Title
Sparsity promoting optimization in quantum mechanical signal processing

Author
Compton, Ryan

Publication Date
2012

Peer reviewed|Thesis/dissertation

University of California
Los Angeles

Sparsity promoting optimization in quantum mechanical signal processing

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Ryan Compton

2013

Abstract of the Dissertation

Sparsity promoting optimization in quantum mechanical signal processing

by

Ryan Compton

Doctor of Philosophy in Mathematics

University of California, Los Angeles, 2013

Professor Chris Anderson, Chair

Signals describing the energy levels of quantum mechanical systems are, by definition, sparse in the energy domain. Processing these signals via sparsity promoting methods is thus reasonable and, as this dissertation argues, valuable.

Quantum mechanical energy levels are determined experimentally through NMR spectroscopy where noise, peak blurring, and long experiment times impede progress. We show how $l1$-penalized optimization can lead to improved signal quality and reduce data acquisition time in NMR spectroscopy.

Quantum mechanical signal processing is central to MRI reconstruction. MRI data acquisition and reconstruction is highly time-consuming and expensive. We provide a fast converging algorithm based on minimizing a combination total variation and framelet norms which produces high-quality images from undersampled MRI data.

In the field of numerical analysis, all the eigenvalues of a Hermitian matrix may be computed by simulating a fictitious quantum dynamical system with Hamiltonian corresponding to the matrix of interest and then determining the energy levels of this fictitious system. By determining the energy levels with $l1$-penalized optimization we show that the number of simulation steps can be significantly reduced.

Quantum mechanical systems have spatial components as well. When the spatial do-

main is partitioned according to the location of potential wells, one often finds low-energy wavefunctions tend to localize within the confines of each partition. For a given partition, the energy levels of its corresponding localized wavefunctions often make up only a small fraction of the complete range of energy levels. For situations where only a few eigenpairs are sought we introduce a "projection-correction" method allowing us to efficiently compute only the low-energy eigenpairs which localize within a given spatial partition. In contrast to standard methods for eigenvalue computation which specify only a part of the spectrum, our method also allows one to isolate regions of space where prior information on eigenfunction locality is known.

The dissertation of Ryan Compton is approved.

Luminita Vese

Stan Osher

Louis Bouchard

Chris Anderson, Committee Chair

University of California, Los Angeles

2013

Table of Contents

List of Figures

List of Tables

Vita

2006	B.A. (Mathematics/Physics), New College of Florida.
2008	M.S. (Mathematics), UCLA, Los Angeles, California.
2006–2011	Teaching Assistant, Department of Mathematics, UCLA.
2011–2012	Research Assistant, Department of Mathematics, UCLA.

Publications

Ryan Compton, Nanette Jarenwattananon, Louis Bouchard and Stanley Osher, "Model based compressed sensing reconstruction of nonuniformly sampled NMR signals", *to be submitted*.

Ryan Compton, Stanley Osher and Louis Bouchard, "Hybrid regularization for MRI reconstruction with static field inhomogeneity correction", *IEEE International Symposium on Biomedical Imaging*, May 2012 (Journal version recently accepted into "Inverse Problems and Imaging").

Ryan Compton, Hankyu Moon and Tsai-Ching Lu, "Catastrophe prediction via estimated network autocorrelation", *WIN Workshop on Information in Networks*, September 2011.

Ryan Compton, "A sparse spectral method for Hamiltonian eigensystems", *technical report*, January 2011

CHAPTER 1

Introduction

In recent years, the realization that many signals of interest can be well-represented with only a few nonzero values has become central to a wide range of signal processing applications. In the field of quantum mechanics, for example, theory guarantees us that the energy levels of Hamiltonians in bounded domains are confined to a few discrete values out of a continuous range of possibilities. Obtaining signals describing quantum energy levels is of fundamental importance to much of physics. In this thesis, we will apply a commonly used sparsity promoting optimization technique, $l1$-minimization, to analyze and reconstruct signals originating from quantum mechanical systems.

Our first applications focus on signals generated thorough nuclear magnetic resonance (NMR). NMR spectroscopy experiments allow chemists to determine properties of organic molecules by examining differences in energy levels between spin up and spin down states. In a standard NMR experiment energy-domain information is obtained via a Fourier transform of an observed time-domain free induction decay (FID) signal. While commonly used and widely applicable, the Fourier transform based approach to NMR signal reconstruction disregards the fact that NMR signals are typically composed of a few Lorentzian lines when represented in the energy domain. In this thesis, we take advantage of this addtional structural knowledge by replacing the Fourier transform with a sparsity promoting, $l1$-penalized, optimization.

To be specific, we advocate a change of basis of the form:

$$\underset{u,a}{argmin}\ |u|_1\ subject\ to\ \|e^{-t/a}R\mathcal{F}^{-1}u - f\| \leq \mu \tag{1.1}$$

where f is the observed time-domain signal, $\mathcal{F}^{-1}$ is an inverse DFT matrix, R is a diagonal

matrix with 1s and 0s placed uniformly at random on the diagonal, $e^{-t/a}$ models spin-spin relaxation time, and μ is the standard deviation of additive Gaussian receiver noise. We will demonstrate that solving eq. (1.1) in place of a Fourier transform leads to an improved signal to noise ratio and allows us to recover signals from undersampled time-domain data.

Previous work on $l1$-penalized minimization for NMR reconstruction exists. Novel in this thesis is the inclusion of the damping factor, $e^{-t/a}$, in eq. (1.1) which is neccessary for an accurate model of spin decay. We find empirically that eq. (1.1) can lead to improved signal quality when compared with an approach where the damping is neglected. We demonstrate these results on real and simulated NMR datasets of varying dimensionality in chapter 2.

Related to NMR spectroscopy, magnetic resonance imaging (MRI) classically requires a discrete Fourier transform to produce images from recorded k-space data. Problematic with this approach is that k-space must be fully sampled and off-resonance effects are ignored when an image is reconstructed with a discrete Fourier transform. Recording all of k-space requires that patients remain in a confined scanner for extended periods of time and reduces the number of scans that can be taken in one day. Overcoming this limitation is a major achievement of $l1$-penalized image reconstruction algorithms.

Current compressed sensing MRI reconstruction approaches, however, neglect the fact that magnetic fields used for imaging are often imperfect due to differences in the magnetic suseptibility of tissues being imaged. Abrupt changes in magentic suseptibility are often found near air/tissue and gray/white matter interfaces. Failing to correct for the resulting magnetic field inhomogenities leads to image blurring and distortion. Modeling field inhomogeneites during the reconstruction stage leads to higher quality images but requires that we solve an optimization problem involving a non-Fourier operator. The field-corrected MRI reconstruction problem is then:

$$\underset{u}{argmin}\ J(u)\ subject\ to\ \|ER\mathcal{F}^{-1}u - f\| \leq \mu \tag{1.2}$$

where E models the inhomogeneity effects and J is some form of sparsity-promoting regularization. In this thesis, we study and advocate a combined regularization based on total

variation and framelets:

$$J(u) = |\nabla u| + |Fu| \tag{1.3}$$

When reconstructing MR images with such a J we find that high quality images can be produced more rapidly than with a regularization based on framelets or total variation alone. This is the topic of chapter 3.

Outside of experimental settings, purely Fouier sampling matrices are readliy available. Simulating quantum dynamical systems allows one to efficiently compute eigenvalues of Hermitian matricies by analyzing signals generated from the simulation. Referred to as "spectral methods" for eigenvalue computations, these techniques traditionally require a Fourier transform of the time autocorrelation of the simulated wave function. The transformed signal is known to be sparse, with peak locations determining eigenvalues of the matrix at hand.

Simulating the wave function is the major computational burden. The burden is heavier when high-resoultion long-duration simulations are needed. Diffculties arise because high-resoultion long-duration simulations are needed to avoiding aliasing effects in the final spectrum. Similar to our approach for NMR spectroscopy, we propose a reduction in the number of time data points collected by replacing the Fourier transform with an $l1$ minimization and invoking compressed sensing theory. We demonstrate exact spectra reconstruction from a small subset of time domain data in chapter 4.

Quantum mechanical systems have spatial components as well. When the spatial domain is partitioned according to the location of potential wells, one often finds low-energy wavefunctions tend to localize within the confines of each partition. For a given partition, the energy levels of its corresponding localized wavefunctions often make up only a small fraction of the complete range of energy levels. In chapter 5 we focus on efficiently computing only the low-energy eigenpairs which localize within a given spatial partition. In contrast to standard methods for eigenvalue computation which specify only a part of the spectrum, our method also allows one to isolate regions of space where prior information on eigenfunction locality is known.

The approach is essentially a two-step procedure. Given a Hermitian matrix, H, and a spatial region, $\Omega \subset \mathbb{R}^2$, where eigenfunctions are believed to localize, we first determine approximations to eigenfunctions localizing in Ω by finding eigenfunctions of the modified operator, $I - P_\Omega \frac{H-\lambda_{max}}{\lambda_{min}-\lambda_{min}} P_\Omega$, where P_Ω is a linear projection onto Ω. The accuracy of these approximations is problem dependant and varies with the choice of Ω and shape of the potential. This approximation is often sufficient when modest accuracy is acceptable, however, when highly accurate eigenpairs are sought or Ω was chosen poorly we refine our eigenpairs by solving a nonlinear "correction equation" via Newton's method or a Jacobi-Davidson approach. We apply our method to the computation of localized wavefunctions of Schrodinger operators with potentials originating in quantum dot simulations and find that significant computational savings are possible with this approach.

CHAPTER 2

Model based compressed sensing reconstruction of nonuniformly sampled NMR signals

The apparent sparsity and long acquisition times in NMR spectroscopic datasets suggest that experiment time may be reduced with an application of compressed sensing. In recent years, traditional, Fourier-based, compressed sensing has been the approach taken by several research groups [SF11] [KO11]. However, this presents a problem in that a purely Fourier-based approach to compressed sensing NMR spectra reconstruction neglects the line broadening which often results from short T_2-relaxation times or static field inhomogeneities.

In this chapter, we compensate for line broadening by simultaneously solving a semi-blind deconvolution in addition to reconstruction. We show that recovering NMR spectra from nonuniformly sampled data by solving the optimization,

$$\underset{u,a}{argmin}\ |u|_1 \ subject\ to\ \|e^{-t/a}R\mathcal{F}^{-1}u - f\| \leq \mu \tag{2.1}$$

where f is the observed time-domain signal, $\mathcal{F}^{-1}$ is an inverse DFT matrix, R is a diagonal matrix with 1s and 0s placed uniformly at random on the diagonal, $e^{-t/a}$ models spin-spin relaxation time, and μ is the standard deviation of additive Gaussian receiver noise. We will demonstrate that solving eq. (1.1) in place of a Fourier transform leads to an improved signal to noise ratio and allows us to recover signals from undersampled time-domain data.

This chapter is joint work with Louis Bouchard, Nanette Jarenwattananon, and Stanley Osher

2.1 Introduction

Reconstructing NMR spectra from noisy, nonuniformly sampled, time domain signals allows one to drastically reduce experiment time on existing magnetic hardware. Classically, NMR spectra are reconstructed from an observed free induction decay (FID) via a discrete Fourier transform. Direct use of the discrete Fourier transform, however, imposes strict requirements on the amount of FID data collected and neglects any prior knowledge one has of spectral structure. Uncertainty principles place limitations on the collected FID data in two ways: resolving closely spaced peaks requires the FID be recorded to a high final acquisition time, T, and obtaining the entire spectral range of interest requires that the spacing between points in the time domain, Δt, is small.

Taking into account the fact that the number of peaks in an NMR spectra is often much smaller than the number of frequency bins makes it possible to avoid these constraints and obtain accurate NMR spectra from undersampled FID data. Overcoming the constraint on T is possible when a desired spectral range is known in advance and measurements can be taken at high SNR by employing the method of filter-diagonalization [WN95] [Man01]. Outside of this situation, one may opt to record an FID for large T while keeping the number of recorded points small by nonuniformly varying Δt. Reconstruction of NMR spectra from noisy, nonuniformly sampled, FID data has been researched for several years now [MSK06]. Approaches based on maximum entropy regularization [MH08], multidimensional decomposition [JI06], and, recently, compressed sensing [SDH07] have been employed.

While the original theory of compressed sensing guarantees exact reconstruction of a signal which is sparse in an orthonormal basis, NMR signals are not properly sparse in any orthonormal basis as they are most accurately modeled as a linear combination of a few damped sinusoids [HS96]. In order to compensate for the damping factors when seeking a sparse representation, one must relax the orthogonality constraint and work with redundant dictionaries. Extensions of compressed sensing to redundant dictionaries exist [RS08] [CEN11]. However, these methods are not immediately applicable as the damping factors

(ie. the choice dictionary) are also unknown prior to the reconstruction process.

A well-known analogous problem occurs in field-corrected MRI [Fes10]. During MRI acquistion, static magnetic field inhomogeneities cause the recorded k-space data to be scaled undesireably. Images reconstructed from the corrupted k-space data are thus blurred and distorted. Modeling the inhomogeneities and iteratively solving a deconvolution in addition to the reconstruction problem leads to sharper final images [COB12].

In NMR, acquisition is done in the time domain and sparsity is often enforced in the frequency domain [QGC11], or in a wavelet domain [Dro07]. The time domain signal model is:

$$f(t) = \sum_{k=1}^{n_s} a_k e^{-t/T_{2,k}} e^{-i\gamma_k t} + \sigma(t) \tag{2.2}$$

The unknown quantities in eq. (2.2) are: the relative complex amplitudes, a_k, the decoherence times, $T_{2,k}$, and γ_k, which is proportional to the magnitude of the background field, the gyromagnetic ratio, and the chemical shift of the kth nucleus [Gra07]. The additive noise, $\sigma(t)$, is assumed to be white and Gaussian.

Intuitively, one would expect the spectrum to be sparse as the number of terms, n_s, is small and the acquisition time, T, is chosen so that the decoherence times, $T_{2,k}$, are relatively large. The resulting optimization problem for spectrum reconstruction is:

$$\underset{u}{argmin}\ |u|_1 \ subject\ to\ \|R\mathcal{F}^{-1}u - f\| \leq \mu \tag{2.3}$$

where R accounts for random undersampling [CRT06].

The problem with eq. (2.3) on NMR data is that the $l1$-norm will seek sparse approximations to Lorentzian lines and much of the effort in solving the optimization will go towards approximating side lobes. Current approaches to compressed sensing NMR take this approach [KO11] [HMW12] [Dro07] [SF11] [HLC] [QGC11].

Reconstructing using a basis of damped sinusoids, as in eq. (2.1), remedies this problem and leads to sparser solutions. The drawback to eq. (2.1) is increased computational cost and loss of convexity as both a and u must be solved for. In the following sections, we will

compare NMR signal reconstruction using eq. (2.1) and eq. (2.3). Our results indicate that compressed sensing approaches to NMR may be improved by including the damping factor in eq. (2.1) in the reconstruction.

2.2 Method

2.2.1 Signal reconstruction

Our goal is to reconstruct a complex-valued n-point NMR spectrum, $u(\xi)$, at high resolution from $m < n$ noisy measurements of an FID, $f(t)$. We draw m points from the time axis uniformly at random as this gives us a high probability of success when $l1$-based procedures are employed [CRT05].

The applicability of compressed sensing relies on how accurately we can approximate the spectra with a sparse signal. In the spectral domain, NMR signals are convolutions of delta functions with Lorentzians. The decay rate of the coefficients is thus asymptotically power-law with exponent 2,

$$|u|_{(k)} \leq C_{T_2} k^{-2} \tag{2.4}$$

where $|u|_{(k)}$ is the kth largest coefficient of $|u|$ and the coefficient depends on the decoherence time. The NMR spectra is thus *compressible* and by [CR06] we have the sparse approximation estimate for the solution, $u^\sharp$, to eq. (2.3):

$$\|u^\sharp - u_0\|_2 \leq C_{1,n_s}\mu + C_{2,n_s} C'_{T_2} n_s^{-3/2} \tag{2.5}$$

where u_0 is the exact, n_s-sparse, signal.

While this guarantees that compressed sensing NMR is possible, the coefficient C_{T_2} (which grows as the decoherence times shrink) may be so large that the error in eq. (2.5) is unacceptable. Empirically, we find that the line broadening can be so severe that a direct sparse approximation is inappropriate. A similar situation occurs in image processing where it has been observed that the wavelet coefficients of natural images decay according to a power law unless the resolution is exceedingly high [Gou01].

Sharpening the Lorentzian lines is possible by solving a semi-blind sparse deconvolution problem in addition to the reconstruction. In its most general form, we must solve

$$\underset{u,a(t)}{argmin}\ |u|_1 + \phi(a)\ subject\ to\ \|diag(e^{-t/a(t)})R\mathcal{F}^{-1}u - f\| \leq \mu \tag{2.6}$$

where ϕ takes into account prior knowledge on the decoherence times. The deconvolution in eq. (2.6) is highly ill-posed as we are required to learn the deconvolution kernel in addition to the spectrum from undersampled data. Reconstructing $a(t)$ as well as u thus requires more time domain data than solving for u alone. We have found that a simplification of homogeneous line broadening (i.e. $a(t) = const.$), allows us to sharpen reconstructed spectra and avoids many of the difficulties present when one is solving for a more general $a(t)$.

This leads us to the optimization

$$\underset{u,a}{argmin}\ |u|_1\ subject\ to\ \|e^{-t/a}R\mathcal{F}^{-1}u - f\| \leq \mu \tag{2.7}$$

which we solve by alternating between minimizations on u and a in alg. 1.

Algorithm 1: Alternating minimization for simultaneous reconstruction/deconvolution

Initialize: $u_0 = \vec{0}$, $a_0 > T$

while $\frac{|a_k - a_{k-1}|}{|a_k|} > tol$ **do**

$u_{k+1} = \underset{u}{argmin}\ |u|_1\ subject\ to\ \|e^{-t/a_k}R\mathcal{F}^{-1}u - f\| \leq \mu$

$a_{k+1} = \underset{a}{argmin}\ \|e^{-t/a}R\mathcal{F}^{-1}u_{k+1} - f\| \leq \mu$

end

While the a_{k+1} update is a differentiable single variable optimization problem, the u_{k+1} update is nontrivial. Several methods are available for solving the basis pursuit problem in the u_{k+1} update of alg. 1 [SDS01] [TG06] [Li]. The computational bottleneck is often finding the solution of a linear system involving the system matrix, $e^{-t/a}\mathcal{F}^{-1}$. This matrix has condition number $\kappa = e^{T/a}$ which may be large enough to be problematic. Recent work on conjugate-gradient based approaches to basis pursuit problems has led to methods which

are insensitive to the condition number of the system matrix [GS]. These newer optimization algorithms markedly reduce computation time when compared against the standard iterative soft threholding approach that is commonly found in the compressed sensing NMR literature [GS].

2.2.1.1 CGIST for basis pursuit

A standard method for solving the basis pursuit problem,

$$\mu|u| + \frac{1}{2}\|Au - f\|^2 \tag{2.8}$$

is "forward-backward splitting" (FBS) [Lions], also referred to in the literature as "iterative shrinkage/thresholding". While minimizing only the quadratic term in eq. (2.8) is straightforward, adding in the $l1$ term in eq. (2.8) leads to a non-differentiable optimization. FBS handles this by alternately taking a gradient descent step on the quadratic term and then seeking a sparse approximation to the result of the gradient update.

Algorithm 2: FBS for basis pursuit

Initialize: $u_0 = \vec{0}$

while $\frac{|u_{k+1}-u_k|}{|u_k|} > tol$ **do**

$\bar{u}_k = u_k + tA^T(Au_k - f)$

$u_{k+1} = \underset{u}{argmin}\ \mu|u| + \frac{1}{2t}\|u - \bar{u}_k\|$

end

As the matrix A has been removed from the update in line 2 of alg. 2 the problem becomes separable in each index of u and we have the closed form solution for u_{k+1}:

$$u_{k+1} = \frac{\bar{u}_k}{|\bar{u}_k|}\max\{\bar{u}_k - t\mu, 0\} = shrink(\bar{u}_k, t\mu). \tag{2.9}$$

When $sign(u_k) = sign(u_{k+1})$ it can be shown that one step of FBS is equivalent to a gradient descent step of eq. (2.8) [GS]. This is problematic as gradient descent may require many applications of A to converge when A is poorly conditioned. However, when the

signs of u_k and u_{k+1} agree we can rewrite the non-differentiable optimization eq. (2.8) as a constrained differentiable problem,

$$u_{k+1} = \underset{u}{argmin}\ \frac{1}{2}\|Au - f\|^2 + \langle u, s\rangle \ subject\ to\ Du = 0 \tag{2.10}$$

where D is a diagonal matrix such that $D_{ii} = 1$ if $u_{k,i} \neq 0$, and $D_{ii} = 0$ otherwise and $s = \mu sign(u_k)$. When the constraint in eq. (2.10) is satisfied we have $Du = u$ and we can rewrite the optimization as the unconstrained problem:

$$u_{k+1} = \underset{u}{argmin}\ \frac{1}{2}\|ADu - f\|^2 + \langle u, s\rangle. \tag{2.11}$$

The optimization eq. (2.11) is differentiable and, since u is expected to be sparse, has a low-rank system matrix. The format of eq. (2.11) is amenable to the "conjugate gradient partan" method of [Partan]. This method is guaranteed to converge when the number of iterations reaches the rank of AD leading to fast updates.

When the signs of u_k and u_{k+1} are not equal we update using FBS with an optimally chosen step size.

2.2.2 2D nonuniform sampling

The design of 2D NMR experiments makes it possible to reduce acquisition time by reducing the sampling rate in the indirect dimension only. The time required to collect data in the direct dimension is independent of the sampling rate.

Recall that a 2D NMR experiment is composed of repetitions of 1D experiments. Each 1D experiment can be broken down into four stages: excitation, evolution, mixing, and acquisition. During excitation, all nuclei are simultaneously tipped into the x-y plane with a 90° pulse. The nuclei are allowed to evolve for a period of time, t_1, until the mixing pulse is applied. Variations in the mixing stage account for type of NMR experiment (e.g. COSY, NOESY, etc.). Following the mixing pulse, the FID is recorded at n_2 points with sampling rate Δt_2. Repeating this experiment n_1 times provides us with a matrix, $f(t_1, t_2) \in \mathbb{R}^{n_1 \times n_2}$, whose DFT is composed of Lorentzian peaks and noise.

Assuming homogeneous line broadening, our 2D signal model is

$$f(t_1, t_2) = e^{-t_1/a_1} e^{-t_2/a_2} u(t_1, t_2)\mu(t_1) + \sigma(t_1, t_2) \tag{2.12}$$

where u is a sum of complex sinusoids and μ is t_1 noise. The t_1 noise is multiplicative in the time domain and thus convolutive in the spectral domain; it is also non-white and non-Gaussian [Gra07]. This bodes poorly for compressed sensing as the spectral domain blurring induced by $\mu(t_1)$ reduces sparsity. Furthermore, we have no parametric form for $\mu(t_1)$ so we can not include this information in our dictionary as was possible with the Lorentzian line broadening. We compensate for this by requiring more measurements than the original papers on compressed sensing advocated [CRT05].

