Abstract

New Strategies for Accelerated Spatial Encoding with Quadratic Fields in Magnetic Resonance Imaging

Jason Peter Stockmann

2012

Spatial encoding with nonlinear magnetic fields has recently drawn attention for its potential to achieve faster gradient field switching within safety limits, tailored resolution in regions of interest, and improved parallel imaging using encoding fields that complement the sensitivity profiles of RF receive arrays. Proposed methods can broadly be divided into those that employ phase encoding (PatLoc and COGNAC) and those that acquire nonlinear projections (O-Space, Null Space imaging, and 4D-RIO). Projection methods typically reconstruct images using iterative algorithms to backproject the data by exploiting the full encoding matrix. For this reason, they are more sensitive than phase encoding methods to systematic errors introduced by imperfect knowledge of the encoding trajectory.

In the present work, voxel-wise phase evolution is mapped at each acquired point in an O-Space trajectory using a variant of chemical shift imaging, capturing all spin dynamics caused by encoding fields, eddy currents, and pulse timing. Phase map calibration is then applied to data acquired from a high-power, 12 cm, Z2 insert coil with a home-built 8-channel RF transmit-receive array on a 3T human scanner. We show the first experimental proof-of-concept O-Space images on in vivo and phantom samples, paving the way for more in-depth exploration of O-Space and other nonlinear projection imaging methods. We also use quadratic phase preparation with a Cartesian pulse sequence to image form a localized field-of-view within the sample. This extends a previously presented method ("GradLoc") by creating localization with a 3-D encoding field and also by combining it with parallel imaging for scan acceleration.

New Strategies for Accelerated Spatial Encoding with Quadratic Fields in Magnetic Resonance Imaging

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by

Jason Peter Stockmann

Dissertation Director: R. Todd Constable

May, 2012

In memory of Mitchell Dubey, 1988-2011

Acknowledgements

This work is indebted to such a multitude of people that I feel more like its steward than owner.

First, I would like first to thank my faculty committee – Todd Constable, Richard Carson, and Douglas Rothman – for their guidance and their thoughtful feedback on this thesis. As my thesis advisor, Dr. Constable has always been extremely supportive of my work, encouraging me to take intellectual and experimental risks as we explored a very novel field of imaging research. I would also like to thank the outside reader of my thesis, John Pauly of Stanford University.

The members of Dr. Constable's laboratory have been involved in every facet of this research. The experimental results shown here would not have been possible without the expertise and perseverance of Gigi Galiana, whose flair for pulse sequences and spin gymnastics helped us overcome countless obstacles on the scanner. Leo Tam also played a big role in ramping up the nonlinear imaging experiments and offered many insights into the theory of nonlinear gradient imaging. Nolwenn Caillet, Dana Peters, and Maolin Qui have each also helped keep my experiments running on track. Pelin Aksit Ciris developed an important early computer simulation of O-Space image reconstruction, paving the way for this thesis. Jitendra Bhawnani was tireless in his Sisyphean effort to keep all the MRRC's computers running, including my own.

This research relied heavily upon a wide array of custom instrumentation. Piotr Starewicz and Bill Punchard of Resonance Research, Inc. assembled the powerful and extremely dependable Z^2 gradient insert that lies at the heart of all my experiments. Terry Nixon and Scott McIntyre at the MRRC provided the amplifiers, temperature sensors, digital electronics, and control software needed to interface the gradient insert with the Siemens 3T scanner. Peter Brown taught me many of the fundamentals of RF coil design and patiently helped me with the fabrication of the 8-channel transmit-receive coil array.

I also grateful to Robin de Graaf and Christoph Juchem for teaching me about field mapping and shimming, without which my experiments would have been impossible. Similarly, I owe the success of my O-Space image reconstructions to Hemant Tagare, who suggested the use of the Kaczmarz method. Dr. Tagare's deep expertise and mathematical observations were a tremendous resource in conducting this research.

This work also would not have been possible without generous support from the National Institutes of Health in the form of a Business Research Partnership Grant (NIH BRP

R01 EB012289) and Neuroengineering Training Grant (5T90DK070068-05).

I would also like to thank my collaborators at other institutions:

– The PatLoc team at Freiburg were pioneers in the field of nonlinear gradient encoding and I have learned a great deal from their group, notably Gerrit Schultz, Anna Welz, Daniel Gallichan, Sebastian Littin, Jakob Asslaender, Hans Weber, Chris Cocosco, Maxim Zaitsev, and Juergen Hennig. I am grateful to Dr. Zaitsev and Dr. Hennig for allowing me to visit Freiburg for three months under the auspices of the Baden-Wuerttemberg Exchange. Special thanks go to Anna Welz, for helping me get settled in Freiburg and for lending me her bicycle, and to Sebastian Littin, for hosting me in his apartment. The monsooned Malabar espresso that Sebastian brews is still the best I've ever tasted. I also want to thank Gerrit Schultz for his mathematical insights and for his expert proofreading of several sections of this thesis.

– Vicente Parot and Pablo Irarrazaval, my collaborators at the Catholic University of Santiago, Chile, for many invigorating exchanges on the subject of the fractional Fourier transform (FrFT). Their mathematical framework laid the basis for much of the material presented in Chapter 4 of this thesis, and for future work yet to come. I also thank Haldun Ozaktas of Bilkent University, Ankara, Turkey for illuminating discussions of the FrFT.

– Larry Wald of Massachusetts General Hospital for his guidance on RF coil array design and fabrication, and for his generous loan of pre-amplifiers and circuit boards for use in the array.

– Daniel Sodickson of New York University and Scott King of the Canadian National Research Council for helpful discussions about parallel imaging and RF coil design.

– My colleagues from the 18 months I spent before grad school at the Jet Propulsion Lab, especially Aluizio Prata, Suzanne Spitz, Richard Hodges, Mark Zawadzki, Neil Chamberlain, Jeff Harrell, Joe Vacchione, and Benny Venkatesan. Everyone in the antenna design group was a true pleasure to work with, and I have many fond memories both of JPL and of mountain biking excursions through the nearby San Gabriels.

I would also like to acknowledge OpenWetWare for the time-saving Ph.D. thesis template used to generate this thesis in LaTeX.

I also want to thank the many people who, while not collaborators, were just as indispensable in getting me through grad school, and in my upbringing:

– Alp Kucukelbir for getting me started on LaTeX and for many stimulating discussions. The passion and intelligence that Alp brings to his work and life are an inspiration to all who know him.

– Aimee Kanzler, Moran Furman, Nicha Chitphakdithai, Patricia Bimboese, Becky Bruccoleri, Farin Loeb, and Aaron Mertz, for sharing their musical energy and for giving me the chance to play chamber music on the piano.

– Randolph Miles and Laura Satveit Miles for giving me the opportunity – and the encouragement – to build a harpsichord from a kit!

– The many friends I have made through Yale Swing & Blues and Elm City Waltz. Thanks especially to Richard Scalzo, Jessica Chen Weiss, Valarie Kaur, Susan de Guardiola, Andrea Hodgins-Davis, Tyler Ames, Pete Molfese, Mattias Gustavsson, Dhruv Bhargava, Sara Packman, Jenn Santiago, Christian Siller, Jenn Bennett, Rosalind Diaz, Nora Ng, Elaine Lee, Eleanor Avrunin, Randi Cohen, Chris Oksanen, Kaya Pierce-Grove, Alex Xenakis, Priya Ghandi, Mary Christensen, and Emily Freed.

– Elizabeth Wytychak for her inspiring dancing and her sense of humor, and for having the courage to hotwire my car's hybrid battery system with me, resurrecting it from a near-death experience.

– Everyone I have had the pleasure of working with through ElmCityCycling, including Matthew Feiner, Mark Abraham, Tom Harned, Richard Stowe, David Streever, Melinda Tuhus, Rob Rocke, Liana Martingano, Mark Scott, Paul Proulx, Laura Burrone, and Patrina Pellet. It has been gratifying to see our advocacy help achieve bicycle improvements in New Haven, including bike lanes, bike hooks on Metro North trains, and bike parking at Union Station, to name only a few. More people than ever are using bicycles to get around town. ElmCityCycling is not only an advocacy group, but also a circle of friends, and it has been a true pleasure getting to know everyone involved in it.

– Scott Aaronson & Dana Moshkovitz, for their humor and intellectual *joie de vivre*, which brims and runs over, spreading to all who know them. They are inspiring and deeply valued friends.

– Matt Siegler & Rita Economos, for their warm friendship, their sense of mischief, and for enlivening the years I spent living in Pasadena.

– Henryk Jaronowski, for his wit and effervescence, and his exuberant words, both sung and spoken.

– Joe Tanniru, for understanding me, for his love of music, and for challenging my political assumptions and keeping me on my feet.

– Rose Radin and Jacob Hyman, for their empathy, kindness, and always thoughtful perspectives on wide-ranging topics.

– Alex & Kenny Liu, for many years of urban exploration, tennis matches, board gaming, and mayhem various and sundry.

– Mark, Libby, & Almanzo Abraham, for being the absolute best of friends, for teaching me a great deal about how to live well, and for too many other things to mention.

– The Websters – Gordon, Gloria, David, Diana, Daniel, Tera, Adeline, and Silvie – for a lifetime of friendship, fellowship, and love. They are truly family to me.

– Elizabeth Weinfield, for being my muse. I am forever enriched by the beautiful music and art that she has brought into my life.

– Everyone in my extended family, most especially Helen Popcun, at 87 still the greatest baker in Syracuse. I am so lucky for the warmth and love everyone in the family has shown me over my whole lifetime, including James, Elaine, and Joseph Popcun, Karin Laakso, Wanda & John Sever, Kristen & Pat Foley, Lauren & Bob Young, Michelle Popcun & Kevin Stack, Maryann & Dennis Kravitz, Ollie & Patti Popcun, Chris & Sue Sever, Lori & Jeff Krakoff, Vicki & Rick Wuest, and Brenda & Kenny Stone.

– My brother Matthew, with whom I can talk about anything, and my parents, Kathy and Peter, for their endless patience, support, encouragement, and love. My parents have lent me many insights from their careers in nursing and electrical engineering. It can't be pure coincidence that my work lies near the intersection of these two fields. My parents have always found their work deeply fulfilling and say they would still do it even if they weren't getting paid. They have encouraged to settle for nothing less in my own career.

Organization

This thesis is organized in five chapters. Chapter 1 offers a very abbreviated review of the fundamentals of spatial encoding in MRI as described in the existing literature. Special attention is given to parallel imaging methods that allow for reduced scan times by combining data from multiple RF receive coils. While some of the equations are expressed in a very general form, the assumption in Chapter 1 is of conventional spatial encoding performed using linear magnetic fields.

Chapter 2 introduces several encoding schemes that use nonlinear encoding fields, notably O-Space imaging, which is the primary subject of the thesis. The bulk of the chapter, Sections 2.1-2.4, is an elaboration of a paper published by the author in 2009 in *Magnetic Resonance in Medicine* (MRM). Section 2.4 briefly outlines a center placement optimization study was conducted by Pelin Aksit Ciris and was reported in greater detail at the *International Symposium for Magnetic Resonance in Medicine* (ISMRM) in 2009. Section 2.6 discusses a sensitivity analysis performed by the author on simulated O-Space images in the presence of sinus and ear cavity inhomogeneity. This work was presented by the author at ISMRM in 2010.

Chapter 3 goes on to describe experimental implementation of the O-Space encoding using a 12-cm diameter insert coil. Sections 3.1-3.6 are described in a journal paper submitted to MRM that is now in review. Section 3.7 presents additional analysis of the O-Space image reconstruction algorithm and its convergence properties that could not be included in the aforementioned MRM article due to length limitations. The material on through-plane dephasing in Section 3.4 is adapted from the work of Gigi Galiana, as published in MRM in 2011.

Chapter 4 describes three additional topics that involve the use of the quadratic field insert coil. Section 4.2 is based on an abstract presented by the author at ISMRM 2011 showing the equivalence of the Fresnel transform and fractional Fourier transforms for achieving scalable-FOV reconstructions in the presence of quadratic phase-scrambling fields (a technique that had previously been demonstrated by several authors). Section 4.3 describes an extension of a method published by another research group called "GradLoc" that uses quadratic phase-scrambling fields to achieve localized region-of-interest acquisitions without aliasing. In the present study, GradLoc is combined with parallel imaging

techniques to reduce the acquisition time, an approach that has not been previously published. The section also introduces quadratic RF pulse phase pre-compensation as a way to use three-dimensional encoding fields for GradLoc. Section 4.3 forms the basis of an abstract submitted to ISMRM 2012 now awaiting review. Section 4.4 revisits the question of O-Space image reconstruction, this time performed using an analytical approach based on the variable-order fractional Fourier transform (VO-FrFT). This work was described in an abstract presented by the author at ISMRM 2011. It builds on previous work from another research group that uses the VO-FrFT to correct for field inhomogeneities in conventional MR imaging.

Chapter 5 is a brief discussion of new work by other investigators who are studying nonlinear encoding fields. All material in this chapter has been presented at a conference or appeared in print within the last year, representing an emerging method.

Contents

List of Figures

List of Tables

Glossary

B_0 : static background field. 2

B_1 : transmitted radio frequency field applied by a coil, usually to generate spin flip angle. 3

E : encoding matrix used as the forward model of the signal acquisition process, evaluated on a discrete grid of source voxels in the object. Typically has size $n_K n_c \times n_\rho^2$. 11

F : reconstruction matrix that operates on the acquired data to produce an image. Typically of size $n_\rho \times n_K n_c$. 11

N : image size; the number of voxels along each dimension of a presumably square image.. 14

N_s : number of sampled acquired by the ADC, at bandwidth BW, during readout. xviii, 10

R : acceleration factor used to describe the acceleration of an MR acquisition due to undersampling of the data. In the context of parallel imaging, acquisitions are typically performed in $1/R$ of the time required for a fully-sampled (Nyquist) acquisition. 16

T_1 : longitudinal (spin-lattice) relaxation, the recovery of the Z-component of nuclear spin magnetization towards its thermal equilibrium value; after T_1 has elapsed, the longitudinal magnetization recovers 67% of its equilibrium value. xviii

T_2 : transverse (spin-spin) relaxation time, governing decay of the magnetization component transverse to B_0; after T_2 has elapsed, the transverse magnetization drops to 37% of its initial value. xviii

T_s : length of the readout acquisition window. 11

X : The X-axis is defined as parallel to the floor and extending rightward from the scanner isocenter if you are facing the front of the scanner. 23

Y : The Y-axis is defined as extending upward toward the ceiling from the scanner isocenter. 23

Z : The Z-axis is defined as originating at the scanner isocenter and then extending outward through the front of the scanner bore. 23

i_n : an index known as the *cyclic control* for iterative algorithms. Used to index through the rows of the encoding matrix, E, during iterative reconstruction performed using the Kaczmarz method. This index varies between 1 and $n_K n_c$ in any desired order. 43

n_K : number of acquired k-space data points. 11

n_c : number of receive coils used in an acquisition. 5, 11

s : readout signal acquired during MR imaging. Indexed according to the q^{th} coil element used for detection, the l^{th} readout performed, and the t^{th} time point in each readout. Each readout has N_s samples. In the entire experiment, the total number of acquired samples is $n_K n_c$, where n_K is the equal to N_s multiplied by the number of readouts acquired. 5

***g*-factor** : geometry factor, the spatially-varying map of noise amplification in SENSE reconstructions. The g-factor tracks the condition number of the unaliasing matrix for each voxel in the reduced FOV. 17, 48

$m(\mathbf{x})$: underlying spin density of the object that finds representation in an image; in this thesis ρ is also taken to incorporate the effects of T_1 recovery and T_2 decay. 1

ADC : analog-to-digital converter, an important part of the scanner's RF receiver system that is used to collect discrete samples from a time-varying voltage detected by receive coils. Typically each receive coil has a dedicated receiver channel and ADC. xvii, xviii, 10, 11

α : used in this thesis to denote the angle of rotation α in the time-frequency domain caused by the fractional Fourier transform (FrFT) operator of order α. 112

N_{PE_x} : number of phase encoding steps used in the X direction during phase mapping calibration of a pulse sequence. Equal to the resulting resolution of the phase map along X. 83

N_{PE_y} : number of phase encoding steps used in the Y direction during phase mapping calibration of a pulse sequence. Equal to the resulting resolution of the phase map along Y. 83

BW : Bandwidth. xvii

CP : center placement. The center of the quadratic SEM used for spatial encoding in O-Space imaging. The location of the SEM remains fixed during each readout, but changes between readouts. The location of the quadratic field offset is determined by the strength of X and Y SEMs played during readout. 35

DFT : discrete Fourier transform. 7, 13

EPI : echo planar imaging. 8

FFT : fast Fourier transform. xx, 7

flip angle : the rotation angle introduced to the net magnetization vector by a transverse RF pulse. xvii, 3

FOV : field of view. xviii, xxi, 28

Fresnel transform : Fresnel transform. xx, 112

FrFT : fractional Fourier transform. xviii, xx, 110, 112

FSE : fast spin echo. 8

Gibbs ringing : Oscillatory distortions introduced into a function when its Fourier spectrum is truncated, providing only finite spectral support for the function. Ringing is most pronounced near sharp edges in an image. In the context of MRI, Gibbs ringing is introduced by the (usually rectangular) window on the acquired readout data. Application of a tapered weighting function to the spectral data is used to suppress Gibbs ringing, at the expense of introducing a modest blur. 43

GradLoc : GRAdient LOCalization. This method uses a phase scrambling quadratic SEM pulse prior to readout to acquire signal only from a desired square window within the FOV, permitting localized, accelerated imaging without aliasing.. 110, 120

GRAPPA : generalized autocalibrating partially parallel acquisitions. 21

gyromagnetic ratio : The ratio, denoted γ, between the frequency of precession for a particular isotope and the B_0 field strength (42.57 MHz/Tesla for ^1H). 2

Kaczmarz method : an iterative row-action backprojection algorithm used in this thesis to obtain an image estimate when the forward model is described by a very large system of equations in the form of an encoding matrix, E. This method is particularly useful when the reconstruction matrix, F, can not be practically computed due to the large

size of E. When suitable value is chosen for the relaxation parameter, λ, the Kaczmarz method has been shown to converge to the minimum-norm least squares estimator. xviii, xx, 43

λ : relaxation parameter used to control the rate of convergence when performing iterative image reconstruction using the Kaczmarz method. xx, 43, 102

Larmor frequency : the frequency, ω_0, at which nuclear spins precess around the static background field, B_0. 2

LCT : linear canonical transform, a generalization of the Fourier transform. Special cases of this operator include the Fourier transform (FFT), fractional Fourier transform FrFT, and Fresnel transform. 110

LSQR : a conjugate gradients-style matrix solver intended for use with sparse matrices. 41, 57

MRI : magnetic resonance imaging. xix, xxi

NMR : nuclear magnetic resonance. xxi, 2

noise matrix : a covariance matrix whose diagonal elements specify the noise variance in each voxel of a reconstructed image using a given encoding matrix and reconstruction matrix, and whose off-diagonal terms specify noise correlations between voxels. xxi, 11

PatLoc : **PA**rallel imaging **T**echnique using **LOC**al gradients. A scheme that uses pairs of higher-order, multi-polar SEMs for spatial encoding, with one SEM played as read-out and the other as phase encoder. The inherent non-bijectivity of the SEMs is resolved using receive coil sensitivities, which permit unwrapping of aliased signals from points lying on the same frequency isocontour. In the simplest case, PatLoc image reconstruction is a generalization of SENSE. 26

PSF : point spread function, which describes the contribution of the source in a given voxel to the image magnitude and phase in all other reconstructed voxels in the FOV. Compare to SRF. xxi

Ψ : receiver noise covariance matrix, incorporating the noise variance in each receive channel as well as correlations across all available channels. Do not confuse with the image noise matrix. 11

q : used to index coil elements from an array of n_c channels. 6

radial : k-space trajectory consisting of spokes of data spaced every π/N radians, when N readouts are acquired. 8

receive coil : one or more radio-frequency probes used to detect signals from precessing spins in a sample. The precessing magnetization of the spins induces a time-varying voltage in each receive element. A receive coil often includes multiple elements arrayed so as to provide localized sensitivities around the ROI being imaged. 3

RF : radio frequency. xix, 3

ROI : the region of interest within an object or image. xxi, 22, 103, 120

SEM : spatial encoding magnetic field. A more general term for "gradient" field, which some consider appropriate only for describing fields that vary linearly over the FOV. xix, 4

SENSE : sensitivity encoding, a method for performing parallel imaging in the image domain with multiple receiver coils. xviii, xx, 16

SMASH : simultaneous acquisition of spatial harmonics. 20

spiral : spiral k-space trajectory. 9

SRF : spatial response function, which describes the contribution of all sources in the FOV to the image magnitude and phase at a given reconstructed voxel. Compare to PSF. xx, 12

static background field : the fixed magnetic field in an NMR experiment, typically ranging in strength between 1.5 Tesla and 7 Tesla for human MRI applications, with higher field strengths under investigation as of this writing. xx, 2

SVD singular value decomposition. 49

1

Spatial encoding and parallel MRI

1.1 The physics of spatial encoding in Magnetic Resonance Imaging

Note: Wherever possible, I have chosen to adopt the notation used in (1)

In magnetic resonance imaging, the signal originating from an object of interest, $m(\mathbf{x})$, depends on the density of water protons in the object, the longitudinal relaxation rate T_1, and the spin-spin dephasing rate T_2. These properties may vary spatially within an object or tissue, imparting useful information both about its bulk properties and the local chemical environment of the water. The resulting image contrast also depends on the repetition time TR and echo time TE of the pulse sequence used. Magnetic Resonance Imaging, or MRI, has a rich history of tailoring pulse sequences to emphasize one or more sources of contrast. The most exhaustive inventory to date is available in the *Handbook of Pulse Sequences* (2). This thesis will treat $m(\mathbf{x})$ as point of departure, focusing instead on ways to localize this signal during the process of image acquisition and reconstruction.

