

UCLA

UCLA Electronic Theses and Dissertations

Title

Nonuniform and Non-Cartesian Sampling in Multidimensional Magnetic Resonance Spectroscopic Imaging

Permalink

<https://escholarship.org/uc/item/7sk8h15x>

Author

Wilson, Neil

Publication Date

2015

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

**Nonuniform and Non-Cartesian Sampling in
Multidimensional Magnetic Resonance
Spectroscopic Imaging**

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

by

Neil Wilson

2015

© Copyright by
Neil Wilson
2015

ABSTRACT OF THE DISSERTATION

Nonuniform and Non-Cartesian Sampling in Multidimensional Magnetic Resonance Spectroscopic Imaging

by

Neil Wilson

Doctor of Philosophy in Biomedical Physics

University of California, Los Angeles, 2015

Professor M. Albert Thomas, Chair

Magnetic resonance spectroscopy (MRS) is used to obtain localized biochemical information noninvasively based on the principles of nuclear magnetic resonance. ^1H in vivo spectra consist of a large number of metabolites in a relatively small spectral range, making identification difficult. Multidimensional MRS incorporates a variable evolution period to enhance the information content and increase spectral dispersion. Recently, multidimensional MRS has been combined with echo planar gradient readout techniques to produce multidimensional magnetic resonance spectroscopic imaging (MRSI). Despite the fast imaging acquisitions, these scans are long for in vivo studies, so more efficiently sampling strategies were investigated.

The first strategy consisted of nonuniform undersampling (NUS) of the volume spanned by the phase-encoded spatial dimensions and the indirect spectral dimension in 5 dimensional (3 spatial + 2 spectral) MRSI. Nonlinear reconstruction was performed according to the theory of compressed sensing (CS) using the split Bregman framework. Formulations that promoted sparsity of the data and its spatial finite differences in 5D J-resolved brain studies were applied, and results were compared favorably to a time-equivalent single slice J-resolved scan.

In 5D correlated MRSI calf studies, reconstruction minimized the group sparsity of nearby points, which produced much better results than reconstruction that minimized the overall sparsity of the data.

The second strategy used concentrically circular k-space trajectories instead of the conventional rectilinear ones. Concentric circles have the advantages of reduced hardware demands, higher achievable spectral bandwidth, less sensitivity to motion, and faster k-space coverage. Single slice 4D (2 spatial + 2 spectral) correlated MRSI using concentrically circular trajectories was compared to a rectilinear counterpart and showed similar data quality. An improved single slice J-resolved MRSI sequence was presented. The new sequence used adiabatic refocusing pulses that are less sensitive to RF field inhomogeneity and result in reduced chemical shift displacement error compared to conventional pulses. Comparison was made to the nonadiabatic sequence with the same echo time as well as with its minimum echo time.

The dissertation of Neil Wilson is approved.

Luminita Vese

Michael McNitt-Gray

Peng Hu

Daniel Ennis

M. Albert Thomas, Committee Chair

University of California, Los Angeles

2015

For my girls

TABLE OF CONTENTS

1	Introduction	1
1.1	Outline	2
2	Background	4
2.1	The NMR phenomenon	4
2.1.1	Nuclear spin	4
2.1.2	The wavefunction	5
2.1.3	Time dependent Schrodinger equation	6
2.1.4	External spin Hamiltonian terms	7
2.1.5	Rotating frame of reference	8
2.1.6	Internal spin Hamiltonian terms	11
2.1.7	Spin ensembles and the density operator	12
2.1.8	Populations and coherences	15
2.1.9	Thermal equilibrium	17
2.1.10	Product operators	19
2.1.11	Observable operators	20
2.2	Sampling overview	24
2.2.1	Cartesian sampling	24
2.2.2	Non-Cartesian polar sampling	27
2.2.3	Sensitivity encoding	29
2.2.4	Signal-to-noise	32
3	Iterative Nonlinear Reconstruction using Split Bregman Algo-	

rithms	35
3.1 Introduction	35
3.2 Background into ℓ_1 solvers	35
3.3 Constrained optimization with Bregman iteration	37
3.4 Split Bregman algorithm and variants	39
3.4.1 Anisotropic TV -regularized	39
3.4.2 Isotropic TV -regularized	43
3.4.3 ℓ_1 -regularized	44
3.4.4 Multiple regularizers	45
3.4.5 Intermediate updating of Bregman variables	46
3.5 Extensions to multicoil data	47
3.5.1 Constructing matrix operators	49
3.5.2 SBTV SENSE	55
3.5.3 SB ℓ_1 SENSE	58
3.5.4 Self sparse SB ℓ_1 SENSE	60
3.6 Group Sparsity	61
3.6.1 Single channel	62
3.6.2 Multi channel with sensitivity maps	63
4 Accelerated Five-dimensional Echo Planar J-resolved Spectroscopic Imaging: Implementation and Pilot Validation in Human Brain	69
4.1 Abstract	69
4.2 Introduction	70
4.3 Methods	72

4.3.1	Phantom	72
4.3.2	in vivo	73
4.3.3	NUS masks	73
4.3.4	Reconstruction	75
4.3.5	Additional processing	76
4.4	Results	77
4.4.1	Phantom	77
4.4.2	in vivo	88
4.5	Discussion	94
4.6	Conclusion	98

5 3D Correlated Spectroscopic Imaging of Calf Muscle Using Group Sparse Reconstruction of Undersampled Single and Multichannel

Data	99
5.1	Introduction	99
5.2	Theory	101
5.2.1	Single channel GS reconstruction	101
5.2.2	Multi channel GS reconstruction	102
5.3	Methods	104
5.3.1	Sequence	104
5.3.2	Apodization and NUS masks	105
5.3.3	Reconstruction	106
5.3.4	Coil compression	109
5.4	Results	111
5.5	Discussion	119

5.6	Conclusion	121
6	Correlated Spectroscopic Imaging using Concentrically Circular k-space Trajectories	122
6.1	Introduction	122
6.2	Methods	126
6.2.1	Sequence design	126
6.2.2	Gridding	128
6.2.3	Experimental	129
6.3	Results	130
6.4	Discussion	132
6.5	Conclusion	135
6.6	Appendix	136
7	Improved J-resolved Spectroscopic Imaging with Adiabatic Re- focusing Pulses and Concentrically Circular Echo Planar k-space Trajectories	138
7.1	Introduction	138
7.2	Methods	141
7.3	Results	144
7.4	Discussion	146
7.5	Conclusion	149
7.6	Appendix	149
7.6.1	Siemens Pulsetool parameters: <i>Refgrad</i> and the BWTP	149
8	Quantitative Localized Correlated Spectroscopy using Prior Knowl-	

edge Fitting Including Macromolecules: Comparison with J-resolved Spectroscopy and Results in HIV-infected Children	151
8.1 Introduction	151
8.2 Methods	154
8.2.1 Phantom Scans	154
8.2.2 in vivo Scans	155
8.2.3 Processing and Fitting	156
8.3 Results	159
8.3.1 Phantom Scans	159
8.3.2 HIV patients in vivo	161
8.4 Discussion	163
8.5 Conclusion	167
9 Conclusion	168
9.1 Summary of significance	168
9.2 Future Directions and Recommendations	169
9.2.1 NUS and reconstruction	169
9.2.2 Concentric circular sampling	170
References	173

LIST OF FIGURES

2.1	Figure representing the magnetic moments of nuclear spins in an unpolarized sample. Reproduced with permission from [75].	15
2.2	Illustration of the magnetization at thermal equilibrium. Spins are oriented almost isotropically in three dimensions, but there is a <i>small</i> net polarization in the vertical direction indicated by the thick green arrow. All spins precess around the z -axis as indicated by the red arrow. Reproduced with permission from [75].	18
2.3	(a) Periodic sampling at the Nyquist rate. (b) Periodic sampling at half the Nyquist rate in the vertical direction results in coherent aliasing. (c-d) Point spread functions (PSF) of (a-b). PSFs were computed from the Fourier transform of the sampling pattern. . .	25
2.4	Incoherent aliasing artifacts with random, uniform density NUS (a) and variable density NUS (b). (c-d) PSFs of (a-b).	26
2.5	The top row shows fully sampled (left), radially undersampled (center), and circularly undersampled (right) polar data. The bottom row is the respective PSFs.	29
2.6	2x undersampling in the vertical direction results in two pixels on the left summing to form the aliased pixel on the right. Here, there is no way to separate out the pixels from the image on the right. .	30
2.7	Sensitivity weighted images from a coil above the image (top) and one below the image (bottom). In the top image, the top pixel contributes more to the aliased pixel, while in the bottom image, it is the bottom pixel that contributes more.	31
2.8	Geometry factor map for the undersampled coil images shown in Fig ???. The color scale is from 1 to 2.	34

3.1	Split Bregman algorithm for anisotropic TV regularization	42
3.2	Split Bregman algorithm for isotropic TV regularization	44
3.3	Split Bregman algorithm for ℓ_1 regularization	45
3.4	Split Bregman algorithm for ℓ_1 regularization with intermediate updating of the Bregman variable. Note the additional computation between the u subproblem and the w subproblem. Though it is applied twice, Wu^{k+1} only needs to be computed once.	47
3.5	Effects of coil combination on NAA metabolite map (a) and COSY spectrum of a central voxel (b) . In each case, the coil combination using complex sensitivity maps is shown on the left while sum-of-squares combination is shown on the right.	48
3.6	Split Bregman algorithm for isotropic TV regularization from multicoil acquisition	58
3.7	Split Bregman algorithm for ℓ_1 -norm regularization from multicoil acquisition	59
3.8	Simple split Bregman algorithm for ℓ_1 regularized self sparse multicoil acquisition with intermediate variable updating	61
3.9	SBGS SENSE I algorithm	66
3.10	SBGS SENSE II algorithm	68
4.1	8x nonuniform sampling mask over the (k_y, k_z, t_1) volume used for phantom and in vivo scans. White squares indicate sampled points.	74
4.2	J-resolved spectrum from the same central voxel in the phantom for different undersampling rates. (a) fully sampled, (b) 4x NUS with minimum energy reconstruction (missing data points were filled in with zeros), (c)-(f) ℓ_1 reconstructions for 4x, 8x, 12x, and 16x NUS, respectively.	77

4.3	Glx metabolite map for a typical phantom scan with various retrospective undersamplings. Slices are stacked vertically, and the readout direction is horizontal. Reconstruction was performed according to Eq. [??], minimizing the TV -norm. ‡Reconstruction did not include spectral masking.	78
4.4	Localized three plane image of phantom (left) and NAA metabolite map for a typical phantom scan with various retrospective undersamplings (right). The phase encoding direction is shown vertically, the readout is horizontal, and the slices are stacked. Reconstruction was performed according to Eq. (??), minimizing the ℓ_1 -norm except where noted. †Minimum energy reconstruction	80
4.5	The effect of spectral masking before reconstruction. J-resolved spectrum from the same central voxel in the phantom for different undersamplings with TV reconstruction. (a)-(c) are for 8x, 12x, and 16x NUS, respectively and include spectral masking. (d)-(f) use the same NUS and reconstruction except without any spectral masking; the contour level has been scaled down by a factor of two to partially compensate for the decreased signal intensity. In all spectra, the F_2 axis is extended to show the residual water peak.	81
4.6	Peak integrals reconstructed from different undersamplings for five metabolites. Integrals were normalized by the mean of Cr in the fully sampled data. Reconstruction was performed according to Eq. [??] (top) and Eq. [??] (bottom).	86
4.7	Metabolites maps (left) for a 29 year old healthy adult scanned with 8x undersampling and reconstructed using Eq. [??]. Each metabolite was normalized to its maximum so that they appear on the same gray scale. Spectra from the highlighted voxels are shown at right.	89

4.8	Regional peak integrals of major metabolites in ten healthy adults with minimum TV reconstruction. The green bars are from 8x NUS 5D data reconstructed using Eq. [??] for voxels in the frontal lobe (top) and in the occipito-parietal lobes (bottom). The orange bar is from a time-equivalent, fully-sampled 4D data set for comparison.	90
4.9	Regional peak integrals of major metabolites in ten healthy adults with minimum ℓ_1 reconstruction. The green bars are from 8x NUS 5D data reconstructed using Eq. (??) for voxels in the frontal lobe (top) and in the occipito-parietal lobes (bottom). The orange bar is from a time-equivalent, fully-sampled 4D data set for comparison.	91
5.1	8x (a), 12x (b), and 16x (c) nonuniform sampling masks over the (k_y, k_z, t_1) volume used to acquire accelerated EP-COSI data. White squares indicate sampled points. Each mask follows the sampling density given in Eq. (??).	107
5.2	Illustration of the difference between two dimensional groups that are non-overlapping and those with 50% overlap. The overlapping groups were used to combine coefficients in the spectral plane (F_2, F_1) .	109
5.3	Magnitude of sensitivity maps from 15-channel data taken directly from a water reference scan (a), after using an x -dependent geometric transformation matrix $A(x)$ (Eq. (??)) (b), and after coil compression using a single transformation matrix A (Eq. (??)). .	112

5.4	Metabolite maps from an 8x undersampled scan of a diabetic calf. Shown are the maps from the creatine and fat peak at 3.0 ppm (Cr3.0), the creatine resonance at 3.9 ppm (Cr3.9), and the lower (Low) and upper (Upp) IMCL and EMCL lipid cross peaks. Within each metabolite group, the minimum energy reconstruction is shown with the CS and GS reconstructions, respectively. Minimum energy reconstructions are scaled by the square root of the undersampling factor. Localization is shown on the left.	113
5.5	COSY spectra from the highlighted voxels in a central slice (a) of the same scan as Fig ???. Yellow spectra (b,c) are from the tibial marrow, red spectra (d,e,f) are from the soleus muscle, and blue spectra (g,h,i) are from the tibialis anterior muscle. The first column shows raw spectra with the minimum energy reconstruction (d,g), the second column shows CS-reconstructed spectra (b,e,h), and the third column shows GS-reconstructed spectra (c,f,i). Contour levels are the same for each reconstruction for a given voxel with the minimum energy reconstructions being scaled by the square root of the undersampling factor.	114
5.6	Metabolite maps from a 16x undersampled scan. Shown are the maps from the creatine and fat peak at 3.0 ppm (Cr3.0), the creatine resonance at 3.9 ppm (Cr3.9), and the lower (Low) and upper (Upp) IMCL and EMCL lipid cross peaks. Within each metabolite group, the minimum energy reconstruction is shown with the ℓ_1 and GS reconstructions. Minimum energy reconstructions are scaled by the square root of the undersampling factor. Localization is shown on the left.	116

5.7	<p>COSY spectra from the highlighted voxels shown in Fig ?? . Yellow spectra (a,b,c) are from the tibial marrow, red spectra (d,e,f) are from the soleus muscle, and blue spectra (g,h,i) are from the tibialis anterior muscle. The first column shows raw spectra with the minimum energy reconstruction (a,d,g), the second column shows CS-reconstructed spectra (b,e,h), and the third column shows GS-reconstructed spectra (c,f,i).</p>	117
5.8	<p>Axial localization with 15-channel knee coil (a). Contour plot of the distribution of the peaks at 3.9 ppm for a single slice MEEP-COSI scan (b) and the central three slices from 12x undersampled 3D EP-COSI (c-e) in the same healthy volunteer. The undersampled 3D EP-COSI scan took 14 minutes, and the MEEP-COSI scan took 13 minutes.</p>	118
6.1	<p>COSI-CONCEPT schematic showing the L-COSY localization with indirect spectral increment t_1 and sinusoidal readout gradients. The incremented gradient in the y direction plays the role of a typical phase encoding gradient so that readout does not begin at the origin of k-space. Spoilers (not shown) are placed symmetrically around the second and third pulses to remove signal due to unwanted coherences.</p>	125
6.2	<p>(a) Example of a concentric k-space trajectory, starting in the origin and traversing from $a \rightarrow b$, and then counter-clockwise around the circle from $b \rightarrow b'$. (b) The individual $k_x(t)$ and $k_y(t)$ trajectories as shown in (a), showing the individual steps from $a \rightarrow b$ as well as $b \rightarrow b'$.</p>	126

6.3	The first row shows highlighted voxels from the tibial bone marrow (labeled in red in the anatomical T2-weighted MRI insert) (a) , tibialis anterior (labeled in yellow) (b) , and soleus (labeled in blue) (c) , from a COSI-CONCEPT scan of a diabetic calf. The second row shows the corresponding voxels from the EP-COSI scan, (d) , (e) , and (f) . Note that spectra are not all scaled the same.	131
6.4	Spatial distribution of the creatine 3.9 ppm peak from a COSI-CONCEPT scan overlayed on top off the anatomic T2-weighted MRI.	132
7.1	JRESI-CONCEPT pulse sequence schematic with semi-LASER localization. The 90° pulse is non-adiabatic and slice selective. The 180° pulses are all adiabatic and slice selective. The first two select the <i>y</i> -slice, and the last two select the <i>x</i> -slice. Sinusoidal gradients are applied in <i>x</i> and <i>y</i> to traverse k-space circularly. The indirect spectral increment t_1 was placed before the final pulse only for maximum echo sampling. Crusher gradients surrounding each refocusing pulse and a WET water suppression module preceding the localization are not shown, and the sequence timing of the readout gradients is not drawn to scale.	142
7.2	Metabolites maps for NAA 2.0 ppm singlet (left), creatine 3.0 ppm (center), and creatine 3.9 ppm (right) acquired with sLASER-JRESI-CONCEPT (top) and PRESS-JRESI-CONCEPT (bottom) acquired at the same echo time (TE = 37 ms) in a healthy volunteer. Maps are scaled to be of similar intensities.	144
7.3	J-resolved spectra from the highlighted voxel in the <i>parietal</i> lobe with sLASER localization (left) and PRESS localization (right) in the same healthy volunteer as Fig ??	145

7.4	J-resolved spectra from the highlighted voxel in the <i>occipital</i> lobe with sLASER localization (left) and PRESS localization (right) from the same scan as Fig ??.	145
7.5	Metabolite maps from a coronal slice of a second healthy adult showing creatine 3.0 ppm (top) and Glx (bottom) for minimum TE sLASER localization (left), same TE PRESS localization (center), and minimum TE PRESS localization (right). Maps are scaled the same between scans but differently between metabolites.	146
7.6	J-resolved spectra from the highlighted voxel in a coronal slice in the healthy volunteer as Fig ?? acquired with sLASER localization at minimum TE (left), PRESS localization at the same TE (center), and PRESS localization at its minimum TE (right).	147
8.1	L-COSY basis set used for phantom scans (a) as simulated by GAMMA and (b) after a frequency-dependent phase factor was applied. The two Cr singlets were simulated and fit independently as described in the text.	159
8.2	An example in vivo spectrum (top) from an HIV patient, its fit (middle), and its residual (bottom), plotted on a logarithmic scale.	163
8.3	Normalized variance matrices for (a) the L-COSY basis set and (b) the JPRESS basis set. The variance matrix is the inverse of the Fisher matrix, and the CRLBs are the diagonal elements.	164
9.1	k-space trajectories (left) and x gradient waveforms. Blue rings collect inner k-space at high SBW, while red and green rings collect outer k-space at half the SBW and are temporally interleaved.	172

LIST OF TABLES

4.1	Spectral ranges used to measure peak integrals. Cho is total choline; Cr30 is the creatine singlet at 3.0 ppm; Cr39 is the creatine singlet at 3.9 ppm; Glx is glutamate and glutamine combined; mI is myo-inositol; NAA is N-acetylaspartate. Wat and Tail are used as a measure of baseline effects due to residual water contamination. Fat is a marker of lipid signal bleed. Sig is a marker of overall signal level. Noise is taken from a region that is devoid of any signal, desired or nuisance.	83
4.2	Mean metabolite peak RMSEs	84
4.3	Phantom peak integrals	87
4.4	in vivo peak integrals occipito-parietal	92
4.5	in vivo peak integrals frontal	93
8.1	Comparison of the ProFit results in L-COSY and JPRESS in phantom. 37 L-COSY spectra and 38 JPRESS spectra were acquired. Ratios are shown with respect to the Cr singlet at 3.03ppm. % Fit is the percent of spectra that were fit with CRLB<20%. Only those voxels were included in the computation of mean ratios, mean CRLBs, and coefficients of variation (CV). CVs are expressed as percentages. tNAA = NAA+NAAG, Glx = Glu+Gln, tCho = PCh+Cho.	160

8.2 ProFit L-COSY results of frontal white/gray brain spectra in healthy controls and patients. 14 healthy controls and 15 HIV patients were scanned. Ratios are shown with respect to the Cr singlet at 3.03ppm. SD is the standard deviation of the ratios. N is the number of voxels that were fit with CRLB<20%. Only those voxels were included in the computation of mean ratios and SDs. p-values indicating statistically significant differences at the 5% level are highlighted in bold. tNAA = NAA+NAAG, Glx = Glu+Gln, tCho = GPC+PCh+Cho. 162

ACKNOWLEDGMENTS

To my advisor, Dr Thomas, I want to thank you for encouraging me to take on challenging projects and to learn them top to bottom. Thank you for being patient and giving me the freedom to work through all the inevitable difficulties and setbacks to grow as a scientist. Your excitement for new ideas was contagious. And thank you for trusting me to prepare my own slides for your class.

I wish to thank my committee for challenging me to create the best version of my work that I could. Your input and varying perspectives were highly appreciated and valued.

Jon and Brian, I was so fortunate to have you two as peers and mentors in the group. Through our talks and whiteboard sessions, I learned so much about NMR (Jon) and signal processing (Brian). More importantly, I learned to be more rigorous with my own work and address weaknesses, which ultimately lead to better, more forward-looking research. And Jon, thank you for introducing me to L^AT_EX.

To all the members of the Thomas group, your camaraderie and encouragement helped keep me going through the ups and downs of Phd research. Scott, thank you for bringing me on to this group and showing me the ropes. You encouraged me to learn all my own coding, so I would not have to rely on anyone else. It was truly sound advice. Zohaib, thank you for helping me organize and finish up projects when I had too much going on to keep them straight.

To my classmates and friends in Biomedical Physics, thank you for making our department feel like a family even after we moved out of our little corner in the basement, and our research paths and office locations diverged.

To my parents, thank you for being supportive of all my academic endeavors. You have inspired in me an intellectual curiosity and given me the confidence to aim high with my goals.

Most of all, I thank you, Elaine, for partnering me through this journey every step of the way...moving cross country, coming with me to present my first poster in Hawaii, consoling me when my work was rejected, celebrating with me when it was accepted, being the primary breadwinner, and of course, starting a family with a poor graduate student. I could not have done any of this without you and forgive you for never getting scanned.

VITA

2007 Bachelor of Science in Physics and Mathematics,
Rutgers University.

PUBLICATIONS

1. **NE Wilson**, Zohaib Iqbal, Brian L. Burns, Margaret Keller, and M. Albert Thomas. Accelerated five-dimensional echo planar j-resolved spectroscopic imaging: Implementation and pilot validation in human brain. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2015.
2. Brian L. Burns, **NE Wilson**, and M. Albert Thomas. Group sparse reconstruction of multi-dimensional spectroscopic imaging in human brain in vivo. *Algorithms*, 7(3):276–294, 2014.
3. Brian Burns, **NE Wilson**, Jon K. Furuyama, and M. Albert Thomas. Non-uniformly under-sampled multi-dimensional spectroscopic imaging in vivo: maximum entropy versus compressed sensing reconstruction. *NMR in Biomedicine*, 27(2):191–201, 2014.
4. Jon K. Furuyama, **NE Wilson**, Brian L. Burns, Rajakumar Nagarajan, Daniel J. Margolis, and M. Albert Thomas. Application of compressed sensing to multidimensional spectroscopic imaging in human prostate. *Magnetic Resonance in Medicine*, 67(6):1499–1505, 2012.
5. Jon K. Furuyama, **NE Wilson**, and M. Albert Thomas. Spectroscopic

imaging using concentrically circular echo-planar trajectories in vivo. *Magnetic Resonance in Medicine*, 67(6):1515–1522, 2012.

6. Preethi Srikanthan, Aparna Singhal, Cathy C. Lee, Rajakumar Nagarajan, **NE Wilson**, Christian K. Roberts, Theodore J. Hahn, and M. Albert Thomas. Characterization of intra-myocellular lipids using 2d localized correlated spectroscopy and abdominal fat using mri in type 2 diabetes. *Magnetic Resonance Insights*, 5:29–36, 10 2012.
7. Whitney Pope, Robert Prins, M. Albert Thomas, Rajakumar Nagarajan, Katharine Yen, Mark Bittinger, Noriko Salamon, Arthur Chou, William Yong, Horacio Soto, **NE Wilson**, Edward Driggers, Hyun Jang, Shinsan Su, David Schenkein, Albert Lai, Timothy Cloughesy, Harley Kornblum, Hong Wu, Valeria Fantin, and Linda Liau. Non-invasive detection of 2-hydroxyglutarate and other metabolites in idh1 mutant glioma patients using magnetic resonance spectroscopy. *Journal of Neuro-Oncology*, 107:197–205, 2012.
8. Jon K. Furuyama, Brian L. Burns, **NE Wilson**, and M. Albert Thomas. Multi-echo-based echo-planar spectroscopic imaging using a 3t mri scanner. *Materials*, 4(10):1818–1834, 2011.
9. Manoj Sarma, Amir Huda, Rajakumar Nagarajan, Charles Hinkin, **NE Wilson**, Rakesh Gupta, Enrique Frias-Martinez, James Sayre, Barry Guze, Steven-Huy Han, and M. Thomas. Multi-dimensional mr spectroscopy: towards a better understanding of hepatic encephalopathy. *Metabolic Brain Disease*, 26:173–184, 2011.

CHAPTER 1

Introduction

Magnetic resonance imaging (MRI) and magnetic resonance spectroscopy (MRS) are, loosely speaking, the *in vivo* applications of the broader field of nuclear magnetic resonance (NMR) and as such, are fundamentally based on the interactions of atomic nuclei with magnetic fields. Both are, at the heart of it, spectroscopic techniques. In MRS, the range of frequencies is provided by the differences in chemical environment of different nuclei, while in MRI, the range of frequencies is provided by the application of a spatially dependent external magnetic field. Because of that, spectral or spatial resolution is not determined by the energy of the electromagnetic radiation. For the magnetic fields in use, energy is in the radio frequency (RF) range, which is non-ionizing and generally considered safe.

Pulse sequences control the application of the RF energy and the spatially dependent magnetic fields, and they can be sensitized to a wide range of physiologic parameters. Multidimensional MRS pulse sequences utilize a variable evolution time to separate out nuclear interactions spectrally and can be combined with spatially dependent gradient fields in multidimensional magnetic resonance spectroscopic imaging (MRSI). However, fully sampling the data over all the spatial and spectral dimensions is often prohibitively long, even with advanced echo planar techniques. This thesis seeks to address the sampling of multidimensional MRSI in two distinct ways.

The first uses the theory of compressed sensing (CS) and the inherent compressibility of the MRSI data to accelerate acquisition by acquiring many fewer

measurements. This is done by nonuniformly undersampling (NUS) the incrementally acquired spatial and spectral dimensions and reconstructing the data using sparsity-promoting, iterative nonlinear algorithms.

In the second way, the means in which the spatial frequency data is collected shifts from a typical Cartesian grid to non-Cartesian concentric circles. Circular sampling covers the spatial frequencies twice as fast as rectilinear sampling. Other advantages of circular sampling will be discussed with specific application to MRSI.

These two methods are entirely distinct in their philosophy and application and can be considered complementary.

1.1 Outline

Chapter 2 gives some background on the theory of NMR and sampling in MRI. A quantum mechanical treatment is given as necessary to explain the two-dimensional J-resolved and correlated spectroscopy sequences even though most of the imaging aspects can be explained using classical theory applied to the bulk magnetization.

Iterative algorithms must be applied to reconstruct the NUS data in compressed sensing. With the large dimensionality of the multidimensional MRSI data from multichannel receivers, fast and efficient solvers must be used. Chapter 3 describes one such algorithm known as split Bregman and its applicability to the various problem constructs.

Chapter 4 highlights the application of two standard split Bregman techniques and applies them to five-dimensional (three spatial and two spectral dimensions) J-resolved MRSI data. Comparisons are made between an algorithm that promotes the self-sparsity of the reconstructed data and one that promotes the sparsity of the finite differences of the data. Phantom scans with various retrospective NUS are compared, and in vivo scans of the human brain are presented that acquired

only 12% of the fully sampled data. This is the first application of any such five-dimensional MRSI scan.

Chapter 5 uses a version of the split Bregman algorithm to minimize the sparsity of grouped data points in a five-dimensional correlated MRSI pulse sequence. The algorithm was able to reconstruct scans sampled with only 6% of the data in calf muscle. Also, a new coil transformation was applied to reduce the 15 physical coils to five virtual ones without loss in signal, speeding up the algorithm and reducing its computational burden.

In Chapter 6, a correlated spectroscopic imaging with concentric circular echo planar trajectories (COSI-CONCEPT) sequence is presented and compared with a rectilinear echo planar correlated spectroscopic imaging (EP-COSI) sequence in calf muscle. Some of the theoretical benefits of circular sampling are discussed. This is the first application of circular sampling with multidimensional MRSI.

Chapter 7 takes a J-resolved concentric circular (JRESI-CONCEPT) sequence and replaces the traditional refocusing pulses with pairs of adiabatic ones that have higher bandwidth and are less susceptible to artifacts due to RF field inhomogeneity. The sequence is compared both to a non-adiabatic JRESI-CONCEPT with the same echo time, as well as to one with its minimum echo time. Adiabatic refocusing pulses can be used to improve all the other pulse sequences described in this dissertation.

Multidimensional MRSI provides more spectral information than single dimension MRSI. However, actual metabolite quantification can only be improved if this information is used maximally through prior knowledge-based fitting. Chapter 8 illustrates prior knowledge fitting of single voxel correlated and J-resolved spectroscopy to look for metabolic differences in pediatric HIV patients.

Lastly, Chapter 9 summarizes this dissertation and provides insight into future work and applications to further develop these techniques.

CHAPTER 2

Background

2.1 The NMR phenomenon

NMR is a result of certain nuclei interacting with magnetic fields. While this interaction is most definitely quantum mechanical on the individual atom scale, the large number of identical atoms (referred to as an *ensemble*) can be treated classically for nearly all MRI applications and many MRS applications. Nevertheless, for a thorough understanding of 2D MRS experiments, a quantum mechanical presentation must be considered and is presented here. More detailed explanations can be found elsewhere [1, 157, 94, 88, 31, 44].

2.1.1 Nuclear spin

Spin is a quantum mechanical form of angular momentum, similar to the classical orbital angular momentum of rotating bodies. In NMR, it is the net nuclear spin only that we are concerned with. Each nucleon (proton or neutron) has spin- $1/2$. ^1H , which is the nuclei of interest in this dissertation, has a single proton nucleus and therefore is also spin- $1/2$. In nuclei with multiple nucleon, the Pauli exclusion principle must be considered (i.e. two protons will pair antiparallel to give a net nuclear spin of 0). However, a single proton and a single neutron are different particles and are not bound by the Pauli principle, and they align parallel, producing a nonzero net nuclear spin. Thus, nuclei with an *odd* number of nucleons are spin- $1/2$, nuclei with *both* an *even* number of protons and an *even*

number of neutrons are spin-0, and nuclei with *both* an *odd* number of protons and neutrons are integer spin > 0 .

The magnetic moment vector $\boldsymbol{\mu}$ of a nucleus is proportional to its nuclear spin \mathbf{I}

$$\boldsymbol{\mu} = \gamma \mathbf{I} \quad (2.1)$$

with the constant of proportionality being equal to the gyromagnetic ratio $2.67 \times 10^8 \text{ s}^{-1} \cdot \text{T}^{-1}$ (often written as $\gamma = 42.6 \text{ MHz/T}$ in frequency units).

A nucleus with spin I has associated substates described by the spin quantum number $m_I \in \{-I, -I + 1, \dots, I\}$. In the absence of a magnetic field, these $2I + 1$ substates have degenerate energy. In the presence of a static magnetic field \mathbf{B} , this degeneracy is broken through the Zeeman effect

$$E = -\boldsymbol{\mu} \cdot \mathbf{B} \quad (2.2)$$

As the spin is quantized, the magnetic moment and also the energy levels are quantized. By convention, the external field is defined to be pointing in the z direction ($\mathbf{B} = B_0 \hat{z}$), so the energy levels are given by

$$E = -\gamma B_0 I_z \quad (2.3)$$

2.1.2 The wavefunction

Quantum mechanically, a measurement is calculated by applying an operator to the *wavefunction* of the particle or system. The wavefunction Ψ is a mathematical construct that captures all the physically relevant information. For instance, to measure the z component of spin of a nucleus, we apply I_z to Ψ , with the result being one of the values of m_I .

$$I_z |\Psi\rangle = m_I \hbar |\Psi'\rangle \quad (2.4)$$

For spin-1/2, there are two value of m_I and therefore, there are two eigenvalue equations for the operator I_z

$$\begin{aligned} I_z |\alpha\rangle &= +\frac{1}{2}\hbar |\alpha\rangle \\ I_z |\beta\rangle &= -\frac{1}{2}\hbar |\beta\rangle \end{aligned} \tag{2.5}$$

where \hbar is Planck's constant divided by 2π . These eigenstates $|\alpha\rangle$ and $|\beta\rangle$ of I_z have definite z components of angular momentum and are commonly referred to as *spin up* and *spin down*, respectively. Here, it should be noted that the operators I_x and I_y have distinct eigenstates and that application of either operator to the spin up or spin down states is fundamentally unpredictable whether the result will be $+1/2 \hbar$ or $-1/2 \hbar$.

The eigenstates form an orthonormal basis for the general wavefunction. In the same way that any 2D plane vector can be written as a sum of its x and y components, any general wavefunction can be written as a sum of its eigenstate components, which for spin-1/2 can be

$$|\Psi\rangle = c_\alpha |\alpha\rangle + c_\beta |\beta\rangle \tag{2.6}$$

where $|c_\alpha|^2$ and $|c_\beta|^2$ give the probabilities of measuring the z angular momentum to be $+1/2 \hbar$ and $-1/2 \hbar$, respectively.

2.1.3 Time dependent Schrodinger equation

The evolution of the wavefunction is governed by the time dependent Schrodinger equation (TDSE)

$$\frac{d}{dt} |\Psi\rangle = -\frac{i}{\hbar} H |\Psi\rangle \tag{2.7}$$

where H is the Hamiltonian (energy operator). From this point forward, we will consider natural units in which $\hbar = 1$ for convenience except where explicitly noted. For a time independent Hamiltonian, the solution to Eq. (2.7) is given by

$$|\Psi(t)\rangle = \exp(-iHt) |\Psi(0)\rangle \tag{2.8}$$

2.1.4 External spin Hamiltonian terms

Equation (2.7) expresses the full quantum state of the wavefunction and its evolution due to a Hamiltonian that comprises all the interactions in the entire system. In principle, though, we can restrict this equation to one of the *nuclear spin state* evolution under the influence of the *nuclear spin Hamiltonian* only, since the time-scales between nuclear and electronic motions are so different, and nuclear energies are generally too small to appreciably affect electron or molecular motion.

There are three prominent external magnetic fields in MRI/MRSI:

1. The large static and homogeneous field, B_0
2. Smaller, position dependent linear gradient fields $B_{x,y,z}$ (e.g $B_x = G_x x$)
3. An oscillating field in the RF range, B_1

2.1.4.1 Main magnetic field

As previously mentioned, the static field is conventionally given by $B_0 = B_0 \hat{z}$, and its Hamiltonian is given by

$$H_0 = -\gamma B_0 I_z \quad (2.9)$$

Since H is proportional to I_z , they commute with each other and have the same eigenstates, $|\alpha\rangle$ and $|\beta\rangle$. The general states then evolve according to $\exp(i\gamma B_0 t I_z)$, which is a rotation around the z axis at an angular frequency of

$$\omega_0 = -\gamma B_0 \quad (2.10)$$

known as the *Larmor frequency*.

2.1.4.2 Gradient fields

The gradient fields are spatially dependent with the following Hamiltonian representations

$$\begin{aligned}H_x &= -\gamma G_x x I_z \\H_y &= -\gamma G_y y I_z \\H_z &= -\gamma G_z z I_z\end{aligned}\tag{2.11}$$

Though the spatial dependence can be either x , y , or z , the field is always in the z direction, so gradient fields act similarly to the main magnetic field and cause precession around the z -axis at a frequency of

$$\omega(x, y, z) = \omega_0 - (G_x x + G_y y + G_z z)\tag{2.12}$$

2.1.4.3 RF field

The RF field is different from the previous two external field types in that it acts in the transverse plane and has a frequency component (i.e. is time dependent). The RF Hamiltonian is given by

$$H_1(t) = \omega_1 [\cos(\omega_{rf}t + \phi)I_x + \sin(\omega_{rf}t + \phi)I_y]\tag{2.13}$$

where ω_{rf} is the carrier frequency, ϕ is the phase of the pulse (0 for x pulse, $\pi/2$ for y pulse), and ω_1 is the nutation frequency that is proportional to the B_1 amplitude.

2.1.5 Rotating frame of reference

In order to calculate the effect of the time dependent RF Hamiltonian, the system is shifted to a rotating frame of reference in which the Hamiltonian appears stationary. At that point, the TDSE in Eq. (2.7) can be applied as before.

Consider a reference frame rotating around the z -axis at frequency ω_{ref} . A

general Hamiltonian in this frame is given by

$$H^r = \exp(i\omega_{ref}tI_z)H \exp(-i\omega_{ref}tI_z) - \omega_{ref}I_z \quad (2.14)$$

The TDSE is the same in the rotating frame provided that the Hamiltonian is given by Eq. (2.14) and the wavefunction is also transformed to the rotating frame, $|\Psi^r\rangle = \exp(i\omega_{ref}tI_z)|\Psi\rangle$.

Before continuing, we pause to consider the first term in Eq. (2.14) that consists of an operator sandwiched by two rotation operators of opposite sign. This type of expression appears routinely in NMR and is often simplified with the following relationship given without proof

$$\exp(-i\theta A)B \exp(i\theta A) = B \cos(\theta) + C \sin(\theta) \quad (2.15)$$

where $C = -i[A, B]$.

2.1.5.1 Free precession under B_0

The rotating frame Hamiltonian of the main magnetic field is given by plugging Eq. (2.9) and (2.10) into Eq. (2.14)

$$H_0^r = \omega_0 \exp(i\omega_{ref}tI_z)I_z \exp(-i\omega_{ref}tI_z) - \omega_{ref}I_z \quad (2.16)$$

I_z and the rotation operator $\exp(i\omega_{ref}tI_z)$ commute with each other, so the order of the product can be flipped, leading to the cancellation of the rotator terms and a particularly simple Hamiltonian

$$H_0^r = (\omega_0 - \omega_{ref})I_z = \Omega I_z \quad (2.17)$$

where Ω is known as the *offset frequency*. If the frame rotates at the resonant frequency, the offset $\Omega = 0$, and the Hamiltonian H_0^r is also 0, indicating no evolution, as expected. Otherwise, the states undergo apparent precession at the offset frequency.

2.1.5.2 RF fields in the rotating frame

First, using the sandwich formula in Eq. (2.15), rewrite Eq. (2.13) as

$$H_1(t) = \omega_1 \exp[-i(\omega_{rf}t + \phi)I_z]I_x \exp[i(\omega_{rf}t + \phi)I_z] \quad (2.18)$$

Then, plugging it into Eq. (2.14) yields

$$H_1^r = \omega_1 \exp(i\omega_{ref}t) \exp[-i(\omega_{rf}t + \phi)I_z]I_x \exp[i(\omega_{rf}t + \phi)I_z] \exp(i\omega_{ref}t) - \omega_{ref}I_z \quad (2.19)$$

If the frame of reference rotates at the same frequency as the RF field oscillates ($\omega_{ref} = \omega_{rf}$), then the Hamiltonian simplifies to

$$H_1^r = \omega_1 \exp(-i\phi I_z)I_x \exp(i\phi I_z) - \omega_{ref}I_z \quad (2.20)$$

and again to

$$H_1^r = \omega_1(I_x \cos \phi + I_y \sin \phi) - \omega_{ref}I_z \quad (2.21)$$

after applying the sandwich formula once more. This rotating frame Hamiltonian has now been rendered time independent and can be applied in the TDSE.

In actuality, the main magnetic field is present even during an RF pulse, and the spins do precess as they nutate. Therefore, the Hamiltonian should also include a term for H_0 , which simply modifies the Hamiltonian with the offset frequency

$$H_{0,1}^r = \omega_1(I_x \cos \phi + I_y \sin \phi) + \Omega I_z \quad (2.22)$$

When the offset is 0, the axis of rotation is in the transverse plane, and the *flip angle* is given by $\omega_1 t$. However, when the offset is not 0, the axis of rotation is at an angle

$$\theta = \arctan\left(\frac{\omega_1}{\Omega}\right) \quad (2.23)$$

from the z -axis, and the flip angle is given by

$$\text{flip angle} = \sqrt{\omega_1^2 + \Omega^2} t \quad (2.24)$$

Interestingly, the flip angle actually gets *larger* with greater offset. However, the rotation is less effective as it is around an oblique axis.

2.1.6 Internal spin Hamiltonian terms

The previous section talked about the applied fields and their subsequent Hamiltonians. However, the nuclei also experience significant magnetic fields that originate from the molecules themselves. In NMR/MRS, these internal interactions are much weaker than the interaction with the main magnetic field and represent perturbations to those results.

2.1.6.1 Chemical shift

In addition to interacting with atomic nuclei, the main magnetic field interacts with the electron clouds, causing currents surrounding the nuclei. These currents then induce their own magnetic fields which are usually opposite B_0 . The precessional frequency of the nuclei due to the sum of these fields with the main field is written as

$$\omega = \omega_0(1 + \delta) \quad (2.25)$$

since the effective shielding scales linear with the main field strength. The *chemical shift* δ is much smaller than 1 and is usually expressed as parts per million (ppm). By convention, chemical shift is defined with reference to 2,2-dimethyl-2-silapentane-5-sulfonate (DSS) and calculated for each resonance as

$$\delta = \frac{\omega - \omega_{DSS}}{\omega_{DSS}} \cdot 10^6 \quad (2.26)$$

As with the gradient fields mentioned in the previous section, the chemical shift acts as a perturbation to the main magnetic field that somewhat changes the precessional frequency. Unlike gradients, chemical shift is not spatial dependent in isotropic solution, and chemical shift can not be turned off or reversed.