We propose a reduction in acquisition time by selecting points from the t_1 dimension of $f(t_1, t_2)$ uniformly at random. The reconstruction procedure is then:

$$\underset{u,a_1,a_2}{argmin}\ |u|_1 \ subject\ to\ \|e^{-t_1/a_1} e^{-t_2/a_2} R\mathcal{F}^{-1} u - f\| \leq \mu \tag{2.13}$$

which can be solved with an alternating directions method as in alg. 1. It should be noted that, in principle, we can reconstruct the spectrum by taking a DFT in t_2 and then solving n_2 one-dimensional compressed sensing problems to fill in the missing t_1 data. However, since the spectrum is approximately sparse in both t_1 and t_2 an l_1 penalty on all dimensions is appropriate and leads to higher quality results as t_2-sparsity is completely neglected in the one-dimensional approach.

2.3 Numerical Results

Our reconstruction method is tested on real and simulated NMR datasets.

The CGIST library [GS] was used to solve eq. (2.3) and eq. (2.7) in all experiments. We set $\mu = 0.2\|f\|$ and stopped the outer loop of alg. 1 with $tol = .05$. We accelerate each u_{k+1} update by making an initial guess for the solution using u_k.

Numerical experiments were run in 64-bit Matlab on twelve cores of a dual hex core system compromised of two 2.67 GHz Intel Xeon CPUs (each with 12 MB of level 2 cache)

and 50 GB of RAM. Before plotting, all data is set to unit norm in order to make the comparisons clear.

2.3.1 1D experiments

Our simulated 1D dataset consisted of a 5-peak 1182-point spectra with unit time increments (i.e. $T = 1182$). The homogeneous decoherence time was $a = 236.4$. Uniform random downsampling to 20% of time domain data was performed before reconstruction. Computation time for a CS-based reconstruction was 0.2342 seconds, computation time for eq. (2.7) was 0.2943 seconds.

In fig. 2.1 we plot our simulated 1D dataset and its FFT. Signal decay is observable in fig. 2.1a which contributes to the line broadening in fig. 2.1b. The line broadening is not significant enough to disregard sparse signal processing methods. The result of uniform random downsampling in the time domain can be seen in fig. 2.1c. An FFT of the downsampled data leads to the poor spectrum found in fig. 2.1d.

In fig. 2.3 we plot the results of $l1$-based reconstruction experiments on our simulated 1D data. Traditional compressed sensing based methods (red) lead to an improvement over an FFT. The proposed method, however, leads to a significant improvement over compressed sensing. This is apparent in the close up of fig. 2.3b. Improved quality is to be expected as the proposed method takes into account all information that went into the simulation while traditional compressed sensing ignores line broadening and attempts to approximate the sidelobes as well as the peaks.

Our real 1D dataset consisted of a 49020-point C13 MERCAPTO 24C DMSO signal retrospectively downsampled to 10% of time domain data. Computation time for a CS-based reconstruction was 11.6120 seconds, computation time for eq. (2.7) was 82.7267 seconds.

In fig. 2.2 we plot the acquired data. Unlike our simulated data, this dataset shows substantial decay in the time domain. Line broadening is visible in fig. 2.2b. After downsampling, an FFT yields poor results in fig. 2.2d.

In fig. 2.4 we plot the results of $l1$-based reconstruction experiments on the actual NMR data. Here, $l1$-based reconstruction shows great improvement over the FFT, however, due to a non-zero baseline in the actual data, $l1$ methods struggle away from peak locations and produce several small spurious peaks rather than a non-zero baseline. Near peak locations, both methods perform more sensibly. Examnining the peak cluster near 4400 we can see again how the proposed reconstruction outperforms standard compressed sensing by not approximating sidelobes.

2.3.2 2D experiments

Our simulated 2D dataset consisted of a 10-peak 600×600-point spectra with unit time increments. The homogeneous decoherence times were $a_1 = 190$ and $a_2 = 600$ leading to a significantly blurred spectrum. Data was undersampled by a factor of 3 in the indirect dimension before reconstruction cf. fig. 2.5a.

The spectra from complete simulated data is shown in fig. 2.6a, peak broadening is more significant along the t_1-axis. Practical NMR downsampling is only possible along the t_1-axis. We adopt uniform random sampling along the t_1-axis according to the schedule in fig. 2.5a. As the t_1 sidelobes are more pronounced, traditional compressed sensing performs poorly along this axis fig. 2.6b. The proposed method alleviates this problem, fig. 2.6c shows reconstruction results with no sidelobe reconstruction problems.

For real data, we acquired a 2D COSY NMR spectrum of 2-ethyl-1-indanone at 2048×957 resolution and retrospectively downsampled the t_1-axis prior to our reconstruction experiments.

This data is contaminated with t_1-noise which is multiplicative in the time domain and thus convolutive in the spectral domain [Gra07]. The convolution leads to line broadening beyond what can be captured with an exponentially damped basis of sinusoids. In fig. 2.8 we can see that both reconstruction methods have a tendency to overestimate the t_1-noise. This worsens as we remove more data. Incorporating some form t_1-noise reduction into the

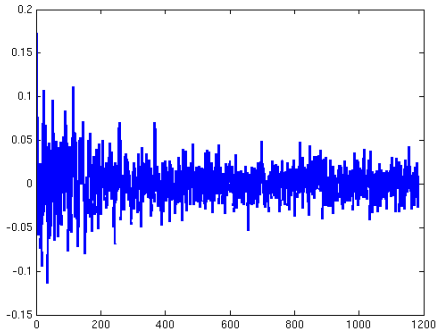

(a) Complete time domain data, real part

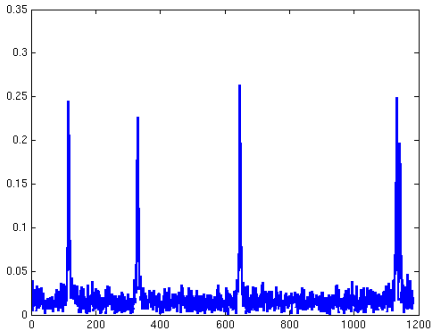

(b) FFT from complete data, magnitude

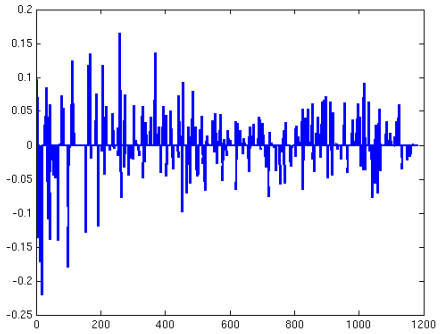

(c) Undersampled time domain data, real part

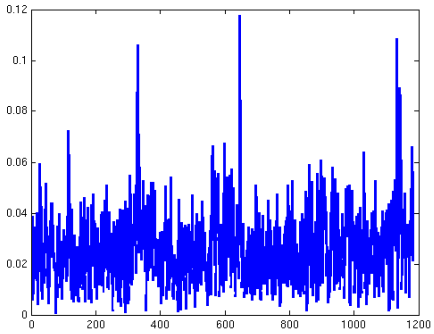

(d) FFT from undersampled data, magnitude

Figure 2.1: Simulated data used in our reconstruction experiment. The FID in fig. 2.1a decays slowly enough that the spectrum in fig. 2.1b is compressible. We downsample by 5x in fig. 2.1c by setting unrecorded coefficients to zero. An FFT of the undersampled data produced poor results, fig. 2.1d

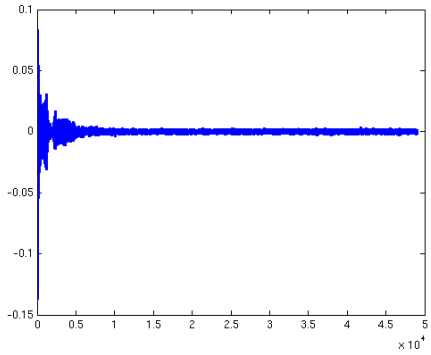

(a) Complete time domain data, real part

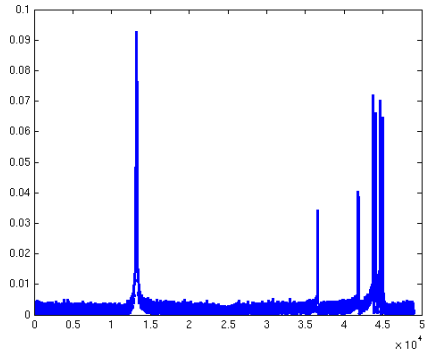

(b) FFT from complete data, magnitude

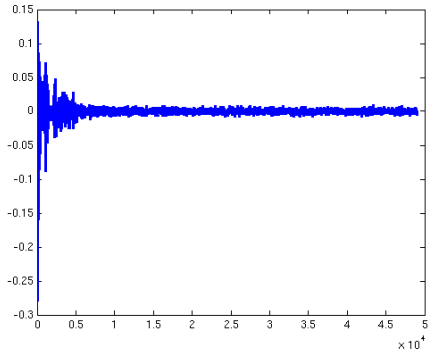

(c) Undersampled time domain data, real part

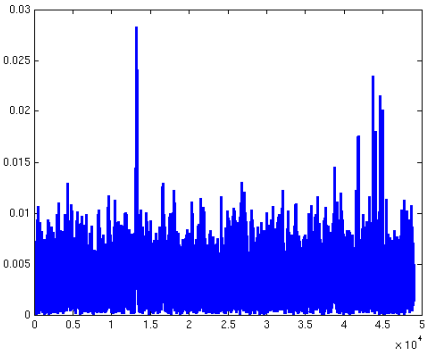

(d) FFT from undersampled data, magnitude

Figure 2.2: Actual data used in our reconstruction experiment. The FID in fig. 2.2a decays rather quickly, however the spectrum in fig. 2.2b still appears to be compressible. We downsample by 10x in fig. 2.2c by setting unrecorded coefficients to zero. An FFT of the undersampled data produced poor results, fig. 2.2d.

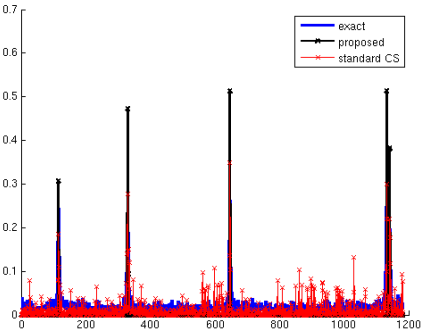

(a) Complete spectra

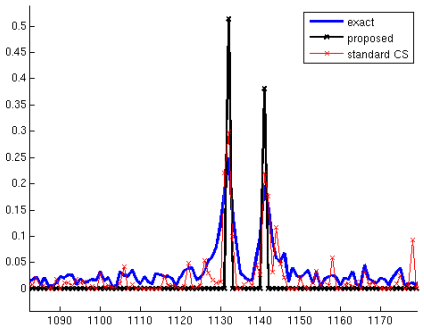

(b) Close up of the peak cluster near 1150

Figure 2.3: Reconstructions from undersampled data using eq. (2.3) (red) and eq. (2.7) (black). Due to low SNR and line broadening, the exact signal (blue) rarely reaches 0. Approximating this noise and broadening leads to poor signal reconstruction fig. 2.3. Zooming in on clustered peaks at the right side of the spectrum makes it easier to see how the standard compressed reconstruction approximates the side lobes while eq. (2.7) avoids this problem.

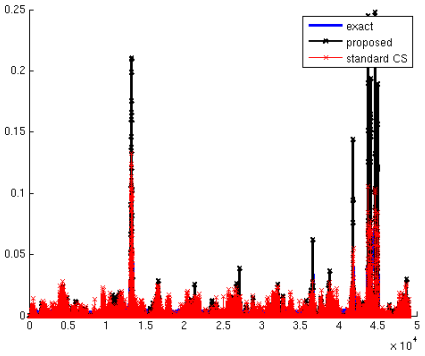

(a) Complete spectra

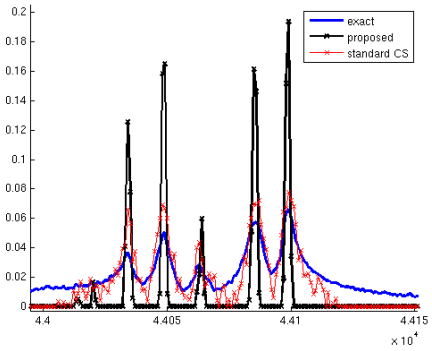

(b) Close up of the peak cluster near 4400

Figure 2.4: Reconstruction experiment on real C13 MERCAPTO 24C DMSO NMR data. Line broadening on this data is more severe than in the simulated example of fig. 2.3 and both algorithms struggle away from the peaks (fig. 2.4a). The clustered peaks on the right side of the spectrum are well approximated by eq. (2.7) while standard compressed sensing approximates the side lobes (fig. 2.4b). Both methods produce passable results from 10% of the time domain data.

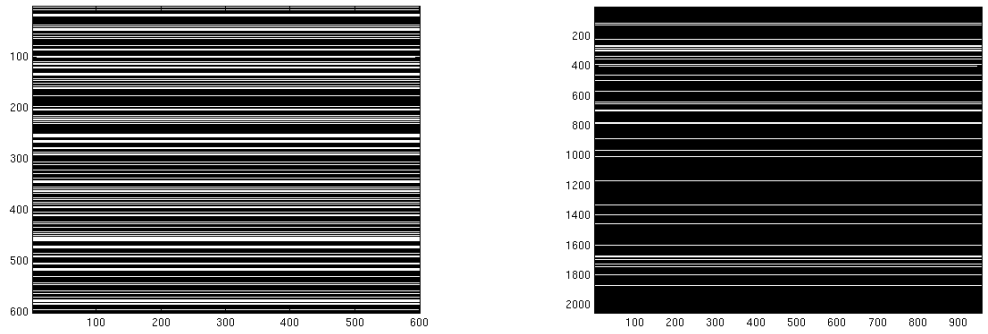

(a) $3x$ t_1 undersampling used in simulated data reconstruction experiment

(b) $10x$ t_1 undersampling used in real data reconstruction experiment

Figure 2.5: Sampling schedules for undersampled reconstruction, white lines correspond to recorded $t1$ points. Collecting fewer FIDs reduces acquisition time and leads to an undersampled t_1-axis. Reducing the number of points recorded along the t_2-axis does not lead to a faster experiment as each FID must be recorded until time T.

reconstruction process is a direction for future work.

Our real spectrum is highly sparse and has a high dynamic range cf. fig. 2.7. The shortest peaks of interest are roughly 10^7 times shorter than the tallest. In this setting, both methods are able to accurately identify peaks as the Lorentizian line broadening is minimal cf. fig. 2.9. Reconstructing with a damped sinusoid basis, however, requires more computation time. Due to the increase in computational cost of the damped sinusoidal basis it may be advisable to reconstruct spectra with standard compressed sensing when prior knowledge of line broadening is available.

2.4 Conclusion

We have presented a methodology for improving resolution in compressed sensing NMR. By taking into account the peak broadening that is inherent in all NMR spectra, we have produced an algorithm for NMR spectra reconstruction which simultaneously deconvolves Lorentzian lines and accounts for missing data in a measured FID. While our method is more

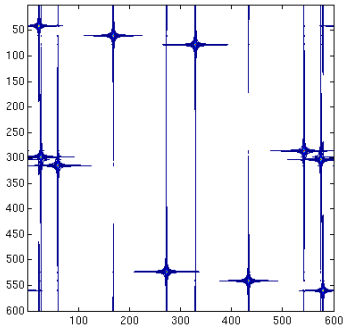

(a) Complete spectra with peak broadening.

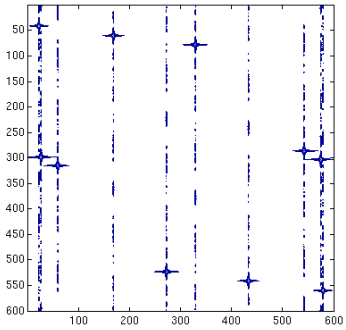

(b) Reconstruction from 33% t_1 data using standard compressed sensing. Computation time: 1.4288 seconds.

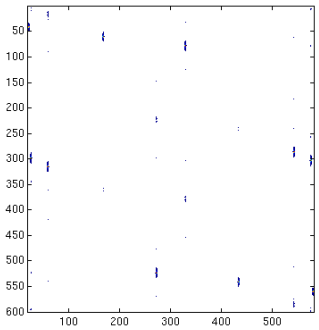

(c) Reconstruction from 33% t_1 data using a damped sinusoid basis. Computation time: 76.3435 seconds.

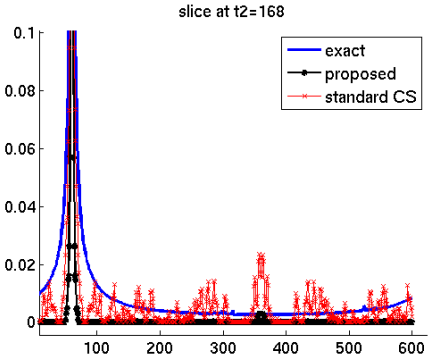

(d) Comparison of results at fixed t_2.

Figure 2.6: 2D simulated data reconstruction. In this test, severe line broadening leads to a non-sparse spectrum which is less amenable to traditional compressed sensing methods. In fig. 2.6d we can see that the true spectrum never reaches 0 and is difficult to approximate with sparsity promoting method. The damped sinusoid basis handles this situation better but has a higher computational cost.

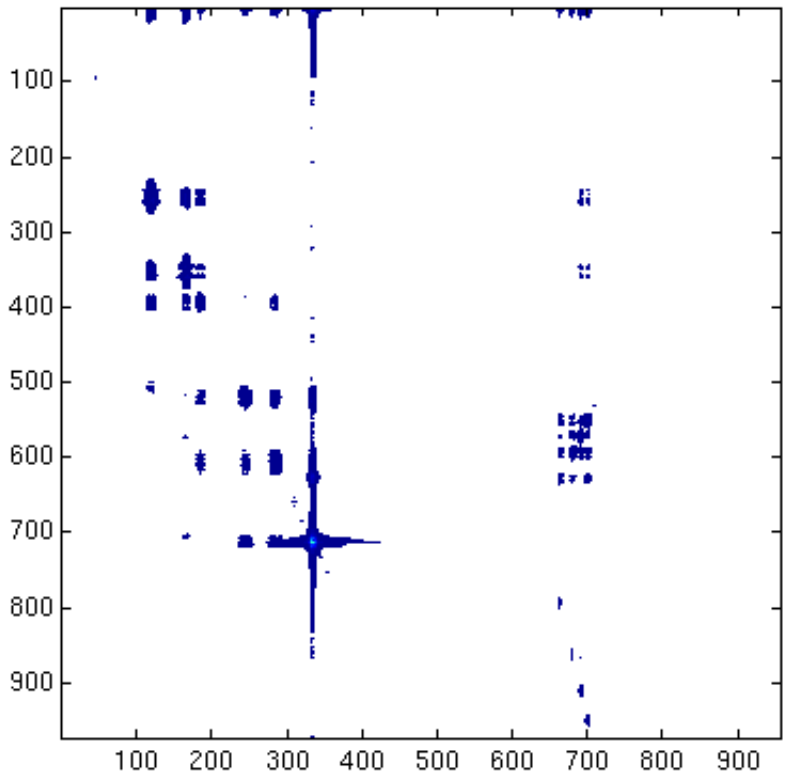

Figure 2.7: Real data used in reconstruction experiments

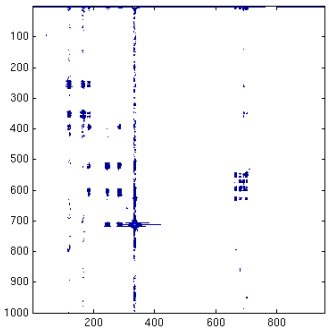

(a) Reconstruction from 33% t_1 data using standard compressed sensing. Computation time: 12.4488 seconds.

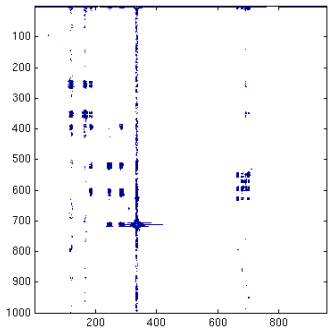

(b) Reconstruction from 33% t_1 data using a damped sinusoid basis. Computation time: 147.9138 seconds.

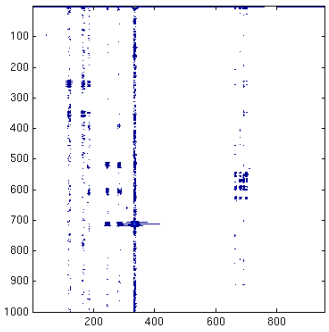

(a) Reconstruction from 20% t_1 data using standard compressed sensing. Computation time: 16.4423 seconds.

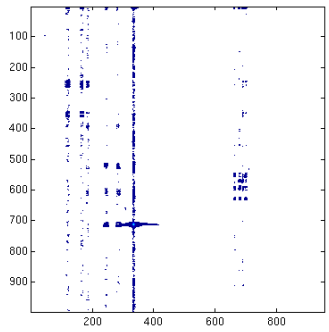

(b) Reconstruction from 20% t_1 data using a damped sinusoid basis. Computation time: 35.4590 seconds.

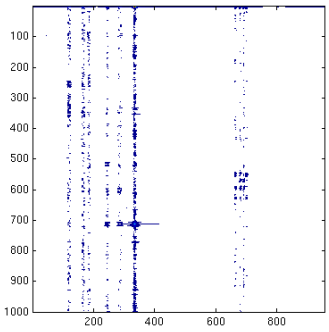

(a) Reconstruction from 10% t_1 data using standard compressed sensing. Computation time: 25.5009 seconds.

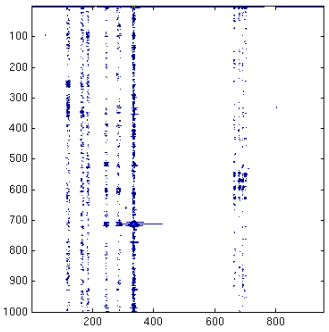

(b) Reconstruction from 10% t_1 data using a damped sinusoid basis. Computation time: 47.6557 seconds.

Figure 2.8: 2D real data reconstruction. In this test, our primary difficulties are t_1-noise and a high dynamic range of interest. The exponential damping factor alone is not enough to compensate for the level of t_1-noise and thus a sparse approximation to this noise is sought by the reconstruction algorithm. Further work is needed to address the problem of t_1 in compressed sensing reconstruction.