The quantum mechanical basis of nuclear magnetic resonance is described elsewhere in the context of imaging, so I will here give only the most cursory overview. The conventional MRI signal is obtained from the nuclear magnetic moments in the hydrogen protons

in the water molecules of the organism being imaged. For simple NMR experiments, spin behavior may be described by a classical model consisting of small magnetic moments subject to torques from applied magnetic fields. The spins precess at the Larmor frequency, $\omega_0 = \gamma B_0$ around the constant, uniform static background field, B_0, oriented in the Z-direction. The gyromagnetic ratio, γ, is the ratio between the frequency of precession for a particular isotope and the B_0 field strength (267.47×10^6 radians/Tesla for ^1H). The B_0 field is sustained by liquid helium-cooled superconducting current loops in the core of the magnet system.

Protons are spin $1/2$ particles, thus representing a spin-1/2 system. In the absence of a magnetic field the spins are statistically distributed between the two states, with equal probability of occupying each state. NMR and MRI exploit isotopes with odd number of nuclei, since this leaves one spin unpaired, thus imparting the nucleus with a net magnetic moment. In the presence of a strong magnetic field, a slight excess of spins align themselves along the field, with the surplus fraction assuming a Boltzmann distribution, $e^{-\frac{h\nu}{kT}}$, where k is the Boltzmann constant, T is temperature, and $h\nu$ is the quantized energy gap between the two spin states. In this way, a net magnetization is said to exist along the direction of B_0 in thermal equilibrium. Excitation by radio frequency pulses at frequency ω_0 results in a certain fraction of the spins becoming excited into a superposition of spin up and spin down states. This manifests itself as a precessing transverse magnetic moment, which induces a voltage in nearby pick-up loop, providing the NMR signal. Without some mechanism of spatial localization, the acquired signal can supply only spectral localization according to the chemical shift of each molecule in the sample for a nucleus resonating at the chosen frequency (3).

Signals arising from unpaired spins can in principle be detected from any odd-numbered isotope, and ^{13}C, ^{31}P, ^{17}O, ^{23}Na are commonly used in nuclear magnetic resonance spectroscopy, and increasingly in spectroscopic imaging. However, this thesis will only conven-

tional MRI, which relies on ^1H, which provides by far the strongest signal of any isotope *in vivo*, given the abundance of water in the human body.

The Bloch equations (2) govern the evolution of the magnetization at each location in the object as a function of the applied magnetic fields. Spin evolution is usually described in a coordinate system that is rotating at ω_0 relative to the lab frame. Positive or negative applied magnetic fields then cause spins to lead or lag the Larmor frequency, resulting in simple rotation about the Z-axis in the rotating frame. Spins can also be manipulated with radio frequency (RF) pulses whose frequency is at or near the Larmor frequency. RF pulses result in rotation of the net magnetization around the axis of the transmitted field, known as B_1. To tip spins into the transverse plane, the applied RF pulse must have a B-field component lying in the transverse plane.

Transmit and receive radio-frequency probes are known as RF coils. Typically the transmit coil is designed to generate a circularly polarized B-field, which appears fixed in the rotating frame. Linearly-polarized coils can also be used, since a linear polarization can be decomposed into a left-handed and right-handed circularly polarized components, but because the right-handed polarization does not contribute to the flip angle, linear coils are less efficient. The same coil is often used both to transmit and receive, in which the opposite circular polarization is used. By the principle of reciprocity, the coil's transmit field and receive sensitivity are complex conjugates of one another, to a first-order approximation. For a circularly-polarized coil, the receive sensitivity will be described by the right-handed circularly polarized B_1 mode. The precessing magnetization of each spin in the sample induces a voltage in the receive coil(s). The magnitude and phase of the voltage reflects the density of spins, the state of spin relaxation, the phase of the spin precession, the proximity of the spins to the given receive coil, and the sensitivity of the coil.

To improve signal-to-noise ratio and provide more independent measurements of the sample, multiple small surface coils are used for reception, forming a phased array phased

array (4). Many combinations of transmit and receiver coils exist, including the transverse electromagnetic mode (TEM) coil (5), the Bollinger coil (6), the microstrip array (7). Sometimes, multiple coil elements are used during transmission to provide more degrees of freedom in shaping the transmit B_1 profile, particularly at ultra high field strengths (such as 7 Tesla), at which the RF wavelength in the tissue approaches the size of the human head. The resulting dielectric resonances cause inhomogeneous B_1 fields and spatially-varying intensity in the resulting image. Adjusting the amplitude and phase of the elements to provide a desired B_1 distribution (usually a uniform mode) is known as RF shimming (8). As will be discussed later in this chapter, multiple RF elements can also be used during signal reception, providing improved sensitivity as well as spatial information about the location of each signal source in the FOV.

Spatial encoding in MRI is achieved primarily through the use of spatially-varying magnetic fields that are pulsed on and off during an acquisition. The great variety of spatial encoding approaches used in MRI, from the original projection-type sequence (9) onward, stems from the tremendous flexibility for the user to apply multiple fields in different orderings, with freedom to vary the pulse shape of each field (i.e., trapezoidal, ramped, sinusoidal, etc.). Pulsed, spatially-varying magnetic fields are applied using separate coils nested within the scanner bore. Because this thesis treats spatial encoding by both linear and nonlinear fields, I will follow (10) in using the term *spatial encoding magnetic field*, or SEM, instead of the conventional phrase *gradient field*, which some investigators feel should be reserved for fields with linear spatial variation[1].

In a pulse sequence, one of the SEMs is typically played during the RF pulse for slice selection. Based on the center frequency and bandwidth of the RF pulse, spins residing

[1]Merriam-Webster's dictionary defines *gradient* as: *(1) (a) the rate of regular or graded ascent or descent : inclination (b) a part sloping upward or downward (2) change in the value of a quantity (as temperature, pressure, or concentration) with change in a given variable and especially per unit distance in a specified direction.* Given these definitions, there is a case to be made for using *gradient field* only in reference to fields having linear spatial variation.

within the range of frequencies in the BW will be excited. The strength of the SEM (or the BW of the RF pulse) are used to control slice thickness. For spatial encoding in-plane, the other two linear SEMs are used. At least one of the SEMs is typically played before the readout the dephase spins. The readout pulse is then played with opposite polarity to refocus the spins, forming an echo at *TE* when all the spins are contributing signals with the same phase. Another SEM is then sometimes used as for phase encoding along a different cardinal direction. The phase encoding pulses vary in amplitude between SEMs, providing the same spatial information as a readout, except spaced over multiple *TR*s.

In the most general case, the temporal signal in an MRI acquisition can be expressed as

$$s_l(t) = \int m(\mathbf{x}) e^{i\phi_l(x,t)} d\mathbf{x} \tag{1.1.1}$$

for the l^{th} phase function, which corresponds to the applied SEM field shape. The l^{th} acquired readout signal is represented by s. The phase evolves over time according to the amplitude of the applied SEM, which may be constant, ramped, sinusoidal, or an arbitrary waveform. The vector $\mathbf{x}$ represents a point in 3-D space. With the addition of surface coil encoding, described by $C_q(\mathbf{x})$, this expression becomes

$$s_{l,q}(t) = \int C_q(\mathbf{x}) m(\mathbf{x}) e^{i\phi_l(t)} d\mathbf{x} \tag{1.1.2}$$

where q denotes the coil index out of a set of n_c receive coils. The accumulated phase $\phi_l(t)$ depends primarily on the history of the applied SEM pulse, $G(t)$,

$$\phi_l(\mathbf{x}, t) = \gamma \int_0^t \mathbf{x} \cdot \mathbf{G}(t') dt' \tag{1.1.3}$$

where γ is the gyromagnetic ratio. A useful and entirely apt analogy is to imagine the spins being "driven" through the frequency domain, or "k-space", by the encoding mag-

netic fields. The k-space location is analogous to spatial displacement, while the slope of the encoding field is analogous to velocity.

The two classic methods for spatial localization, radial projection imaging (9) and phase encoding (11), both employ linear SEMs oriented along X, Y, and Z, creating a Fourier basis for spatial encoding. In radial trajectories, each acquired readout corresponds to a "spoke" in k-space, while for Cartesian trajectories, each readout is a line. Radial datasets may be reconstructed using the radon transform (12), also known as filtered backprojection, while Cartesian data can usually be reconstructed using the 2-D Discrete Fourier transform (13).

The purpose of this thesis is to investigate what happens when the Fourier basis for spatial encoding is replaced with fields generated using arbitrary combination of spherical harmonics. Specifically, I would like to know what stands to be gained by:

- Complementing RF coil profiles
- Tailor spatially-varying resolution to region of interest within the FOV

Phase encoding was the first nonlinear encoding scheme to be explored, using pairs second-order multipolar magnetic fields (10). The present author's work was the first to explore projection imaging with nonlinear fields (14). But before delving into the promise and pitfalls of using nonlinear encoding fields, a review of parallel imaging and its potential benefits is in order. Ignoring spin relaxation, the l^{th} readout acquired with the q^{th} receive coil in a Cartesian phase encoded sequence takes the form of a Fourier Transform,

$$s_{q,l}(t) = \int C_q(\mathbf{x}) m(\mathbf{x}) e^{i2\pi f(\mathbf{x})t} df \tag{1.1.4}$$

Considering the case of 2-D phase encoding in the X-Y (axial) plane explicitly,

$$s_q(k_x, \Delta k_y) = \int C_q(x,y) m(x,y) e^{i(k_x x + \Delta k_y y)} dx \tag{1.1.5}$$

This expression shows that in the presence of linear SEMs scanned across a Cartesian trajectory, the MRI signal reduces to a 2-D Fourier transform weighted by the receive coil profiles. When a single, uniform receive coil is used, the signal reduces to a pure Fourier transform, permitting image reconstruction via inverse Fourier transformation.

Since discrete time samples are measured during readout, the acquired data actually corresponds to the discrete Fourier transform of the object (DFT)[1]. In this case, the shape of the acquisition window causes the object to be convolved with a blurring kernel, most commonly a 2-D sinc function corresponding to the Fourier transform of the 2-D rectangular function that bounds the acquisition window. Truncation of k-space is responsible for the Gibbs ringing that commonly occurs in MR images, corresponding to missing high frequency components necessary to exactly reproduce the edge. Ringing is most often suppressed by windowing k-space, trading ringing artifacts for a small amount of blurring.

More generally, the signal is a function of the multi-dimensional k-space vectors $\mathbf{k}$ and $\mathbf{G}$, where $\mathbf{G}$ describes the amplitude of the components of the local SEM in units of Hz/cm, (15)

$$s_q(\mathbf{k}(t)) = \int m(\mathbf{x}) C_q(\mathbf{x}) e^{i\mathbf{k(t)}^T \mathbf{G(x)}} d\mathbf{x} \tag{1.1.6}$$

1.2 Speed is of the essence

As compared with modalities such as CT or ultrasound, MRI acquisition times are comparatively slow, with clinical protocols typically running between 30 and 60 minutes. While

[1] When the number of data points is a power of 2, the DFT is usually evaluated using the computationally-efficient fast Fourier transform (FFT) (13)

this raises concerns over patient comfort and economic cost, the bottleneck becomes an even more acute concern for certain imaging applications. Consider just a partial list:

- Imaging of the beating heart (16), its perfusion (17), strain tracking (18) in the myocardium;
- Abdominal imaging (19) requiring repeated breath-holds of 20 seconds or longer, stressing subjects, particularly the elderly;
- Vascular bolus tracking (20);
- Whole-brain rapid functional MRI of the blood oxygenation level dependent (BOLD) response (21);
- Imaging of children without use of anesthesia (22).

An additional concern when long signal readouts are used is B_0 inhomogeneity, which causes voxel displacements along the phase encode direction (23).

1.2.1 Fundamental limits on acquisition speed

The cause of the bottleneck in MRI is the need to acquire k-space points serially in successive TRs. A whole menagerie of hat tricks have been used to speed up this process. They may roughly be divided into two categories, each with distinct benefits and drawbacks:

- Acquiring multiple lines of Cartesian k-space in a single TR;
  - Echo planar imaging (EPI) (24).
  - Fast spin echo (FSE) (25) and Turbo spin echo.
  - Interleaved slice selection.
- Quickly acquiring much or all of k-space in one TR using a non-Cartesian trajectory;
  - radial projection trajectories that frequently sample the center of k-space (9), improving robustness to subject motion .

- spiral trajectories that quickly cover k-space with lower slew rates than Cartesian trapezoidal pulses.
- Rosette (26) and other trajectory shapes.

Non-Cartesian trajectories also enjoy other benefits in addition to their speed. For instance, radial trajectories permit imaging with very short TE times and reduced sensitivity to motion. As compared with Cartesian trajectories, spiral trajectories can in principle cover k-space in equivalent time with reduced SEM field slew rates.

This list, while not intended to be exhaustive, highlights the two practical limits to conventional MRI. The methods that collect multiple lines per TR have been intensively developed and are reaching their performance limits due to bounds on (a) SEM switching limits and (b.) specific absorption ratio, or SAR. Each of these is regulated by the Food and Drug Administration (27). SEM field switching must not induce any peripheral nerve stimulation or otherwise cause the subject discomfort. This imposes practical limits on the slew rate of imaging pulses (28). Additionally, all SEM coils have finite inductances, limiting the rate at which fields can be switched. SAR imposes limits on the performance of sequences such as fast spin echo, which require multiple 180◦ RF pulses in rapid succession. Furthermore, SAR becomes much more severe at ultra-high field strengths at or above 7T that are now being investigated for clinical application.

Non-Cartesian trajectories have been the subject of extensive research over the past 20 years, but have found only very limited adoption on clinical scanners. The primary reason for this is the difficulty of accurately measuring the k-space trajectory, which depends on the exact pulse timing, eddy currents, coupling to other SEM coils, and subtle effects such as concomitant fields (29, 30). Image reconstruction of multi-coil datasets is also more time consuming than Cartesian reconstruction via the FFT since it requires use of (a) regridding (31) with k-space density compensation (32), (b) the non-uniform FFT (33), or (c) iterative algorithms such as Conjugate Gradients (34, 35), or some combination of

these approaches. Non-Cartesian trajectories also typically require more data points to be acquired in order to satisfy the Nyquist criterion. For instance, in radial acquisitions, in order to inscribe the usual Cartesian rectangle within the circle spanned by the spokes and still satisfy the Nyquist criterion, an additional $\pi/2$ spokes must be acquired, each with $\sqrt{2}N_s$ readout samples, where N_s samples would be required for the Cartesian trajectory.

One solution to the serial readout bottleneck is to acquire data simultaneously using multiple sensors - in this case, RF receive coils. Signals received in parallel by an array of independent RF receivers can then be combined to compensate for missing spatial encoding in undersample k-space datasets. This approach, known as parallel imaging, will be discussed in detail later in this chapter.

1.3 Image reconstruction fundamentals

Before discussing parallel imaging methods, a brief review of MR image reconstruction is in order. I do not offer an exhaustive survey, but rather a selective presentation of the material most relevant to "conventional" parallel imaging techniques. I have modeled my notation after that used in (1).

The signal equation, Eq. (1.1.6), can be represented as the integral of the product of the underlying magnetization function with an encoding function,

$$s_{q,\kappa} = \int m(\mathbf{x}) C_q(\mathbf{x}) e^{i\mathbf{k}^T \mathbf{G}(\mathbf{x})} d\mathbf{x} \tag{1.3.1}$$

Because the signal integral is a linear operator, it may be recast as a matrix operating on a vectorized, discretized version of the source object. Let the forward model be denoted $s = Em+n$, where n is zero-mean additive Gaussian noise that may be correlated across the receive channels. The scanner's analog-to-digital converter, or ADC, acquires N_s discrete

samples during a rectangular acquisition window of duration T_s.

1.3.1 Weighted least squares estimator of image

Consider a general encoding matrix, E, evaluated on a discrete grid of voxels in the FOV. E has size $n_K n_c \times N^2$, where n_K is total number of acquired k-space points, n_c is the number of coils, and N is the number of voxels being represented along each dimension of the source object (and also the reconstructed image). Square image grids are assumed for simplicity. A receiver noise covariance matrix, Ψ, may be obtained by placing the receive coils in proximity to the sample, opening up the ADC to take samples of the noise in each receive channel, and then computing all the cross-correlations between these noise vectors: $\Psi = E\left(nn^H\right)$, for a matrix n, each row of which holds a series noise measurements from a given coil channel. To use Ψ in combination with the full encoding matrix, E, Ψ must be tiled along the diagonal using the outer product, forming a block-diagonal matrix: $\tilde{\Psi} = \Psi \otimes I_{n_K}$ where I_{n_K} is the identity matrix of size n_K. Define a reconstruction matrix that is optimal in a least-squares sense, F, which operates on the acquired MRI data to generate an image. Using F, E, and $\tilde{\Psi}$, the noise properties of the experiment can be derived by calculating the covariance matrix.

The image covariance matrix, known in the MR literature simply as the noise matrix, describes both the expected variance σ_ρ within each voxel ρ (diagonal elements) and the noise correlations between voxels (off-diagonal elements). The covariance matrix can be derived by taking the expected value of the reconstructed image estimator, $\hat{m}_{LS}$,(36)

$$\begin{aligned} \hat{m}_{LS} &= F\left(Em+n\right) \\ cov(\hat{m}_{LS}) &= \left[\left(\hat{m}_{LS}-m\right)\left(\hat{m}_{LS}-m\right)^{H}\right] \\ &= \left[\left(m+Fn-m\right)\left(m+Fn-m\right)^{H}\right] \\ &= F\tilde{\Psi}F^{H}. \end{aligned} \tag{1.3.2}$$

One way to compute a reconstruction matrix F is obtained using a Lagrange multiplier. Following the treatment in (1), one seeks to minimize the total image noise variance, defined as the sum of the diagonal elements in the noise matrix, $Tr\left(F\tilde{\Psi}F^{H}\right)$. This minimization is also subject to the so-called "weak" constraint of MR image reconstruction; namely, that the spatial response (SRF) function of each voxel must equal unity at the center of the voxel and zero at the center of all other voxels. Since the SRF of a given voxel with index ρ is specified in the ρ^{th} column of the matrix product FE, the weak voxel function constraint is enforced by minimizing $FE-I$. The problem of finding F thus becomes one of constrained optimization, and its form is treatable using a Lagrange multiplier. The two conditions are written as a cost function, L, with one term weighted by the Lagrange multiplier, Λ,

$$L = Tr\left(F\tilde{\Psi}F^{H}\right) + \Lambda\left(FE-I\right)$$

The derivative is taken with respect to both Λ and the variable for which an optimum expression is sought, F, and the system of equations is solved for Λ and then F:

$$
\begin{aligned}
\frac{\partial L}{\partial \Lambda} &= (FE)^H - I = E^H F^H - I = 0 \\
\frac{\partial L}{\partial F} &= \frac{\partial}{\partial F} Tr\left(F\tilde{\Psi}F^H\right) + \frac{\partial}{\partial F}\Lambda\left(FE - I\right) = 0 \\
&= 2\tilde{\Psi}F^H + E\Lambda^H = 0 \\
E\Lambda^H &= -2\tilde{\Psi}F^H \\
E^H\tilde{\Psi}^{-1}E\Lambda^H &= -2E^HF^H = -2 \\
\Lambda^H &= -2\left(E^H\tilde{\Psi}^{-1}E\right)^{-1} \\
2\tilde{\Psi}F^H &= 2E\left(E^H\tilde{\Psi}^{-1}E\right)^{-1} \\
F &= \left(\tilde{\Psi}^{-1}E\left(E^H\tilde{\Psi}^{-1}E\right)^{-1}\right)^H \\
&= \left(E^H\tilde{\Psi}^{-1}E\right)^{-1}E^H\tilde{\Psi}^{-1}.
\end{aligned}
\tag{1.3.3}
$$

For Cartesian encoding with no surface coil encoding and completely uncorrelated noise, F is simply the discrete Fourier transform (DFT). In this case, the unitarity of the DFT implies that $F = E^H$.