2.1.6.2 J-coupling

Scalar or isotropic J-coupling is the result of energy changes in a system of bonded nuclei due to the pairing of the electron spins. In the absence of bonding, it is slightly favorable energetically for electrons to polarize in the same direction as a nearby nucleus because they have opposite gyromagnetic ratios. However, the Pauli exclusion principle requires bonded electrons to have opposite spin quantum numbers, which can force the spin system into a slightly higher or lower energy state. Only nuclei separated by a few bonds or less are close enough to experience J-coupling.

The Hamiltonian for this interaction is given by

$$\begin{aligned} H_J &= 2\pi J \mathbf{I}_1 \cdot \mathbf{I}_2 \\ &= 2\pi J I_{1z} I_{2z} \quad (\text{weak coupling}) \end{aligned} \tag{2.27}$$

The value of J is typically given in Hz and does not vary with field strength.

2.1.6.3 Dipole-dipole coupling

Dipole-dipole coupling is the direct influence of the magnetic field generated by a nuclear spin on spins that are nearby. Dipole-dipole coupling only depends on the proximity between the coupled spins and therefore can be intra or intermolecular. In isotropic solution, motional averaging causes the dipole-dipole Hamiltonian H_{DD} to be approximately 0. Nevertheless, these interactions are primarily responsible for the slight field perturbations that cause relaxation and result in a nonzero thermal equilibrium magnetization described later.

2.1.7 Spin ensembles and the density operator

The preceding sections detailed the effects of the MRI/MRS Hamiltonians on single wavefunctions. However, direct application to each wavefunction would

require calculations over all of the spins in a tissue or sample, which is highly impractical. As mentioned in the previous section, the main magnetic field and the RF field greatly dominate the equations of motion, and therefore, the spins in the tissue can be considered an *ensemble* of identical, independent spins. Since all the measurements we are interested in are macroscopic, calculations can be done on the quantum state of the entire ensemble without worrying about the individual spins.

To see how this is done, we first consider the expectation value of an observable operator Q for a spin with a single wavefunction.

$$\begin{aligned}
\langle Q \rangle &= \langle \Psi | Q | \Psi \rangle \\
&= \sum_{j=\alpha,\beta} \sum_{j'=\alpha,\beta} \langle \psi_{j'} | c_{j'}^* Q_{j',j} c_j | \psi_j \rangle \\
&= \sum_{j=\alpha,\beta} \sum_{j'=\alpha,\beta} c_{j'}^* c_j Q_{j',j} \langle \psi_{j'} | \psi_j \rangle \\
&= \sum_{j=\alpha,\beta} \sum_{j'=\alpha,\beta} c_{j'}^* c_j Q_{j',j} \\
&= \text{Tr} \left\{ |\Psi\rangle \langle \Psi| Q \right\}
\end{aligned} \tag{2.28}$$

where ψ_j are the basis states, and Tr is the trace. For an ensemble of spins, this calculation is summed over all the different spin states weighted by their fractional populations w_i

$$\begin{aligned}
\langle Q \rangle &= \sum_i w_i \langle \Psi_i | Q | \Psi_i \rangle \\
&= \sum_i \text{Tr} \left\{ w_i |\Psi_i\rangle \langle \Psi_i| Q \right\} \\
&= \text{Tr} \left\{ \sum_i (w_i |\Psi_i\rangle \langle \Psi_i|) Q \right\} \\
&= \text{Tr} \{ \rho Q \}
\end{aligned} \tag{2.29}$$

where the *density operator* is defined as

$$\rho = \sum_i (w_i |\Psi_i\rangle \langle \Psi_i|) \tag{2.30}$$

The time evolution of Eq. (2.30) is given by

$$\begin{aligned}
\frac{d}{dt}\rho &= \frac{d}{dt} \sum_i w_i |\Psi_i\rangle \langle \Psi_i| \\
&= \sum_i w_i \frac{d}{dt} (|\Psi_i\rangle \langle \Psi_i|) \\
&= \sum_i w_i \left(\frac{d}{dt} |\Psi_i\rangle \langle \Psi_i| + |\Psi_i\rangle \frac{d}{dt} \langle \Psi_i| \right) \\
&= \sum_i w_i \left(-iH |\Psi_i\rangle \langle \Psi_i| + |\Psi_i\rangle \langle \Psi_i| iH \right) \\
&= -i \left(H \sum_i w_i |\Psi_i\rangle \langle \Psi_i| - \sum_i w_i |\Psi_i\rangle \langle \Psi_i| H \right) \\
&= -i [H, \rho]
\end{aligned} \tag{2.31}$$

which is called the Liouville-von Neumann equation (LVNE). The solution to the LVNE for a time independent Hamiltonian is given by

$$\begin{aligned}
\rho(t) &= \sum_i (w_i |\Psi_i(t)\rangle \langle \Psi_i(t)|) \\
&= \sum_i \left(w_i \exp(-iHt) |\Psi_i(0)\rangle \langle \Psi_i(0)| \exp(iHt) \right) \\
&= \exp(-iHt) \left(\sum_i w_i |\Psi_i(0)\rangle \langle \Psi_i(0)| \right) \exp(iHt) \\
&= \exp(-iHt) \rho(0) \exp(iHt)
\end{aligned} \tag{2.32}$$

Though the density operator carries all the information necessary for any macroscopic measurements on the ensemble, it is not uniquely specified by the ensemble. Entirely different spin systems can have the same density operator provided that these spin systems cannot be distinguished by macroscopic measurements. For instance, an unpolarized ensemble could be composed of 50% spin up wavefunctions and 50% spin down wavefunctions or of general wavefunctions whose moments point in arbitrary directions as shown in Fig 2.1, with the latter being the correct interpretation physically.




Figure 2.1: Figure representing the magnetic moments of nuclear spins in an unpolarized sample. Reproduced with permission from [75].

2.1.7.1 Liouville-von Neumann equation in the rotating frame

The LVNE equation can be solved in the rotating frame using a similar equation to (2.31)

$$\frac{d}{dt}\rho^r = -i [H^r, \rho^r] \quad (2.33)$$

provided that the Hamiltonian is transformed to the rotating frame using Eq. (2.14) and the density operator is also transformed by the following

$$\rho^r = \exp(i\omega_{ref}tI_z)\rho \exp(-i\omega_{ref}tI_z) \quad (2.34)$$

2.1.8 Populations and coherences

Diagonal elements of the density operator matrix are called *populations*, while off-diagonal elements are called *coherences*. The populations are real, positive, and normalized so that

$$\rho_\alpha + \rho_\beta = c_\alpha c_\alpha^* + c_\beta c_\beta^* = 1 \quad (2.35)$$

It is tempting to look at ρ_α as the fraction of spins that are spin up and ρ_β as the fraction of spins that are spin down. However, we know that the individual spin states are in a superposition of those two states. What ρ_α actually represents is the fraction of measurements of z angular momentum would give $+\hbar/2$, while ρ_β is the fraction of measurements that would give $-\hbar/2$. Alternatively, ρ_α gives the fraction of energy measurements that would give $-\hbar\omega_0/2$, and ρ_β gives the fraction that would give $+\hbar\omega_0/2$. The difference between the populations gives the net longitudinal polarization.

The coherences, on the other hand, are complex numbers associated with transverse magnetization. Coherence exists between two eigenstates and is characterized by its order number, which is the difference in angular momentum between those eigenstates. For the non-interacting spin-1/2 particles described here, there are only two coherences with two orders, ± 1 -quantum coherences. The $+1$ -quantum coherence is given by

$$\rho_+ = \overline{c_\alpha c_\beta^*} \quad (2.36)$$

and the -1 -quantum coherence is given by

$$\rho_- = \overline{c_\beta c_\alpha^*} \quad (2.37)$$

The \pm -quantum coherences are conjugates of each other as seen from (2.36) and (2.37).

Though populations and coherences seem distinct, they are really two of the same thing, representing polarization. The population difference described above in the basis of eigenstates of z angular momentum would appear as coherence in another basis.

As coherences are actually a type of phase coherence between spins, gradients can be used to dephase them by fanning out the spins over a tissue volume. Dephasing is proportional to both the gradient moment and the coherence order, so gradients are a means to select a coherence transfer pathway in a single shot

as opposed to phase cycling.

2.1.9 Thermal equilibrium

Thermal equilibrium is the state of the ensemble after a long time without external probing. In thermal equilibrium, the coherences, and therefore the transverse magnetization, are zero, and the populations are given by the Boltzmann distribution

$$\rho_j = \frac{\exp(-m_j \hbar \omega_0 / k_B T)}{\sum_{j'=\alpha,\beta} \exp(-m_{j'} \hbar \omega_0 / k_B T)} \quad (2.38)$$

Since $\hbar \omega_0 / k_B T \equiv \mathcal{B} \sim 10^{-4}$, there is very little difference in the populations. Using the first two terms in the power series expansion of the exponentials shows that

$$\begin{aligned} \rho_\alpha &= \frac{1}{2} + \frac{\mathcal{B}}{4} \\ \rho_\beta &= \frac{1}{2} - \frac{\mathcal{B}}{4} \end{aligned} \quad (2.39)$$

All the coherences are 0 at thermal equilibrium because of the random transverse orientations of the magnetic moments. These thermal equilibrium conditions are illustrated in Fig. 2.2.

For any spin that is in a large B_0 field, the solution to the TDSE in Eq. (2.8) shows that the magnetic moment precesses around the z -axis at the Larmor frequency. This is true of any wavefunction and continues indefinitely without any change. It is reasonable then to question where the population difference at equilibrium comes from. While it's true that the main magnetic field dominates the inter and intramolecular interactions, those internal interactions contribute to the overall magnetic field and can alter its direction on the order of thousandths of a degree. These alterations fluctuate very quickly and act in random directions, yet they are able to shift the precession cone of a spin over time. Because of the slight energy preference for spins to align with the field, more spins are shifted up than down.




Figure 2.2: Illustration of the magnetization at thermal equilibrium. Spins are oriented almost isotropically in three dimensions, but there is a *small* net polarization in the vertical direction indicated by the thick green arrow. All spins precess around the z -axis as indicated by the red arrow. Reproduced with permission from [75].

The macroscopic magnetization vectors are related to the elements of the density operator by

$$\begin{aligned} M_z &= \frac{2}{\mathcal{B}}(\rho_\alpha - \rho_\beta) \\ M_x &= \frac{4}{\mathcal{B}}\text{Re}(\rho_-) \\ M_y &= \frac{4}{\mathcal{B}}\text{Im}(\rho_-) \end{aligned} \quad (2.40)$$

2.1.10 Product operators

Just as the wavefunction in Eq. (2.6) can be written as a linear combination of orthonormal basis states, the density operator can be written as a linear combination of orthonormal basis operators A .

$$\rho(t) = \sum_{k=1}^{4^N} a_k(t) A_k \quad (2.41)$$

where N is the number of independent spins in the system. For instance, I_x , I_y , and I_z make up a basis along with the identity matrix for an ensemble of independent ($N = 1$) spins. Writing the density operator as a sum naturally extends the expectation value (Eq. (2.29)) and LVNE (Eqs. (2.31) and (2.32)) to

$$\langle Q \rangle = \sum_k a_k \text{Tr}\{A_k Q\} \quad (2.42a)$$

$$\frac{d}{dt}\rho(t) = -i \sum_k a_k(t) [H, A_k] \quad (2.42b)$$

$$\rho(t) = \sum_k a_k(0) \exp(-iHt) A_k \exp(iHt) \quad (2.42c)$$

There are several advantages to writing the density operator this way. First, the magnetization is easily read off from the coefficients without having to compute the sum

$$M_z = a_z, \quad M_x = a_x, \quad M_y = a_y \quad (2.43)$$

(ignoring physically meaningless scale factors). Secondly, we only have to consider a finite number of evolutions in Eq. (2.42c) since we are only concerned with a few

Hamiltonians and the number of basis operators is limited to N . Lastly, this expansion scales easily to higher order spin systems. By contrast, the density matrix is $4^N \times 4^N$ and becomes very large as the number of spins increases. Calculating the basis operator matrices for a higher order spin system is done by taking the Kronecker product of all the combinations of 1-spin operator matrices (hence the name *product operators*). For example, a 2-spin system has basis operators like $2I_{1x}I_{2y}$, $2I_{1z}I_{2z}$, and 1_1I_{2x} . (The 2 is simply for normalization purposes, and the 1 is the identity matrix.) For multiple spin systems, operators with more than one transverse component represent multiple quantum coherences, while those with exactly one transverse component represent single quantum coherences.

Product operator calculations routinely make use of the relationship in Eq. (2.15). Shown below are example calculations under different Hamiltonians on the I_x operator

$$\begin{aligned}
I_x &\xrightarrow{\theta I_y} \cos \theta I_x - \sin \theta I_z && \text{(RF pulse)} \\
I_x &\xrightarrow{\Omega t I_z} \cos(\Omega t) I_x + \sin(\Omega t) I_y && \text{(chemical shift)} \\
I_x &\xrightarrow{2\pi J_{12} t I_{1z} I_{2z}} \cos(\pi J_{12} t) I_x + \sin(\pi J_{12} t) 2I_{1y} I_{2z} && \text{(weak J-coupling)}
\end{aligned} \tag{2.44}$$

For product operators on multiple spin systems, the same rules can be applied to the active part only. For example,

$$I_{1,x} I_{2,z} \xrightarrow{\theta I_{1,y}} (\cos \theta I_{1,x} - \sin \theta I_{1,z}) I_{2,z} \quad \text{(RF pulse)} \tag{2.45}$$

The product operator formalism allows intuitive and relatively simple calculation of the effects of various pulse sequences. One drawback though is that the J-coupling rules are strictly valid only for the weak coupling regime.

2.1.11 Observable operators

From Eq (2.43), it is clear that operators I_x and I_y give rise to the observable signal for a single spin system, while the operator I_z does not. The situation is

less obvious for a coupled spin system though, so a more general approach will be shown.

The observation operator representing the complex magnetization is given by $I^+ = I_x + iI_y$, so the measured signal is given by

$$\begin{aligned}\langle I^+ \rangle(t) &= \text{Tr}\{\rho(t)I^+\} \\ &= \text{Tr}\{\rho(t)I_x\} + i \text{Tr}\{\rho(t)I_y\} \\ &= a_x(t) + i a_y(t)\end{aligned}\tag{2.46}$$

since the basis operators are all orthogonal and again ignoring meaningless scaling. This result is general and shows that the observable signal is *always* the instantaneous sum of the I_x and I_y coefficients.

2.1.11.1 In phase terms

Consider a experiment in which $\rho = I_x$ at the start of acquisition. For a single spin system, $a_x(t) = \cos(\Omega t)$ and $a_y(t) = \sin(\Omega t)$, and the signal precesses around the complex unit circle with frequency Ω in the rotating frame

$$\exp(i\Omega t)\tag{2.47}$$

as known classically.

If the system consists of two weakly coupled spins, the free precession Hamiltonian must be the sum of chemical shift terms *for each spin* and a coupling term between spins

$$H^r = \Omega_1 I_{1z} + \Omega_2 I_{2z} + 2\pi J I_{1z} I_{2z}\tag{2.48}$$

These terms commute with one another and can therefore be taken in any order. If we consider the signal that evolves from an I_{1x} operator, the I_{2z} chemical shift Hamiltonian has no influence. The I_{1z} Hamiltonian leads to

$$I_{1x} \xrightarrow{\Omega_1 t I_{1z}} I_{1x} \cos(\Omega_1 t) + I_{1y} \sin(\Omega_1 t)\tag{2.49}$$

Now, consider the J coupling Hamiltonian term acting on the result

$$\begin{aligned}
& I_{1x} \cos(\Omega_1 t) + I_{1y} \sin(\Omega_1 t) \xrightarrow{2\pi J t I_{1z} I_{2z}} \\
& I_{1x} \cos(\Omega_1 t) \cos(\pi J t) + 2I_{1y} I_{2z} \cos(\Omega_1 t) \sin(\pi J t) \\
& + I_{1y} \sin(\Omega_1 t) \cos(\pi J t) - 2I_{1x} I_{2z} \sin(\Omega_1 t) \sin(\pi J t)
\end{aligned} \tag{2.50}$$

The observable signal is then read off as

$$\begin{aligned}
\cos(\Omega_1 t) \cos(\pi J t) + i \sin(\Omega_1 t) \cos(\pi J t) &= \exp(i\Omega_2 t) \cos(\pi J t) \\
&= \exp(i\Omega_2 t) \frac{1}{2} \{ \exp(i\pi J t) + \exp(-i\pi J t) \} \\
&= \frac{1}{2} \{ \exp(i[\Omega_1 + \pi J]t) + \exp(i[\Omega_1 - \pi J]t) \}
\end{aligned} \tag{2.51}$$

Three things are noteworthy in Eq (2.51):

1. J coupling caused the single resonance frequency in Eq (2.47) to split in units of πJ (the separation between peaks is $2\pi J$ or simply J in units of Hz)
2. Each peak is reduced in amplitude by half
3. Both peaks have the same polarity

This doublet is referred to as *in phase*. Similar results are obtained from starting operators I_{1y} , I_{2x} , and I_{2y} .

2.1.11.2 Antiphase terms

We showed in the last sections that only I_x and I_y operators yield observable magnetization and that a readout that starts at one of them results in an in phase doublet for a two spin system. Now, we show what happens during readout to an operator such as $2I_{1x}I_{2z}$, considering the same Hamiltonian in Eq (2.48)

$$2I_{1x}I_{2z} \xrightarrow{\Omega_1 t I_{1z}} 2I_{1x}I_{2z} \cos(\Omega_1 t) + 2I_{1y}I_{2z} \sin(\Omega_1 t) \tag{2.52}$$

Since the spin two contributions to the operators are in the z direction only, the I_{2z} chemical shift Hamiltonian does not evolve the system. The J coupling Hamiltonian evolution gives rise to

$$\begin{aligned}
& 2I_{1x}I_{2z} \cos(\Omega_1 t) + 2I_{1y}I_{2z} \sin(\Omega_1 t) \xrightarrow{\Omega_2 t I_{2z}} \\
& 2I_{1x}I_{2z} \cos(\Omega_1 t) \cos(\pi J t) + I_{1y} \cos(\Omega_1 t) \sin(\pi J t) \quad (2.53) \\
& + 2I_{1y}I_{2z} \sin(\Omega_1 t) \cos(\pi J t) - I_{1x} \sin(\Omega_1 t) \sin(\pi J t)
\end{aligned}$$

and the signal is therefore

$$\begin{aligned}
-\sin(\Omega_1 t) \sin(\pi J t) + i \cos(\Omega_1 t) \sin(\pi J t) &= i \exp(i\Omega_1 t) \sin(\pi J t) \\
&= i \exp(i\Omega_1 t) \frac{1}{2i} \{ \exp(i\pi J t) - \exp(-i\pi J t) \} \\
&= \frac{1}{2} \{ \exp(i[\Omega_1 + \pi J]t) - \exp(i[\Omega_1 - \pi J]t) \} \quad (2.54)
\end{aligned}$$

The peaks in Eq (2.54) have the same splitting and reduction in amplitude as those in Eq (2.51). However, the polarity of the peaks are now opposite one another. This doublet is referred to as *antiphase*. Similar results are obtained starting with the operators $2I_{1y}I_{2z}$, $2I_{1z}I_{2x}$, and $2I_{1z}I_{2y}$.

Note that this antiphase doublet is unrelated to the phase of the Lorentzian lineshape and cannot be removed by any phase correction. Antiphase multiplets are problematic in vivo where the broad linewidths blend the two peaks, and signal cancellation occurs. In contrast, broad linewidths make in phase multiplets appear as large singlets.

The operators that give rise to observable magnetization are single quantum coherences, either directly (in phase) or through evolution (antiphase). Zero and multiple quantum coherences are not observable.

2.2 Sampling overview

2.2.1 Cartesian sampling

Conventionally, sampling is a periodic measurement of a continuous signal. Because of the discrete nature of k-space sampling, the Fourier transform results in a periodic repetition of the image. For a bandwidth (BW) limited signal, if the sampling is done faster than the Nyquist rate

$$\Delta t = \frac{1}{2 \text{ BW}} \quad (2.55)$$

the image repetitions will appear outside the object. If sampling does not satisfy the Nyquist rate, image repetitions will occur on top of the object, and aliasing occurs. This is illustrated in Fig 2.3, comparing a Nyquist sampled image with one from 2x undersampling in the vertical direction. The point spread function (PSF) of each sampling pattern is shown, indicating the periodicity. For the fully sampled data, the PSF in Fig 2.3c shows a spike in the center with no other spike within the field of view (*FOV*). For the undersampled PSF in Fig 2.3d, the spikes occur at a distance of half the *FOV* apart.

The theory of compressed sensing shows that exact image reconstruction can be obtained for sub-Nyquist sampled data using nonuniform sampling (NUS) along with iterative reconstructions that promote sparsity [28, 48]. As NUS does not sample at the Nyquist rate, aliasing artifacts do occur. However, they are incoherent and give the appearance of noise. An image that was 2x undersampled randomly with uniform density is shown in Fig 2.4a with an image that was undersampled with a variable density in Fig 2.4b. The periodicity of the images are the same as the fully sampled image in 2.3a. Both PSFs are similarly incoherent, yet the variable density NUS samples more of the high SNR data and therefore, is more recognizable.

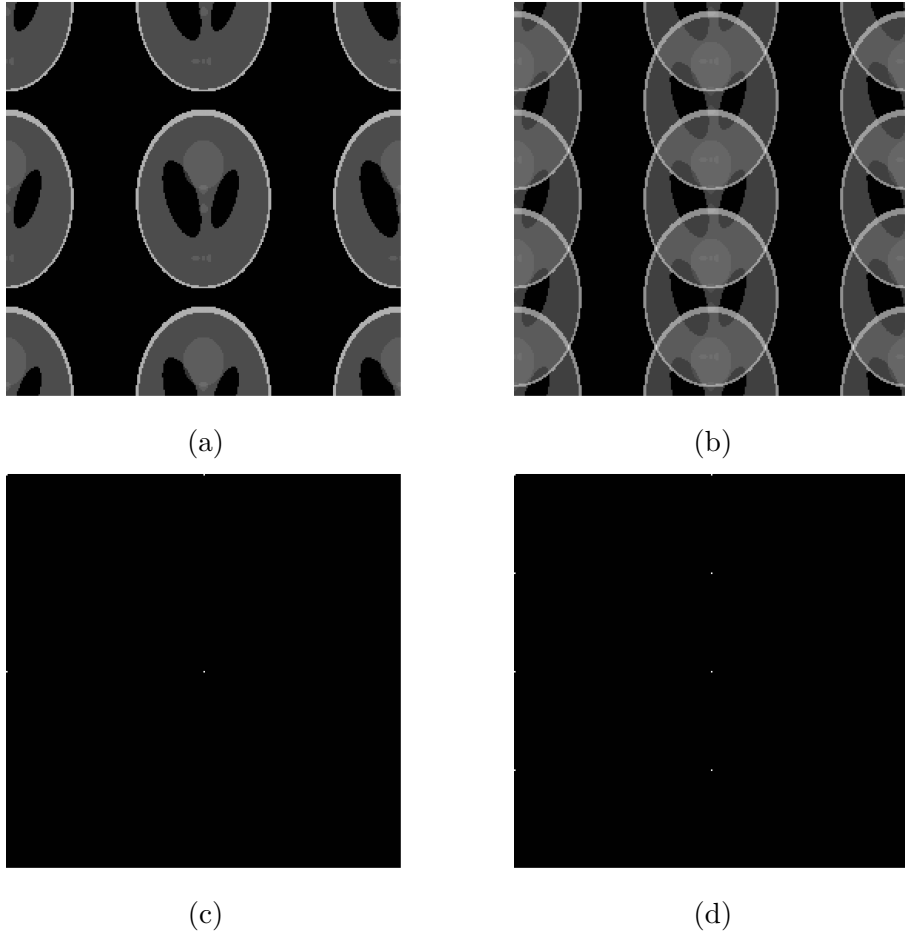


Figure 2.3: (a) Periodic sampling at the Nyquist rate. (b) Periodic sampling at half the Nyquist rate in the vertical direction results in coherent aliasing. (c-d) Point spread functions (PSF) of (a-b). PSFs were computed from the Fourier transform of the sampling pattern.

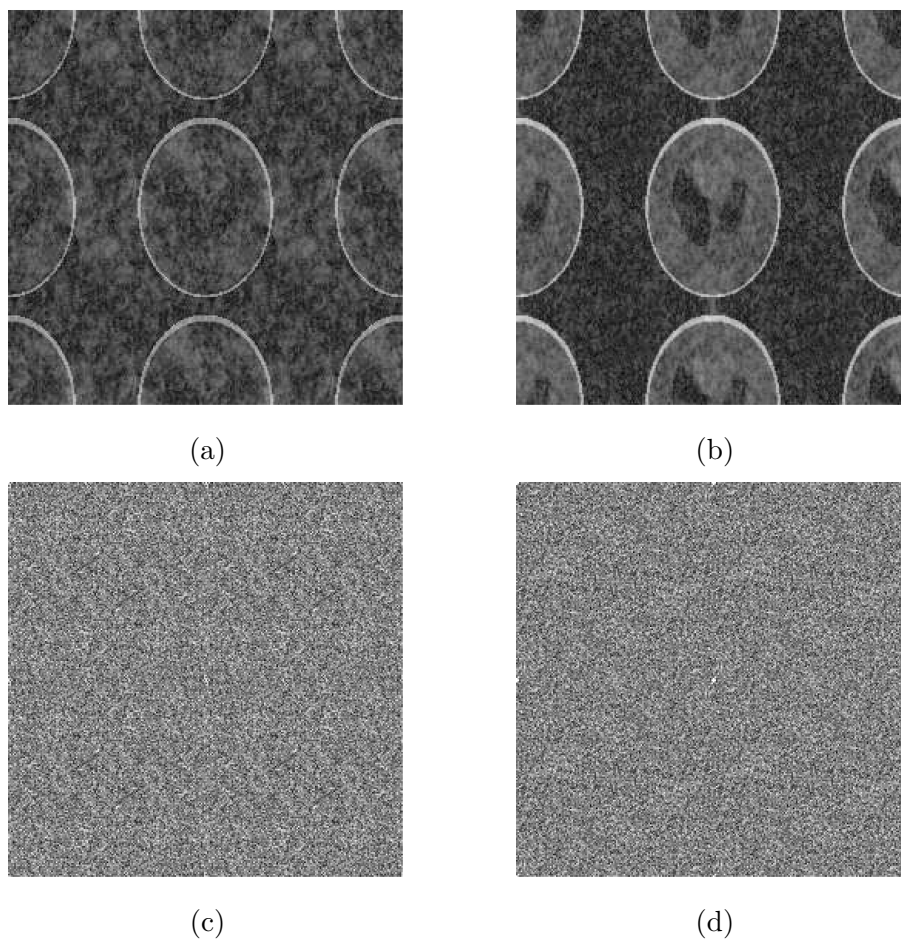


Figure 2.4: Incoherent aliasing artifacts with random, uniform density NUS (a) and variable density NUS (b). (c-d) PSFs of (a-b).

2.2.2 Non-Cartesian polar sampling

The conditions to satisfy the Nyquist criteria for polar sampling schemes are derived. In each case, it is assumed that the FOV of the final image is square (i.e. $FOV_x = FOV_y$) and that the k-space location furthest from the origin is k_{max} .

Figure 2.5 shows reconstruction from fully sampled polar data as well as from radially and circularly undersampled. It also shows the corresponding PSFs, highlighting the relative incoherence of radial undersampling compared to circular, though both are more incoherent than periodic Cartesian undersampling shown in Fig 2.3d.

2.2.2.1 Concentric circles

Here, it is envisioned that k-space has been sampled in concentric circles. In particular, each ring covers all four quadrants with an angular range $[0, 2\pi]$. The natural definition of the field of view uses the radial component

$$FOV = \frac{1}{\Delta k_r} \quad (2.56)$$

The largest ring satisfies

$$k_{max} = N_r \Delta k_r = \frac{N_r}{FOV} \quad (2.57)$$

where N_r is the number of rings, and Δk_r is the spacing of the rings. The number of angular samples N_θ satisfies

$$N_\theta \Delta \theta = 2\pi \quad (2.58)$$

for full circular sampling.

The incremental k-space sampling in the angular direction is dependent on the ring

$$\Delta k_\theta = k_r \Delta \theta \leq k_{max} \Delta \theta \quad (2.59)$$

From this point on, we will only consider the maximal angular spacing, and will require $k_{max}\Delta\theta \leq FOV^{-1}$ which can be rearranged to

$$\Delta\theta \leq 1/(FOV \cdot k_{max}) \quad (2.60)$$

Substituting in Eq (2.58) results in

$$2\pi/N_\theta \leq 1/(FOV \cdot k_{max}) \quad (2.61)$$

which can be rearranged to give

$$N_\theta \geq 2\pi \cdot FOV \cdot k_{max} \quad (2.62)$$

Substituting in Eq (2.57) gives the final result

$$N_\theta \geq 2\pi N_r \quad (2.63)$$

2.2.2.2 Radial spokes

Here, it is envisioned that k-space has been sampled radially, line-by-line. In particular, each line cuts through the origin from $[-k_{max}, +k_{max}]$, with a range of projections from $[0, \pi]$. Most of the ideas are the same as the previous section except now Eq (2.57) and (2.58) are replaced by

$$k_{max} = \frac{N_r}{2} \frac{1}{FOV} \quad (2.64)$$

and

$$N_\theta \Delta\theta = \pi \quad (2.65)$$

In essence, $N_{\theta,rad} \rightarrow N_{\theta,cir}/2$ and $N_{r,rad} \rightarrow 2N_{r,cir}$. Substituting these values into Eq (2.63) gives the relationship

$$2N_\theta \geq 2\pi \frac{N_r}{2} \quad (2.66)$$

which can be rewritten in final form as

$$N_\theta \geq \frac{\pi}{2} N_r \quad (2.67)$$




Figure 2.5: The top row shows fully sampled (left), radially undersampled (center), and circularly undersampled (right) polar data. The bottom row is the respective PSFs.

This relationship is equivalent to Eq (2.63) in the previous section as long as the correct definitions for N_θ and N_r are applied.

2.2.3 Sensitivity encoding

In certain cases, the incoherent aliasing caused by nonuniform sampling described in the previous section can be removed through iterative reconstruction. Those types of sparsity-promoting reconstructions cannot remove coherent aliasing though. This problem is illustrated in Fig 2.6 where there is no way to determine the relative contributions of the separate pixels from the left to the aliased pixel on the right.

Fortunately, there are other reconstruction techniques that can remove coherent aliasing artifacts specific to MRI by taking into account the receive coil sensitivities in a multichannel array [132, 161, 70, 103]. In order to visualize how

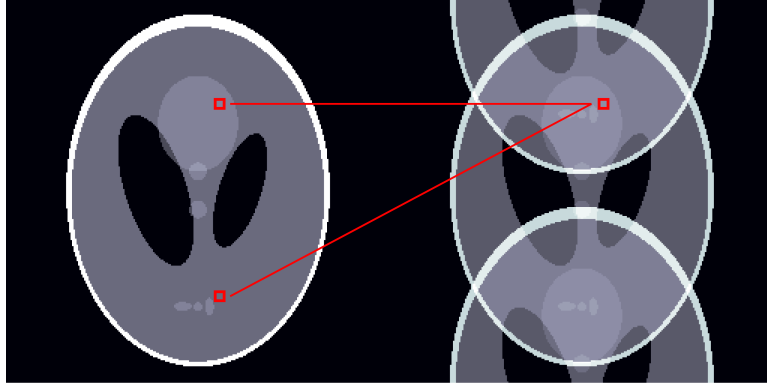


Figure 2.6: 2x undersampling in the vertical direction results in two pixels on the left summing to form the aliased pixel on the right. Here, there is no way to separate out the pixels from the image on the right.

this works in sensitivity encoding (SENSE) [132], consider a two coil array in which the first coil is located on the top of the image, and the second coil is located on the bottom of the image. Each coil will then see the actual image modulated by its sensitivity profile. Pixels near the top will appear brighter in the first coil, and pixels near the bottom will appear brighter in the second as illustrated in Fig 2.7. For N_x -fold undersampling in the x direction and N_y -fold undersampling in the y direction, the aliased image pixels, $a(x, y)$, can be expressed as a weighted sum of the pixels from the true image $u(x, y)$ for each coil c

$$a_c(x, y) = \sum_{n_x=0}^{N_x-1} \sum_{n_y=0}^{N_y-1} S_c\left(x + n_x \frac{FOV_x}{2}, y + n_y \frac{FOV_y}{2}\right) u\left(x + n_x \frac{FOV_x}{2}, y + n_y \frac{FOV_y}{2}\right) \quad (2.68)$$

where S_c are the coil sensitivity profiles. The specific pixel illustrated in Fig 2.7 can therefore be written as the system of equations

$$\begin{aligned} a_1 &= S_1(\text{low}) u(\text{low}) + S_1(\text{upp}) u(\text{upp}) \\ a_2 &= S_2(\text{low}) u(\text{low}) + S_2(\text{upp}) u(\text{upp}) \end{aligned} \quad (2.69)$$

where “low” is the lower pixel, and “upp” is the upper pixel. This illustrates how



Figure 2.7: Sensitivity weighted images from a coil above the image (top) and one below the image (bottom). In the top image, the top pixel contributes more to the aliased pixel, while in the bottom image, it is the bottom pixel that contributes more.

the SENSE formalism can be written as a matrix inversion

$$\mathbf{u} = (S'S)^{-1}S'\mathbf{a} \quad (2.70)$$

For maximal image SNR, a coil noise correlation matrix, Ψ , can be included to whiten the data if the system is overdetermined

$$\mathbf{u} = (S'\Psi^{-1}S)^{-1}S'\Psi^{-1}\mathbf{a} \quad (2.71)$$

2.2.4 Signal-to-noise

There are many factors that affect the signal-to-noise (SNR) of the final reconstructed data. Some are physical (field strength, coils, temperature, etc.), others are sequence-based (timing, flip angles, voxel size, resolution, sampling, etc.), while others are processing dependent (apodization, regularization, etc.). In this section, we will restrict the discussion to those due to undersampling.

For any N -fold undersampling, there is an SNR penalty proportional to \sqrt{N} simply due to the reduction in acquired data. This goes hand-in-hand with an N -fold acceleration in the same way that halving the number of averages also cuts the SNR by $\sqrt{2}$. There are also additional SNR enhancements or penalties depending on the reconstruction.

2.2.4.1 Nonuniform undersampling

With NUS, we have already seen from Fig 2.4 that the particular sampling density can affect the SNR when the higher signal points are preferentially sampled. However, the biggest influence on the SNR of NUS data is the reconstruction parameters. Generally speaking, compressed sensing reconstructions consist of a minimization problem featuring regularization terms and a data consistency term. Sparsity-promoting regularization terms such as total variation are effective at denoising [146], so CS reconstructions are also associated with an increase in the apparent SNR.

Some CS reconstructions solve an unconstrained optimization problem. In that case, the relative weighting of the data consistency term is crucial. An appropriate level of denoising enhances the image appearance, but the wrong weighting can lead to noisy images with residual artifacts if the data consistency term is overweighted. The problems with an underweighted data consistency term can be more subtle though, since the images produced will have the highest SNR

and look very clean. However, they are also accompanied by a loss of resolution and the removal of low amplitude signal that is potentially diagnostic. Therefore, it is best to be cautious and err towards overweighting the data consistency term.

Other reconstruction algorithms enforce a certain level of data consistency. Projection onto convex sets (POCS) is one that enforces strict data consistency in every iteration which can slow its convergence. The split Bregman algorithm [66] detailed in the next chapter solves the constrained problem by terminating as the data consistency term monotonically decreases to the desired level, allowing some flexibility in the level of denoising. Again, it is best to be cautious in deciding when to terminate the algorithm.

2.2.4.2 Sensitivity encoding

With SENSE reconstruction, there is a noise correlation between pixels as the encoding operator is no longer unitary [132], resulting in spatially dependent noise amplification. The reconstructed SNR is given by

$$\text{SNR} = \frac{\text{SNR}_{full}}{g\sqrt{N}} \quad (2.72)$$

where $g \geq 1$ is the additional noise penalty beyond that of undersampling and is known as the *geometry factor* since it is highly dependent on coil geometry. The geometry factor at pixel p is given by

$$g_p = \sqrt{[(S'\Psi^{-1}S)^{-1}]_{p,p} (S'\Psi^{-1}S)_{p,p}} \quad (2.73)$$

and is shown in Fig 2.8 for the example shown in Fig 2.7.




Figure 2.8: Geometry factor map for the undersampled coil images shown in Fig 2.7. The color scale is from 1 to 2.

CHAPTER 3

Iterative Nonlinear Reconstruction using Split Bregman Algorithms

3.1 Introduction

Formulating the compressed sensing problem is only a beginning step in the reconstruction. There are a nearly uncountable number of solvers, some general, some specific to certain regularizers or acquisition domains. Some solve the unconstrained problem only. For multidimensional spectroscopic imaging, the choice of solver is nontrivial as the datasets are much larger than those typically encountered in imaging and tend to require a higher number of iterations. In this dissertation, the split Bregman algorithm of Goldstein and Osher [66] was chosen for the majority of the processing due to its computational speed, adaptability to a variety of different regularizers, and scalability to large datasets. The split Bregman algorithm is from the class of Alternating Direction Method of Multiplier (ADMM) methods that replace a difficult to solve constrained problem with a sequence of simpler unconstrained subproblems [16]. It can handle multiple regularizers with ℓ_1 or TV norms.

3.2 Background into ℓ_1 solvers

The CS problem seeks to reconstruct data with maximal sparsity (i.e. minimum ℓ_1) in some transform domain. This sparsity requirement is what separates CS

from other regularized reconstructions and allows for sub Nyquist sampling. It is also what makes the problem more difficult to solve, having no closed form solution except in the most basic of cases.

In contrast, Tikhonov regularization is based on minimizing the ℓ_2 (or Euclidean)-norm ($\|u\|_2 = \sqrt{\sum_i u_i^2}$) and formulates the simplest unconstrained problem as

$$\min_u \frac{1}{2} \|u\|_2^2 + \frac{\lambda}{2} \|u - f\|_2^2 \quad (3.1)$$

where u is the reconstructed data and f is original data. Here, the square of the ℓ_2 norm has been used for simplicity as it has the same solution as the unsquared norm (with proper weighting of λ), and the original data is in the same domain as the reconstructed for simplicity. Differentiation and equating to zero then yields

$$u = \frac{f}{1 + \lambda} \quad (3.2)$$

Even with a more complicated sensing/blurring matrix A , the Tikhonov-regularized problem

$$\min_u \frac{1}{2} \|u\|_2^2 + \frac{\lambda}{2} \|Au - f\|_2^2 \quad (3.3)$$

still has closed form solution

$$u = \left(A'A + \frac{\lambda}{2} I \right)^{-1} A'f \quad (3.4)$$

where A' is the Hermitian conjugate of A , and I is the identity matrix.

The fundamental ℓ_1 minimization problem can be written in its unconstrained form as

$$\min_u \|u\|_1 + \frac{\lambda}{2} \|u - f\|_2^2 \quad (3.5)$$

In this simplest case, the u_i variables are independent and can be solved individually. This leads to a closed form solution

$$u = \frac{f}{|f|} \max(|f| - 1/\lambda, 0) \quad (3.6)$$

known as the *shrink* or *soft thresholding* function. However, in the more complicated problems

$$\begin{aligned} \min_u \|\Psi u\|_1 + \frac{\lambda}{2}\|u - f\|_2^2 \\ \min_u \|u\|_1 + \frac{\lambda}{2}\|Au - f\|_2^2 \\ \min_u \|\Psi u\|_1 + \frac{\lambda}{2}\|Au - f\|_2^2 \end{aligned}$$

the coupling between the ℓ_1 and ℓ_2 terms precludes a closed form solution like *shrink* and makes these require iterative solvers.

3.3 Constrained optimization with Bregman iteration

Consider a convex, nondifferentiable function $E(u)$ and a linear operator A . A generalized constrained optimization problem can be formulated as

$$\min_u E(u) \quad \text{s.t. } Au = f \quad (3.7)$$

A first step in order to solve (3.7) is to rewrite it as an unconstrained problem

$$\min_u E(u) + \frac{\lambda}{2}\|Au - f\|_2^2 \quad (3.8)$$

The unconstrained problem (3.8) is much easier to solve. However, its solution does not solve the constrained problem (3.7) unless $\lambda \rightarrow \infty$. Simply choosing an extremely large λ or using continuation methods that increase $\lambda_k \rightarrow \infty$ as $k \rightarrow \infty$ also lead to the condition number of the Hessian approaching infinity, making the solver less efficient and stable. An alternative way of solving the constrained problem as a series of unconstrained problems uses Bregman iteration [101].