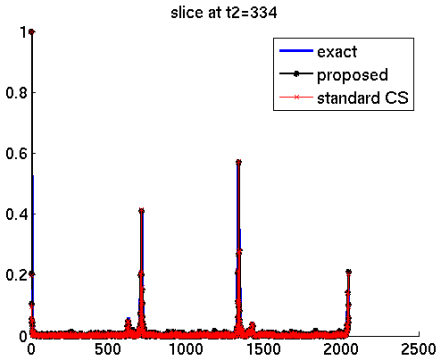

(a) Slice at $t_2 = 334$, both methods accurately recover peaks from a sparse spectrum

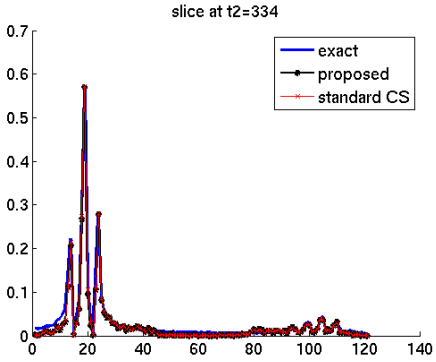

(b) Zoom of fig. 2.9a near a peak

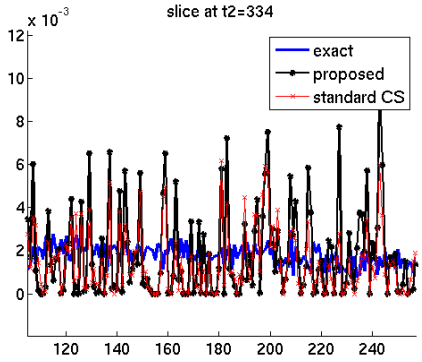

(c) Zoom of fig. 2.9a away from peaks, both methods have difficulty in the presence of t_1-noise

Figure 2.9: Slice along constant t_2 of real data reconstruction experiments from 33% data.

expensive computationally than a traditional compressed sensing reconstruction, numerical results indicate that modeling the damping factors can lead to much sharper final peaks from the same FID data. How well our method compares against approaches based on maximum entropy or multidimensional decomposition as well as combining the $l1$-term with a entropy promoting term are directions for future work.

Algorithm 3: CGIST for basis pursuit

Initialize: $u_0 = \vec{0}$, $e_0 = Au_0 - f$, $g_0 = A^T e_0$, $D_0 = (u_0 \neq 0)$,

$r_0 D_0 g_0 + \mu sign(u_0) + (1 - D_0) shrink(g_0, \mu)$, $r_1 = \frac{\|r_0\|^2}{\|Ar_0\|}$ and

$u_1, e_1, g_1, D_1, r_1, \alpha_1$ from one iteration of FBS.

while $\frac{|u_{k+1} - u_k|}{|u_k|} > tol$ **do**

Compute the step size:

$r_k = D_k g_k + \mu sign(u_k) + (1 - D_k) shrink(g_k, \mu)$

$\alpha_k = \frac{\|r_k\|^2}{\|Ar^k\|^2}$

Make one proximal step on u:

$\bar{u}_k = shrink(u_k - \alpha_k g_k, \alpha\mu)$

if $sign(\bar{u}_k) == sign(u_k) == sign(u_{k-1})$ **then**

Update u_k with conjugate gradient partan:

$\bar{e}_k = e_k - \alpha A r_k$

$\bar{g}_k = A^T \bar{e}_k$

$\bar{r}_k = r_k + D_k(\bar{g}_k + g_k)$

$\bar{\beta}_k = \frac{\langle \bar{r}_k, r_{k-1} \rangle}{\|r_{k-1}\|^2}$

$\beta_k = \min(\bar{\beta}_k, D_k \bar{u}_k / u_{k-1})$

$u_{k+1} = \frac{\bar{u}_k - \beta_k u_{k-1}}{1 - \beta_k}$

$e_{k+1} = \frac{\bar{e}_k - \beta_k e_{k-1}}{1 - \beta_k}$

$D_{k+1} = (u_{k+1} \neq 0)$

$g_{k+1} = \frac{\bar{g}_k - \beta_k g_{k-1}}{1 - \beta_k}$

$r_{k+1} = \frac{\bar{r}_k - \beta_k r_{k-1}}{1 - \beta_k}$

end

else

FBS update:

$u_{k+1} = \bar{u}_k$

$D_{k+1} = (u_k \neq 0)$

$e_{k+1} = Au_{k+1} - f$

$g_{k+1} = A^T e_{k+1}$

$r_{k+1} = D_{k+1} g_{k+1} + \mu sign(u_{k+1}) + (1 - D_{k+1}) shrink(g_{k+1}, \mu)$

end

end

CHAPTER 3

Hybrid regularization for MRI reconstruction with static field inhomogeneity correction

Rapid acquisition of magnetic resonance (MR) images via reconstruction from undersampled k-space data has the potential to greatly decrease MRI scan time on existing medical hardware. To this end, iterative image reconstruction based on the technique of compressed sensing has become the method choice for many researchers [LDP07]. However, while conventional compressed sensing relies on random measurements from a discrete Fourier transform, actual MR scans often suffer from off-resonance effects and thus generate data by way of a non-Fourier operator [FOS05a]. Correcting for these effects requires more sophisticated image reconstruction methods which come with additional computational bottlenecks that are not encountered in traditional compressed sensing.

In this work, we demonstrate how one may accelerate the convergence of algorithms for solving the image reconstruction problem,

$$\underset{\rho}{argmin}\ J(\rho) \text{ subject to } A\rho = s \tag{3.1}$$

by opting for a sparsity promoting regularization function, $J(\rho)$, of the form:

$$J(\rho) = |\nabla\rho| + \nu|F\rho| \tag{3.2}$$

when F is a tight frame and A is approximately a Fourier transform. In our experiments, reconstructing field-corrected MR images with the hybrid regularization of eq. (3.2) provides a speedup of roughly one order of magnitude when compared with an approach based solely

This chapter is joint work with Louis Bouchard and Stanley Osher and has been submitted to Inverse Problems in Imaging

on total-variation and can produce higher quality images than an approach based solely on tight frames.

3.1 Introduction

While the problem of image reconstruction in magnetic resonance imaging (MRI) has a long history, the techniques employed usually assume the static magnetic field is homogeneous and the applied magnetic-field gradients are unidirectional with constant magnitude [Man82]. In recent years, the prospects of generating MRI images using portable sensors have been explored [Pra03][Per04]. Although some efforts have focused on generating remote homogeneous fields [PBG08], most portable sensors are single-sided and characterized by inhomogeneous fields [BBE98] [MMC06]. The inhomogeneous field can be used to select a slice for imaging [Pra03]. In other instances, so-called "shimming" radio-frequency (RF) pulses have been used to prepare the spins and impart a phase that compensates for the effects of an inhomogeneous field [FDM09][TMS04][MSH01]. In the limit of low magnetic fields, the magnetic-field gradients are no longer unidirectional according to Maxwell's equations in the absence of a rotating-wave approximation. MRI is still feasible with the use of coherent averaging techniques [MST05][Kel09], oscillating fields [Bou06] or spatial shimming [BA07]. Each of these techniques have limitations. For coherent averaging and RF shimming techniques, the stroboscopic sequences required are not practical for *in vivo* use because of RF heating and the short T_2 relaxation times in tissues. In this chapter, we treat the problem of image reconstruction in high-field MRI, in the presence of inhomogeneous static and gradient magnetic fields. The direct reconstruction (as opposed to the use of compensating techniques) may be the preferred approach in clinical applications, where rapid imaging is required.

The most general MRI reconstruction problem assumes no rotating-wave approximation. Spins precess about a field $\vec{B}(\mathbf{x}, t)$ which may include time-dependent gradients. Its direction and magnitude may be a function of space. Because of the time-dependence of the local field,

the rotations are non-commutative and a time-ordered product must be used to describe its effect. For spin $I = 1/2$, the rotations are described by the spin operators $I_i = \sigma_i/2$, $i = x, y, z$, where σ_i are the Pauli matrices. The rotation operator is written in terms of Dyson time-ordering $\vec{T}$:

$$U(\mathbf{x}, t) = \vec{T} e^{i\gamma \int_0^t [I_x B_x(\mathbf{x},t) + I_y B_y(\mathbf{x},t) + I_z B_z(\mathbf{x},t)] dt}$$

The initial state of the nuclear magnetization density can be expressed as $\rho(\mathbf{x}) = I_x \rho_x(\mathbf{x}) + I_y \rho_y(\mathbf{x}) + I_z \rho_z(\mathbf{x})$. It is evolved in time according to the unitary transformation $\rho(\mathbf{x}, T) = U^\dagger(\mathbf{x}, T) \rho(\mathbf{x}) U(\mathbf{x}, T)$. Detection is done along a given axis defined by a field $\vec{B}_1(\mathbf{x})$. Neglecting relaxation effects, the NMR signal is therefore proportional to the volume integral of the Hilbert space trace (Tr):

$$s(t) = \int_V \mathrm{Tr}[\{B_{1x}(\mathbf{x}) I_x + B_{1y}(\mathbf{x}) I_y + B_{1z}(\mathbf{x}) I_z\} U^\dagger(\mathbf{x}, t) \rho(\mathbf{x}) U(\mathbf{x}, t)] d^3\mathbf{x}. \tag{3.3}$$

The image reconstruction problem consists of solving for the magnetization distribution $\rho(\mathbf{x})$ by measuring $s(t)$ at different times t while time-modulating the gradients.

If a strong static magnetic field is applied, only the components of the magnetic field gradients parallel to the static magnetic field are preserved. The perpendicular components of the gradient fields rapidly average to zero. This rotating-wave approximation leads to a local magnetic field whose direction remains fixed. Only its magnitude can change with time. Therefore, time-ordered rotations are no longer needed. The signal equation reduces to a simpler expression:

$$s(t) = \int_V \rho(\mathbf{x}) e^{-z(\mathbf{x}) t} e^{-2\pi i \mathbf{k}(t) \cdot \mathbf{x}} d\mathbf{x} \tag{3.4}$$

where $s(t_i)$ is the recorded signal, $z(\mathbf{x})$ describes the off-resonance effects (imaginary part). Relaxation effects can be included in $z(\mathbf{x})$ (real part). The magnetization density, $\rho(\mathbf{x})$, is a scalar-valued field. The goal of MRI reconstruction is to solve for the proton density map, $\rho(\mathbf{x})$, in the signal equation. The k-space coordinates, $\mathbf{k}(t)$, which are defined by the applied

gradients, are known only on a nonuniform and possibly undersampled grid [HB99]. The gradients are assumed to be uniform. For a 2D slice, V is a subset of $\mathbb{R}^2$.

A n-point discretization of the physical space integral at m different values of t leads to the $m \times n$ linear system

$$s = A\rho \tag{3.5}$$

with system matrix

$$A_{ij} = e^{-z(\mathbf{x}_j)t_i} e^{-2\pi i \mathbf{k}(t_i)\cdot \mathbf{x}_j}. \tag{3.6}$$

Efficiently inverting the linear system (3.5) in a way that produces high quality images is central to MRI.

With conventional MRI methods the background field is assumed perfectly homogeneous and the phase accrual due to off resonant frequency, $e^{-z(\mathbf{x}_j)t_i}$, is then ignored. Here, A becomes a Fourier encoding matrix and image reconstruction may be accomplished by directly inverting A or, in the event that our system is underdetermined, solving a regularized inverse problem using methods from compressed sensing [Rom07]. Regularization terms based on total variation [LDP07], wavelets [GHP11], and shearlets [ALG11] have been successful. In the more general case of parallel imaging these regularizations still provide quality images [CPB11].

While simple computationally, failing to accommodate for off-resonance effects often leads to blurring and image distortion in many types of MR scans. These effects are often due to abrupt spatial variations in the magnetic susceptibility of biological tissues being imaged. Near interfaces between regions of different susceptibility, the static background field, B_0, is perturbed by an additional, weaker, magnetic field, $\Delta B(\mathbf{x})$, whose strength is proportional to the strength of the background field and a function depending on the tissue, i.e. $\Delta B(\mathbf{x}) = \chi_m(\mathbf{x})B_0$ [YSW06]. The final background field, $B_0 + \Delta B(\mathbf{x})$, thus varies in space.

The value of χ_m is often small (for grey matter, $\chi_m \approx 8.97 \times 10^{-6}$ [YSW06]). However,

new developments in MRI technology are leading to stronger B_0 fields where the off-resonance effects are more substantial [TCS04]. In ultra-high field MRI (i.e. $\geq$ 7 T), static field inhomogeneities place notable limitations on the images a device can produce [TCS04].

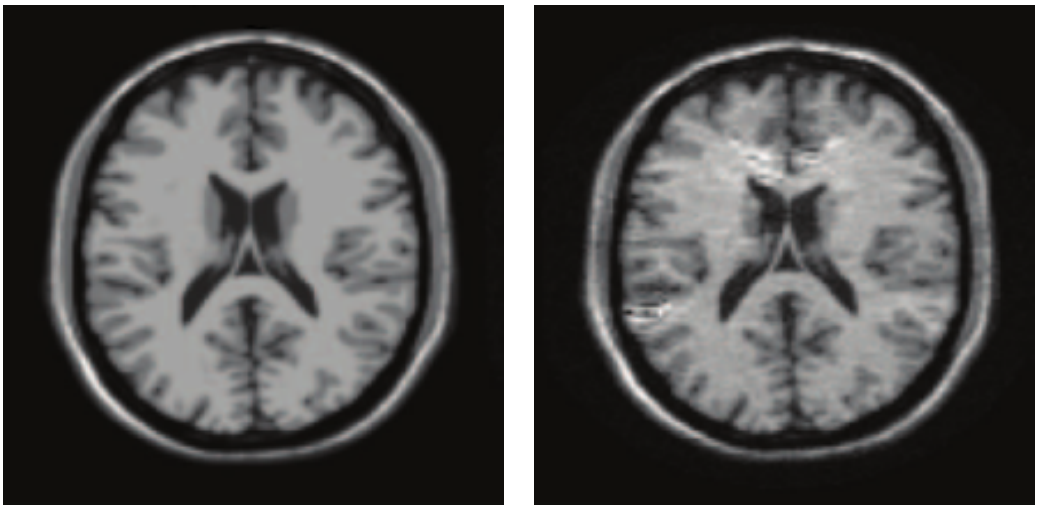

Figure 3.1: True and corrupted brain images, taken from [Fes10]. Image distortions near the nasal cavity result when off-resonance effects are ignored during reconstruction (an FFT in this case). The off-resonance modeling factor, $e^{-z(\mathbf{x})t}$, is multiplied in k-space, leading to a convolution in physical space. Accurate MR image formation thus requires deconvolution in addition to reconstruction.

Standard examples of images corrupted by off-resonance effects are found in cranial MRI scans near air/tissue interfaces fig. 3.1. Here, the differences in magnetic susceptibility between air and water are responsible for the creation of the perturbing field [RVY97]. The situation occurs again when imaging regions containing gray/white matter boundaries [YSW06]. As a result, correcting for off-resonance effects is of specific importance for surgical planning near the nasal sinuses, auditory canals, and cerebral cortex [MBB95] [NAG05].

Nonstandard examples occur when imaging patients who use hair products containing iron oxide or cobalt particles [MJ04]. Proper maintenance of "twists" or "dreadlocks" mandates that hair is saturated with beeswax (or colored beeswax) near the scalp in order to prevent essential knots from coming undone. Magnetization of the hair product leads to

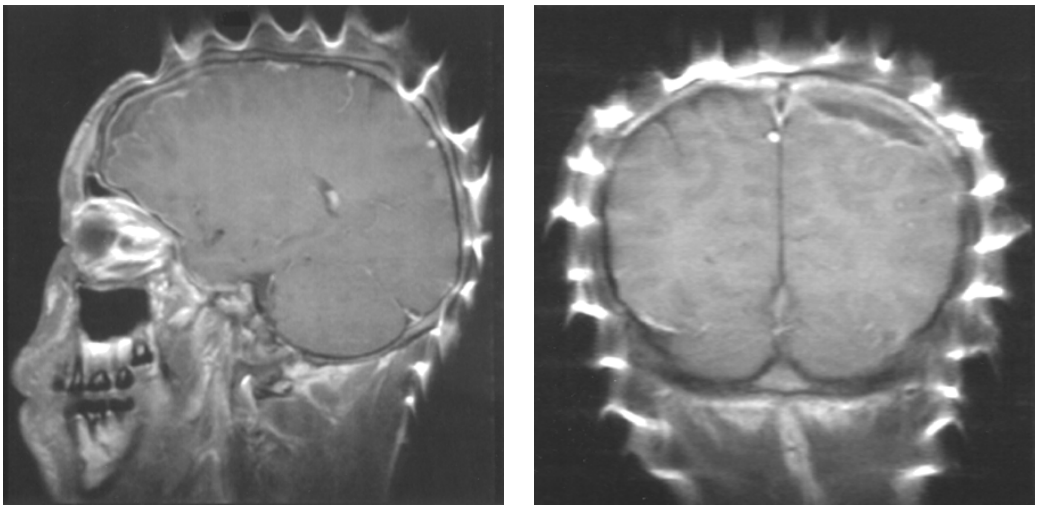

Figure 3.2: MRI of 15-year-old boy with acute sinusitis and subdural empyemas, taken from [MJ04]. Sagittal (left image) and coronal (right image) T_1-weighted images show subdural empyema as a dark grey region separating the brain and skull above the eye in the sagittal image and above the (anatomical) left hemisphere in the coronal image. Diagnosis is complicated due to image distortion from susceptibility artifacts caused by iron oxide particles suspended in beeswax dressing in patient's hair.

strong static field inhomogeneities and thus highly distorted images fig. 3.2. A similar effect occurs when scanning patients wearing colored eye makeup [Dun01].

To compensate for these distortions, we first note that the off-resonance factor, $e^{-z(\mathbf{x})t}$, in the image model eq. (3.4) is multiplied with the proton density map in k-space. This leads to a convolution in physical space. The reconstruction problem thus involves deconvolution in addition to Fourier decoding.

We focus on the case where $e^{-z(\mathbf{x})t}$ is known beforehand (i.e. non-blind deconvolution). There is a large and growing body of work aimed at determining $z(\mathbf{x})$. Current methods make use of many ideas that overlap with techniques used in general image processing. Some examples include determination of a field map via MAP estimation [BFS10] as well as $l1$-penalized optimization [KLS10].

Knowing the field map, however, does not immediately provide one with a prescription for high quality image formation as the matrix

$$E_{ij} = e^{-z(\mathbf{x}_j)t_i} \tag{3.7}$$

is often so large that storage is impossible on most modern desktops [FMS03]. As an example, to produce a 512×512 image the (dense) E matrix takes on dimensions $512^2 \times 512^2$. Storing E in 32-bit floating point precision requires 256GB of space. Consequently, directly forming the system matrix, A, is infeasible for practical image sizes and some form of operator compression must be employed in the reconstruction procedure [Fes10] [SMN03].

Once a procedure to apply E is known and k-space has been fully sampled one may reconstruct ρ using a variety of techniques [BUF07] [RF10] [WBP11]. Non-iterative approaches for field correction, such as the conjugate phase methods of [CSM08], [Sch99] and [NFS05] are fast and commonly used. However, in recent years, iterative approaches based on regularization have come into favor as they ignore this assumption and produce higher quality images. The drawback to iterative methods tends to be an increase in computational cost [Fes10].

One unexpected advantage of iterative reconstruction methods is the almost exact overlap they have with the recent and prolific theory of compressed sensing. As a result, recent years have seen tremendous progress in fast algorithms for iterative reconstruction [GO09]. The canonical problem in both methods is:

$$\underset{\rho}{argmin}\, J(\rho) \text{ subject to } A\rho = s \tag{3.8}$$

where J represents some form of sparsity promoting regularization.

Traditional iterative image reconstruction methods are known to be very reliable when a large amount of k-space has been acquired [SMN03]. In the absence of off-resonance effects, it is known from compressed sensing that exact image reconstruction is possible from extremely undersampled k-space data [Rom07].

In this work, we propose solutions when these two issues are combined, providing a

method for image recovery with a generalized operator when k-space is undersampled. Furthermore, we introduce a regularization term based on framelets [COS10] and establish a significant reduction in computational cost when this term is included.

3.2 Method

3.2.1 Low-rank inhomogeneity correction

Before we can successfully carry out any iterative solution procedure we must first be able to apply our system matrix. Writing A as the pointwise matrix product,

$$A = E \circ \mathcal{F} \tag{3.9}$$

we can see that a low-storage approximation to E will yield an approximation to A which fits in memory as the Fourier operator requires no storage.

Given a target rank, r, we construct a low-rank approximation to the off-resonance matrix,

$$E \approx BC \tag{3.10}$$

where $B \in \mathbb{R}^{m \times r}$ and $C \in \mathbb{R}^{r \times n}$ are thin (ie $r \ll m$ and $r \ll n$). Now,

$$\begin{aligned} s_i &= \sum_{j=1}^{n} A_{ij}\rho_j && (3.11)\\ &= \sum_{j=1}^{m} E_{ij} \circ \mathcal{F}_{ij}\rho_j && (3.12)\\ &\approx \sum_{j=1}^{m}\sum_{l=1}^{r} B_{il}C_{lj}\mathcal{F}_{ij}\rho_j && (3.13)\\ &= \sum_{l=1}^{r}\sum_{j=1}^{m} B_{il}C_{lj}\mathcal{F}_{ij}\rho_j && (3.14)\\ &= \sum_{l=1}^{r} B_{il}\sum_{j=1}^{m} \mathcal{F}_{ij}C_{lj}\rho_j && (3.15) \end{aligned}$$

leading to,

$$A \approx \sum_{l=1}^{r} diag(B_{(:,l)})\mathcal{F}diag(C_{(l,:)}) \tag{3.16}$$

where $diag(B_{(:,l)})$ and $diag(C_{(l,:)})$ are diagonal matrices with entries taken from the lth column and row of B and C, respectively. Provided that the singular values of E decay quickly enough that the low-rank approximation eq. (3.10) is valid, we may rapidly apply A with little storage overhead using eq. (3.16). It turns out that this is indeed the case for empirically observed field maps [FOS05b]. As the inhomogeneities become stronger, higher rank approximations are needed to accurately describe E.

When working with nonuniformly or undersampled k-space data we only need to replace the Fourier matrix with a discrete nonuniform Fourier transform operator (NUFFT) of type 2 [GL04]. In which case we write

$$A \approx \sum_{l=1}^{r} B_l \mathcal{G} C_l. \tag{3.17}$$

3.2.2 Forming the approximation

For the matrix decomposition, $E \approx BC$, several methods can be found in recent literature on MRI reconstruction [FOS05b] [IMN96] [MPM96] [Nol91] [MDL03]. The common foundation on which all these methods are built is approximation of the exponential,

$$e^{-z(\mathbf{x}_j)t_i} \approx \sum_{l=1}^{r} B_{il} C_{lj}. \tag{3.18}$$

Strategies for forming eq. (3.18) can be split into roughly two camps: those that focus on functional approximations of the form

$$e^{-z(\mathbf{x})t} = \sum_{l=1}^{r} b_l(t) c_l(\mathbf{x}) \tag{3.19}$$

and those that treat the decomposition as a general low-rank matrix approximation problem.

Matrix-based approaches typically lead to the most accurate approximations for a given rank, but require substantial computational overhead as the matrix, E, must be examined in full. Recent developments in randomized algorithms originating within the numerical linear algebra community may somewhat alleviate this burden [RST09] [MRT11] [HMT11].

3.2.2.1 Interpolative Decompositions

For a purely algebraic approach to the decomposition, we introduce a recently developed matrix factorization [LWM07]:

Definition 1 (Interpolative Decomposition). Let $E \in \mathbb{R}^{n\times n}$ have rank at most r. The decomposition

$$E_{n\times n} = B_{m\times r} C_{r\times n} \tag{3.20}$$

where the columns of B make up a subset of the columns of E, and each entry in C had magnitude at most two, is referred to as the *interpolative decomposition* of E.

Thus, we form an approximation to E using a bounded linear combination of a few columns of E. The existence of an interpolative decomposition for $|C_{ij}| \leq 1$ is established in [PTR98]. Computation of such a C, however, turns out to be NP-hard [cM09].