1.4 Phase array image reconstruction

It has been well documented that arrays of multiple, localized receive coils can improve SNR, with the gains greatest at the periphery of the object (4, 37, 38, 39, 40). Commercial MRI scanners have moved to using ever-greater numbers of receive coils placed close to the tissue being imaged in order to exploit the high sensitivity of the coils to signals originating in their vicinity. However, when multiple coils are used, it is no longer obvious that the FFT provides the optimal image quality. Single coil Fourier reconstructions have the advantage of zero noise correlation between voxels, i.e., all off-diagonal elements in the noise matrix are zero. When multiple coils are used, their sensitivities overlap such that

thermal fluctuations in the object at a given location induce noise signals in two or more coils. The noise covariance matrix between coil channels can be readily computed from noise samples. As was shown in the previous section, this matrix can then be incorporated into image reconstruction to weight noisy coils less and to account for correlations between signals.

Let matrix C hold the n_c complex-valued coil sensitivities represented on an $N \times N$ grid, where N is the desired number of voxels (image resolution) in each dimension. Further let $\hat{s}$ hold the FFTs of the Cartesian-sampled k-space data acquired by all k-space points, i.e., the intermediate images provided by each receive coil. Both of these matrices are of size and vectorized into a matrix of size $N^2 \times n_c$. The optimal reconstruction and its corresponding SNR have been shown to be given by (4)

$$\begin{aligned} \hat{m}_{opt} &= \sqrt{\hat{s}^H \Psi C} \\ SNR_{opt} &= diag\left(\sqrt{C^H \Psi^{-1} C}\right) \end{aligned} \tag{1.4.1}$$

where $diag$ extracts the diagonal elements. For coil arrays with modest noise correlations, Ψ approaches a diagonal matrix and the SNR-optimal reconstruction is approximated by the sum-of-squares reconstruction $\hat{m}_{opt} = \sqrt{\hat{s}^H s}$, which can be evaluated with no knowledge of the coil profiles or their noise covariance[1].

If instead of optimal SNR, one instead desires an image with uniform intensity, the coil profiles can be combined into a single coil with uniform sensitivity using this expression: (4)

$$\hat{m}_{sensitivity} = \frac{\hat{s}^H \Psi^{-1} C}{C^H \Psi^{-1} C} \tag{1.4.2}$$

[1]The eigenvalues and eigenvectors of the noise covariance matrix provide linear combinations of coil sensitivities that form orthogonal "virtual" coil modes that have uncorrelated noise. For symmetric, encircling arrays of coil profiles, these modes resemble the orthogonal modes of a birdcage coil, of which only the uniform mode is typically used in practice.

Finally, to obtain an image with a uniform noise level, the above expression 1.4.2 is modified slightly to keep the noise variance the same in every voxel,

$$\hat{m}_{noise} = \frac{\hat{s}^H \Psi^{-1} C}{\sqrt{C^H \Psi^{-1} C}} \tag{1.4.3}$$

1.4.1 Estimating coil sensitivities

While the sum-of-squares approach described above does not require knowledge of the coil sensitivities, many reconstructions require accurate coil maps, particularly when reconstructing from undersampled data. The difficulty with obtaining coil maps is that coil-wise images include weightings from both the transmit and receive coils, which must be distinguished from one another. If the same coil elements are used to transmit and receive, a simple method to obtain the receive profiles is to take the square root of the coil-wise images. Another method is to divide each coil image by the sum-of-squares image formed using all available coil data. The approach used in this thesis, due to Walsh (41), provides estimates of the coil sensitivities as well as a combined image that exploits the signal and noise statistics of the data, without the need for pre-acquiring the coil profiles or the noise covariance matrix. Instead, signal (Ψ_s) and noise (Ψ_n) correlation matrices are created using local regions within each coil image's FOV that are dominated by signal or noise, respectively. The SNR-optimal combined image is then provided using the weights provided by the eigenvector of the matrix $P = \Psi_n^{-1}\Psi_s$ that has the largest eigenvalue.

1.5 Parallel imaging methods

Up to this point, discussion has centered on coil datasets that are Nyquist-sampled in k-space. The question, then, is how to combine multiple, undersampled coil datasets in a way that exploits coil sensitivity spatial encoding to "fill in" the missing k-space data. A

host of "parallel imaging" methods have been introduced to address this problem. I will review only a small subset of the most popular parallel methods below. They may broadly be divided into those that operate in the image domain (SENSE (1)) and those that directly fill in missing k-space data (SMASH (42), GRAPPA (43), and variations on these). Hybrid image domain/k-space methods also exist, including SPACE-RIP (44), GSMASH (45), and others. Methods also exist to address temporal undersampling in time series data (46) as well as slice-direction undersampling (47), but for the purposes of this thesis I will treat only time-independent datasets.

1.5.1 SENSE

A simple way to accelerate Cartesian MR scans is to skip the acquisition of R phase encoded k-space lines for each acquired line, leading to scan acceleration by a factor of R. Regular undersampling increases the spacing between lines from Δk to $R\Delta k$, in turn reducing the FOV in the imaging domain from FOV to FOV/R. Object features falling outside this reduced FOV "alias" or wrap around into the other end of the FOV, causing voxel overlap within the object. The most popular image-domain method for unwrapping this aliasing is known as "Sensitivity Encoding", or SENSE (1). This approach uses independent coil sensitivities to disambiguate signals in the reduced FOV, unwrapping them back into their source voxels in the full FOV. The SENSE unwrapping problem can be expressed using small aliasing matrices on a voxel-wise basis in the reduced FOV, leading to efficient computation. With coil sensitivities contained in a matrix of size $n_c \times R$, the aliased signal intensities for all coils in the reduced FOV, a, are expressed as

$$a = Sv \tag{1.5.1}$$

where v is a $R \times 1$ vector of signal intensities in the source voxels of the full FOV and a is $n_c \times 1$ in the reduced FOV. An unwrapping matrix $S^{\dagger}$ is then used to provide an estimate of the source voxel intensities in the full FOV,

$$\hat{v} = S^{\dagger} a \tag{1.5.2}$$

The least squares estimate of the unwrapping matrix has been shown (1) to be the pseudoinverse weighted by the coil noise covariance matrix:

$$S^{\dagger} = \left(S^{H}\Psi^{-1}S\right)^{-1} S^{H}\Psi^{-1} \tag{1.5.3}$$

where S is the $n_c \times R$ matrix of coil profiles at each voxel in the full FOV mapping onto a particular voxel in the reduced FOV. This procedure is repeated at each voxel in the reduced FOV, producing a composite unaliased image in the full FOV.

SENSE parallel imaging typically removes aliasing artifacts quite well, but brings two SNR penalties. The condition number of matrix $S^{\dagger}$ depends on the spatial encoding provided by the magnitude and phase variation of the coil sensitivities at each point in the FOV. It is possible to compute a spatial map of coil encoding performance by explicitly calculating the ratio between noise levels in fully-sampled and undersampled SENSE reconstructions. The noise matrix at each voxel has been shown to be $X = \frac{1}{n_K}\left(S^{H}\Psi^{-1}S\right)^{-1}$ where n_K is the number of acquired k-space points. The ratio of noise levels in the ρ^{th} voxel is therefore

$$\frac{X^{red}_{\rho,\rho}}{X^{full}_{\rho,\rho}} = \sqrt{R}\sqrt{\left[\left(S^{H}\Psi^{-1}S\right)^{-1}\right]_{\rho,\rho}\left(S^{H}\Psi^{-1}S\right)_{\rho,\rho}} \tag{1.5.4}$$

This expression shows the two sources of SNR loss in SENSE reconstructions. The first penalty is a reduction of $\sqrt{R}$ in SNR simply because only $1/R$ of the usual k-space points are being acquired. The second penalty is noise amplification due to the ill-conditioned matrix pseudoinverse required to compute $S^{\dagger}$. SENSE noise amplification is described in the form of a geometry factor map, or *g*-factor map, over the full FOV:

$$g_{\rho,\rho} = \sqrt{\left[\left(S^{H}\Psi^{-1}S\right)^{-1}\right]_{\rho,\rho}\left(S^{H}\Psi^{-1}S\right)_{\rho,\rho}} \tag{1.5.5}$$

It is straightforward to generalize SENSE to the case of undersampling along multiple phase-encode directions, permitting highly-accelerated 3-D acquisitions provided there is adequate spatial encoding from the receive coils. Typically, as the number of coil elements grows, or as the coils become more spatially distinct, the g-factor decreases. Simulations of SENSE reconstructions and g-factors at various acceleration factors using 8 receive coils[1] are shown in 1.1. As expected, the g-factor grows dramatically with the acceleration factor, overwhelming the image with noise as R approaches the number of coils. The g-factor is also highest in regions of the reduced FOV with the most overlapping voxels.

The most common way to improve SENSE parallel imaging performance is to move to ever-greater numbers of receive coils (48). However, this approach requires the use of additional coil elements and associated pre-amplifiers, adding considerable cost and complexity. An alternative way to control noise amplification is to use Tikhonov regularization (49). In this approach, the SVD of the unfolding matrix is obtained and singular values below a certain threshold are truncated, limiting the propagation of noise components that are co-linear with the corresponding high-frequency singular vectors. This causes a trade-off between noise levels and artifact levels in the resulting image.

The poor performance of SENSE at high acceleration factors has motivated the search for alternative ways to accelerate MRI scans, including the nonlinear SEM methods described in the next two chapters.

1.5.2 SMASH

An alternative to unwrapping aliasing in the image domain is to directly fill in the skipped lines of k-space. This was first achieved using a method called "SiMultaneous Acquisition

[1] The coil profiles belong to an 8-channel commercial knee array (*In vivo* Corp., Ganesville, Florida) with elements distributed circumferentially around the sample.

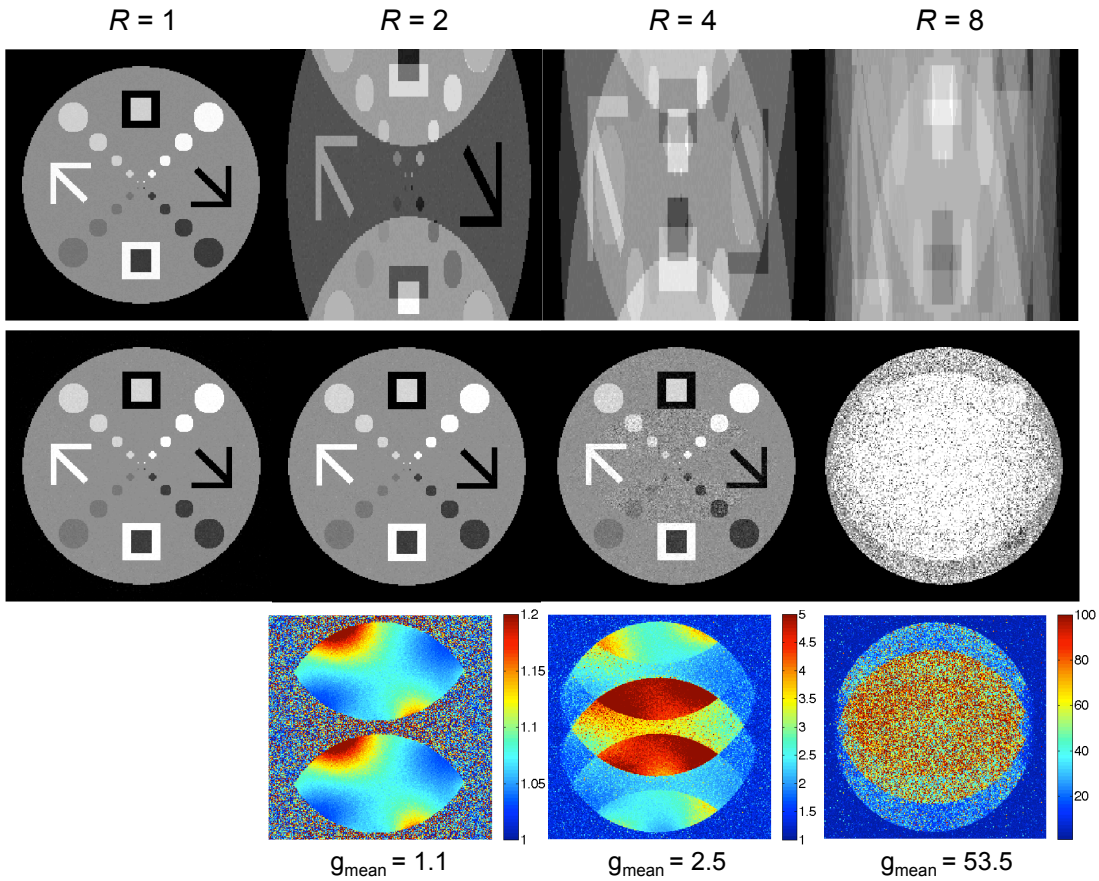

Figure 1.1: Simulations of SENSE reconstructions based on 8 RF receive coil profiles and a numerical phantom with lesion-like features. The top row shows the aliasing that occurs when undersampled k-space data are directly reconstructed via the FFT. The middle row shows SENSE multi-coil reconstructions at each acceleration factors. SENSE removes the aliasing but introduces spatially-varying noise amplification, as described by the g-factor maps (bottom row). In regions with high degrees of aliasing (multiple overlapping voxels in the reduced FOV), the g-factor grows quickly with acceleration factor, degrading image quality.

of Spatial Harmonics", SMASH (42), that uses the frequency shift property of the Fourier transform. In this approach, spatial harmonics along the undersampled dimension are approximated in the image domain using linear combinations of the available coil profiles to form a composite coil,

$$C^{comp}(x,y) = \sum_j w_j C_j(x,y) = exp(iM\Delta k_y y) \tag{1.5.6}$$

where $\Delta k_y = 2\pi/FOV$ and M is an integer corresponding to the undersampling factor in k-space. Missing k-space lines spaced M phase encode steps away from the preceding acquired line are then filled in by combining all n_c coil k-space datasets using the weights w_j used to build the associated complex spatial harmonic.

$$\begin{aligned} S(k_x,k_y) &= \int\int C^{comp}(x,y)\, m(x,y)\, exp(-ik_x x - ik_y y)\, dxdy \\ &= \int\int m(x,y)\, exp(-ik_x x - i(k_y - m\Delta k_y)\, y) \\ &= \hat{m}(k_x, k_y - m\Delta k_y) \end{aligned} \tag{1.5.7}$$

To form composite k-space data at the R=1 lines that have already been acquired, weights can be calculated such that the composite coil profile approximates a flat sensitivity. For the R=2 lines adjacent to each acquired line, the coil profiles must approximate $\exp(i2\pi y/FOV)$, corresponding to one cycle of a cosine and quadrature sine function over the FOV. For R=3 lines, the coil profiles must be able to approximate sinusoids that complete two cycles within the FOV. As the acceleration factor increases, more and more localized coil sensitivities are required to form the high frequency variation of the harmonics.

1.5.3 GRAPPA

In its original formulation, SMASH requires explicit knowledge of the coil sensitivities in order to approximate spatial harmonics. A method called AUTO-SMASH circumvents this by including a few autocalibration (ACS) lines in the acquisition that would otherwise be skipped (45, 50, 51). The weights needed to fill in missing lines elsewhere in k-space can then be determined by calculating the weights necessary to create an ACS lines using the nearest acquired lines from all coil datasets. In this way, only R-1 additional lines need to be acquired with the undersampled k-space data instead of a whole additional reference scan of the coil profiles.

A refinement of this approach known as "Generalized autocalibrating partially parallel acquisitions", or GRAPPA (43), has found widespread use in the MR community. In GRAPPA, the SMASH algorithm is generalized to include the use of convolution kernels in k-space that use multiple acquired lines to fill in each missing line. Also, weights are found for fitting each coil's ACS lines using acquired lines from all n_c coil datasets, permitting reconstruction of coil-weighted images, which can then be combined by sum-of-squares or another approach. GRAPPA is computationally efficient. While a clean analytical expression for the g-factor does not exist for GRAPPA, as it does for SENSE, significant progress has recently been made toward estimating GRAPPA voxel-wise noise amplification (52).

1.5.4 Non-Cartesian parallel imaging

Substantial effort has been invested in recent years to achieve reliable parallel reconstruction of undersample non-Cartesian trajectories. While a complete review of this work is beyond the scope of this thesis, it is worth noting a few algorithms that are relevant in the present context. SENSE has been extended (34) to treat non-Cartesian encoding matrices by using the FFT and k-space re-gridding to achieve a fast implementation of the conjugate gradients algorithm (53) that does not require explicit calculation of the full encoding matrix. The

GRAPPA kernel has also been generalized to fill in k-space points that do lie on a Cartesian grid (54). Undersampled radial are a particularly relevant comparison for the Kaczmarz reconstructions of O-Space (and radial) datasets to be shown in future chapters of this thesis. In one method (35), undersampled radial streaking artifacts are suppressed by combining fast conjugate gradients algorithm in (34) with an L1-norm penalty to minimize noise and artifacts while preserving object edges.

1.5.5 Parallel transmission

It should be noted that the same principles used in parallel reception can also be used to reduce the length of RF pulses used to achieve a given in-slice excitation profile (55, 56, 57). When accompanied by time-varying SEMs, RF pulses traverse the excitation k-space of the desired object profile. When the RF pulse duration is reduced - for instance, to reduce heat deposition in the subject - the resulting excitation profile will experience aliasing. Using appropriately-tailored RF pulse shapes and SEM waveforms, a parallel transmit array of coil elements can supply this missing spatial information and permit reduced RF pulse durations without creating aliased excitation profiles.

Parallel transmission is useful not only for exciting spins in a desired region of interest (ROI), but also for creating a more even excitation at ultra-high field strengths. When the RF wavelength is comparable in size to the anatomy being imaged, dielectric resonances can perturb the B_1 field and produce an uneven flip angle.

1.6 Ultimate SNR

Regardless of the reconstruction approach used, the performance of parallel imaging is ultimately bounded by the spatial information available in the receive coils. For undersampled Cartesian datasets, the ultimate SNR has been derived as a function of the B_0 field

strength and the acceleration factor, assuming that an arbitrarily large number of physically-realizable coil eigenmodes are available for performing signal detection (58). In this analysis, Wiesinger chooses vector spherical harmonics are taken to form the coil eigenbasis, since they satisfy Maxwell's equations and provide mutually independent information about the object, forming a basis set. But even supposing that an infinite number of such coils were available, an upper bound exists on parallel imaging performance (and a lower bound on the achievable g-factor).

1.7 Coil eigenmodes and SNR

In addition to the bounds on parallel imaging performance, there are diminishing returns to image SNR from using ever-greater numbers of local receive coils. Eigenmode analysis shows that as surface coils in an array become smaller and more numerous, gains in SNR are relegated to ever-thinner bands near the periphery of the FOV, close to the coils (39). Higher-order eigenmodes do not contribute signal at the center of the FOV, where all of the signal comes from the uniform "birdcage" mode. In theory, then, large arrays can never perform better at the center of the FOV than a single-channel volume resonator like the birdcage coil.

This calculated result agrees with experimental measurements of SNR in large arrays, which show experimental comparisons of transverse ($X-Y$ plane[1]) images acquired using 12, 32, and 96-channel loop coil head arrays (48). While moving from 12 to 32 channels produces a roughly 4-fold improvement in SNR at the edge of the brain, moving from 32 to 96 elements provides less than a 2-fold improvement. Furthermore, while the 32 elements outperform the 12 elements by a substantial margin at points halfway between the center and edge of the brain, the 96-channel array only offers significant improvements over the 32-channel version at the very outermost 1/6 of the brain's diameter.

[1] This thesis adopts the MRI community's most common definition of the Z-axis as emerging from the front of the magnet bore. The X-axis is defined to be parallel to the floor, pointing to the right as you face the front of the scanner bore. The Y-axis points to the ceiling.

2

Imaging with nonlinear encoding fields

The great strength of parallel reception is that multiple signals are acquired at the same time, potentially permitting dramatic reductions in imaging time. However, as ever-greater numbers of receive coils are used, several limitations emerge:

- Coil array cost and complexity grows;
- Inter-element coupling becomes harder to compensate or suppress;
- *G*-factor reductions with increasing array size are ultimately constrained to a lower bound imposed by the B_0 field strength and by Maxwell's equations;
- SNR gains are increasingly confined to the periphery as the number of elements grows.

These bounds on performance motivate a reappraisal of parallel imaging from its fundamentals on up. What if the question of how to optimize performance is turned on its head? Namely, instead of trying to design coil arrays so as to optimize parallel imaging with linear SEMs, why not take a typical coil array as a point of departure, and instead try to choose optimal SEM fields from a nonlinear basis set? This question invites us to revisit the fundamentals of spatial encoding in MRI and to envision hybrid encoding functions that most efficiently pairs coil encoding and SEMs.