Bregman iteration says that function E can be replaced in the unconstrained problem with its Bregman distance provided that the Bregman parameter (defined later) is also updated iteratively. The Bregman distance for the function $E(u)$ at the point u^k is:

$$D_E^p(u, u^k) = E(u) - E(u^k) - \langle p_u^k, u - u^k \rangle \quad (3.9)$$

where the Bregman parameter p_u^k is the subgradient of $E(u)$ at u^k , and the problem can then be reformulated as

$$\begin{aligned} u^{k+1} &= \min_u D_E^p(u, u^k) + \frac{\lambda}{2} \|Au - f\|_2^2 \\ p^{k+1} &= p^k - \lambda A'(Au^{k+1} - f) \end{aligned} \tag{3.10}$$

In [119], Osher *et al.* proved that $\|Au^k - f\|_2^2 \rightarrow 0$ as $k \rightarrow \infty$ under fairly weak assumptions, and therefore, the constrained problem is solved when infinite iterations are used. In addition, they proved that the decrease in the ℓ_2 -norm is monotonic. (Actually, their proof was for a more general function than the ℓ_2 -norm data consistency constraint, but this dissertation is only concerned with that particular form.) Compared to a projection onto convex sets (POCS) method that enforces strict data consistency at every iteration, Bregman iteration only achieves strict data consistency as $k \rightarrow \infty$. While this may appear to be a disadvantage, real data is typically noisy, and a constraint of the form $\|Au - f\|_2^2 < \sigma^2$ is often more appropriate. In that case, the monotonic decrease of the data consistency term in Bregman iteration gives an additional stopping criteria of the algorithm.

In [180], it was shown that the steps in (3.10) are equivalent to the following simpler steps provided that A is linear

$$\begin{aligned} u^{k+1} &= \min_u E(u) + \frac{\lambda}{2} \|Au - f^k\|_2^2 \\ f^{k+1} &= f^k + f - Au^{k+1} \end{aligned} \tag{3.11}$$

Equation (3.11) tells us that the constrained problem in (3.7) can be solved by solving the simpler unconstrained problem (3.8) and updating the data term by “adding the noise back in”.

Bregman iterations typically converge quickly if $E(u)$ is an ℓ_1 based term compared to continuation methods. In addition, the value of the regularization parameter λ does not change

3.4 Split Bregman algorithm and variants

In this section, I will review the split Bregman algorithm proposed in [66] with isotropic and anisotropic total variation regularization for single coil acquisitions. I will show a variant of the algorithm that minimizes an ℓ_1 -norm term only. I will then describe how these algorithms can be extended to problems with multiple, mixed regularizers.

3.4.1 Anisotropic TV -regularized

Total variation (TV) was shown to be an effective regularizer for denoising in the seminal paper by Rudin, Osher, and Fatemi [146]. However, the TV -regularized problem is difficult for some solvers since the gradient operator is not invertible like the wavelet or discrete cosine transforms.

First, consider anisotropic TV over the i points of u

$$TV_{ani}(u) = \sum_i |\nabla_x u|_i + |\nabla_y u|_i + |\nabla_z u|_i \quad (3.12)$$

Anisotropic TV is somewhat easier to solve than isotropic TV , which is generally considered the superior regularizer. However, TV_{ani} could be appropriate for volume-based localization methods in MRSI in which the excitation volume is more or less a rectangular solid.

The anisotropic problem is then given by

$$\min_u |\nabla_x u| + |\nabla_y u| + |\nabla_z u| \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (3.13)$$

where R is the sampling operator and \mathcal{F} is the Fourier transform. Here we use the shorthand notation $|\nabla u| = \sum_i |\nabla u|_i$. We know from the previous section that this problem can be solved by iteratively solving the unconstrained version and

adding the noise back in to f .

$$u^{k+1} = \min_u |\nabla_x u| + |\nabla_y u| + |\nabla_z u| + \frac{\mu}{2} \|R\mathcal{F}u - f^k\|_2^2 \quad (3.14a)$$

$$f^{k+1} = f^k + f - Au^{k+1} \quad (3.14b)$$

The second step is an explicit computation and easily performed, so we will turn our attention to the unconstrained optimization of the first step. This problem does not have a closed form solution as written due to the coupling between the ℓ_1 and ℓ_2 terms.

In order to decouple the terms, auxiliary variables $d_{x,y,z} = \nabla_{x,y,z}u$ are introduced. Then the original *unconstrained* problem can be written as an augmented *constrained* problem

$$\min_u \sum_{i=x,y,z} |d_i| + \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 \quad \text{s.t. } d_x = \nabla_x u, d_y = \nabla_y u, d_z = \nabla_z u \quad (3.15)$$

The Bregman iteration procedure can thus be applied to this augmented constrained problem.

$$(u^{k+1}, d_x^{k+1}, d_y^{k+1}, d_z^{k+1}) = \min_{u, d_x, d_y, d_z} \sum_{i=x,y,z} |d_i| + \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 + \frac{\lambda}{2} \|d_i - \nabla_i u - b_i^k\|_2^2 \quad (3.16a)$$

$$b_x^{k+1} = b_x^k + \nabla_x u^{k+1} - d_x^{k+1} \quad (3.16b)$$

$$b_y^{k+1} = b_y^k + \nabla_y u^{k+1} - d_y^{k+1} \quad (3.16c)$$

$$b_z^{k+1} = b_z^k + \nabla_z u^{k+1} - d_z^{k+1} \quad (3.16d)$$

Though the unconstrained problem in (3.16a) has additional terms compared to the original unconstrained problem in (3.14a), it can be decoupled and solved for u, d_x, d_y , and d_z sequentially where in each step, the uninvolved variables are

held constant.

$$u^{k+1} = \min_u \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 + \sum_{i=x,y,z} \frac{\lambda}{2} \|d_i^k - \nabla_i u - b_i^k\|_2^2 \quad (3.17a)$$

$$d_x^{k+1} = \min_{d_x} |d_x| + \frac{\lambda}{2} \|d_x - \nabla_x u^k - b_x^k\|_2^2 \quad (3.17b)$$

$$d_y^{k+1} = \min_{d_y} |d_y| + \frac{\lambda}{2} \|d_y - \nabla_y u^k - b_y^k\|_2^2 \quad (3.17c)$$

$$d_z^{k+1} = \min_{d_z} |d_z| + \frac{\lambda}{2} \|d_z - \nabla_z u^k - b_z^k\|_2^2 \quad (3.17d)$$

The d optimizations in (3.17) are of the form of (3.5) and can therefore be solved explicitly

$$d_{i=x,y,z} = \mathit{shrink}(\nabla_i u^{k+1} + b_i^k, 1/\lambda) \quad (3.18)$$

The u subproblem in (3.17a) is composed of multiple ℓ_2 -norm terms and is differentiable with derivative given by

$$\mu \mathcal{F}' R' (R\mathcal{F}u - f) - \sum_{i=x,y,z} \lambda \nabla'_i (d_i^k - \nabla_i u - b_i^k)$$

Setting the derivative equal to 0 and grouping the u terms yields

$$\left(\mu \mathcal{F}' R' R\mathcal{F} + \sum_{i=x,y,z} \lambda \nabla'_i \nabla_i \right) u = \mu \mathcal{F}' R' f + \sum_{i=x,y,z} \lambda \nabla'_i (d_i^k - b_i^k) \quad (3.19)$$

A key aspect to the speed of the split Bregman algorithm is that the multiplier of u is block circulant with circulant blocks (BCCB) and therefore diagonalizable using Fourier transforms. The left side of (3.19) can then be rewritten as

$$\mathcal{F}' K \mathcal{F} u$$

where $K = (\mu R' R - \lambda \mathcal{F} \Delta \mathcal{F}')$ with $\Delta = \sum_i \nabla'_i \nabla_i$ being the Laplacian operator. The u subproblem solution is then

$$u = \mathcal{F}' K^{-1} \mathcal{F} \left(\mu \mathcal{F}' R' f + \sum_{i=x,y,z} \lambda \nabla'_i (d_i^k - b_i^k) \right) \quad (3.20)$$

which is only a matter of some simple computation, Fourier transforms, and point-wise division since K is diagonal.

The entire algorithm is setup as an inner loop to satisfy the augmented constrained problem in (3.15) and an outer loop to satisfy the data consistency constraint of the original problem in (3.13). In [66], the authors recommend defining the inner loop with a fixed number of iterations as opposed to a requirement of full convergence in order to save computational time. This is especially true when σ^2 is relatively large. In our implementations, we used three stopping criteria for the outer loop: (1) a maximum number of iterations, (2) desired data consistency, (3) convergence of solution. Tolerances were generally set to ensure that (2) was the most influential and that (3) was unimportant.

```

Precompute  $K^{-1} = (\mu R'R - \lambda \mathcal{F} \Delta \mathcal{F}')^{-1}$ 
repeat
  for  $k = 1$  to  $N$  do
    Compute  $u^{k+1}$  using (3.20).
     $d_x^{k+1} = \mathit{shrink}(\nabla_x u^{k+1} + b_x, 1/\lambda)$ 
     $d_y^{k+1} = \mathit{shrink}(\nabla_y u^{k+1} + b_y, 1/\lambda)$ 
     $d_z^{k+1} = \mathit{shrink}(\nabla_z u^{k+1} + b_z, 1/\lambda)$ 
     $b_x \leftarrow b_x + \nabla_x u^{k+1} - d_x^{k+1}$ 
     $b_y \leftarrow b_y + \nabla_y u^{k+1} - d_y^{k+1}$ 
     $b_z \leftarrow b_z + \nabla_z u^{k+1} - d_z^{k+1}$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + f - R\mathcal{F}u^{k+1}$ 
until  $\|R\mathcal{F}u^{k+1} - f\|_2^2 < \sigma^2$ 

```

Figure 3.1: Split Bregman algorithm for anisotropic TV regularization

3.4.2 Isotropic TV -regularized

The isotropic total variation is given by

$$TV_{iso}(u) = \sum_i \sqrt{|\nabla_x u|_i^2 + |\nabla_y u|_i^2 + |\nabla_z u|_i^2} \quad (3.21)$$

and gives an overall measure of variation at a given point independent of direction. It is the more appropriate regularizer for objects that are not aligned with the coordinate axis in all three dimensions typical of MRI and large FOV MRSI scans.

The algorithm to solve the isotropic TV -regularized problem

$$\min_u TV_{iso}(u) \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (3.22)$$

is similar to the anisotropic case except that the gradient terms are coupled. Nevertheless, the same auxiliary variables $d_{x,y,z} = \nabla_{x,y,z}u$ are introduced and the augmented constrained problem can be written as

$$\min_{u,d_x,d_y,d_z} \|(d_x, d_y, d_z)\|_2 + \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 \quad \text{s.t.} \quad d_x = \nabla_x u, d_y = \nabla_y u, d_z = \nabla_z u \quad (3.23)$$

where

$$\|(d_x, d_y, d_z)\|_2 = \sum_i \sqrt{d_{x,i}^2 + d_{y,i}^2 + d_{z,i}^2}$$

and the sum is taken over all i data points. This obviously changes how the d subproblems must be solved, but less obvious is that it *does not* change the solution to the u subproblem because the gradient terms remain uncoupled in those terms related only to the variable substitution.

Fortunately, the d subproblems, though coupled, can still be solved explicitly using a generalized shrinkage formula

$$(x_s, y_s, z_s) = \mathit{shrink3}(x, y, z, 1/\lambda) \quad (3.24)$$

that computes

$$\{x, y, z\}_s = \frac{\{x, y, z\}}{\sqrt{x^2 + y^2 + z^2}} \max(\sqrt{x^2 + y^2 + z^2} - 1/\lambda, 0)$$

```

Precompute  $K^{-1} = (\mu R'R - \lambda \mathcal{F} \Delta \mathcal{F}')^{-1}$ 
repeat
  for  $k = 1$  to  $N$  do
    Compute  $u^{k+1}$  using (3.20).
     $(d_x^{k+1}, d_y^{k+1}, d_z^{k+1}) = \text{shrink3}(\nabla_x u^{k+1} + b_x, \nabla_y u^{k+1} + b_y, \nabla_z u^{k+1} + b_z, 1/\lambda)$ 
     $b_x \leftarrow b_x + \nabla_x u^{k+1} - d_x^{k+1}$ 
     $b_y \leftarrow b_y + \nabla_y u^{k+1} - d_y^{k+1}$ 
     $b_z \leftarrow b_z + \nabla_z u^{k+1} - d_z^{k+1}$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + f - R\mathcal{F}u^{k+1}$ 
until  $\|R\mathcal{F}u^{k+1} - f\|_2^2 < \sigma^2$ 

```

Figure 3.2: Split Bregman algorithm for isotropic TV regularization

3.4.3 ℓ_1 -regularized

There are many solvers for ℓ_1 -regularized compressed sensing problems that work when the regularizer is a pure ℓ_1 -norm or those with an invertible transform. Some are even faster than the split Bregman algorithm. However, split Bregman is still a fairly efficient solver for these types of problems as described below.

The ℓ_1 -regularized problem can be written as

$$\min_u |Wu|_1 \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (3.25)$$

where W is some invertible transform. Here, the only auxiliary variable we need to introduce is $w = Wu$ which leads to the unconstrained augmented problem

$$(u^{k+1}, w^{k+1}) = \min_{u,w} |w|_1 + \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 + \frac{\lambda}{2} \|w - Wu - b^k\|_2^2$$

that can be broken up into the sequential subproblems

$$u^{k+1} = \min_u \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 + \frac{\lambda}{2} \|w^k - Wu - b^k\|_2^2 \quad (3.26a)$$

$$w^{k+1} = \min_w |w|_1 + \frac{\lambda}{2} \|w - Wu^k - b^k\|_2^2 \quad (3.26b)$$

The u subproblem in (3.26a) is again differentiable with explicit solution

$$u = \mathcal{F}'K^{-1}\mathcal{F}(\mu\mathcal{F}'R'f + \lambda W'(w^k - b^k)) \quad (3.27)$$

where K is now $\mu R'R + \lambda$ since it was previously assumed that $W'W$ is the identity. Clearly, K is still diagonal and simple to invert. The w subproblem just requires straightforward application of the *shrink* function.

If the reconstructed data is self sparse, W can simply be the identity transform.

```

Precompute  $K^{-1} = (\mu R'R + \lambda)^{-1}$ 
repeat
  for  $k = 1$  to  $N$  do
    Compute  $u^{k+1}$  using (3.27).
     $w^{k+1} = \mathit{shrink}(Wu^{k+1} + b, 1/\lambda)$ 
     $b \leftarrow b + Wu^{k+1} - w^{k+1}$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + f - R\mathcal{F}u^{k+1}$ 
until  $\|R\mathcal{F}u^{k+1} - f\|_2^2 < \sigma^2$ 

```

Figure 3.3: Split Bregman algorithm for ℓ_1 regularization

3.4.4 Multiple regularizers

In some contexts, having multiple regularizer terms could be ideal. For instance, a combination of wavelet and TV regularizers was presented in Lustig's original

paper [102]. The split Bregman algorithm is particularly well suited for such problems with little added computational cost per term. Each term simply requires an additional auxiliary variable and Bregman parameter, and usually, the auxiliary variable subproblem can be solved with some soft thresholding function. In fact, we can think of the anisotropic TV algorithm in Figure 3.1 as having three individual regularizers (variation in x , variation in y , and variation in z), but a mixture of total variation terms and ℓ_1 terms can be constructed in the same way.

As mentioned above, the particular combination of wavelet and TV regularization has been successfully applied to MRI data. The added wavelet term can help reduce the sometimes “blocky” reconstruction performance of TV .

This dissertation focuses on the identity-transformed ℓ_1 and TV regularizers. While these two can of course be used together, there is little justification to. The reason is that the TV can be thought of as the ℓ_1 -norm of the finite differences which are typically applied over only some of the dimensions. Therefore, the self sparsity of the remaining dimensions is already implied, and a second regularizer would be redundant. A potential exception would be if the finite differences were taken across *all* the dimensions, but the data is also self sparse.

3.4.5 Intermediate updating of Bregman variables

In the preceding algorithms, the Bregman variables b were updated once per inner loop after the image variable and the auxiliary variables were updated. (The order in which the variables and Bregman variables are updated may make *small* differences in convergence rate but is generally not considered crucial.) Allison *et.al.* showed with a similar ADMM algorithm that the convergence rate can be improved by updating the Bregman variables after *each* variable is updated [6]. This type of intermediate updating is computationally inexpensive and easy to implement. Below is what the split Bregman algorithm for ℓ_1 regularization

shown in Figure 3.3 would look like with intermediate updating of its Bregman variable.

```

Precompute  $K^{-1} = (\mu R'R + \lambda)^{-1}$ 
repeat
  for  $k = 1$  to  $N$  do
    Compute  $u^{k+1}$  using (3.27).
     $b \leftarrow b + Wu^{k+1} - w^k$ 
     $w^{k+1} = \text{shrink}(Wu^{k+1} + b, 1/\lambda)$ 
     $b \leftarrow b + Wu^{k+1} - w^{k+1}$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + f - R\mathcal{F}u^{k+1}$ 
until  $\|R\mathcal{F}u^{k+1} - f\|_2^2 < \sigma^2$ 

```

Figure 3.4: Split Bregman algorithm for ℓ_1 regularization with intermediate updating of the Bregman variable. Note the additional computation between the u subproblem and the w subproblem. Though it is applied twice, Wu^{k+1} only needs to be computed once.

3.5 Extensions to multicoil data

The algorithms described in the previous section are best suited for single coil data or data that has been retrospectively undersampled after the coils have been combined. They can be applied to multicoil data by reconstructing each coil separately and combining the coils at the end as a sum-of-squares as was done in [61, 178]. While this method is straightforward and simple, there are numerous reasons to avoid doing it this way.

From a signal perspective, Roemer *et. al.* showed that images reconstructed




Figure 3.5: Effects of coil combination on NAA metabolite map **(a)** and COSY spectrum of a central voxel **(b)**. In each case, the coil combination using complex sensitivity maps is shown on the left while sum-of-squares combination is shown on the right.

using coil sensitivities (or B_1 field maps) have slightly higher SNR than the simple sum-of-squares [141]. This is illustrated with 4D EP-COSI data in Figure 3.5. In addition, complex phase information is utilized for fitting of spectra in both 1D [129, 130] and 2D [153, 59], but phase is lost in a sum-of-squares combination.

From an algorithmic perspective, the separate coil method is computationally inefficient unless the coils can be reconstructed with parallel processors. While multicoil reconstructions do take longer per iteration (how much depends on the specific algorithm), for optimized algorithms, the factor at which each iteration scales is usually notably less than the number of coils. Also, reconstructing the coils individually does not utilize any of the information between coils. Perhaps even more important is that the choice of regularizers are usually based on the final image and may not be ideal for the individual coil images. As an example,

minimizing TV will only properly reconstruct piecewise constant images. Though this is appropriate for many MRI/MRSI scans, the individual coil images have an approximately polynomial profile superimposed, and piecewise constancy is violated.

To handle multicoil data, the coil sensitivity variations are taken into account. While this can be done using k -space-based methods, the image-based SENSE method [132, 131] is a much more natural choice due to its similarities to the present data consistency constraint. SENSE using Tikhonov regularization even predates compressed sensing [97]. The data consistency constraint in regularized SENSE is expressed through the following ℓ_2 norm

$$\|R\mathcal{F}Su - f\|_2$$

where S is coil sensitivity operator.

Now the discussion up until this point has described operators R , \mathcal{F} , and ∇ that act on the sampled data f and reconstructed data u without specifying their numerical structures. This has not really mattered because none of the operators changes the size of the data they operate on, and f and u are the same size. However, in the upcoming case, f contains samples from all the coils while u is a single image. So in this next section, we will take a brief step back from the split Bregman algorithm to go through the numeric structures of these operators.

3.5.1 Constructing matrix operators

For single coil imaging, u and f are the same 4D ($n_x \times n_y \times n_2 \times n_1$) or 5D ($n_x \times n_y \times n_z \times n_2 \times n_1$) matrices with (n_x, n_y, n_z) representing the number of points in the spatial dimensions, and (n_2, n_1) representing the number of spectral or temporal points. To apply the operator notation, u and f must be vectorized into $N \times 1$ column vectors using standard linear indexing, where N is the total number of points $n_x \cdot n_y \cdot n_z \cdot n_2 \cdot n_1$.

In practice, the following operators are not always written as explicit matrix multiplications, and the data is not always vectorized. This is due to the speed with which functional forms can be computed (e.g. FFT in Matlab) and the very large, memory hogging matrices required for multiplication. (Using “sparse” matrices does alleviate much of the memory requirements.) Nevertheless, for proper derivations, the effects of these operators acting sequentially must be determined, and this effect cannot be determined if the operator is only understood as a function acting on the data.

3.5.1.1 Sampling matrix

The sampling matrix R is easily constructed as an $N \times N$ diagonal matrix where the i th diagonal element is a 1 or 0 depending whether the i th data point was sampled or not. For example the following matrix indicates that the first, second, and third points were sampled, while the fourth and last were not.

$$R = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 \end{pmatrix}$$

A few properties of the sampling matrix are worth mentioning. First, since R is diagonal and logical, it can be thought of as a type of projection from the entire data space to the sampled data space only, and $R^2 = R$. It is also Hermitian with $R' = R$, so the product $R'R$ that appears in the u subproblem simply equals R itself. This simplification has not been written in the algorithms here for clarity but was implemented in the Matlab code for efficiency.

3.5.1.2 nD Fourier matrix

Next, the Fourier transform operator \mathcal{F} is detailed. The $N \times N$ 1D Fourier matrix is the Vandermonde matrix of the N^{th} complex roots of unity.

$$\mathcal{F}_{1D,N} = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & w & w^2 & w^3 & \cdots & w^{N-1} \\ 1 & w^2 & w^4 & w^6 & \cdots & w^{2(N-1)} \\ 1 & w^3 & w^6 & w^9 & \cdots & w^{3(N-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & w^{N-1} & w^{2(N-1)} & w^{3(N-1)} & \cdots & w^{(N-1)(N-1)} \end{pmatrix}$$

where $w = e^{-\frac{2\pi i}{N}}$. After application to an $N \times 1$ vector u , $F_{1D}u$ is the Fourier transform of u assuming the points of u all lie in the same dimension.

However, u is not one dimensional but 4 or 5 dimensional. Therefore, we need a 4 or 5D Fourier transform operator that respects the original dimensionality of u . A general multidimensional Fourier operator can be constructed as a series of Kronecker products of the 1D operators $\mathcal{F}_{1D,n}$

$$\mathcal{F}_{5D,N} = \mathcal{F}_{1D,n_x} \otimes \mathcal{F}_{1D,n_y} \otimes \mathcal{F}_{1D,n_z} \otimes \mathcal{F}_{1D,n_2} \otimes \mathcal{F}_{1D,n_1}$$

Note that the order of the products must follow the 5D order of the dimensions which is why the n_2 term is before the n_1 and that the dimensions of $\mathcal{F}_{5D,N}$ are the product of the dimensions of each of the 1D operators as explained below.

The Kronecker product between matrices A and B results in a block diagonal matrix constructed from the product of each element of A with the entire matrix

B.

$$A \otimes B = \begin{pmatrix} a_{11}b_{11} & a_{11}b_{12} & \cdots & a_{11}b_{1q} & \cdots & \cdots & a_{1n}b_{11} & a_{1n}b_{12} & \cdots & a_{1n}b_{1q} \\ a_{11}b_{21} & a_{11}b_{22} & \cdots & a_{11}b_{2q} & \cdots & \cdots & a_{1n}b_{21} & a_{1n}b_{22} & \cdots & a_{1n}b_{2q} \\ \vdots & \vdots & \ddots & \vdots & & & \vdots & \vdots & \ddots & \vdots \\ a_{11}b_{p1} & a_{11}b_{p2} & \cdots & a_{11}b_{pq} & \cdots & \cdots & a_{1n}b_{p1} & a_{1n}b_{p2} & \cdots & a_{1n}b_{pq} \\ \vdots & \vdots & & \vdots & \ddots & & \vdots & \vdots & & \vdots \\ \vdots & \vdots & & \vdots & & \ddots & \vdots & \vdots & & \vdots \\ a_{m1}b_{11} & a_{m1}b_{12} & \cdots & a_{m1}b_{1q} & \cdots & \cdots & a_{mn}b_{11} & a_{mn}b_{12} & \cdots & a_{mn}b_{1q} \\ a_{m1}b_{21} & a_{m1}b_{22} & \cdots & a_{m1}b_{2q} & \cdots & \cdots & a_{mn}b_{21} & a_{mn}b_{22} & \cdots & a_{mn}b_{2q} \\ \vdots & \vdots & \ddots & \vdots & & & \vdots & \vdots & \ddots & \vdots \\ a_{m1}b_{p1} & a_{m1}b_{p2} & \cdots & a_{m1}b_{pq} & \cdots & \cdots & a_{mn}b_{p1} & a_{mn}b_{p2} & \cdots & a_{mn}b_{pq} \end{pmatrix}$$

So a 2D Fourier matrix would look like

$$\mathcal{F}_{2D} = \begin{pmatrix} 1 & 1 & \cdots & 1 & 1 & 1 & \cdots & \cdots & 1 \\ 1 & w & \cdots & w^{n-1} & 1 & w & \cdots & \cdots & w^{n-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & & & \vdots \\ 1 & w^{n-1} & \cdots & w^{(n-1)(n-1)} & 1 & w^{n-1} & \cdots & \cdots & w^{(n-1)(n-1)} \\ 1 & 1 & \cdots & 1 & w & w & \cdots & \cdots & w^{n-1} \\ 1 & w & \cdots & w^{n-1} & w & w^2 & \cdots & \cdots & w^{2(n-1)} \\ \vdots & \vdots & & \vdots & \vdots & \vdots & \ddots & & \vdots \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \ddots & \vdots \\ 1 & w^{n-1} & \cdots & w^{(n-1)(n-1)} & w^{n-1} & w^{2(n-1)} & \cdots & \cdots & w^{2(n-1)(n-1)} \end{pmatrix}$$

Sometimes it is desirable to Fourier transform only over specific dimensions. In that case, the operator is constructed using the identity matrix in the Kronecker product. For example, in the split Bregman algorithm, it is only necessary to Fourier transform over the undersampled dimensions during each iteration

$$\mathcal{F}_{5D,us} = I_{n_x} \otimes \mathcal{F}_{1D,n_y} \otimes \mathcal{F}_{1D,n_z} \otimes I_{n_2} \otimes \mathcal{F}_{1D,n_1}$$

where I_n is the $n \times n$ identity matrix.

In each case, the inverse Fourier transform is simply given by the Hermitian transpose \mathcal{F}' .

3.5.1.3 Sensitivity operator matrix

Before the numeric structure of S is discussed, we must first understand intuitively what the operator does. In brief, S is an operator that when applied to the single image u results in multiple images of u weighted by the various coil sensitivity profiles. Therefore, if u is $n_x \times n_y \times n_z \times n_2 \times n_1$, then Su is $n_x \times n_y \times n_z \times n_2 \times n_1 \times n_c$, where n_c is the number of coils. Su is also the same size as f since we know that R and \mathcal{F} operators do not change the data size.

As before, consider u being vectorized into an $N \times 1$ column matrix, then f 's vectorized size must be $Nn_c \times 1$. Sensitivity maps are logically acquired as $n_x \times n_y \times n_z \times n_c$ matrices, so they must be expanded to $n_x \times n_y \times n_z \times n_2 \times n_1 \times n_c$ to include the temporal domains. Each $n_x \times n_y \times n_z \times n_2 \times n_1$ sensitivity matrix is then reshaped to an $N \times N$ diagonal matrix with the complex sensitivities along the diagonal.

$$S_c = \begin{pmatrix} s_{(1,1,1,1,1,c)} & 0 & 0 & \cdots & 0 \\ 0 & s_{(2,1,1,1,1,c)} & 0 & \cdots & 0 \\ 0 & 0 & s_{(3,1,1,1,1,c)} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & s_{(n_x, n_y, n_z, n_2, n_1, c)} \end{pmatrix}$$

The full matrix S is formed by stacking the individual S_c matrices

$$S = \begin{pmatrix} S_1 \\ S_2 \\ \vdots \\ S_{n_c} \end{pmatrix}$$

and is therefore $Nn_c \times N$.

The adjoint operator S' is $N \times Nn_c$ and when applied has the effect of multiplying each coil image by the conjugate of the corresponding coil sensitivity and adding up over all the coils. Thus, the product $S'S$ is a diagonal $N \times N$ matrix with each element on the diagonal being equal to the sum over the coils of the square magnitude of each point.

$$S'S = \begin{pmatrix} \sum_c s^* s_{(1,1,1,1,1,c)} & 0 & 0 & \cdots & 0 \\ 0 & \sum_c s^* s_{(2,1,1,1,1,c)} & 0 & \cdots & 0 \\ 0 & 0 & \sum_c s^* s_{(3,1,1,1,1,c)} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \sum_c s^* s_{(n_x, n_y, n_z, n_2, n_1, c)} \end{pmatrix}$$

3.5.1.4 Gradient operator as a finite differences matrix

For either the anisotropic or isotropic TV case, the variation in the data is found by approximating the effect of the gradient operator ∇ on the data. Typically, the effect of ∇ is approximated by taking first order finite differences between two consecutive data points in different directions (though more than two points can be used). In particular, ∇_{row} takes differences between consecutive rows, while ∇_{col} takes differences between consecutive columns of a matrix. Both use backwards differences here. In matrix form, we have the $N \times N$

$$\nabla_{row, N} = \begin{pmatrix} 1 & 0 & 0 & \cdots & -1 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

to take row differences on a vector or another matrix when *left* multiplied. The -1 in the first row satisfies the boundary conditions and is required for reasons that will soon become clear.

Since matrix multiplication is a row-by-column operation, there is no way to

take column differences (i.e. y gradient) of a matrix through left multiplication with another matrix. Instead column differences can be taken by *right* multiplication only. However, we are not really interested in taking column differences of a matrix, as the image data u is already vectorized.

Instead, the gradient operator must take a directional derivative of the vectorized data. This can be accomplished using the Kronecker product as was done in constructing the n D Fourier operator. In keeping with that notation,

$$\begin{aligned}\nabla_x &= I_{n_x} \otimes \nabla_{row,n_y} \otimes I_{n_z} \otimes I_{n_2} \otimes I_{n_1} \\ \nabla_y &= \nabla_{row,n_x} \otimes I_{n_y} \otimes I_{n_z} \otimes I_{n_2} \otimes I_{n_1} \\ \nabla_z &= I_{n_x} \otimes I_{n_y} \otimes \nabla_{row,n_z} \otimes I_{n_2} \otimes I_{n_1}\end{aligned}$$

where each of the gradient operators are $N \times N$. For each of these matrices, the product of the adjoint and itself is block circulant with circulant blocks (BCCB) and therefore, diagonalizable by the Fourier transform (i.e. $\mathcal{F}'_{5D} \nabla'_{x,y,z} \nabla_{x,y,z} \mathcal{F}_{5D}$ is diagonal) for the boundary conditions imposed on ∇_{row} .

3.5.2 SBTV SENSE

With a better understanding of the relevant operators, we are now ready to return to the problem of using split Bregman to solve the regularized, multicoil SENSE problem. First we consider the SENSE extension to the isotropic TV -regularized problem in (3.22). For brevity, only this isotropic version will be presented, though an anisotropic one could be constructed similarly. The problem can be formulated as

$$\min_u TV_{iso}(u) \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (3.28)$$

which highly resembles (3.22) with only the addition of the S operator that carries the coil sensitivity information.

The auxiliary variables $q = u$, $v = Su$, and $d_{x,y,z} = \nabla_{x,y,z}q$ are introduced, and

the constrained problem (3.28) is relaxed into the augmented constrained problem

$$\begin{aligned} \min_{u, d_x, d_y, d_z, q, v} \quad & \|(d_x, d_y, d_z)\|_2 + \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 \quad \text{s.t. } q = u, v = Su \\ & d_x = \nabla_x q, d_y = \nabla_y q, d_z = \nabla_z q \end{aligned} \quad (3.29)$$

The problem in (3.29) is again relaxed with the addition of the Bregman variables b_x, b_y, b_z, b_v , and b_q , with the resulting unconstrained problem decoupled into

$$v^{k+1} = \min_v \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 + \frac{\lambda_v}{2} \|v - Su - b_v\|_2^2 \quad (3.30a)$$

$$(d_x^{k+1}, d_y^{k+1}, d_z^{k+1}) = \min_{d_x, d_y, d_z} \|(d_x, d_y, d_z)\|_2^2 + \sum_{i=x,y,z} \frac{\lambda_d}{2} \|d_i - \nabla_i q - b_i\|_2^2 \quad (3.30b)$$

$$u^{k+1} = \min_u \frac{\lambda_q}{2} \|v - Su - b_v\|_2^2 + \frac{\lambda_q}{2} \|q - u - b_q\|_2^2 \quad (3.30c)$$

$$q^{k+1} = \min_q \frac{\lambda_q}{2} \|q - u - b_q\|_2^2 + \sum_{i=x,y,z} \frac{\lambda_d}{2} \|d_i - \nabla_i q - b_i\|_2^2 \quad (3.30d)$$

The d subproblem (3.30b) is the same as in the single coil isotropic TV case (3.23) and is therefore also solved with *shrink3*.

The v subproblem (3.30a) is differentiable, leading to the following solution after setting the derivative to zero and grouping

$$v = \mathcal{F}'K^{-1}\mathcal{F}(\mu\mathcal{F}'R'f + \lambda_v(Su + b_v)) \quad (3.31)$$

with $K = \mu R'R + \lambda_v$ being diagonal.

The u subproblem (3.30c) is also a sum of ℓ_2 -norms and is therefore differentiable, with solution given by

$$u = L^{-1}(\lambda_v S'(v - b_v) + \lambda_q(q - b_q)) \quad (3.32)$$

where L is diagonal and given by $\lambda_v S'S + \lambda_q$.

Lastly, the q subproblem (3.30d) is also differentiable, and its solution is

$$q = \mathcal{F}'M^{-1}\mathcal{F}\left(\lambda_q(u + b_q) + \sum_{i=x,y,z} \lambda_d \nabla'_i(d_i - b_i)\right) \quad (3.33)$$

with $M = \lambda_q - \lambda_d \mathcal{F}' \Delta \mathcal{F}$.

Note that like the previous algorithms, the differentiable subproblems are quickly solvable because the required inversion can be performed on a diagonal matrix. Again, this particular characteristic is why this family of algorithms are so fast and inexpensive.

However, there is a small difference from previous problems in the u subproblem (3.32). There, the matrix L that multiplies u was diagonal itself, since it is composed of a combination of the identity matrix and the matrix $S'S$. By contrast, the matrices that multiply v and q in (3.31) and (3.33) are BCCB and only diagonalized following right and left multiplication of the Fourier transform and its inverse. In particular, the matrix multiplying v is BCCB because of the $\mathcal{F}'R'R\mathcal{F}$ term, while the matrix multiplying q is BCCB because of the Laplacian term Δ . All three subproblems have a matrix term proportional to the identity (which is the only matrix that is both diagonal and circulant). That $S'S$ is already diagonal is the reason for the seemingly excessive auxiliary variable q .

The entire algorithm is presented below without including intermediate updating of the Bregman variables for clarity.

```

Precompute  $K^{-1} = (\mu R'R + \lambda_v)^{-1}, L^{-1} = (\lambda_v S'S + \lambda_q)^{-1}, M^{-1} =$ 
 $(\lambda_q - \lambda_d \mathcal{F}' \Delta \mathcal{F})^{-1}$ 
repeat
  for  $k = 1$  to  $N$  do
    Compute  $u^{k+1}$  using (3.32).
    Compute  $v^{k+1}$  using (3.31).
    Compute  $q^{k+1}$  using (3.33).
     $(d_x^{k+1}, d_y^{k+1}, d_z^{k+1}) = \text{shrink3}(\nabla_x q^{k+1} + b_x, \nabla_y q^{k+1} + b_y, \nabla_z q^{k+1} + b_z, 1/\lambda_d)$ 
     $b_q \leftarrow b_q + u^{k+1} - q^{k+1}$ 
     $b_x \leftarrow b_x + \nabla_x q^{k+1} - d_x^{k+1}$ 
     $b_y \leftarrow b_y + \nabla_y q^{k+1} - d_y^{k+1}$ 
     $b_z \leftarrow b_z + \nabla_z q^{k+1} - d_z^{k+1}$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + f - R\mathcal{F}Su^{k+1}$ 
until  $\|R\mathcal{F}Su^{k+1} - f\|_2^2 < \sigma^2$ 

```

Figure 3.6: Split Bregman algorithm for isotropic TV regularization from multicoil acquisition

3.5.3 SB ℓ_1 SENSE

The ℓ_1 -regularized version of the SENSE problem can be formulated as

$$\min_u |Wu|_1 \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (3.34)$$

The substitutions $v = Su$ and $q = Wu$ are applied, resulting in the constrained augmented problem

$$\min_{u,q,v} |q|_1 + \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 \quad \text{s.t.} \quad q = Wu, v = Su \quad (3.35)$$

Bregman variables b_q and b_v are included, and (3.35) is relaxed and separated into the following subproblems

$$v^{k+1} = \min_v \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 + \frac{\lambda_v}{2} \|v - Su - b_v\|_2^2 \quad (3.36a)$$

$$q^{k+1} = \min_q |q|_1 + \frac{\lambda_q}{2} \|q - Wu - b_q\|_2^2 \quad (3.36b)$$

$$u^{k+1} = \min_u \frac{\lambda_v}{2} \|v - Su - b_v\|_2^2 + \frac{\lambda_q}{2} \|q - Wu - b_q\|_2^2 \quad (3.36c)$$

The v subproblem in (3.36a) is the same as (3.30a) in the previous TV problem and is also solved by (3.31). The q subproblem (3.36b) is solved by *shrink*, and the u subproblem (3.36c) has solution given by (3.32) provided that $W'W$ is equal to the identity.

Precompute $K^{-1} = (\mu R'R + \lambda_v)^{-1}$, $L^{-1} = (\lambda_v S'S + \lambda_q)^{-1}$, $M^{-1} = (\lambda_q - \lambda_d \mathcal{F}' \Delta \mathcal{F})^{-1}$

repeat

for $k = 1$ to N **do**

Compute u^{k+1} using (3.32).

$b_v \leftarrow b_v + Su^{k+1} - v^k$

$b_q \leftarrow b_q + Wu^{k+1} - q^k$

Compute v^{k+1} using (3.31).

$q = \textit{shrink}(Wu + b_q, 1/\lambda_q)$

$b_v \leftarrow b_v + Su^{k+1} - v^{k+1}$

$b_q \leftarrow b_q + Wu^{k+1} - q^{k+1}$

end for

$i \leftarrow i + 1$

$f^{i+1} = f^i + f - R\mathcal{F}Su^{k+1}$

until $\|R\mathcal{F}Su^{k+1} - f\|_2^2 < \sigma^2$

Figure 3.7: Split Bregman algorithm for ℓ_1 -norm regularization from multicoil acquisition

3.5.4 Self sparse SB ℓ_1 SENSE

Alternatively, if the reconstructed data u is in fact self sparse (which it is assumed to be in this dissertation), we recognize that the individual coil images Su must be at least as sparse. In this case, the problem can be formulated in a way with a particularly simple solution.

$$\min_u |Su|_1 \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (3.37)$$

Applying the variable substitution $v = Su$ results in the problem

$$\min_u |v|_1 \quad \text{s.t.} \quad \|R\mathcal{F}v - f\|_2^2 < \sigma^2, \quad v = Su \quad (3.38)$$

Note that the substitution has completely decoupled u from the rest of the optimization, and therefore, the problem can be solved first for v before final transformation to u . Solving for v is identical to solving for u in the single coil problem (3.25) if W is equal to the identity.

Once v has been determined, u is simply computed as

$$u = (S'S)^{-1} S'v \quad (3.39)$$

where the matrix inversion is again simple since $S'S$ is diagonal. For optimal SNR, the noise correlation between individual coils Ψ can be included in the inversion

$$u = (S'\Psi^{-1}S)^{-1} S'\Psi^{-1}v$$

However, in many SENSE applications with modern coil geometries, it is assumed that Ψ is the identity.

```

Precompute  $K^{-1} = (\mu R'R + \lambda)^{-1}$ 
repeat
  for  $k = 1$  to  $N$  do
    Compute  $v^{k+1}$  using (3.27).
     $b \leftarrow b + v^{k+1} - w^k$ 
     $w^{k+1} = \text{shrink}(v^{k+1} + b, 1/\lambda)$ 
     $b \leftarrow b + v^{k+1} - w^{k+1}$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + f - R\mathcal{F}v^{k+1}$ 
until  $\|R\mathcal{F}v^{k+1} - f\|_2^2 < \sigma^2$ 
 $u = (S'S)^{-1} S'v^{k+1}$  or  $u = (S'\Psi^{-1}S)^{-1} S'\Psi^{-1}v^{k+1}$ 

```

Figure 3.8: Simple split Bregman algorithm for ℓ_1 regularized self sparse multicoil acquisition with intermediate variable updating

3.6 Group Sparsity

The group sparsity (GS) regularizer replaces the ℓ_1 -norm with an $\ell_{2,1}$ -norm [181]. The $\ell_{2,1}$ -norm is found by calculating the ℓ_1 -norm of the ℓ_2 -norms of grouped coefficients.

$$\|u\|_{2,1} = \|u_{g_1}\|_2 + \|u_{g_2}\|_2 + \cdots + \|u_{g_n}\|_2 \quad (3.40)$$

where the g_i 's are the n groups of u . By grouping certain coefficients before promoting sparsity through the ℓ_1 -norm, GS can exploit the tendency of large coefficients to be clustered and increases the structural sparsity of the groups.

There is a lot of flexibility in how the groupings can be defined. They can be across a single dimension as in the case of joint sparsity taken across the coils in MRI reconstruction [51], or they can span multiple dimensions. Groups can

partition the data space, or they can overlap with each other.

3.6.1 Single channel

The unweighted, single channel GS problem is formulated as

$$\min_u \|Gu\|_{2,1} \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (3.41)$$

where $G \in \mathbb{R}^{n \times N}$ is the matrix that groups coefficients of u . This problem was solved in [46] using the alternate direction method of multipliers (ADMM) and similarly, in [27] using the split Bregman algorithm [66].

