATLAB's `imregtform` to estimate rigid motion parameters for all frames. We compare it with conjugate phase reconstruction using a system operator $\mathbf{E}(\mathbf{T})$ with motion compensation.

We compare the two reconstructions against an image sequence corrected by SPM. Fig. 6.2.1 shows an image slice at different four time frames across the scan and their difference map to the SPM-correct images. We observe a visible motion effect towards back of the brain. Fig. 6.2.2 shows the signal change across time for a voxel in the brain region with motion effects; the motion-corrected timeseries resembles the SPM-correct timeseries, while the uncorrected timeseries exhibits some discrepancy.

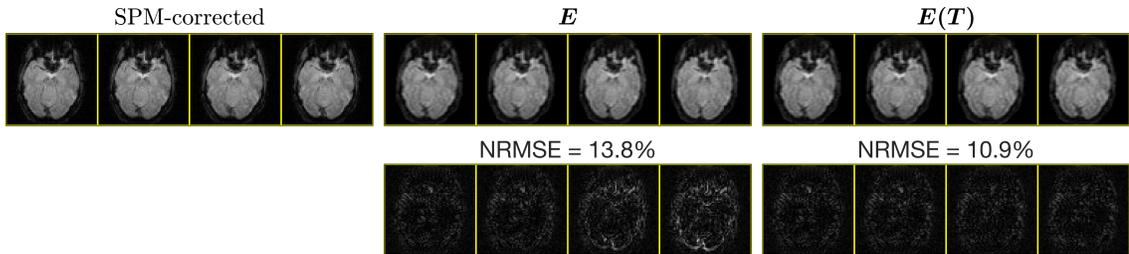


Figure 6.2.1: SPM motion-corrected and conjugate phase reconstructed images without and with motion correction, for the 1st, 61st, 121st, and 181st time frames.

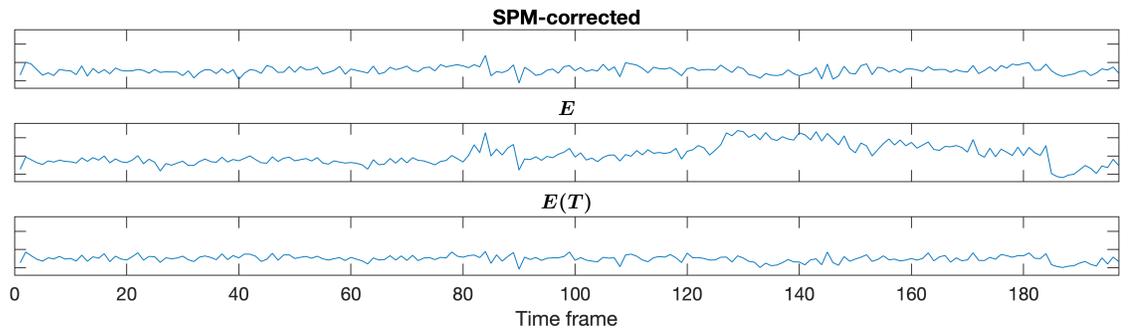


Figure 6.2.2: Timeseries of a voxel in SPM motion-corrected image, and conjugate phase reconstructed image without and with motion correction.

6.3 Summary: a general framework

This section summarizes the overarching story of this dissertation by putting all the components together in a general model-based reconstruction framework. This framework considers undersampled multi-coil dynamic MRI with the option of field inhomogeneity and head motion correction.

Fig. 6.3.1 shows the framework in a flowchart. In data acquisition, we collect (dynamic) k -space data \mathbf{d} for the image sequence, and (static) multi-echo data for field inhomogeneity estimation. We compute the Fourier operator Q with undersampling Ω using the (known) k -space trajectory, as discussed in Section 2.2.1. Using an initial reconstruction X_0 (e.g., by conjugate phase) on \mathbf{d} , we can estimate sensitivity maps S and motion parameters across time frames for a rigid transformation \mathbf{T} , as discussed in Sections 2.2.1 and 6.2. With a reconstruction \mathbf{y} on the multi-echo data, we can estimate field map ω and use its difference as gradient \mathbf{g} , as in Chapter 4 and Section 6.1. For efficient implementation we approximate the intravoxel field inhomogeneity by basis expansion, as discussed in Sections 2.2.1 and 6.1. We consider these components in a system operator \mathbf{E} as in Section 2.2.1, and formulate an optimization problem with regularization $R(X)$ that considers assumptions on images X . Chapter 5 is one example of a regularized optimization problem in the context of task-based functional MRI. With a specified cost function, we also explore efficient algorithms to solve for X , and Chapter 3 is an example of accelerated algorithms for the low-rank plus sparse decomposition in dynamic MRI reconstruction. After reconstructing images \hat{X} , we evaluate and analyze the result based on our goals in reconstruction, such as spatial quality and convergence as in Chapter 3, or timeseries and correlation maps as in Chapter 5.

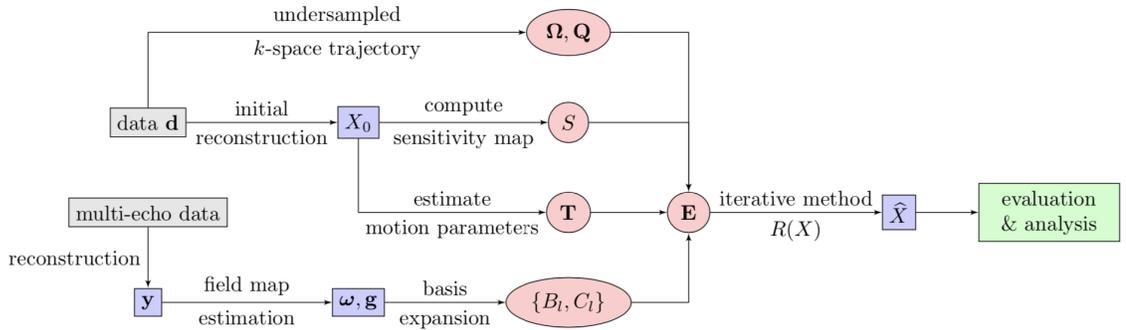


Figure 6.3.1: A general model-based reconstruction framework.