2.1 Improved RF coil-SEM complementarity in parallel imaging

For axial, single-slice imaging, the most common parallel imaging approach is to undersample the phase encode direction, shown in Fig. 2.1. To optimally recover this lost encoding using an RF array, one would need to lay out the coils along the undersampled direction. While such geometries could be practical for a small subset of clinical applications, such as spine imaging, they do not lend themselves to typical MRI scans of the body or head. For head imaging in particular, the most efficient way to place receiver loop coils is in a "soccer-ball geometry" (40), providing an encircling ring of coils in any particular slice. For this geometry, most of the spatial localization provided by the coils is in the azimuthal direction. If Cartesian imaging is performed a the phase encode direction is undersampled, only those coils oriented along this particular direction can supplant the missing phase encodes with localized coil encoding. By contrast, if projection imaging is performed with a radially-varying SEM, each encircling coil would be well positioned to disambiguate signals in the azimuthal direction while the radial SEM provided dense encoding along the radial direction.

2.1.1 Match encoding fields to the available SNR of RF receive coils

Conventional imaging with Cartesian SEMs provides the same resolution everywhere in the FOV. For a Cartesian grid, every voxel has the same two-dimensional sinc point spread function[1]. Given that surface coils in multi-channel arrays are always the most sensitive at the periphery (39, 40, 60), it stands to reason that finer features could be resolved in this region using reconstructions that have spatially-varying resolution. Nonlinear SEM shapes are a natural way to provide spatially-varying resolution. The final image may be

[1]When multiple RF coils are used, the exact sinc function may vary slightly over the FOV due to the spatially-varying impact of the coil profiles, but the PSF shape is still dominated by the SEM encoding

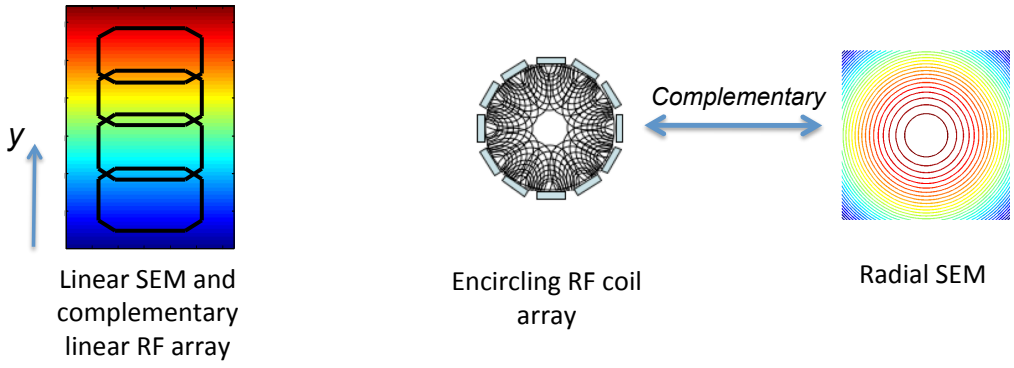

Figure 2.1: At left, a linear phase encoding SEM is shown along y, the direction in which k-space is often undersampled for parallel imaging. The color gradient represents the linear variation in the amplitude of the SEM (and corresponding spin frequency). When RF coils are arranged so as to provide optimal spatial encoding for unwrapping aliasing in the undersampled acquisition, the result is a linear array along y, which is not a practical geometry for most applications. RF coils are mor often arranged in an encircling array (shown at center), with their localized sensitivity isocontours drawn in black. Such an array disambiguates signals in the azimuthal direction, suggesting the use of a radially-symmetric SEM such as that shown at right. In this way, the chosen shape of the SEM complements the layout of the RF coils. Figure adapted from (59).

reconstructed onto a grid of identically-sized voxels, but the true resolution - the ability to resolve small adjacent object features - will be higher at the periphery, leading to a "sharper" image in this region and a "blurrier" image in the center of the FOV. Unfortunately, the resolution in the center depends almost entirely on the linear SEMs, so it is not clear how to improve resolution in this region without taking more Cartesian data. But if extra encoding can be provided at the periphery without compromising resolution at the center, and if parallel imaging performance benefits from improved complementarity between SEMs and RF coils, then nonlinear SEMs may be superior to linear SEMs for many applications.

2.2 PatLoc

Arguably the earliest framework for nonlinear spatial encoding was known as "PArallel imaging Technique with LOCalized gradients", or PatLoc. The first proposed PatLoc encoding scheme used non-bijective, curvilinear SEMs in an otherwise conventional Cartesian pulse sequence. In this version of PatLoc, one SEM is played as the readout and the other

field is played as the phase encode. The method potentially enjoys the advantages of (a) faster gradient switching without peripheral nerve stimulation, (b) improved spatial encoding in peripheral high-SNR region, and (c) excellent parallel imaging performance for circular RF coil arrays (10, 61). If two orthogonal multipolar SEMs are used for image encoding, then spatial encoding can be described using the real and imaginary parts of the conformal mapping, $f : \mathbb{C} \rightarrow \mathbb{C}$,

$$f(z) = z^L = (x + iy)^L = u(x, y) + iv(x, y) \tag{2.2.1}$$

Because this mapping preserves the local angle between the isocontours of the inputs x and y, the vector gradients of fields u and v are everywhere orthogonal ($\nabla u \cdot \nabla v = 0$), permitting use of the two fields as frequency and phase-encoding SEMs. This is equivalent to stating that u and v satisfy LaPlace's equation

$$\nabla^2 B(x, y) = 0 = \nabla^2 \left(u\left(x, y\right) + iv\left(x, y\right)\right) \tag{2.2.2}$$

where $B(x, y)$ is the magnetic encoding field formed by a superposition of u and v. Since all polynomial functions of the complex variable $z = x+iy$ are analytic and satisfy Laplace's equation, the shapes u and v are physically realizable for all values of L(62).

The resulting spherical harmonic fields of order L vary in polarity with angular position and grow as r^n, where r is the distance from the center. In the most general case, SEMs of any order can be used to perform PatLoc imaging. For purposes of illustration, I will describe the second-order case here in more detail, as discussed in (61). For L=2, the non-bijective fields u and v assume the hyperbolic forms $u(x, y) = x^2 - y^2$ and $v(x, y) = 2xy$, respectively, as shown in Fig. 2.2.

In this second-order instance of PatLoc imaging, the two curvilinear SEMs are used

in place of the usual readout and phase encoding fields in a Cartesian pulse sequence[1]. For each coil, $N/(R_{pe}L)$ k-space lines are acquired, where R_{pe} is the k-space undersampling factor, for a net acceleration factor of $R = R_{pe}L$ as compared with a fully sampled conventional acquisition.

The reduced B-field excursion over the bore of the scanner may permit faster SEM switching times for the same dB/dt as compared with linear SEMs, potentially allowing for faster imaging without violating safety limits for peripheral nerve stimulation[2]. Additionally, PatLoc fields are a natural complement to an encircling array of local receiver coil profiles. Such coils are well positioned to disambiguate aliased signals arising from different voxels that are mapped onto the same point in the non-bijective data domain. Also, the coils naturally provide higher SNR at the periphery, where PatLoc resolution is the greatest. However, the relatively flat frequency isocontours at the center of the FOV cause pronounced blurring and Gibbs ringing in this region. This principal limitation of the PatLoc encoding strategy described here can be overcome by combining curvilinear SEMs with linear SEMs in order to resolve features at the center of the FOV. This is the general approach used in schemes such as O-Space imaging, to be introduced shortly, and in emerging methods such as 4D-RIO (15).

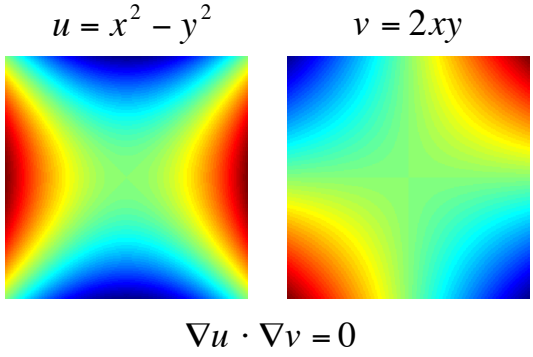

Figure 2.2: Hyperbolic PatLoc fields generated by the L=2 conformal mapping described in Eq. (2.2.1). Note that the local gradient of the two fields is everywhere orthogonal.

[1] Imaging with two multipolar SEMs and a Cartesian trajectory is just one particular instance of PatLoc imaging, and wide flexibility exists to alter the encoding scheme. For instance, a version of PatLoc has recently been proposed in which two multipolar fields traverse a radial k-space trajectory (63).

[2] This argument assumes that linear SEMs, when aligned along the human body, have the ability to create larger potential differences and induce stronger currents than comparable second-order fields, which would have a parabolic field shape over the same limb or segment of torso. In this way, while the local dB/dt may be just as high as for the linear SEMs, the nonlinear SEMs will induce smaller currents due to time-varying potential differences over extended regions.

2.2.1 Cartesian-trajectory PatLoc image-space reconstruction

As described in (61), when Cartesian trajectories are traversed by a pair of orthogonal multipolar PatLoc SEMs, reconstruction is readily achieved in the image domain using a generalization of SENSE (1). First, the k-space data in each RF coil channel are Fourier transformed to yield n_c distorted and possibly aliased images in the encoding space defined by curvilinear coordinates. For full-sampled data the images in encoding space are of size $N/L \times N$, where the desired size of the final image is $N \times N$. Because a given frequency occurs L times in of the pair of the multipolar SEMs, a non-bijective mapping occurs from image space to encoding space, causing L-fold voxel overlap in the curvilinear encoding space. During reconstruction, these L source points must be "unwrapped" using RF coils, each of which is preferentially sensitive to one of the L source voxel signals. This results in L separate images in encoding space. The unwrapping problem is analogous to the SENSE unwrapping described in the previous chapter and may be solved by simple matrix inversion, as for the case of SENSE reconstruction (61). Care must be taken to map the coil profiles from the Cartesian image domain representation into the encoding domain, interpolating when necessary.

At this point, the L intermediate images must be intensity corrected. The conformal mapping used to generate PatLoc fields, $f(z) = z^L$, maps $\mathbb{C}$ into $\mathbb{C}$ and has a well-defined Jacobian determinant. For the particular case of L=2, the determinant is

$$J = det \begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix} = 4x^2 + 4y^2 \tag{2.2.3}$$

The Jacobian determinant describes the amount by which the mapping $f(z)$ shrinks or dilates areas at each point in the FOV. This provides a simple and approximate way to describe the spatially-varying resolution of a PatLoc image[1]. The inverse of the Jacobian determinant

[1]Note that for the identity mapping, L=1, the Jacobian determinant is simply 1, demonstrating that

provides a simple way to correct for the intensity of unwrapped PatLoc images to prevent the voxels that are large in the PatLoc domain from appearing bright in the Cartesian domain.

The final step in reconstruction is to perform the inverse of the coordinate mapping $f(z) = z^L$, thereby unwarping the distorted image. This function may be inverted analytically or, equivalently, the inverse mapping may be performed mathematically. Careful interpolation must again be performed to obtain image values on the desired Cartesian grid.

PatLoc acquisitions can be accelerated by simply skipping phase encode lines, just as in conventional parallel imaging. This causes the intermediate images in encoding space to be aliased. Unwrapping of the aliased signals can be performed on each of these L images just as in the case of SENSE. Simulations of undersampled L=2 PatLoc imaging are shown below in Fig. 2.12. The Jacobian determinant vanishes at the origin, giving rise to images with very sharp resolution at the periphery and minimal ability to resolve features at the center of the FOV.

2.3 O-Space imaging

In O-Space imaging (14), the explicit goal from the outset is to optimize parallel imaging performance by using combinations of multiple SEMs to form field shapes tailored to the spatial information contained in the available coil profiles. In principle, different SEM shapes can be chosen for each successive echo to obtain suitable projections of the object. In the most general case, an array of surface coil profiles is selected and its spatial encoding properties are assessed. An encoding scheme is then designed so as to optimally complement the spatial information provided by the array. For purposes of this thesis, only the axial plane is considered. In this plane the RF coils typically encircle the object being im-

Cartesian encoding has the same resolution at every voxel

aged. The coil profiles vary smoothly throughout the FOV, with the area of peak sensitivity localized to the angular region subtended by each coil. A circumferential array therefore provides more encoding in the angular direction than in the radial direction, a fact that has not been exploited by past encoding schemes based on linear SEMs alone. By using both linear and higher-order SEMs, O-Space imaging attempts to benefit from this properties of nonlinear SEMs while preserving resolution at the center of the FOV.

Radially-varying fields are a natural choice for providing encoding that complements an encircling array of RF coils. However, many of the shapes that come to mind, such as $f(r) = |r|$, have discontinuous spatial derivatives and do not satisfy Laplace's equation, which governs all magnetostatic fields. One promising set of functions that do satisfy Laplace's equation are the spherical harmonics. The quadrupolar fields described in section 2.2 for one instance of Cartesian-trajectory PatLoc imaging are second-order spherical harmonics. Combinations of spherical harmonics are also widely employed in MRI and MRS as a basis set for "shimming" out inhomogeneities in the B_0 field. These harmonics, listed up to fourth order in Table 2.1, form a rich palette for designing nonlinear encoding schemes. Each Cartesian harmonic is formed using a linear combination of the conventional spherical harmonics, expressed in spherical coordinates (r, ϕ, θ), with a spherical-to-Cartesian change of coordinates.

For proof of concept in this study, the encoding problem was restricted to axial plane imaging only. The first 9 spherical harmonics that are non-degenerate in the axial plane are shown in Fig. 2.3. The subset of fields chosen for O-Space imaging are presented in the first row. The combination of X, Y, and Z^2 SEMs provides a natural point of departure for an exploration of combined first and second-order projection imaging. This choice of SEMs is motivated by two factors: (1) the ability of the Z^2 SEM to provide excellent spatial encoding along the radial direction, where circumferential coil arrays provide the least encoding; and (2) the ready availability of coil designs for producing the Z^2 spherical har-

Shim name	$f(x,y)$	Shim name	$f(x,y)$
X	x	XYZ	xyz
Y	y	X^3 or $C3$	$x\left(x^2-3y^2\right)$
Z	z	Y^3 or $S3$	$y\left(3x^2-y^2\right)$
Z^2	$z^2\frac{1}{2}\left(x^2+y^2\right)$	Z^4	$z^4-3z^2\left(x^2+y^2\right)+\frac{3}{8}\left(x^2+y^2\right)^2$
XZ	xz	XZ^3	$xz^3-\frac{3}{4}zx\left(x^2+y^2\right)$
YZ	yz	YZ^3	$yz^3-\frac{3}{4}zy\left(x^2+y^2\right)$
X^2-Y^2 or $C2$	x^2-y^2	$Z2C2$	$z^2\left(x^2-y^2\right)+\frac{1}{6}\left(x^2-y^2\right)\left(x^2+y^2\right)^2$
XY or $S2$	xy	$Z2S2$	$xyz^2-\frac{1}{3}xy\left(x^2+y^2\right)$
Z^3	$z^2-\frac{3}{2}z\left(x^2+y^2\right)$	$ZC3$	$z\left(x^3-3xy^2\right)$
XZ^2	$x\left(4z^2-\left(x^2+y^2\right)\right)$	$ZS3$	$z\left(3yx^2-y^3\right)$
YZ^2	$y\left(4z^2-\left(x^2+y^2\right)\right)$	$C4$	$x^4+6x^2y^2+y^4$
$Z(X^2-Y^2)$	$z\left(x^2-y^2\right)$	$S4$	$4\left(yx^3-xy^3\right)$

Table 2.1: The spherical harmonics up to fourth order, as expressed in Cartesian coordinates. Reproduced from (64)

monic (60, 65). Additionally, the quadratic field varies more steeply near the periphery, where surface RF coil sensitivity is typically the highest (40),(39).

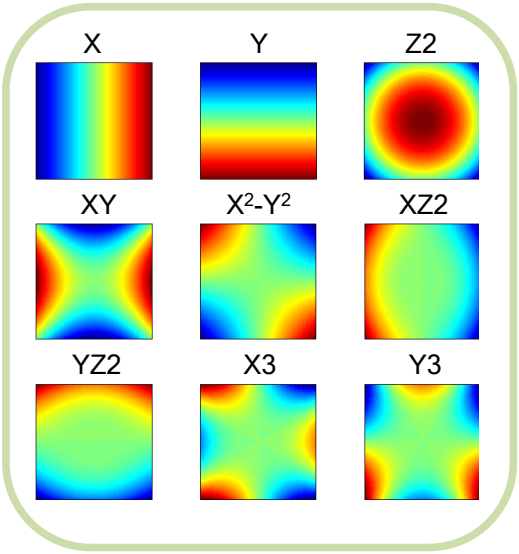

Figure 2.3: Candidate spherical harmonics for nonlinear encoding.

The principal drawback of this combination of fields, it should be noted, is that while they are orthogonal over three-dimensions, they are not orthogonal over the entire two-dimensional axial plane ($\nabla B_a \cdot \nabla B_b \neq 0$). The central hypothesis of this thesis is that the loss in encoding efficiency is offset by the fact that O-Space imaging preserves the high pe-

ripheral resolution and coil-field complementarity of PatLoc imaging while also providing resolution at the center of the FOV using the linear SEMs. While O-Space can in principle never do better than conventional encoding methods like spiral, radial, or Cartesian trajectories at the center of the FOV[1], it may substantially outperform them elsewhere in the FOV.

O-Space imaging uses the quadratic axial field of the Z^2 harmonic in combination with linear SEMs to provide spatial encoding throughout the FOV. The difference between O-Space and conventional imaging is illustrated in Fig. 2.4 for a simplified one-dimensional case along the x-axis. The figure depicts two RF receive coils placed on opposite sides of the object to be imaged. Imaging with a linear X SEM is depicted in the top diagram. Because the image resolution varies with the slope of the SEM, and the X SEM has constant slope, this field provides the same resolution at every position along the x-axis. By contrast, the quadratic SEM shown in the middle diagram has a linear slope that is large at the periphery near the surface coils and passes through zero at x=0. This provides excellent resolution at the periphery but no spatial encoding at the center.

To remedy this problem, the X SEM is switched on during the second readout, causing the quadratic shape to shift from x=0 to x=$-x_0$. The position of x_0 depends on the strength of the linear SEM played. Note the regions shown in yellow near Coil 2 in which the quadratic SEM slope exceeds that of the linear SEM, providing superior resolution. When the quadratic field is shifted, the yellow region grows considerably, greatly improving the resolution near Coil 2. In this way, the linear and quadratic SEMs can be combined in successive readouts to provide the desired resolution throughout the FOV.

Generalizing to 2-D acquisitions, conventional phase encoding is discarded and replaced by projection acquisitions with the center of the Z^2 function shifted off center using the X and Y SEMs, as shown in Fig. 2.5. With each acquired echo, the object is pro-

[1]Because nonlinear SEMs provide no spatial encoding at the center point of the FOV, there is no way to acquire more spatial information at this location with linear and nonlinear SEMs than that which could be obtained from linear SEMs alone.