The solution is found by first writing the problem as an unconstrained minimization,

$$\min_u \|Gu\|_{2,1} + \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 \quad (3.42)$$

defining an auxiliary variable $z = Gu$, splitting the unconstrained problem, and using Bregman iterations to enforce the equality of the substitution. The unconstrained problem can then be solved iteratively as follows

$$\begin{aligned} u^{k+1} &= \min_u \frac{\mu}{2} \|R\mathcal{F}u - f\|_2^2 + \frac{\lambda}{2} \|z^k - Gu - b_z^k\|_2^2 \\ z^{k+1} &= \min_z \|z\|_{2,1} + \frac{\lambda}{2} \|z - Gu^{k+1} - b_z^k\|_2^2 \\ b_z^{k+1} &= b_z^k + (Gu^{k+1} - z^{k+1}) \end{aligned} \quad (3.43)$$

which is Eq. [12] in [27]. Here, μ and λ are regularization parameters, and b_z is the Bregman parameter enforcing the equality in the variable substitution. The original constrained problem (Eq. (5.3)) is solved by applying an outer loop of Bregman updates on the sampled data f until the ℓ_2 -norm of the data consistency term falls below σ . The b_z update is simply algebra, and the z update is solved using the group-wise shrinkage function

$$gshrink(Gu, \lambda) = \max\left(0, 1 - \frac{\lambda}{\|Gu\|_{2,1}}\right) \cdot Gu \quad (3.44)$$

The u subproblem is differentiable and can be solved very quickly since $\mu\mathcal{F}'R'R\mathcal{F} + \lambda G'G$ is circulant as long as all points are in the same number of groups. $G'G = nI$

is simply the number of groups each point is in n multiplied by the identity matrix I . (Here, and for the rest of the paper, primes denote the Hermitian conjugate.)

If prior knowledge of the coefficients in each group is known, it can be incorporated into the $\ell_{2,1}$ -norm by assigning each group a weight w_i as illustrated by

$$\|Gu\|_{w,2,1} = \sum_i^n w_i \|u_{g_i}\|_2 \quad (3.45)$$

Equation 3.45 is referred to as the $\ell_{w,2,1}$ -norm and slightly modifies the objective of the z subproblem, which must then be solved using the weighted group-wise shrinkage function

$$gshrink_w(Gu, \lambda, w) = \max \left(0, 1 - w \cdot \frac{\lambda}{\|Gu\|_{2,1}} \right) \cdot Gu \quad (3.46)$$

The remaining iterations are unchanged.

3.6.2 Multi channel with sensitivity maps

The method described above applies to single channel data or can be used to reconstruct multi channel data channel-by-channel. Alternatively, if coil sensitivity information is available, the multi channel problem can be formalized as

$$\min_u \|Gu\|_{w,2,1} \quad \text{s.t.} \quad \|RFSu - f\|_2^2 < \sigma^2 \quad (3.47)$$

where S is the sensitivity profile operator [132]. While Eq. (3.47) resembles to Eq. (5.3), the presence of S complicates the matter.

The split Bregman algorithm achieves such high efficiency in part because the u subproblem (which is the most expensive step) involves inverting a circulant matrix that can be made diagonal with application of a Fourier Transform. In general, fast solvers based on ADMM or the split Bregman algorithm share this characteristic. But once the sensitivity profiles are included, $\mu S' \mathcal{F}' R' RFS + \lambda G' G$ is no longer circulant. In that case, the u subproblem can be approximately solved

using a few iterations of the conjugate gradient method [66, 39], but this can still be expensive even with preconditioning.

Alternatively, variable splitting can be applied to the data fidelity term in the same way it is applied to the objective term [136, 37]. This has the advantage of separating the image space operator S from the k-space operators R and \mathcal{F} and allows for exact solution of each subproblem. With the addition of the auxiliary variables, the problem in Eq. (3.47) can be written as

$$\min_{u,v,z} \|z\|_{w,2,1} \quad \text{s.t.} \quad \|R\mathcal{F}v - f\|_2^2 < \sigma^2, \quad z = Gu, \quad v = Su \quad (3.48)$$

As before, the constraints are relaxed, and the Bregman formulation is applied, resulting in an objective function of

$$\min_{u,v,z} \|z\|_{w,2,1} + \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 + \frac{\lambda_z}{2} \|z - Gu - b_z\|_2^2 + \frac{\lambda_v}{2} \|v - Su - b_v\|_2^2 \quad (3.49)$$

where μ , λ_z , and λ_v are regularization parameters and b_z and b_v are Bregman variables.

The objective function is then split into minimization subproblems for each variable. The v subproblem is given by

$$\min_v \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 + \frac{\lambda_v}{2} \|v - Su - b_v\|_2^2 \quad (3.50)$$

and resembles the u subproblem in Eq. (3.43). Differentiating with respect to v (holding u constant) and setting the result equal to zero yields

$$(\mu\mathcal{F}'R'R\mathcal{F} + \lambda_v)v = \mu\mathcal{F}'R'f + \lambda_v(Su + b_v) \quad (3.51)$$

The term in parenthesis on the left is circulant, so the solution is given by

$$v = \mathcal{F}'K^{-1}\mathcal{F}(\mu\mathcal{F}'R'f + \lambda_v(Su + b_v)) \quad (3.52)$$

where $K = (\mu R'R + \lambda_v)$ is diagonal.

Next, the z subproblem is given by

$$\min_z \|z\|_{w,2,1} + \frac{\lambda_z}{2} \|z - Gu - b_z\|_2^2 \quad (3.53)$$

which is solved by Eq. (3.46) with

$$z = \mathit{ghrink}_w(Gu + b_z, 1/\lambda_z, w) \quad (3.54)$$

as before.

The u subproblem is written as

$$\min_u \frac{\lambda_v}{2} \|v - Su - b_v\|_2^2 + \frac{\lambda_z}{2} \|z - Gu - b_z\|_2^2 \quad (3.55)$$

which is also differentiable with solution given by

$$u = M^{-1} (\lambda_v S' (v - b_v) + \lambda_z G' (z - b_z)) \quad (3.56)$$

where $M = (\lambda_v S' S + \lambda_z G' G)$. Both $S' S$ and $G' G$ are diagonal, so M is easily inverted.

The remainder of the algorithm consists of Bregman updates to the parameters b_v , b_z , and b_u in each inner loop and Bregman updating of the sampled data f in an outer loop to enforce the original constrained problem in Eq. (3.47). We use intermediate updating of the Bregman parameters to help speed up convergence at minimal added computational cost [6].

```

Precompute  $K^{-1} = (\mu R'R + \lambda_v)^{-1}$  and  $M^{-1} = (\lambda_v S'S + \lambda_z G'G)^{-1}$ 
repeat
  for  $k = 1$  to 15 do
    Compute  $v^{k+1}$  using (3.52).
     $b_v \leftarrow b_v + (Su^k - v^{k+1})$ 
    Compute  $z^{k+1}$  using (3.54).
     $b_z \leftarrow b_z + (Gu^k - z^{k+1})$ 
    Compute  $u^{k+1}$  using (3.56).
     $b_v \leftarrow b_v + (Su^{k+1} - v^{k+1})$ 
     $b_z \leftarrow b_z + (Gu^{k+1} - z^{k+1})$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + (f - R\mathcal{F}Su^{k+1})$ 
until  $\|R\mathcal{F}Su^{k+1} - f\|_2^2 < \sigma^2$ 

```

Figure 3.9: SBGS SENSE I algorithm

Alternatively, if the final image/spectral data u is in fact self sparse (which is the premise for using the $\ell_{w,2,1}$ -norm of groups of the reconstructed data itself with no additional transform), we recognize that the individual coil image/spectral data Su must be at least as sparse. The multi channel problem can then be formulated as

$$\min_u \|GSu\|_{w,2,1} \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (3.57)$$

This problem is actually easier to solve than Eq. (3.47) as it requires one less auxiliary variable to be iteratively updated. This can be seen by applying the substitution $v = Su$ to get

$$\min_{u,v} \|Gv\|_{w,2,1} \quad \text{s.t.} \quad \|R\mathcal{F}v - f\|_2^2 < \sigma^2, \quad v = Su \quad (3.58)$$

Since u has been completely decoupled from the rest of the optimization, the

problem can be solved iteratively for v first before transforming to u .

Ignoring the $v = Su$ constraint, the remaining problem is simply the problem for single channel data (Eq. (5.3)) with v replacing u . It is solved as before by relaxing the constraint, making the substitution $z = Gv$, using the Bregman formulation, and iteratively solving each subproblem while holding the other variable constant. In this case, the v subproblem is

$$\min_v \frac{\mu}{2} \|R\mathcal{F}v - f\|_2^2 + \frac{\lambda}{2} \|z^k - Gv - b_z^k\|_2^2 \quad (3.59)$$

which is differentiable with solution

$$v = \mathcal{F}'L^{-1}\mathcal{F}(\mu\mathcal{F}'R'f + \lambda G'(z - b_z)) \quad (3.60)$$

where $L = \mu R'R + n\lambda$ is diagonal as already discussed. The z subproblem is therefore

$$\min_z \|z\|_{w,2,1} + \frac{\lambda}{2} \|z - Gv^{k+1} - b_z^k\|_2^2 \quad (3.61)$$

and is again solved by group-wise shrinkage of Eq. (3.46).

Bregman updates are applied, including over an outer loop on f until the problem is accurately solved for v . At that point u is given from v by

$$u = (S'\Psi^{-1}S)^{-1} S'\Psi^{-1}v \quad (3.62)$$

where Ψ is the noise correlation matrix between the coils and is used to achieve optimal SNR [132].

```

Determine  $\Psi$  or set  $\Psi = I$ .
Precompute  $L^{-1} = (\mu R'R + n\lambda)^{-1}$ 
repeat
  for  $k = 1$  to 15 do
    Compute  $v^{k+1}$  using (3.60).
     $b_z \leftarrow b_z + (Gv^{k+1} - z^k)$ 
    Compute  $z^{k+1}$  using (3.54).
     $b_z \leftarrow b_z + (Gv^{k+1} - z^{k+1})$ 
  end for
   $i \leftarrow i + 1$ 
   $f^{i+1} = f^i + (f - R\mathcal{F}Su^{k+1})$ 
until  $\|R\mathcal{F}v^{k+1} - f\|_2^2 < \sigma^2$ 
Compute  $u$  using (5.9).

```

Figure 3.10: SBGS SENSE II algorithm

CHAPTER 4

Accelerated Five-dimensional Echo Planar J-resolved Spectroscopic Imaging: Implementation and Pilot Validation in Human Brain

4.1 Abstract

Purpose: To implement an accelerated five dimensional (5D) echo-planar J -resolved spectroscopic imaging (EP-JRESI) sequence combining 3 spatial and 2 spectral encoding dimensions and to apply the sequence in human brain.

Methods: An echo planar readout was used to acquire a single spatial and a single spectral dimension during one readout. Nonuniform sampling (NUS) was applied to the two phase encoded spatial directions and the indirect spectral dimension. Nonlinear reconstruction was used to minimize the ℓ_1 -norm or the total variation (TV) and included a spectral mask to enhance sparsity. Retrospective reconstructions at multiple undersamplings were performed in phantom. Ten healthy volunteers were scanned with 8x undersampling and compared to a fully sampled single slice scan.

Results: Retrospective reconstruction of fully sampled phantom data showed excellent quality at 4x, 8x, 12x, and 16x undersampling using either reconstruction method. Reconstruction of prospectively acquired in vivo scans with 8x undersampling showed excellent quality in the occipito-parietal lobes and good quality

in the frontal lobe, consistent with the fully sampled single slice scan.

Conclusion: By utilizing NUS with nonlinear reconstruction, 2D J -resolved spectra can be acquired over a 3D spatial volume with a total scan time of 20 minutes, which is reasonable for in vivo studies.

4.2 Introduction

Spectroscopic imaging (SI) allows the acquisition of spatially mapped spectra over two or three dimensions [24]. Compared to single voxel techniques, SI offers much larger spatial coverage and better resolution as well as the ability to discriminate spectra at different locations. SI can be an important tool for mapping the spatial extent of diseased tissue as underlying metabolic changes can occur before any visible anatomic variation.

The scan time for classical SI sequences is proportional to the number of phase encoding increments and therefore increases with the resolution and number of dimensions. Mansfield proposed using an echo planar readout to interleave the collection of one of the spatial dimensions along with the directly-acquired spectral dimension [107]. This type of spatial-spectral readout has been applied in vivo with the originally-proposed Cartesian k-space trajectory [127] as well as with spiral-based [2] and concentric circular-based trajectories [62].

A major issue with proton (^1H) spectroscopy in the human brain is that nearly all of the measurable resonances lie in a relatively small bandwidth covering approximately 4 ppm, resulting in spectra that suffer from severe overlap and crowding. Spectral dispersion can be increased using two dimensional (2D) MR Spectroscopy (MRS) [9]. Localized versions of some of the simpler 2D MRS sequences have been applied in single voxel techniques in vivo [147, 171, 8]. In particular, localized J -resolved spectroscopy (JPRESS) [147] is achieved by adding incremental delays symmetrically around the last refocusing pulse in a PRESS [20] localiza-

tion. At each increment, the chemical shift is refocused at readout. However, scalar coupling is not refocused by a 180° pulse and therefore, evolves differently for each increment. Two dimensional Fourier transformation of the time domain data results in a spectrum in which peak splitting of coupled metabolites is in a separate dimension from the chemical shift.

Multidimensional spectroscopic imaging refers to SI scans that collect two spectral dimensions, such as with J-resolved [3] or correlated spectroscopy [99, 7]. The additional information present in the second spectral dimension comes at the expense of greatly increased scan time, limiting the application of multidimensional SI to relatively low resolution single slice acquisitions even with an echo planar readout.

Techniques used to accelerate MRI such as partial Fourier acquisition [104] or parallel imaging [132, 70] have been applied to acquire fewer phase encoding steps in spectroscopic imaging as well [113, 52, 98, 121]. However, the potential acceleration is limited for low resolution SI, and these methods do not accelerate the acquisition of the indirect spectral dimension.

Alternatively, nonuniform sampling (NUS) with compressed sensing (CS) reconstruction offers greater potential acceleration as the undersampling can be spread over multiple dimensions [102]. Hu *et al.* used a blipped phase encode during readout to nonuniformly sample the spatial and spectral dimensions in hyperpolarized C-13 spectroscopic imaging [81]. Our group has previously shown the potential for 4-fold acceleration of single-slice echo planar J-resolved spectroscopic imaging (EP-JRESI) by undersampling the (k_y, t_1) plane [61].

Here, we extend those principles to 5D (3 spatial + 2 spectral) spectroscopic imaging and apply NUS to the (k_y, k_z, t_1) volume. The third spatial dimension can provide nearly full coverage of the brain in a single scan. This aspect is critical when locations of interest (e.g. tumors) are not apparent in an MRI scan or require detailed mapping in three dimensions [145]. In addition, regional metabo-

lite differences have been found from whole brain SI in varied pathologies such as traumatic brain injury [67] and ALS [165, 177]. The third spatial dimension also enhances the self sparsity of data, and spreading the undersampling over an additional dimension increases the incoherence of the artifacts [102]. These factors improve the performance of CS reconstruction and allow us to acquire data with an 8-fold acceleration, which is the same time required for a fully sampled single-slice data set at the resolution shown here. It should be mentioned that previous applications of CS to SI were used to accelerate already existing protocols, whereas in our case, 5D SI is untenable in vivo without significant acceleration.

Existing fast 3D SI methods all acquire data with a single spectral dimension and have typically focused on high resolution mapping of the singlet signals of NAA, creatine, and choline [121, 71, 111]. Dreher *et al.* used a steady-state free precession EPSI to additionally map myo-Inositol, though their technique is not well suited for detecting glutamate/glutamine (Glx) resonances [49]. Recently, reliable short TE whole brain maps of Glx were also shown [47]. However, none of these techniques offer much potential for separately resolving highly overlapped metabolites such as glutamate and glutamine.

4.3 Methods

4.3.1 Phantom

All scans were acquired on a Siemens 3T Trio scanner, with TR/minimum TE = 1200/30 ms. A phantom with physiological concentrations [68] of 16 grey matter brain metabolites was scanned 12 times to measure spectral quality and reproducibility. The fully-sampled 5D EP-JRESI sequence acquired a k-space volume with 16 points in the readout direction (k_x) and 16 and 8 points in the two phase encoding directions (k_y and k_z , respectively). The field of view (FOV) was $16 \times 16 \times 12$ cm³. Each k_x line was repeatedly acquired for 256 t_2 points per

TR, and 64 TE increments (t_1 points) were acquired. A maximum echo sampling scheme was applied in which t_1 is incremented only before the last refocusing pulse [106, 154]. The second increment that leads to refocused chemical shift is applied as a frequency-dependent, linear phase correction in postprocessing for each acquisition. The reconstructed bandwidths in the direct (F_2) and indirect (F_1) spectral dimensions were 1190 and 500 Hz, respectively. Water suppression was applied before each scan [117]. Scans were recorded with a single channel transmit/receive extremity coil mostly, though three scans were recorded with the same head receive coil used in the in vivo experiments. Total scan time was 2h 44m. Different NUS rates were applied retrospectively.

4.3.2 in vivo

Ten healthy volunteers (mean age 25.0 years old) were scanned using the same parameters as the phantom scans except that the FOV was increased to $24 \times 24 \times 16$ cm³, and 8x NUS was prospectively applied for a scan time of around 20 minutes. Outer volume saturation bands were included outside the PRESS volume of interest (VOI). An eight channel head receive coil was used. One scan was excluded due to poor water suppression as a result of subject motion. For a time-equivalent comparison, a fully-sampled 4D single slice scan with the same slice thickness and in-plane resolution was also acquired in one healthy 23 year old volunteer.

4.3.3 NUS masks

Incoherent aliasing was achieved by randomly undersampling the (k_y, k_z, t_1) volume. As previously described in Ref [61], a nonuniform sampling density function that favors the higher SNR data points was considered and was given by

$$\rho(k_y, k_z, t_1) = \exp \left\{ -\frac{|k_y|}{a} - \frac{|k_z|}{b} - \frac{t_1}{c} \right\} \quad (4.1)$$




Figure 4.1: 8x nonuniform sampling mask over the (k_y, k_z, t_1) volume used for phantom and in vivo scans. White squares indicate sampled points.

where $\rho(k_y, k_z, t_1)$ is the probability of sampling point (k_y, k_z, t_1) , $k_y \in \{-8, -7, \dots, 7\}$, $k_z \in \{-4, -3, \dots, 3\}$, $t_1 \in \{0, 1, \dots, 63\}$, and a , b , and c give the relative decay in each dimension. This method is similar in principle and motivation to that described by Barna *et al.* [12]. In the 5D EP-JRESI sequence as described above, there is less sampling in the spatial dimensions than the indirect spectral dimension, so in order to ensure adequate sampling of that indirect spectral dimension, we have set $a = b = 2$ and $c = 1$. The 8x mask used in phantom retrospectively and in vivo prospectively is shown in Fig 4.1. Independent masks were generated for 4x, 8x, 12x, and 16x undersampling. In each case, 50 masks were generated offline, and the one with the smallest maximum off center peak of the point spread function was implemented on the scanner.

4.3.4 Reconstruction

Two versions of the CS reconstruction problem were formulated. In the first, the reconstructed data is considered self sparse, due largely to the sparsity of 2D JPRESS spectra, and the optimization is written as

$$\min_u \|u\|_1 \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (4.2)$$

where $u = u(x, y, z, F_2, F_1)$ is the reconstructed data, $f = f(x, k_y, k_z, F_2, t_1)$ is the undersampled data Fourier transformed across the fully sampled dimensions, σ^2 is an estimate of the noise variance of the sampled data divided by 100, R is the NUS masking operator, and \mathcal{F} is the Fourier transform operator across the undersampled dimensions only. The objective function is the ℓ_1 -norm of the reconstructed data, and the constraint expresses data consistency through the ℓ_2 -norm of the difference between the acquired data and the subsampled reconstructed data. Each iteration of the solver requires transformation from the f domain to the u domain and back. For such large datasets and with a high number of iterations, a significant amount of reconstruction time is spent on Fourier transforms. Transforming the (k_x, t_2) dimensions only once before reconstruction therefore saves time and is permissible since $\mathcal{F}^{-1}R\mathcal{F}$ equals identity when the dimension is fully sampled ($R = \vec{1}$).

The second CS reconstruction is based on the assumption of piecewise constant spatial profiles and seeks to minimize the total variation (TV) of the reconstructed data with the same data consistency constraint. It is given by

$$\min_u TV(u) \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (4.3)$$

and was similar to the method used in Ref [61]. The TV -norm was given isotropically as

$$TV(u) = \sum_{i=1}^N \sqrt{(\nabla_x u)^2 + (\nabla_y u)^2 + (\nabla_z u)^2} \quad (4.4)$$

where N is the number of acquired data points. In this experiment, this problem took about three times longer per iteration than that of Eq. [4.2].

Both Eqs. [4.2] and [4.3] were solved using the split Bregman algorithm [66]. The split Bregman algorithm seeks to minimize the Bregman distance of the objective function and uses operator splitting [16] to simplify the computations. It is equivalent to the alternating direction method of multipliers (ADMM) with one iteration [58]. Each reconstruction had the same stopping criteria, namely: (1) a maximum number of 200 iterations and (2) accepted level of deviation from the acquired data as expressed through the parameter σ .

4.3.5 Additional processing

After Fourier transformation across t_2 but before reconstruction, a linear phase correction was applied to effectively refocus chemical shift as described above. At the same time, a spectral mask was applied in F_2 to select only the region between 1.2 and 4.3 ppm, which is where nearly all of the brain metabolites resonate. This constrains the dynamic range of the data and effectively removes the large residual water signal and most lipid contamination. A similar approach was applied by Eslami *et. al.* to SI data with one spectral dimension [57].

The reconstruction was performed over each coil separately, and the individual coil reconstructions were combined as a sum-of-squares. Metabolite maps were found by integrating the reconstructed spectra over the peak ranges given in Table 4.1. An average noise power for each sized region was subtracted from the integral to compensate for the bias introduced by nonzero mean noise in the absolute magnitude spectra [72].

For the quantitative comparison, voxels were excluded based on any of three criteria: (1) low SNR, (2) undue influence of the residual water tail, (3) large fat signal bleed. The spectral regions used to perform these tests are also given in

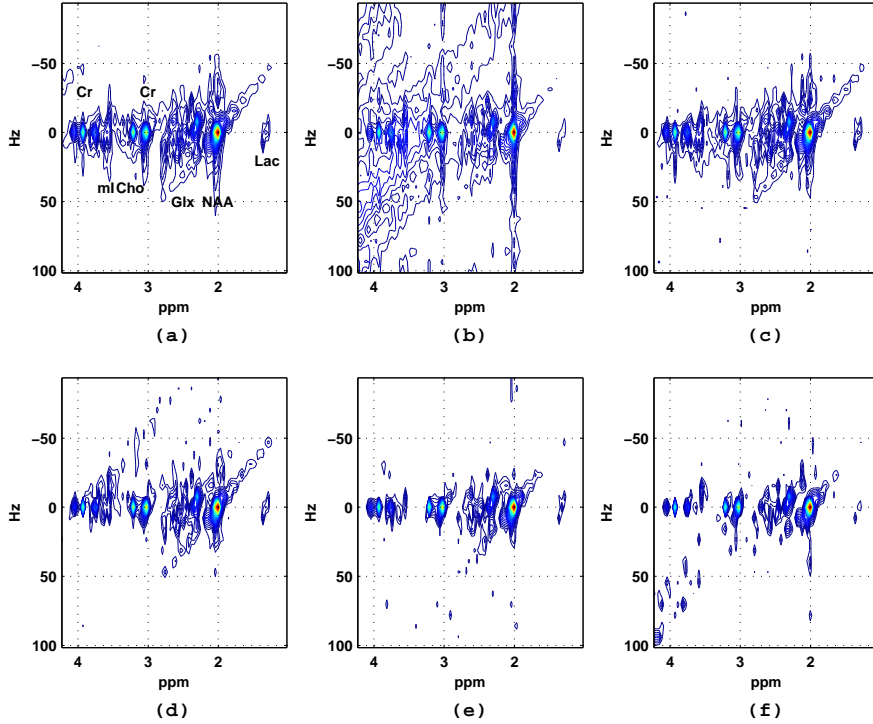


Figure 4.2: J-resolved spectrum from the same central voxel in the phantom for different undersampling rates. (a) fully sampled, (b) 4x NUS with minimum energy reconstruction (missing data points were filled in with zeros), (c)-(f) ℓ_1 reconstructions for 4x, 8x, 12x, and 16x NUS, respectively.

Table 4.1. The first condition effectively excluded voxels outside the VOI, while the second was important for voxels in the frontal lobe. The third condition is important for voxels near the skull marrow. Acceptance criteria was manually confirmed with visual inspection of the NAA, Cr, and tCho singlets.

4.4 Results

4.4.1 Phantom

Spectra from a central voxel in a typical phantom scan are shown with various retrospective undersamplings in Figure 4.2 along with the same fully sampled spec-




Figure 4.3: Glx metabolite map for a typical phantom scan with various retrospective undersamplings. Slices are stacked vertically, and the readout direction is horizontal. Reconstruction was performed according to Eq. [4.3], minimizing the TV -norm. ‡Reconstruction did not include spectral masking.

trum. The minimum energy reconstruction in which the unacquired data points are filled in with zeros is shown (Fig 4.2b) for the same 4x undersampling as that reconstructed in Fig 4.2c with Eq (4.2). Note that the spectral aliasing artifacts appear tilted along the F_1 dimension due to the maximum echo sampling scheme and postprocessing. The CS reconstructed 4x spectra show excellent spectral quality with respect to the fully sampled. As expected, a decrease in spectral quality exemplified by residual aliasing and the removal of low lying peaks is observed with progressively larger undersampling. A metabolite map of the integrated volume of the NAA singlet is presented in Fig 4.4, showing the spatial aliasing artifacts in the minimum energy reconstruction in both the y - and z -directions. Metabolite maps for the remaining metabolites show similar quality.

Figure 4.3 shows the metabolite maps of Glx from the same phantom scan reconstructed using TV minimization in Eq. (4.3) undersampled by 8, 12, and 16, along with the fully sampled reconstruction. For each undersampling, the reconstructions are shown with and without prior spectral masking as noted. Full spectra from the central voxel of each reconstruction are shown in Fig 4.5 with the display extended to show the residual water peak. Without spectral masking, the residual water peak is much higher, and the metabolite peaks are reconstructed at lower intensities. Also, comparing Fig 4.5(a)-(c) to Fig 4.2(d)-(f) qualitatively suggests that the TV -norm objective performs better than the ℓ_1 -norm, especially at greater undersampling factors.

Quantitative comparisons between the various undersampling factors and reconstructions over all the phantom scans are shown in Table 4.2 and Figure 4.6. In Table 4.2, peak root-mean-square-errors (RMSEs) are computed between the fully sampled and reconstructed data and are reported in decibels (dB). The same spectral region used for peak integration was used for each metabolite RMSE. For proper comparison, the minimum energy reconstructions are also included. Within each reconstruction, there is a clear trend of higher RMSEs with greater under-




Figure 4.4: Localized three plane image of phantom (left) and NAA metabolite map for a typical phantom scan with various retrospective undersamplings (right). The phase encoding direction is shown vertically, the readout is horizontal, and the slices are stacked. Reconstruction was performed according to Eq. (4.2), minimizing the ℓ_1 -norm except where noted. †Minimum energy reconstruction

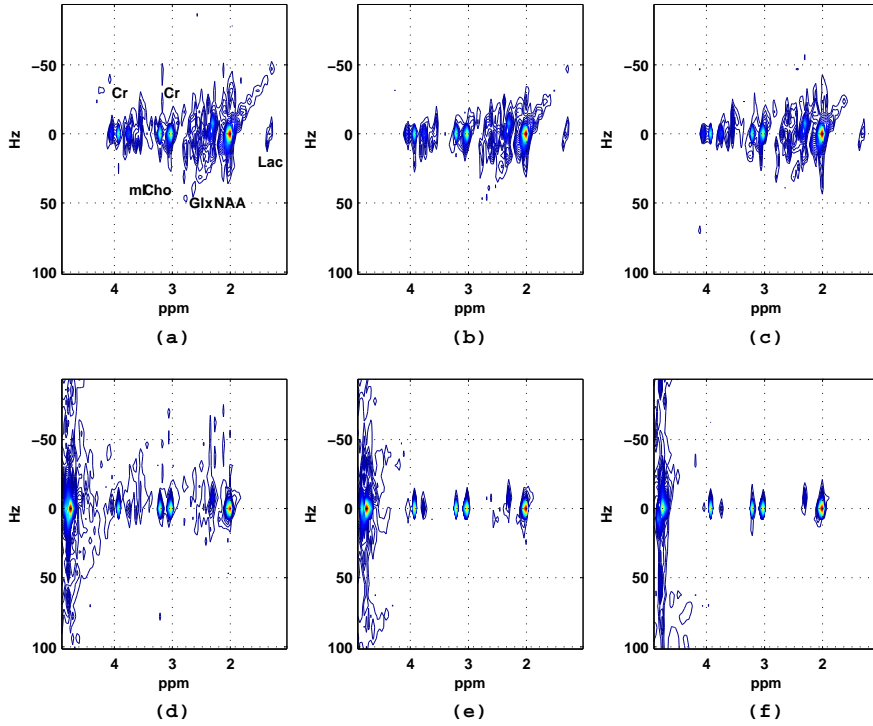


Figure 4.5: The effect of spectral masking before reconstruction. J-resolved spectrum from the same central voxel in the phantom for different undersamplings with TV reconstruction. (a)-(c) are for 8x, 12x, and 16x NUS, respectively and include spectral masking. (d)-(f) use the same NUS and reconstruction except without any spectral masking; the contour level has been scaled down by a factor of two to partially compensate for the decreased signal intensity. In all spectra, the F_2 axis is extended to show the residual water peak.

sampling factors for each metabolite. Exceptions occur in the TV reconstruction between 4x and 8x undersampling and in the ℓ_1 reconstruction of Cr39 between 12x and 16x. RMSEs are also presented for 8x undersampled data reconstructed with ℓ_1 minimization without prior spectral masking. Though the NAA results are surprisingly close to the fully sampled, most of the other metabolites are much less similar and are even worse than the minimum energy reconstruction in some instances, indicating that the spectral mask improves the reconstruction of small to medium peaks. Trends between the ℓ_1 and TV reconstructions are less obvious with ℓ_1 exhibiting lower RMSEs at lesser undersampling but TV exhibiting lower ones at greater undersampling.

	$F_2(ppm)$	$F_1(Hz)$
NAA	(1.8, 2.2)	(-15, 15)
Glx	(2.2, 2.6)	(-15, 15)
Cr30	(2.9, 3.1)	(-15, 15)
Cho	(3.1, 3.3)	(-15, 15)
mI	(3.4, 3.8)	(-15, 15)
Cr39	(3.8, 4.0)	(-15, 15)
Wat	(4.0, 4.3)	(-15, 15)
Tail	(2.0, 4.3)	(-250, 47)
Fat	(0.7, 1.6)	(-15, 15)
Sig	(2.4, 3.1)	(-15, 15)
Noise	(1.3, 2.5)	(120, 240)

Table 4.1: Spectral ranges used to measure peak integrals. Cho is total choline; Cr30 is the creatine singlet at 3.0 ppm; Cr39 is the creatine singlet at 3.9 ppm; Glx is glutamate and glutamine combined; mI is myo-inositol; NAA is N-acetylaspartate. Wat and Tail are used as a measure of baseline effects due to residual water contamination. Fat is a marker of lipid signal bleed. Sig is a marker of overall signal level. Noise is taken from a region that is devoid of any signal, desired or nuisance.

	<i>Min En</i>				ℓ_1					<i>TV</i>			
	4x	8x	12x	16x	4x	8x [‡]	8x	12x	16x	4x	8x	12x	16x
NAA	-10.23	-7.22	-5.32	-4.22	-14.71	-14.95	-14.65	-14.53	-14.47	-14.59	-14.67	-14.61	-14.51
Glx	-9.22	-6.64	-4.52	-3.17	-13.98	-6.53	-13.72	-13.67	-13.58	-13.81	-13.67	-13.68	-13.47
Cr30	-9.64	-6.84	-5.10	-3.81	-14.20	-10.74	-14.22	-13.79	-13.76	-14.13	-14.21	-13.99	-13.85
tCho	-9.30	-6.54	-5.08	-3.54	-13.82	-7.09	-13.77	-13.57	-13.41	-13.82	-13.84	-13.68	-13.59
mI	-7.90	-5.78	-4.87	-4.84	-10.96	-4.67	-10.81	-10.72	-10.61	-10.77	-10.49	-10.20	-10.22
Cr39	-9.56	-7.12	-5.50	-4.55	-14.09	-8.03	-13.84	-13.57	-13.71	-14.07	-14.09	-13.81	-13.75

Table 4.2: Mean metabolite peak RMSEs in phantom in decibel scale. *Min En* is the minimum energy reconstruction.

‡84

‡Spectral masking was not applied prior to reconstruction.

Figure 4.6 compares the reconstructions by looking at peak integrals of NAA, Glx, Cr, tCho, and mI. In each case, the peak integrals tend to decrease with greater undersampling with the differences being smaller for *TV* reconstruction. The standard deviations of the fully sampled data are similar to those for each reconstruction with those of mI being the highest due to its closer proximity to the residual water tail. Table 4.3 gives the numbers used for this figure.

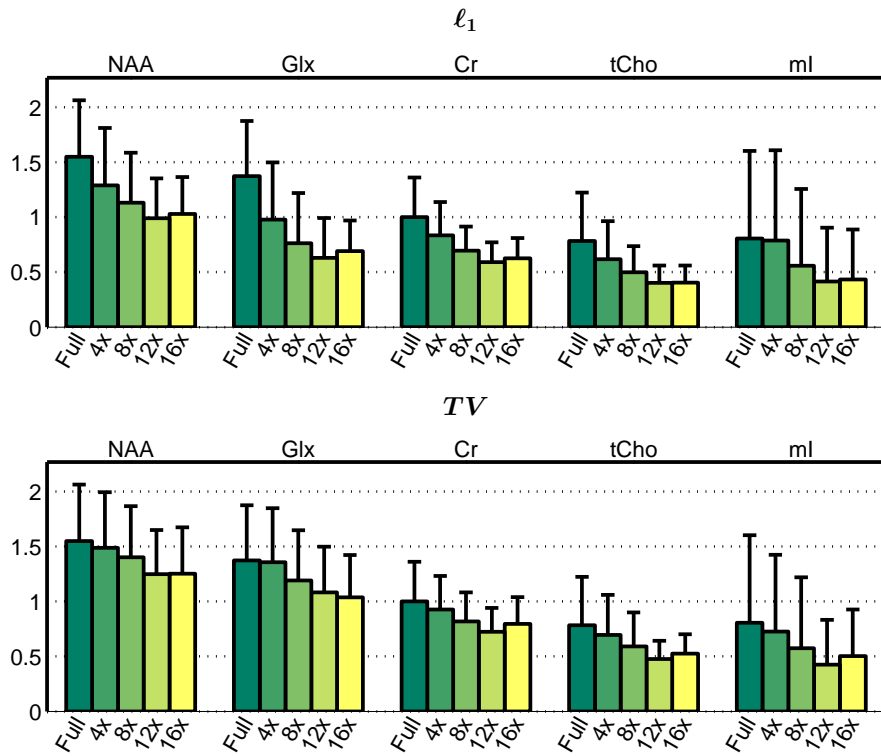


Figure 4.6: Peak integrals reconstructed from different undersamplings for five metabolites. Integrals were normalized by the mean of Cr in the fully sampled data. Reconstruction was performed according to Eq. [4.2] (top) and Eq. [4.3] (bottom).

	ℓ_1					TV			
	Full	4x	8x	12x	16x	4x	8x	12x	16x
NAA	1.55 (0.51)	1.29 (0.52)	1.13 (0.46)	0.99 (0.36)	1.03 (0.34)	1.49 (0.50)	1.40 (0.47)	1.25 (0.40)	1.25 (0.42)
Glx	1.37 (0.50)	0.98 (0.52)	0.76 (0.46)	0.63 (0.36)	0.69 (0.28)	1.36 (0.49)	1.19 (0.46)	1.08 (0.42)	1.04 (0.39)
Cr30	1.00 (0.36)	0.84 (0.30)	0.70 (0.22)	0.59 (0.18)	0.62 (0.19)	0.93 (0.31)	0.82 (0.27)	0.72 (0.22)	0.80 (0.24)
tCho	0.78 (0.44)	0.62 (0.35)	0.50 (0.24)	0.40 (0.16)	0.41 (0.15)	0.70 (0.36)	0.59 (0.31)	0.48 (0.17)	0.52 (0.18)
mI	0.81 (0.80)	0.79 (0.82)	0.56 (0.70)	0.41 (0.49)	0.43 (0.45)	0.73 (0.70)	0.57 (0.64)	0.43 (0.41)	0.50 (0.42)

Table 4.3: Mean metabolite peak integrals in phantom with standard deviations in parentheses.

4.4.2 in vivo

Due to obvious time constraints, 5D in vivo data was prospectively undersampled only, so there is no fully-sampled data to compare with. Quality metrics such as RMSE cannot therefore be calculated. Instead, peak integrals are used to show reproducibility and that values are within the physiological range for healthy adults. Figure 4.7a shows metabolite maps from a healthy 29 year old adult using 8x undersampling and TV reconstruction. Each metabolite was scaled to its maximum so that they all can appear with the same grey scale despite different overall intensities. Figure 4.7b shows the full spectra from the voxels highlighted in Fig 4.7a. The red and blue voxels are from different sides of the occipital lobe, the yellow voxel is from the parietal lobe, and the green voxel is from the frontal lobe. The occipital voxels tend to have the best water suppression and exhibit the highest quality spectra. This is consistent with our previous experience both in single voxel and single slice spectroscopic imaging. Fat leakage into the occipital voxels is present but small enough not to disturb the NAA singlet. Voxels in the frontal lobe show worse water suppression due to susceptibility related effects and difficulties in shimming, and the residual water tail can affect the integrals of the Cr39 singlet and the mI multiplet.

Figure 4.8 shows the mean metabolite peak integrals over nine healthy adults using TV minimization. Voxels in the frontal lobe (top) were separated from those in the occipital and parietal lobes (bottom). This is because of the markedly increased standard deviations in the frontal lobe in scans with poorer water suppression. A time-equivalent, 4D fully-sampled data set is included and shows similar means and standard deviations as the reconstructed 5D data. Figure 4.9 shows the data reconstructed minimizing the ℓ_1 -norm. Both figures are similar, though the TV reconstruction results in a tighter spread of intersubject averages and slightly smaller intrasubject standard deviations as well. The numeric data for all 10 healthy volunteers along with the number of acceptable voxels for each

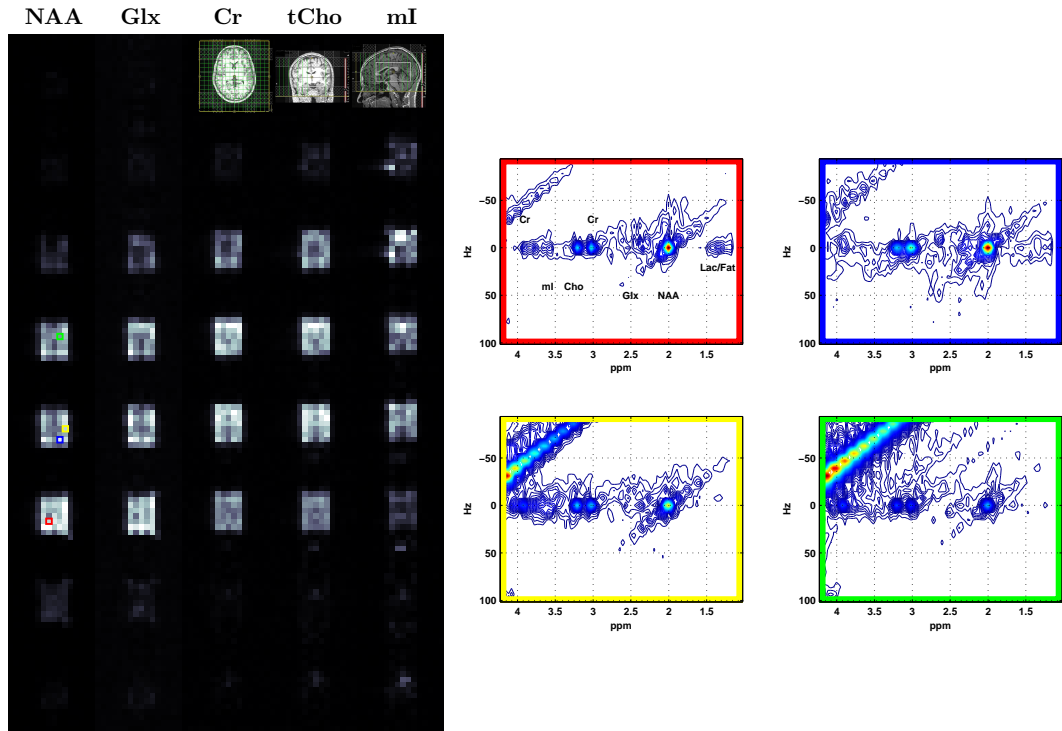


Figure 4.7: Metabolites maps (left) for a 29 year old healthy adult scanned with 8x undersampling and reconstructed using Eq. [4.3]. Each metabolite was normalized to its maximum so that they appear on the same gray scale. Spectra from the highlighted voxels are shown at right.