Gray box: k -space data, blue box: images, red ellipse: system operator components.

Chapter 7

Future Work

7.1 Physiological noise correction

A type of artifact correction not addressed explicitly in this dissertation is physiological noise including cardiac and respiratory rhythms, a known confounding factor in fMRI signals [70]. While conventional fMRI attempts to remove physiological noise in postprocessing, we desire to address it during reconstruction. The formulation in Chapter 5 simply absorbs such effects into the background image sequence B and possibly the phase sequence P , without using any reference waveforms.

One category of correction involves external measurement such as electrocardiography (EKG), pulse oximetry, and respiratory belts to record physiological signals. After preprocessing the measurements [71,77,99], a direct extension to the spatial-temporal model in Chapter 5 is to account for the physiological timeseries as part of the temporal basis, and discard its spatiotemporal component after reconstruction.

Due to the added technical complexity in acquiring external physiological signals, data-driven approaches have been proposed to estimate physiological noise directly from fMRI images [18,40,153,187]. One possible direction to account for it during reconstruction is to apply similar techniques on an initial reconstruction from under-sampled (e.g., center of k -space) data, then build it into a temporal basis as mentioned above. Another direction is to include physiological noise estimation as part of the optimization problem. [131] estimates both spatial and temporal bases during reconstruction using a temporal Fourier sparsity regularizer. One can further specify regularization by focusing on the range of plausible physiological frequencies.

7.2 Other models

The design of our task-based fMRI model in Chapter 5 is driven by its goal to identify task-activated brain regions. In the case of resting state fMRI, we also desire

a reconstruction model that accounts for signal properties and artifact corrections. A comprehensive model evaluation includes comparison to contemporary resting state reconstruction models such as [9,39,199] and functional connectivity analysis. It also remains to be explored how to effectively combine handcrafted regularizers with machine learning in fMRI reconstruction.

7.3 Model-driven data acquisition

This dissertation focuses on reconstructing high spatiotemporal resolution images given undersampled k -space data. In the other direction, our reconstruction model also helps motivate data acquisition techniques. One example is task-based fMRI, where the analysis of activation maps is much more emphasized over the spatial resolution of an image sequence. With a known task waveform, therefore, our goal is to estimate the spatial parameters corresponding to a given temporal basis. This goal, specific to task fMRI, can drive the design of faster data acquisition techniques while maintaining (or improving) activation detection. One might also consider applying magnetic resonance fingerprinting (MRF) [126] to dynamic MRI with a specific model structure. Another direction is joint trajectory design and image reconstruction; to this end, extension from the 2D setting [196] to 3D dynamic MRI is desired. After all, collaboration between data acquisition and image reconstruction is crucial to fast and high quality magnetic resonance imaging.

Bibliography

- [1] 1200 Subjects Data Release. <https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release>. Accessed: 2020-11-17.
- [2] Fat-Water Toolbox. <http://ismrm.org/workshops/FatWater12/data.htm>. Accessed: 2020-03-07.
- [3] Matlab codes for k-t SLR. <https://research.engineering.uiowa.edu/cbig/content/matlab-codes-k-t-slr>. Accessed: 2018-05-20.
- [4] Michigan Image Reconstruction Toolbox (MIRT). <https://web.eecs.umich.edu/~fessler/code>. Accessed: 2020-02-20.
- [5] Poblano Toolbox for MATLAB. <https://github.com/sandialabs>. Accessed: 2020-08-16.
- [6] Statistical Parametric Mapping. <https://www.fil.ion.ucl.ac.uk/spm>. Accessed: 2020-10-8.
- [7] D. Adolf, S. Weston, S. Baecke, M. Luchtmann, J. Bernarding, and S. Kropf. Increasing the reliability of data analysis of functional magnetic resonance imaging by applying a new blockwise permutation method. *Frontiers in neuroinformatics*, 8:72, 2014.
- [8] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo. An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems. *IEEE Transactions on Image Processing*, 20(3):681–695, 2010.
- [9] P. Aggarwal and A. Gupta. Double temporal sparsity based accelerated reconstruction of compressively sensed resting-state fMRI. *Computers in biology and medicine*, 91:255–266, 2017.
- [10] M. J. Allison and J. A. Fessler. Accelerated computation of regularized field map estimates. In *Proc. Intl. Soc. Mag. Res. Med*, page 0413, 2012.
- [11] S. K. Arja, Z. Feng, Z. Chen, A. Caprihan, K. A. Kiehl, T. Adali, and V. D. Calhoun. Changes in fMRI magnitude data and phase data observed in block-design and event-related tasks. *NeuroImage*, 49(4):3149–3160, 2010.

- [12] S. P. Awate and E. V. DiBella. Spatiotemporal dictionary learning for undersampled dynamic MRI reconstruction via joint frame-based and dictionary-based sparsity. In *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 318–321. IEEE, 2012.
- [13] R. Bammer, M. Aksoy, and C. Liu. Augmented generalized SENSE reconstruction to correct for rigid body motion. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 57(1):90–102, 2007.
- [14] H. H. Bauschke, P. L. Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- [15] A. Beck and M. Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE transactions on image processing*, 18(11):2419–2434, 2009.
- [16] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- [17] C. F. Beckmann, M. DeLuca, J. T. Devlin, and S. M. Smith. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1457):1001–1013, 2005.
- [18] Y. Behzadi, K. Restom, J. Liau, and T. T. Liu. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage*, 37(1):90–101, 2007.
- [19] C. Bilen, I. W. Selesnick, Y. Wang, R. Otazo, D. Kim, L. Axel, and D. K. Sodickson. On compressed sensing in parallel MRI of cardiac perfusion using temporal wavelet and TV regularization. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 630–633. IEEE, 2010.
- [20] Ç. Bilen, Y. Wang, and I. W. Selesnick. High-speed compressed sensing reconstruction in dynamic parallel MRI using augmented Lagrangian and parallel processing. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(3):370–379, 2012.
- [21] J. M. Bioucas-Dias and G. Valadao. Phase unwrapping via graph cuts. *IEEE Transactions on Image processing*, 16(3):698–709, 2007.
- [22] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic resonance in medicine*, 34(4):537–541, 1995.
- [23] S. Biswas, H. K. Aggarwal, and M. Jacob. Dynamic MRI using model-based deep learning and SToRM priors: MoDL-SToRM. *Magnetic resonance in medicine*, 82(1):485–494, 2019.