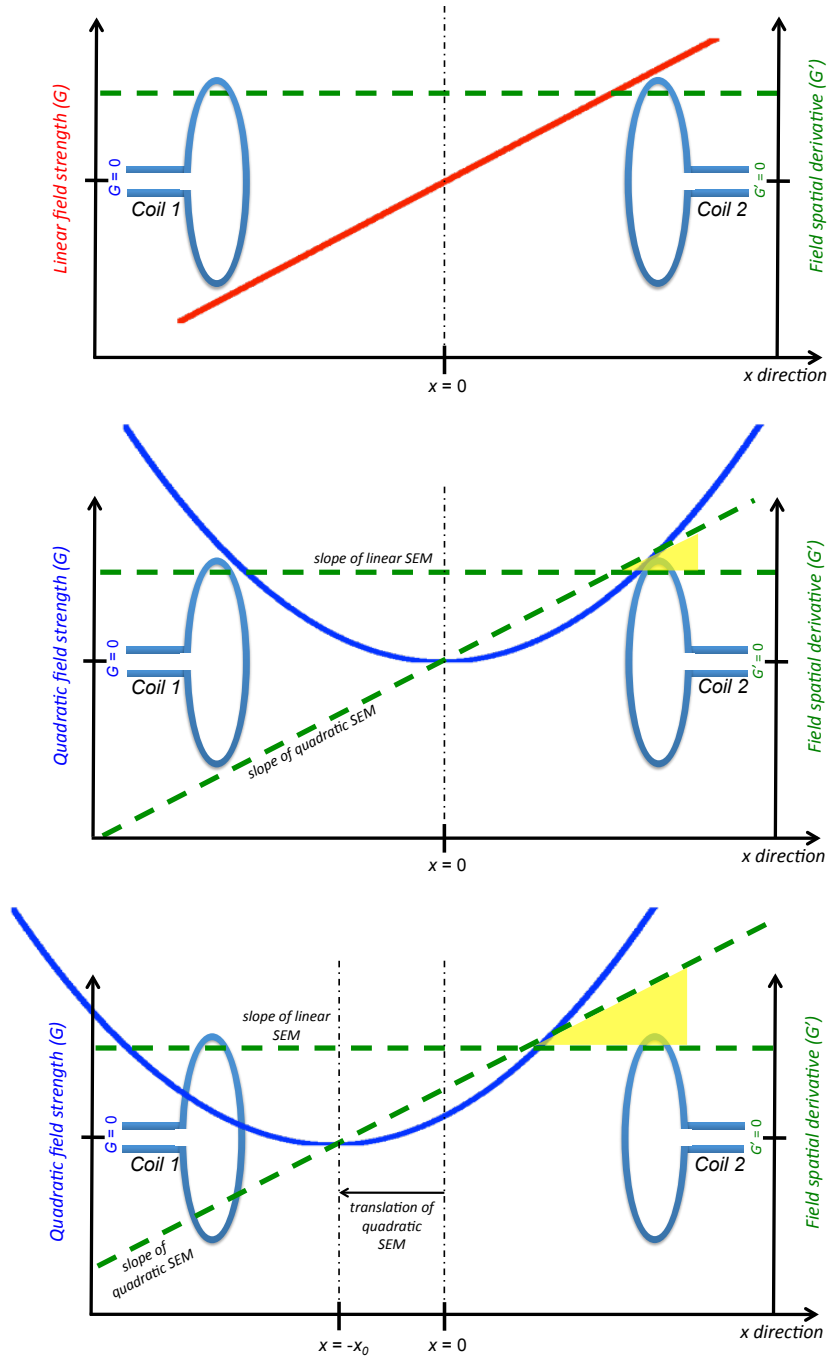

Figure 2.4: One-dimensional illustration of O-Space encoding. Encoding with a linear SEM (top diagram) has constant slope, providing equal resolution at all points in the FOV. By contrast, a quadratic SEM (middle) has linear slope, providing dense encoding at the periphery and no encoding at the center. To recover resolution at the center, the linear and quadratic SEMs are played simultaneously during a successive signal readout (bottom). The regions in yellow show improvements in resolution provided by the quadratic field near Coil 2. See text for further discussion.

jected onto a set of frequency isocontour rings that are concentric about a different center placement (CP) in the FOV, suggesting the term "O-space imaging". By shifting the Z^2 quadratic shape off center, it is ensured that there are enough overlapping isocontours from different projections at the center of the FOV to resolve features in this region. At the same time, the curvilinear isocontours of the quadratic field hold the potential to better complement an encircling array of RF coils, particularly when the shift away from isocenter is small. The O-Space pulse sequence 2.5 is essentially a conventional radial sequence to which a Z^2 SEM pulse has been added.

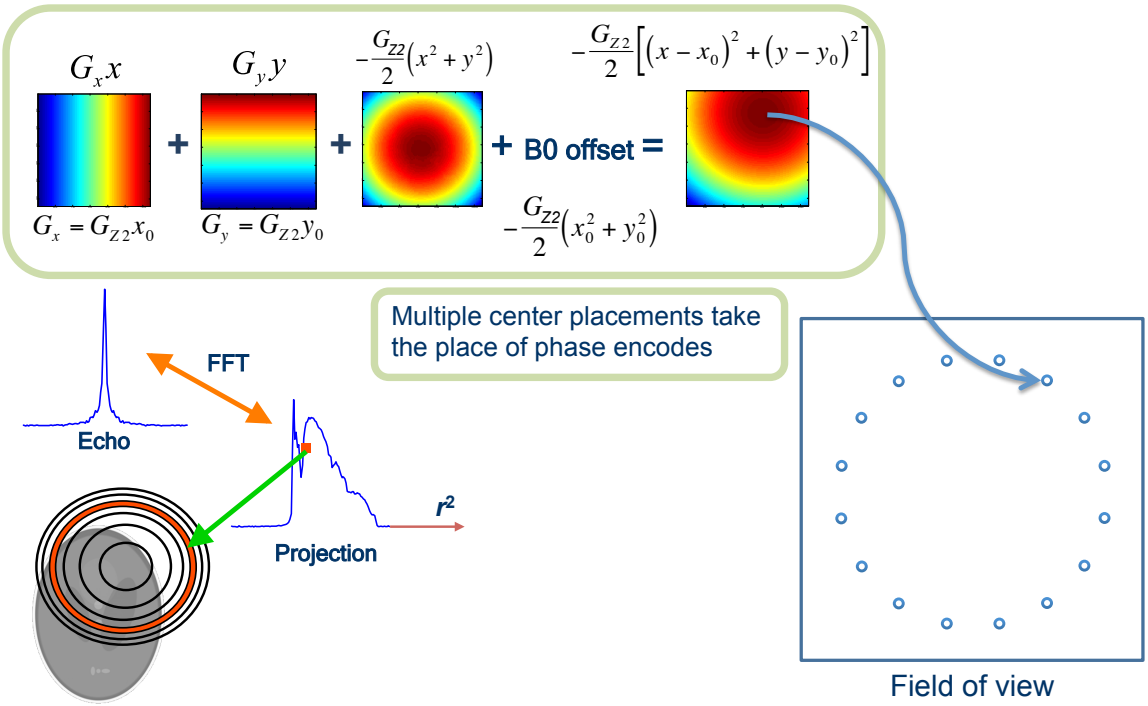

Figure 2.5: In O-Space imaging, a quadratic SEM is translated to different center placements (CPs) between successive readouts using different combinations of linear SEMs. The FFT of the readout corresponds to projections of the object along (complex-valued) ring isocontours. The CPs shown here reduce to a radial k-space trajectory in the absence of the Z^2 field.

The Fourier transform of an echo obtained in the presence of a radially symmetric SEM yields a projection of the object onto a set of concentric rings. With radial localization provided by the SEMs, the surface coils are ideally positioned to provide spatial localization in the angular direction. Furthermore, as will be shown in section 2.5.4, since the readout

SEM provides spatial encoding in two dimensions, rather than just one as in Cartesian trajectories, additional encoding is provided by increasing the SEM strength and sampling the echo more densely, with essentially no impact on the imaging time. In Cartesian parallel imaging, densely sampling the echo increases resolution in the readout direction but does not reduce the amount of aliasing in the phase-encode direction.

2.3.1 Spatially-varying resolution that can be tailored to the ROI

When nonlinear SEMs are played during an MRI readout, spins in each voxel are advanced along a different k-space trajectory. This fact has motivated the use of "local k-space" plots to compare the k-space paths traversed by voxels in different parts of the FOV (15, 66). Spins in each voxel are "driven" through k-space based on the local "gradient" that they experience in the magnetic field,

$$\phi(\mathbf{x}, t) = \gamma \int_0^t \mathbf{x} \cdot \mathbf{G}(t')dt'. \quad (2.3.1)$$

It is useful to Taylor expand the spatial component of the phase function around a given point $\mathbf{x_0}$ as follows (15):

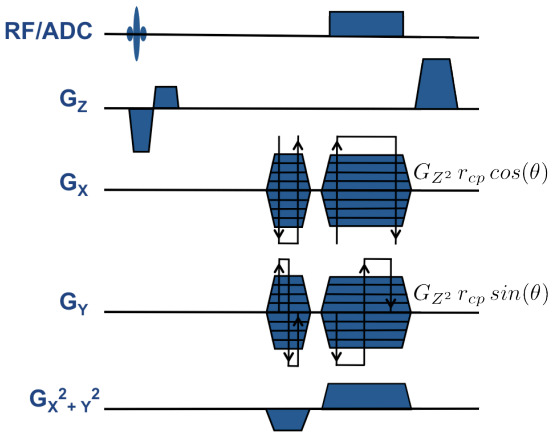

Figure 2.6: The O-Space imaging sequence (left) consists of a radial trajectory modified by the addition of the Z^2 SEM. For the encoding scheme shown in Fig. 2.5, the linear SEMs are varied while the same Z^2 amplitude is repeated during each readout. The angle of the linear k-space spokes, θ, varies between 0 and 2π. The amplitude of the Z^2 is chosen to position the ring of center placements at r_{cp} in the FOV according to the relation $G_{Z^2} = G_{x_{max}}/r_{cp} = G_{y_{max}}/r_{cp}$. The linear SEMs shift the encoding field to a point along the chosen ring of center placements based on the amplitude of the Z^2 field, G_{Z^2}. In the absence of the Z^2 SEM the pulse sequence reverts back to a conventional radial sequence.

$$\phi(\mathbf{x}) \approx \phi(\mathbf{x_0}) + \nabla\phi(\mathbf{x_0})\,(\mathbf{x} - \mathbf{x_0}) + \ldots \quad (2.3.2)$$

Figure 2.7: Local k-space trajectories for a representative 5×5 grid of voxels from different sections of the FOV. The trajectories correspond to the O-Space readouts for the center placements shown in Fig. 2.5. At left is a conventional radial k-space trajectory corresponding to the center of the FOV, where the Z^2 SEM has zero amplitude, imparting no $\mathbf{k}_{shift}$. Elsewhere in the FOV, the Z^2 applies a $\mathbf{k}_{shift}$ vector, altering the shape of k-space (and the resulting resolution) on a voxel-wise basis. The k-space plots do not include the effects of RF coil encoding.

The zeroth-order term describes the conventional k-space trajectory traversed by the linear SEMs. The first-order term in the expansion approximates the spatially-varying k-space shift vector applied by the nonlinear SEM,

$$\mathbf{k_{shift}}(\mathbf{x_0}, t) \approx \nabla\phi(\mathbf{x_0}, t) \tag{2.3.3}$$

The local k-space trajectory is the sum of the zeroth-order term and the spatially-varying first-order term. For the CP scheme shown in Fig. 2.5, the local k-space trajectory traversed by representative subset of voxels is shown in Fig. 2.7. Local k-space plots clearly show that different spatial encoding occurs at the center and at the periphery in O-Space imaging. Note that for the voxel at the center of the FOV, local k-space reduces to a conventional radial trajectory with 256 spokes, showing that only linear SEMs provide significant encoding in this region. For voxels near the periphery, the quadratic SEM extends k-space coverage, causing the trajectory assumes a "barbell" shape. It should be noted, though, that some regions are not covered by the barbell k-space trajectories.

Local k-space could conceivably be used as a guide in crafting new trajectories from

a given set of SEMs. But at least two issues must first be resolved. First, where the effects of the O-Space nonlinear SEM are strongest, it is not obvious what information can be gleaned from the trajectories by themselves. In the barbell local k-space plots, some parts of k-space are sampled out to higher spatial frequencies, suggesting higher resolution has been achieved, while other areas of k-space are skipped altogether. Perhaps this allows higher resolution at the periphery owing to the conjugate symmetry of k-space quadrants for real-valued objects?

A second limitation of the approach as presented is its inability to model the effects of RF coils on spatial encoding. When linear SEMs are used, it is straightforward to calculate the effect of a coil profile on a given k-space data set. One simply convolves the k-space data with the Fourier transform of the coil sensitivity. However, when considering the local k-space trajectory for a single voxel, which is represented by one voxel in a coil sensitivity map, it is not clear how to use this coil information (or its spatial derivative) to calculate local k-space trajectories on a coil-by-coil basis.

When trajectories are designed with coil complementarity and parallel imaging in mind, the differences between each coil's k-space data are important to consider. The concept of O-Space imaging arises from image-domain consideration of coil and SEM encoding in polar coordinates, but additional insights could be gleaned from comparisons of local k-space trajectories if they could be made to reflect coil sensitivity information. This remains an open question of research.

2.3.2 O-Space image reconstruction

The Z^2 SEM is translated to the desired CP in the FOV using linear SEMs and a amplitude of static field offset to complete the square (Fig. 2.5). After slice selection, the X, Y, and Z^2 SEMs are used to dephase and rephase the spins, as in conventional projection imaging.

The signal equation for the l^{th} acquired echo is

$$s_{l,q}(t) = \int m(\mathbf{x}) C_q(\mathbf{x}) e^{i2\pi \bar{\gamma} G_{Z^2} \frac{1}{2}((x-x_l)^2+(y-y_l)^2)t} d\mathbf{x} \tag{2.3.4}$$

where $\bar{\gamma} = \frac{\gamma}{2\pi}$, $m(x,y)$ is the object magnetization (including the effects of T_1, T_2, and other contrast mechanism), $C_q(x,y)$ is the q^{th} receive coil sensitivity, (x_l, y_l) specifies the l^{th} CP, and $\bar{\gamma} G_{Z^2}$ is the strength of the Z^2 spherical harmonic in Hz/cm^2. SEM strengths G_X and G_Y are chosen such that $G_X = G_{Z^2} x_l$ and $G_Y = G_{Z^2} y_l$ in Hz/cm. Echoes formed using different CPs during successive TRs comprise a dataset from which the image is reconstructed. The number of CPs used to form an image is equivalent in terms of acquisition time to the number of phase encodes used in a conventional Cartesian sampled acquisition. In the discrete case, the integral kernel is represented as a projection matrix $E_{t,l,q}$ where the rows describe time point t, CP l, and coil q, while the columns correspond to voxels in the object. The object is vectorized and the echoes and encoding functions from multiple CPs and coils are stacked to produce a single matrix equation:

$$s = Em \tag{2.3.5}$$

2.3.2.1 Approximate reconstruction based on ring isocontours

If radius r_l is defined relative to each CP, the integral may be recast in polar coordinates:

$$s_{l,q}(t) = \int \int m(r_l, \phi) C_q(r_l, \phi) \exp\left(-i2\pi \bar{\gamma} G_{Z^2} \frac{1}{2} r_l^2 t\right) r_l dr_l d\phi \tag{2.3.6}$$

With the choice of $u_l = \frac{1}{2} r_l^2$, the radial integral in 2.3.6 becomes a Fourier transform:

$$s_{l,q}(t) = \int \int m(u_l, \phi) C_q(r_l, \phi) \exp\left(-i2\pi \bar{\gamma} G_{Z^2} u_l t\right) du_l d\phi \tag{2.3.7}$$

The inverse Fourier transform of each echo now yields $P_{l,q}(u)$, the projection of the object along the isocontours encircling the l^th CP. The projection specifies the amount of energy in the coil profile-weighted object that is smeared around ringlike regions that decrease in width with increasing r_l (as in Fig. 2.5).

Because the encoding function is not in the form of a Fourier integral kernel, the data do not reside in k-space. Consequently, image reconstruction cannot be achieved using k-space density compensation and re-gridding approaches similar to those employed in non-Cartesian imaging with linear SEMs. Due to this fact, image reconstruction is performed by directly solving the matrix equation $s = Em$ using one of two methods to be discussed below: a spatial domain algorithm based on the projections and a frequency domain algorithm based on the echoes.

When projections are obtained using the discrete Fourier transform, each point in the projection $P_{l,q}[u]$ corresponds to the sum of the object intensity at all voxels lying within a band that can be approximated as the u^{th} isofrequency ring. If N_s samples are acquired during readout, then there exist N_s isofrequency rings. The radius of the outermost ring is specified by

$$r_{max} = \sqrt{BW/G_{Z^2}} \tag{2.3.8}$$

where BW is the readout sampling bandwidth. The sum over all voxels lying within a given ring is weighted by q^{th} receive coil profile at each point within the ring,

$$P_{l,q}[u] = \sum_{x,y \in ring\ u} C_q[x,y]m[x,y]. \tag{2.3.9}$$

When the object and coil profiles are represented in vector form, the set of all ring-domain equations may be vertically concatenated to form a single matrix equation accounting for

all N_s rings, L CPs, and n_c coils in the element-wise product of C and W:

$$P[u] = [C \circ W]\, m = E\, m \tag{2.3.10}$$

where C contains the coil profiles and W is a sparse matrix whose u^{th} row weights each voxel according to its contribution to the u^{th} ring of a given CP. The simplest version of W contains ones for each voxel lying within the u^{th} ring and zeroes elsewhere. For an $N \times N$ reconstruction, the encoding matrix E is of size $[N_s \times L \times n_c, N \times N]$. Direct inversion of this matrix is challenging for practical imaging applications, but the sparsity of the matrix can be exploited by a conjugate gradient-type algorithm known as LSQR (67) that is available as a function call in MatLab (MathWorks, Natick, MA). LSQR was selected for its ability to quickly solve sparse, non-square, complex-valued matrix equations.

2.4 Center placement optimization

For simplicity, it was decided that the Z^2 SEM would play at the same amplitude during each O-Space readout. The question remained, then, how to set the X and Y SEM amplitudes at successive readouts; in other words, a sequence of center placements was sought. In the LSQR simulations, 128×128 reconstructions from 128-point echoes were used to determine a highly efficient CP scheme within the FOV for datasets composed of 32 and 16 echoes (59). By analogy to Cartesian parallel imaging, this corresponds to 4-fold and 8-fold undersampling, respectively. A variety of coil geometries were also considered, ranging from 8 to 32 circumferentially distributed loop coils, for which B-fields can be approximated in the magnetostatic limit using the exact analytical field expression for circular loops (68). For simplicity and computational efficiency, a ring approximation to the true PSF was used in which each voxel lying between two frequency isocontours was blurred

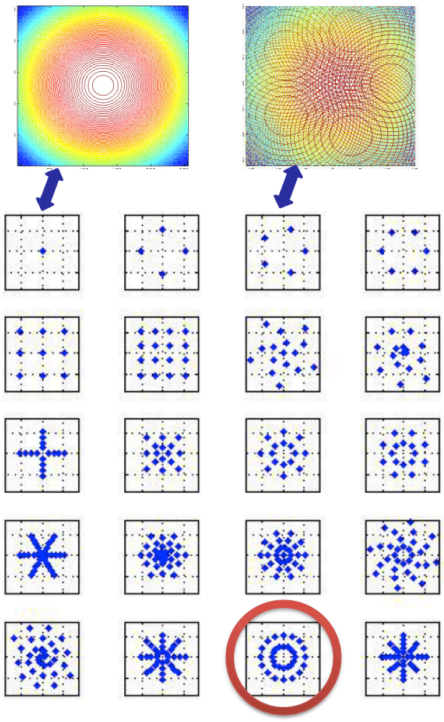

Figure 2.8: Some of the center placement schemes simulated for parallel imaging performance using the LSQR algorithm. The CP scheme circled in red was found to provide low mean squared error for a variety of coil array sizes and undersampling factors. This scheme corresponds to two interleaved sets of undersampled radial k-space trajectories traversed by the X and Y SEMs. Figure reprinted from (59).

evenly over all voxels enclosed between the two contours.

The CP scheme yielding the minimum mean squared error reconstruction in LSQR simulations is shown in Fig. 2.8. In the absence of the Z^2 SEM, this scheme reduces to two interleaved, undersampled radial k-space trajectories played on the linear SEMs. This is a natural result, since radial trajectories have long been valued for the graceful degradation of their noise and resolution performance when they are undersampled. In the context of O-Space imaging, the linear SEMs provide all the resolution at the center of the FOV, with the quadratic field enhancing resolution near the periphery.

2.5 Reconstruction using the full encoding matrix

The spatial-domain LSQR reconstruction amounts to backprojecting points in each projection onto the corresponding rings in the image. For a fast search across many potential encoding schemes, the approximation of real-valued ring frequency isocontours was con-

sidered adequate. However, for more precise image reconstruction, this approximation poses several difficulties. First, it does not fully describe the spatially varying, complex-valued point spread function of each applied SEM shape. Second, care must be taken in defining the boundary between rings. Third, radial Gibbs ringing of the PSF due to convolution with the Fourier transform of the acquisition window (typically rectangular) must also be considered. Although a sparse encoding matrix comprised of distinct rings is computationally more tractable, it ignores the fact that each PSF has sidelobes that overlap with other ring isocontours.

These obstacles can be surmounted by directly solving Eq. (2.3.4) in the frequency (echo) domain using the Kaczmarz method, also known as the algebraic reconstruction technique, an iterative row-action backprojection algorithm that has been applied to computed tomography and cryo-electron microscopy (69),(70). This algorithm compares each echo time point with the inner product of the appropriate row of the projection matrix, denoted $E_{t,l,q}$ and the n^{th} iterate of the image estimator. The difference between these scalars weights the amount of basis function $E_{t,l,q}$ that is added to the estimator going into the next iteration:

$$\hat{m}^{(n+1)} = \hat{m}^{(n)} + \lambda \frac{s_{i_n} - \left\langle E_{i_n}, \hat{m}^{(n)} \right\rangle}{\| E_{i_n} \|^2} E_{i_n}^*, \qquad 1 < i_n < n_K n_c \tag{2.5.1}$$

where $*$ denotes complex conjugation and λ is a relaxation parameter[1] used to control the rate of convergence ($0 < \lambda < 2$). The n^{th} row backprojection that is performed is referred to as the n^{th} *step*, as indicated in the superscript of m. Once the algorithm has stepped through all chosen rows of the encoding matrix, one *iteration* has been completed. The index i_n, known as the *cyclic control*, typically increments from 1 to $n_K n_c$ during each iteration through the rows of the encoding matrix and the corresponding points of acquired signal data, s. But the rows of E may be backprojected in any order and in the most general case

[1] Choice of values for λ will be discussed further in section 3.7

some rows may even be omitted to speed reconstruction[1]. For the simulations and empirical images shown in this thesis, a straightforward cyclic control is chosen in which all rows of E are backprojected and i_n steps directly from 1 to $n_K n_c$.