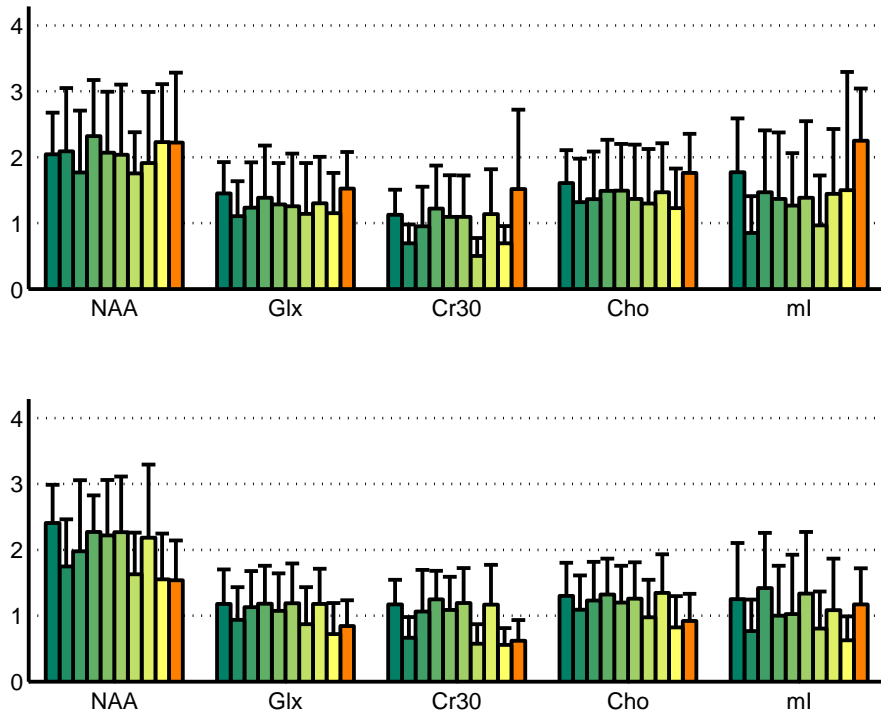


Figure 4.8: Regional peak integrals of major metabolites in ten healthy adults with minimum TV reconstruction. The green bars are from 8x NUS 5D data reconstructed using Eq. [4.3] for voxels in the frontal lobe (top) and in the occipito-parietal lobes (bottom). The orange bar is from a time-equivalent, fully-sampled 4D data set for comparison.

scan are given in Table 4.4 and Table 4.5 for voxels in the occipito-parietal and frontal lobes, respectively.




Figure 4.9: Regional peak integrals of major metabolites in ten healthy adults with minimum ℓ_1 reconstruction. The green bars are from 8x NUS 5D data reconstructed using Eq. (4.2) for voxels in the frontal lobe (top) and in the occipito-parietal lobes (bottom). The orange bar is from a time-equivalent, fully-sampled 4D data set for comparison.

	Scan1	Scan2	Scan3	Scan4	Scan5	Scan6	Scan7	Scan8	Scan9	Scan10	4D Full	
ℓ_1	NAA	1.72 (0.68)	2.78 (0.66)	1.50 (0.62)	2.22 (1.29)	2.53 (0.69)	2.41 (0.94)	2.56 (1.02)	1.51 (0.65)	2.50 (1.34)	1.43 (0.68)	1.54 (0.60)
	Glx	1.09 (0.45)	1.38 (0.65)	0.83 (0.41)	1.23 (0.63)	1.34 (0.72)	1.21 (0.68)	1.38 (0.78)	0.78 (0.49)	1.36 (0.67)	0.79 (0.42)	0.84 (0.39)
	Cr30	1.26 (0.39)	1.32 (0.49)	0.74 (0.34)	1.03 (0.69)	1.13 (0.50)	1.03 (0.55)	1.14 (0.56)	0.57 (0.23)	1.26 (0.71)	0.72 (0.28)	0.62 (0.32)
	tCho	1.13 (0.55)	1.52 (0.64)	0.93 (0.45)	1.36 (0.70)	1.55 (0.68)	1.33 (0.67)	1.46 (0.72)	0.96 (0.68)	1.56 (0.71)	0.87 (0.47)	0.92 (0.41)
	mI	1.41 (0.58)	1.49 (1.16)	0.90 (0.39)	1.55 (1.01)	1.10 (0.95)	1.14 (1.07)	1.50 (1.23)	0.80 (0.42)	1.28 (0.99)	0.92 (0.36)	1.17 (0.55)
TV	NAA	1.98 (0.71)	2.43 (0.59)	1.77 (0.72)	2.00 (1.09)	2.29 (0.56)	2.24 (0.85)	2.29 (0.85)	1.64 (0.64)	2.21 (1.12)	1.57 (0.70)	1.54 (0.60)
	Glx	1.06 (0.49)	1.19 (0.53)	0.95 (0.50)	1.14 (0.55)	1.20 (0.58)	1.08 (0.58)	1.20 (0.61)	0.88 (0.57)	1.19 (0.54)	0.73 (0.48)	0.84 (0.39)
	Cr30	0.83 (0.28)	1.18 (0.37)	0.67 (0.32)	1.07 (0.64)	1.26 (0.44)	1.10 (0.51)	1.21 (0.53)	0.58 (0.30)	1.18 (0.61)	0.56 (0.26)	0.62 (0.32)
	tCho	1.13 (0.56)	1.31 (0.51)	1.10 (0.53)	1.24 (0.59)	1.34 (0.55)	1.21 (0.57)	1.27 (0.56)	0.99 (0.57)	1.36 (0.59)	0.83 (0.48)	0.92 (0.41)
	mI	0.93 (0.38)	1.27 (0.86)	0.77 (0.48)	1.43 (0.85)	1.01 (0.77)	1.04 (0.91)	1.35 (0.94)	0.81 (0.57)	1.10 (0.79)	0.63 (0.37)	1.17 (0.55)
N_{vox}	38	37	40	60	79	55	61	65	63	38	23	

Table 4.4: Mean metabolite peak integrals from the occipito-parietal lobes in vivo of 10 healthy volunteers. Standard deviations are in parentheses, and the number of included voxels are given for each scan. The top half shows the reconstruction minimizing the ℓ_1 -norm, while the bottom half is the result from minimizing the TV -norm. The last column shows the values from a fully sampled single slice (4D) scan.

	Scan1	Scan2	Scan3	Scan4	Scan5	Scan6	Scan7	Scan8	Scan9	Scan10	4D Full	
T_1	NAA	1.94 (0.88)	2.42 (0.77)	1.69 (1.09)	1.94 (1.10)	2.56 (1.01)	2.25 (1.10)	2.26 (1.26)	1.73 (0.83)	2.20 (1.30)	1.71 (0.94)	2.22 (1.06)
	Glx	1.57 (0.82)	1.67 (0.61)	0.88 (0.38)	1.39 (0.87)	1.59 (0.93)	1.55 (0.79)	1.47 (1.02)	1.15 (0.96)	1.53 (0.88)	0.95 (0.42)	1.52 (0.55)
	Cr30	1.26 (0.41)	1.28 (0.48)	0.87 (0.31)	0.74 (0.51)	1.06 (0.68)	0.84 (0.60)	0.91 (0.64)	0.60 (0.27)	1.15 (0.80)	0.88 (0.36)	1.52 (1.21)
	Cho	1.37 (0.88)	1.89 (0.62)	1.03 (0.49)	1.56 (0.89)	1.74 (0.97)	1.79 (0.89)	1.65 (1.07)	1.19 (0.84)	1.71 (0.91)	0.97 (0.45)	1.76 (0.59)
	mI	6.40 (7.75)	2.05 (0.96)	1.05 (0.48)	1.59 (1.08)	1.59 (1.17)	1.51 (0.96)	1.61 (1.52)	0.89 (0.55)	1.63 (1.24)	2.02 (1.74)	2.25 (0.79)
T_2	NAA	2.27 (0.71)	2.07 (0.64)	2.11 (0.97)	1.78 (0.95)	2.35 (0.86)	2.09 (0.93)	2.06 (1.08)	1.77 (0.63)	1.93 (1.09)	2.25 (0.89)	2.22 (1.06)
	Glx	1.64 (0.67)	1.47 (0.48)	1.12 (0.54)	1.25 (0.69)	1.40 (0.80)	1.30 (0.63)	1.27 (0.81)	1.15 (0.78)	1.31 (0.71)	1.16 (0.62)	1.52 (0.55)
	Cr30	0.89 (0.36)	1.14 (0.38)	0.70 (0.29)	0.96 (0.61)	1.23 (0.66)	1.11 (0.64)	1.11 (0.64)	0.51 (0.27)	1.15 (0.69)	0.70 (0.26)	1.52 (1.21)
	Cho	1.75 (0.96)	1.62 (0.51)	1.33 (0.67)	1.38 (0.73)	1.51 (0.78)	1.51 (0.72)	1.38 (0.83)	1.31 (0.83)	1.48 (0.75)	1.24 (0.61)	1.76 (0.59)
	mI	5.25 (4.40)	1.79 (0.82)	0.86 (0.57)	1.48 (0.95)	1.38 (1.02)	1.28 (0.80)	1.40 (1.17)	0.97 (0.77)	1.46 (0.99)	1.52 (1.81)	2.25 (0.79)
N_{vox}	16	27	19	52	43	22	36	21	58	20	17	

Table 4.5: Same as Table 4.4 for voxels in the frontal lobe.

4.5 Discussion

The main benefit of the technique described here is the potential for J-resolved spectroscopic imaging in three spatial dimensions in vivo. However, the challenges faced in conventional 3D SI are also present here. We have chosen to use a volume-localized (PRESS) EP-JRESI as opposed to spin echo or slice-localized sequence, despite the better coverage and shorter echo times afforded by the latter. Volume-localized techniques permit better shimming, which leads to taller, narrower peaks suited for CS reconstruction and better water suppression. Also, slice-localized techniques suffer from skull marrow lipid contamination due to ringing of the high amplitude, low resolution point spread function. Contamination can be partially compensated with inversion recovery [162] or postprocessing techniques involving a high resolution lipid image [78]. Outer volume suppression pulses effectively remove all signal (metabolites included) but are difficult to position over the skull marrow in a 3D scan without significantly suppressing cortical signal near the surface of the brain. The chosen spectral mask further ensured minimization of any residual lipid artifacts but would be inadequate to remove the ringing caused by unsuppressed signal from the skull marrow by itself. Volumetric localization using adiabatic RF pulses [63] or adiabatic refocusing pulses [150] can be used to greatly reduce chemical shift displacement error and allow ROIs very near the skull marrow. Comparing the metabolite maps in Figure 4.7a, there is about a one pixel shift per ppm in the y -direction and a 1/2 slice shift per ppm using the “Mao” refocusing pulses [109] in PRESS.

In this work, we have used two methods for quantitative comparison of reconstructed data. The RMSE is applicable when a reference signal is available as was the case with phantom data only. It is a point-by-point method, and as such, it is sensitive to differences in amplitude, linewidth, and peak shape. However, it is also influenced by noise and can be biased by the nonlinear reconstruction meth-

ods used here. To minimize the influence of noise and better characterize where discrepancies are occurring, RMSEs were constrained to the individual metabolite peak areas. As an additional quantitative method, peak integration simply sums the volume in a given region and therefore does not directly describe differences in amplitude, linewidth, or peak shape. No reference data is required, but over- or underestimation can be determined if it is present. From a philosophical standpoint, peak integration provides a relative measure of metabolite concentration which is the penultimate goal in spectroscopy studies anyway. Both methods can be considered complementary.

The decrease in peak volumes from fully sampled to 12x undersampling shown in Fig 4.6 is consistent with the corresponding increase in RMSE values shown in Table 4.2 and is most likely the result of the nonlinear reconstruction itself. In either ℓ_1 or TV minimization, soft thresholding tends to shrink the amplitudes of the signal in the transform domain (identity or finite differences), reducing the peak integral. The data consistency term is what mitigates this effect. However, the larger the undersampling factor, the less influence the data consistency term has (assuming that the value of σ does not depend on R). Simply put, there are fewer data points to constrain.

Previous work on CS-reconstructed single slice J-resolved spectroscopic imaging data has suggested that TV reconstruction should perform better than ℓ_1 [61]. In this case though, adding a third spatial dimension increases the self sparsity of the data as the ROI is about half the FOV in the z -direction, and several slices do not contribute any signal. This effectively narrows the gap between the reconstructions to the point that it is not definitive if there is a significant difference. The reconstruction in Eq. [4.2] takes about 10 minutes per coil, while that of Eq. [4.3] takes about half an hour on a 64-bit octa core 3.1 GHz Intel Xeon E5-2687W processor with 128 GB RAM, so there is strong incentive to choose the faster reconstruction method when the number of coils is large. Coil compression methods

could be applied to the acquired data to reduce the number of virtual coils that need to be reconstructed [25]. It should be noted that a scan with a tighter FOV around the VOI would be less self sparse though and could negatively influence the ℓ_1 reconstruction. However, the TV reconstruction should be relatively immune to the ROI/FOV ratio as long as the object profile is not wildly varying.

A relatively large indirect spectral bandwidth of 500 Hz was used in this paper. Often, J-resolved spectra can be acquired with a bandwidth less than 100 Hz due to the small coupling constants found in most metabolites [68]. The primary reason for this choice was to keep the residual water tail away from the metabolite signals, as we have found it difficult to achieve consistent water suppression over the whole brain. A further consequence of this choice is that the 2D spectra are highly sparse, which benefits the reconstruction. However, a large indirect spectral bandwidth is not a prerequisite for the technique described. The same principles of sampling (Eq. (4.1)) and reconstruction can be applied with minimal modification to low indirect bandwidth data with the similar caveat that the ℓ_1 reconstruction will be more negatively affected than the TV .

Figures 4.3 and 4.5 show how important a simple spectral masking is prior to the reconstruction. In either reconstruction, the data fidelity term will make sure the largest peaks are reconstructed most accurately, while low-lying peaks are “cleaned up” as aliasing and largely removed. Even with good water suppression, the residual water is typically the largest peak in brain spectra. Individual peaks of multiplets are smaller and reconstructed with much less intensity as shown in the above mentioned figures. On the other hand, the NAA singlet peak has the second highest amplitude and is well separated from the residual water tail. It is reconstructed accurately with or without masking in Table 4.2, but the errors in reconstruction for Glx and mI are much more pronounced. As an alternative to the simple spectral masking presented here, other postprocessing methods to remove the residual water signal can be applied instead, for instance that based on

wavelet transformation [73]. This method would be advantageous when requiring accurate quantitation of metabolites very near the water peak, such as glutathione. The same idea applies to the lipid peaks and their effect on lactate and alanine.

Table 4.2 and Figure 4.5 suggest that undersampling factors up to 16x can be achieved with excellent reconstruction quality in phantom. However, we have limited ourselves to 8x undersampling in vivo. The reason for this conservative choice is in consideration of the motivation to do multidimensional spectroscopy in the first place. The major singlet resonances in brain are well resolved in single dimensional spectroscopy with long echo times. The advantage of using multidimensional spectroscopy lies in its ability to resolve low amplitude, J -coupled resonances acquired with short echo times. Without accurate measurements of these additional resonances, there is no reason not to choose the much faster 1D technique. As the undersampling factor increases, more of the low amplitude resonances are buried in the noise and residual aliasing and will not be reconstructed accurately. Future work will look at prior knowledge fitting as an improved method of quantification of up to 20 brain metabolites [153].

Further improvements may be made by considering additional NUS schemes and reconstruction objectives. In particular, the Poisson gap sampling scheme shows reduced artifacts compared to random undersampling [83]. Maximum entropy has been used as an objective in 2D NMR and might offer some improvement compared to ℓ_1 -norm minimization [166, 26], while group sparsity uses the $\ell_{2,1}$ -norm objective to consider the proximity of nonzero coefficients [27].

Although this paper has focused on 2D J-resolved spectroscopy, the techniques can be applied with suitable modification to imaging analogs of other 2D acquisitions such as correlated spectroscopy, TE-averaged PRESS [82], or CT-PRESS [50]. Despite the common practice of displaying the latter two as single dimension spectra, the acquisition requires two spectral dimensions and can be undersampled and reconstructed as 2D spectra with the methods described here before taking

the appropriate projections.

4.6 Conclusion

Nonuniform sampling with compressed sensing reconstruction can be used to acquire 5D (3 spatial + 2 spectral) spectroscopic images in human brain in a clinically feasible scan time. Conservatively, an 8-fold acceleration was applied in vivo for a 20 minute scan time, while a 16-fold acceleration was applied to phantom data. Reconstruction can be performed either by minimizing the ℓ_1 -norm or the TV -norm.

CHAPTER 5

3D Correlated Spectroscopic Imaging of Calf Muscle Using Group Sparse Reconstruction of Undersampled Single and Multichannel Data

5.1 Introduction

Correlated spectroscopy (COSY) was the first proposed two dimensional (2D) spectroscopy technique [9]. A COSY spectrum is characterized by the presence of cross peak multiplets indicating scalar coupling between resonances and is achieved following the application of a 90° coherence transfer pulse to single quantum coherences and incrementing the echo time. Due to its relative simplicity, COSY was also one of the first 2D spectroscopy sequences successfully applied in vivo [22, 171]. While the spectral resolution in vivo is T2 limited and not high enough to fully resolve the individual peaks that comprise each cross peak multiplet, the overall cross peaks showcase the clear presence of coupled partners and where those partners resonate. Compared to single dimensional (1D) spectra, COSY spectra are less dense with improved spectral dispersion.

However, applications of COSY as a single voxel technique are greatly limited by the lack of coverage and the large voxel size. Spectroscopic imaging (SI) [24] solves these problems while acquiring spectra from multiple locations simultaneously. Standard CSI techniques use phase-encoding in each spatial direction and offer no improvement in scan time efficiency compared to single voxel techniques.

However, by interleaving the acquisition of an entire k-space line with the temporal or spectral information in a single excitation, an echo planar readout [108] can reduce the scan time by at least an order of magnitude. The remaining spatial dimension(s) are phase encoded [126].

Recently, COSY has been combined with an echo planar readout for correlated spectroscopic imaging (EP-COSI) [99, 7]. These sequences have been limited in coverage to a single slice because of the scan times required to incrementally acquire a phase-encoded spatial direction (k_y) and the indirect spectral dimension (t_1). A fully sampled 3D volume scan would take 2-3 hours even with the echo planar readout.

Nonuniform sampling (NUS) with compressed sensing reconstruction has been used to accelerate the acquisition of these single slice correlated SI sequences by subsampling the plane spanned by the phase-encoded direction (k_y) and t_1 [26]. Here we extend those techniques to subsample and reconstruct the volume spanned by two phase-encoded directions (k_y and k_z) and t_1 by utilizing prior knowledge that the COSY spectrum is self-sparse (i.e. the 2D spectra are composed of relatively few non-zero peaks surrounded by low level noise). We also utilize the broadness of the peaks to improve the reconstruction by grouping coefficients that are in close proximity in the spectral plane. Qualitative comparison is performed between a mixed-norm group sparse (GS) reconstruction and a typical compressed sensing ℓ_1 -norm (CS) reconstruction using different undersampling factors with single channel and multi channel data. Reconstructions of prospectively undersampled whole calf correlated spectroscopic imaging data with three spatial and two spectral dimensions are compared.

5.2 Theory

Previous implementations of compressed sensing to multidimensional SI have been formulated based on the signal equation for a single channel. Multicoil scans were reconstructed channel by channel and combined as a sum-of-squares [61, 178]. This single channel problem is formulated as a minimization of an ℓ_1 -norm term while maintaining data consistency through an ℓ_2 -norm term and can be written as

$$\min_u \|\phi(u)\|_1 \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (5.1)$$

where $u = u(x, y, z, F_2, F_1)$ is the reconstructed data, $f = f(k_x, k_y, k_z, t_2, t_1)$ is the undersampled data, σ^2 is an estimate of the noise variance of the sampled data, R is the sampling mask, \mathcal{F} is the Fourier transform, and $\phi(u)$ is a function that transforms u to a sparse domain. The simplest choice is the identity transform $\phi(u) = u$, which is applicable only when the reconstructed data is self sparse. Alternatively, ϕ can be a wavelet transform or a combination finite difference operator as in total variation $\phi(u) = TV(u)$.

Unlike conventional ℓ_1 -norm-based CS reconstructions that consider each transform coefficient independently, GS [181] exploits the tendency of large transform coefficients to be clustered. In order to do this, certain coefficients are grouped and reconstructed as a unit. This grouping allows points in a group to influence each other as a model of signal correlation and has been shown to offer improved results compared to ℓ_1 minimization in the context of MR [174, 128, 27]. Large peaks in an otherwise sparse COSY spectrum are an ideal situation to apply group sparsity.

5.2.1 Single channel GS reconstruction

The objective of a GS reconstruction problem is an $\ell_{2,1}$ mixed-norm in which the ℓ_1 -norm of the ℓ_2 -norms of each group is minimized.

$$\|Gu\|_{2,1} = \sum_i^n \|u_{g_i}\|_2 \quad (5.2)$$

where $G \in \mathbb{R}^{n \times N}$ is the matrix operator that places N vectorized points of u into n different groups. There is a lot of freedom in the choice of grouping. Groups can be identically sized or varying, and they can partition the entire region or overlap with each other.

In [27], the single channel GS problem

$$\min_u \|Gu\|_{2,1} \quad \text{s.t.} \quad \|R\mathcal{F}u - f\|_2^2 < \sigma^2 \quad (5.3)$$

was solved using the split Bregman algorithm [66], and similarly in [46], it was solved using the alternate direction method of multipliers (ADMM).

The solution was found by first writing the problem as an unconstrained minimization, defining an auxiliary variable to split the terms, and using Bregman iterations to enforce the equality of the substitution as well as the original constraint. The $\ell_{2,1}$ -norm subproblem was solved using the group-wise shrinkage function

$$gshrink(Gu, \lambda) = \max\left(0, 1 - \frac{\lambda}{\|Gu\|_{2,1}}\right) \cdot Gu \quad (5.4)$$

5.2.2 Multi channel GS reconstruction

The multicoil CS problem can easily be formulated as a regularized SENSE [132] reconstruction [18, 96]. If the coil sensitivity profiles S are known, the multicoil GS problem is formulated as

$$\min_u \|Gu\|_{2,1} \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (5.5)$$

While Eq. (5.5) resembles Eq. (5.3), the presence of S complicates the matter. This is because split Bregman and ADMM achieve such high efficiency because the most expensive subproblem involves inverting a circulant matrix that can be made diagonal with application of Fourier transforms. But once the sensitivity profiles are included, the matrix that must be inverted ($\mu S' \mathcal{F}' R' R \mathcal{F} S + \lambda G' G$) is no longer circulant. In that case, the subproblem can be approximately solved using a few iterations of the conjugate gradient method [66, 39], but this can be expensive even with preconditioning. Alternatively, variable splitting can be applied to the data fidelity term in the same way it is applied to the objective term [136, 37] at the expense of an additional parameter and without the same guarantees of convergence.

As the reconstructed image/spectral data, u , is self sparse (since the problem seeks to minimize the $\ell_{2,1}$ -norm of groups with no additional transform), the individual coil image/spectral data Su must also be. Instead of (5.5), the SENSE-regularized multicoil GS problem can be formulated as

$$\min_u \|GSu\|_{2,1} \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (5.6)$$

This problem is actually easier to solve than Eq. (5.5) as it requires one less auxiliary variable to be iteratively updated. Applying the substitution $v = Su$, gives

$$\min_{u,v} \|Gv\|_{2,1} \quad \text{s.t.} \quad \|R\mathcal{F}v - f\|_2^2 < \sigma^2, \quad v = Su \quad (5.7)$$

where $v = v(x, y, z, F_2, F_1, c)$ is a six dimensional matrix with the inclusion of the channel dimension c . Since u has been completely decoupled from the rest of the optimization, the problem can be solved iteratively for v first before transforming back to u .

Ignoring the $v = Su$ constraint, the remaining problem is simply the single-channel GS formulation in Eq. (5.3) with v replacing u . It is solved exactly as before by relaxing the constraint, making the substitution $z = Gv$, using the

Bregman formulation, and iteratively solving each subproblem while holding the other variable constant. Bregman updates are applied, including over an outer loop on f until the problem is accurately solved for v . The intermediate solution v is basically the reconstructed coil images similar to those that would be calculated individually in a channel-by-channel reconstruction except stacked as one large matrix.

At this point, u is calculated from v and S by

$$u = (S'S)^{-1} S'v \quad (5.8)$$

For maximal SNR, the reconstructed data can be whitened in order to remove any noise correlations between coils

$$u = (S'\Psi^{-1}S)^{-1} S'\Psi^{-1}v \quad (5.9)$$

where Ψ is the noise correlation matrix between the coils described in detail in [132].

5.3 Methods

5.3.1 Sequence

All scans were acquired on a Siemens 3T Trio scanner. The fully sampled 3D EP-COSI sequence contained a k-space volume with 16 points in the readout direction (k_x) and 16 and 8 points in the two phase-encoding directions (k_y and k_z , respectively). The field of view (FOV) was $16 \times 16 \times 12$ cm³. Each k_x line was repeatedly acquired for 256 t_2 points per TR, and 64 t_1 points (TE increments) were acquired. The (k_y, k_z, t_1) volume was prospectively nonuniformly undersampled by a factor of 8, 12, or 16 as described in the following section. The bandwidths in the direct (F_2) and indirect (F_1) spectral dimensions were 1190 and 1250 Hz, respectively. The 8x and 12x undersampled scans were acquired with minimum

TE/TR = 30/1200 ms for scan times of 20 min and 14 min, respectively. The 16x undersampled scans were acquired with minimum TE/TR = 30/1500 ms for a scan time of 13 min. Water suppression was applied [117]. Scans were acquired on a single channel transmit/receive extremity coil with 8x and 12x undersampling and on a 15-channel knee coil with 12x undersampling. A separate scan with full spatial sampling and a single TE was acquired without water suppression before each EP-COSI scan for reference to determine the complex coil sensitivities. This reference scan took 2 min 28 s with TR = 1200 ms and 3 min 18 s with TR = 1500 ms. When time permitted, a fully sampled, single-slice multi-echo EP-COSI (MEEP-COSI) [60] scan was also acquired through the central slice of the 3D volume with slice thickness of 2 cm. The rest of the applicable parameters were the same, and the single slice scan took 13 minutes with TR = 1.5 s or 10 minutes with TR = 1.2 s.

5.3.2 Apodization and NUS masks

Due to T2 line broadening and the mixed phase characteristics of COSY cross peaks, their amplitudes are usually up to an order of magnitude lower than the diagonal peaks. As the cross peaks are the most interesting and descriptive signals in a COSY spectrum, apodization is applied in t_2 and t_1 to enhance them. The filters used were squared sine bells in each dimension, and they were applied *before* reconstruction to enhance the coherence transfer signal envelope. The unapodized data has much higher diagonal peak to cross peak dynamic range, and reconstruction would therefore favor the diagonal peaks at the expense of cross peaks. Post reconstruction apodization would not be able to restore the cross peaks once the minimization of the $\ell_{2,1}$ -norm has effectively denoised them.

Incoherent aliasing was achieved by randomly undersampling the (k_y, k_z, t_1) volume. In [178], the nonuniform sampling density followed a decaying exponential so that the highest SNR data points were adequately sampled. Here, we follow

the same philosophy except that we wish to ensure the highest SNR data points *after apodization* were subsampled. For that reason the sampling density function was given by

$$\rho(k_y, k_z, t_1) = \exp \left\{ -\frac{|k_y|}{n_y} - \frac{|k_z|}{n_z} \right\} \cdot \sin^2 \left\{ \left(\frac{t_1}{n_1} \right)^{sk} \pi \right\} \quad (5.10)$$

where $\rho(k_y, k_z, t_1)$ is the probability of sampling point (k_y, k_z, t_1) , $sk = 0.5$ is the skew parameter, and n_y , n_z , and n_1 give the number of points in the y , z , and t_1 dimensions, respectively. The skewed squared sine bell density function approximates the combined effect of squared sine bell apodization on a decaying exponential curve. The NUS masks used are shown in Fig 5.1.

5.3.3 Reconstruction

The coil sensitivities S were determined from the non-water-suppressed reference scan by taking the first time point and dividing each channel by the sum-of-squares from all channels in order to determine the proportional contribution of each channel. This ensures that the product $S'S$ in Eq. (5.8) is invertible. The first time point was chosen because each resonance is in phase at the echo time. A region of interest (ROI) was also estimated from the reference sum-of-squares image where the intensities were greater than twice the mean.

As the v optimization in Eq. (5.7) is equivalent to the u optimization in Eq. (5.3), it can similarly be solved for the entire v matrix at once; the $\ell_{2,1}$ -norm is a point-by-point method regardless of the dimensionality of the matrix. Prior to solving this problem, the signal from each channel should be scaled to the same noise level so that a single value of σ is appropriate for the entire dataset. If data are not properly scaled, the reconstruction will be closest to the highest intensity coil regardless of its SNR, whereas a proper reconstruction should favor the highest SNR coils at each point.

This method is also more computationally intensive than solving Eq. (5.3)




Figure 5.1: 8x (a), 12x (b), and 16x (c) nonuniform sampling masks over the (k_y, k_z, t_1) volume used to acquire accelerated EP-COSI data. White squares indicate sampled points. Each mask follows the sampling density given in Eq. (5.10).

because the number of points in the matrix v is larger than the number of points in u by a factor of the number of channels n_c . In particular, each temporary matrix and the large sparse matrix $G \in \mathbb{R}^{nn_c \times Nn_c}$ require much more storage, and computations with them are slower. But since the norm is point by point only, each coil image can be computed separately and concatenated for the final v matrix, similar to the previous sum-of-square techniques. In this case though, the coils were combined after deconvolving the complex coil sensitivities. Since the coil sensitivities were derived from a reference water scan here, the deconvolution is similar to the method of Klose’s eddy current correction [89] applied to CSI data [140], and it is possible to deconvolve with the temporal phases of the reference over all the time points instead of just using the first time point. However, there is strong lipid signal in the calf muscle even without water suppression, so the temporal phase over each time point is not reflective of the water signal only, and the lipid signal must first be removed to apply that correction method. Therefore, we used only the first time point phase for correction. Generally, coil sensitivities can be estimated in other ways with faster imaging-based sequences as well.

Coefficients were grouped in the spectral plane (F_2, F_1) with 50% overlap between adjacent groups in each direction and with each group consisting of 8×4 points [27]. This grouping strategy is motivated by the sparsity of the 2D COSY spectra except for the presence of large, dominant peaks and is illustrated in Fig 5.2.

For comparison, undersampled data were also reconstructed with a similar self-sparse SENSE-based ℓ_1 -norm minimization

$$\min_u \|Su\|_1 \quad \text{s.t.} \quad \|R\mathcal{F}Su - f\|_2^2 < \sigma^2 \quad (5.11)$$

This problem was also solved with the split Bregman algorithm using similar auxiliary variable substitutions as in Eq. (5.6). The regularization parameters were set to $\mu = 1$ for both algorithms, $\lambda = 1/2$ for the CS problem, and $\lambda = 1/(2 \cdot$

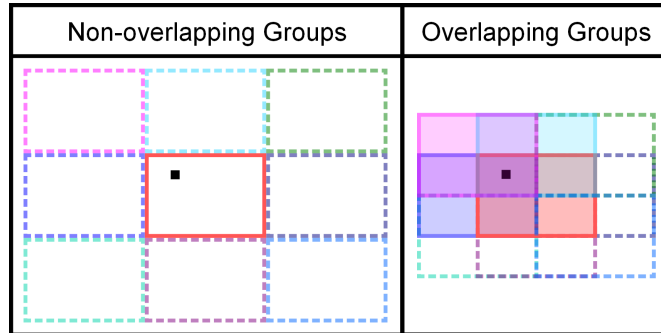


Figure 5.2: Illustration of the difference between two dimensional groups that are non-overlapping and those with 50% overlap. The overlapping groups were used to combine coefficients in the spectral plane (F_2, F_1) .

32) for the GS problem, where μ is the parameter that relaxes the data consistency constraint and λ is the parameter that relaxes the variable substitution constraint. The solution from the split Bregman algorithm is relatively insensitive to these parameter choices [66], but the time to convergence can vary. The CS problem had a maximum of 100 outer Bregman iterations, the GS problem had a maximum of 50 outer iterations, and each problem had 15 inner Bregman iterations per outer iteration. The reconstructions were performed on a 64-bit octa core 3.1 GHz Intel Xeon E5-2687W processor with 128 GB RAM.

5.3.4 Coil compression

Whether the multichannel data are reconstructed coil-by-coil or as one large matrix, the reconstruction time is longer than for a single coil. For the coil-by-coil reconstruction, this is basically linear with the number of coils. For the single matrix reconstruction, the speed depends more on how well the software and processors handle large data matrices, but we found that it becomes impractically slow for the full 15-channel data here. In order to reduce the dimensionality of the problem and speed up the reconstruction, the 15-channel coil array was compressed into five virtual coils using a custom hybrid method based on [25] and

[182].

In [25], the coil transformation is expressed as

$$f'_k = Af_k \quad (5.12)$$

where f_k is a $n_c \times 1$ vector of the real channel elements for each k-space point k , A is a $n_{c_{vir}} \times n_c$ matrix that takes a linear combination of the original coils, f'_k is a $n_{c_{vir}} \times 1$ vector of the truncated virtual coil elements, and $n_{c_{vir}} \leq n_c$ is the truncated number of virtual coils. Matrix A is calculated to maximize the SNR in a region of interest using the SENSE formalism. It is given by

$$A = CU'T \quad (5.13)$$

where $C = (I_{n_{c_{vir}}} | 0)$ is the $n_{c_{vir}}$ row selector matrix, U is the unitary matrix from the singular value decomposition of a matrix P that is given by $\sum_{p \in ROI} S_p S_p^{-1}$, and T transforms the noise covariance between coils into the identity (and must also be applied to S). Once A is determined, it is applied to the raw data and the original sensitivity maps, as the linear transformation works the same in k-space or spatial coordinates and is applied over all the spectral points. This method allows SNR optimization over an explicit region of interest when the sensitivity maps are available and is therefore well-suited for the relatively sparse volumes of interest (VOI) here.

Zhang *et. al.* used a data-based coil compression method that does not rely on explicit coil sensitivities for 3D imaging [182]. There, they took advantage of the fact that at least one spatial dimension is fully sampled and refer to the method as geometric coil compression. For the fully sampled dimension (x in this paper), the raw k-space data can be transformed *before* reconstruction, and a coil compression matrix can be created with different linear combinations at each x location. This permits the use of far fewer virtual coils than if a single linear combination was required for the entire 3D volume. (As described in [178], the

x and F_2 dimensions are Fourier transformed once before reconstruction to save time transforming back and forth during each iteration in the reconstruction.)


Since we already require explicit coil sensitivity maps and have a fully sampled spatial dimension, our method basically uses the transformation described by Buehrer *et. al.* [25] in Eq. (5.12) with an A that is x -dependent

$$f(x)'_k = A(x)f(x)_k \quad (5.14)$$

Figure 5.3 shows the beneficial effect of using Eq. (5.14) as opposed to Eq. (5.12) to create the virtual coils. Figure 5.3b has only five coils with intense signal while Fig 5.3c has 11-12 significant coils compared to the original 15 in Fig 5.3a.

5.4 Results

Figure 5.4 shows the 3D localization and metabolite maps from an 8x undersampled scan of a diabetic calf taken with a single channel extremity coil. Metabolite maps were computed by integrating over the regions around the creatine/fat peak at 3.0 ppm, the creatine peak(s) at 3.9 ppm, the lower extra- and intramyocellular lipid (E/IMCL) cross peaks, and the upper E/IMCL cross peaks, respectively. Within each metabolite grouping, the minimum energy reconstruction (missing data points filled in with zeros), the CS reconstruction (using Eq. (5.11)), and the GS reconstruction (using Eq. (5.6)) are shown. The minimum energy reconstruction is scaled by the square root of the undersampling factor to produce similar intensities to the other reconstructions. Note the aliasing in both phase-encoding directions in the minimum energy reconstruction. The CS and GS maps look very similar, with the GS maps being slightly noisier. Figure 5.5 shows select COSY spectra from the tibial bone marrow and the soleus and tibialis anterior muscles. From the full COSY spectra, the t_1 aliasing is apparent. The CS reconstructions provide good denoising and remove the aliasing of many of the peaks, but the aliasing due to the two dominant lipid peaks and residual



(a)



(b)



(c)

Figure 5.3: Magnitude of sensitivity maps from 15-channel data taken directly from a water reference scan (a), after using an x -dependent geometric transformation matrix $A(x)$ (Eq. (5.14)) (b), and after coil compression using a single transformation matrix A (Eq. (5.12)).




Figure 5.4: Metabolite maps from an 8x undersampled scan of a diabetic calf. Shown are the maps from the creatine and fat peak at 3.0 ppm (Cr3.0), the creatine resonance at 3.9 ppm (Cr3.9), and the lower (Low) and upper (Upp) IMCL and EMCL lipid cross peaks. Within each metabolite group, the minimum energy reconstruction is shown with the CS and GS reconstructions, respectively. Minimum energy reconstructions are scaled by the square root of the undersampling factor. Localization is shown on the left.




Figure 5.5: COSY spectra from the highlighted voxels in a central slice (a) of the same scan as Fig 5.4. Yellow spectra (b,c) are from the tibial marrow, red spectra (d,e,f) are from the soleus muscle, and blue spectra (g,h,i) are from the tibialis anterior muscle. The first column shows raw spectra with the minimum energy reconstruction (d,g), the second column shows CS-reconstructed spectra (b,e,h), and the third column shows GS-reconstructed spectra (c,f,i). Contour levels are the same for each reconstruction for a given voxel with the minimum energy reconstructions being scaled by the square root of the undersampling factor.

water remains. This aliasing nearly completely obscures the lower E/IMCL cross peaks, as only the GS reconstruction faithfully shows the splitting of the EMCL and IMCL components in the lower cross peaks in the soleus and tibialis anterior. Each reconstruction faithfully reconstructs the spatial locations as can be seen by looking at the spectrum around the 3.9 ppm diagonal, as creatine is absent in the marrow, a strong singlet in the soleus, and a split doublet in the tibialis anterior [90]. For the reconstruction shown in Fig 5.4 and 5.5, the GS converged after only five outer iterations in about 15 minutes, while the CS reconstruction did not fully converge in 100 outer iterations and took almost an hour.

Figures 5.6 and 5.7 show metabolite maps and full spectra from a 16x prospectively undersampled scan of a healthy adult. Aliasing is even more prevalent with the higher degree undersampling. Again, GS reconstruction appears slightly noisier but has a much greater reduction in aliasing compared to CS. In all cases, the spectra are not quite as clean as the 8x undersampled, with some smaller peaks and the separation between EMCL and IMCL being difficult to resolve. Here, GS converged after 16 outer iterations in about 45 minutes, while CS again took 100 outer iterations and just under an hour.

Contour plots of the creatine signal at 3.9 ppm are shown for single slice MEEP-COSI in Fig 5.8b and the three central slices of a 12x undersampled 3D EP-COSI in Fig 5.8c, 5.8d, and 5.8e. The MEEP-COSI reconstruction used all 15 acquired channels, while the 3D EP-COSI reconstruction used only the first five virtual channels. On this dataset, the CS reconstruction took an average of 73 iterations and 47 minutes per coil (almost 4 hours total), and the GS reconstruction took an average of 15 iterations and 54 minutes per coil (about 4.5 hours total). Each shows similar localization with the single, strong peak in the soleus and gastrocnemius muscles, a split doublet in the tibialis anterior, and little creatine signal in the marrow. The MEEP-COSI scan has somewhat higher SNR and a cleaner profile, but the thinner slices of the 3D EP-COSI better show the location




Figure 5.6: Metabolite maps from a 16x undersampled scan. Shown are the maps from the creatine and fat peak at 3.0 ppm (Cr3.0), the creatine resonance at 3.9 ppm (Cr3.9), and the lower (Low) and upper (Upp) IMCL and EMCL lipid cross peaks. Within each metabolite group, the minimum energy reconstruction is shown with the ℓ_1 and GS reconstructions. Minimum energy reconstructions are scaled by the square root of the undersampling factor. Localization is shown on the left.




Figure 5.7: COSY spectra from the highlighted voxels shown in Fig 5.6. Yellow spectra (a,b,c) are from the tibial marrow, red spectra (d,e,f) are from the soleus muscle, and blue spectra (g,h,i) are from the tibialis anterior muscle. The first column shows raw spectra with the minimum energy reconstruction (a,d,g), the second column shows CS-reconstructed spectra (b,e,h), and the third column shows GS-reconstructed spectra (c,f,i).

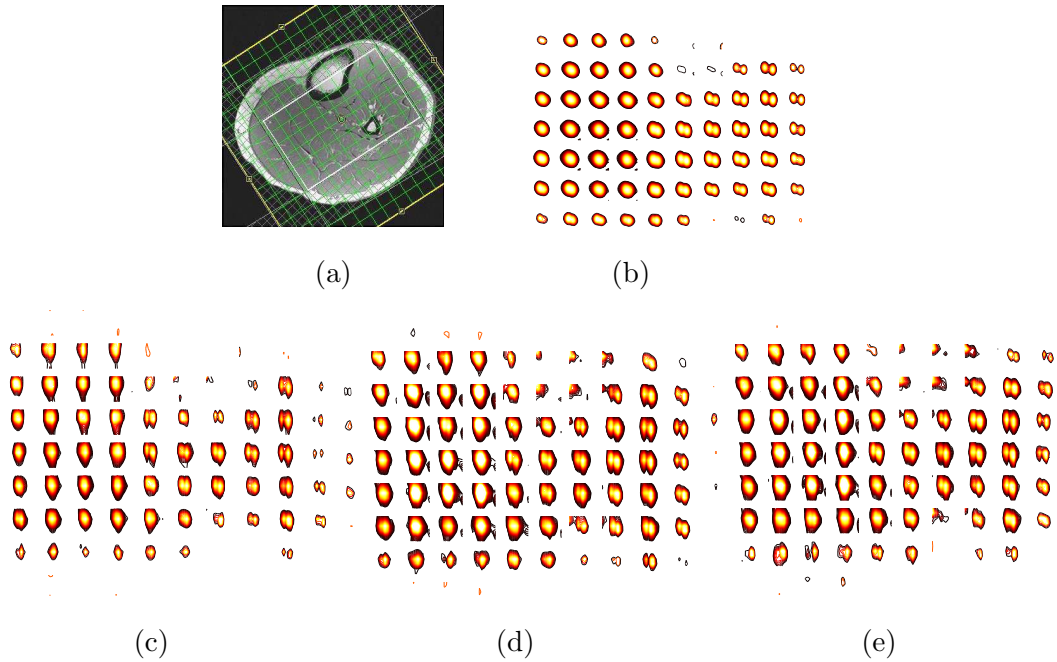


Figure 5.8: Axial localization with 15-channel knee coil (a). Contour plot of the distribution of the peaks at 3.9 ppm for a single slice MEEP-COSI scan (b) and the central three slices from 12x undersampled 3D EP-COSI (c-e) in the same healthy volunteer. The undersampled 3D EP-COSI scan took 14 minutes, and the MEEP-COSI scan took 13 minutes.

of the fibular marrow at the bottom right of the VOI.