- [24] T. A. Bley, O. Wieben, C. J. François, J. H. Brittain, and S. B. Reeder. Fat and water magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 31(1):4–18, 2010.
- [25] S. Boyd, N. Parikh, and E. Chu. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- [26] M. Bydder, D. J. Larkman, and J. V. Hajnal. Detection and elimination of motion artifacts by regeneration of k -space. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 47(4):677–686, 2002.
- [27] M. Bydder, T. Yokoo, G. Hamilton, M. S. Middleton, A. D. Chavez, J. B. Schwimmer, J. E. Lavine, and C. B. Sirlin. Relaxation effects in the quantification of fat using gradient echo imaging. *Magnetic resonance imaging*, 26(3):347–359, 2008.
- [28] N. Cai, W. Xie, Z. Su, S. Wang, and D. Liang. Sparse parallel MRI based on accelerated operator splitting schemes. *Computational and mathematical methods in medicine*, 2016, 2016.
- [29] V. D. Calhoun, T. Adalı, G. Pearlson, P. Van Zijl, and J. J. Pekar. Independent component analysis of fMRI data in the complex domain. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 48(1):180–192, 2002.
- [30] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):1–37, 2011.
- [31] L. Chaari, P. Ciuciu, S. Mériaux, and J.-C. Pesquet. Spatio-temporal wavelet regularization for parallel MRI reconstruction: application to functional MRI. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 27(6):509–529, 2014.
- [32] R. H. Chan and M. K. Ng. Conjugate gradient methods for Toeplitz systems. *SIAM review*, 38(3):427–482, 1996.
- [33] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st International Conference on Image Processing*, volume 2, pages 168–172. IEEE, 1994.
- [34] C. Chavarrias, J. Abascal, P. Montesinos, and M. Desco. Exploitation of temporal redundancy in compressed sensing reconstruction of fMRI studies with a prior-based algorithm (PICCS). *Medical physics*, 42(7):3814–3821, 2015.
- [35] Y. Chen, W. Hager, F. Huang, D. Phan, X. Ye, and W. Yin. Fast algorithms for image reconstruction with application to partially parallel MR imaging. *SIAM Journal on Imaging Sciences*, 5(1):90–118, 2012.

- [36] M. Chiew, N. N. Graedel, J. A. McNab, S. M. Smith, and K. L. Miller. Accelerating functional MRI using fixed-rank approximations and radial-cartesian sampling. *Magnetic resonance in medicine*, 76(6):1825–1836, 2016.
- [37] M. Chiew, N. N. Graedel, and K. L. Miller. Recovering task fMRI signals from highly under-sampled data with low-rank and temporal subspace constraints. *NeuroImage*, 174:97–110, 2018.
- [38] M. Chiew and K. L. Miller. Improved statistical efficiency of simultaneous multi-slice fMRI by reconstruction with spatially adaptive temporal smoothing. *NeuroImage*, 203:116165, 2019.
- [39] M. Chiew, S. M. Smith, P. J. Koopmans, N. N. Graedel, T. Blumensath, and K. L. Miller. k-t faster: acceleration of functional mri data acquisition using low rank constraints. *Magnetic resonance in medicine*, 74(2):353–364, 2015.
- [40] N. W. Churchill and S. C. Strother. PHYCAA+: An optimized, adaptive procedure for measuring and controlling physiological noise in BOLD fMRI. *Neuroimage*, 82:306–325, 2013.
- [41] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- [42] L. Cordero-Grande, R. P. A. Teixeira, E. J. Hughes, J. Hutter, A. N. Price, and J. V. Hajnal. Sensitivity encoding for aligned multishot magnetic resonance reconstruction. *IEEE Transactions on Computational Imaging*, 2(3):266–280, 2016.
- [43] Y.-H. Dai and Y. Yuan. A nonlinear conjugate gradient method with a strong global convergence property. *SIAM Journal on optimization*, 10(1):177–182, 1999.
- [44] J. S. Damoiseaux, S. Rombouts, F. Barkhof, P. Scheltens, C. J. Stam, S. M. Smith, and C. F. Beckmann. Consistent resting-state networks across healthy subjects. *Proceedings of the national academy of sciences*, 103(37):13848–13853, 2006.
- [45] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 57(11):1413–1457, 2004.
- [46] J. A. de Zwart, P. van Gelderen, P. Kellman, and J. H. Duyn. Application of sensitivity-encoded echo-planar imaging for blood oxygen level-dependent functional brain imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 48(6):1011–1020, 2002.

- [47] R. Deichmann, O. Josephs, C. Hutton, D. Corfield, and R. Turner. Compensation of susceptibility-induced BOLD sensitivity losses in echo-planar fMRI imaging. *Neuroimage*, 15(1):120–135, 2002.
- [48] M. N. Diefenbach, S. Ruschke, H. Eggers, J. Meineke, E. J. Rummeny, and D. C. Karampinos. Improving chemical shift encoding-based water-fat separation based on a detailed consideration of magnetic field contributions. *Magnetic resonance in medicine*, 80(3):990–1004, 2018.
- [49] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, and M. Lustig. Compressed sensing for chemical shift-based water-fat separation. *Magnetic resonance in medicine*, 64(6):1749–1759, 2010.
- [50] Y. Drori and M. Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. *Mathematical Programming*, 145(1-2):451–482, 2014.
- [51] J. Eckstein. Parallel alternating direction multiplier decomposition of convex programs. *Journal of Optimization Theory and Applications*, 80(1):39–62, 1994.
- [52] W. F. Eddy, M. Fitzgerald, and D. C. Noll. Improved image registration by using Fourier interpolation. *Magnetic resonance in medicine*, 36(6):923–931, 1996.
- [53] R. C. Fair. On the robust estimation of econometric models. *Annals of Economic and Social Measurement*, 2:667–677, 1974.
- [54] Z. Fang, N. Van Le, M. Choy, and J. H. Lee. High spatial resolution compressed sensing (HSPARSE) functional MRI. *Magnetic resonance in medicine*, 76(2):440–455, 2016.
- [55] J. Fessler and B. Sutton. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing*, 51:560–574, 2003.
- [56] J. A. Fessler. Model-based image reconstruction for MRI. *IEEE signal processing magazine*, 27(4):81–89, 2010.
- [57] J. A. Fessler. Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms. *IEEE Signal Processing Magazine*, 37(1):33–40, 2020.
- [58] J. A. Fessler and S. D. Booth. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE transactions on image processing*, 8(5):688–699, 1999.
- [59] J. A. Fessler, S. Lee, V. T. Olafsson, H. R. Shi, and D. C. Noll. Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity. *IEEE Transactions on Signal Processing*, 53(9):3393–3402, 2005.