Computing an interpolative decomposition is fairly straightforward when a capable linear algebra library is available. Additional implementation is needed when one does not have access to a pivoted-QR decomposition or linear system solvers.

Algorithm 4: Interpolative Decomposition for low-rank E.

Form $Q_{n\times r} R_{r\times n} \Pi_{n\times n} = E_{n\times n}$ via a pivoted-QR decomposition [GL96]

Define $S_{r\times r}$ and $T_{r\times(n-r)}$ by $\left(S_{r\times r} \mid T_{r\times(n-l)}\right) = R_{r\times n}$

Compute $B_{n\times r} = Q_{n\times r} S_{r\times r}$

Compute $C_{r\times n} = \left(1_{r\times r} \mid (S^{-1})_{r\times r} T_{r\times(n-r)}\right) \Pi_{n\times n}$

The columns of B are a subset of the columns of E since they are determined by multiplying Q with columns of the R matrix found in the QR decomposition of E. The matrix S is small in our applications (e.g. $S \in \mathbb{R}^{10\times 10}$) so that the inversion is easily computed.

The decomposition in alg. 4 is exact when the rank of E is exactly r. In the general setting, where E is high-rank and we seek a low-rank approximation, we have the estimate for each r:

$$\|E_{n\times n} - B_{n\times r} C_{r\times n}\| \lesssim \sigma_{r+1} \tag{3.21}$$

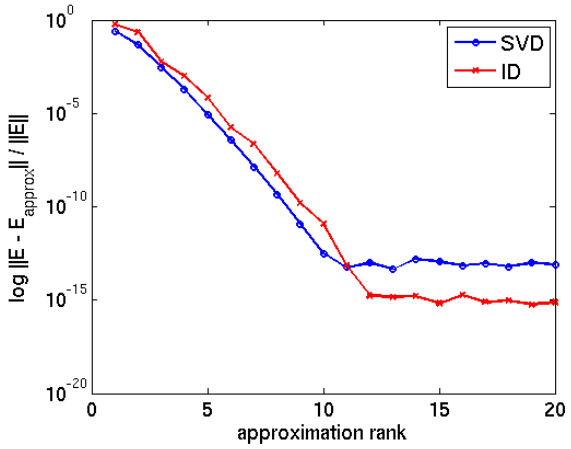

Figure 3.3: Relative error, $\log\left(\frac{\|E-E_{approx}\|}{\|E\|}\right)$, for varying rank approximations using a partial SVD (blue) and randomized interpolative decompositions (red). The field map is 64×64 leading to a 4096×4096 E matrix. Machine precision is 10^{-15}. High accuracy approximations at low rank are possible with a single pass over the data.

where σ_{r+1} is the $r+1$st greatest singular value of E [GE96].

For matrix compression tasks, we premultiply E with a Gaussian random matrix, $G \in \mathbb{R}^{r\times m}$, before factoring [LWM07]. Applying E to a small random matrix allows us to work in a reduced subspace which approximates the range of E. Only one read of E is needed to form the product. When a few buffer vectors are added to G (leading to $G \in \mathbb{R}^{p\times m}$ with $p \geq r$) the accuracy of this approximation is notably improved. Sharp bounds on the error eq. (3.21) are a topic of current research [HMT11].

Algorithm 5: Randomized Interpolative Decomposition for rank-r approximation to E.

Compute $Y_{p\times n} = G_{p\times n}E_{n\times n}$ for an integer $p \geq r$

Form $Z_{p\times r}C_{r\times n} \approx Y_{r\times n}$ using alg. 4

The columns of Z correspond to columns of Y. That is, for $j = 1 \ldots l$ we can find $i_1 \ldots i_l$ such that the jth column of Z is the i_jth column of Y. Form $B_{m\times r}$ by selecting columns $i_1 \ldots i_r$ from $E_{m\times n}$.

3.2.2.2 Randomized Partial SVDs

For a given rank, randomized interpolative decompositions provide approximations with accuracy asymptotically equivalent to a partial SVD [MRT11] and fig. 3.3. While sometimes less accurate in practice, interpolative decompositions have the advantage that the columns of B are columns of E which may be helpful in certain algorithms for off-resonance correction [FOS05b].

Fast computation of partial SVDs may also be performed via a similar, randomized, algorithm [HMT11]. When a linear algebra library is available, computing randomized SVDs is straightforward as well.

The core of randomized SVD computations consists of finding a matrix, $Q \in \mathbb{R}^{m \times p}$, such that the range of Q is a good approximation to the range of E. That is, we construct Q such that

$$\|E - QQ^*E\| \approx \min_{rank(A)=l} \|A - E\|. \tag{3.22}$$

This can be done, for example, using a QR-decomposition of the matrix Y in alg. 5. With this choice of Y, a bound on the expected value of the error is known [HMT11]:

$$\mathbb{E}\|E - Q^*QE\| \leq \left(1 + \frac{4\sqrt{p}}{p-l-1}\sqrt{n}\right)\sigma_{l+1}. \tag{3.23}$$

When Y is formed by applying E several times, the error section 3.2.2.2 may be reduced [RST09]. This is analogous with convergence of traditional power iterations with the difference that the bound takes into account the number of buffer vectors used before forming the approximation.

Once Q is obtained we compute the SVD of E by multiplying against Q, taking an SVD in the reduced subspace, and finally projecting back into the full space.

This gives us the approximation,

$$E \approx U\Sigma V^*. \tag{3.24}$$

Algorithm 6: Randomized SVD

Multiply $Y = GE^q$

Form $QR = Y$

Multiply $D = Q^*E$

Form $\tilde{U}\Sigma V^* = D$

Replace $U = Q\tilde{U}$

The final error in alg. 6 is

$$\mathbb{E}\|E - U\Sigma V^*\| \leq \left(1 + \frac{4\sqrt{p}}{p-l-1}\sqrt{n}\right)^{\frac{1}{4q+1}} \sigma_{l+1}. \tag{3.25}$$

where q is the number of power iterations used to form Y. Errors from randomized SVDs are comparable (though slightly better) than errors in the randomized ID approximation.

3.2.2.3 Functional decompositions

An alternative way to form the approximation eq. (3.10) is by explicitly approximation the exponential in the form of eq. (3.19). Standard approaches begin by selecting a subset of points uniformly from either the time or "frequency" domain [MPM96].

Segmenting in the time domain is accomplished by first choosing a set of points $\{\tilde{t}_l\}_{l=1}^r$ and a single value, $\bar{z}$, to approximate the field map. One then has the formulas

$$B_{il} = b_l(t_i)e^{-\bar{z}t_i} \tag{3.26}$$

$$C_{lj} = e^{-(z(\mathbf{x}_j)-\bar{z})\tilde{t}_l}. \tag{3.27}$$

Interpolating functions, $b_l(t)$, are used to evaluate the coefficients between each pair of $\tilde{t}_l$. The method used to interpolate the time domain has a large effect on the accuracy of this approach. Current methods favor a "min-max" approach [SMN03].

Selecting $\bar{z} = 0$ corresponds to standard unidirectional interpolation of the function $e^{-z(\mathbf{x})t}$. Choosing an optimal $\bar{z}$ improves accuracy when the $\tilde{t}_l$ are fixed in advance, optimizing the $\tilde{t}_l$ can lead to situations where the choice of $\bar{z}$ is irrelevant [FMS03]. A comparison of

many currently available methods for approximations of the form eq. (3.19) can be found in [FOS05b].

3.2.3 Restricted isometry constraints

A difficulty with the corrected system arises when we attempt to make the extreme reductions in k-space data that were suggested in the original compressed sensing literature. The modified system matrix obeys uncertainty properties differing from the pure Fourier case and no longer have the same guarantees on exact signal reconstruction [CR06].

Recall the concept of a *restricted isometry constant*, δ_s^A, of a matrix A for an integer s as the smallest number such that

$$(1-\delta_s^A)\|x\|_{l2}^2 \leq \|Ax\|_{l2}^2 \leq (1+\delta_s^A)\|x\|_{l2}^2 \tag{3.28}$$

where s is the number of nonzero entries in x. When such a constant exists A is said to satisfy the *restricted isometry property*. This property guarantees exact image reconstruction from vastly undersampled k-space with overwhelmingly high probability. In the case of an undersampled Fourier matrix, much work has been devoted to showing that $\delta_s^{\mathcal{G}}$ is small [CRT06].

In our corrected system, we can infer that each term in the sum eq. (3.18) has less than ideal isometric properties by examining

$$\|B_l\mathcal{G}C_l x\|_2 \leq \|B_l\|\|\mathcal{G}C_l x\|_2 \tag{3.29}$$

$$\leq \max_i(b_{il})(1+\delta_s^{\mathcal{G}})\max_j(c_{lj})\|x\|_2 \tag{3.30}$$

By combining eq. (3.30) with the analogous lower bound a crude estimate for the restricted isometry constant for an $r=1$ correction to our inhomogeneity can be found

$$\delta_s^{B_l\mathcal{G}C_l} \leq \min\Bigg(1-\min_i(b_{il})\min_j(c_{lj})+\min_i(b_{il})\min_j(c_{lj})\delta_s^{\mathcal{G}},\ 1-\max_i(b_{il})\max_j(c_{lj}) + \max_i(b_{il})\max_j(c_{lj})\delta_s^{\mathcal{G}}\Bigg). \tag{3.31}$$

Larger values of r increase the corresponding isometry constants and as a result more samples are required for accurate image recovery than in the pure Fourier case. Smaller values of r result in poorer approximations to the correction matrix. In this work, we choose r such that E is well approximated and accept that more k-space samples must be acquired.

3.2.4 Sparse recovery via $l1$ minimization

We propose a reduction in scan time by appealing to the theory of compressed sensing, downsampling our Fourier matrix uniformly at random and accounting for the missing data with sparsity promoting regularization.

While MR images are not sparse in the image domain, they are sparse in an appropriately chosen transform domain [LDP07]. Representation of images in bases of wavelets, and by extension, framelets yield sparse collections of coefficients [COS10]. Similarly, reconstruction methods based on the total variation norm well as its nonlocal counterpart have been shown to accurately reproduce detailed images from sparse frequency data [LWY09].

We advocate a composite of total variation and framelet regularization,

$$J(\rho) = |\nabla\rho| + \nu|F\rho| \tag{3.32}$$

where the first term is a total variation norm and F is the discrete framelet decomposition [COS10]. The total variation term enhances edges in our reconstructed image while the inclusion of the framelet term allows us to reconstruct smooth images. The parameter ν is chosen by the user and controls the relative impotance of edges versus smoothness in the sought after image.

It has been found previously that hybrid regularization based on wavelets often improves image quality [GO09]. Our framelet based regularization extends this work to redundant orthogonal bases where natural images can be represented more accurately.

3.2.5 Split Bregman iterations for image recovery

Our image is the solution to the constrained optimization 3.1 which we solve with the Split Bregman method of [GO09]. We begin by converting the constrained optimization into a sequence of unconstrained problems via Bregman iteration

$$\begin{cases} \rho^{k+1} &= \min_\rho |\nabla\rho| + \nu|F\rho| + \frac{\mu}{2}\|A\rho - s^k\|^2 \\ s^{k+1} &= s^k + s - A\rho^{k+1} \end{cases} \tag{3.33}$$

where the parameter μ affects the convergence rate and is chosen by the user [OBG05]. Typical values of μ range between 0.5 and 1. High values of μ better enforce the constraint at each iteration but have the drawback that the ρ^{k+1} update becomes harder to solve [GO09]. Alternately updating ρ^k and s^k produces a sequence $\rho^k \to \rho$, the solution to the constrained problem.

Updating s^k is straightforward. Minimization of the unconstrained step in eq. (3.33) is done by introducing auxiliary variables, $d_x = \nabla_x\rho$, $d_y = \nabla_y\rho$, and $w = F\rho$, allowing us to rewrite our ρ^k update in the equivalent split form

$$\rho^{k+1} = \min_\rho |\nabla d| + \nu|w| + \frac{\mu}{2}\|A\rho - s^k\|^2 \tag{3.34}$$

$$\text{subject to} \begin{cases} (d_x, d_y) = (\nabla_x\rho, \nabla_y\rho) \\ w = F\rho \end{cases} \tag{3.35}$$

This constrained optimization is then converted to a sequence of unconstrained problems via a second Bregman iteration leading us to the following algorithm:

Algorithm 7: Split Bregman iteration for constrained optimization

Initialize: $\rho^0 = A^t s$, and $d_x^0 = d_y^0 = w^0 = b_x^0 = b_y^0 = b_w^0 = 0$

while $\|A\rho^k - s\|_2^2 > tol$ **do**

for $i = 1$ *to* n_{inner} **do**

$\rho^{k+1} = \min_\rho \frac{\mu}{2}\|A\rho - s^k\|^2 + \frac{\lambda}{2}\|d_x^k - \nabla_x\rho - b_x^k\|^2 + \frac{\lambda}{2}\|d_x^k - \nabla_x\rho - b_x^k\|^2 + \frac{\gamma}{2}\|w^k - F\rho - b_w^k\|^2$

$d_x^{k+1} = shrink(\nabla_x\rho^{k+1} + b_x^k, 1/\lambda)$

$d_y^{k+1} = shrink(\nabla_y\rho^{k+1} + b_y^k, 1/\lambda)$

$w^{k+1} = shrink(F\rho^{k+1} + b_w^k, \nu/\gamma)$

$b_x^{k+1} = b_x^k + (\nabla_x\rho^{k+1} - d_x^{k+1})$

$b_y^{k+1} = b_y^k + (\nabla_y\rho^{k+1} - d_y^{k+1})$

$b_w^{k+1} = b_w^k + (F\rho^{k+1} - w^{k+1})$

end

$s^{k+1} = s^k + s - A\rho^{k+1}$

end

Here, the function *shrink* comes from the wavelet literature and is defined as

$$shrink(x, a) = \frac{x}{|x|}\max(|x| - a, 0). \tag{3.36}$$

The constants λ and γ are chosen by the user and affect the convergence rate.

Computationally, the ρ^{k+1} update is the most expensive part of our algorithm by a wide margin. The speed of our image reconstruction is determined by how fast we can solve this minimization. In the purely Fourier case, an analytic solution exists leading to a notably fast algorithm [GO09]. We have no such formula for the generalized A we work with and instead rely on iterations of the conjugate gradient method to update ρ^{k+1}.

By differentiating with respect to ρ and setting the result to zero we find our ρ^{k+1} update as the solution to:

$$(\mu A^t A + \lambda\nabla_x^t\nabla_x + \lambda + \gamma F^t F)\rho^{k+1} = rhs^k \tag{3.37}$$

where

$$rhs^k = \mu A^t s^k + \lambda\nabla_x^t(d_x^k - b_x) + \lambda\nabla_y^t + \gamma F^t(w - b_w). \tag{3.38}$$

Making use of the identities $\nabla^t\nabla = -\triangle$ and $F^tF = I$, gives the system

$$(\mu A^t A - \lambda\triangle + \gamma I)\rho^{k+1} = rhs^k \tag{3.39}$$

which we solve with conjugate gradient iterations.

A major advantage of our hybrid regularizer is now apparent. Consider the system resulting from a regularization based only on total variation (ie $\gamma = 0$),

$$(\mu A^t A - \lambda\triangle)\rho^{k+1} = rhs^k. \tag{3.40}$$

Denoting maximal and minimal eigenvalues of the matrix in eq. (3.40) by $\lambda_{\max}$ and $\lambda_{\min}$ we can write the condition number of our hybrid regularized system as

$$\frac{\lambda_{\max} + \gamma}{\lambda_{\min} + \gamma} < \frac{\lambda_{\max}}{\lambda_{\min}} \tag{3.41}$$

notably speeding our updates.

3.3 Numerical Results

Our hybrid reconstruction method is compared against three alternate approaches: uncorrected nonuniform-FFT, field corrected total variation regularization, and field corrected framelet regularization.

The nonuniform Fourier transforms are computed with the min-max method of [FMS03] using 6 points of interpolation. Piecewise linear B-spline framelets [COS10] are the framelet of choice. An open source implementation of the framelet transformation is available at http://www.math.ucla.edu/ jegilles/BregmanCookbook.html.

The majority of our computation time is spent in the FFT as each application of A requires only FFTs and diagonal matrix multiplication. High quality, scalable, parallel FFT implementations are readily available [FJ05] and our algorithm thus adapts naturally to parallel architectures.

Numerical experiments were run in 64-bit Matlab on twelve cores of a dual hex core

(a) 128x128 Shepp-Logan phantom

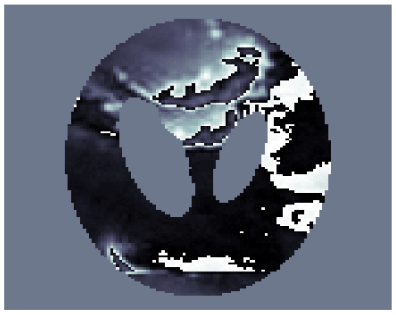

(b) Simulated field map, $z(\mathbf{x})$ (values range between -15 and 15 Hz)

(c) Phantom reconstruction without field map correction

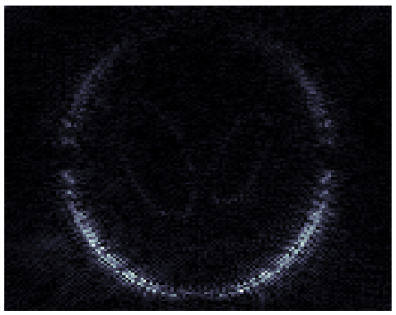

(d) Magnitude of difference between fig. 3.4a and fig. 3.4c

Figure 3.4: Phantom and simulated field map data used in experiments.

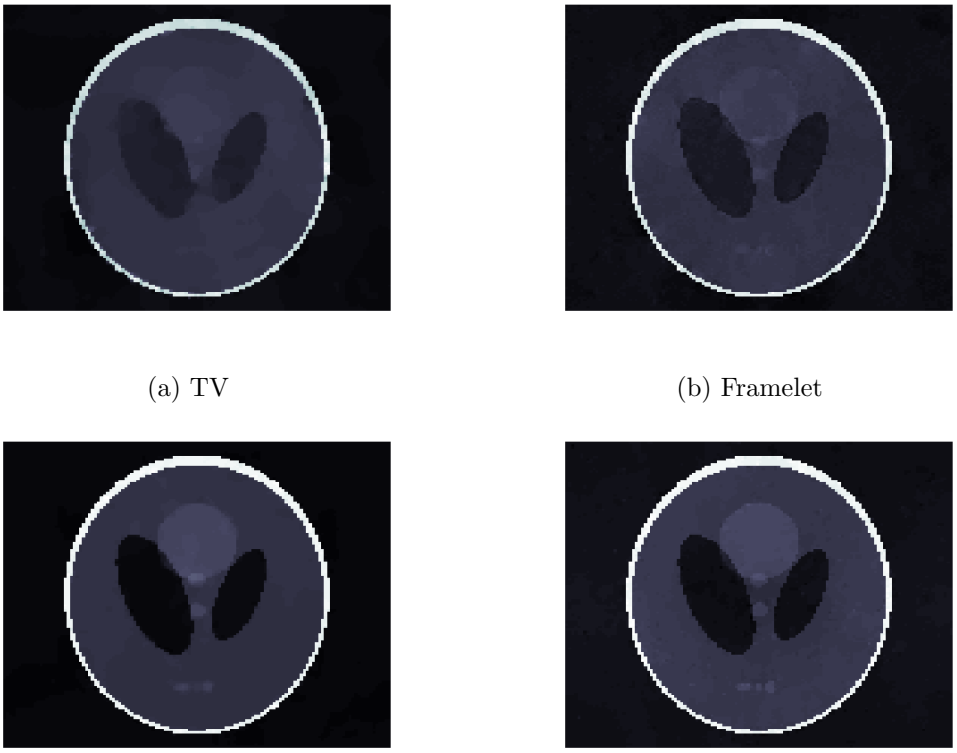

(a) TV

(b) Framelet

(c) Hybrid, $\nu = 1$

(d) Hybrid, $\nu = .01$

Figure 3.5: Comparison of phantom reconstructions from 20% data. All experiments were stopped after 3000 applications of A, independent of convergence. Average computation time was 88.35 seconds. The phantom is exactly piecewise constant and thus amenable to total variation regularization. However, the poor conditioning of eq. (3.40) prevents us from reaching the desired image in time. The highest quality image is fig. 3.5d, likely due to the facts that total variation is a good fit for the phantom and the positive ν allows us to optimize fast.

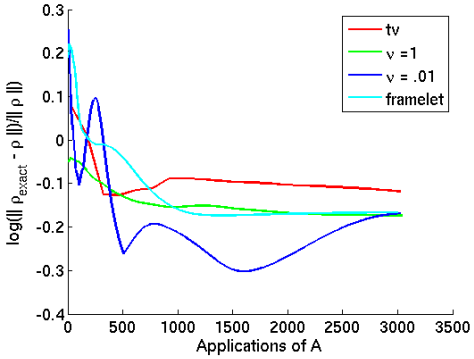

(a) Image domain errors per application of A

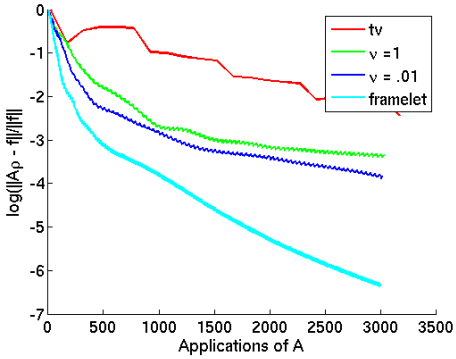

(b) Residual errors per application of A

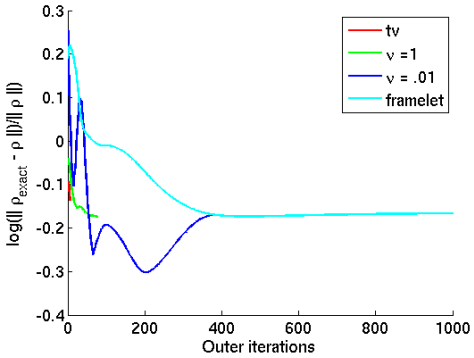

(c) Image domain errors per outer iteration

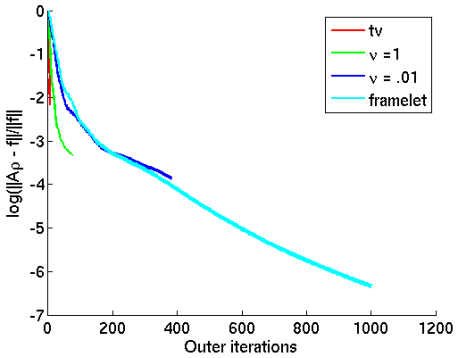

(d) Residual errors per outer iteration

Figure 3.6: Errors during reconstruction of phantom. In the first two rows, errors are recorded for each application of A. We note that hybrid methods attain roughly the same numerical accuracy as framelet based methods. However, visual inspection of fig. 3.5d suggests that methods penalizing total variation result in a higher quality image here. In fig. 3.6c and fig. 3.6d we see that the framelet term allows us to execute more outer iterations before the iteration limit is reached.