5.5 Discussion

Comparing the contour plots in Fig 5.8 suggests that reconstructed data have some residual in plane signal bleed compared to the fully sampled single slice. This is likely due to the small amount of blurring caused by compressed sensing reconstruction, as it is also seen slightly in the indirect spectral dimension. It may also be due to some residual aliasing, but the COSY spectra in Fig 5.5 and 5.7 suggest that spectral aliasing is well cleaned up by the GS reconstruction.

Although the additional computation of grouping and taking the ℓ_2 -norm over each group prolongs each iteration of GS, the algorithm converged in each example and did so with far fewer iterations than the CS method. The number of outer iterations in the CS problem was chosen such that most reconstructions would converge or be close to convergence and that each coil reconstruction would take under an hour. Nevertheless, when the number of outer iterations was not restricted, some of the CS reconstructions took *much* longer before finally converging, yet visual inspection of the data did not show much, if any, improvement beyond the first 100 iterations. Therefore, we feel the limit of 100 outer iterations is justified and is already more than in previous implementations [178, 27]. With those limits, each method seems to take about 45-55 minutes per coil, though some are significantly shorter. We have favored using constant values of the regularization parameters for ease of implementation and reproducibility, but a more adaptive choice could affect the convergence rates of the algorithms and the total reconstruction time. Choice of regularization parameters is an active area of research in compressed sensing [74, 135, 137].

Previous work with 3D-localized J-resolved spectroscopic imaging suggests that minimizing the total variation (TV) performs somewhat better than the CS reconstruction presented here [178]. The difference was not overwhelming, though, and work with single-slice localized EP-COSY shows GS performing much better than either CS or TV [27]. Incorporating a third spatial dimension increases the self-sparsity of the data relative to single slice, so this suggests that applying a sparsifying transform, such as TV, reconstruction to the correlated spectroscopic imaging problem here would not yield better results than GS.

Transforming the coil basis before reconstruction does not result in a loss of information since the aggregate signal over all the coils is still the same, and truncating the number of coils results in minimal loss of information as shown in Fig 5.3. For the coil-by-coil processing done here, there is a direct proportionality

between the number of coils and the total reconstruction time, so compressing the basis from 15 to five channels reduced the reconstruction time by a factor of three. However, the individual sparsity of the coil images is different in the virtual coil basis. In particular, the first virtual channel is less sparse than any of the subsequent ones or any in the original basis, but it is more sparse than the final image, so there should be no question as to the appropriateness of the self-sparsity assumption. This is most likely why the reconstruction of the first coil in a multichannel scan took much longer than the others with GS. Curiously, though, there did not seem to be much of an effect for the CS reconstructions.

In the reconstruction method applied here, grouping was done across the two dimensions of the spectral plane as in [27] to take advantage of the bunching of large coefficients in a highly sparse COSY spectrum. However, grouping does not need to be restricted to the spectral plane and could take advantage of the proximity of large coefficients in the spatial dimensions as well. The algorithm simply requires the number of points in each group to be the same. Joint sparsity [51, 176] across the channels is a form of group sparsity and could also be applied here, except that in the joint sparse case, the channels can no longer be reconstructed one at a time, so the computational demands are much greater as previously discussed.

The EP-COSI and MEEP-COSI sequences used here both used a PRESS volumetric slice selection module in which each of three pulses selects an orthogonal slice [20], and the intersection of the slices is the VOI. This method results in a clean, box-like profile in each dimension and lends itself to ROI definition in the coil compression matrix of Eq. (5.13). A shorter version of the sequence similar to the original COSY can be constructed using two 90° pulses that select a single slice only. The two-pulse sequence would have a shorter minimum echo time and reduced SAR, but excitation cannot be limited to a VOI. At 3T, the SAR was not much of a concern. Even using the three-pulse sequence, a small amount of

the marrow was included. Although the calf muscles have a large amount of lipid signal, the marrow has more, which results in signal bleed as seen in Fig 5.4 and 5.6. Also, many muscle spectroscopy studies look at overweight and obese patients whose calves often have a large amount of subcutaneous EMCL. Unlike techniques in slice-based spectroscopic imaging in brain that use inversion recovery (IR) to selectively suppress the lipid signal [162], the lipid signal is often of primary interest in the muscle, and it has been shown that the IMCL peak has longer T2 and is better resolved at long TE [139]. Outer volume suppression is an alternative that can be used to ensure proper VOI excitation [156] but is limited by the number of saturation bands that can be applied.

5.6 Conclusion

A 3D correlated spectroscopic imaging sequence is presented here that is achievable only by highly undersampling the volume of the two phase-encoding dimensions and one indirect spectral dimension. Reconstruction of the data using group sparsity minimizes the mixed $\ell_{2,1}$ -norm of the data and takes advantage of the proximity of high amplitude coefficients found in spectroscopy peaks. Group sparsity performs much better than a standard ℓ_1 -norm minimization algorithm without significant time penalty.

CHAPTER 6

Correlated Spectroscopic Imaging using Concentrically Circular k-space Trajectories

6.1 Introduction

Chemical shift imaging (CSI) [24] (or magnetic resonance spectroscopic imaging [112]) typically achieves spatial encoding by a combination of slice selective RF pulses and phase encoding gradients. As phase encoded dimensions have to be incrementally acquired in CSI, scan times scale quickly with desired resolution. Fast spectroscopic imaging (SI) sequences utilize gradients during readout to repeatedly acquire an arc in k-space [107]. This allows the interleaved collection of one of the spatial dimensions with the direct spectral dimension and cuts the scan time by an order of magnitude. The original fast SI implementations covered k-space with Cartesian sampling by repeatedly acquiring a single line during readout [127, 126] and are referred to here as echo planar spectroscopic imaging (EPSI). Methods tracing out other non-Cartesian k-space arcs, such as spiral [2], circular [62], and radial [138], have also been implemented. The benefits and drawbacks of each will be briefly reviewed.

EPSI uses a bipolar, trapezoidal gradient readout train, which is the easiest to implement both on the scanner and in processing where Fourier transformation can be applied directly over the k-space dimensions. However, this simple processing requires sampling only during the gradient plateaus, which means that the ramp-up and ramp-down times of each gradient are dead time. Ramp sampling

can be applied [120], but processing then requires the additional complexity of the non-Cartesian methods. Minimizing the spectral dwell time to achieve adequate spectral bandwidth (SBW) requires shorter ramp times and higher amplitude gradients that can lead to eddy current distortions and peripheral nerve stimulation, or simply gradient hardware limitations.

Like spiral acquisitions in imaging, the non-Cartesian trajectories in Spiral-SI are able to traverse k-space more quickly and efficiently than Cartesian ones and can cover all of k-space in a single acquisition [2]. Short echo times are also achievable when no prephasing gradient is required. While there is an inherent oversampling of high SNR central k-space, in order to achieve adequate sampling density in outer k-space [151], longer spirals must be used, thereby decreasing the SBW. Multiple identical applications with time-shifted echo times then can be used to increase the effective SBW at the expense of overall scan time in the process of temporal interleaving [110]. Once k-space data has been collected, it must be interpolated onto a Cartesian grid before Fourier transformation by gridding [84].

Radial-SI is similar to Spiral-SI in that it can achieve very short echo times with inherent oversampling of central k-space. It also requires gridding to Cartesian coordinates. Radial trajectories use similar gradient waveforms to EPSI and can achieve similar SBW. However, repeated gradients must be played over multiple gradient directions, which can easily reach hardware limitations when scanning oblique slices. Nevertheless, this is less important in 3D imaging, and most 2D SI applications use slices corresponding to the physical gradient axes. Radial sampling schemes also generally require more time than Cartesian ones to fully satisfy the Nyquist criteria, but artifacts from undersampled data are more benign and do not necessarily need to be corrected [124].

Spectroscopic imaging using concentrically circular echo planar trajectories (SI-CONCEPT) is similar to the other non-Cartesian sequences in its inherent

oversampling of central k-space, additional processing requirement of gridding, and potential to run into hardware limitations for oblique slices. Compared to EPSI, SI-CONCEPT samples twice as efficiently by sampling from all four k-space quadrants. Circular k-space sampling is also more robust to gradient timing errors and eddy currents than either spiral or Cartesian sampling. These types of errors lead to spatial blurring in spiral and spatial/spectral ghosting in Cartesian trajectories due to the temporal asymmetry between the forward arc and the negative rewind arc. (Flyback EPSI only acquires during one gradient lobe and therefore avoids ghosting artifacts at the expense of a lower SBW [42].) By contrast, these errors in circular k-space encoding manifest as image rotations [179].

There are also other subtle benefits of SI-CONCEPT that allows it to achieve high SBW more efficiently. Because each circle starts and ends at the same location in k-space, there is no need to rewind the gradient. This rewinding can be in the form of a bipolar gradient that continues to sample but increases the spectral dwell time and must be time reversed in processing [127, 138], or it can be a fast, high amplitude blip that is dead time [2, 42]. Either way, the lack of a rewinding gradient allows SI-CONCEPT to sample temporally more efficiently than any of the other sequences.

These benefits of circular sampling all carry over to multidimensional spectroscopic imaging. In this paper, concentric circular readout is applied to the localized correlated spectroscopy (L-COSY) [171] excitation scheme in which the final slice selective 180° refocusing pulse is replaced by a slice selective 90° refocusing/coherence transfer pulse, and the spacing between the second and third pulse is incremented. This correlated spectroscopic imaging using concentrically circular echo planar trajectories (COSI-CONCEPT) sequence is illustrated schematically in Figure 6.1. The rate of sampling in the direct spectral dimension t_2 is drawn as dots in the RF axis, showing that the spectral dwell time is the same as the




Figure 6.1: COSI-CONCEPT schematic showing the L-COSY localization with indirect spectral increment t_1 and sinusoidal readout gradients. The incremented gradient in the y direction plays the role of a typical phase encoding gradient so that readout does not begin at the origin of k-space. Spoilers (not shown) are placed symmetrically around the second and third pulses to remove signal due to unwanted coherences.

period of oscillation of the gradient waveforms.

6.2 Methods

6.2.1 Sequence design

The circular k-space trajectory is shown in Figure 6.2a and broken up into its sinusoidal k_x and k_y components in Fig 6.2b. Point b represents the start of acquisition, which occurs at TE. Therefore, the k-space trajectories during readout are given by

$$\begin{aligned}
 k_x(t) &= -k_n \cdot \sin \left[\frac{2\pi}{T}(t - \text{TE}) \right] \\
 k_y(t) &= +k_n \cdot \cos \left[\frac{2\pi}{T}(t - \text{TE}) \right]
 \end{aligned} \tag{6.1}$$

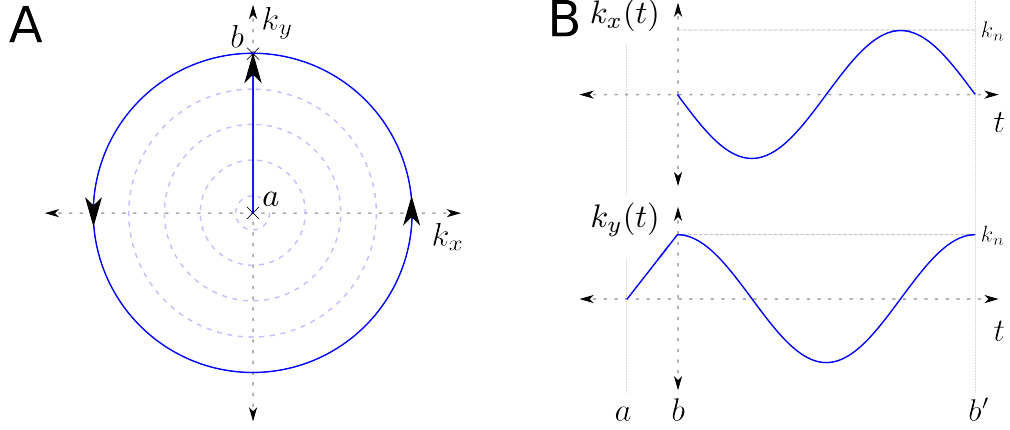


Figure 6.2: **(a)** Example of a concentric k -space trajectory, starting in the origin and traversing from $a \rightarrow b$, and then counter-clockwise around the circle from $b \rightarrow b'$. **(b)** The individual $k_x(t)$ and $k_y(t)$ trajectories as shown in (a), showing the individual steps from $a \rightarrow b$ as well as $b \rightarrow b'$.

where t is the time from excitation, T is the spectral dwell time to go from $b \rightarrow b'$, and k_n is the k -space radius of the n th circle.

The gradient waveforms can be calculated using the partial derivative relationship

$$\vec{G}(t) = \frac{2\pi}{\gamma} \frac{\partial \vec{k}(t)}{\partial t} \quad (6.2)$$

where γ is the gyromagnetic ratio [100]. Putting Eq (6.1) into Eq (6.2) yields

$$\begin{aligned} G_x(t) &= -\frac{4\pi^2 k_n}{\gamma T} \cdot \cos \left[\frac{2\pi}{T} (t - \text{TE}) \right] \\ G_y(t) &= -\frac{4\pi^2 k_n}{\gamma T} \cdot \sin \left[\frac{2\pi}{T} (t - \text{TE}) \right]. \end{aligned} \quad (6.3)$$

There is some flexibility in how to choose the k -space radii, as k_n can be written as

$$k_n = (n + f) \Delta k \quad (6.4)$$

where $\Delta k = \text{FOV}^{-1}$ and $f \in [0, 1]$. Choosing $f = 0$ ensures acquisition of the most central k -space with the highest signal. However, as discussed in the next section, density compensation when gridding is proportional to the distance from

the origin, so a 0 radius circle would be given a weighting of 0. As a result, the ring for $n = 0$ needs to have a special correctional factor to avoid nulling the data collected at the origin [92]. Using $f > 0$ eliminates the need for the correction factor in the gridding and also allows for greater k-space coverage with the larger maximum radius. However, the center of k-space is not directly sampled and is only interpolated. Here, we have chosen to use $f = 1/2$.

Note from Equation (6.3) that $G_x(\text{TE}) \neq 0$ which means that the gradient must be ramped up before TE. This ramp creates an additional gradient moment that must be precompensated, which is done through a slight modification of the last spoiler gradient after the second 90° pulse. This is distinct from the “phase encoding” gradient of Fig 6.1 that takes k-space from $a \rightarrow b$ in Fig 6.2 and does impart a desirable gradient moment. Both gradients depend on the radius of the circle in k-space and scale appropriately.

Water suppression was applied using a three pulse WET sequence before signal excitation [117] but is not shown in Fig 6.1 for simplicity.

6.2.2 Gridding

As mentioned in the previous section, non-Cartesian k-space data must be gridded to a Cartesian plane before application of the Fourier transform, and this is no different for COSI-CONCEPT.

Gridding consists of convolution of the non-Cartesian data with an interpolating kernel and resampling along the standard Cartesian trajectory. Convolution in the Fourier domain is equivalent to multiplication in the image domain, so the ideal gridding kernel is a *sinc* function whose transform is a box function. However, *sinc* has infinite extent and is therefore replaced by a compactly supported function; here we use a Kaiser-Bessel kernel given by

$$KB(k) = \frac{N_g}{W} I_0 \left(\beta \sqrt{1 - (2N_g k/W)^2} \right) \square (2N_g k/W) \quad (6.5)$$

where N_g is the number of grid points, W is the kernel width (in grid units), β is a factor that shapes the kernel, I_0 is the zero-order modified Bessel function of the first kind, and \square represents the box function [84].

A kernel width $W = 3$ was used. The number of grid points was defined to be twice the number of points of the image in each dimension to reduce aliasing artifacts in a process referred to as *overgridding*. The final image then has twice the desired field of view (FOV) and was cropped to the appropriate size. To compensate for the convolution, the image must also be divided by the Fourier transform of Eq. (6.5). Beatty *et. al.* [14] gives the optimal value of β in terms of the kernel width and the overgridding factor α as

$$\beta = \pi \sqrt{\frac{W^2}{\alpha^2} (\alpha - 1/2)^2 - 0.8} \quad (6.6)$$

For a Cartesian grid, each sampled point covers the same area in k -space, but for non-Cartesian trajectories, the sampling density is not the same at each point. When the gradients are larger, k -space is traversed quicker and each sampled point represents a larger area. Without compensating for different densities, the convolution would give relatively higher values where the points are more densely sampled. Therefore, the data must be preweighted by the area it covers in k -space. For circular sampling, the preweighting is given by the radius of the ring. Once the data has been resampled onto a Cartesian plane and density compensated, standard FT processing can be performed.

The resampling part of the gridding process can be performed as a matrix multiplication

$$k_{car} = Gk_{nonc} \quad (6.7)$$

However, in imaging, the resampling matrix G is often too large to store. For example, in gridding to a 256×256 image from a radial acquisition of 402 *spokes* \times 256 *samples per spoke*, G would be 262144×102912 including overgridding. Therefore, gridding is typically computed by performing the convolution point-by-point

and summing up the results. For faster computation, a look-up-table is precalculated, and the value of the convolution kernel is approximated from the table. On the other hand, in spectroscopic imaging, the matrix sizes are much smaller. For a typical COSI-CONCEPT acquisition, 8 *rings* \times 64 *points per ring* are gridded onto a 32×32 image matrix, resulting in a G that is a much more manageable 1024×512 . Therefore, no approximate kernel values need to be taken, and the entire gridding matrix can be precomputed. This is especially useful as G must be applied to each time point of the spectroscopic imaging sequence.

6.2.3 Experimental

COSI-CONCEPT was compared with its Cartesian correlated spectroscopic imaging analog (EP-COSI) [99] in human calf muscle at 3T. Each sequence had the following parameters: TR/minimum TE = 1500/30 ms, FOV = 16×16 cm² with slice thickness 2 cm, 512 direct and 50 indirect spectral points, $\Delta t_1 = 0.8$ ms yielding $SBW_1 = 1250$ Hz. EP-COSI had 16 phase encodings. With 16 points per readout and a 50 kHz readout BW, $SBW_2 = 1190$ Hz. One average was taken. COSI-CONCEPT acquired 8 circles with 64 points per circle to ensure azimuthal oversampling. (The polar Nyquist criteria for circular sampling requires $N_\theta \geq 2\pi N_{rings}$ [151].) The readout BW was 80kHz, yielding a 1250 Hz SBW_2 . Two averages were taken to produce time equivalent scans of 20 minutes. Pixel resolution was 1×1 cm² in plane for each sequence.

For each sequence, a reference scan was collected to use for zero and first order phase correction due to eddy currents and complex coil sensitivity profiles [89, 140]. The reference scan had identical parameters to each SI scan, except that water suppression was turned off and the number of indirect spectral points was only 1. Each reference scan took only 30 s including prep scans, so the added time was minimal. Since the gradients and coils are identical with the reference scan, it is implicitly assumed that the phase errors are as well. This assumption

is not strictly true when applied over the t_1 increments though, as the sequence timing is slightly different. However, it is reasonable to believe that most of the eddy current distortions are due to the echo planar readout gradients and not due to the spoilers, in which case the reference phases are applicable.

In postprocessing, temporal data was zero filled from 512×50 to 1024×100 before apodization with a skewed squared sine bell in t_2 (skew parameter 0.5) and a squared sine bell in t_1 (skew parameter 1). No spatial zero filling or apodization was applied.

6.3 Results

As the EP-COSI and COSI-CONCEPT scans were taken sequentially on the same volunteer, they share the same anatomical reference images. Corresponding voxels from the soleus, tibialis anterior (TA), and bone marrow from an EP-COSI scan and a COSI-CONCEPT scan taken sequentially on a single healthy volunteer are shown in Figure 6.3. Note the similarities in the voxels between the two scans and the clear presence of the lipid cross peaks above and below the diagonal between the resonances at 5.5 ppm with those at 2 ppm and 3 ppm.

In particular, the marrow voxels in Fig 6.3a and 6.3d show no discernible creatine signal in the vicinity of 3.9 ppm, and the large signal at 3.0 ppm is, in this case, due to lipids as evidenced by its coupling connectivity with the lipid peak around 5.5 ppm. The soleus muscle voxels in Fig 6.3c and 6.3f show the splitting of the intra/extramyocellular lipids (I/EMCL) that can be seen with the cross peaks from 5.3-5.5 ppm [134]. This splitting is not resolvable in the corresponding diagonal peak.

Due to coherent muscle fiber orientation in the tibialis anterior in human calf, there exists residual dipolar coupling between creatine peaks that manifests as a doublet splitting of the peak at 3.9 ppm [19]. Splitting is not observed at 3.0




Figure 6.3: The first row shows highlighted voxels from the tibial bone marrow (labeled in red in the anatomical T2-weighted MRI insert) **(a)**, tibialis anterior (labeled in yellow) **(b)**, and soleus (labeled in blue) **(c)**, from a COSI-CONCEPT scan of a diabetic calf. The second row shows the corresponding voxels from the EP-COSI scan, **(d)**, **(e)**, and **(f)**. Note that spectra are not all scaled the same.




Figure 6.4: Spatial distribution of the creatine 3.9 ppm peak from a COSI-CONCEPT scan overlaid on top of the anatomic T2-weighted MRI.

ppm in general because of the overlapping lipid peak. This phenomenon is clearly observed in Fig 6.3b and 6.3e.

Figure 6.4 shows the spatial distribution of the creatine 3.9 ppm peak overlaid on the anatomic MRI. The dipolar peak splitting is readily observed on the top right side, while the tibial bone marrow is nearly devoid of any creatine save for a small amount due to partial voluming and the voxel PSF. The soleus muscle has a single creatine 3.9 ppm peak that is also more intense as its signal was not split.

6.4 Discussion

The spectral quality and spatial localization of COSI-CONCEPT compare well to that of EP-COSI, though there are some minor differences. First, it appears from Figure 6.3 that the SNR of the EP-COSI voxels is slightly higher. This is consistent with the 12% difference in SNR efficiency reported in [62] and the 13% difference reported in [85]. The biggest reason for this is that signal acquired during a time-varying readout gradient has higher variance compared to signal read out with a constant or no gradient [105, 125]. Theoretically, based on the formulas in [125], the COSI-CONCEPT sequence used here has about twice the signal variance as the EP-COSI version. Since COSI-CONCEPT traverses k-space twice as quickly, two averages can be taken for a time-equivalent scan, and the SNR efficiency should be about the same. The small added difference could be due to parameter choices in the gridding or the fractional shift of the k-space rings f in Eq (6.4).

The high SNR in EP-COSI is the result of effectively taking two averages per scan by combining the echoes formed during the positive bipolar gradient lobes with those formed during the negative. The negative gradient lobes also serve to rewind k-space, and as previously mentioned, this rewinding also indirectly reduces the SBW. The interlaced Fourier transform [114] can be applied to process

the bipolar echoes as separate time points, doubling the SBW. However, it is highly susceptible to phase errors and results in a large drop in SNR with some residual aliasing [53], so it is more common to combine the echoes as was done here.

For the choice of $f = 1/2$ used here, the circular k-space trajectories are equivalent to radial spokes, as each sampled point is equidistant in the radial direction. That means that reconstruction can be performed using the inverse Radon transform instead of gridding. This method is used in computed tomography (CT), where it is referred to as backprojection. Although it is largely equivalent to gridding in terms of reconstruction quality in MRI in practice, there are some differences in terms of aliasing and the point spread function (PSF) when using backprojection with a filter or undersampling, as discussed in detail in [93]. However, no spatial apodization was used in this paper, and the sampling rate satisfies the polar Nyquist criteria, so the two reconstructions would be expected to be effectively the same.

Additional differences exist between the PSFs of COSI-CONCEPT and EP-COSI due specifically to the sampling. For the parameters used here, the main lobe of COSI-CONCEPT's PSF was about 5% broader, resulting in slightly lower effective spatial resolution. As the EP-COSI-equivalent resolution COSI-CONCEPT scan is definitely not hardware limited, one additional circle could be acquired, thereby reducing the width of the main lobe below that of EP-COSI. On the other hand, EP-COSI is often run at the hardware limits, so acquiring a different resolution requires a nontrivial modifying of a number of parameters. Differences in the PSFs can also be reduced through apodization.

Currently, gradient slew rates are FDA-limited to be 200 mT/m/s and scanner-limited to be 180 mT/m/ms on the particular clinical scanner used here. These slew rate limits are the main bottleneck limiting the achievable SBW in echo planar type SI sequences. For the parameters used here, the EP-COSI sequence had a maximum slew rate over 90% of the allowable limit, whereas the COSI-

CONCEPT sequence had a maximum readout slew rate less than half that while still achieving a slightly higher SBW. (The actual maximum slew rate occurred during the ramps of the spoilers and was closer to that of EP-COSI, but this value is not parameter dependent.) This means that COSI-CONCEPT can achieve much higher resolution or SBW than EP-COSI when pushed to the allowable limit, which is of great importance for high field scanning where the increased peak dispersion requires higher SBW to avoid spectral aliasing.

Generally, another benefit of lower slew rates is reduced eddy current distortions. However, the actual reduction in eddy currents is somewhat mitigated by gradient preemphasis that most vendors automatically perform for trapezoidal waveforms. Yet preemphasis is not applied for arbitrary gradient waveforms unless supplied by the user. Nevertheless, the eddy current distortions in COSI-CONCEPT do seem to be less because the central k-space circles use very low amplitude gradients with almost no eddy current effect. By contrast, each line of EP-COSI, whether through the center of k-space or the periphery, has the same high slew rate.

The use of concentric circular readout was applied with L-COSY localization, though it could easily be included with other 2D spectroscopy sequences such as JPRESS [147], L-EXSY [170], or localized TOCSY [8]. Each of these 2D sequences would benefit from reduced eddy current distortions and high SBW. In particular, artifacts due to residual eddy currents can be particularly problematic for the phase-sensitive multiplet resonances even after using Klose's correction method [89] in postprocessing.

Phase encoded CSI scans do not acquire data with a readout gradient playing and are therefore not gradient-limited in terms of SBW. For this reason, nearly all high field SI sequences are CSI-based. Overall scan time limits the achievable resolution, but quality spectra can be obtained in a reasonable amount of time. Correlated CSI with two spectral dimensions, though, is not feasible in vivo with-

out large acceleration, and methods such as parallel imaging [52] or compressed sensing [65, 29] do not offer enough acceleration to rectify that.

6.5 Conclusion

A fast correlated spectroscopic imaging sequence that traverses k-space in concentric circles was presented here and compared with the more common fast sequence that traverse k-space rectilinearly. It was shown that for similar acquisition parameters, the two sequences have about the same SNR efficiency, with COSI-CONCEPT sampling k-space twice as fast. The lower gradient slew rates in circular sampling result in a reduction in eddy current distortion and allows for scanning at much higher spectral bandwidths, thereby increasing the potential applicability of correlated spectroscopic imaging at high field strengths.

6.6 Appendix

6.6.0.1 COSI-CONCEPT slew rates

In COSI-CONCEPT, the slew rates are given by the time derivatives of Eq. (6.3)

$$\begin{aligned} G'_x(t) &= +\frac{8\pi^3 k_n}{\gamma T^2} \cdot \sin\left[\frac{2\pi}{T}(t - \text{TE})\right] \\ G'_y(t) &= -\frac{8\pi^3 k_n}{\gamma T^2} \cdot \cos\left[\frac{2\pi}{T}(t - \text{TE})\right] \end{aligned} \quad (6.8)$$

Consider k-space points that exactly lie on the axes. These points are the same whether k-space is collected rectilinearly or circularly. In particular, if Δk is the spacing between rings, $\text{FOV} = \Delta k^{-1}$. Making use of the relationships $k_n = (n_{rings} + f) \cdot \text{FOV}^{-1}$ and $T^{-1} = \text{SBW}_c$ and Eq. (6.8), the maximum slew rate for COSI-CONCEPT is given by

$$G'_{max,concept} = \frac{8\pi^3(n_{rings} + f) \cdot \text{SBW}_c^2}{\gamma \cdot \text{FOV}} \quad (6.9)$$

where it can be seen to scale with the square of the spectral bandwidth.

In EP-COSI, a repeating bipolar trapezoidal gradient waveform is applied across the readout direction and has 0 slew rate during the flat tops and maximal slew rate during ramping. The gradient amplitude during the flat top is given by

$$G = \frac{2\pi}{\gamma} \frac{\Delta k_{ro}}{\Delta t} \quad (6.10)$$

where Δt is the readout dwell time, and the subscript *ro* refers to the *readout* direction. Calling the time to ramp the gradient between its maximum positive and negative amplitudes t_{ramp} , the maximum slew rate for the EP-COSI sequence is

$$G'_{max,epsi} = \frac{2G}{t_{ramp}} = \frac{4\pi\Delta k_{ro}}{t_{ramp}\gamma\Delta t} \quad (6.11)$$

Here, the spectral bandwidth is given by $SBW_e = [2(n_{ro}\Delta t + t_{ramp})]^{-1}$ which can be used to rewrite Eq. (6.11) as

$$G'_{max,epsi} = \frac{8\pi\Delta k_{ro}}{t_{ramp}\gamma} \left(\frac{n_{ro} \cdot SBW_e}{1 - 2t_{ramp} \cdot SBW_e} \right) \quad (6.12)$$

Taking the ratio of Eq. (6.9) to Eq. (6.12) and assuming the FOV in the EP-COSI readout direction is the same as the FOV in each direction of COSI-CONCEPT, the ratio of the maximum slew rates of COSI-CONCEPT to EP-COSI is found to be

$$\frac{G'_{max,concept}}{G'_{max,epsi}} = \frac{\pi^2(n_{rings} + f)}{n_{ro}} [t_{ramp} \cdot BW (1 - 2t_{ramp} \cdot BW)] \quad (6.13)$$

The term in brackets is parabolic as a function of the ramp time-bandwidth product with maximum value of 1/8 when $t_{ramp} \cdot BW = 1/4$. Typically, the largest ring in COSI-CONCEPT is chosen such that the full extent of k-space coverage is a circle inscribed in that of the square formed from EP-COSI. In that case, the number of rings is half the number of points in the EP-COSI readout direction, $n_{rings} = n_{ro}/2$, which highlights the increased efficiency of circular sampling relative to rectilinear. Therefore, the ratio of maximum slew rates satisfies

$$\frac{G'_{max,concept}}{G'_{max,epsi}} \leq \frac{\pi^2}{2} \cdot \frac{1}{8} \cdot \frac{(n_{rings} + f)}{n_{rings}} \quad (6.14)$$

which is < 1 when $(n_{rings} + f)/n_{rings} < 16/\pi^2$. As the number of rings increases, $(n_{rings} + f)/n_{rings}$ decreases asymptotically to 1. Since $f \leq 1$, that ratio will be less than 1.5 for $n_{rings} > 1$, and the slew rate ratio in Eq. (6.14) will be less than 1. Therefore, COSI-CONCEPT has lower slew rates than EP-COSI for any realistic sampling scenario.

CHAPTER 7

Improved J-resolved Spectroscopic Imaging with Adiabatic Refocusing Pulses and Concentrically Circular Echo Planar k-space Trajectories

7.1 Introduction

In the early days of continuous wave (CW) NMR [17, 133], spectra were collected from nuclei in a constantly oscillating RF field (B_1) by varying the amplitude of the polarizing B_0 field. Because the B_0 field varied slowly, the bulk magnetization would remain collinear with the net effective magnetic field. This method was able to achieve excitation over a wide band of resonant frequencies while being robust to inhomogeneity in the RF field. However, Fourier transform (FT) NMR largely replaced CW NMR due to its orders of magnitude improvement in sensitivity that allowed for much quicker experiments [55]. FT NMR uses a static B_0 field and applies short, high intensity RF pulses to produce a *free induction decay*, and Fourier transform of the free induction decay gives the spectrum. However, the effect of the RF pulse on the magnetization is highly dependent on the B_1 field strength.

Adiabatic RF pulses [155, 13, 76, 15, 172, 173, 64, 45] use a frequency or phase sweep of the RF field to accomplish adiabatic fast passage of the bulk magnetization. (Fast, in this case, is with respect to the relaxation constants T_1 and T_2 .) Just as in the CW NMR experiment, the magnetization follows the

effective magnetic field. The pulses act as part of an FT NMR experiment with the added CW benefits of broad band excitation and insensitivity to inhomogeneous RF fields. They can be thought of as the limit of a continuously varying composite pulse [95] with far greater tolerance for B_1 variation. Fundamentally, these pulses accomplish one of two tasks: either excitation ($M_z \rightarrow M_x$) or inversion ($M_z \rightarrow -M_z$). Those that excite are known as adiabatic half passage (AHP) pulses, and those that invert are known as adiabatic full passage (AFP). The reason for these two primary usages are that adiabatic pulses do not cause plane rotations like typical RF pulses, and the end result is dependent on the initial magnetization. Still, combinations of AHP and time-reversed AHP pulses can be used to achieve plane rotation [172, 164].

AHP and adiabatic plane rotation pulse are not typically used for spatial selectivity due to off-resonance effects brought about by the gradient [86]. AFP pulses, on the other hand, can be used efficiently for slice selective inversion [155], but when attempting to use the AFP pulse for slice selective refocusing, a nonlinear phase dispersion results through the slice, leading to destructive cancellation of the signal. Conolly *et. al.* showed that this phase dispersion can be exactly refocused by application of a second, identical slice-selective AFP pulse [40]. Symmetric, non-adiabatic pulses self-refocus this phase dispersion and, therefore, do not require a paired pulse.

This realization that adiabatic pulses could be used for general refocusing with relative ease opened the door to their use in localized spectroscopy experiments. Fully adiabatic volumetric localization can be achieved following application of a nonselective AHP pulse and three *pairs* of slice selective AFP pulses as in SADLOVE [158] or LASER [63]. It was shown that LASER VOIs had sharper edges than their PRESS [20] counterparts, leading to a reduction in signal bleed, which is especially relevant for voxels near the skull marrow.

Chemical shift displacement error (CSDE) is another problem that confounds

spectroscopy and spectroscopic imaging as it results in shifts in apparent localization of different resonances. CSDE occurs in spectroscopy because a slice-selective gradient spreads frequencies over a spatial direction, and different metabolites have different frequencies. The fractional amount of the error is defined as the ratio of the physical displacement δz to the slice thickness Δz

$$\text{CSDE} = \frac{\delta z}{\Delta z} \quad (7.1)$$

Since the slice thickness is determined by the bandwidth of the pulse (BW) and the strength of the gradient (G)

$$\Delta z = \frac{\text{BW}}{\gamma G}, \quad (7.2)$$

and the physical displacement is dependent on the chemical shift δf between the resonances and also on the gradient strength

$$\delta z = \frac{\delta f}{\gamma G}, \quad (7.3)$$

the CSDE can be written very simply in terms of the frequencies

$$\text{CSDE} = \frac{\delta f}{\text{BW}} \quad (7.4)$$

Equation (7.4) very clearly shows the inverse proportionality between the RF pulse bandwidth and the CSDE that makes using high bandwidth adiabatic pulses appealing.

Nevertheless, because of the seven total pulses required of LASER, the sequence is not applicable for short T2 experiments. Also, the individual adiabatic pulses typically have higher SAR than traditional or SLR-optimized [123] pulses, so LASER is a relatively high SAR sequence compared to its alternatives. Scheenen *et al.* sought to address those issues with LASER by replacing the initial nonselective AHP pulse and the first pair of slice-selective AFP pulses with a single, non-adiabatic slice-selective pulse and calling the sequence semi-LASER (or

sLASER) [150]. sLASER has less SAR and a shorter minimum echo time than LASER. However, the sequence is no longer fully adiabatic.

For coupled metabolites, in addition to the CSDE and unwanted coherences generated by improper flip angles that plague singlet signals, pulse imperfections can also cause a spatially dependent J-evolution [54]. This latter issue is particularly troublesome for localized J-resolved spectroscopy (JPRESS) [147] in which refocusing pulses are expected to affect each coupled partner equally. When both resonances are refocused, J-coupling evolves for the entire TE, leading to the expected spectral splitting in the indirect dimension. However, when the refocusing pulse only affects one of the coupling partners, J-evolution can actually be refocused. The resulting spectrum is then a weighted combination of all the different refocusing possibilities, which makes it lower in sensitivity and exceedingly hard to quantify.

Here, we apply semi-LASER localization to a fast J-resolved spectroscopic imaging sequence that traverses k-space in concentric circles (JRESI-CONCEPT) [62]. Individual spectra and metabolite maps are compared with a PRESS-based JRESI-CONCEPT.

7.2 Methods

The semi-LASER JRESI-CONCEPT sequence schematic is illustrated in Figure 7.1. A 90° hamming-filtered sinc pulse is applied with a slice-selective gradient along the z axis. A phase encoding gradient is applied in y to shift k-space to the appropriate radius. Two consecutive AFP pulses were played with a y gradient with two more that are slice selective in x following. Sinusoidal gradients are applied in x and y to traverse circles in k-space. Though not shown, spoiler gradients are placed symmetrically around each AFP. A precompensatory gradient is included in one of the x spoilers to negate the moment as the x gradient ramps




Figure 7.1: JRESI-CONCEPT pulse sequence schematic with semi-LASER localization. The 90° pulse is non-adiabatic and slice selective. The 180° pulses are all adiabatic and slice selective. The first two select the y -slice, and the last two select the x -slice. Sinusoidal gradients are applied in x and y to traverse k -space circularly. The indirect spectral increment t_1 was placed before the final pulse only for maximum echo sampling. Crusher gradients surrounding each refocusing pulse and a WET water suppression module preceding the localization are not shown, and the sequence timing of the readout gradients is not drawn to scale.

up. The echo delay time t_1 was applied before the last AFP. As opposed to the traditional J-resolved experiment in which the increment is applied symmetrically around the refocusing pulse, here t_1 was only applied before the refocusing pulse as part of a maximum echo sampling scheme [106, 154].

The AFP refocusing pulses used in sLASER JRESI-CONCEPT were identical hyperbolic secant HS1 pulses [155] with bandwidth-time product (BWTP) 25.60. SAR was not really an issue at 3T, so the pulses were applied for the hardware minimum duration (due to B_1 amplitude) of 5.6 ms, resulting in a 4.6 kHz RF bandwidth. The conventional PRESS sequence (and our PRESS-JRESI-CONCEPT) uses optimized “Mao” pulses [109] that have a 6.00 BWTP. These pulses were also applied at their minimum duration of 5.2 ms, resulting in a 1.1 kHz BW.

Each sequence collected eight rings with 64 points per ring at a direct spectral bandwidth of 1250 Hz. For each, the Δt_1 increment was 1 ms, resulting in a reconstructed indirect spectral bandwidth of 500 Hz after frequency-dependent, linear phase correction in postprocessing. Forty t_1 increments were collected with TR = 1.5 s, for a total scan time of eight minutes. The circular k-space points were gridded to a 16×16 Cartesian plane using a Kaiser-Bessel kernel [84]. The field of view was 16×16 cm² in plane with a 1.5 cm slice thickness.

The minimum echo time for the sLASER-based sequence was 37 ms including all the crusher gradients (not shown in Fig 7.1) and additional time delays to balance each echo. The minimum echo time for the PRESS-based sequence was 29 ms. Though the PRESS-based version has two pulses of 5.6 ms each less (and the remaining two pulses are almost half a second shorter), the crusher gradients were slightly larger to ensure adequate dephasing of signal from unwanted coherences. Scans comparing the two sequences at the same TE are shown, as well as scans acquired at the minimum TE for each.

Each sequence localization was preceded by WET water suppression [117] (also




Figure 7.2: Metabolites maps for NAA 2.0 ppm singlet (left), creatine 3.0 ppm (center), and creatine 3.9 ppm (right) acquired with sLASER-JRESI-CONCEPT (top) and PRESS-JRESI-CONCEPT (bottom) acquired at the same echo time ($TE = 37$ ms) in a healthy volunteer. Maps are scaled to be of similar intensities.

not shown in Fig 7.1). All scans were taken on a 3T Trio-TIM with an 8 channel head coil. Coils were combined as a sum-of-squares as no reference water scan for phase correction [89, 140] was acquired. Metabolite maps were found by integrating the appropriate peak volumes.

7.3 Results

Figure 7.2 shows metabolite maps acquired with each localization at the same echo time ($TE = 37$ ms) in a healthy adult volunteer. NAA 2.0 ppm singlet is on the left, creatine 3.0 ppm is in the center, and creatine 3.9 ppm is on the right. The top row was acquired using sLASER localization, while the bottom row used PRESS. Each metabolite's map is scaled differently to give the same overall intensity, but the scaling is the same between localizations. Figure 7.3 shows the J-resolved spectra from each scan for the same voxel in the parietal lobe. Figure 7.4 shows the same thing for a voxel on the edge of the VOI from the occipital

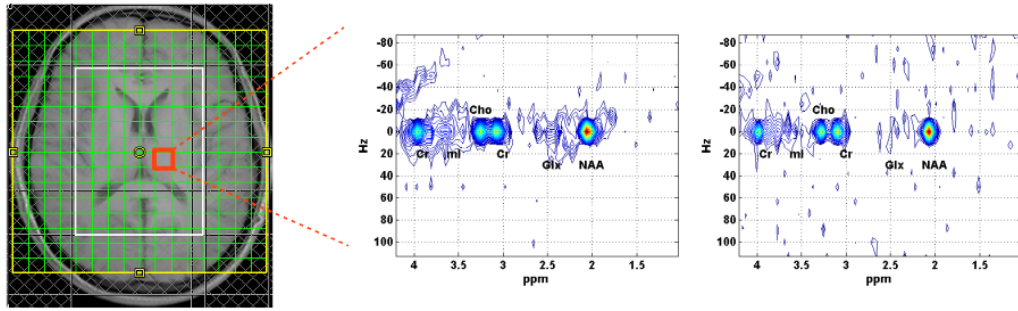


Figure 7.3: J-resolved spectra from the highlighted voxel in the *parietal* lobe with sLASER localization (left) and PRESS localization (right) in the same healthy volunteer as Fig 7.2.

lobe.