- [60] J. A. Fessler and D. C. Noll. Iterative image reconstruction in MRI with separate magnitude and phase regularization. In *2004 2nd IEEE International Symposium on Biomedical Imaging: Nano to Macro (IEEE Cat No. 04EX821)*, pages 209–212. IEEE, 2004.
- [61] J. A. Fessler and D. C. Noll. Model-based MR image reconstruction with compensation for through-plane field inhomogeneity. In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 920–923. IEEE, 2007.
- [62] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs. *IEEE Transactions on Image Processing*, 5(9):1346–1358, 1996.
- [63] R. Fletcher and C. M. Reeves. Function minimization by conjugate gradients. *The computer journal*, 7(2):149–154, 1964.
- [64] K. J. Friston, A. P. Holmes, K. J. Worsley, J.-P. Poline, C. D. Frith, and R. S. Frackowiak. Statistical parametric maps in functional imaging: a general linear approach. *Human brain mapping*, 2(4):189–210, 1994.
- [65] A. K. Funai, J. A. Fessler, D. T. Yeo, V. T. Olafsson, and D. C. Noll. Regularized field map estimation in MRI. *IEEE transactions on medical imaging*, 27(10):1484–1494, 2008.
- [66] C. R. Genovese, N. A. Lazar, and T. Nichols. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15(4):870–878, 2002.
- [67] C. R. Genovese, D. C. Noll, and W. F. Eddy. Estimating test-retest reliability in functional MR imaging I: Statistical methodology. *Magnetic Resonance in Medicine*, 38(3):497–507, 1997.
- [68] G. Glover and E. Schneider. Three-point Dixon technique for true water/-fat decomposition with B_0 inhomogeneity correction. *Magnetic resonance in medicine*, 18(2):371–383, 1991.
- [69] G. H. Glover and S. Lai. Self-navigated spiral fMRI: interleaved versus single-shot. *Magnetic resonance in medicine*, 39(3):361–368, 1998.
- [70] G. H. Glover and A. T. Lee. Motion artifacts in fMRI: comparison of 2DFT with PR and spiral scan methods. *Magnetic resonance in medicine*, 33(5):624–635, 1995.
- [71] G. H. Glover, T.-Q. Li, and D. Ress. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 44(1):162–167, 2000.

- [72] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.
- [73] M. I. Grivich and D. P. Jackson. The magnetic field of current-carrying polygons: An application of vector field rotations. *American Journal of Physics*, 68(5):469–474, 2000.
- [74] S. Guo, J. A. Fessler, and D. C. Noll. High-resolution oscillating steady-state fMRI using patch-tensor low-rank reconstruction. *IEEE Transactions on Medical Imaging*, 39(12):4357–4368, 2020.
- [75] J. P. Haldar, D. Hernando, and Z.-P. Liang. Compressed-sensing MRI with random encoding. *IEEE transactions on Medical Imaging*, 30(4):893–903, 2010.
- [76] T. B. Harshbarger and D. B. Twieg. Iterative reconstruction of single-shot spiral MRI with off resonance. *IEEE transactions on medical imaging*, 18(3):196–205, 1999.
- [77] A. K. Harvey, K. T. Pattinson, J. C. Brooks, S. D. Mayhew, M. Jenkinson, and R. G. Wise. Brainstem functional magnetic resonance imaging: disentangling signal from physiological noise. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 28(6):1337–1344, 2008.
- [78] M. W. Haskell, S. F. Cauley, and L. L. Wald. Targeted motion estimation and reduction (TAMER): data consistency based motion mitigation for MRI using a reduced model joint optimization. *IEEE transactions on medical imaging*, 37(5):1253–1265, 2018.
- [79] N. He, R. Wang, and Y. Wang. Dynamic MRI reconstruction exploiting blind compressed sensing combined transform learning regularization. *Neurocomputing*, 392:160–167, 2020.
- [80] R. M. Heidemann, M. A. Griswold, N. Seiberlich, G. Krüger, S. A. Kannengiesser, B. Kiefer, G. Wiggins, L. L. Wald, and P. M. Jakob. Direct parallel image reconstructions for spiral trajectories using GRAPPA. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 56(2):317–326, 2006.
- [81] D. Hernando, J. Haldar, B. Sutton, J. Ma, P. Kellman, and Z.-P. Liang. Joint estimation of water/fat images and field inhomogeneity map. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 59(3):571–580, 2008.