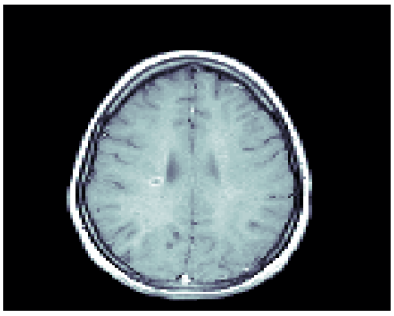

(a) 128x128 clean axial MRI

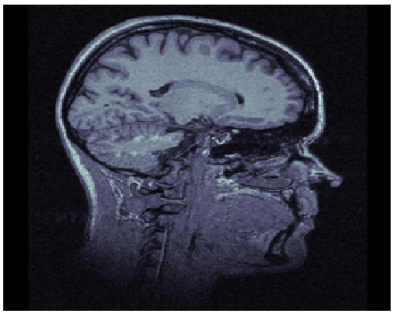

(b) 256x256 noisy sagittal MRI

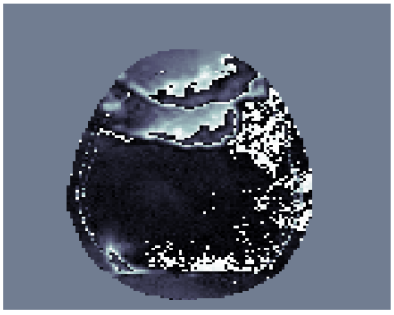

(c) Field map used with fig. 3.7a, values range from −15 to 15 Hz.

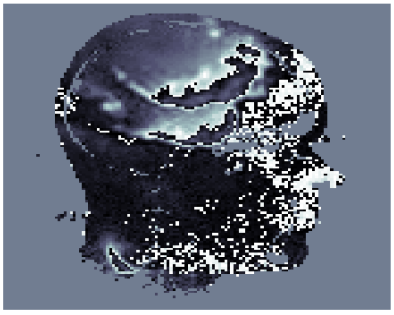

(d) Field map used with fig. 3.7b, values range from −15 to 15 Hz.

Figure 3.7: Exact images and field maps used in experiments.

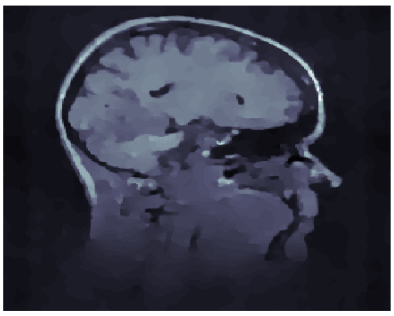

(a) TV

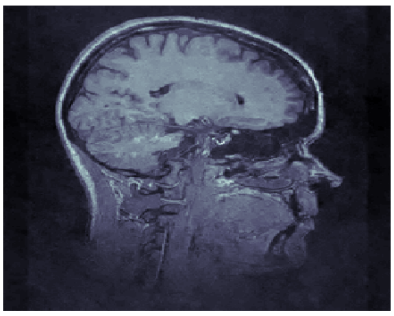

(b) Framelet

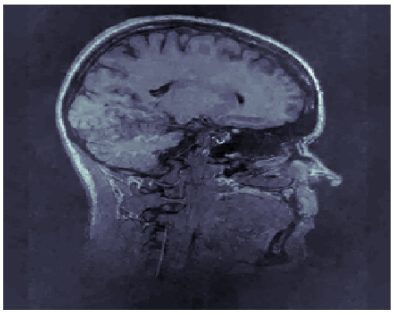

(c) Hybrid, $\nu = 1$

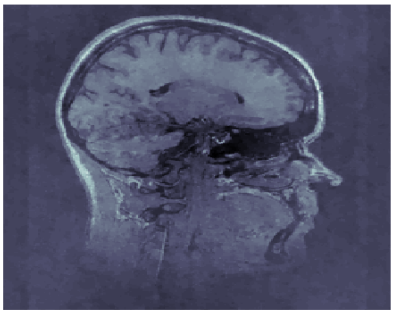

(d) Hybrid, $\nu = .01$

Figure 3.8: Comparison of image reconstructions from 40% data. All experiments were stopped after 3000 applications of A. Average computation time was 535.69 seconds. Total variation has removed noise but has blurred important features. The framelet based regularization in fig. 3.8b produced the highest quality image here.

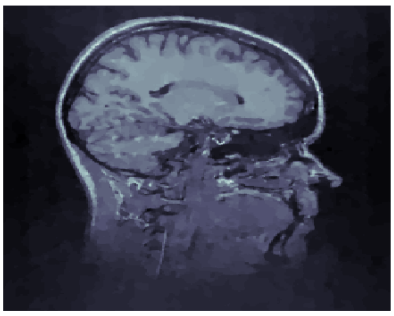

(a) TV

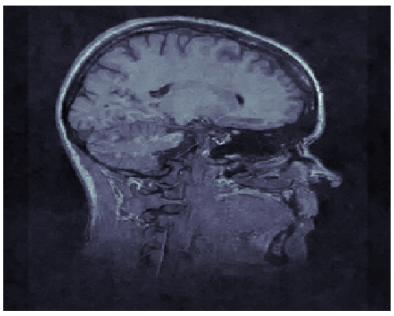

(b) Framelet

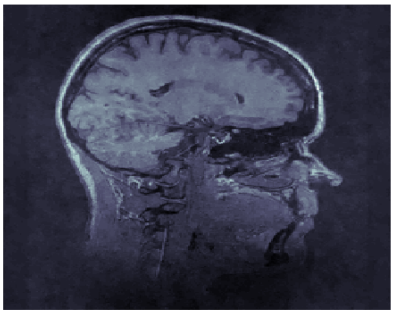

(c) Hybrid, $\nu = 1$

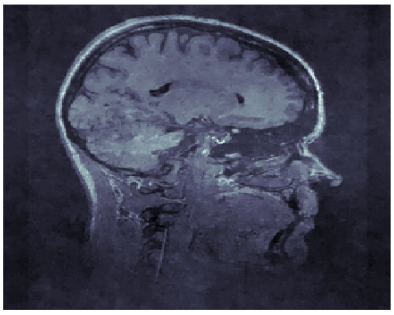

(d) Hybrid, $\nu = .01$

Figure 3.9: Comparison of image reconstructions from 40% data. All experiments were stopped after 10000 applications of A. Average computation time was 1767.2 seconds. Total variation has had time to produce a more detailed image than in fig. 3.8a. The regularizations involving framelets have produced similiar images.

system compromised of two 2.67 GHz Intel Xeon CPUs with 12 MB of level 2 cache each and 50 GB of RAM.

In every image reconstruction experiment we fixed $\mu = 0.5$. When a total variation term was present we set $\lambda = 0.5$. When a framelet term was present we set $\gamma = 5.0$. The rank of the low-rank approximation to E is fixed at 10 and computed using a randomized interpolative decomposition.

In our three-dimensional experiment we set $\mu = 1$, $\lambda = 0.5$, and $\gamma = 5.0$. We approximated E with rank a 6 approximation computed using time interpolation.

For two-dimensional datasets, we used a Shepp-Logan phantom (cf. fig. 3.4 and fig. 3.7a) and the standard Matlab MRI data set as our clean images. To validate our method on realistic and noisy data, we obtained a sagittal MRI of the head courtesy of Dr. Rohan Dharmakumar from Mount Sinai Hospital (cf. fig. 3.7).

For our three-dimensional dataset we used a T1-weighted volume scan of a healthy 22 year old male taken on a Brucker BioSpin at 4.3T with 128x128x128 resolution obtained from the LONI Image Database at https://ida.loni.ucla.edu (cf. fig. 3.11). We retrospectively undersampled k-space uniformly down to 40% before our volume reconstruction experiment. For the hybrid regularization in three dimensions, we used a Haar wavelet term. While it is known that Haar wavelets are suboptimal for MRI reconstruction [CPB11], a three dimensional framelet library is not yet available. Even with a suboptimal regularization, the additional term speeds computation enough that a Hybrid regularization still outperforms the total variation regularization, fig. 3.12.

Model field maps were obtained by scaling a smoothed and noisy image to values comparable with empirically observed maps [FOS05a]. In practice, the complete field map is often unattainable and one must estimate it from multiple scans [FFY08]. For each experiment, two acquisitions at echo times differing by $\Delta_1 = 2$msec were simulated to produce images:

$$y_1 = \rho + \sigma_1 \tag{3.42}$$

$$y_2 = e^{iz(\mathbf{x})\Delta_1}\rho + \sigma_2 \tag{3.43}$$

allowing us to obtain an estimate for the field map:

$$z(\mathbf{x}) = \mathrm{Arg}(y_1^* y_2)/\Delta_1 \tag{3.44}$$

For all imaging experiments the same background field was used, the estimate, however, was recomputed for each image.In our volume reconstruction experiment we stacked the background field image along the z-axis and repeated the procedure in eq. (3.44).

Plotting errors per iteration in fig. 3.6c and fig. 3.6d we see that total variation is the most efficient per outer loop iteration. This suggests that total variation alone may be advisable when the ρ^{k+1} update can be done analytically. However, when the number of matrix multiplications is taken into account, total variation is outperformed by other regularization strategies (cf. fig. 3.10 and fig. 3.5). When the data contains noise reconstruction is more difficult (cf. fig. 3.8 and fig. 3.9).

3.4 Conclusion

We have presented an efficient and fast-converging algorithm for MRI image reconstruction in the presence of inhomogeneous fields. This requires knowledge of the k-space trajectory and field inhomogeneity profile. The latter could be obtained in real time using B_0 mapping techniques. The former is prescribed by the imaging pulse sequence. The technique should be applicable to the class of problems where the signal is a linear function of the spin density, $s = A\rho$, and the matrix A can be decomposed into inhomogeneity and spatial encoding matrices. We note that the case of low-field MRI, the signal equation (3.3) is also of the form $s = A\rho$, where A can be separated into encoding and inhomogeneity matrices. One difference is that computational overhead is increased due to the requirement for computing products of non-commutative rotation matrices. A more important difference is that the encoding matrix is no longer of type NUFFT. It is unclear at this stage if such non-NUFFT matrices would lead to convergence of our algorithm. Further investigations would be needed.

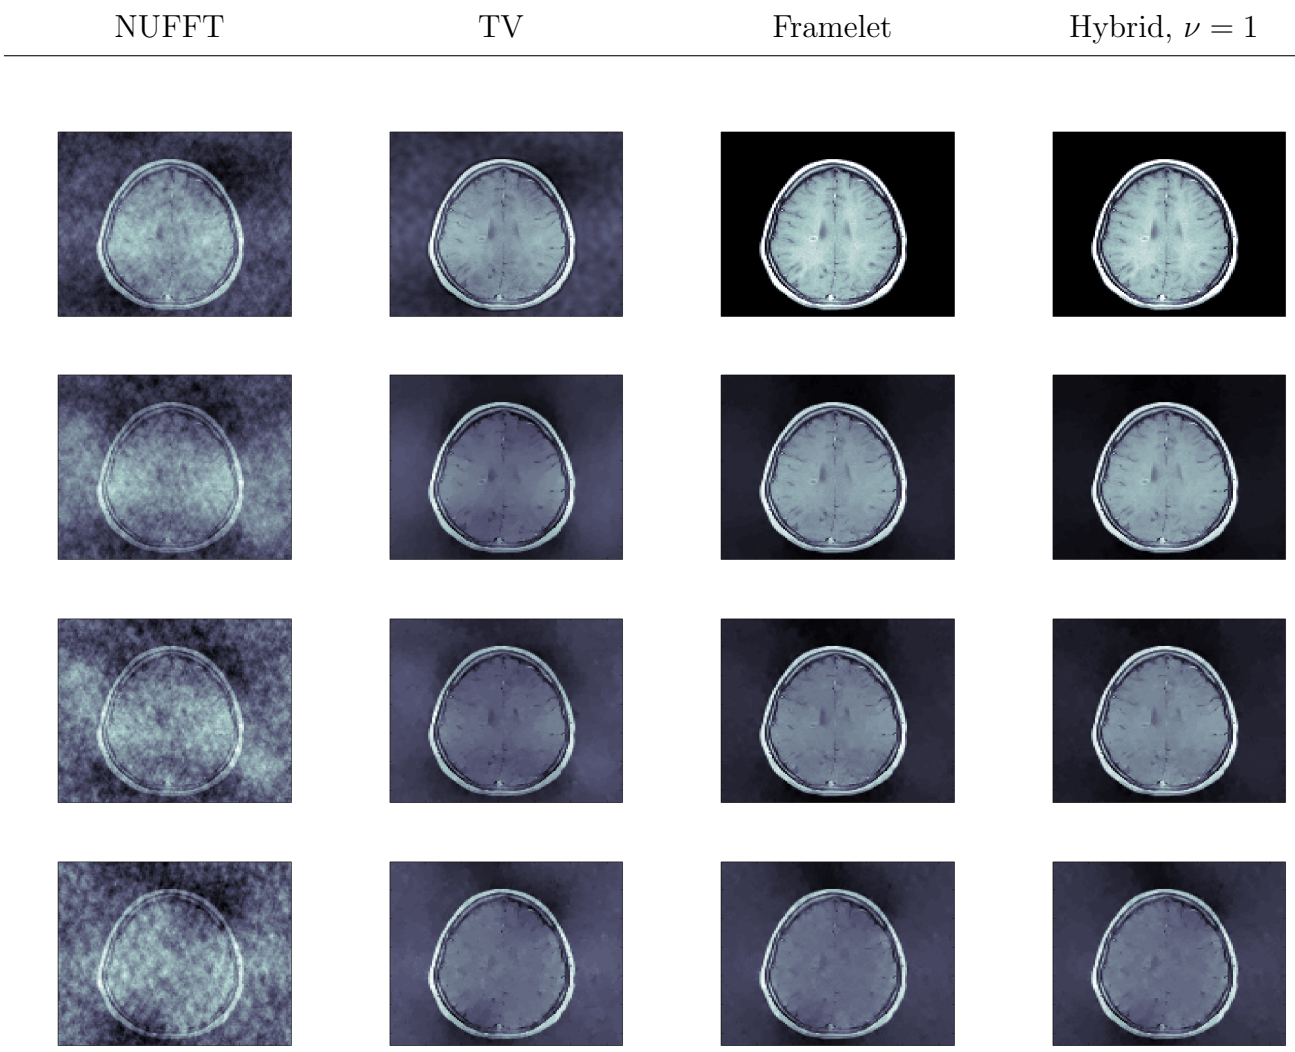

Figure 3.10: Clean brain image reconstruction. Rows correspond to downsample factors of 66%, 40%, 29% and 22%. All iterative methods were stopped after 3000 applications of A. Average reconstruction time was 147.83 seconds.

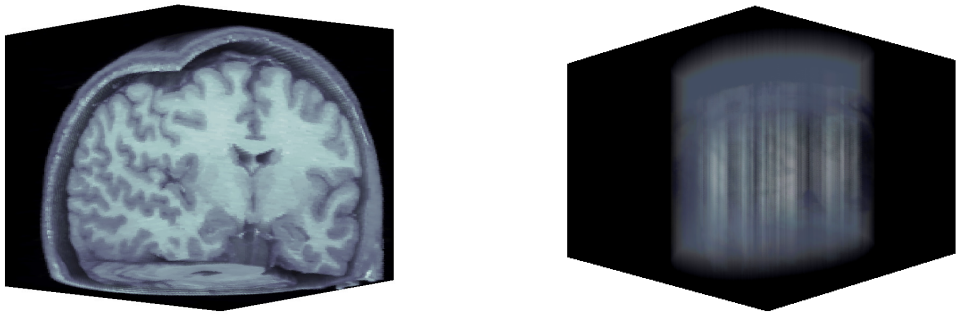

(a) Exact

(b) Field map, values range from −15 to 15 Hz.

Figure 3.11: Exact data used in volume reconstruction experiment. The volume dataset size was 128x128x128. Field map values range from −15 to 15 Hz, alpha mapping in fig. 3.11b has been reduced to make the most distorted parts of the volume visible.

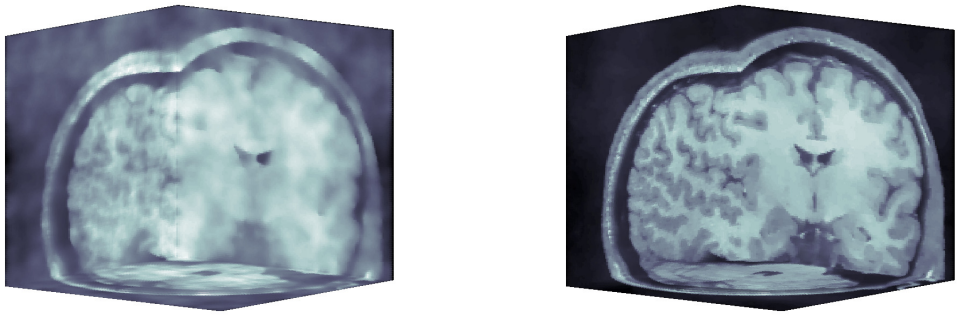

(a) TV

(b) Hybrid, $\nu = 1$

Figure 3.12: Volume reconstruction of a brain scan at 40% undersampling. All iterative methods are stopped after 5500 applications of A. Average computation time was 2981.13 seconds. The average number of applications of A required to update ρ^{k+1} was 124.22 when a total variation regularizer was used. This was reduced to 9.79 with the hybrid method.

CHAPTER 4

A sparse spectral method for Hermitian eigensystems

We study "spectral methods" [FFS82] for the calculation of eigenvalues of a Hermitian matrix, H. When computing eigenvalues with a traditional spectral method, one invents a fictitious quantum dynamical system with Hamilitonian H, simulates this system over time, and uses the data generated from the simulation to compute the eigenvalues of H. Classically, the computational cost of the method is directly related to the time step size used in the simulation and the time step size is restricted by the spectral radius of H.

In this chapter, we demonstrate how one may overcome time step size restriction by advancing the simulation with randomly varying time step sizes. The results of this chapter broaden the applicability of traditional spectral methods to to situations where the time step size is limited by the spectral radius of H rather than the accuracy of the propagation scheme used in the simulation.

4.1 Introduction

Computing eigenvalues of Hermitian matrices from time dependent quantum mechanical simulations has proven to be an efficient and robust methodology for nearly thirty years [Kos94]. First advocated by Feit, Fleck and Steiger in 1982, these so-called "spectral methods" have become fundamental in the development of many practical quantum mechanical and linear algebraic algorithms [Fei83][Mar91][FFS82].

The goal of spectral methods is to find solutions of the Hermitian eigenvalue problem,

$$H\psi = \lambda\psi \tag{4.1}$$

where H is time-independent. The idea behind spectral methods is to use solutions to a related, time-dependent, Schrodinger equation to construct solutions to eq. (4.1). Consider the following time-dependent equation,

$$-i\frac{d}{dt}\psi(t) = H\psi(t). \tag{4.2}$$

For an initial wavefunction, ψ_0, solutions to eq. (4.2) at are expandable in a basis of eigenfunctions of H,

$$\begin{aligned} \psi(t) &= e^{-iHt}\psi_0 && (4.3)\\ &= e^{-iHt}\sum_{n=1}^{d} A_n\phi_n && (4.4)\\ &= \sum_{n=1}^{d} A_n e^{-i\lambda_n t}\phi_n. && (4.5) \end{aligned}$$

Here, $\{\phi_n\}_{n=1}^{d}$ are the orthonormal eigenfunctions at time zero with eigenvalues λ_n and the coefficients, A_n, are defined by $A_n = \langle\phi_n, \psi_0\rangle$. Form the autocorrelation function of $\psi(t)$,

$$\mathcal{P}(t) = \langle\psi_0, \psi(t)\rangle \tag{4.6}$$

and Fourier transform to obtain:

$$\begin{aligned} \hat{\mathcal{P}}(\lambda) &= 2\pi\int_{-\infty}^{\infty}\mathcal{P}(t)e^{-i\lambda t}dt && (4.7)\\ &= \sum_{n=1}^{d}|A_n|^2\delta(\lambda-\lambda_n). && (4.8) \end{aligned}$$

The eigenvalues are now readily obtainable from eq. (4.8) by locating the spikes in $\hat{\mathcal{P}}(\lambda)$.

In practice, the energy domain representation of the autocorrelation function, $\hat{\mathcal{P}}(\lambda)$, can not be known with perfect resolution and aliasing effects impede our ability to accurately determine eigenvalues. The accuracy of the energy domain representation is limited by the uncertainty principle in two ways: a large final propagation time, T, is needed to obtain a high resolution $\hat{\mathcal{P}}(\lambda)$, and a small time step size, Δt, is needed to avoid aliasing errors in $\hat{\mathcal{P}}(\lambda)$.

Explicitly, the lower bound on the density of states that can be resolved from a signal $\mathcal{P}(t)$ is $\frac{2\pi}{T}$ where T. An upper bound on the time step Δt is controlled by $\Delta t < \frac{2\pi}{\Delta E}$ where ΔE is the spectral radius [BH95].

Overcoming the large T uncertainty principle and allowing for short propagation times while still resolving closely spaced eigenvalues is the achievement of the Filter Diagonalization method introduced by Neuhauser in 1990 [WN95] [Neu94] [Man01] [Neu90].

The small Δt bound is traditionally dealt with by truncating the potential above some cutoff value and advancing ψ over a uniform grid spaced by $\frac{1}{\Delta E_{cutoff}}$. The purpose of this work is to avoid the small Δt bound by exploiting the sparse structure of $\hat{\mathcal{P}}(\lambda)$.

Specifically, we propose recovery of $\hat{\mathcal{P}}(\lambda)$, from a few random measurements of $\mathcal{P}(t)$ by replacing the Fourier transform in eq. (4.8) with an $l1$-penalized optimization. $l1$-penalized optimization does not require uniformity in the time-step size, allowing us to run our simulation of the time-dependent system over substantially fewer time-domain points in order to reach the same T. By invoking the central theorems of compressed sensing [CRT05], our results guarantee that the eigenvalues of H are exactly obtained with overwhelmingly high probability when a randomized time stepping procedure is employed.

4.2 Method

Recall that if the Schrodinger equation in eq. (4.2) is solved over $[0, T]$ with step size Δt then the computed spectrum is in the range $[-\pi/\Delta t, \pi/\Delta t]$ with resolution $\Delta\lambda = 2\pi/T$. By the Nyquist-Shannon sampling theorem, capturing the full spectrum enforces the time step constraint,

$$\lambda_{max} - \lambda_{min} < \frac{2\pi}{\Delta t}. \tag{4.9}$$

Failure to satisfy eq. (4.9) will introduce "aliasing" errors and impede our ability to accurately determine λ_{max} and λ_{min}. This presents a problem as sampling many data points is computationally expensive fig. 4.1.

If an FFT is used to transform the autocorrelation function into the energy domain it is necessary to sample $\mathcal{P}(t)$ at all N grid points in the time domain. Note however that the number of nonzero entries in $\hat{\mathcal{P}}(\lambda)$, d, is typically much smaller than N [CHE96]. In this situation we may take advantage of new techniques in sparse signal processing known collectively as "compressed sensing" allowing us to sample $\mathcal{P}(t)$ at significantly fewer points by replacing the FFT with a convex optimization problem [CRT05].

4.2.1 Compressed sensing

Our key result rests on the fact that the time/energy conversion may be accomplished with much fewer data points than Feit had originally advocated. That this is possible is the central result of "compressed sensing" [Can06] which we briefly review in this section.