For each cycle through the encoding matrix, the row corresponding to each time point (for each CP) for a particular receive coil is backprojected once. The row is treated as a basis function that could be "missing" from the current image estimate. If the dot product between the given row and the current image estimate $\hat{m}_n$ differs from the corresponding data point s_n, this difference is used to weight the row as it is added to the $(n+1)^{th}$ image estimate. In theory, convergence is achieved when this difference becomes arbitrarily close to zero. In practice, a stopping criterion must be defined based on image statistics.

The entire projection matrix is usually too large to fit in memory, precluding the use of LSQR as a Matlab call. But with the Kaczmarz approach, only one data point is treated at a time, permitting individual basis functions (rows) to be recomputed on the fly or loaded from the hard drive. For sufficiently small values of λ, under-relaxed Kaczmarz reconstructions have been shown to converge to the minimum-norm least-squares estimate for $\hat{m}$ (72), equivalent to that obtained using the pseudoinverse of E.

The structure of the encoding matrix is shown in Fig. 2.9. At top is a pictorial representation of how each line of the encoding matrix, E, is calculated using an example RF coil profile and O-Space SEM (ignoring the shim map). Below this a simplified encoding matrix is drawn corresponding to a hypothetical O-Space acquisition having only 2 RF coils, 2 center placements (CPs), 4 readout time points, and 12 voxels in the FOV. The encoding matrix is populated by advancing t to calculate the spin phase evolution for all voxels in the FOV at each readout time point of a particular CP. This phasor is weighted by the RF coil

[1] Some investigators (71) have achieved faster convergence using randomly-selected rows of the encoding matrix for each step, but this has only been conclusively shown for the case of consistent systems of equations. In the presence of noise and systematic error, O-Space equations are inconsistent, so it is not clear that random sampling speeds convergence.

magnitude and phase and then vectorized to become a row of the encoding matrix. Once the rows for the first CP are populated, the O-Space SEM is shifted to a second CP, and the process repeats for the time points in this second readout. The SEMs for both CPs are then weighted by the other RF coil profile and the second half of the encoding matrix is populated. The encoding matrix is then used in combination with the acquired data, s, to reconstruct the object, m, using the Kaczmarz method.

There is no single unique way to populate the encoding matrix. It is also possible to interleave data from coil channels, such that the algorithm steps through all coils at a given time point before advancing to the next time point for a given center placement. In this case, care must be taken to rearrange the entries in the data vector s accordingly.

2.5.1 Comparison by simulation

Simulations were used to investigate O-space imaging in more detail. First, LSQR image-domain reconstructions at low resolution were used to quickly explore a variety of CP schemes to identify one that provides efficient O-Space encoding. Second, once a CP scheme had been chosen, Kaczmarz frequency-domain reconstructions at high resolution were used to compare O-Space with SENSE, PatLoc, and radial reconstructions over a wide range of acceleration factors. Third, Kaczmarz reconstructions were used to investigate the degradation of O-Space and SENSE reconstructions in the presence of increasing amounts of noise. Fourth, Kaczmarz simulations were used to explore the effect of increased ring density on the resolution of O-space images.

Two phantoms were used for each Kaczmarz simulation: an axial brain image (Fig. 2.10) (a) obtained using a conventional acquisition and a numerical phantom (b) designed to illustrate the spatially varying resolution and contrast properties of O-space encoding SEMs. The numerical phantom incorporates small lesion-like features at four contrast lev-

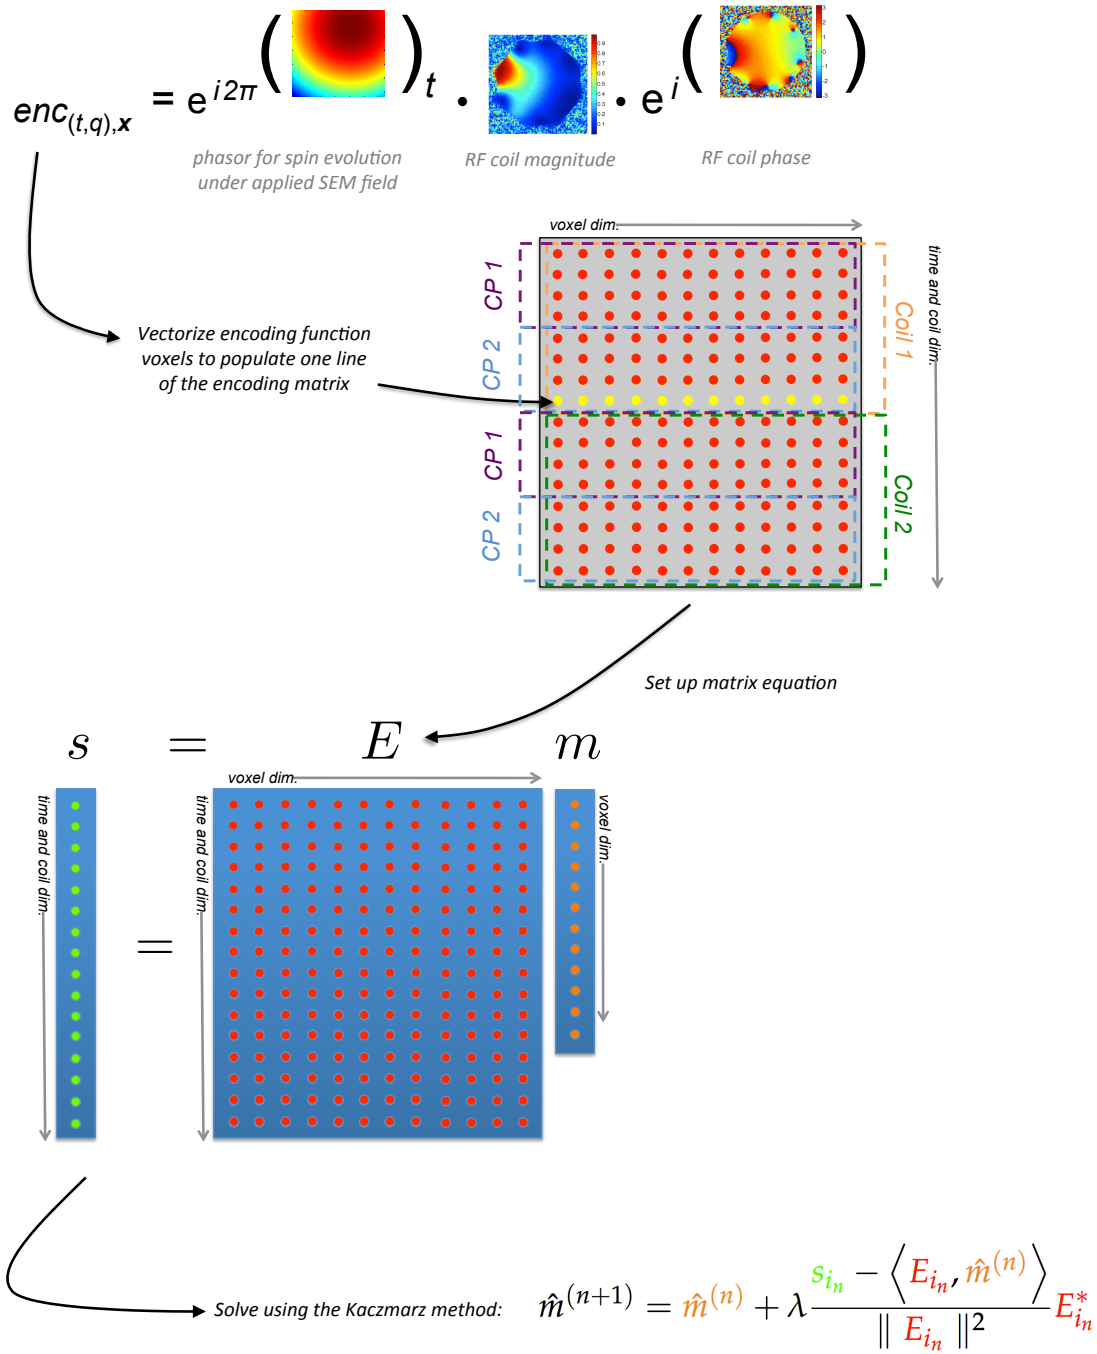

Figure 2.9: Anatomy of the encoding matrix. At top is a diagram of the calculation performed to populate each row of the encoding matrix. The encoding function is calculated over all voxels in the FOV for a given readout time point and then vectorized for insertion into the encoding matrix, E. The matrix is organized in blocks corresponding to different RF coil profiles and center placements. Once the matrix is populated, the equation $s = Em$ may be solved to obtain an image estimate, $\hat{m}$, via the Kaczmarz method. See text for more detail.

els, as well as sharp edges in the X and Y directions.

Once a highly efficient CP scheme was chosen, the Kaczmarz algorithm was used to perform reconstructions by directly solving the frequency-domain matrix representation of the signal equation for simulated 512-point echoes. O-space 256×256 reconstructions were compared to time-equivalent SENSE, quadrupolar PatLoc (Cartesian trajectory), and radial reconstructions for R={4, 8, 16}. Since the scope of this study is limited to two-dimensional imaging, SENSE acceleration was only performed along one phase encode direction. PatLoc reconstructions (61) were performed using $L = 2$ multipolar fields with undersampling of the phase-encode k-space by factors of 2, 4, and 8, yielding net acceleration factors of R=4, 8, 16.

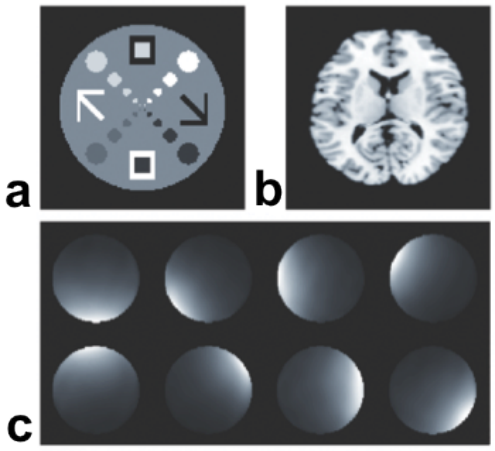

Figure 2.10: Coil profiles and phantoms used for O-Space simulations with the Kaczmarz algorithm.

Previous studies have displayed high-quality parallel reconstructions from undersampled multi-coil radial k-space data (35). If the Z^2 SEM is turned off during an O-space acquisition, what remains in the case of a circular arrangement of CPs closely resembles a radial acquisition. In this vein, a comparison between O-space and undersampled radial reconstructions is useful for isolating the effect of the extra encoding provided by the addition of the Z^2 SEM. For radial reconstructions, the Nyquist criterion requires $N\pi/2$ spokes for an image of size N. However, for ease of comparison against the SENSE, PatLoc, and O-space images, radial reconstructions using N/R spokes were considered to be "R undersampled," neglecting the extra factor of $\pi/2$ in the undersampling factor. Radial reconstructions were performed using $N\sqrt{2}$ readout points per spoke, consistent with the

Nyquist criterion for a fully sampled radial dataset. For further ease of comparison, the Kaczmarz algorithm was used to reconstruct the radial images based on the full encoding matrix.

Gaussian noise was added to each point of the simulated echoes. The noise standard deviation was scaled relative to the mean intensity in the phantom and then further scaled by the phantom dimension N to account for two-dimensional Fourier transformation into the frequency domain. Noise correlations between the coil channels were neglected for the purposes of this study and will be treated in future work. It is expected that these correlations will similarly impact both conventional Cartesian SENSE reconstruction and the proposed O-Space approach.

To explore the effects of ring density on resolution, the readout window was held constant while SEM strength and the number of readout points (N_s) were incrementally increased. Extra channel noise was injected into the echoes to model the increased sampling bandwidth. Care was taken not to exceed practical SEM strengths for a Z^2 coil design.

The phantoms and coil profiles (12 cm FOV) used in the high-resolution Kaczmarz simulations are shown in Fig. 2.10. For reference, the geometry factor, or *g*-factor, was computed (1) for SENSE for several acceleration factors (Fig. 2.11).

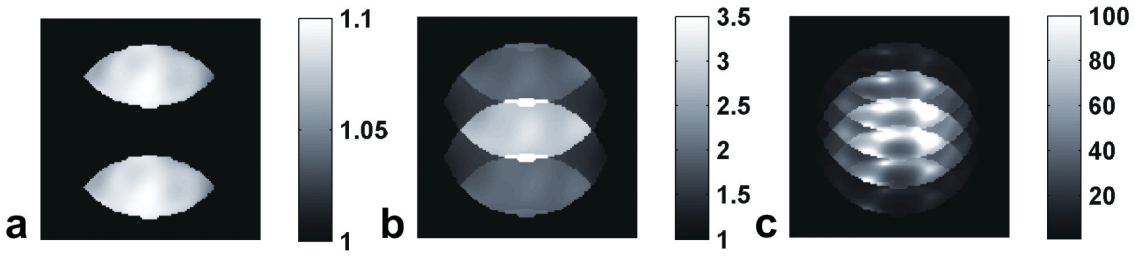

Figure 2.11: The geometry factor, or *g*-factor, which maps spatially-varying noise amplification, for the coil array used in O-Space Kaczmarz simulations, shown for R=2 (a), R=4 (b), and R=8 (c). The *g*-factor assumes regular k-space undersampling along the vertical (Y) dimension.

2.5.1.1 Regularized SENSE reconstructions

For the R=4 case, O-space and SENSE reconstructions were compared in the presence of varying amounts of noise. Noise amplification in the SENSE reconstructions is mitigated using Tikhonov regularization (73). This approach minimizes a functional that includes the data consistency term as well as a weighted L_2-norm on the unaliased voxels v:

$$\| Sv - a \|^2 + \| \Gamma v \|^2 \tag{2.5.2}$$

where v are the source voxels, S is the aliasing matrix from Eq. (1.5.1), a holds the aliased voxel for each coil's aliased image. For Tikhonov regularization, $\Gamma = \beta I$, for $n_c \times n_c$ identity matrix I and regularization parameter β. The solution that minimizes this functional is given by

$$\hat{v} = \left(S^H S + \Gamma^H \Gamma\right)^{-1} S^H a \tag{2.5.3}$$

For $\beta = 0$, minimizing this functional yields the least squares estimator. As β grows, the reconstruction increasingly favors images with minimum norms, limiting the propagation of noise. A more intuitive way to implement Tikhonov regularization is by truncating the SVD of the aliasing matrix (73). Decomposing the aliasing matrix into its singular value decomposition (SVD), we obtain

$$\begin{aligned} S &= U\Sigma V^H \\ a &= Sv = U\Sigma V^H v \end{aligned} \tag{2.5.4}$$

where Σ hold the singular values and U and V hold the orthonormal singular vectors in the range and domain spaces, respective. The unwrapping matrix, $S^\dagger$, can be represented using the SVD matrices as $V\Sigma^{-1}U^H$, permitting the R source voxels to be calculated as

$$\hat{v} = V\Sigma^{-1}U^H a \tag{2.5.5}$$

The condition number is the ratio of the largest to smallest singular values in Σ. For high acceleration factors, the problem becomes "ill-condition" and the condition number grows. The larger the condition number, the more susceptible the image will be to small perturbations, i.e. noise, in the data. To limit the propagation of noise in an ill-condition, highly-accelerated SENSE reconstruction, the smallest singular values in Σ may be truncated, preventing the corresponding entries in Σ^{-1} from growing too large. Regularization is performed by replacing Σ^{-1} with a matrix D,

$$D_{ii} = \frac{\sigma_i}{\sigma_i^2 + \beta^2}, \quad D_{i \neq j} = 0, \tag{2.5.6}$$

with a suitably-chosen regularization parameter β. The effect of this block components of the data lying along the singular vectors corresponding to the truncated singular values. In ill-conditioned problems it may be beneficial to block these components, since the unwrapping matrix would other only pass noise while "missing" object features.

The strength of the regularization is a trade-off between overwhelming noise amplification in the unregularized case and excessive aliasing in the heavily regularized case. As β tends toward infinity, the unwrapping matrix ceases to resolve aliased voxels and the resulting image converges toward the image one would get by direct Fourier transformation of undersampled sampled k-space data. This is a classic trade-off between noise and bias (aliasing); the optimum amount of regularization depends upon the imaging task. In the context of SENSE, the "optimal" β is a function of the acceleration factor and the coil array; the parameter is chosen *ad hoc* by gradually increasing β to suppress noise propagation until the point where aliasing artifacts began to reemerge in the image. In iterative methods such

as the Kaczmarz algorithm, the trade-off between noise and bias can be effected either via a regularization parameters or by varying the number of iterations through the encoding matrix, with each successive iteration tending to reduce bias and increase noise.

The preceding analysis can be adapted to include the effects of the noise covariance matrix, Ψ, by pre-whitening the data and the aliasing matrix (73). This is achieved using the eigen-decomposition of the noise covariance matrix,

$$\begin{aligned} \Psi &= M\Lambda M^H \\ \tilde{S} &= \Gamma^{-1/2} M^H S \\ \tilde{a} &= \Gamma^{-1/2} M^H a \end{aligned} \tag{2.5.7}$$

where the eigenvalues of Ψ are contained in Λ and the eigenvectors form the columns of M. But for purposes of the present simulations, the coil noise covariance information is ignored.

2.5.2 Acceleration comparison

In the high-resolution R=4 reconstructions with 5% noise, O-space images appear qualitatively similar to the SENSE images (Fig. 2.12). At R=8, however, the SENSE reconstruction with eight coils becomes ill-conditioned and either entirely overwhelmed by noise or plagued by severe residual artifacts, depending on the Tikhonov regularization parameter selected. By contrast, the O-space reconstruction at R=8 shows a mild increase in noise and blurring. At R=16, an acceleration factor for which SENSE reconstructions are impossible with only eight coils, O-space images suffer some blurring of small features but retain impressive image quality for images based on only 16 projections. Notably, the features near the center of the brain phantom appear less distorted in the O-Space image than in the radial image. This is surprising because most gains from the addition of the quadratic SEM

were expected to occur closer to the periphery.

PatLoc reconstructions at R=4 show low levels of noise amplification, indicating that the encircling surface coils are complementary to multipolar PatLoc encoding fields; that is to say, the coils are well positioned to localize signal from each acquired point in encoding space back onto the source voxels in the two bijective regions of image space. As the net acceleration factor approaches the number of coils, however, noise amplification grows severe in parts of each bijective region. As expected, resolution degrades near the center of the FOV as the spatial derivative of the encoding fields approaches zero. In this region, the coil profiles do not provide adequate localization to make up for the lost SEM encoding.

O-space and radial reconstructions performed comparably at low acceleration factors, each displaying excellent resolution at R=4. At R=8, the streaking artifacts in radial images become more pronounced, while the O-Space artifacts remain incoherent enough to avoid clear manifestation in the image. The O-Space images do exhibit slightly higher noise levels, in part because of the higher readout bandwidth of the simulated 512-point O-Space echoes. At R=16, the O-Space brain image preserves more detail at the center of the FOV.

The projections were ordered in the encoding matrix according to Fig. 2.9. The time points for a given center placement formed a block of rows. The set of all center placements in turn formed larger blocks corresponding to one RF coil dataset. However, when projections were reordered such that the RF coil data were interleaved on adjacent rows of the matrix, no difference in either the O-Space and radial images was observed in the final converged image.