Figure 7.5 shows the creatine 3.0 ppm and Glx metabolite maps from a coronal slice of a second healthy volunteer. The left maps were acquired with sLASER localization at $TE = 37$ ms, the middle maps were acquired with PRESS localization at the same echo time, and the right maps were acquired with PRESS localization at its minimum echo time of 29 ms. Figure 7.6 shows a parietal lobe voxel from each of the three scans. The minimum echo time PRESS has more signal for all resonances compared to the longer echo time version, but this is especially true of the coupled metabolites due to the combination of T2 decay and antiphase evolution. Nevertheless, the coupled metabolites still have less signal than the sLASER localization with longer echo time.

7.4 Discussion

Two things are immediately evident from Figure 7.2. The first is that the sLASER-localized JRESI-CONCEPT has higher SNR than the PRESS-localized version. The is most likely due to imperfections in the flip angle of the “Mao” pulse such that spins are not receiving a true 180° refocusing. Furthermore, this would ex-




Figure 7.4: J-resolved spectra from the highlighted voxel in the *occipital* lobe with sLASER localization (left) and PRESS localization (right) from the same scan as Fig 7.3.




Figure 7.5: Metabolite maps from a coronal slice of a second healthy adult showing creatine 3.0 ppm (top) and Glx (bottom) for minimum TE sLASER localization (left), same TE PRESS localization (center), and minimum TE PRESS localization (right). Maps are scaled the same between scans but differently between metabolites.




Figure 7.6: J-resolved spectra from the highlighted voxel in a coronal slice in the healthy volunteer as Fig 7.5 acquired with sLASER localization at minimum TE (left), PRESS localization at the same TE (center), and PRESS localization at its minimum TE (right).

plain why the PRESS-localized version required larger gradient crushers to remove the unwanted coherences. CSDE in the slice direction could also contribute to this and would be exacerbated with the presence of outer volume suppression (OVS) pulses that are often used in volumetric localizations to prevent signal bleed from outside the VOI. The adiabatic HS1 pulses have better slice profiles and therefore do not require OVS, though.

Also clear from Figure 7.2 is the CSDE with the PRESS localization. At 3T, the frequency per ppm is 123.23 Hz, so using Eq. (7.4), we find that the HS1 pulses result in a CSDE of 2.7% per ppm while the “Mao” pulses have a CSDE of 10.7% per ppm. Since the VOI was about 12 cm in the anterior-posterior direction, the CSDE with “Mao” pulse is a little over 1 cm/ppm, which is roughly the size of a voxel. For the HS1 pulses, the CSDE is less than 1/3 of the voxel size. It should be noted that CSDE does also occur in the readout dimension. However, because the readout BW is much larger (on the order of 100 kHz), the error is negligible.

For comparison, the BWTP of the hamming-filtered sinc excitation pulse was 8.75, and it was played for 2.2 ms, resulting in a BW of 3.98 kHz and CSDE of 3.1% per ppm. The left-right VOI was about 8 cm, so the excitation pulse resulted in a CSDE of about 1/4 voxel. The non-adiabatic excitation pulse has much higher BW and lower CSDE than the non-adiabatic “Mao” pulse. In fact,

it is comparable in both to the adiabatic HS1 pulses. This is mainly due to its shorter pulse duration, which is achievable because a 90° flip angle takes less energy than a 180° one. That is why semi-LASER is just as effective as LASER at reducing CSDE in this case. However, there are other advantages to only using adiabatic pulses that semi-LASER localization lacks.

Figure 7.4 also highlights the problem with attempting to look for regional changes in spectroscopic images that have large CSDE. In the sLASER-localized image, all the major resonances are clearly present, but in the PRESS-localized image, NAA is greatly diminished, and Glx is unidentifiable. If this experiment were being done to test for some disease or pathology with the PRESS-based localization, one would associate this voxel with decreased neuronal activity based on the relative NAA. However, the real reason for the low NAA signal is the one voxel shift compared to creatine seen in the maps in Fig 7.2. Fitting programs for metabolite quantitation [129, 153] typically use creatine as an internal standard of reference and would therefore be highly susceptible to this problem.

From a parietal lobe voxel, we see in Fig 7.3 that the intensities of the singlet peaks (NAA, creatine, and choline) are about the same with either localization. On the other hand, coupled metabolites like Glx and myo-inositol (mI) show much higher signal in the sLASER version. This is also clear from Figures 7.5 and 7.6, where the Glx signal from sLASER is much larger than from PRESS even when the latter is acquired at a shorter echo time. The reason behind this is that the additional 180° refocusing pulses act as a Carr-Purcell train [30] and effectively freeze out the antiphase coherence in J-evolution of strongly coupled spin systems [5]. This same idea was used by Hennig *et. al.* where additional (non-adiabatic) refocusing pulses were added to a standard PRESS localization [79]. Therefore, although T2 decay is still greater for the longer echo times with sLASER localization, the signal “decay” associated with antiphase coherences in coupled spin systems is actually much less.

The sLASER localization scheme could be added to any fast spectroscopic imaging sequence that is currently based on PRESS, such as EP-COSY [99]. Yet, there are a number of reasons to prefer the choice of concentrically circular k-space sampling. First, circular sampling is more efficient at covering k-space and requires half the number of phase encodes as rectilinear sampling. Eddy current distortions are also less. Another reason to choose circular sampling is the higher achievable spectral bandwidth [62] that is required for high field imaging. At high fields, though, the CSDE becomes worse according to Eq. (7.4) as the spectral dispersion δf increases. And while it was not as big of an issue here, RF homogeneity is more difficult to achieve at high field strength. For those reasons, the use of sLASER localization with SI-CONCEPT readout shows great promise for high field applications.

7.5 Conclusion

A volumetric localized fast J-resolved spectroscopic imaging using concentrically circular k-space trajectories can be improved upon by replacing the standard 180° refocusing pulses with adiabatic full passage ones. This improves errors due to chemical shift displacement and RF inhomogeneity and is well suited for coupled metabolites despite the longer minimum echo time.

7.6 Appendix

7.6.1 Siemens Pulsetool parameters: *Refgrad* and the BWTP

An RF pulse can be characterized by its bandwidth-time product (BWTP = $BW \cdot T_p$) which is dependent only on the overall pulse definition and not the specific units of measure. Shortening/lengthening the pulse duration will cause the amplitude to scale in order to deliver the same amount of RF energy and the BW

to increase/decrease by the same factor.

Knowing the BWTP of a pulse allows one to specify its BW by applying it for the appropriate time. From Eq. (7.2), the BW of a pulse is also related to the slice thickness achieved when played with a given gradient. As most pulses are used for slice selection, this gives a more physically intuitive way to understand the pulse BW in imaging.

The Siemens sequence programming environment (IDEA) utilizes most of its RF pulses from external *.pta* files that specify the scaled magnitude and phase of each point defining the pulse shape along with a couple other parameters. The BWTP is not one of the given parameters, with its role being played by parameter *Refgrad*. *Refgrad* is an experimentally determined parameter specifying the gradient strength required to select a 1 cm slice when the given pulse is played for 5.12 ms and is given in units of mT/m.

The BW of the pulse can be calculated off the gradient and slice thickness from Eq. (7.2), so the BWTP is simply

$$\text{BWTP} = \text{BW} \cdot T_p = \gamma G \Delta z \cdot T_p \quad (7.5)$$

Rearranging and plugging in the appropriate reference values of $\gamma = 42576 \text{ Hz/mT}$ (for protons), $T_p = 5.12\text{e-}3 \text{ s}$, and $\Delta z = 0.01 \text{ m}$, yields an equation for *Refgrad* in terms of the BWTP

$$\text{Refgrad} = \frac{\text{BWTP}}{(42576 \text{ Hz/mT}) (0.01 \text{ m}) (5.12\text{e-}3 \text{ s})} = \frac{\text{BWTP}}{2.180 \text{ m/mT}} \quad (7.6)$$

CHAPTER 8

Quantitative Localized Correlated Spectroscopy using Prior Knowledge Fitting Including Macromolecules: Comparison with J-resolved Spectroscopy and Results in HIV-infected Children

8.1 Introduction

Involvement of the central nervous system (CNS) is very common in HIV and clinical CNS-related disease has been documented to be present in approximately 17% of HIV survivors [169]. In contrast to adult patients who acquire HIV with a mature immune system and in whom HIV encephalopathy is usually a later manifestation of HIV, pediatric patients have infection of the CNS early in the infection and continuing effects from the developmental insult of the brain and continuing subclinical or clinical involvement. Many researchers have investigated HIV effects of cerebral metabolites using one-dimensional MRS at 1.5 T field strength and reported changes in NAA, choline, myo-inositol, and creatine [21, 11, 91, 149, 115, 34, 33].

Magnetic resonance spectroscopy (MRS) provides a means to measure metabolic function non-invasively. Proton MRS of the human brain can detect about 20 metabolites in vivo. However, these resonances exist in a relatively small spec-

tral window from about 1 ppm to the broad water peak centered at 4.7 ppm, and many exhibit complicated multiplet structures due to J-coupling. This results in 1D spectra that are overcrowded and dominated by the four singlets of N-acetylaspartate (NAA), creatine (Cr), and total choline (tCho) and possibly nuisance signals do to residual water or lipid contamination. Quantification by peak height or area is unreliable for the singlets due to baseline effects and impossible for overlapping multiplets. Spectral editing techniques [142, 80] been used to highlight a single J coupled metabolite at a time but are highly susceptible to field inhomogeneity artifacts caused by motion. In order to obtain objective quantification of many metabolites, spectral fitting methods that make use of prior knowledge have been developed [175, 129, 159].

These methods, whether in the time domain as in VARPRO [175] or in the frequency domain as in LCModel [129], fit acquired data as a linear combination of metabolite data that compose the *basis set*. In the frequency domain, this relationship can be expressed as

$$S = \sum_m c_m B_m + \epsilon \quad (8.1)$$

where S is the acquired spectrum, B_m is a matrix of the m th metabolite basis spectra, c_m is the concentration of metabolite m , and ϵ is the residual noise. Other parameters such as frequency shift, linewidths, and phase correction are included to improve the overall accuracy of the fit and the expense of computational complexity and time. The quality of the fit can be judged objectively by the Cramér-Rao lower bounds (CRLBs), which take into account both the noisiness of the spectra and the amount of overlap in the basis set [32]. Despite the great improvement in fitting brought about by these prior knowledge-based algorithms, certain metabolites, such as lactate (Lac) and threonine (Thr) or glutamate and glutamine, overlap so much that they still can not be untangled reliably. One way to greatly reduce spectral crowding and unambiguously distinguish highly overlapping peaks is by acquiring frequency information in a second, orthogonal

dimension.

The basic MRS pulse sequence consists of a preparation period in which equilibrium magnetization is excited, followed by an evolution period without RF pulses, a mixing period, and an acquisition period in which a free induction decay (FID) or echo is read out. In a 2D pulse sequence [9], the duration of the evolution period t_1 is incremented each repetition such that the echo time TE becomes increment-dependent, and the magnetization during readout is a function of t_1 . Due to the longer scan times required to acquire the t_1 dimension indirectly in this way, only the simplest 2D sequences have been applied in vivo [148, 171]. The 2D analog of PRESS, called JPRESS [148], encodes chemical shift and J-coupling information along the direct dimension and J-coupling only along the indirect dimension. Complicated multiplet structures that appear in a single dimension in a PRESS spectrum are smeared into the second spectral dimension in JPRESS.

JPRESS greatly reduces the spectral crowding in the presence of coupled spins. However, while it is easy to identify spins that experience coupling, there is no information that connects coupled partners. An alternative sequence that provides this correlation map and enhanced spectral dispersion is the localized analog of Jeener's original proposed 2D sequence dubbed L-COSY for localized correlation spectroscopy [171]. In L-COSY, the final 180° localization/refocusing pulse of JPRESS is changed to a 90° localization/refocusing/coherence transfer pulse, linking the frequency in the evolution period with the frequency during acquisition.

Early attempts to quantify 2D spectra had limited success. As with 1D, peak integration has been used but limits are subjective, and the method is unreliable for crowded and overlapping peaks. De Beer et al. used the 2D Hankel singular value decomposition (SVD) to decompose the time domain data into its resonances and linewidths [43]. However, SVD fits do not make full use of prior knowledge and cannot be constrained to physically plausible values. Another strategy im-

plemented by Slotboom *et al.* took the 1D TDFD fitting routine and extended it as a series of constrained 1D fits, but they found no improvement in fitting of coupled metabolites relative to 1D [159]. Schulte and Boesiger proposed a true 2D fitting algorithm called ProFit (PRiOr knowledge FITting) and showed an increase in the number of quantifiable metabolites using 2D JPRESS compared with 1D PRESS [152].

ProFit is an iterative function relying on both linear and non-linear least squares optimization to minimize a cost function made up of the sum of squares residual plus penalty factors for regularization. It begins by determining the global phase correction and frequency shifts using the dominant singlets. From there, it enters its first iteration, fitting the relative frequency shift and phase non-linearly for the dominant metabolites, then fitting concentrations linearly. The second iteration includes a lineshape distortion to the non-linear fit, and the third iteration includes all the basis metabolites. By dividing the algorithm into a linear and non-linear part, the number of parameters is reduced, leading to faster overall convergence. In this paper, we demonstrate the accuracy and reliability of fitting L-COSY data using the ProFit algorithm. We then use L-COSY with ProFit to identify differences in metabolite concentrations in young HIV patients relative to healthy controls. Two-dimensional L-COSY was evaluated recently using a 1.5 T scanner with a pilot group of HIV children [10]. In that study, quantitation was done using operator-defined peak volumes in the frequency domain without prior knowledge.

8.2 Methods

8.2.1 Phantom Scans

Phantom scans were performed on a gray matter phantom containing the following 16 metabolites at physiological concentrations [69]: aspartate (Asp), N-

acetylaspartate (NAA), N-acetylaspartylglutamate (NAAG), γ -aminobutyric acid (GABA), free choline (Cho), creatine (Cr), glucose (Glc), glutamate (Glu), glutamine (Gln), glutathione (GSH), lactate (Lac), myo-inositol (mI), phosphorylcholine (PCh), phosphorylethanolamine (PE), taurine (Tau), and threonine (Thr).

Spectroscopic sequences were run following a three pulse water suppression module as described by Ogg *et al.* [117]. The 1D PRESS spectra were acquired with TE/TR=30/5000 ms, 1024 averages, voxel size of 27 cm³, and bandwidth of 2000 Hz, resulting in a scan duration of 1hr 25min. The 2D L-COSY spectra were acquired with initial TE/TR=30/5000 ms, 100 t_1 increments with step size $\Delta t_1=0.8$ ms, 8 averages, and a voxel size of 27 cm³. The bandwidth in the direct dimension f_2 was again 2000 Hz, with a bandwidth in the indirect dimension f_1 of 1250 Hz. The 2D JPRESS spectra were acquired with the same parameters except with a bandwidth of 1000 Hz in the indirect dimension. Both scan durations were 1hr 07min. Each spectroscopic sequence was preceded by a short (1 average, 1 t_1 increment) non water-suppressed scan which was used in initial preprocessing for coil combination and zero-order phase correction before the data was fed into ProFit.

8.2.2 *in vivo* Scans

Sixteen children with HIV and 14 healthy volunteers (9 male, 5 female, average age 16) were scanned using the L-COSY sequence with the following parameters: initial TE/TR=30/2000 ms, 100 t_1 increments with step size 0.8 ms, 8 averages, voxel size of 27 cm³, and bandwidths of 2000 Hz in the direct dimension and 1250 Hz in the indirect dimension. A single average, single t_1 increment non water-suppressed scan was also acquired for eddy current correction and coil combination [89]. Localization was performed with a 3 minute axial T_1 -weighted scan. Total scan time was about 30 min. Spectra were acquired from a voxel in the frontal lobe that was composed of predominantly white matter. All scans were performed on

a Siemens 3T Trio-TIM scanner (Siemens Medical Solutions, Erlangen, Germany) using a 12 channel phased array receive coil acquired in triple mode with 4 primary channels. The body coil was used for transmit.

8.2.3 Processing and Fitting

The ProFit algorithm is described in detail in the paper by Schulte and Boesiger in [152]. Here, only a brief overview will be given, highlighting the adaptations made to process L-COSY spectra.

ProFit consists of both non-linear and linear constrained least squares optimizations in order to implement maximal prior knowledge while keeping small the number of degrees of freedom. The following global parameters were optimized non-linearly: zero-order phase ϕ_0 , Gaussian line-broadening in the direct dimension σ_g , frequency shift in the indirect dimension Δ_1 , and biexponential phase decay amplitudes θ_1 and θ_2 . Fitting a biexponential phase decay model corrects for lineshape distortions caused primarily by eddy currents. Frequency shifts were determined from the maximum of the dominant singlets of NAA, Cr30, and PCh in the magnitude spectrum. Phase correction in ProFit is done by optimizing peak shapes and areas of the dominant singlets in both the real and imaginary spectra and therefore, does not require a separate reference scan. In the real spectrum, peak area and peak height are both maximized. In the imaginary spectrum, peak area is minimized to zero, and the highest positive peak is matched to the lowest negative peak. Zero- and first-order terms are linearly fit to the overall phase corrections applied to the singlets.

In addition to relying on adjustment of global parameters for the experimental data, a successful fit is achieved by making small adjustments to individual basis spectra. For each metabolite, exponential line-broadening σ_e is added in both spectral dimensions to replicate T_2^* losses in the acquired data, and a small

frequency shift in the direct dimension Δ_2 is included to compensate for pH or temperature dependent effects. The individual metabolite concentrations c_m are determined linearly by the product of the Moore-Penrose pseudoinverse of the matrix of basis spectra and the experimental spectrum. A nonlinear optimization is applied to minimize a cost function that is made up of the residual as well as regularization terms for the line broadening and chemical shifts and is given by

$$f = \sum_{k,l} \left[\Re \{S_{k,l}(\Delta_1, \phi_0, \theta_1, \theta_2)\} - \sum_m c_m \Re \{B_{k,l,m}(\Delta_{2,m}, \sigma_{e,m}, \sigma_g)\} \right]^2 + \rho_\sigma \sum_m c_m (\sigma_{e,m} - \sigma_{e,typ})^2 + \rho_\Delta \text{var}(\Delta_{2,m}) \quad (8.2)$$

where the regularization factors $\rho_\sigma = (\sum_m c_m)^{-1}$ and $\rho_\Delta = 1$, $\sigma_{e,typ}$ was 1, and var indicates the variance. The summation indices k and l were taken over the all the rows and columns of a truncated fit region. All processing was done offline using MATLAB (The Mathworks, Natick MA, USA).

Metabolite basis spectra were simulated numerically with GAMMA [160], using chemical shifts and J-couplings from the literature [69]. The basis sets used for the phantom scans contained the same 16 metabolites as the phantom itself, while the basis set used for the in vivo scans also contained alanine (Ala), ascorbic acid (Asc), glycerophosphorylcholine (GPC), glycine (Gly), iso-leucine (Ile), leucine (Leu), scyllo-inositol (sI), and valine (Val). The two singlets of Cr were fit as separate spectra (Cr30 and Cr39) to offer a quick check of fit accuracy, as their ratios should always be 1:1. Cr39 ratios less than 0.6 or greater than 1.3 were used as exclusion criteria for the entire spectrum.

Individual metabolite fit quality was determined qualitatively by visual inspection of the residual and quantitatively using the CRLBs. The CRLB of each metabolite is calculated from the diagonal element of the Fisher information ma-

trix F as

$$\text{CRLB}_m = \sqrt{F_{m,m}^{-1}} \quad (8.3)$$

with F given by

$$F = \frac{1}{\sigma^2} \Re(B^T B) \quad (8.4)$$

where σ^2 is the noise variance of the experimental spectrum, B is the basis matrix consisting of the simulated spectra, and \Re indicates the real part of the matrix.

In order to be detected, signal following coherence transfer must evolve from an antiphase operator term such as $2I_x S_z$ to an in phase operator term such as I_y , which modulates the signal by a factor of $\sin(\pi J t_2)$ in the direct dimension. A similar modulation occurs in the indirect dimension, as an in phase operator must evolve to an antiphase operator. This causes relatively less signal for these “cross peaks” compared to signal that did not evolve from coherence transfer and is cosine modulated (i.e. “diagonal peaks”) due to T_2^* relaxation. In addition, T_2^* causes severe line broadening in vivo that is larger than the J-coupling constant. Adjacent peaks in the multiplet with opposite signs will overlap and lead to signal cancellation. Therefore, before being fed into ProFit, L-COSY data was filtered using a skewed square sine bell (skew parameter=0.5) in the direct dimension and a square sine bell in the indirect dimensions to enhance the presence of cross peaks in the 2D spectrum. The same filtering was applied to the L-COSY basis set as well so that the relative contributions of cross peaks and diagonal peaks would be consistent with the experimental data.

The “cross peaks” in a JPRESS spectrum do not have their own multiplet structure and have the same orientation, so partial signal cancellation is not an issue. Also, they have maximum signal at the echo just as singlet peaks, and therefore, filtering is not required.

A frequency-dependent phase accumulation factor was applied in post-processing to simulate the effect of beginning acquisition at the echo maximum, whereas ac-




Figure 8.1: L-COSY basis set used for phantom scans (a) as simulated by GAMMA and (b) after a frequency-dependent phase factor was applied. The two Cr singlets were simulated and fit independently as described in the text.

tual spectra were acquired with acquisition beginning immediately after the last crusher gradient for all TE values [106, 154]. This gives the L-COSY spectrum the appearance of a spin-echo correlation spectroscopy (SECSY) [116] spectrum and the JPRESS spectrum the appearance of a traditional J-resolved spectrum in which the diagonal peaks are refocused and appear horizontally along f_2 , and the frequency differences are scaled by a factor of 1/2 in the f_1 direction. By processing the spectra this way, a smaller spectral region can be fit which should speed up computation time. This effect is illustrated on the L-COSY basis set before and after the phase factor was applied in Fig.8.1a and 8.1b, respectively.

8.3 Results

8.3.1 Phantom Scans

	L-COSY				JPRESS				Actual
	Mean Ratio	CV	Mean CRLB	% Fit	Mean Ratio	CV	Mean CRLB	% Fit	Phantom Ratio
Cr30	1	0	0.242	100	1	0	0.464	100	1
Cr39	0.728	9.202	0.435	100	0.884	3.197	0.793	100	1.000
NAA	1.263	5.800	0.181	100	1.281	5.273	0.358	100	1.271
PCh	0.299	9.533	0.539	100	0.126	8.997	1.474	100	0.086
Asp	0.284	29.173	3.390	100	0.192	21.800	9.572	97.4	0.300
Cho	0.171	13.235	0.686	100	0.127	8.016	1.414	100	0.129
GABA	0.072	169.560	10.304	13.5	0.182	37.279	5.872	100	0.100
Glc	0.129	28.503	7.464	94.6	0.106	58.744	14.999	15.8	0.143
Gln	0.106	29.374	8.905	70.3	0.125	15.344	10.087	73.7	0.357
Glu	0.941	24.393	1.439	100	1.398	6.173	1.104	100	1.786
GSH	0.057	58.912	7.131	100	0.141	11.501	3.879	100	0.290
Lac	0.128	29.670	5.277	100	0.086	12.389	13.896	89.5	0.143
mI	0.895	11.131	0.770	100	0.870	16.072	1.570	100	0.629
NAAG	0.144	40.939	1.892	97.3	0.211	26.628	2.178	84.2	0.073
PE	0.051	244.220	14.127	16.2	0.118	109.493	12.961	60.5	0.143
Tau	0.348	17.693	1.748	100	0.223	19.141	6.015	100	0.257
Thr	0.035	19.868	14.971	21.6	-	-	-	0	0.043
tNAA	1.404	4.837	0.234	100	1.459	3.498	0.354	100	1.344
Glx	1.020	19.091	1.664	100	1.504	6.300	1.266	100	2.143
tCho	0.470	6.135	0.372	100	0.252	4.579	0.657	100	0.214

Table 8.1: Comparison of the ProFit results in L-COSY and JPRESS in phantom. 37 L-COSY spectra and 38 JPRESS spectra were acquired. Ratios are shown with respect to the Cr singlet at 3.03ppm. % Fit is the percent of spectra that were fit with CRLB<20%. Only those voxels were included in the computation of mean ratios, mean CRLBs, and coefficients of variation (CV). CVs are expressed as percentages. tNAA = NAA+NAAG, Glx = Glu+Gln, tCho = PCh+Cho.

Table 8.1 shows the fit accuracy and reproducibility in the L-COSY phantom scans and compares to those in JPRESS scans. Each sequence was able to reliably fit 12 of the 16 metabolites with coefficients of variation (CV) under 30%, with 5 and 7 metabolites showing CVs under 15% for L-COSY and JPRESS, respectively. Also apparent is that the different information provided by the two sequences results in differences in fit quality for different metabolites with L-COSY performing better for Asp, Glc, Lac, and Thr and JPRESS performing better for PCh, Cho, GABA, and Tau to name a few. Certain metabolites are consistently and predictably underestimated by both sequences such as Cr39, Gln, and Glu, while others such as mI and NAAG are consistently overestimated.

It must be noted though that because the L-COSY data was filtered, the noise variance was reduced, and the reported CRLBs are not tight bounds on the standard deviations of the ratios. For that reason, the reader is urged to consider the CV as the primary measure of reproducibility and discouraged from comparing mean CRLBs between L-COSY and JPRESS. Nevertheless, the CRLBs in the L-COSY sequence still give a relative measure of confidence in the fitted concentrations, and it is for that reason the authors have included them in Table 8.1 and determined fit success based on a CRLB threshold of 20%. The unfiltered CRLBs in JPRESS also highly underestimate the true variation.

8.3.2 HIV patients in vivo

A typical fit with residual is shown graphically in Fig. 8.2b for the voxel location shown in Fig. 8.2a. The fit was restricted to the spectral region enclosed by the white box. From the spectrum, it is clear that lipid contamination was minimal despite the voxel's proximity to the skull, and from the white box, residual water was mostly excluded from the fit.

The different fitted concentrations of healthy controls and HIV patients are

	Healthy Controls		HIV Patients		p-value
	Mean Ratio±SD	N	Mean Ratio ± SD	N	
Cr30	1.000±0.000	14	1.000±0.000	15	-
Cr39	0.904±0.172	13	0.906±0.107	15	0.966
NAA	1.300±0.181	14	1.388±0.181	15	0.201
GPC	0.196±0.115	14	0.197±0.061	15	0.973
Ala	0.055±0.044	14	0.121±0.023	2	0.065
Asc	0.404±0.329	11	0.342±0.206	13	0.596
Asp	0.433±0.102	14	0.510±0.087	15	0.038
Cho	0.173±0.066	14	0.152±0.046	15	0.327
Gln	0.062±0.053	10	0.266	1	-
Glu	1.672±0.221	14	1.886±0.224	15	0.015
GABA	0.270±0.182	14	0.381±0.115	4	0.259
Glc	0.205±0.142	13	0.364±0.117	13	0.005
GSH	0.344±0.137	12	0.287±0.151	15	0.311
Ile	0.172±0.119	12	0.272±0.053	11	0.019
Lac	0.171±0.090	13	0.185±0.064	14	0.638
Leu	0.018±0.008	9	0.108±0.021	2	0.095
mI	1.187±0.385	14	1.444±0.262	15	0.048
NAAG	0.005±0.003	4	0.079±0.031	2	0.181
PCh	0.166±0.025	14	0.162±0.024	15	0.685
PE	0.535±0.224	14	0.524±0.140	15	0.874
Scy	0.064±0.017	14	0.084±0.023	15	0.011
Tau	0.282±0.071	14	0.272±0.088	15	0.737
Thr	0.123±0.059	13	0.172±0.061	7	0.108
Val	0.093±0.058	14	0.037±0.036	3	0.087
tNAA	1.333±0.139	14	1.388±0.173	15	0.351
Glx	1.741±0.236	14	1.947±0.263	15	0.035
tCho	0.491±0.049	14	0.511±0.093	15	0.488
Ile+Leu	0.193±0.103	14	0.247±0.107	15	0.172
Lac+Thr	0.294±0.137	13	0.317±0.111	15	0.628

Table 8.2: ProFit L-COSY results of frontal white/gray brain spectra in healthy controls and patients. 14 healthy controls and 15 HIV patients were scanned. Ratios are shown with respect to the Cr singlet at 3.03ppm. SD is the standard deviation of the ratios. N is the number of voxels that were fit with CRLB<20%. Only those voxels were included in the computation of mean ratios and SDs. p-values indicating statistically significant differences at the 5% level are highlighted in bold. tNAA = NAA+NAAG, Glx = Glu+Gln, tCho = GPC+PCh+Cho.




Figure 8.2: An example in vivo spectrum (top) from an HIV patient, its fit (middle), and its residual (bottom), plotted on a logarithmic scale.

shown in Table 8.2. Fourteen healthy controls and 16 HIV patients were scanned. One of the HIV patient scans was of poor quality and excluded from the analysis based on a Cr39 ratio that was below 0.6 or above 1.3. A two-tailed, heteroscedastic (unequal variances) two-sample t-test was performed to look for statistically significant differences between the two groups. At the 5% level, significant differences existed in Asp, Glu, Glc, Ile, mI, and Scy concentration ratios which are highlighted in bold in the table.

8.4 Discussion

From Table 8.1, it is seen that while both sequences accurately fit many brain metabolites, the JPRESS sequence tended to achieve fits with lower CVs for most. This is not surprising and is due to the higher overall sensitivity of JPRESS, as the final 90° pulse in L-COSY only refocuses half the signal. The CRLBs in L-COSY were for the most part lower, but this was primarily because the L-COSY




Figure 8.3: Normalized variance matrices for (a) the L-COSY basis set and (b) the JPRESS basis set. The variance matrix is the inverse of the Fisher matrix, and the CRLBs are the diagonal elements.

data was apodized. Though the reason for the apodization was to enhance the cross peaks and not the sensitivity, the filters still reduced the noise variance and therefore the calculated CRLBs.

What is more interesting is that L-COSY results in better fit accuracy for metabolites with weakly coupled spins such as Lac and Thr. The spectral dispersion of cross peaks from the diagonal in JPRESS is the size of the J-coupling constant (5-10 Hz), whereas the spectral dispersion of cross peaks in L-COSY is the difference in the chemical shifts between the coupled spins (375 Hz for Lac/Thr at 3T). In the case of weak coupling, the chemical shift difference greatly exceeds the J-value, and L-COSY has much larger spectral dispersion in the indirect dimension relative to JPRESS. Asparate is an ABX spin system with both strongly and weakly coupled spins. For typical in vivo conditions, though, the resonances at 2.80 ppm and 2.65 ppm highly overlap and appear as a broad resonance weakly coupled to the spin at 3.89 ppm. The in vivo spectrum of the strongly coupled seven spin system Glc is similarly simplified to two broad resonances around 3.4

and 3.8 ppm.

The improved spectral dispersion of L-COSY results in greater orthogonality in the basis set, which can be seen in Figure 8.3 showing the variance matrices for an L-COSY scan in 8.3a and a JPRESS scan in 8.3b. The variance matrix is the inverse of the Fisher information matrix given in Eq. (8.4). To compare the basis orthogonality only and provide similar scaling, the variance matrices were normalized by the noise. It is clear from normalized variance matrices that there is less correlation between metabolites in the L-COSY basis set than in the JPRESS basis set. In particular, note the reduced correlation between Lac and Thr, Cr39 and Asp, and Cr30 and GABA. On the other hand, the correlation between Glc and Tau is higher in the L-COSY basis set which is somewhat surprising since Glc was fit better in L-COSY than JPRESS.

Both sequences slightly underestimated the Cr39 singlet due mainly to the effects of imperfect water suppression pulses which had a suppression bandwidth of 40 Hz (about 0.3 ppm) but nevertheless, slightly suppress signals farther away from water. This is despite a large amplitude signal that sits on top of the residual water tail. In addition, the excitation frequency was set to be at the middle of the spectrum at 2.7 ppm, so the Cr30 singlet was more on resonance than the Cr39 singlet. For Glx, the situation is more complicated, as there is no singlet resonance. The reason for this is most likely do to the relatively short T_2 values compared to Cr [38]. Using a short echo sequence mitigates signal losses due to T_2 relaxation, but the increased line broadening pushes the peak tails into the noise floor and results in underestimated concentrations. This reason is further suggested by the in vivo data, where fitted concentrations of Glx are closer to those reported in the literature [122, 163]. In that case, T_2^* dominates T_2 relaxation and is the same for each metabolite, so the only effect of differing T_2 values is in the signal amplitude at the beginning of readout, and these are close for short echo sequences.

This study was conducted entirely at 3T, but ProFit can be easily adapted for any field strength as long as the basis set is obtained appropriately. At lower field strength, the sensitivity of 2D experiments becomes more of a limiting factor in spectral quality, so JPRESS would be expected to perform better. At higher field strength though, the sensitivity is greater, more spin systems become more weakly coupled, and the difference in spectral dispersion between L-COSY and JPRESS becomes greater. Therefore, L-COSY would be expected to perform better.

There are conflicting reports about changes in Glu in HIV, whether it increases possibly due to damage to neuroglial elements at earlier stages of infection that antedate decreases in NAA [77] or decreases due to reduced astrocytic reuptake of GLU, secondary excitotoxicity, and mitochondrial toxicity from antiretroviral treatments [56]. In [77], mI was shown to be increased in HIV patients relative to controls but that the mI concentration decreased with age in HIV patients. That study only looked at adults patients though, so recognizing similar trends in pediatric patients can help establish a more thorough understanding of progression. Aspartate and isoleucine are very difficult to reliably detect with 1D spectroscopy techniques, so there is very little information about their roles in HIV in vivo, but aspartate does play a role in HIV-1 protease inhibitor drugs [23]. In HIV-related dementia studies, glucose has been shown to decrease for severe cases [144] but increase for mild cases [143]. None of the pediatric patients in this study had any symptoms of dementia. Increased scyllo-inositol has been correlated with increased myo-inositol as measures of negative brain health [87].

From Table 8.2, it appears that more of the metabolites are showing increase in ratios of patients relative to healthy controls. This could actually be attributable to a slight decrease in the creatine signal. Typically, creatine is used as an internal reference standard because its concentration is relatively stable across multiple pathologies. However, there are reports of creatine decrease in the frontal lobe and basal ganglia in HIV [36, 35]. Further research needs to be done with non-

metabolite-based reference standards, such as those using electronic referencing [4].

8.5 Conclusion

Two dimensional prior knowledge fitting was applied to a localized correlated spectroscopy sequence and compared favorably to fits obtained with a J-resolved spectroscopy sequence. Significant increases in aspartate, glucose, glutamine, isoleucine, and myo- and scyllo-inositol were observed.

CHAPTER 9

Conclusion

9.1 Summary of significance

In this dissertation, we have expanded on the use of NUS in multidimensional MRSI to achieve full 3D volume localization along with two spectral dimensions. We showed in Chapter 4 that 5D EP-JRESI could be acquired in brain in about 20 minutes with reliable results comparable to a time-equivalent single-slice scan using either an ℓ_1 -norm minimization or total variation minimization. By masking the spectra in the direct spectral dimension, the influence of large nuisance signals, such as residual water and lipids, can be minimized during the reconstruction.

In Chapter 5, 5D EP-COSI was shown in human calf muscle. Improved results were obtained by exploiting the group sparsity of the COSY spectral plane, compared to traditional point-by-point ℓ_1 -norm minimization used previously. To handle acquired data from many receiver coils, a version of the algorithm with computational complexity similar to the single coil case was developed. In addition, a coil array compression technique that maximizes the SNR based on the sensitivity maps while simultaneously allowing different linear weightings for compression was used to improve the speed of the algorithm and reduce its computational burden.

The next chapters concerned circular k-space sampling in multidimensional MRSI. Circular sampling is well suited for these scans as k-space can be sampled more efficiently with reduced gradient slew rates. COSI-CONCEPT calf mus-

cle scans were compared to time-equivalent EP-COSY scans, showing the same localization and spectral quality in Chapter 6.

In Chapter 7, the conventional refocusing pulses used in JRESI-CONCEPT were replaced with adiabatic ones. Adiabatic pulses are higher in bandwidth and therefore exhibit reduced chemical shift displacement error, and they are insensitive to RF field inhomogeneities. In order to use the adiabatic pulses for refocusing, a pair had to replace each original pulse, increasing the minimum echo time. The new sLASER-based JRESI-CONCEPT was compared to the original in the case of equal echo times and in the case of minimum echo times. Both showed improved localization and spectral quality compared to the original, especially for strongly coupled metabolites such as Glx.

Chapter 8 shows the application of prior knowledge fitting that includes simulated macromolecule resonances to COSY data in pediatric HIV patients and compares phantom results to JPRESS data, highlighting the benefits of each sequence.

9.2 Future Directions and Recommendations

9.2.1 NUS and reconstruction

The results in Chapter 4 show a proof of concept that five-dimensional scans can be acquired in vivo. Chapter 5 extends that by showing improved reconstruction with group sparsity and considering multiple coils. Nevertheless, there is room for improvement, and more validation needs to be done.

1. Despite the crowding in 1D spectroscopy, prior knowledge-fitting programs like LCModel can still separate many overlapping metabolites. Current versions of quantitation use peak integration, which may not actually be any better than single dimensional spectroscopy with LCModel. Profit has

been shown to offer improved fitting compared to LCModel and should be applied to this NUS data. Preliminary work seems to show consistent fitting of reconstructed data. However, one confounding factor is the effect of the non-linear reconstruction on the noise level and therefore the acceptability criteria (based on CRLBs) of the fits.

2. Weighted shrinkage operators can be used to include additional prior knowledge on the reconstructions for either ℓ_1 or GS reconstructions. These weights could be spatial and obtained from the sensitivity maps, or they could be spectral and obtained from known chemical shifts and J-couplings. Preliminary work suggests that weighting can enhance the cross peak reconstruction in undersampled EP-COSY, but there can be problems with convergence as the sampled data remains unchanged.
3. Group sparsity was applied with groups taken in the spectral plane. Grouping could be extended over to the spatial dimensions and the coil dimension, which should offer some improvement in the spatial SNR of the GS reconstructions.
4. Chapter 4 suggested that applying finite differences to the spatial dimensions (as in total variation minimization) performed better than the simple ℓ_1 -norm of the data itself. While grouping greatly improved this self-sparse reconstruction in Chapter 4, a reconstruction that combines grouping with TV could perform even better. Grouping could be done on the finite differences, or finite differences could be taken on the groups.

9.2.2 Concentric circular sampling

Perhaps the greatest advantage to concentric circular sampling is its ability to acquire higher spectral bandwidths. Despite that, due to available hardware, all scans in this dissertation were taken at 3T, where the required SBW is lower.

Ultimately, gradient amplitudes and slew rates limit the achievable SBW per shot. At 7T, the increased spectral dispersion means that the SBW must be higher as well. This has been a big problem for EPSI or Spiral-SI at high fields.

1. However, there are other issues besides spectral bandwidth that make 7T spectroscopic imaging difficult. Namely, the increased field inhomogeneity, the shorter T2 decay times, and the increased energy deposition (SAR) are of particular concern. Using adiabatic refocusing pulses as demonstrated in Chapter 7 addresses the RF inhomogeneity issue while at the same time increasing the SAR and minimum TE and therefore, exacerbating the latter two issues. In order to stay within SAR limits at 7T, the adiabatic pulses must be elongated, further increasing the minimum TE and increasing the CSDE. In order to combat the SAR problem, methods such as variable-rate selective excitation (VERSE) [41], frequency offset corrected inversion (FOCI) [118], and gradient-modulated offset-independent adiabaticity (GOIA) [167, 168] can be used to distribute the RF amplitude more uniformly over the pulse by modulating the slice-selective gradient, thereby lowering the SAR without any reduction in bandwidth.
2. While the SI-CONCEPT sequence can achieve higher SBW than either EPSI or Spiral-SI in a single shot, it is not without limits. The amplitude and slew rates scale with the k-space radius, and as a result, higher resolution scans are more demanding on the gradient hardware. Temporal interleaving will eventually be required even for SI-CONCEPT. At that point, another, subtle advantage of circular sampling can be exploited. Unlike spiral or rectilinear trajectories, the gradient hardware demands are different for each shot. This means that the limits are reached (and temporal interleaving is required) only for certain shots in outer k-space, and only those shots need to be interleaved. This idea is illustrated in Fig. 9.1.

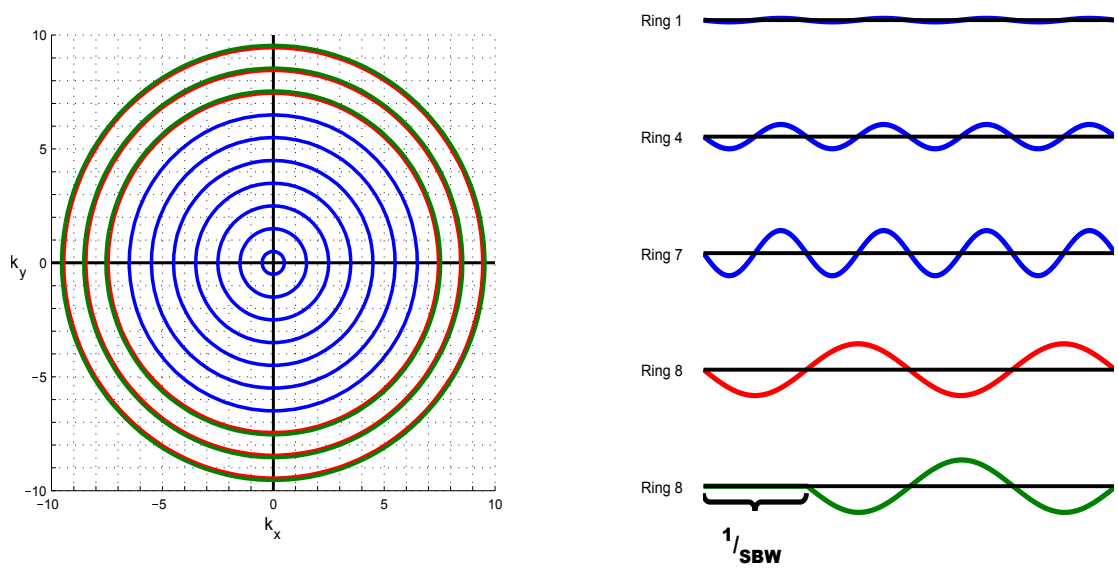


Figure 9.1: k-space trajectories (left) and x gradient waveforms. Blue rings collect inner k-space at high SBW, while red and green rings collect outer k-space at half the SBW and are temporally interleaved.