Our target data is contained in the length N signal,

$$\hat{\mathcal{P}}(\lambda) = \sum_{n=1}^{d} \delta(\lambda_n - \lambda). \tag{4.10}$$

If we discretize $\hat{\mathcal{P}}(\lambda)$ at N points the Nyquist-Shannon sampling theorem asserts that the time domain autocorrelation function, $\mathcal{P}(t)$, must be sampled at $\frac{N}{2}$ points if one is to obtain the d peaks from an FFT. Compressed sensing exploits the fact that $d \ll N$ by replacing the FFT with a convex optimization allowing us to extract energy levels from $cd\log(\frac{N}{d})$ randomly located samples of $\mathcal{P}(t)$ where c is a small constant. Explicitly, we solve

$$\min |u|_1 \; \textit{subject to} \; R\mathcal{F}^{-1}u = \mathcal{P}(t_n) \tag{4.11}$$

where $R\mathcal{F}^{-1}$ is a randomly subsampled IDFT matrix. Central results of compressed sensing state that the solution of the $l1$ optimization problem, eq. (4.11), is identical to the solution of the associated NP-hard $l0$ optimization problem with overwhelmingly high probability [CRT06]. This guarantees that the number of peaks in the computed $\hat{\mathcal{P}}(\lambda)$ is as small as possible, and, since we know $\hat{\mathcal{P}}(\lambda)$ is necessarily a sum of Dirac delta functions, the sparsest $\hat{\mathcal{P}}(\lambda)$ that fits the observed $\mathcal{P}(t)$ is also the true solution. The amount of observed data required to make this assertion is known from compressed sensing theory to be $O(d\log(N))$

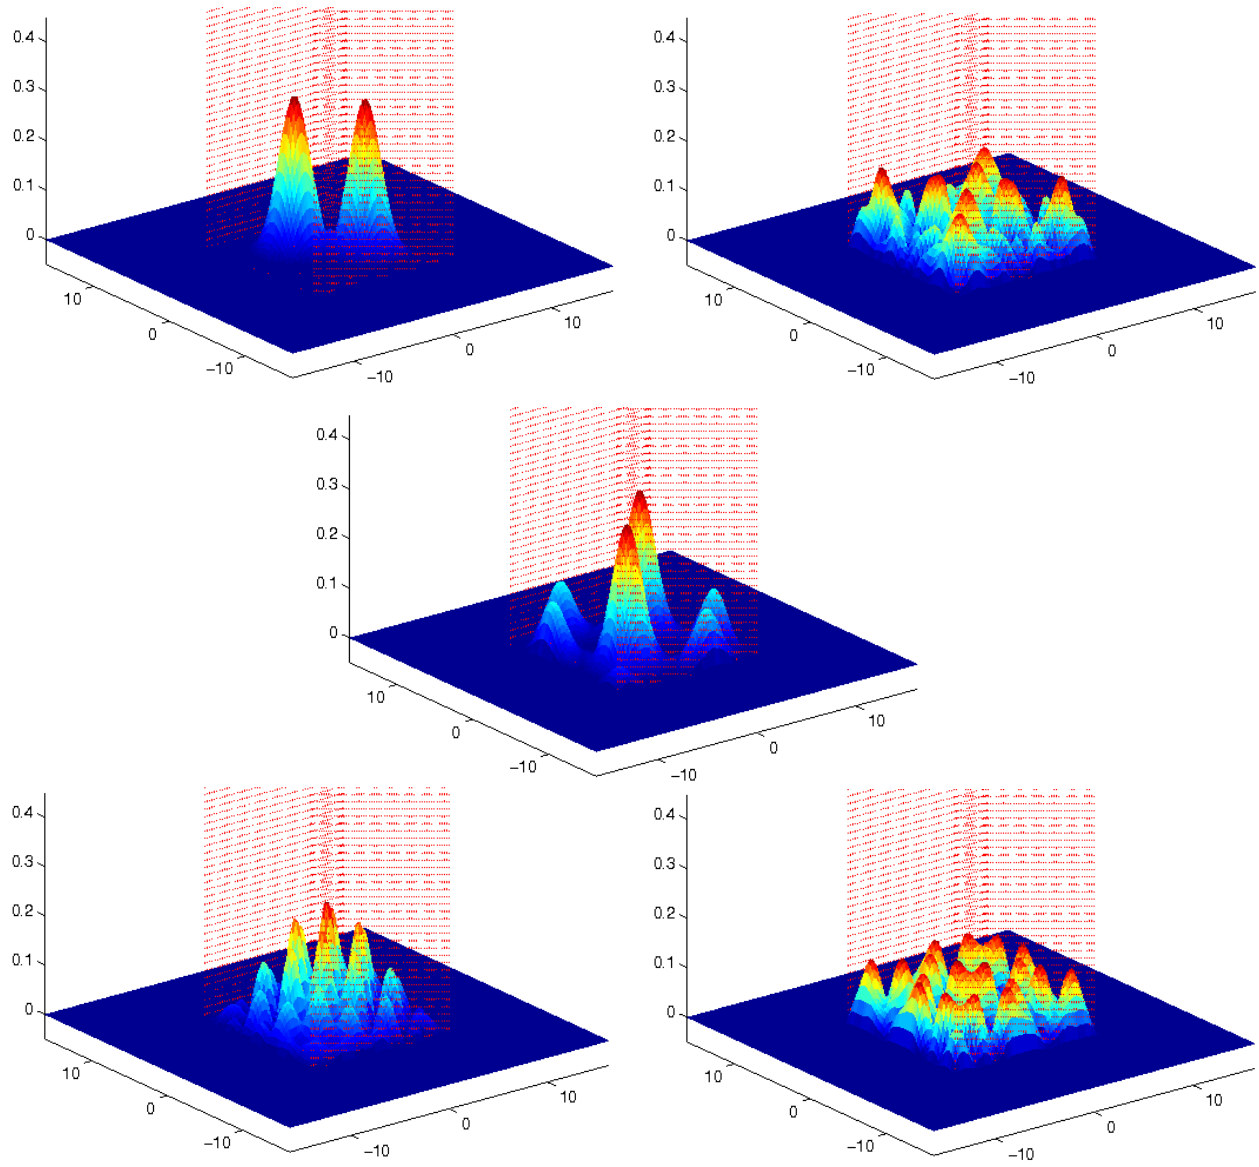

Figure 4.1: Snapshots of evolution of $\psi(\vec{x}, t)$ in a 2D box. Traditional spectral methods extract spectral information from $\mathcal{P}(t) = \int \psi(t)\psi_0^* dx$ via the FFT. Our approach obtains the same information via $l1$ minimization when significantly less time data is collected.

[CRT06].

We physically interpret the autocorrelation function,

$$\mathcal{P}(t) = \sum_{n=1}^{d} |A_n|^2 e^{-i\lambda_n t} \tag{4.12}$$

as a coherent superposition of components with various energy levels [CHE96]. Resolving the energy levels when many different superpositions have been calculated is possible via Fourier transform, however only a few λ_n are interesting and extraction of these energies from the coherent state data via the proposed method is possible from much fewer evaluations of the autocorrelation function.

We remark that sparse signal processing has been applied to coherent state representations in the past with great success. The MP/SOFT framework for multiparticle dynamics makes use of the "matching pursuit" algorithm [Jie08]. MP/SOFT optimizes over an overcomplete dictionary of coherent states, generated by importance sampling Monte Carlo, in order to identify a small Hamiltonian system where the eigenvalues can be cheaply found [WB04] [SC08]. Our method samples the coherent state data by advancing ψ across a nonuniform grid.

4.2.2 A randomized time stepping scheme

We propose propagation of $\psi(t)$ via a randomized time stepping scheme. Given a division of $[0, T]$ into N points spaced Δt apart we draw $M \in \mathcal{O}(d \log(N))$ points uniformly at random where $\psi(t)$ will be evaluated. Denote by Δt_n the spacing between the subsampled grid points, our method proceeds by applying propagators of varying step sizes to $\psi(t)$ and recording $\mathcal{P}(t_n)$ at each iteration. We finally move into the energy domain with $l1$ minimization.

Solving the $l1$ optimization problem is not trivial due to the non-differentiability of the $l1$ norm. However, in recent years, algorithms for $l1$ minimization have seen a surge in development in recent years with new codes outperforming classical methods by large factors [GO] [GS]. This bodes well for our approach as a robust an efficient $l1$ minimization tool is an essential ingredient. We employ the Split-Bregman algorithm in this work [GO]. With this

Algorithm 8: Sparse Spectral Method.

Initialize: $\psi_0 = \texttt{randn}(nx, ny, nz)$, $R\mathcal{F}^{-1}$ a randomly subsampled IDFT matrix

for $n = 1 \dots M$ **do**

$\psi(t_n) = e^{-iH\Delta t_n}\psi(t_{n-1})$

$\mathcal{P}(t_n) = \int \psi(t_n)\psi_0^* dx$

end

$\hat{\mathcal{P}}(\lambda) = \min\{|u|_1 s.t. R\mathcal{F}^{-1}u = \mathcal{P}\}$.

algorithm, the computational cost incurred by replacing the FFT with an $l1$ minimization is negligible when compared with wavefunction propagation.

A brief remark is in order here. For the $l1$ minimization to succeed in obtaining the correct signal it is necessary that $R\mathcal{F}^{-1}$ satisfy the "restricted isometry property"[Can06]. This can be guaranteed with high probability provided that our time domain samples are selected uniformly at random [CRT06].

When the t_i are drawn uniformly with probability p then Δt_i takes on a geometric distribution with parameter p. The cumulative distribution function for each Δt_i is then

$$P(\Delta t_i < k) = 1 - (1-p)^{k+1} \tag{4.13}$$

and for a propagation across M points we can estimate the largest time step using the Mth order statistic

$$P(\max\{\Delta t_i\} < k) = (1 - (1-p)^{k+1})^M \tag{4.14}$$

For $p = \frac{1}{5}$ and $M = 3882$ the probability that all gaps traverse less than 20 steps is 4.416×10^{-6}. Advancing ψ through a second order method (eg SOD [KK83], Strang Splitting [FFS82]) will therefore increase the L^∞ error $||\psi_{approx}(t) - \psi_{true}(t)||_\infty$ by a factor of at least 400 and likely introduce stability problems.

If we are to propagate with an approximate method we must overlay our random grid with a uniformly spaced grid of acceptable step size to avoid the large excursions taken by Δt_i. Alternatively, we may advance ψ with more expensive yet spectrally accurate methods.

4.2.3 Choice of propagator

Accurate computation of $e^{-iH\Delta t}$ for varying Δt is necessary for successful determination of $\mathcal{P}(t)$. The calculation of $e^{-iH\Delta t}$ is a well studied problem in approximation theory which has been approached by Strang splitting [FFS82], central difference schemes [KK83], iterative Lanczos reduction [PL86], expansion in Hermite, Newton and Faber polynomials [HKH94][Kos94][VWB99] as well as a host of Krylov subspace methods [ML03].

We favor calculation of $e^{-iH\Delta t}$ via an expansion in Chebyshev polynomials [TK84]. Expansion in Chebyshev polynomials is a well known technique providing spectral accuracy and well as computational efficiency. It also provides sufficient robustness to advance $\psi(t)$ across our randomized grid.

Specifically, the Chebyshev polynomial expansion of our complex exponential is:

$$e^{-iH\Delta t} \approx \sum_{k=0}^{K} a_k \rho_k(-iH\Delta t). \tag{4.15}$$

where $\rho_k(\omega) = T_k(-i\omega)$, $\omega \in [-i, i]$ are the complex Chebyshev polynomials. The Chebyshev polynomials of the first kind $T_k(x)$ are defined by the recurrence

$$\begin{aligned} T_1(x) &= 1 && (4.16)\\ T_2(x) &= x && (4.17)\\ T_k(x) &= 2xT_{k-1}(x) + T_{k-2}(x) && (4.18) \end{aligned}$$

To understand bounds on K we consider the problem of approximating e^z for $z \in [i\lambda_{min}\Delta t, i\lambda_{max}\Delta t]$. Introduce notation

$$\begin{aligned} R &= \frac{\Delta t}{2}(\lambda_{max} - \lambda_{min}) && (4.19)\\ G &= \Delta t \lambda_{min} && (4.20)\\ \omega &= \frac{1}{R}(z - i(R+G)) && (4.21)\\ & && (4.22) \end{aligned}$$

Note that $\omega \in [-i, i]$ and write

$$e^z = e^{i(R+G)} e^{R\omega} \tag{4.23}$$

$$= e^{i(R+G)} \sum_{k=0}^{K} C_k J_k(R) \rho_k(\omega) \tag{4.24}$$

where $C_k \in \{1, 2\}$. Note that the Bessel functions of the first kind, $J_k(R)$, are driven exponentially towards zero for $k > R$ yet vanish only on a set of measure zero for $k < R$. We therefore take $K = \alpha R$ with $\alpha > 1$ to ensure convergence of the series.

Regrettably, the computational cost as measured by the number of calls to the Hamiltonian of applying a Chebyshev polynomial expansion scales linearly with respect to both the time step size and the spectral radius of H. We remark however that the purpose of this work is to demonstrate a reduction in the number of time steps, the development of a propagator with asymptotically decreasing computational cost per unit time step is deferred to future work.

4.3 Numerical Results

4.3.1 Finite square well potential

The first validation of our approach is done on a finite square well potential:

$$V(x) = 2.8 \cdot \mathbb{1}_{\{x<L/16\}} + 2.8 \cdot \mathbb{1}_{\{x>L/16\}} \tag{4.25}$$

for $x \in [-L/2, L/2]$ with $L = 48$, and $M = 25$. We discretize at 512 points in space and compute the action of Laplacian term in $H = \frac{-1}{2M}\triangle + V(x)$ by multiplication in the spectral domain. Our initial wavefunction is taken as

$$\psi(0) = e^{-(5(x-0.7))^2} \tag{4.26}$$

With a final time of $T = 1941$ and the smallest $\Delta t = 0.5$ the time domain is divided into 3882 points. We downsample this domain by factors of 2,4,6 and 8 and record the relative

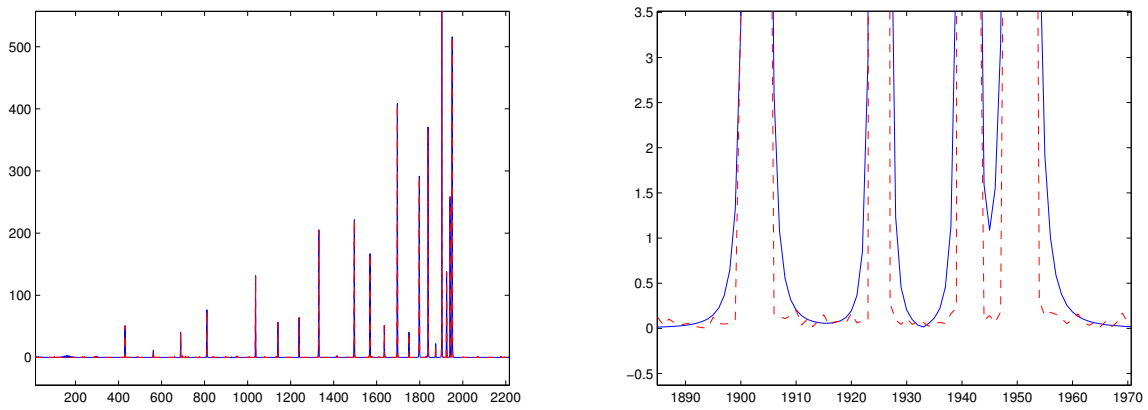

Figure 4.2: Plot of $\hat{\mathcal{P}}(\lambda)$ for the double well potential constructed from the full data using an FFT in blue and the $l1$ reconstruction in red. The $l1$ reconstruction was done with 8x less points taken from $\mathcal{P}(t)$. Interestingly, the reconstructed data tends to favor a sparser solution and is thus closer to the ground truth signal.

error in $\hat{\mathcal{P}}(\lambda)$ as $\frac{||\hat{\mathcal{P}}_{l1}(\lambda)-\hat{\mathcal{P}}_{FFT}(\lambda)||_2}{||\hat{\mathcal{P}}_{FFT}(\lambda)||_2}$. Computation of the eigenvalues done using Matlab's `findpeaks`.

In fig. 4.2 we plot the energy-domain autocorrelation function, $\hat{\mathcal{P}}(\lambda)$, constructed from an FFT of the fully-sampled autocorrelation function (blue) and from the output of alg. 8 with 8x undersampling. Identifying eigenvalues of H amounts to finding peak locations in $\hat{\mathcal{P}}(\lambda)$. Larger values of T lead to sparser $\hat{\mathcal{P}}(\lambda)$ allowing us to identify eigenvalues more easily. We note that in fig. 4.2 the version of $\hat{\mathcal{P}}(\lambda)$ constructed from $l1$-minimization is sparser than the FFT-based reconstruction indicating that we get closer to the true signal with alg. 8.

4.3.2 Double well potential

We run our experiment again on the double well potential originally studied by Feit [FFS82]. The double well shape is described by:

$$V(x) = k_0 - k_2x^2 + k_3x^3 + k_4x^4 \tag{4.27}$$

# subsampled time steps	relative error in $\hat{\mathcal{P}}(\lambda)$
1941	2.7109e-4
971	.0062
647	.0138
486	.0234

Table 4.1: Difference in the $l1$ and FFT based reconstructions for downsample factors of 2, 4, 6, and 8 on the finite square well potential. The difference between the proposed method and the classical method is negligible even for 8x undersampling. Futhermore, as fig. 4.2 shows, the difference is primarily in the sidelobes and not peak locations.

with $k_0 = -132.707$, $k_2 = 7$, $k_3 = 0.5$, $k_4 = 1$, $x_1 = 3.813$, $x_2 = -4.112$. Our spatial discretization is $dx = 0.0825$ across 512 points.

Unbounded potential wells lead to unbounded energy spectra which require an infinitesimally small time step restriction via eq. (4.9). Again, following [FFS82], we truncate the potential to 0 outside of $[x1, x2]$ and take our smallest allowable time step as $\Delta t = 1.25$.

A final time of $T = 2560$ is used in all double well experiments. We uniform randomly downsample the time domain by factors of 2, 4 and 6 prior to wavefunction propagation and compare the resulting $\hat{\mathcal{P}}(\lambda)$ with one obtained from an FFT of the complete data.

4.4 Conclusions and future work

In this work we have developed and analyzed a novel spectral method for the computation of Hamiltonian eigenvalues which considerably reduces the number of time steps required to accurately transform into the energy domain. Specifically, our work shows that the natural sparse structure present in quantum eigensystems may be exploited with $l1$ minimization allowing us to circumvent the small Δt uncertainty principle.

One drawback to our work is that existing spectrally accurate methods of approximation

# subsampled time steps	relative error in $\hat{\mathcal{P}}(\lambda)$
2048	2.7109e-4
1024	.0062
512	.0138
256	.0216

Table 4.2: Difference in the $l1$ and FFT based reconstructions for downsample factors of 2, 4, 6, and 8 on the double well potential. Again, the difference between the proposed method and the classical method is negligible even for 8x undersampling.

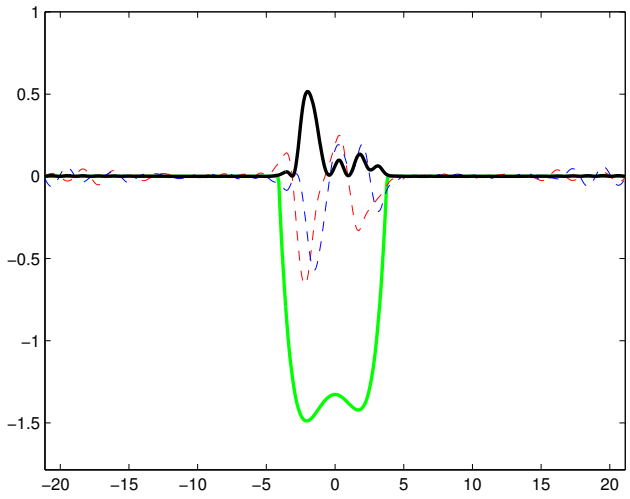

Figure 4.3: Wavefunction evolution in the double well potential used by Feit Fleck and Steiger in the original implementation of the spectral method. The potential is green, the probability density, $|\psi(x)|^2$, is black, the real and imaginary parts of ψ are labeled red and blue.

for $e^{-iH\Delta t}$ have a linearly scaling computational cost with regard to both Δt and ΔE. However, approximate methods (eg SOD, Strang splitting) show no increase in computation time with a large Δt and we expect identification of quantum mechanical systems amenable to approximate methods will yield profitable numerical algorithms. Higher order approximate methods as well as hybrid approximate/spectral methods are also planned for the future.

CHAPTER 5

A projection-correction method for the computation of a few localized eigenvectors

5.1 Introduction

Eigenfunction localization has been the focus of a large body of theoretical and applied science for several decades. Localized eigenfunctions are observed in a variety of research including ground states of Schrodinger operators [Gri04], billiard dynamics [HS10], quantum dot simulation [And], and Turing patterns on random networks [NM11].

One specific application associated with the design of quantum dot devices is to determine whether or not a localized state exists in a pre-determined sub-region of the domain, and to compute that state if it does exist [And].

For eigensystem computations of large problems it is prohibitively expensive to compute all the eigenfunctions and then scan for those localized in a given region. For large problems one typically uses eigensystem computation procedures that focus on a specific range of the spectrum. However, the use of these methods is problematic because it ignores any properties of spatial localization. Since the spatial location of an eigenfunction may not be determined from its eigenvalue alone, one cannot efficiently use routines which target only a given part of the spectrum when one's interest is focused on a small region of space as well.

In this chapter, we propose a two-step solution to this problem based on projection operators and nonlinear eigenvalue correction equations. Given a hermitian matrix, H, corresponding to a discretization of a differential operator, and a spatial region of interest,

This chapter is joint work with Chris Anderson

$\Omega \subset \mathbb{R}^n$, we provide a method to compute only the low-energy eigenfunctions which localize within Ω by restricting H to Ω and then compensating for errors introduced by this restriction.

Central to the first step of our method is a restriction of H onto Ω which allows us to compute approximate, localized, low-energy eigenvectors. Since the restriction of H to Ω is low rank, directly computing the lowest eigenpairs of $H|_\Omega$ with standard iterative methods is problematic as the induced degeneracies at 0 can contaminate the spectrum near regions of interest. The solution to this difficulty is scaling and reordering spectrum of $H|_\Omega$ before computing eigenpairs. By shifting nonlocalized eigenpairs from 0 to 1 it becomes to possible to compute only the eigenpairs of interest with standard methods. The shifting is essential as such a reordering of the spectrum can not be accomplished with projections alone.

When H is a discrete Schrodinger operator, our projection method is similar to the idea of using a confining potential to isolate wavefunctions inside Ω in that it notably reduces the number of candidate wavefunctions one must iterate over when seeking localized ground states. Our approach, however, is purely linear algebraic allowing applications to non-Schrodinger systems as well.

Depending on the specific problem at hand, the eigenpairs computed in the first step may be rather crude approximations to true eigenpairs of H. If higher accuracy is desired, the eigenpairs can be refined by an additional step. There are a variety of ways that the refinement can be carried out, in the investigation at hand we solve a nonlinear correction equation with Newton's method or a Jacobi-Davidson type correction.