- [82] D. Hernando, P. Kellman, J. Haldar, and Z.-P. Liang. Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 63(1):79–90, 2010.
- [83] P. W. Holland and R. E. Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics-theory and Methods*, 6(9):813–827, 1977.
- [84] C. Hu, S. Reeves, D. C. Peters, and D. Twieg. An efficient reconstruction algorithm based on the alternating direction method of multipliers for joint estimation of R_2^* and off-resonance in fMRI. *IEEE transactions on medical imaging*, 36(6):1326–1336, 2017.
- [85] J.-p. Huang, L.-k. Zhu, L.-h. Wang, and W.-l. Song. Compressed sensing MRI using sparsity averaging and FISTA. *Applied Magnetic Resonance*, 48(8):749–760, 2017.
- [86] P. J. Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992.
- [87] P. J. Huber. *Robust statistics*, volume 523. John Wiley & Sons, 2004.
- [88] W. Huh, J. Fessler, and A. Samsonov. Water-fat decomposition with regularized field map. In *Proc. Intl. Soc. Mag. Reson. Med.*, volume 16, page 1382, 2008.
- [89] P. Irarrazabal, C. H. Meyer, D. G. Nishimura, and A. Macovski. Inhomogeneity correction using an estimated linear field map. *Magnetic resonance in medicine*, 35(2):278–282, 1996.
- [90] M. Jacob and B. P. Sutton. Algebraic decomposition of fat and water in MRI. *IEEE transactions on medical imaging*, 28(2):173–184, 2008.
- [91] M. Jenkinson, P. Bannister, M. Brady, and S. Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841, 2002.
- [92] P. Jezzard and R. S. Balaban. Correction for geometric distortion in echo planar images from B_0 field variations. *Magnetic resonance in medicine*, 34(1):65–73, 1995.
- [93] R. Jones, O. Haraldseth, T. Müller, P. Rinck, and A. Øksendal. K-space substitution: a novel dynamic imaging technique. *Magnetic resonance in medicine*, 29(6):830–834, 1993.
- [94] H. Jung, K. Sung, K. S. Nayak, E. Y. Kim, and J. C. Ye. k-t FOCUSS: a general compressed sensing framework for high resolution dynamic MRI. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 61(1):103–116, 2009.

- [95] H. Jung and J. C. Ye. Performance evaluation of accelerated functional MRI acquisition using compressed sensing. In *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 702–705. IEEE, 2009.
- [96] Y. M. Kadah and X. Hu. Algebraic reconstruction for magnetic resonance imaging under B_0 inhomogeneity. *IEEE Transactions on Medical Imaging*, 17(3):362–370, 1998.
- [97] S. Kamesh Iyer, T. Tasdizen, D. Likhite, and E. DiBella. Split bregman multicoil accelerated reconstruction technique: A new framework for rapid reconstruction of cardiac perfusion MRI. *Medical physics*, 43(4):1969–1981, 2016.
- [98] N. Kang, J. Zhang, and E. S. Carlson. Performance of ILU preconditioning techniques in simulating anisotropic diffusion in the human brain. *Future Generation Computer Systems*, 20(4):687–698, 2004.
- [99] L. Kasper, S. Bollmann, A. O. Diaconescu, C. Hutton, J. Heinzle, S. Iglesias, T. U. Hauser, M. Sebold, Z.-M. Manjaly, K. P. Pruessmann, et al. The PhysIO toolbox for modeling physiological noise in fMRI data. *Journal of neuroscience methods*, 276:56–72, 2017.
- [100] D. B. Keator, T. G. van Erp, J. A. Turner, G. H. Glover, B. A. Mueller, T. T. Liu, J. T. Voyvodic, J. Rasmussen, V. D. Calhoun, H. J. Lee, et al. The function biomedical informatics research network data repository. *Neuroimage*, 124:1074–1079, 2016.
- [101] D. S. Kershaw. The incomplete Cholesky – conjugate gradient method for the iterative solution of systems of linear equations. *Journal of computational physics*, 26(1):43–65, 1978.
- [102] D. Kim and J. A. Fessler. Optimized first-order methods for smooth convex minimization. *Mathematical programming*, 159(1-2):81–107, 2016.
- [103] D. Kim and J. A. Fessler. On the convergence analysis of the optimized gradient method. *Journal of optimization theory and applications*, 172(1):187–205, 2017.
- [104] D. Kim and J. A. Fessler. Adaptive restart of the optimized gradient method for convex optimization. *Journal of Optimization Theory and Applications*, 178(1):240–263, 2018.
- [105] K. Koolstra, J. Van Gemert, P. Börnert, A. Webb, and R. Remis. Accelerating compressed sensing in parallel imaging reconstructions using an efficient circulant preconditioner for cartesian trajectories. *Magnetic resonance in medicine*, 81(1):670–685, 2019.
- [106] L.-D. Kuang, Q.-H. Lin, X.-F. Gong, F. Cong, Y.-P. Wang, and V. D. Calhoun. Shift-invariant canonical polyadic decomposition of complex-valued multi-subject fMRI data with a phase sparsity constraint. *IEEE transactions on medical imaging*, 39(4):844–853, 2019.

- [107] S.-H. Lai and M. Fang. A dual image approach for bias field correction in magnetic resonance imaging. *Magnetic resonance imaging*, 21(2):121–125, 2003.
- [108] F. Lam and B. P. Sutton. Intravoxel B_0 inhomogeneity corrected reconstruction using a low-rank encoding operator. *Magnetic Resonance in Medicine*, 84(2):885–894, 2020.
- [109] F. Lam, B. Zhao, Y. Liu, Z.-P. Liang, M. Weiner, and N. Schuff. Accelerated fMRI using low-rank model and sparsity constraints. *Proc. Int. Soc. Magn. Reson. Med. Salt Lake City*, page 2013, 2013.
- [110] K. Lange, D. R. Hunter, and I. Yang. Optimization transfer using surrogate objective functions. *Journal of computational and graphical statistics*, 9(1):1–20, 2000.
- [111] M. Le and J. A. Fessler. Efficient, convergent SENSE MRI reconstruction for nonperiodic boundary conditions via tridiagonal solvers. *IEEE transactions on computational imaging*, 3(1):11–21, 2016.
- [112] M. H. Lee, C. D. Smyser, and J. S. Shimony. Resting-state fMRI: a review of methods and clinical applications. *American Journal of Neuroradiology*, 34(10):1866–1872, 2013.
- [113] X. Li, X. Ma, L. Li, Z. Zhang, X. Zhang, Y. Tong, L. Wang, S. Song, and H. Guo. Dual-TRACER: high resolution fMRI with constrained evolution reconstruction. *Neuroimage*, 164:172–182, 2018.
- [114] C. Y. Lin and J. A. Fessler. Accelerated methods for low-rank plus sparse image reconstruction. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 48–51. IEEE, 2018.
- [115] C. Y. Lin and J. A. Fessler. Efficient dynamic parallel mri reconstruction for the low-rank plus sparse model. *IEEE transactions on computational imaging*, 5(1):17–26, 2018.
- [116] C. Y. Lin and J. A. Fessler. Efficient regularized field map estimation in 3D MRI. *IEEE Transactions on Computational Imaging*, 6:1451–1458, 2020.
- [117] C. Y. Lin, D. C. Noll, and J. A. Fessler. A temporal model for task-based functional MR images. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 1035–1038. IEEE, 2020.
- [118] M. A. Lindquist, J. M. Loh, L. Y. Atlas, and T. D. Wager. Modeling the hemodynamic response function in fMRI: efficiency, bias and mis-modeling. *Neuroimage*, 45(1):S187–S198, 2009.
- [119] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob. Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR. *IEEE transactions on medical imaging*, 30(5):1042–1054, 2011.