REFERENCES

- [1] A. Abragam. *The Principles of Nuclear Magnetism*. International series of monographs on physics. Clarendon Press, 1961.
- [2] Elfar Adalsteinsson, Pablo Irarrazabal, Simon Topp, Craig Meyer, Albert Macovski, and Daniel M. Spielman. Volumetric spectroscopic imaging with spiral-based k-space trajectories. *Magnetic Resonance in Medicine*, 39(6):889–898, 1998.
- [3] Elfar Adalsteinsson and Daniel M. Spielman. Spatially resolved two-dimensional spectroscopy. *Magnetic Resonance in Medicine*, 41(1):8–12, 1999.
- [4] Mark J. Albers, Thomas N. Butler, Iman Rahwa, Nguyen Bao, Kayvan R. Keshari, Mark G. Swanson, and John Kurhanewicz. Evaluation of the eretic method as an improved quantitative reference for 1h hr-mas spectroscopy of prostate tissue. *Magnetic Resonance in Medicine*, 61(3):525–532, 2009.
- [5] Adam Allerhand. Analysis of carrpurcell spinecho nmr experiments on multiplespin systems. i. the effect of homonuclear coupling. *The Journal of Chemical Physics*, 44(1), 1966.
- [6] M.J. Allison, S. Ramani, and J.A. Fessler. Accelerated regularized estimation of mr coil sensitivities using augmented lagrangian methods. *Medical Imaging, IEEE Transactions on*, 32(3):556–564, March 2013.
- [7] Ovidiu C. Andronesi, Borjan A. Gagoski, Elfar Adalsteinsson, and A. Gregory Sorensen. Correlation chemical shift imaging with low-power adiabatic pulses and constant-density spiral trajectories. *NMR in Biomedicine*, 25(2):195–209, 2012.
- [8] Ovidiu C. Andronesi, Saadallah Ramadan, Carolyn E. Mountford, and A.Gregory Sorensen. Low-power adiabatic sequences for in vivo localized two-dimensional chemical shift correlated mr spectroscopy. *Magnetic Resonance in Medicine*, 64(6):1542–1556, 2010.
- [9] W. P. Aue, E. Bartholdi, and R. R. Ernst. Two-dimensional spectroscopy. application to nuclear magnetic resonance. *The Journal of Chemical Physics*, 64(5):2229–2246, 1976.
- [10] S Banakar, MA Thomas, A Deveikis, J Watzl, J Hayes, and MA Keller. Two-dimensional 1h mr spectroscopy of the brain in human immunodeficiency virus (hiv)-infected children. *Journal of Magnetic Resonance Imaging*, 27:710–717, 2008.

- [11] PB Barker, RR Lee, and JC McArthur. Aids dementia complex: evaluation with proton mr spectroscopic imaging. *Radiology*, 195:58–64, 1995.
- [12] J.C.J Barna, E.D Laue, M.R Mayger, J Skilling, and S.J.P Worrall. Exponential sampling, an alternative method for sampling in two-dimensional nmr experiments. *Journal of Magnetic Resonance*, 73(1):69 – 77, 1987.
- [13] J. Baum, R. Tycko, and A. Pines. Broadband and adiabatic inversion of a two-level system by phase-modulated pulses. *Phys. Rev. A*, 32:3435–3447, Dec 1985.
- [14] P.J. Beatty, D.G. Nishimura, and J.M. Pauly. Rapid gridding reconstruction with a minimal oversampling ratio. *Medical Imaging, IEEE Transactions on*, 24(6):799 –808, june 2005.
- [15] M.Robin Bendall and David T Pegg. Uniform sample excitation with surface coils for in vivo spectroscopy by adiabatic rapid half passage. *Journal of Magnetic Resonance*, 67(2):376 – 381, 1986.
- [16] Dimitri P. Bertsekas. Multiplier methods: A survey. *Automatica*, 12(2):133 – 145, 1976.
- [17] F. Bloch. Nuclear induction. *Phys. Rev.*, 70:460–474, Oct 1946.
- [18] 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):1086–1098, 2007.
- [19] Chris Boesch and Roland Kreis. Dipolar coupling and ordering effects observed in magnetic resonance spectra of skeletal muscle. *NMR in Biomedicine*, 14(2):140–148, 2001.
- [20] Paul A. Bottomley. Spatial localization in nmr spectroscopy in vivo. *Annals of the New York Academy of Sciences*, 508(1):333–348, 1987.
- [21] A Brand, C Richter-Landsberg, and D Leibfritz. Multinuclear nmr studies on the energy metabolism of glial and neuronal cells. *Developmental Neuroscience*, 15:289–298, 1993.
- [22] Ian M. Brereton, Graham J. Galloway, Stephen E. Rose, and David M. Doddrell. Localized two-dimensional shift correlated spectroscopy in humans at 2 tesla. *Magnetic Resonance in Medicine*, 32(2):251–257, 1994.
- [23] Ashraf Brik and Chi-Huey Wong. Hiv-1 protease: mechanism and drug discovery. *Org. Biomol. Chem.*, 1:5–14, 2003.
- [24] T R Brown, B M Kincaid, and K Ugurbil. Nmr chemical shift imaging in three dimensions. *Proceedings of the National Academy of Sciences*, 79(11):3523–3526, 1982.

- [25] Martin Buehrer, Klaas P. Pruessmann, Peter Boesiger, and Sebastian Kozerke. Array compression for mri with large coil arrays. *Magnetic Resonance in Medicine*, 57(6):1131–1139, 2007.
- [26] Brian Burns, Neil E. Wilson, Jon K. Furuyama, and M. Albert Thomas. Non-uniformly under-sampled multi-dimensional spectroscopic imaging in vivo: maximum entropy versus compressed sensing reconstruction. *NMR in Biomedicine*, 27(2):191–201, 2014.
- [27] Brian L. Burns, Neil E. Wilson, and M. Albert Thomas. Group sparse reconstruction of multi-dimensional spectroscopic imaging in human brain in vivo. *Algorithms*, 7(3):276–294, 2014.
- [28] E.J. Candés, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489 – 509, feb. 2006.
- [29] Peng Cao and Ed X. Wu. Accelerating phase-encoded proton mr spectroscopic imaging by compressed sensing. *Journal of Magnetic Resonance Imaging*, 41(2):487–495, 2015.
- [30] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Phys. Rev.*, 94:630–638, May 1954.
- [31] John Cavanagh, Wayne J Fairbrother, Arthur G Palmer III, and Nicholas J Skelton. *Protein NMR spectroscopy: principles and practice*. Academic Press, 1995.
- [32] S Cavassila, S Deval, C Huegen, D van Ormondt, and D Graveron-Demilly. Cramér-rao bounds: an evaluation tool for quantitation. *NMR in Biomed.*, 14:278–283, 2001.
- [33] KM Cecil and BV Jones. Magnetic resonance spectroscopy of the pediatric brain. *Topics in Magnetic Resonance Imaging*, 12:435–452, 2001.
- [34] L Chang, T Ernst, M Leonido-Yee, I Walot, and E Singer. Cerebral metabolite abnormalities correlate with clinical severity of hiv-1 cognitive motor complex. *Neurology*, 52:100–108, 1999.
- [35] L. Chang, P.L. Lee, C.T. Yiannoutsos, T. Ernst, C.M. Marra, T. Richards, D. Kolson, G. Schifitto, J.G. Jarvik, E.N. Miller, R. Lenkinski, G. Gonzalez, and B.A. Navia. A multicenter in vivo proton-mrs study of hiv-associated dementia and its relationship to age. *NeuroImage*, 23(4):1336 – 1347, 2004.
- [36] Linda Chang, Thomas Ernst, Mallory D Witt, Nina Ames, Megan Gaiefsky, and Eric Miller. Relationships among brain metabolites, cognitive function, and viral loads in antiretroviral-naïve {HIV} patients. *NeuroImage*, 17(3):1638 – 1648, 2002.

- [37] Y. Chen, W. Hager, F. Huang, D. Phan, X. Ye, and W. Yin. Fast algorithms for image reconstruction with application to partially parallel mr imaging. *SIAM Journal on Imaging Sciences*, 5:90, 2012.
- [38] Changho Choi, Nicholas J. Coupland, Paramjit P. Bhardwaj, Sanjay Kalra, Colin A. Casault, Kim Reid, and Peter S. Allen. T2 measurement and quantification of glutamate in human brain in vivo. *Magnetic Resonance in Medicine*, 56(5):971–977, 2006.
- [39] R. Compton, S. Osher, and L. Bouchard. Hybrid regularization for mri reconstruction with static field inhomogeneity correction. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 650–655, May 2012.
- [40] Steven Conolly, Gary Glover, Dwight Nishimura, and Albert Macovski. A reduced power selective adiabatic spin-echo pulse sequence. *Magnetic Resonance in Medicine*, 18(1):28–38, 1991.
- [41] Steven Conolly, Dwight Nishimura, Albert Macovski, and Gary Glover. Variable-rate selective excitation. *Journal of Magnetic Resonance*, 78(3):440 – 458, 1988.
- [42] Charles H. Cunningham, Daniel B. Vigneron, Albert P. Chen, Duan Xu, Sarah J. Nelson, Ralph E. Hurd, Douglas A. Kelley, and John M. Pauly. Design of flyback echo-planar readout gradients for magnetic resonance spectroscopic imaging. *Magnetic Resonance in Medicine*, 54(5):1286–1289, 2005.
- [43] R. de Beer, A. van den Boogaart, D. van Ormondt, W. W. F. Pijnappel, J. A. den Hollander, A. J. H. Marien, and P. R. Luyten. Application of time-domain fitting in the quantification of in vivo 1h spectroscopic imaging data sets. *NMR in Biomedicine*, 5(4):171–178, 1992.
- [44] R.A. de Graaf. *In Vivo NMR Spectroscopy: Principles and Techniques*. Wiley, 1998.
- [45] Robin A. De Graaf and Klaas Nicolay. Adiabatic rf pulses: Applications to in vivo nmr. *Concepts in Magnetic Resonance*, 9(4):247–268, 1997.
- [46] Wei Deng, Wotao Yin, and Yin Zhang. Group sparse optimization by alternating direction method. In *Proc. SPIE*, volume 8858, pages 88580R–88580R–15, 2013.
- [47] Xiao-Qi Ding, Andrew A. Maudsley, Mohammad Sabati, Sulaiman Sheriff, Paulo R. Dellani, and Heinrich Lanfermann. Reproducibility and reliability of short-te whole-brain mr spectroscopic imaging of human brain at 3t. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2014.

- [48] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, april 2006.
- [49] Wolfgang Dreher, Peter Erhard, and Dieter Leibfritz. Fast three-dimensional proton spectroscopic imaging of the human brain at 3 t by combining spectroscopic missing pulse steady-state free precession and echo planar spectroscopic imaging. *Magnetic Resonance in Medicine*, 66(6):1518–1525, 2011.
- [50] Wolfgang Dreher and Dieter Leibfritz. Detection of homonuclear decoupled in vivo proton nmr spectra using constant time chemical shift encoding: Ct-press. *Magnetic Resonance Imaging*, 17(1):141–150, 1999.
- [51] M.F. Duarte, Shriram Sarvotham, D. Baron, M.B. Wakin, and R.G. Baraniuk. Distributed compressed sensing of jointly sparse signals. In *Signals, Systems and Computers, 2005. Conference Record of the Thirty-Ninth Asilomar Conference on*, pages 1537–1541, October 2005.
- [52] Ulrike Dydak, Markus Weiger, Klaas P. Pruessmann, Dieter Meier, and Peter Boesiger. Sensitivity-encoded spectroscopic imaging. *Magnetic Resonance in Medicine*, 46(4):713–722, 2001.
- [53] Andreas Ebel, Andrew A. Maudsley, Michael W. Weiner, and Norbert Schuff. Achieving sufficient spectral bandwidth for volumetric 1h echo-planar spectroscopic imaging at 4 tesla. *Magnetic Resonance in Medicine*, 54(3):697–701, 2005.
- [54] Richard A. E. Edden and Peter B. Barker. If j doesn't evolve, it won't j-resolve: J-press with bandwidth-limited refocusing pulses. *Magnetic Resonance in Medicine*, 65(6):1509–1514, 2011.
- [55] R. R. Ernst and W. A. Anderson. Application of fourier transform spectroscopy to magnetic resonance. *Review of Scientific Instruments*, 37(1), 1966.
- [56] Thomas Ernst, Caroline S. Jiang, Helenna Nakama, Steven Buchthal, and Linda Chang. Lower brain glutamate is associated with cognitive deficits in hiv patients: A new mechanism for hiv-associated neurocognitive disorder. *Journal of Magnetic Resonance Imaging*, 32(5):1045–1053, 2010.
- [57] R. Eslami and M. Jacob. Robust reconstruction of mrsi data using a sparse spectral model and high resolution mri priors. *Medical Imaging, IEEE Transactions on*, 29(6):1297–1309, June 2010.
- [58] Ernie Esser. Applications of lagrangian-based alternating direction methods and connections to split bregman. *CAM report*, 9:31, 2009.

- [59] Alexander Fuchs, Peter Boesiger, Rolf F. Schulte, and Anke Henning. Profit revisited. *Magnetic Resonance in Medicine*, 71(2):458–468, 2014.
- [60] Jon K. Furuyama, Rajakumar Nagarajan, Christian K. Roberts, Cathy C. Lee, Theodore J. Hahn, and M. Albert Thomas. A pilot validation of multi-echo based echo-planar correlated spectroscopic imaging in human calf muscles. *NMR in Biomedicine*, 27(10):1176–1183, 2014.
- [61] Jon K. Furuyama, Neil E. Wilson, Brian L. Burns, Rajakumar Nagarajan, Daniel J. Margolis, and M. Albert Thomas. Application of compressed sensing to multidimensional spectroscopic imaging in human prostate. *Magnetic Resonance in Medicine*, 67(6):1499–1505, 2012.
- [62] Jon K. Furuyama, Neil E. Wilson, and M. Albert Thomas. Spectroscopic imaging using concentrically circular echo-planar trajectories in vivo. *Magnetic Resonance in Medicine*, 67(6):1515–1522, 2012.
- [63] Michael Garwood and Lance DelaBarre. The return of the frequency sweep: Designing adiabatic pulses for contemporary nmr. *Journal of Magnetic Resonance*, 153(2):155 – 177, 2001.
- [64] Michael Garwood and Yong Ke. Symmetric pulses to induce arbitrary flip angles with compensation for rf inhomogeneity and resonance offsets. *Journal of Magnetic Resonance*, 94(3):511 – 525, 1991.
- [65] Sairam Geethanath, Hyeon-Man Baek, Sandeep K. Ganji, Yao Ding, Elizabeth A. Maher, Robert D. Sims, Changho Choi, Matthew A. Lewis, and Vikram D. Kodibagkar. Compressive sensing could accelerate 1h mr metabolic imaging in the clinic. *Radiology*, 262(3):985–994, 2012. PMID: 22357898.
- [66] T. Goldstein and S. Osher. The split bregman method for l1 regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.
- [67] Varan Govind, Stuart Gold, Krithica Kaliannan, Gaurav Saigal, Steven Falcone, Kristopher L Arheart, Leo Harris, Jonathan Jagid, and Andrew A Maudsley. Whole-brain proton mr spectroscopic imaging of mild-to-moderate traumatic brain injury and correlation with neuropsychological deficits. *Journal of neurotrauma*, 27(3):483496, March 2010.
- [68] Varanavasi Govindaraju, Karl Young, and Andrew A. Maudsley. Proton nmr chemical shifts and coupling constants for brain metabolites. *NMR in Biomedicine*, 13(3):129–153, 2000.
- [69] Varanavasi Govindaraju, Karl Young, and Andrew A. Maudsley. Proton nmr chemical shifts and coupling constants for brain metabolites. *NMR in Biomedicine*, 13(3):129–153, 2000.

- [70] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.
- [71] Meng Gu, Dong-Hyun Kim, Dirk Mayer, Edith V. Sullivan, Adolf Pfefferbaum, and Daniel M. Spielman. Reproducibility study of whole-brain 1h spectroscopic imaging with automated quantification. *Magnetic Resonance in Medicine*, 60(3):542–547, 2008.
- [72] HKon Gudbjartsson and Samuel Patz. The rician distribution of noisy mri data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995.
- [73] Ulrich L. Günther, Christian Ludwig, and Heinz Rüterjans. Wavewat-improved solvent suppression in nmr spectra employing wavelet transforms. *Journal of Magnetic Resonance*, 156(1):19 – 25, 2002.
- [74] Per Christian Hansen and Dianne Prost OLeary. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 14(6):1487–1503, 1993.
- [75] Lars G. Hanson. Is quantum mechanics necessary for understanding magnetic resonance? *Concepts in Magnetic Resonance Part A*, 32A(5):329–340, 2008.
- [76] C.J Hardy, W.A Edelstein, and D Vatis. Efficient adiabatic fast passage for nmr population inversion in the presence of radiofrequency field inhomogeneity and frequency offsets. *Journal of Magnetic Resonance*, 66(3):470 – 482, 1986.
- [77] Jarek Harezlak, Steven Buchthal, Michael Taylor, Giovanni Schifitto, Jianhui Zhong, Eric Daar, Jeffrey Alger, Elyse Singer, Thomas Campbell, Constantin Yiannoutsos, et al. Persistence of hiv-associated cognitive impairment, inflammation, and neuronal injury in era of highly active antiretroviral treatment. *Aids*, 25(5):625–633, 2011.
- [78] Christine I. Haupt, Norbert Schuff, Michael W. Weiner, and Andrew A. Maudsley. Removal of lipid artifacts in 1h spectroscopic imaging by data extrapolation. *Magnetic Resonance in Medicine*, 35(5):678–687, 1996.
- [79] Jrgen Hennig, Thorsten Thiel, and Oliver Speck. Improved sensitivity to overlapping multiplet signals in in vivo proton spectroscopy using a multiecho volume selective (cpress) experiment. *Magnetic Resonance in Medicine*, 37(6):816–820, 1997.
- [80] H P Hetherington, M J Avison, and R G Shulman. 1h homonuclear editing of rat brain using semiselective pulses. *Proceedings of the National Academy of Sciences*, 82(10):3115–3118, 1985.

- [81] Simon Hu, Michael Lustig, Albert P. Chen, Jason Crane, Adam Kerr, Douglas A.C. Kelley, Ralph Hurd, John Kurhanewicz, Sarah J. Nelson, John M. Pauly, and Daniel B. Vigneron. Compressed sensing for resolution enhancement of hyperpolarized ^{13}C flyback 3d-mrsi. *Journal of Magnetic Resonance*, 192(2):258 – 264, 2008.
- [82] Ralph Hurd, Napapon Sailasuta, Radhika Srinivasan, Daniel B. Vigneron, Daniel Pelletier, and Sarah J. Nelson. Measurement of brain glutamate using te-averaged press at 3t. *Magnetic Resonance in Medicine*, 51(3):435–440, 2004.
- [83] Sven G. Hyberts, Koh Takeuchi, and Gerhard Wagner. Poisson-gap sampling and forward maximum entropy reconstruction for enhancing the resolution and sensitivity of protein nmr data. *Journal of the American Chemical Society*, 132(7):2145–2147, 2010. PMID: 20121194.
- [84] J.I. Jackson, C.H. Meyer, D.G. Nishimura, and A. Macovski. Selection of a convolution function for fourier inversion using gridding. *Medical Imaging, IEEE Transactions on*, 10(3):473 –478, sep 1991.
- [85] Wenwen Jiang, Michael Lustig, and Peder E. Z. Larson. Concentric rings k-space trajectory for hyperpolarized ^{13}C mr spectroscopic imaging. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2014.
- [86] Anton J Johnson, Michael Garwood, and Kmil Uurbil. Slice selection with gradient-modulated adiabatic excitation despite the presence of large b_1 inhomogeneities. *Journal of Magnetic Resonance*, 81(3):653 – 660, 1989.
- [87] Lana G. Kaiser, Norbert Schuff, Nathan Cashdollar, and Michael W. Weiner. Scyllo-inositol in normal aging human brain: ^1H magnetic resonance spectroscopy study at 4 tesla. *NMR in Biomedicine*, 18(1):51–55, 2005.
- [88] James Keeler. *Understanding NMR spectroscopy*. John Wiley & Sons, 2011.
- [89] Uwe Klose. In vivo proton spectroscopy in presence of eddy currents. *Magnetic Resonance in Medicine*, 14(1):26–30, 1990.
- [90] Roland Kreis and Chris Boesch. Liquid-crystal-like structures of human muscle demonstrated by in vivo observation of direct dipolar coupling in localized proton magnetic resonance spectroscopy. *Journal of Magnetic Resonance, Series B*, 104(2):189–192, 1994.
- [91] J Laubenberger, D Haussinger, S Bayer, S Thielemann, B Schneider, A Mundinger, J Hennig, and M Langer. Hiv-related metabolic abnormalities in the brain: depiction with proton mr spectroscopy with short echo times. *Radiology*, 199:805–810, 1996.

- [92] M. Louis Lauzon and Brian K. Rutt. Effects of polar sampling in k-space. *Magnetic Resonance in Medicine*, 36(6):940–949, 1996.
- [93] M. Louis Lauzon and Brian K. Rutt. Polar sampling in k-space: Reconstruction effects. *Magnetic Resonance in Medicine*, 40(5):769–782, 1998.
- [94] Malcolm H Levitt. *Spin dynamics: basics of nuclear magnetic resonance*. John Wiley & Sons, 2001.
- [95] Malcolm H. Levitt and Ray Freeman. Nmr population inversion using a composite pulse. *Journal of Magnetic Resonance*, 33(2):473 – 476, 1979.
- [96] Dong Liang, Bo Liu, Jiunjie Wang, and Leslie Ying. Accelerating sense using compressed sensing. *Magn Reson Med*, 62(6):1574–84, Dec 2009.
- [97] Fa-Hsuan Lin, Kenneth K. Kwong, John W. Belliveau, and Lawrence L. Wald. Parallel imaging reconstruction using automatic regularization. *Magnetic Resonance in Medicine*, 51(3):559–567, 2004.
- [98] Fa-Hsuan Lin, Shang-Yueh Tsai, Ricardo Otazo, Arvind Caprihan, Lawrence L. Wald, John W. Belliveau, and Stefan Posse. Sensitivity-encoded (sense) proton echo-planar spectroscopic imaging (pepsi) in the human brain. *Magnetic Resonance in Medicine*, 57(2):249–257, 2007.
- [99] Scott Lipnick, Gaurav Verma, Saadallah Ramadan, Jon Furuyama, and M. Albert Thomas. Echo planar correlated spectroscopic imaging: Implementation and pilot evaluation in human calf in vivo. *Magnetic Resonance in Medicine*, 64(4):947–956, 2010.
- [100] Stig Ljunggren. A simple graphical representation of fourier-based imaging methods. *Journal of Magnetic Resonance*, 54(2):338 – 343, 1983.
- [101] L.M. and Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.
- [102] Michael Lustig, David Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- [103] Michael Lustig and John M. Pauly. Spirit: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic Resonance in Medicine*, 64(2):457–471, 2010.
- [104] James R. MacFall, Norbert J. Pelc, and Robert M. Vavrek. Correction of spatially dependent phase shifts for partial fourier imaging. *Magnetic Resonance Imaging*, 6(2):143 – 155, 1988.

- [105] A Macovski. Noise in mri. *Magn. Reson. Med.*, 36:494–497, 1996.
- [106] Slobodon Macura and Larry R Brown. Improved sensitivity and resolution in two-dimensional homonuclear j-resolved nmr spectroscopy of macromolecules. *Journal of Magnetic Resonance*, 53(3):529 – 535, 1983.
- [107] P Mansfield. Spatial mapping of the chemical shift in nmr. *J. Phys. D. Appl. Phys.*, 16:L235–L238, 1983.
- [108] P. Mansfield. Spatial mapping of the chemical shift in nmr. *Magnetic Resonance in Medicine*, 1(3):370–386, 1984.
- [109] Jintong Mao, T.H Mareci, and E.R Andrew. Experimental study of optimal selective 180 radiofrequency pulses. *Journal of Magnetic Resonance*, 79(1):1 – 10, 1988.
- [110] S Matsui, K Sekihara, and H Kohno. Spatially resolved nmr spectroscopy using phase-modulated spin-echo trains. *Journal of Magnetic Resonance (1969)*, 67(3):476 – 490, 1986.
- [111] A.A. Maudsley, C. Domenig, V. Govind, A. Darkazanli, C. Studholme, K. Arheart, and C. Bloomer. Mapping of brain metabolite distributions by volumetric proton mr spectroscopic imaging (mrsi). *Magnetic Resonance in Medicine*, 61(3):548–559, 2009.
- [112] A.A Maudsley, S.K Hilal, W.H Perman, and H.E Simon. Spatially resolved high resolution spectroscopy by four-dimensional {NMR}. *Journal of Magnetic Resonance (1969)*, 51(1):147 – 152, 1983.
- [113] Andrew A. Maudsley, G. B. Matson, J. W. Hugg, and M. W. Weiner. Reduced phase encoding in spectroscopic imaging. *Magnetic Resonance in Medicine*, 31(6):645–651, 1994.
- [114] Gregory Metzger and Xiaoping Hu. Application of interlaced fourier transform to echo-planar spectroscopic imaging. *Journal of Magnetic Resonance*, 125(1):166–170, 1997.
- [115] D Meyerhoff, C Bloomer, V Cardenas, D Norman, M Weiner, and G Fein. Elevated subcortical choline metabolites in cognitively and clinically asymptomatic hiv patients. *Neurology*, 52:995–1003, 1999.
- [116] K. Nagayama, K. Wthrich, and R.R. Ernst. Two-dimensional spin echo correlated spectroscopy (secsy) for 1h nmr studies of biological macromolecules. *Biochemical and Biophysical Research Communications*, 90(1):305 – 311, 1979.

- [117] R.J. Ogg, R.B. Kingsley, and J.S. Taylor. Wet, a t1- and b1-insensitive water-suppression method for in vivo localized 1h nmr spectroscopy. *Journal of Magnetic Resonance, Series B*, 104(1):1 – 10, 1994.
- [118] Roger J. Ordidge, Marzena Wylezinska, James W. Hugg, Edward Butterworth, and Florence Franconi. Frequency offset corrected inversion (foci) pulses for use in localized spectroscopy. *Magnetic Resonance in Medicine*, 36(4):562–566, 1996.
- [119] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Modeling and Simulation*, 4(2):460–489, 2005.
- [120] Ricardo Otazo, Bryon Mueller, Kamil Ugurbil, Lawrence Wald, and Stefan Posse. Signal-to-noise ratio and spectral linewidth improvements between 1.5 and 7 tesla in proton echo-planar spectroscopic imaging. *Magnetic Resonance in Medicine*, 56(6):1200–1210, 2006.
- [121] Ricardo Otazo, Shang-Yueh Tsai, Fa-Hsuan Lin, and Stefan Posse. Accelerated short-te 3d proton echo-planar spectroscopic imaging using 2d-sense with a 32-channel array coil. *Magnetic Resonance in Medicine*, 58(6):1107–1116, 2007.
- [122] Jullie W. Pan, Graeme F. Mason, Gerald M. Pohost, and Hoby P. Hetherington. Spectroscopic imaging of human brain glutamate by water-suppressed j-refocused coherence transfer at 4.1 t. *Magnetic Resonance in Medicine*, 36(1):7–12, 1996.
- [123] J. Pauly, P. Le Roux, Dwight Nishimura, and A. Macovski. Parameter relations for the shinnar-le roux selective excitation pulse design algorithm. *Medical Imaging, IEEE Transactions on*, 10(1):53–65, Mar 1991.
- [124] Dana C. Peters, Frank R. Korosec, Thomas M. Grist, Walter F. Block, James E. Holden, Karl K. Vigen, and Charles A. Mistretta. Undersampled projection reconstruction applied to mr angiography. *Magnetic Resonance in Medicine*, 43(1):91–101, 2000.
- [125] James Pipe Mr and Jeffrey L. Duerk. Analytical resolution and noise characteristics of linearly reconstructed magnetic resonance data with arbitraryk-space sampling. *Magnetic Resonance in Medicine*, 34(2):170–178, 1995.
- [126] S Posse, C DeCarli, and D Le Bihan. Three-dimensional echo-planar mr spectroscopic imaging at short echo times in the human brain. *Radiology*, 192(3):733–738, 1994. PMID: 8058941.
- [127] Stefan Posse, Gioacchino Tedeschi, Robert Risinger, Robert Ogg, and Denis Le Bihan. High speed 1h spectroscopic imaging in human brain by

- echo planar spatial-spectral encoding. *Magnetic Resonance in Medicine*, 33(1):34–40, 1995.
- [128] C. Prieto, M. Usman, J. M. Wild, S. Kozerke, P. G. Batchelor, and T. Schaeffter. Group sparse reconstruction using intensity-based clustering. *Magnetic Resonance in Medicine*, 69(4):1169–1179, 2013.
- [129] Stephen W. Provencher. Estimation of metabolite concentrations from localized in vivo proton nmr spectra. *Magnetic Resonance in Medicine*, 30(6):672–679, 1993.
- [130] Stephen W. Provencher. Automatic quantitation of localized in vivo 1h spectra with lmodel. *NMR in Biomedicine*, 14(4):260–264, 2001.
- [131] Klaas P. Pruessmann, Markus Weiger, Peter Börnert, and Peter Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine*, 46(4):638–651, 2001.
- [132] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, and Peter Boesiger. Sense: Sensitivity encoding for fast mri. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999.
- [133] E. M. Purcell, H. C. Torrey, and R. V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Phys. Rev.*, 69:37–38, Jan 1946.
- [134] Saadallah Ramadan and Carolyn E. Mountford. Adiabatic localized correlation spectroscopy (al-cosy): Application in muscle and brain. *Journal of Magnetic Resonance Imaging*, 33(6):1447–1455, 2011.
- [135] S. Ramani, T. Blu, and M. Unser. Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms. *Image Processing, IEEE Transactions on*, 17(9):1540–1554, Sept 2008.
- [136] S. Ramani and J.A. Fessler. Parallel mr image reconstruction using augmented lagrangian methods. *Medical Imaging, IEEE Transactions on*, 30(3):694–706, March 2011.
- [137] S. Ramani, Zhihao Liu, J. Rosen, J. Nielsen, and J.A. Fessler. Regularization parameter selection for nonlinear iterative image restoration and mri reconstruction using gcv and sure-based methods. *Image Processing, IEEE Transactions on*, 21(8):3659–3672, Aug 2012.
- [138] Marc S. Ramirez, Jaehyuk Lee, Christopher M. Walker, Vlad C. Sandulache, Franciszek Hennel, Stephen Y. Lai, and James A. Bankson. Radial spectroscopic mri of hyperpolarized [1-13c] pyruvate at 7 tesla. *Magnetic Resonance in Medicine*, 72(4):986–995, 2014.

- [139] Jimin Ren, A. Dean Sherry, and Craig R. Malloy. 1h mrs of intramyocellular lipids in soleus muscle at 7 t: Spectral simplification by using long echo times without water suppression. *Magnetic Resonance in Medicine*, 64(3):662–671, 2010.
- [140] Joseph R. Roebuck, David O. Hearshen, Matthew O’Donnell, and Tom Raidy. Correction of phase effects produced by eddy currents in solvent suppressed 1h-csi. *Magnetic Resonance in Medicine*, 30(3):277–282, 1993.
- [141] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The nmr phased array. *Magnetic Resonance in Medicine*, 16(2):192–225, 1990.
- [142] D L Rothman, K L Behar, H P Hetherington, and R G Shulman. Homonuclear 1h double-resonance difference spectroscopy of the rat brain in vivo. *Proceedings of the National Academy of Sciences*, 81(20):6330–6334, 1984.
- [143] DA Rottenberg, JR Moeller, SC Strother, JJ Sidtis, BA Navia, V Dhawan, JZ Ginos, and RW Price. The metabolic pathology of the aids dementia complex. *Annals of neurology*, 22(6):700–706, 1987.
- [144] David A. Rottenberg, John J. Sidtis, Stephen C. Strother, Kirt A. Schaper, Jon R. Anderson, Mary J. Nelson, and Richard W. Price. Abnormal cerebral glucose metabolism in hiv-1 seropositive subjects with and without dementia. *Journal of Nuclear Medicine*, 37(7):1133–1141, 1996.
- [145] Bhaswati Roy, RakeshK. Gupta, AndrewA. Maudsley, Rishi Awasthi, Sulaiman Sheriff, Meng Gu, Nuzhat Husain, Sudipta Mohakud, Sanjay Behari, ChandraM. Pandey, RamK.S. Rathore, DanielM. Spielman, and JeffryR. Alger. Utility of multiparametric 3-t mri for glioma characterization. *Neuroradiology*, 55(5):603–613, 2013.
- [146] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(14):259 – 268, 1992.
- [147] Lawrence N. Ryner, James A. Sorenson, and M.Albert Thomas. Localized 2d j-resolved 1h mr spectroscopy: Strong coupling effects in vitro and in vivo. *Magnetic Resonance Imaging*, 13(6):853 – 869, 1995.
- [148] LN Ryner, JA Sorenson, and MA Thomas. 3d localized 2d nmr spectroscopy on an mri scanner. *J Magn. Reson. B*, 107:126–137, 1995.
- [149] T Scarabino, T Popolizio, A Bertolino, and U Salvolini. Proton magnetic resonance spectroscopy of the brain in pediatric patients. *European Journal of Radiology*, 30:142–153, 1999.

- [150] Tom W. J. Scheenen, Dennis W. J. Klomp, Jannie P. Wijnen, and Arend Heerschap. Short echo time 1h-mrsi of the human brain at 3t with minimal chemical shift displacement errors using adiabatic refocusing pulses. *Magnetic Resonance in Medicine*, 59(1):1–6, 2008.
- [151] Klaus Scheffler and Jrgen Hennig. Reduced circular field-of-view imaging. *Magnetic Resonance in Medicine*, 40(3):474–480, 1998.
- [152] RF Schulte and P Boesiger. Profit: two-dimensional prior-knowledge fitting of j-resolved spectra. *NMR in Biomedicine*, 19:255–263, 2006.
- [153] Rolf F. Schulte and Peter Boesiger. Profit: two-dimensional prior-knowledge fitting of j-resolved spectra. *NMR in Biomedicine*, 19(2):255–263, 2006.
- [154] Rolf F. Schulte, Thomas Lange, Johannes Beck, Dieter Meier, and Peter Boesiger. Improved two-dimensional j-resolved spectroscopy. *NMR in Biomedicine*, 19(2):264–270, 2006.
- [155] M.S Silver, R.I Joseph, and D.I Hoult. Highly selective 2 and pulse generation. *Journal of Magnetic Resonance*, 59(2):347 – 351, 1984.
- [156] Saryu Singh, Brian K Rutt, and R Mark Henkelman. Projection presaturation: a fast and accurate technique for multidimensional spatial localization. *Journal of Magnetic Resonance (1969)*, 87(3):567–583, 1990.
- [157] Charles P Slichter. *Principles of magnetic resonance*, volume 1. Springer Science & Business Media, 1990.
- [158] J Slotboom, A.F Mehlkopf, and W.M.M.J Bove. A single-shot localization pulse sequence suited for coils with inhomogeneous rf fields using adiabatic slice-selective rf pulses. *Journal of Magnetic Resonance*, 95(2):396 – 404, 1991.
- [159] Johannes Slotboom, Chris Boesch, and Roland Kreis. Versatile frequency domain fitting using time domain models and prior knowledge. *Magnetic Resonance in Medicine*, 39(6):899–911, 1998.
- [160] S.A. Smith, T.O. Levante, B.H. Meier, and R.R. Ernst. Computer simulations in magnetic resonance. an object-oriented programming approach. *Journal of Magnetic Resonance, Series A*, 106(1):75 – 105, 1994.
- [161] Daniel K. Sodickson and Warren J. Manning. Simultaneous acquisition of spatial harmonics (smash): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38(4):591–603, 1997.
- [162] Daniel M. Spielman, John M. Pauly, Albert Macovski, Gary H. Glover, and Dieter R. Enzmann. Lipid-suppressed single-and multisection proton

- spectroscopic imaging of the human brain. *Journal of Magnetic Resonance Imaging*, 2(3):253–262, 1992.
- [163] Radhika Srinivasan, Charles Cunningham, Albert Chen, Daniel Vigneron, Ralph Hurd, Sarah Nelson, and Daniel Pelletier. Te-averaged two-dimensional proton spectroscopic imaging of glutamate at 3 t. *NeuroImage*, 30(4):1171 – 1178, 2006.
- [164] RS Staewen, AJ Johnson, BD Ross, T Parrish, H Merkle, and M Garwood. 3-d flash imaging using a single surface coil and a new adiabatic pulse, bir-4. *Investigative radiology*, 25(5):559–567, May 1990.
- [165] Charlotte J. Stagg, Steven Knight, Kevin Talbot, Mark Jenkinson, Andrew A. Maudsley, and Martin R. Turner. Whole-brain magnetic resonance spectroscopic imaging measures are related to disability in als. *Neurology*, 80(7):610–615, 2013.
- [166] Alan S. Stern, David L. Donoho, and Jeffrey C. Hoch. Nmr data processing using iterative thresholding and minimum l1-norm reconstruction. *Journal of Magnetic Resonance*, 188(2):295 – 300, 2007.
- [167] Alberto Tannus and Michael Garwood. Improved performance of frequency-swept pulses using offset-independent adiabaticity. *Journal of Magnetic Resonance, Series A*, 120(1):133 – 137, 1996.
- [168] Alberto Tannus and Michael Garwood. Adiabatic pulses. *NMR in Biomedicine*, 10(8):423–434, 1997.
- [169] M Tardieu, J Le Chenadec, A Persoz, L Meyer, S Blanche, and M Mayaux. Hiv-1-related encephalopathy in infants compared with children and adults. *Neurology*, 54(5):1089–1095, 2000.
- [170] M. Albert Thomas, Hyun-Kyung Chung, and Holly Middlekauff. Localized two-dimensional 1h magnetic resonance exchange spectroscopy: A preliminary evaluation in human muscle. *Magnetic Resonance in Medicine*, 53(3):495–502, 2005.
- [171] M. Albert Thomas, Kenneth Yue, Nader Binesh, Pablo Davanzo, Anand Kumar, Benjamin Siegel, Mark Frye, John Curran, Robert Lufkin, Paul Martin, and Barry Guze. Localized two-dimensional shift correlated mr spectroscopy of human brain. *Magnetic Resonance in Medicine*, 46(1):58–67, 2001.
- [172] Kamil Ugurbil, Michael Garwood, and M Robin Bendall. Amplitude- and frequency-modulated pulses to achieve 90 plane rotations with inhomogeneous b1 fields. *Journal of Magnetic Resonance*, 72(1):177 – 185, 1987.

- [173] Kamil Ugurbil, Michael Garwood, and Alan R Rath. Optimization of modulation functions to improve insensitivity of adiabatic pulses to variations in b1 magnitude. *Journal of Magnetic Resonance*, 80(3):448 – 469, 1988.
- [174] M. Usman, C. Prieto, T. Schaeffter, and P. G. Batchelor. k-t group sparse: A method for accelerating dynamic mri. *Magnetic Resonance in Medicine*, 66(4):1163–1176, 2011.
- [175] J. W. C. van der Veen, R. de Beer, P. R. Luyten, and D. van Ormondt. Accurate quantification of in vivo 31p nmr signals using the variable projection method and prior knowledge. *Magnetic Resonance in Medicine*, 6(1):92–98, 1988.
- [176] S.S. Vasanawala, M.J. Murphy, M.T. Alley, P. Lai, K. Keutzer, J.M. Pauly, and M. Lustig. Practical parallel imaging compressed sensing mri: Summary of two years of experience in accelerating body mri of pediatric patients. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1039–1043, March 2011.
- [177] Gaurav Verma, John H. Woo, Sanjeev Chawla, Sumei Wang, Sulaiman Sheriff, Lauren B. Elman, Leo F. McCluskey, Murray Grossman, Elias R. Melhem, Andrew A. Maudsley, and Harish Poptani. Whole-brain analysis of amyotrophic lateral sclerosis by using echo-planar spectroscopic imaging. *Radiology*, 267(3):851–857, 2013. PMID: 23360740.
- [178] Neil E. Wilson, Zohaib Iqbal, Brian L. Burns, Margaret Keller, and M. Albert Thomas. Accelerated five-dimensional echo planar j-resolved spectroscopic imaging: Implementation and pilot validation in human brain. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2015.
- [179] Hochong H. Wu, Jin Hyung Lee, and Dwight G. Nishimura. Mri using a concentric rings trajectory. *Magnetic Resonance in Medicine*, 59(1):102–112, 2008.
- [180] W. Yin, S. Osher, D. Goldfarb, and J. Darbon. Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing. *SIAM Journal on Imaging Sciences*, 1(1):143–168, 2008.
- [181] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [182] Tao Zhang, John M. Pauly, Shreyas S. Vasanawala, and Michael Lustig. Coil compression for accelerated imaging with cartesian sampling. *Magnetic Resonance in Medicine*, 69(2):571–582, 2013.