Related work may be found in domain decomposition methods developed by the structural mechanics community. Component mode synthesis (CMS) [GN] approximates eigenfunctions in a specified region of space by restricting the original operator to the region of interest and treating the domain of the restricted operator as approximately invariant for the original operator. This subspace is then corrected by enriching it with selected eigenvectors from the complimentary subspace leading to more accurate approximate eigenvectors. Many steps of CMS lead to the automated multilevel substructuring algorithm (AMLS) and higher

accuracies [EV]. Both these approaches may be understood by taking the viewpoint that the eigenvalue problem is a nonlinear system of equations which is amenable to traditional root finding algorithms. This straightforward approach has recently be put forward by Philippe et. al. [Phi06] and is the method that is used here.

5.2 Method

5.2.1 Spectral approximation via projection operators

Our first step consists of finding approximate eigenpairs by restricting to Ω and applying iterative methods. Take $H \in \mathbb{R}^{n\times n}$ to be symmetric positive definite. Let P be projection onto vectors supported in Ω.

When, for example, H is a discretization of a Schrodinger operator, $\triangle + V(x)$, over a uniform grid. P is a diagonal matrix with $P_{i,i} = 1$ for grid points in Ω and $P_{i,i} = 0$ else.

Localized minimal eigenpairs of H may be approximated with minimal eigenpairs of PHP. Working with PHP directly is problematic as PHP has a degenerate subspace at 0. Consequently, we may not simply apply standard iterative procedures to PHP in order to approximate nontrivial localized eigenpairs of H as the inroduced degenerate subspace at 0 contaminates the eigenvalues of interest which may be very close to 0. To circumvent this complication, we normalize and flip the spectrum of H so that the degenerate eigenvalues induced by the projection are distant from 0. Specifically, form

$$Pr(H) = I - P\frac{H - \lambda_{max}}{\lambda_{max} - \lambda_{min}}P \tag{5.1}$$

where λ_{max} and λ_{min} are estimates to the maximal and minimal eigenvalues of H obtained from the Lanczos algorithm [GL96]. The spectrum of $Pr(H)$ lies in $[0, 1]$ and has been flipped and shifted so that $\lambda_{max} \to 1$ and $\lambda_{min} \to 0$. We can obtain minimal eigenpairs of H by rescaling eigenpairs of $Pr(H)$. The flip after the projection is essential to remove degeneracy problems near 0. Seeking minimal eigenpairs of $Pr(H)$ can be done easily as the degeneracy problem has been shifted to 1.

Given the smallest eigenvalues of $Pr(H)$ we obtain our approximate eigenvalues $\{\lambda_{Pi}\}_{i=1}^{k}$ via the relation:

$$\lambda = (1 - \lambda_{Pr(H)})\lambda_{min} + \lambda_{Pr(H)}\lambda_{max}. \tag{5.2}$$

Alternatively, one may return to the original scalings by forming the Rayleigh quotient using H and the computed eigenvectors of $Pr(H)$. Accuracy depends on how well Ω encloses the support of the eigenvectors of H. In the event that Ω is perfectly chosen P is projection onto an invariant subspace of H and for each $i = 1 \dots k$ we have $\lambda_{Pi} = \lambda_j$ for some $j \in 1 \dots n$.

5.2.2 Eigenpair corrections

There are several situations where the first step alone is not sufficient to reach an accurate eigenpair. Imperfections in the choice of Ω, and/or Schrodinger equations with a low particle mass (and this a high degree of tunneling outside of Ω) may require refinement after the initial projection step.

Once we have obtained our approximate eigenpairs we correct by interpreting the eigenproblem as a nonlinear system of algebraic equations. Solving the nonlinear system is possible via Newton's method. In addition, by observing the change associated with this refinement, one can validate the identification of the eigenpair as one that is indeed localized.

It should be noted that correcting an approximate eigenpair in this manner has previously received attention from the numerical linear algebra community under the names of Olsen's Method [Saa96], Trace Minimization [Ahm82], and the Jacobi-Davidson method [SV00].

We write the original eigenvalue problem as a nonlinear system:

$$\begin{cases} (H - \lambda I)\psi &= 0 \\ \psi^*\psi &= 1 \end{cases} \tag{5.3}$$

and note that (5.3) may be solved with Newton's method producing the iteration

$$\begin{pmatrix} \psi^{new} \\ \lambda^{new} \end{pmatrix} = \begin{pmatrix} \psi \\ \lambda \end{pmatrix} - \begin{pmatrix} H - \lambda I & -\psi \\ 2\psi^* & 0 \end{pmatrix}^{-1} \begin{pmatrix} r \\ 0 \end{pmatrix}. \tag{5.4}$$

Where we have defined the residual eigenvector, r, as

$$r = H\psi - \lambda\psi. \tag{5.5}$$

From an initial approximation, (λ_P, ψ_P), iterating (5.4) compensates for any error introduced by the projection P and yields high accuracy eigenvectors of H.

Algorithm 9: Single vector iterative Newton correction

Initialize: Approximate eigenpairs $\{\lambda_{Pi}, \psi_{Pi}\}_{i=1}^{k}$

for $i = 1 \dots k$ **do**

while $\|r\|_2 > tol$ **do**

$$\begin{pmatrix} \psi^{new} \\ \lambda^{new} \end{pmatrix} = \begin{pmatrix} \psi_P \\ \lambda_P \end{pmatrix} - \begin{pmatrix} H - \lambda_P I & -\psi_P \\ 2\psi_{P*} & 0 \end{pmatrix}^{-1} \begin{pmatrix} r \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} \psi_P \\ \lambda_P \end{pmatrix} = \begin{pmatrix} \psi^{new} \\ \lambda^{new} \end{pmatrix}$$

$r = H\psi_P - \lambda_P\psi_P$

end

end

An alternative way to update our eigenpairs is to compute and add in a small correction to our current approximation. Here, we solve a linearized version of eq. (5.3) in the space orthogonal to ψ_P to obtain our correction and repeat until convergence. Related techniques have been explored in several alternative eigenvalue solvers including the Davidson [CPS94], and Jacobi-Davidson [SV00] methods.

Specifically, with an approximate eigenpair, (λ_P, ψ_P) where $\lambda_P = \frac{\psi_P^* H \psi}{\psi_P^* \psi}$, and corresponding residual $r = H\psi_P - \lambda_P\psi_P$, we seek a correction, (η, v), as a solution to the eigenproblem

$$H(\psi_P + v) = (\lambda_P + \eta)(\psi_P + v). \tag{5.6}$$

As our correction is expected to be small we can neglect the second order term ηv in (5.6) to arrive at the linear system

$$\begin{aligned} (H - \lambda_P I)v - \eta\psi_P &= \lambda_P\psi_P - H\psi_P & (5.7) \\ &= -r. & (5.8) \end{aligned}$$

Since η is not given, the system (5.7) has n equations and $n+1$ unknowns. To obtain a unique solution we must enforce additional constraints. A natural choice is

$$(\psi_P + v)^*(\psi_P + v) = 1. \tag{5.9}$$

We expand (5.9), ignore the second order term, v^2, and get a linear constraint,

$$\psi_P^* v = 0. \tag{5.10}$$

Equation (5.10) is interpreted a requirement to search for improved eigenvectors in a direction orthogonal to our current approximation. Since r is orthogonal to ψ_P we can left multiply both sides of (5.7) by $(I - \psi_P\psi_P^*)$ and use the fact that v is orthogonal to ψ_P to rewrite our correction equation as an underdetermined linear system,

$$(I - \psi_P\psi_P^*)(H - \lambda_P I)v = -r. \tag{5.11}$$

The degeneracy is overcome when restricting solutions to lie in the orthogonal compliment of ψ_P. This can be enforced by solving

$$(I - \psi_P\psi_P^*)(H - \lambda_P I)(I - \psi_P\psi_P^*)v = -r. \tag{5.12}$$

as is done in the Jacobi-Davidson method.

Algorithm 10: Linearized correction

Initialize: Approximate localized eigenpairs $\{\lambda_{Pi}, \psi_{Pi}\}_{l=1}^{L}$

Output: Corrected eigenpairs $\{\tilde{\lambda}_i, \tilde{\psi}_i\}_{l=1}^{L}$

for $l = 1 \dots L$ **do**

$\tilde{\psi} = \psi_P$, $\tilde{\lambda} = \lambda_P$

while $\|r\| \geq tol$ **do**

$r = H\tilde{\psi} - \tilde{\lambda}\tilde{\psi}$

Solve for z, $(I - \tilde{\psi}\tilde{\psi}^*)(H - \tilde{\lambda}I)(I - \tilde{\psi}\tilde{\psi}^*)z = -r$.

$\tilde{\psi} = \tilde{\psi} + z$

$\tilde{\lambda} = \tilde{\psi}^* H \tilde{\psi}$

end

end

We remark that the solution to the linear system in Algorithm 10 may be obtained by solving for z in

$$\begin{pmatrix} H - dI & \tilde{\psi} \\ \tilde{\psi}^* H & 1 \end{pmatrix} \begin{pmatrix} z \\ 0 \end{pmatrix} = \begin{pmatrix} -r \\ 0 \end{pmatrix} \tag{5.13}$$

We find that alg. 10 provides highly accurate eigenvectors of H with only a few calls to a linear solver. In large and poorly conditioned problems, the use of a linear solver may be computationally prohibitive and subspace methods may be necessary to handle the eigensystem computations. We plan to investigate subspace methods in future work.

5.3 Numerical results

All tests were run in Matlab on one core of a 1.2 GHz Intel Centrino Core 2 Duo CPU with 2 MB of level 2 cache and 3 GB of RAM. Computations were done in IEEE standard double-precision variables yielding a machine epsilon of 1.1102e-16 and a relative precision of 2.2204e-16.

5.3.1 Nonnegative multiple well potential

Our first test makes use of a Schrodinger operator with a positive multiple well potential on a 64×64 grid (cf figure 5.1)

$$H = -\triangle + V(x). \tag{5.14}$$

Here, low energy eigenfunctions of H tend to localize within each potential well however there is no correlation between their location and energy level.

To obtain our region of interest, Ω, we partition $V(x)$ using a watershed transformation. Originating in image processing, watershed transforms have been adapted to electronic potential segmentation several times in the past [Pin12].

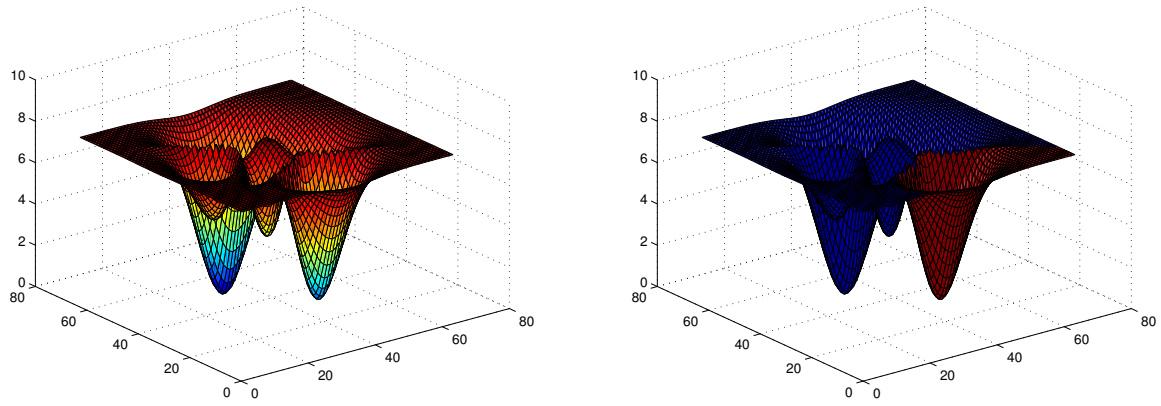

Figure 5.1: Potential function, $V(x)$, used in our first experiment. In the figure on the right we have colored Ω red. Our goal is to compute only those eigenfunctions which are supported inside Ω.

We seek the 4 lowest eigenvectors supported inside Ω. Trial and error shows that we can obtain these by computing the lowest 16 eigenpairs of H and then filtering out eigenvectors supported outside Ω. Computation of all 16 eigenpairs using ARPACK called from Matlab requires 0.586239 seconds of computation time. Computation of the lowest 4 eigenpairs inside Ω by our method requires 0.304272 seconds when the Newton approach is used and 0.254958 seconds with the Jacobi-Davidson approach. Both correction methods were stopped when the norm of the residual became less than 1e-6.

5.3.2 Quantum dot confinement potential

In recent years semiconductor quantum dots have become the focus of a large body of experimental and theoretical research. Numerical models of electronic structure play an essential role in quantum dot design and thus several methods have been developed to solve the resulting equations. Methods for the computation of ground states confined to the quantum dot region are of specific importance to the design of semiconductor quantum dot devices. The ground state of the system being modeled is often not confined to the dot well, prompting us to use the methods described in this chapter for quantum dot simulation. In

exact eigenvalues of H	error with eigenvalues of $Pr(H)$	error after Newton correction	error after Jacobi-Davidson correction
0.8292			
1.5052			
1.7667			
2.1404			
2.3825	1.520642e-06	2.771117e-13	8.437695e-15
2.4206			
2.6279			
2.7365			
3.0336			
3.0851	2.343630e-06	7.323031e-13	3.015366e-13
3.2607			
3.2840	2.608168e-05	9.600720e-11	1.807443e-13
3.2950			
3.4173			
3.6070			
3.7442	8.582280e-06	1.140599e-11	2.375877e-13

Table 5.1: Relative eigenvalue error with multiple well potential. Eigenvalues computed directly from a restriction to Ω are somewhat accurate (second column). However, accuracy improvements of six orders of magnitude are possible by iteratively solving a nonlinear correction equation (third and fourth columns).

fact, in the examples we study, we find that the ground state in the well is the 79th lowest state of the system.

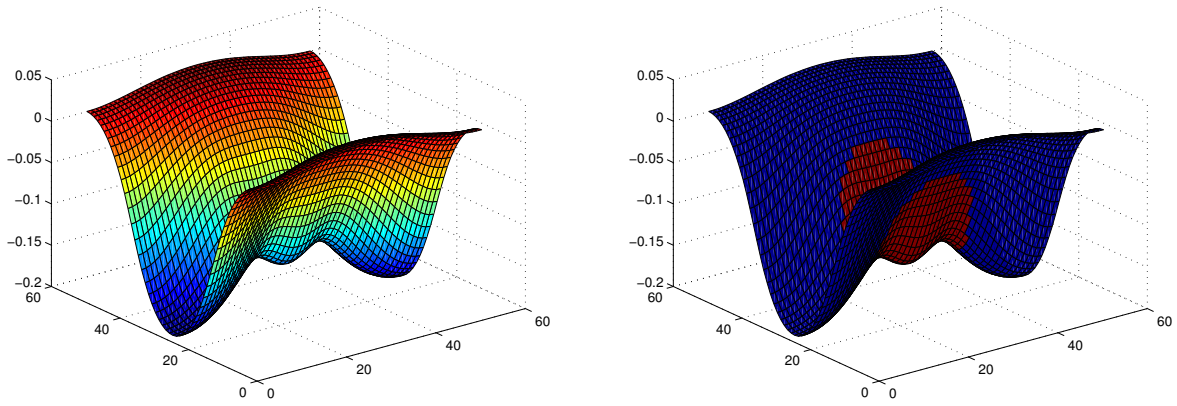

Figure 5.2: Two-dimensional dot confinement potential and manually identified region of interest used in our simulations. Here, the lowest energy eigenfunction contained within the dot is the 79th lowest eigenfunction of H.

Eigenfunctions localizing within the dot region typically have high energies and sidelobes extending into the lower wells. For a Hamiltonian corresponding to the potential in figure 5.2, $H = -\alpha\triangle + V(x)$ with $\alpha = 15e-3$, the lowest energy eigenfunction contained within the dot is the 79th lowest eigenfunction of H. After the ground state in the well, 21 more eigenfunctions localize in the outer region before the first excited state in the well can be observed (cf. figure 5.3).

We seek the 6 lowest energy eigenfunctions inside the dot region pictured in Figure 5.2. We work on a 51×51 uniform grid with unit spacing. Computation with the shift-invert techniques found in Matlab's `eigs()` function require iterations on 139 vectors and requires 7.078678 seconds to compute the states of interest. Newton iterations and Newton-Sylvester iterations respectively took 0.801474 and 0.717489 seconds.

It is worth noting that we had to do a couple of trial-and-error experiments before settling on the 139 figure when using the `eigs()` function. This is because we can not say in advance how many eigenpairs need to be computed before the 6 lowest localized eigenpairs are found.

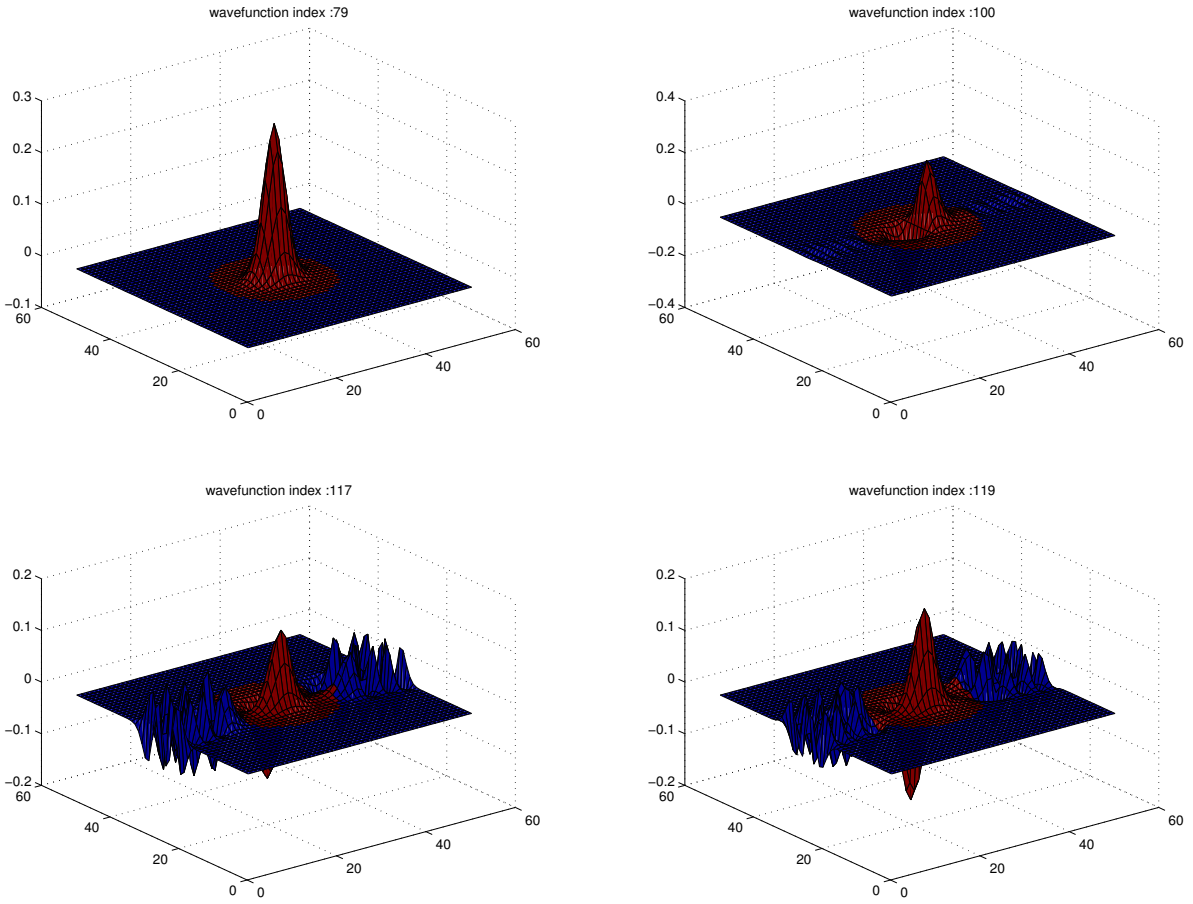

Figure 5.3: Low energy wavefunctions and region of interest for the potential in figure 5.2. Higher energy states tunnel outside the confining region necessitating corrections to restricted eigenvector approximations.

With the method developed in this chapter, however, we need only to specify the region and the number of localized eigenpairs sought.

5.4 Conclusion

We have presented a method for efficiently computing a few of the smallest eigenpairs localized in a given spatial region. Our approach allows one to significantly reduce the number

Index	Eigenvalue	Error with PH	Newton correction error	Newton-Sylvester error
79	-1.369912e-01	8.582711e-10	7.549517e-15	2.775558e-16
100	-1.310642e-01	2.268016e-06	1.170332e-08	2.220446e-16
119	-1.271600e-01	3.659504e-04	4.231893e-13	2.220446e-16
126	-1.264260e-01	5.573836e-07	9.863085e-08	8.769820e-09
130	-1.257553e-01	6.565336e-05	4.065431e-11	4.718448e-16
139	-1.226983e-01	2.959336e-04	2.107283e-06	2.775558e-16

Table 5.2: Comparison of methods for localized eigenvalue computation in the quantum dot potential. Columns correspond to: the eigenvalue index in the complete spectrum, exact value of the sought eigenvalue, error without any correction, error using the Newton scheme, and error using the Newton-Sylvester scheme.

of candidate vectors one must iterate over when solving for states in high-energy potential wells. Most notably, the method is insensitive to prior knowledge of the eigenfunction's geometry as it makes use of correction equations to compensate for an imperfect choice of dividing contour.

CHAPTER 6

Conclusion

In this thesis, we studied the applicability of $l1$-penalized optimization to datasets arising in quantum mechanical experiments and developed methods for efficient computation of eigenpairs in quantum mechanical simulations. We have also demonstrated how $l1$-penalized optimization techniques can be made useful in situations where perfect energy-domain sparsity is unattainable and shown how localized eigenpairs can be efficiently computed with a projection/correction method.

In our first application, NMR spectroscopy, we focused on the problem of energy-domain signal reconstruction from undersampled time-domain FID data. In the past, several researchers have used $l1$-penalized optimization for this purpose. Past research, however, has neglected the fact that NMR energy-domain signals are not perfectly sparse due to the damping effects naturally present in FIDs. We demonstrated how incorporating these effects into the system matrix via a rank-1 exponential correction prior to $l1$-penalized signal reconstruction often leads to higher final SNR albeit at a higher computational cost.

Related to NMR, our second application, field-corrected MRI image reconstruction, suffers from a similar problem. Here, magnetic field inhomogeneities resulting from abrupt changes in magnetic susceptibility of the tissues being imaged can lead to image blurring and distortion when neglected. We investigated correction methods for these inhomogeneous fields in the context of sparsity promoting image reconstruction. Our approach involves modeling the field inhomogeneities by altering the system matrix with a low rank correction matrix prior to sparsity promoting image reconstruction. This alteration leads to high quality images. We have also shown, theoretically and empirically, how incorporating

a framelet-based regularization term into the optimization reduces computational cost.