- [120] S. G. Lingala and M. Jacob. Blind compressive sensing dynamic MRI. *IEEE transactions on medical imaging*, 32(6):1132–1145, 2013.
- [121] S. G. Lingala, M. Nadar, C. Ched’Hotel, L. Zhang, and M. Jacob. Unified reconstruction and motion estimation in cardiac perfusion MRI. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 65–68. IEEE, 2011.
- [122] Y. Liu and J. P. Haldar. NAPALM: An algorithm for MRI reconstruction with separate magnitude and phase regularization. *Proceedings of the International Society for Magnetic Resonance in Medicine, Montreal*, 4764, 2019.
- [123] W. Lu and B. A. Hargreaves. Multiresolution field map estimation using golden section search for water-fat separation. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 60(1):236–244, 2008.
- [124] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- [125] M. Lustig and J. M. Pauly. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic resonance in medicine*, 64(2):457–471, 2010.
- [126] D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A. Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, 2013.
- [127] T. A. Manteuffel. An incomplete factorization technique for positive definite linear systems. *Mathematics of computation*, 34(150):473–497, 1980.
- [128] A. Matakos and J. A. Fessler. Joint estimation of image and fieldmap in parallel MRI using single-shot acquisitions. In *2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 984–987. IEEE, 2010.
- [129] L. B. Montefusco, D. Lazzaro, S. Papi, and C. Guerrini. A fast compressed sensing approach to 3D MR image reconstruction. *IEEE transactions on medical imaging*, 30(5):1064–1075, 2010.
- [130] M. M. Monti. Statistical analysis of fMRI time-series: a critical review of the GLM approach. *Frontiers in human neuroscience*, 5:28, 2011.
- [131] M. J. Muckley, S. J. Peltier, D. C. Noll, and J. A. Fessler. Model-based reconstruction for physiological noise correction in functional MRI. In *Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA*, 2013.

- [132] K. Murphy, R. M. Birn, D. A. Handwerker, T. B. Jones, and P. A. Bandettini. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *Neuroimage*, 44(3):893–905, 2009.
- [133] Y. Nesterov. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Sov. Math. Dokl*, volume 27.
- [134] G.-C. Ngo, J. L. Holtrop, M. Fu, F. Lam, and B. P. Sutton. High temporal resolution functional MRI with partial separability model. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 7482–7485. IEEE, 2015.
- [135] H. Nguyen and G. Glover. Field-corrected imaging for sparsely-sampled fMRI by exploiting low-rank spatiotemporal structure. In *Proceedings of the 22nd Annual Meeting of ISMRM, Milan Italy*, 2014.
- [136] H. M. Nguyen and G. H. Glover. A modified generalized series approach: application to sparsely sampled fMRI. *IEEE Transactions on Biomedical Engineering*, 60(10):2867–2877, 2013.
- [137] T. E. Nichols and A. P. Holmes. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human brain mapping*, 15(1):1–25, 2002.
- [138] J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [139] D. C. Noll, J. A. Fessler, and B. P. Sutton. Conjugate phase MRI reconstruction with spatially variant sample density correction. *IEEE transactions on medical imaging*, 24(3):325–336, 2005.
- [140] D. C. Noll, C. R. Genovese, L. E. Nystrom, A. L. Vazquez, S. D. Forman, W. F. Eddy, and J. D. Cohen. Estimating test-retest reliability in functional MR imaging II: Application to motor and cognitive activation studies. *Magnetic Resonance in Medicine*, 38(3):508–517, 1997.
- [141] D. C. Noll, C. H. Meyer, J. M. Pauly, D. G. Nishimura, and A. Macovski. A homogeneity correction method for magnetic resonance imaging with time-varying gradients. *IEEE transactions on medical imaging*, 10(4):629–637, 1991.
- [142] S. Ogawa, T.-M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *proceedings of the National Academy of Sciences*, 87(24):9868–9872, 1990.
- [143] S. Ogawa, R. Menon, D. Tank, S. Kim, H. Merkle, J. Ellermann, and K. Ugurbil. Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. a comparison of signal characteristics with a biophysical model. *Biophysical journal*, 64(3):803–812, 1993.

- [144] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89(13):5951–5955, 1992.
- [145] V. T. Olafsson, D. C. Noll, and J. A. Fessler. Fast joint reconstruction of dynamic R_2^* and field maps in functional MRI. *IEEE transactions on medical imaging*, 27(9):1177–1188, 2008.
- [146] V. T. Olafsson, D. C. Noll, and J. A. Fessler. Fast spatial resolution analysis of quadratic penalized least-squares image reconstruction with separate real and imaginary roughness penalty: Application to fMRI. *IEEE transactions on medical imaging*, 37(2):604–614, 2017.
- [147] G. Ongie, J. Shi, and J. A. Fessler. Efficient computation of regularized field map estimates in 3D. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 700–703. IEEE, 2017.
- [148] R. Otazo, E. Candes, and D. K. Sodickson. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. *Magnetic resonance in medicine*, 73(3):1125–1136, 2015.
- [149] R. Otazo, A. Franco, J. Chen, C. Marmar, and F. Boada. Low-rank plus sparse (L+ S) decomposition for separation of subsampled physiological noise in fMRI. In *Proceedings of the 21st Annual Meeting of the Organization for Human Brain Mapping (OHBM)*, volume 1690, page 2015, 2015.
- [150] B. O’donoghue and E. Candes. Adaptive restart for accelerated gradient schemes. *Foundations of computational mathematics*, 15(3):715–732, 2015.
- [151] N. Parikh, S. Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- [152] S. Peltier and D. Noll. Systematic noise compensation for simultaneous multislice acquisition using rosette trajectories (SMART). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 41(5):1073–1076, 1999.
- [153] V. Perlberg, P. Bellec, J.-L. Anton, M. Péligrini-Issac, J. Doyon, and H. Benali. CORSICA: correction of structured noise in fMRI by automatic identification of ICA components. *Magnetic resonance imaging*, 25(1):35–46, 2007.
- [154] A. Y. Petrov, M. Herbst, and V. A. Stenger. Improving temporal resolution in fMRI using a 3D spiral acquisition and low rank plus sparse (L+S) reconstruction. *Neuroimage*, 157:660–674, 2017.
- [155] J. G. Pipe. Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging. *Magnetic Resonance in Medicine*:

An Official Journal of the International Society for Magnetic Resonance in Medicine, 42(5):963–969, 1999.