2.5.3 O-Space noise performance

When the acceleration factor is held constant at R=4 and the noise level is varied (Fig. 2.13), the noise in O-Space reconstructions appears to be more spatially uniform than the

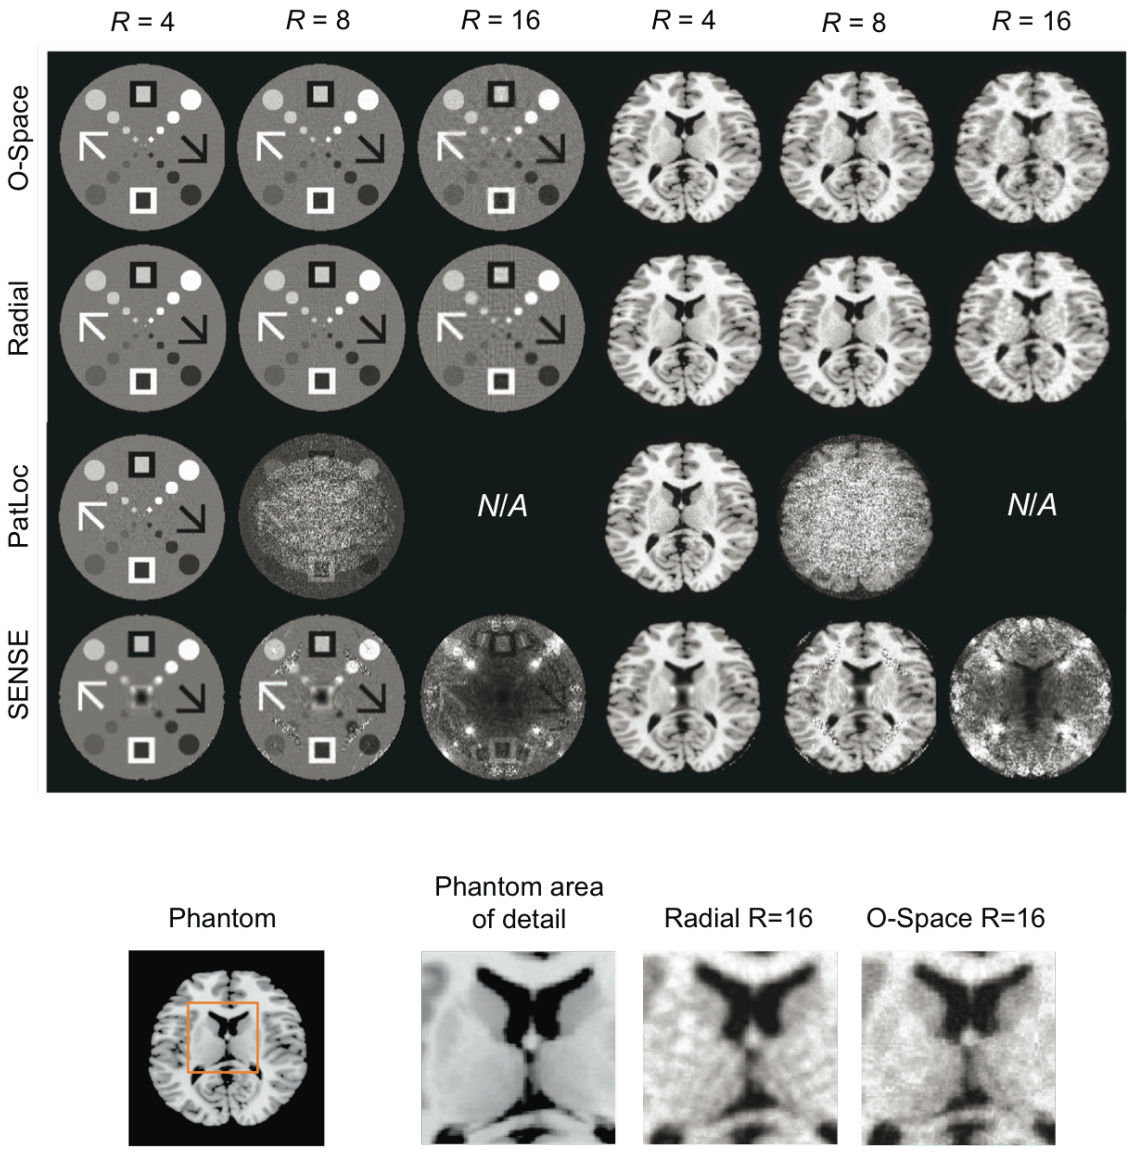

Figure 2.12: Comparison of O-space, radial, Cartesian SENSE, and PatLoc 256×256 reconstructions as a function of acceleration factor for an 8-coil array with moderate noise (5% of mean phantom intensity). As R increases beyond 4, the SENSE and PatLoc reconstructions rapidly deteriorate. By contrast, O-Space performance degrades gradually, displaying only moderate noise amplification and blurring of small features. O-Space reconstructions show promise even when R exceeds the number of coils, a scenario not possible with SENSE. Radial reconstructions perform comparably to O-space reconstructions, but at R=16, the radial images show the expected streaking artifacts (see numerical phantom) and loss of clarity in features (see brain phantom area of detail at bottom). SENSE images are not shown for R=16 because the undersampling factor exceeds the number of RF coils, so a well-conditioned matrix inversion is not possible at each voxel.

noise in SENSE reconstructions, which concentrates in regions of high g-factor. In O-space imaging, each voxel is smeared out along rings according to the spatially varying PSF of each CP, preventing noise (and coherent artifacts) from concentrating in any one region of the image.

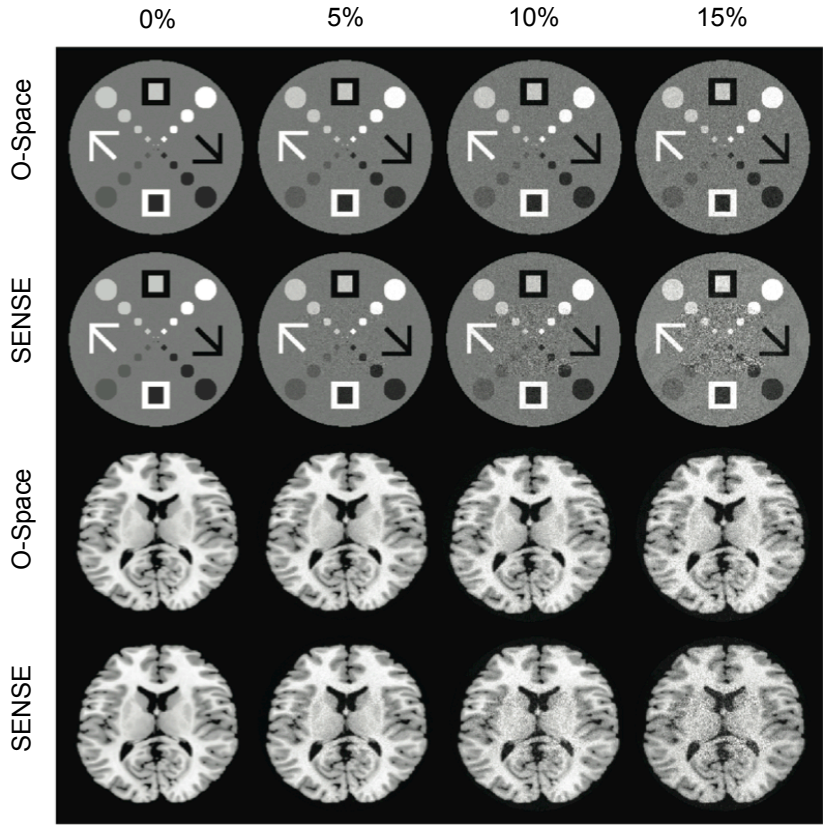

Figure 2.13: Noise levels for eight-coil O-Space and Cartesian SENSE reconstructions at R=4. While SENSE is highly sensitive to noise and requires regularization to prevent noise amplification from dominating the image, O-Space reconstructions degrade gracefully even when the data are highly noisy (a property shared with undersampled radial images). In addition, O-Space noise is distributed more evenly throughout the FOV, while SENSE noise is visibly concentrated in areas of high g-factor.

2.5.4 Increased readout density

As expected, increased ring density within the phantom contributes to significant improvements in resolution. O-Space reconstructions with only 256 rings show substantial blurring and background non-uniformity at R = 16 (Fig. 2.14), suggesting that setting N_s

equal to N does not provide enough rings for highly accelerated acquisitions. But as the number of rings is increased to 512 and 1024, resolution improves noticeably in the R=8 case and substantially in the R=16 case. This stands in contrast to Cartesian acquisitions, where increasing the readout SEM strength provides no reduction in aliasing artifacts along the undersampled direction(s).

In principle, the achievable O-space ring density in the sample is limited only by (a) the diminishing SNR provided by very thin rings and (b) safety regulations on SEM switching rates. It should be emphasized that increasing the number of samples in the readout and sampling faster (higher bandwidth) does not change the overall imaging time, and thus the imaging time along a row in Fig. 2.14 is constant. In summary, increasing the readout SEM strength (within the bounds of the noise floor) benefits O-Space imaging resolution in two dimensions while benefiting Cartesian imaging resolution in only one dimension. This is a significant advantage that is only possible when SEMs vary in more than one dimension.

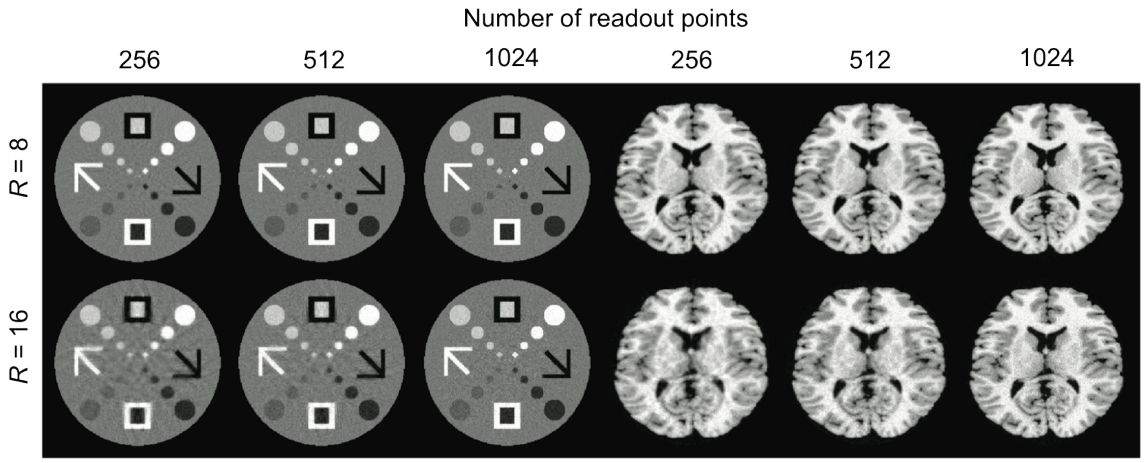

Figure 2.14: Time-equivalent R=8 and R=16 O-space 256×256 phantom reconstructions based on 256 rings, 512 rings, and 1024 rings per echo, along with corresponding brain simulations, demonstrate the resolution improvements created by dense ring spacing within the object. Uncorrelated noise with standard deviation equal to 5% of the mean phantom intensity was added during echo simulation. The noise is then scaled by $\sqrt{2}$ and $\sqrt{4}$ in the 256- and 512-point reconstructions, respectively, to reflect the increased sampling bandwidth.

2.5.5 Applicability of O-Space simulations

The simulations do not include the effects of calibration errors in SEM field strength, timing, eddy currents, or concomitant fields, so in a sense they may be considered to represent the best-case performance that is achievable. Also, the noise scaling was a best estimate of real signal behavior, and in respect it may have been scaled too low.

2.5.6 Properties of the encoding matrix

Approximations to the O-Space PSF were obtained from the matrix product $E^H E$, where E^H denotes the conjugate transpose of the encoding matrix[1]. Each row of EHE is the vectorized PSF for a particular voxel in the FOV. To isolate the effects of the SEM encoding alone, a single coil with uniform sensitivity was used. Representative PSFs for the 16-CP encoding scheme are shown in Fig. 2.14; for each CP, the source voxel blurs along a ring that has the CP at its origin. The source is localized to the point where all 16 rings overlap, with additional localization provided when a surface coil array is used.

To illustrate the radial encoding provided by the Z^2 SEM alone, the encoding matrix for the Z^2 SEM shape was Fourier transformed along the temporal dimension (row-wise), yielding the ringlike shapes that correspond to each point in a Z^2 projection. Magnitude plots of a horizontal profile through the center of the FOV are shown in Fig. 2.15 (c,d) illustrate the radial variation in resolution and the effects of echo truncation. Since the spatial derivative of the Z^2 field shape is zero at the center of the FOV, nearly all SEM encoding in this region comes from the application of the X and Y SEMs to shift the rings off center. Truncation of an echo with a window of duration results in convolution of the rings by $sinc(2u/t)$ in u-space, corresponding to $sinc(2r^2/t)$ in the image domain. This leads

[1] More accurate PSFs can be obtained from FE, but it is intractable in practice to explicitly calculate F for O-Space imaging

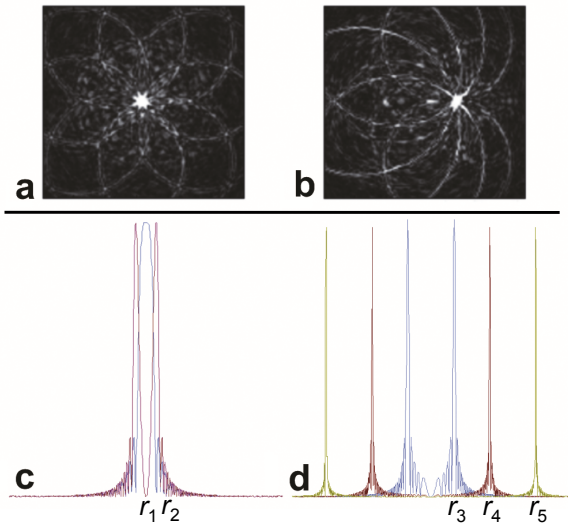

Figure 2.15: Spatially-varying O-Space PSF of the 16-CP scheme used in simulations for a point source of unit intensity located at voxel index $(x, y) = (0, 0)$ (a) and $(x, y) = (75, 0)$ (b) within a 256×256 FOV. The scale is compressed 10-fold to bring out low-level ring features. For projections formed using the Z^2 SEM alone (c,d), horizontal profiles through the ring projections illustrate the poor radial resolution of the two centermost rings at r_1 and r_2 (c) in contrast to the fine resolution of rings at the larger radii r_3, r_4, and r_5 (d).

to widening of the side lobes as r approaches zero.

The LSQR reconstructions described above (for CP optimization) have the advantage of being fast and non-memory intensive, but they do not work on actual data. The simplified encoding matrix assumed that each voxel has a real-valued PSF consisting of discrete rings with no sidelobes. True O-Space PSFs are complex-valued and include Gibbs ringing due to the rectangular window on the readout data. As such, the real-valued approximate encoding matrix is not able to account for the full phase evolution of spins during an O-Space acquisition, leading to phase cancelation during reconstruction.

Fortunately, there are ways to combine the integrity of the full encoding matrix with the speed of using a sparse matrix for reconstruction. The shape of each complex-valued "ring" can be approximated by Fourier transforming each $N_s \times N$ block of the full encoding matrix along the time (row) dimension. These image-domain matrices can be assembled into an equation for the unknown image using the projection data (i.e., FFT of the echoes). This version of the encoding matrix can be sparsified by truncating all values falling below a certain threshold, leaving only the PSF in the vicinity of each ring. This approach was explored in (74) to quicken reconstruction of data encoded using nonlinear SEMs. Using a sparse matrix that fits in memory, LSQR would present an attractive alternative to Kacz-

marz for reconstructing the image, but we reserve this comparison for future work.

Transforming the encoding matrix into a sparse basis could also speed convergence of the Kaczmarz algorithm. When the matrix elements are populated directly in the frequency domain, each element has a magnitude of unity. This means all rows of the encoding matrix have the same L_2-norm. Weighting the rows by RF coil profiles changes the L_2-norm for the rows corresponding to each coil profile, but this only increases the number of distinct L_2-norm values from 1 to n_c. However, if the encoding matrix is then transformed into a sparse basis, row L_2-norms will likely have a wide spread of values. This would permit rows to be selected randomly for each step of the Kaczmarz algorithm, with the probability of selection proportional to the L_2-norm. This approach has been shown to significantly speed convergence (71) as compared with standard Kaczmarz (in which rows are chosen sequentially from first to last), simple randomized Kaczmarz (in which rows are chosen completely at random), and the conjugate gradients algorithm. However, in order to achieve faster convergence with randomized Kaczmarz, there must be a wide spread in the value of the L_2-norm among the rows, as occurs in a sparse matrix[1].

2.6 Off-resonance effects on O-Space imaging

Regions of B_0 inhomogeneity have long been a concern in MR imaging. Perturbations of the background B_0 field arise at the interfaces of tissues having different magnetic susceptibilities, such as bone, fluid, and cortex. Interfaces of the sort are typically accompanied by substantial B_0 inhomogeneity in the human head, particularly near the sinuses and ear canals. This field inhomogeneity can cause distortions in conventional imaging sequences, particularly those with long readouts such as echo planar imaging, which suffers spatial

[1] As noted earlier, improved convergence with random row selection has only been shown for the case of consistent systems of equations and for equations in which the data are contaminated by random noise (75). The convergence properties have not been investigated for equations, such as those involved in O-Space imaging, in which the matrix holding the coefficients is contaminated by systematic error.

translation of voxels along the phase-encode direction. This distortion can either be addressed during image reconstruction for Cartesian trajectories (23), provided that accurate field maps are available. Inhomogeneity can be preempted through highly-accurate B_0 shimming (76). To date, however, no combination of active and passive shimming has succeeded in eliminating all residual inhomogeneity. In theory, higher readout bandwidths can be used to limit the effect of off-resonance spin phase evolution, preserving image quality. But in practice, the readout BW is limited by signal-to-noise considerations, gradient switching limits, and other sequence-dependent factors.

Off-resonance is of additional concern because it is common for a Z^2 SEM coil to produce a Z_0 uniform field cross-term. This global term arises due to the similarity of the conductor windings used to generate the Z^2 and Z_0 fields. It is hard for SEM coil designers to achieve fields that are both strong and pure. If known, this Z_0 component can be incorporated into the encoding matrix during image reconstruction. But it is important to explore the effects of this applied off-resonance field on image quality, in the event that the user has imperfect knowledge of the Z_0 amplitude.

Regional and global off-resonance both pose a significant obstacles to O-Space imaging. Because of the way that frequency isocontours from different CPs overlap to provide signal localization (Fig. 2.16), the method is potentially sensitive to errors in effective field strength. Therefore it is important to investigate via simulation whether O-Space imaging could plausibly be performed in areas like the human head, where sinus and ear cavities introduce substantial local inhomogeneities. The complex, spatially-varying point

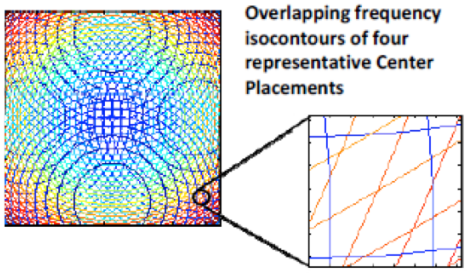

Figure 2.16: Illustration of overlapping frequency isocontours from different center placements. Misalignment of isocontours causes phase cancellation.

spread functions that occur in O-Space imaging invite inquiry into the sensitivity of the method - and non-linear projection imaging more broadly - to local field offsets. In the presence of spatially-varying B_0 inhomogeneity, the signal equation 2.3.4 becomes

$$\begin{aligned} s_{l,q} &= \int\int m(x,y)\, C_q(x,y)\, e^{-i2\pi\left[\gamma G_{Z^2}\frac{1}{2}\left((x-x_l)^2+(y-y_l)^2\right)\right]t} e^{\Delta B_0(x,y)(t+TE-T_s/2)} dxdy \\ &= E_{t,l,q} m. \end{aligned} \tag{2.6.1}$$

where $\gamma = \frac{\gamma}{2\pi}$ and T_s is the readout duration. Field errors cause the matrix equation to become inconsistent. When inconsistency arises from random errors in s due to noise, regularization is an effective way to constrain noise propagation in the reconstruction. For systematic errors in E such as those caused by field errors, regularization has less utility, but can still be used to find a constrained estimate of m when no unique solution to the system of equations exists. To test the ability of regularization to minimize the adverse effects of systematic error from ΔB_0, we use a modified Kaczmarz implementation that incorporates Tikhonov regularization (77). To perform regularization, the reconstruction seeks to minimize the new functional

$$r \parallel Em - s \parallel^2 + \parallel \hat{m} \parallel^2 . \tag{2.6.2}$$

The first term enforces data consistency. The second term favors those estimators for $\hat{m}$ that have the smallest norms from among all possible solutions that satisfy data consistency. This second term therefore regularizes the reconstruction and favors the *minimum norm least squares estimator*, with the strength of the regularization controlled by the parameter r. To find a regularized least-squares solution, $\begin{pmatrix} \hat{z} \\ \hat{m} \end{pmatrix}$, consider the following augmented system of equations, which is always consistent:

$$\begin{pmatrix} I & r^{1/2}E \end{pmatrix} \begin{pmatrix} z \\ x \end{pmatrix} = r^{1/2}s \tag{2.6.3}$$

where E is the $n_K n_c \times N^2$ encoding matrix, I is the $n_K n_c \times n_K n_c$ identity matrix, and z is a vector of length $n_K n_c$. A regularized image estimate can then be obtained by setting $x^0 = r^{1/2}E^H z^0$, with z^0 arbitrary, and using the following modified Kaczmarz algorithm:

$$\begin{aligned} z^{(n+1)} = z^{(n)} + c^{(n)} e_{i_n}, \quad \hat{m}^{(n+1)} = \hat{m}^{(n)} + r^{1/2} c^{(n)} E_{i_n} \\ c^{(n)} = \lambda_n \frac{r^{1/2} s_{i_n} - z_{i_n}^{(n)} - r^{1/2} \left\langle E_{i_n}, \hat{m}^{(n)} \right\rangle}{1 + r \parallel E_{i_n} \parallel^2} \end{aligned} \tag{2.6.4}$$

where e_{i_n} is defined as a vector of length $n_K n_c$ comprised of all zeroes except for the n^{th} entry. At each step in the algorithm, only a single entry of the vector z is updated. For the most straightforward top-to-bottom cycling through the encoding matrix, i_n simply equals n, though in general they may be different (i.e., rows of the encoding matrix can be selected randomly for each step). The scalar parameter r controls the strength of the regularization.