Outside of experimental settings, our final application of $l1$-penalized optimization deals with eigenvalues of Hermitian matrices. Simulating quantum dynamical systems allows one to efficiently compute eigenvalues of Hermitian matricies by analyzing signals generated from the simulation. These signals are sparse, but require a large amount of simulated data to obtain with traditional spectral methods. We have shown that it is possible to reduce the amount of simulated data needed by reconstructing the eigenvalue signal with $l1$-penalized optimization in place of the standard discrete Fourier transform.

Finally, we developed methods for efficient computation of a few eigenvalues of a Hermitian matrix which localize withing a spatial and spectral region of interest. For a given partition, the eigenvalues of its corresponding localized eigenvectors often make up only a small fraction of the complete range of energy levels. We have produced a projection/correction method for the computation of only these eigenpairs. In contrast to standard methods for eigenvalue computation which specify only a part of the spectrum, our method also allows one to isolate regions of space where prior information on eigenvector locality is known.

References

[Ahm82] Ahmed H . Sameh and John A . Wisniewski. "A trace minimization algorithm for the generalized eigenvalue problem." *SIAM Journal on Numerical Analysis*, **19**(6):1243–1259, 1982.

[ALG11] Jan Aelterman, Hiêp Quang Luong, Bart Goossens, Aleksandra Pižurica, and Wilfried Philips. "Augmented Lagrangian based reconstruction of non-uniformly sub-Nyquist sampled MRI data." *Signal Processing*, **91**(12):2731–2742, December 2011.

[And] Christopher R. Anderson. "Efficient solution of the Schroedinger-Poisson equations in layered semiconductor devices." *Journal of Computational Physics*.

[BA07] Louis-S. Bouchard and M. Anwar. "Synthesis of matched magnetic fields for controlled spin precession." *Physical Review B*, **76**(1):1–10, July 2007.

[BBE98] B. Blumich, P. Blumer, G. Eidman, A. Guthausen, R. Haken, U. Schmitz, K. Saito, and G. Zimmer. "The NMR-Mouse: Construction, Excitation, and Applications." *Magnetic resonance imaging*, **16**:479–484, January 1998.

[BFS10] Fabio Baselice, Giampaolo Ferraioli, and Aymen Shabou. "Field map reconstruction in magnetic resonance imaging using Bayesian estimation." *Sensors*, **10**(1):266–79, January 2010.

[BH95] Briggs and Henson. *The DFT Owner's Manuel*. SIAM, 1995.

[Bou06] Louis-S. Bouchard. "Unidirectional magnetic-field gradients and geometric-phase errors during Fourier encoding using orthogonal ac fields." *Physical Review B*, **74**(5):1–11, August 2006.

[BUF07] Kai Tobias Block, Martin Uecker, and Jens Frahm. "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint." *Magnetic resonance in medicine*, **57**(6):1086–98, June 2007.

[Can06] E. J. Candès. "Compressive sampling." *Proceedings of the International Congress of Mathematicians*, **Madrid, Sp**, January 2006.

[CEN11] Emmanuel Candes, Yonina Eldar, Deanna Needell, and Paige Randall. "Compressed sensing with coherent and redundant dictionaries." *Applied and Computational Harmonic Analysis*, pp. 1–21, 2011.

[CHE96] R CHEN. "An accurate spectral method with arbitrarily large time step sizes." *Chemical Physics Letters*, **252**(3-4):201–205, April 1996.

[cM09] Ali Çivril and Malik Magdon-Ismail. "On selecting a maximum volume submatrix of a matrix and related problems." *Theoretical Computer Science*, **410**(47-49):4801–4811, November 2009.

[COB12] Ryan Compton, Stanley Osher, and Louis S Bouchard. "Hybrid regularization for MRI reconstruction with static field inhomogeneity correction." *IEEE ISBI*, pp. 1–16, 2012.

[COS10] J F Cai, S Osher, and Z Shen. "Split Bregman Methods and Frame Based Image Restoration." *Multiscale Modeling & Simulation*, **8**(2):337, 2010.

[CPB11] Lotfi Chaâri, Jean-Christophe Pesquet, Amel Benazza-Benyahia, and Philippe Ciuciu. "A wavelet-based regularized reconstruction algorithm for SENSE parallel MRI with applications to neuroimaging." *Medical image analysis*, **15**(2):185–201, April 2011.

[CPS94] M. Crouzeix, B. Philippe, and M. Sadkane. "The Davidson Method." *SIAM Journal on Scientific Computing*, **15**(1):62, 1994.

[CR06] EJ Candes and JK Romberg. "Stable signal recovery from incomplete and inaccurate measurements." *Communications on pure*, **40698**:1–15, 2006.

[CRT05] Emmanuel Candes, Justin Romberg, and Terence Tao. "Stable Signal Recovery from Incomplete and Inaccurate Measurements." *arXiv : math / 0503066v2 [math . NA] 7 Dec 2005*, **40698**:1–15, 2005.

[CRT06] E.J. Candes, J. Romberg, and T. Tao. "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information." *IEEE Transactions on Information Theory*, **52**(2):489–509, February 2006.

[CSM08] Weitian Chen, Christopher T Sica, and Craig H Meyer. "Fast conjugate phase image reconstruction based on a Chebyshev approximation to correct for B0 field inhomogeneity and concomitant gradients." *Magnetic Resonance in Medicine*, **60**(5):1104–11, November 2008.

[Dro07] Iddo Drori. "Fast 1 Minimization by Iterative Thresholding for Multidimensional NMR Spectroscopy." *EURASIP Journal on Advances in Signal Processing*, **2007**:1–11, 2007.

[Dun01] Ian C. Duncan. "The "Aura" Sign: An Unusual Cultural Variant Affecting MR Imaging." *American Journal of Roentgenology*, (December):1485–1489, 2001.

[EV] K Elssel and H Voss. "AUTOMATED MULTI-LEVEL SUBSTRUCTURING FOR NONLINEAR EIGENPROBLEMS 1 Introduction." **1**(2).

[FDM09] John M Franck, Vasiliki Demas, Rachel W Martin, Louis-s Bouchard, and Alexander Pines. "Shimmed matching pulses : Simultaneous control of rf and static gradients for inhomogeneity correction." *Physics*, pp. 1–16, 2009.

[Fei83] M D Feit. "Solution of the Schrodinger equation by a spectral method II: Vibrational energy levels of triatomic moleculesa)." *The Journal of Chemical Physics*, **78**(1):301, 1983.

[Fes10] J Fessler. "Model-Based Image reconstruction for MRI." *IEEE Signal Processing Magazine*, **936726**(July):81–89, 2010.

[FFS82] M Feit, J Fleck, and A Steiger. "Solution of the Schrödinger equation by a spectral method." *Journal of Computational Physics*, **47**(3):412–433, September 1982.

[FFY08] Amanda K Funai, Jeffrey a Fessler, Desmond T B Yeo, Valur T Olafsson, and Douglas C Noll. "Regularized field map estimation in MRI." *IEEE transactions on medical imaging*, **27**(10):1484–94, October 2008.

[FJ05] Matteo Frigo and Steven˜G. Johnson. "The Design and Implementation of {FFTW3}." *Proceedings of the IEEE*, **93**(2):216–231, 2005.

[Special issue on "Program Generation, Optimization, and Platform Adaptation".]

[FMS03] Jeffrey A Fessler, Senior Member, and Bradley P Sutton. "Nonuniform Fast Fourier Transforms Using Min-Max Interpolation." **51**(2):560–574, 2003.

[FOS05a] J.a. Fessler, 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, September 2005.

[FOS05b] J.a. Fessler, 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, September 2005.

[GE96] Ming Gu and Stanley C Eisenstat. "An Efficient Algorithm for Computing a Strong Rank-Revealing QR Factorization." *SIAM Journal on Scientific Computing*, **17**(4):848–869, 1996.

[GHP11] M Guerquin-Kern, M Häberlin, K P Pruessmann, and M Unser. "A fast wavelet-based reconstruction method for magnetic resonance imaging." *IEEE transactions on medical imaging*, **30**(9):1649–60, September 2011.

[GL96] Gene Golub and Charles Van Loan. *Matrix Computations*. Johns Hopkins University Press, 3rd edition, 1996.

[GL04] Leslie Greengard and June-Yub Lee. "Accelerating the Nonuniform Fast Fourier Transform." *SIAM Review*, **46**(3):443, 2004.

[GN] D A Glasgow and H D Nelson. "Stability analysis of rotor-bearing systems using component mode synthesis.".

[GO] Tom Goldstein and Stanley Osher. "The split bregman method for l1 regularized problems." pp. 1–21.

[GO09] Tom Goldstein and Stanley Osher. "The Split Bregman Method for L1-Regularized Problems." *SIAM Journal on Imaging Sciences*, **2**(2):323, 2009.

[Gou01] Y Gousseau. "Are natural images of bounded variation?" *SIAM journal on mathematical*, 2001.

[Gra07] J. Granwehr. "Multiplicative or t1 Noise in NMR Spectroscopy." *Applied Magnetic Resonance*, **32**(1-2):113–156, August 2007.

[Gri04] David Griffiths. *Introduction to Quantum Mechanics*. 2004.

[GS] Tom Goldstein and Simon Setzer. "High-Order Methods for Basis Pursuit." pp. 1–17.

[HB99] E Mark Haacke and Robert Brown. "Magnetic Resonance Imaging Physical Principles and Sequence Design.", 1999.

[HKH94] Youhong Huang, Donald J Kouri, and David K Hoffman. "General , energy-separable Faber polynomial representation of operator functions : Theory and application in quantum scattering." **101**(December), 1994.

[HLC] Simon Hu, Michael Lustig, Albert P Chen, Jason Crane, Adam Kerr, A C Douglas, Ralph Hurd, John Kurhanewicz, Sarah J Nelson, John M Pauly, and B Daniel. "Compressed Sensing for Resolution Enhancement of Hyperpolarized C Flyback 3D-MRSI." (415).

[HMT11] N Halko, PG Martinsson, and JA Tropp. "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions." *SIAM Review*, **53**(2):217–288, 2011.

[HMW12] Sven G Hyberts, Alexander G Milbradt, Andreas B Wagner, Haribabu Arthanari, and Gerhard Wagner. "Application of iterative soft thresholding for fast reconstruction of NMR data non-uniformly sampled with multidimensional Poisson Gap scheduling." *Journal of biomolecular NMR*, **52**(4):315–27, April 2012.

[HS96] Jeffrey C. Hoch and Alan S. Stern. *NMR Data Processing*. Wiley-Liss, 1996.

[HS10] Steven M. Heilman and Robert S. Strichartz. "Localized eigenfunctions: Here you see them, there you don't." *Notices of the AMS*, **57**(5):624–629, 2010.

[IMN96] 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):278–82, February 1996.

[JI06] Victor Jaravine and Ilgis Ibraghimov. "Removal of a time barrier for high-resolution multidimensional NMR spectroscopy." **3**(8):605–607, 2006.

[Jie08] Quanlin Jie. "Approximating the ground state of fermion systems by multiple determinant states: Matching pursuit approach." *Physical Review E*, **77**(2):1–5, February 2008.

[Kel09] Nathan Dean Kelso. *SQUID-Detected MRI in the Limit of Zero Static Field*. Thesis, Berkeley, 2009.

[KK83] D Kosloff and R Kosloff. "A Fourier Method Solution for the Time Dependent Schrddinger Equation as a Tool in Molecular Dynamics." *Journal of Computational Physics*, **53**:35–53, 1983.

[KLS10] Kressler B, T Liu, P Spincemaille, Q Jiang, and Y Wang. "Nonlinear regularization for per voxel estimation of magnetic susceptibility distributions from MRI field maps." *IEEE transactions on medical imaging*, **29**(2):273–281, February 2010.

[KO11] Krzysztof Kazimierczuk and Vladislav Yu Orekhov. "Accelerated NMR spectroscopy by using compressed sensing." *Angewandte Chemie (International ed. in English)*, **50**(24):5556–9, June 2011.

[Kos94] Ronnie Kosloff. "Propagation Methods for Quantum Molecular Dynamics." (1), 1994.

[LDP07] M Lustig, D Donoho, and J Pauly. "Sparse MRI: The application of compressed sensing for rapid MR imaging." *Magnetic resonance in medicine*, **58**(6):1182–1195, December 2007.

[Li] Yingying Li. "Coordinate Descent Optimization for 1 Minimization with Application to Compressed Sensing ; a Greedy Algorithm Solving the Unconstrained Problem." pp. 1–17.

[LWM07] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. "Randomized algorithms for the low-rank approximation of matrices." *Proceedings of the National Academy of Sciences of the United States of America*, **104**(51):20167–72, December 2007.

[LWY09] Dong Liang, Haifeng Wang, and Leslie Ying. "SENSE reconstruction with nonlocal TV regularization." *Annual International Conference of the IEEE EMBS*, **2009**(3):1032–5, January 2009.

[Man82] P. Mansfield. *NMR Imaging in Biomedicine (Advances in magnetic resonance)*. Academic Press, 1982.

[Man01] V Mandelshtam. "FDM: the filter diagonalization method for data processing in NMR experiments." *Progress in Nuclear Magnetic Resonance Spectroscopy*, **38**(2):159–196, March 2001.

[Mar91] P Maraner. "Spectral methods in computational quantum mechanics." *Transform*, **37**:209–219, 1991.

[MBB95] M a Moerland, R Beersma, R Bhagwandien, H K Wijrdeman, and C J Bakker. "Analysis and correction of geometric distortions in 1.5 T magnetic resonance images for use in radiotherapy treatment planning." *Physics in medicine and biology*, **40**(10):1651–1654, October 1995.

[MDL03] Hisamoto Moriguchi, Brian M Dale, Jonathan S Lewin, and Jeffrey L Duerk. "Block regional off-resonance correction (BRORC): a fast and effective deblurring method for spiral imaging." *Magnetic resonance in medicine*, **50**(3):643–8, September 2003.

[MH08] Mehdi Mobli and Jeffrey C. Hoch. "Maximum entropy spectral reconstruction of nonuniformly sampled data." *Concepts in Magnetic Resonance Part A*, **32A**(6):436–448, November 2008.

[MJ04] Robert C. McKinstry and Delma Y. Jarrett. "Magnetic Susceptibility Artifacts on MRI: A Hairy Situation." *American Journal of Roentgenology*, **182**(February):532–535, 2004.

[ML03] Cleve Moler and Charles Van Loan. "Nineteen Dubious Ways to Compute the Exponential of a Matrix , Twenty-Five Years." *Compute*, **45**(1):801–836, 2003.

[MMC06] Andrew E Marble, Igor V Mastikhin, Bruce G Colpitts, and Bruce J Balcom. "A constant gradient unilateral magnet for near-surface MRI profiling." *Journal of magnetic resonance (San Diego, Calif. : 1997)*, **183**(2):228–34, December 2006.

[MPM96] Lai-chee Man, John M Pauly, and Albert Macovski. "Multifrequency Interpolation for Fast Off-resonance Correction." *Magnetic resonance in medicine*, pp. 785–792, 1996.

[MRT11] Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. "A randomized algorithm for the decomposition of matrices." *Applied and Computational Harmonic Analysis*, **30**(1):47–68, January 2011.

[MSH01] Carlos A Meriles, Dimitris Sakellariou, Henrike Heise, and Alexander Pines. "Approach to High-Resolution ex Situ NMR Spectroscopy." *Science*, **293**(July):82–85, 2001.

[MSK06] MW Maciejewski, AS Stern, and GF King. "Nonuniform sampling in biomolecular NMR." *Modern Magnetic*, pp. 1305–1311, 2006.

[MST05] Carlos a Meriles, Dimitris Sakellariou, Andreas H Trabesinger, Vasiliki Demas, and Alexander Pines. "Zero- to low-field MRI with averaging of concomitant gradient fields." *Proceedings of the National Academy of Sciences of the United States of America*, **102**(6):1840–2, February 2005.

[NAG05] A Neufeld, Y Assaf, M Graif, T Hendler, and G Navon. "Susceptibility-matched envelope for the correction of EPI artifacts." *Magnetic resonance imaging*, **23**(9):947–951, November 2005.

[Neu90] Daniel Neuhauser. "Bound state eigenfunctions from wave packets: Timeenergy resolution." *The Journal of Chemical Physics*, **93**(4):2611, 1990.

[Neu94] Daniel Neuhauser. "Circumventing the Heisenberg principle: A rigorous demonstration of filter-diagonalization on a LiCN model." *The Journal of Chemical Physics*, **100**(7):5076, 1994.

[NFS05] Douglas C Noll, Jeffrey a Fessler, and Bradley P Sutton. "Conjugate phase MRI reconstruction with spatially variant sample density correction." *IEEE Transactions on Medical Imaging*, **24**(3):325–36, March 2005.

[NM11] Hiroya Nakao and Alexander S Mikhailov. "Turing patterns in network-organized activatorinhibitor systems." *Nature*, 2011.

[Nol91] D.C. Noll. *Reconstruction Techniques for Magnetic Reasonance Imaging*. PhD thesis, Stanford University, 1991.

[OBG05] Stanley Osher, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin. "An Iterative Regularization Method for Total Variation-Based Image Restoration." *Multiscale Modeling & Simulation*, **4**(2):460, 2005.

[PBG08] Jeffrey L Paulsen, Louis S Bouchard, Dominic Graziani, and Bernhard Blu. "Volume-selective magnetic resonance imaging using an adjustable , single-sided , portable sensor SCIENCES." *Proc. of the Nat. Acad. of Sciences*, **105**(52):20601–20604, 2008.

[Per04] J Perlo. "3D imaging with a single-sided sensor: an open tomograph." *Journal of Magnetic Resonance*, **166**(2):228–235, February 2004.

[Phi06] Bernard Philippe. "On correction equations and domain decomposition for computing invariant subspaces." *Work*, 2006.

[Pin12] Grigore Pintilie. "Identifying Components in 3D Density Maps of Protein Nanomachines by Multi-scale Segmentation Citation Accessed Citable Link Detailed Terms Identifying Components in 3D Density Maps of Protein Nanomachines by Multi-scale Segmentation." *Evaluation*, 2012.

[PL86] Tae Jun Park and J. C. Light. "Unitary quantum time evolution by iterative Lanczos reduction." *The Journal of Chemical Physics*, **85**(10):5870, 1986.

[Pra03] Pablo J. Prado. "Single sided imaging sensor." *Magnetic Resonance Imaging*, **21**(3-4):397–400, April 2003.

[PTR98] CH Papadimitriou, Hisao Tamaki, and P Raghavan. "Latent semantic indexing: A probabilistic analysis." *Proceedings of the*, pp. 1–18, 1998.

[QGC11] Xiaobo Qu, Di Guo, Xue Cao, Shuhui Cai, and Zhong Chen. "Reconstruction of self-sparse 2D NMR spectra from undersampled data in the indirect dimension." *Sensors (Basel, Switzerland)*, **11**(9):8888–909, January 2011.

[RF10] Sathish Ramani and Jeffrey A Fessler. "An Accelerated Iterative Reweighted Least Squares Algorithm for Compressed Sensing MRI." *IEEE ISBI*, pp. 257–260, 2010.

[Rom07] Justin Romberg. "Imaging via Compressive Sampling." *IEEE Signal Processing Magazine*, (March 2008):14–20, 2007.

[RS08] Holger Rauhut and Karin Schnass. "Compressed sensing and redundant dictionaries." *Information Theory, IEEE*, pp. 1–19, 2008.

[RST09] Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. "A randomized algorithm for PCA." *Siam J matrix anal. appl.*, **31**(3):1100–1124, 2009.

[RVY97] Jurgen Reichenbach, Ramesh Venkatesan, Dmitriy Yablonskiy, Michael R Thompson, and E Mark Haacke. "Theory and Application of Static Field Inhomogeneity Effects in Gradient- Echo Imaging." *Journal of Magnetic Resonance Imaging*, **7**:266–279, 1997.

[Saa96] Y. Saad. *Numerical Methods for Large Eigenvalue Problems*. 1996.

[SC08] Dmitrii V Shalashilin and Mark S Child. "Basis set sampling in the method of coupled coherent states: coherent state swarms, trains, and pancakes." *The Journal of chemical physics*, **128**(5):054102, March 2008.

[Sch99] H Schomberg. "Off-resonance correction of MR images." *IEEE transactions on medical imaging*, **18**(6):481–95, June 1999.

[SDH07] Alan S Stern, David L Donoho, and Jeffrey C Hoch. "NMR data processing using iterative thresholding and minimum l(1)-norm reconstruction." *Journal of magnetic resonance (San Diego, Calif. : 1997)*, **188**(2):295–300, October 2007.

[SDS01] Scott Shaobing, David L Donoho, and Michael A Saunders. "Atomic Decomposition by Basis Pursuit *." *Society*, **43**(1):129–159, 2001.

[SF11] Yoav Shrot and Lucio Frydman. "Compressed sensing and the reconstruction of ultrafast 2D NMR data: Principles and biomolecular applications." *Journal of magnetic resonance (San Diego, Calif. : 1997)*, **209**(2):352–8, April 2011.

[SMN03] Bradley P Sutton, Student Member, Douglas C Noll, Jeffrey A Fessler, and Senior Member. "Fast , Iterative Image Reconstruction for MRI in the Presence of Field Inhomogeneities." *IEEE transactions on medical imaging*, **22**(2):178–188, 2003.

[SV00] Gerard L G Sleijpen and Henk A V A N D E R Vorst. "A Jacobi-Davidson iteration method for linear eigenvalue problems." *SIAM Review*, **17**(2), 2000.

[TCS04] T-k Truong, D W Chakeres, and P Schmalbrock. "Effects of B 0 and B 1 Inhomogeneity in Ultra-High Field MRI." *Proc. Intl. Soc. Mag. Reson. Med.*, **11**:2170, 2004.

[TG06] Joel A Tropp and Anna C Gilbert. "Signal recovery from random measurements via orthogonal matching pursuit." *East*, 2006.

[TK84] H Tal-Ezer and R Kosloff. "Chebyshev Method Solution for the Time Dependent Schrddinger Equation as a Tool in Molecular Dynamics." 1984.

[TMS04] Daniel Topgaard, Rachel W Martin, Dimitris Sakellariou, Carlos A Meriles, and Alexander Pines. ""Shim pulses" for NMR spectroscopy and imaging." *Proceedings of the National Academy of Sciences of the United States of America*, **101**(51):17576–81, December 2004.

[VWB99] Amrendra Vijay, Robert E. Wyatt, and Gert D. Billing. "Time propagation and spectral filters in quantum dynamics: A Hermite polynomial perspective." *The Journal of Chemical Physics*, **111**(24):10794, 1999.

[WB04] Yinghua Wu and Victor S Batista. "Quantum tunneling dynamics in multidimensional systems: a matching-pursuit description." *The Journal of chemical physics*, **121**(4):1676–80, July 2004.

[WBP11] Bertram J Wilm, Christoph Barmet, Matteo Pavan, and Klaas P Pruessmann. "Higher order reconstruction for MRI in the presence of spatiotemporal field perturbations." *Magnetic resonance in medicine*, **1701**:1690–1701, April 2011.

[WN95] Michael R Wall and Daniel Neuhauser. "Extraction, through filter-diagonalization, of general quantum eigenvalues or classical normal mode frequencies from a small number of residues or a short-time segment of a signal. I. Theory and application to a quantum-dynamics model." **102**(20), 1995.

[YSW06] Qing Yang, Michael Smith, and Jianli Wang. "Magnetic susceptibility effects in high field MRI." *Biological Magnetic Resonance*, **26**:249–284, 2006.