- [156] J. G. Pipe and P. Menon. Sampling density compensation in mri: rationale and an iterative numerical solution. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 41(1):179–186, 1999.
- [157] 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*, 3(R1):35–43, 1969.
- [158] J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*, 59(3):2142–2154, 2012.
- [159] J. D. Power, B. L. Schlaggar, and S. E. Petersen. Recent progress and outstanding issues in motion correction in resting state fMRI. *Neuroimage*, 105:536–551, 2015.
- [160] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 42(5):952–962, 1999.
- [161] S. Ramani and J. A. Fessler. Parallel MR image reconstruction using augmented Lagrangian methods. *IEEE Transactions on Medical Imaging*, 30(3):694–706, 2010.
- [162] V. Rasche, R. W. D. Boer, D. Holz, and R. Proksa. Continuous radial data acquisition for dynamic MRI. *Magnetic resonance in medicine*, 34(5):754–761, 1995.
- [163] S. Ravishankar, B. E. Moore, R. R. Nadakuditi, and J. A. Fessler. Low-rank and adaptive sparse signal (LASSI) models for highly accelerated dynamic imaging. *IEEE transactions on medical imaging*, 36(5):1116–1128, 2017.
- [164] S. Ravishankar, J. C. Ye, and J. A. Fessler. Image reconstruction: From sparsity to data-adaptive methods and machine learning. *Proceedings of the IEEE*, 108(1):86–109, 2019.
- [165] 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. *Magnetic Resonance in Medicine*, 39(2):328–330, 1998.
- [166] S. B. Reeder, J. H. Brittain, T. M. Grist, and Y.-F. Yen. Least-squares chemical shift separation for ^{13}C metabolic imaging. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 26(4):1145–1152, 2007.

- [167] S. B. Reeder, A. R. Pineda, Z. Wen, A. Shimakawa, H. Yu, J. H. Brittain, G. E. Gold, C. H. Beaulieu, and N. J. Pelc. Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): application with fast spin-echo imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 54(3):636–644, 2005.
- [168] S. B. Reeder, Z. Wen, H. Yu, A. R. Pineda, G. E. Gold, M. Markl, and N. J. Pelc. Multicoil Dixon chemical species separation with an iterative least-squares estimation method. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 51(1):35–45, 2004.
- [169] S. F. Roohi, D. Zonoobi, A. A. Kassim, and J. L. Jaremko. Multi-dimensional low rank plus sparse decomposition for reconstruction of under-sampled dynamic MRI. *Pattern Recognition*, 63:667–679, 2017.
- [170] D. B. Rowe and B. R. Logan. A complex way to compute fMRI activation. *Neuroimage*, 23(3):1078–1092, 2004.
- [171] A. Samsonov, F. Liu, and J. V. Velikina. Resolving estimation uncertainties of chemical shift encoded fat-water imaging using magnetization transfer effect. *Magnetic resonance in medicine*, 82(1):202–212, 2019.
- [172] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert. A deep cascade of convolutional neural networks for MR image reconstruction. In *International Conference on Information Processing in Medical Imaging*, pages 647–658. Springer, 2017.
- [173] E. Schneider and G. Glover. Rapid in vivo proton shimming. *Magnetic Resonance in Medicine*, 18(2):335–347, 1991.
- [174] K. Sekihara, S. Matsui, and H. Kohno. NMR imaging for magnets with large nonuniformities. *IEEE transactions on medical imaging*, 4(4):193–199, 1985.
- [175] K. Setsompop, B. A. Gagoski, J. R. Polimeni, T. Witzel, V. J. Wedeen, and L. L. Wald. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic resonance in medicine*, 67(5):1210–1224, 2012.
- [176] S. D. Sharma, N. S. Artz, D. Hernando, D. E. Horng, and S. B. Reeder. Improving chemical shift encoded water-fat separation using object-based information of the magnetic field inhomogeneity. *Magnetic Resonance in Medicine*, 73(2):597–604, 2015.
- [177] S. D. Sharma, H. H. Hu, and K. S. Nayak. Accelerated water-fat imaging using restricted subspace field map estimation and compressed sensing. *Magnetic Resonance in Medicine*, 67(3):650–659, 2012.

- [178] S. D. Sharma, H. H. Hu, and K. S. Nayak. Chemical shift encoded water-fat separation using parallel imaging and compressed sensing. *Magnetic resonance in medicine*, 69(2):456–466, 2013.
- [179] Z. Shehzad, A. C. Kelly, P. T. Reiss, D. G. Gee, K. Gotimer, L. Q. Uddin, S. H. Lee, D. S. Margulies, A. K. Roy, B. B. Biswal, et al. The resting brain: unconstrained yet reliable. *Cerebral cortex*, 19(10):2209–2229, 2009.
- [180] V. Singh, A. H. Tewfik, and D. B. Riss. Under-sampled functional MRI using low-rank plus sparse matrix decomposition. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 897–901. IEEE, 2015.
- [181] D. K. Sodickson and W. J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magnetic resonance in medicine*, 38(4):591–603, 1997.
- [182] V. A. Stenger, S. Peltier, F. E. Boada, and D. C. Noll. 3D spiral cardiac/respiratory ordered fMRI data acquisition at 3 tesla. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 41(5):983–991, 1999.
- [183] B. Sutton, D. Noll, and J. Fessler. Compensating for within-voxel susceptibility gradients in BOLD fMRI. In *Proc. Intl. Soc. Mag. Res. Med.*, page 349, 2004.
- [184] B. P. Sutton, D. C. Noll, and J. A. Fessler. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE transactions on medical imaging*, 22(2):178–188, 2003.
- [185] B. P. Sutton, D. C. Noll, and J. A. Fessler. Dynamic field map estimation using a spiral-in/spiral-out acquisition. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 51(6):1194–1204, 2004.
- [186] A. B. Taylor, J. M. Hendrickx, and F. Glineur. Exact worst-case performance of first-order methods for composite convex optimization. *SIAM Journal on Optimization*, 27(3):1283–1313, 2017.
- [187] C. G. Thomas, R. A. Harshman, and R. S. Menon. Noise reduction in BOLD-based fMRI using component analysis. *Neuroimage*, 17(3):1521–1537, 2002.
- [188] B. Trémouhéac, N. Dikaios, D. Atkinson, and S. R. Arridge. Dynamic MR image reconstruction–separation from undersampled (k, t) -space via low-rank plus sparse prior. *IEEE transactions on medical imaging*, 33(8):1689–1701, 2014.
- [189] J. Tsao and Y. Jiang. Hierarchical IDEAL: fast, robust, and multiresolution separation of multiple chemical species from multiple echo times. *Magnetic resonance in medicine*, 70(1):155–159, 2013.

- [190] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magnetic resonance in medicine*, 71(3):990–1001, 2014.
- [191] M. Unser, P. Thevenaz, and L. Yaroslavsky. Convolution-based interpolation for fast, high-quality rotation of images. *IEEE Transactions on image processing*, 4(10):1371–1381, 1995.
- [192] M. Usman, L. Kakkar, A. Matakos, A. Kirkham, S. Arridge, and D. Atkinson. Joint B_0 and image estimation integrated with model based reconstruction for field map update and distortion correction in prostate diffusion MRI. *Magnetic Resonance Imaging*, 65:90–99, 2020.
- [193] G. Vaillant, C. Prieto, C. Kolbitsch, G. Penney, and T. Schaeffter. Retrospective rigid motion correction in k-space for segmented radial MRI. *IEEE transactions on medical imaging*, 33(1):1–10, 2013.
- [194] J. J. Van Vaals, M. E. Brummer, W. Thomas Dixon, H. H. Tuithof, H. Engels, R. C. Nelson, B. M. Gerety, J. L. Chezmar, and J. A. Den Boer. “keyhole” method for accelerating imaging of contrast agent uptake. *Journal of Magnetic Resonance Imaging*, 3(4):671–675, 1993.
- [195] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin. Adaptive reconstruction of phased array MR imagery. *Magnetic Resonance in Medicine*, 43(5):682–690, 2000.
- [196] G. Wang, T. Luo, J.-F. Nielsen, D. C. Noll, and J. A. Fessler. B-spline parameterized joint optimization of reconstruction and k-space trajectories (BJORK) for accelerated 2D MRI. *arXiv preprint arXiv:2101.11369*, 2021.
- [197] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [198] M. Weiger, K. P. Pruessmann, R. Österbauer, P. Börnert, P. Boesiger, and P. Jezzard. Sensitivity-encoded single-shot spiral imaging for reduced susceptibility artifacts in BOLD fMRI. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 48(5):860–866, 2002.
- [199] L. Weizman, K. L. Miller, Y. C. Eldar, and M. Chiew. PEAR: Periodic and fixed rank separation for fast fMRI. *Medical physics*, 44(12):6166–6182, 2017.
- [200] B. Wen, S. Ravishankar, L. Pfister, and Y. Bresler. Transform learning for magnetic resonance image reconstruction: From model-based learning to building neural networks. *IEEE Signal Processing Magazine*, 37(1):41–53, 2020.

- [201] C. Windischberger, S. Robinson, A. Rauscher, M. Barth, and E. Moser. Robust field map generation using a triple-echo acquisition. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 20(4):730–734, 2004.
- [202] B. Wohlberg. ADMM penalty parameter selection by residual balancing. *arXiv preprint arXiv:1704.06209*, 2017.
- [203] Z. Xu, M. Figueiredo, and T. Goldstein. Adaptive ADMM with spectral penalty parameter selection. In *Artificial Intelligence and Statistics*, pages 718–727. PMLR, 2017.
- [204] X. Ye, Y. Chen, W. Lin, and F. Huang. Fast MR image reconstruction for partially parallel imaging with arbitrary k -space trajectories. *IEEE Transactions on Medical Imaging*, 30(3):575–585, 2010.
- [205] D. T. B. Yeo, J. A. Fessler, and B. Kim. Motion robust magnetic susceptibility and field inhomogeneity estimation using regularized image restoration techniques for fMRI. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 991–998. Springer, 2008.
- [206] H. Yu, S. B. Reeder, A. Shimakawa, J. H. Brittain, and N. J. Pelc. Field map estimation with a region growing scheme for iterative 3-point water-fat decomposition. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 54(4):1032–1039, 2005.
- [207] H. Yu, A. Shimakawa, C. A. McKenzie, E. Brodsky, J. H. Brittain, and S. B. Reeder. Multiecho water-fat separation and simultaneous R_2^* estimation with multifrequency fat spectrum modeling. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 60(5):1122–1134, 2008.
- [208] E. Yudilevich and H. Stark. Spiral sampling in magnetic resonance imaging—the effect of inhomogeneities. *IEEE transactions on medical imaging*, 6(4):337–345, 1987.
- [209] F. Zhao, D. C. Noll, J.-F. Nielsen, and J. A. Fessler. Separate magnitude and phase regularization via compressed sensing. *IEEE transactions on medical imaging*, 31(9):1713–1723, 2012.
- [210] K. Zhou, M. Zaitsev, and S. Bao. Reliable two-dimensional phase unwrapping method using region growing and local linear estimation. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 62(4):1085–1090, 2009.
- [211] M. V. Zibetti and A. R. De Pierro. Improving compressive sensing in MRI with separate magnitude and phase priors. *Multidimensional Systems and Signal Processing*, 28(4):1109–1131, 2017.