2.6.1 Simulations of off-resonance effects

A 128×128 numerical phantom is used to model potential feature distortion and noise amplification caused by B_0 inhomogeneity. The field inhomogeneity ΔB_0 of the sinus and ear cavities is modeled as a Gaussian shape[1] rolling off radially from the point of peak field offset (Fig. 2.17),

$$\Delta B_0(\bar{r}) \approx \Delta B_{peak} \exp\left(-\frac{\bar{r} - \bar{r}_0}{FOV/20} \right)^2. \tag{2.6.5}$$

for several positions of r_0 in the FOV. The peak offset ΔB_{peak} is varied between 0 and 500 Hz in simulations. To assess artifact levels and distortion, the mean squared error is

[1] Gaussian/quadratic shapes have been used previously to roughly model B_0 inhomogeneities (78).

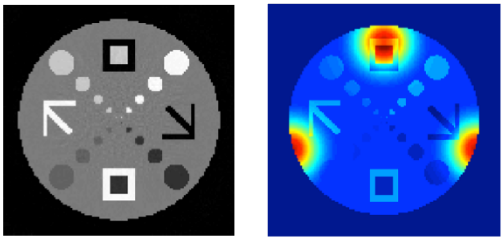

Figure 2.17: A simple model of cavity-induced B_0 inhomogeneity superimposed on a numerical phantom for purposes of simulation. The peak level of inhomogeneity, ΔB_0, is varied in simulations to observe its effect on image quality.

computed between the reconstructed images and a reference image with no B_0 inhomogeneity (Fig. 2.19). Acceleration factors of $R = 4, 8, 16$ are simulated with and without regularization for an array of 8 surface coils (R factors correspond to 32, 16, and 8 echoes, respectively). To speed up simulations, an image grid of 128×128 is used instead of the 256×256 grid used elsewhere in this thesis. A noise floor with standard deviation equal to 5% of the mean phantom intensity is added.

For the images shown in Fig. 2.18 and Fig. 2.19, the acquisition is simulated with a short TE (3 ms) and the readout BW is set to 150 KHz, a typical value in clinical protocols. In Fig. 2.18, artifacts manifest themselves as noise-like graininess, with rippling artifacts in the region of the greatest field deviation (Fig. 2.18). But geometric distortion or warping are not observed near the simulated cavities, as they are in echo planar imaging. Unregularized images show noticeable artifact levels for $\Delta B_{peak} > 100 Hz$.

Simluation results in Fig. 2.19 show that the MSE grows with ΔB_{peak} for all acceleration factors (Fig. 2.19). Visual inspection reveals that artifacts grow more severe with increasing R. Regularized reconstructions, by contrast, degrade gracefully with ΔB_{peak} for all acceleration factors, showing only modest MSE and artifact increases up to $\Delta B_{peak} = 300\, Hz$. Regularization also suppresses noise, as expected, operating as a low-pass filter. As expected, however, regularization comes at the cost of image resolution, particularly at the center of the FOV where the intrinsic O-Space resolution is lowest. In this respect, the MSE is not an adequate metric for assessing the combined effects of noise amplification,

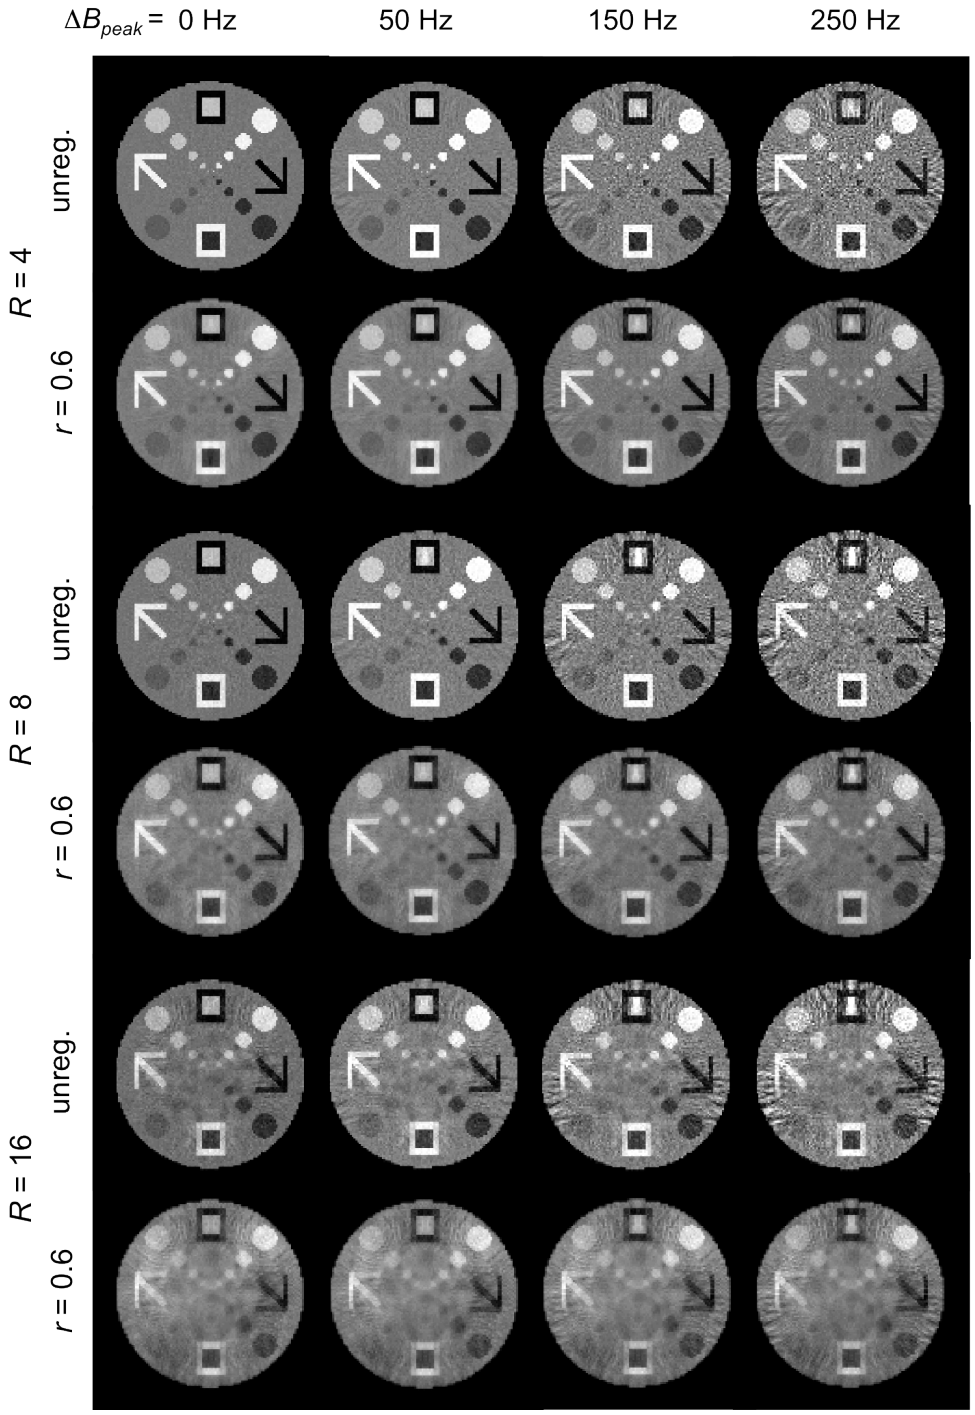

Figure 2.18: Image quality degrades steeply with increasing local B_0 inhomogeneity. While regularization limits the propagation of artifacts, it does so at the expense of blurring the image, obscuring small features. Readout BW is 150 KHz.

artifact propagation, and blurring. Although regularization reduces the MSE, it does not necessarily improve the image quality. The regularization parameter r controls the trade-off between artifact suppression and blurring. For the images shown here, $r = 0.6$, corresponding to regularization that strongly minimizes the norm of the image estimate, as described in Eq. (2.6.2).

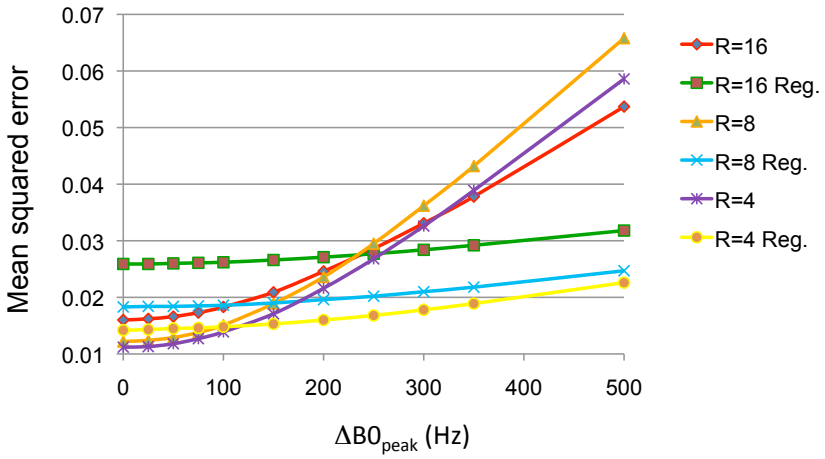

Figure 2.19: Mean squared error for reconstructed images relative to a reference phantom with peak amplitude of unity. While regularization slows the growth in MSE with increasing ΔB_{peak}, it comes at the cost of severe blurring, as shown in Fig. 2.18. Readout BW is 150 KHz.

These simulations suggest that local inhomogeneity could be deleterious to practical O-Space imaging in the human head. However, if accurate field maps are available, the off-resonance effects can be incorporated into image reconstruction. Compensation for off-resonance effects has been shown to reduce voxel-displacement along the phase encode direction in echo planar imaging with Cartesian SEMs (23). Extending this approach into the domain of nonlinear SEMs, field maps can be incorporated directly into the O-Space encoding matrix. As shown in Fig. 2.20, images reconstructed with accurate knowledge

of $\Delta B_0(x, y)$ are almost identical to those simulated in the absence of off-resonance fields. Even if only crude field maps are available, it is beneficial to include the maps in the reconstruction. Even when there is a 20% error in the field map estimate of B_{peak} as compared with the true B_{peak} of the off-resonance field, the images still show a large improvement over those reconstructed without the field map correction. Methods such as FASTMAP can be used to quickly acquire field maps during the pre-scan in a clinical protocol. With this knowledge of the off-resonance fields, it does not appear that local inhomogeneity poses an insurmountable obstacle to the use of O-Space imaging in clinical settings.

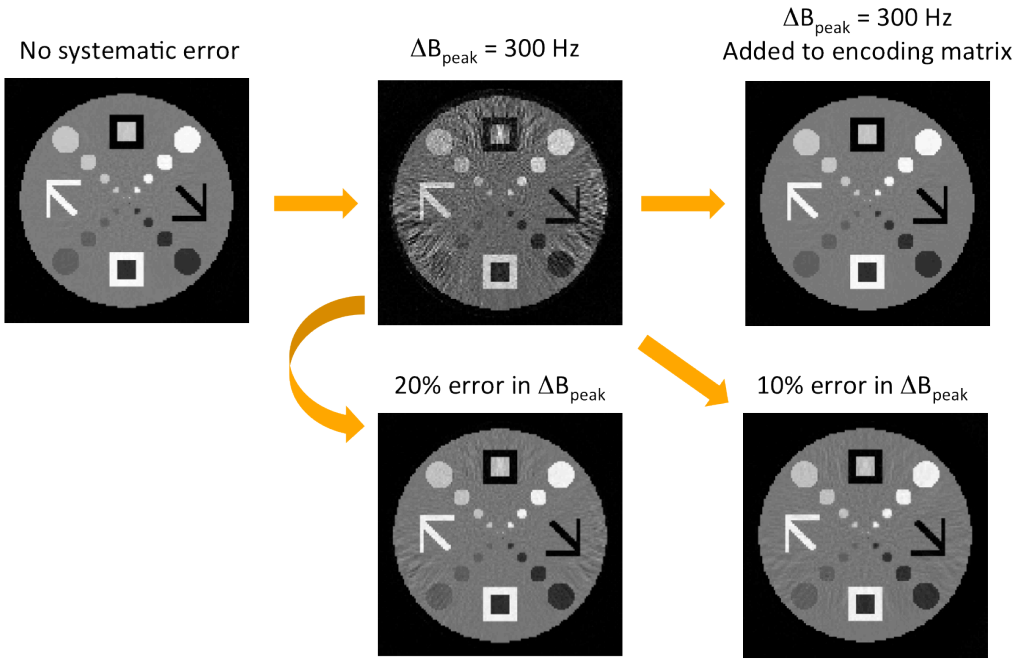

Figure 2.20: Regional inhomogeneities have effects across the reconstructed image (readout BW is 150 KHz). However, if field maps are available, they may be incorporated into the encoding matrix, recovering virtually the same image quality as for a perfectly-shimmed object. Even when there is error in estimating ΔB_{peak}, large improvements are observed over uncorrected images.

3

Experimental proof of concept

3.1 Overview

This chapter begins by describing the hardware and methods used to implement O-Space imaging with a Z^2 insert coil on a Siemens 3T human scanner. Special attention is paid to methods for accurately measuring the magnetic field of the insert coil, which proved to be the principal obstacle for our investigations. Experimental images are then presented. The convergence properties of the Kaczmarz method are explored using experimental data. Finally, some common artifacts and ways to eliminate them are discussed. It is my hope that the chapter can serve as a "recipe" for other investigators interested in performing nonlinear projection imaging in their own laboratories.

3.2 Z^2 SEM insert coil

An actively-shielded, liquid-cooled Z2 insert coil with a 12-cm bore was designed and fabricated by Resonance Research, Inc. (Billerica, MA). Its specifications are displayed in Table 3.1 and a cross-section of the geometry is shown in Fig. 3.2. The insert consists of a cylindrical coil that holds the current-carrying windings along with a support structure (not shown in the schematic). The support structure conforms mechanically with the patient

table of a Siemens Trio 3T scanner (Erlangen, Germany). The zero frequency isocontour of the insert is aligned with the origin of the scanner coordinate system. Teeth cut into the side support struts of the coil fit snugly into grooves cut in the patient table, fixing the position of the insert and minimizing mechanical translation. Because the insert must be removed at the end of each imaging session, repeatable positioning from day to day is extremely important. The insert is further secured in place using a screw-actuated arm with a padded hand that braces against the top of the scanner bore.

The insert is driven by a Techron Model 8606 amplifier (Elkhart, IN) capable of supplying up to 120 amps of current, corresponding to a field strength of 13.6 KHz/cm^2. The amplifier is controlled by a Dynamic Shim Updating system that was originally designed to update shims in between TRs on a slice-by-slice basis (79, 80). The DSU controller updates previously stored field strengths using TTL pulse triggers called during an otherwise conventional radial pulse sequence, producing O-Space center placements along a circle of the desired radius. Temperature probes are placed throughout the insert coil to permit fast shutdown of the amplifiers in the event of overheating from exceeding the insert coil's duty cycle, a short circuit, or other malfunction.

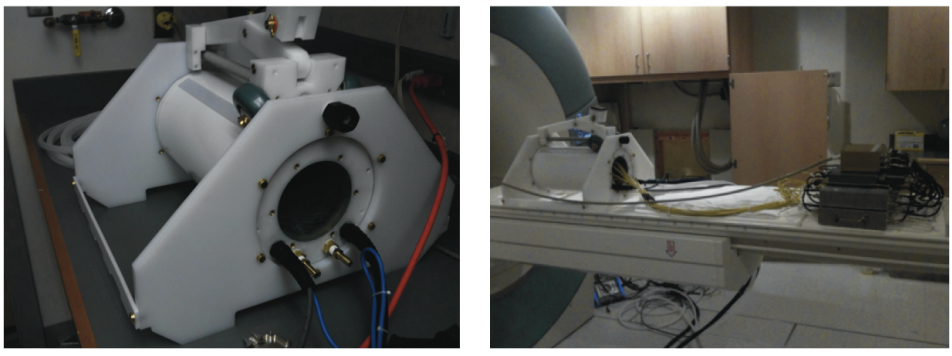

Figure 3.1: The actively-shielded, liquid-cool Z^2 insert coil being tested on the lab bench (left) and mounted on the patient table with accompanying RF hardware.

Specification	Unit	Value
Inner Diameter	mm	120.7
Outer Diameter	mm	199 $\pm$1
Temperature	$^\circ$ C	30.0
Current	Amps	151.1
Resistance	mΩ	69
Power	W	1580
Peak Z^2 Strength	Gauss/cm^2	4.26
Peak Z^2 Strength	Hz/cm^2	18,135
Inductance	mH	80
Rise Time	ms	100
Inductive Voltage	V	121

Table 3.1: Design specifications for the Z2 insert coil (courtesy of Resonance Research, Inc.). Note that the peak Z^2 design strength is only half as strong in the axial plane, where the images shown in this thesis are acquired.

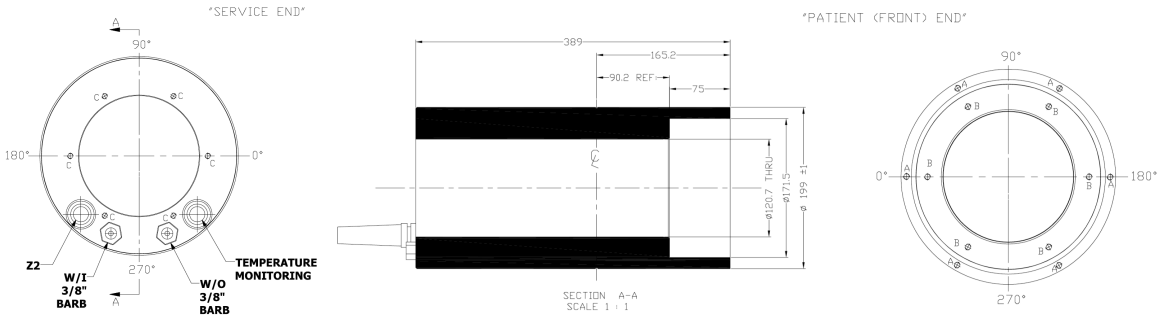

Figure 3.2: Cross section of the Z^2 insert coil geometry. Coolant is supplied through hoses entering on the foot-side of the housing. Note the bolt circles used to mount the collar of the RF coil array to the Z2 housing (schematic courtesy of Resonance Research, Inc.).

3.2.1 Dynamic shim updating and amplifier

The Z^2 imaging system takes advantage of hardware previously developed at Yale (76) to implement dynamic shimming for applications such as improved multi-voxel spectroscopy (81), spectroscopic imaging, and echo-planar imaging. The hardware was repurposed to drive the Z^2 and Z_0 windings of the Resonance Research Z^2 insert coil. A rack of electronics was assembled for this purpose (Fig. 3.3), including the following elements:

- Temperature monitoring circuit to monitor the temperature of the SEM coil and disable the amplifiers in the event of overheating;
- A coolant pump to keep water circulating in the SEM coil during use;
- A box holding digital circuit boards that interface between the the monitoring equipment, the DSI Loader, and the amplifiers;
- A set of three pre-emphasis circuits for both the Z^2 and the Z_0 eddy currents produced by the Z^2 coil;
- An amplifier to drive the Z_0 coil;
- A 120 amp Techron amplifier to drive the Z^2 coil;
- A switch box to turn the cooling and DSU systems on and off.

Additionally, a laptop is set up in the scanner control room to run the DSI Loader, a custom-written program that allows users to update dynamic shim settings (Terry Nixon, Yale Univ. MRRC). DSU settings are entered as a percentage of the amplifier's peak output. The DSI Loader can hold up to 126 independent values for both the Z^2 and the Z_0 term. Because a pair of pulses requires four entries - negative lobe, zero, positive lobe, zero - the system permits O-Space imaging with up to 31 independent echoes[1]. The Siemens scanner provides two 10 microsecond TTL pulses for triggering user-supplied hardware. One

[1]For the experimental images shown in this thesis, the same Z^2 amplitude is played for each readout, requiring use of only 4 DSI registers. However, in future versions of the DSU, the number of independent shim settings will be greatly increased, allowing greater flexibility for imaging experiments.

of these pulses is used to advance to the next DSU setting in the DSI register. The second TTL pulse is played at the end of a TR (or sequence of TRs) to reset the DSI Loader to the initial value.

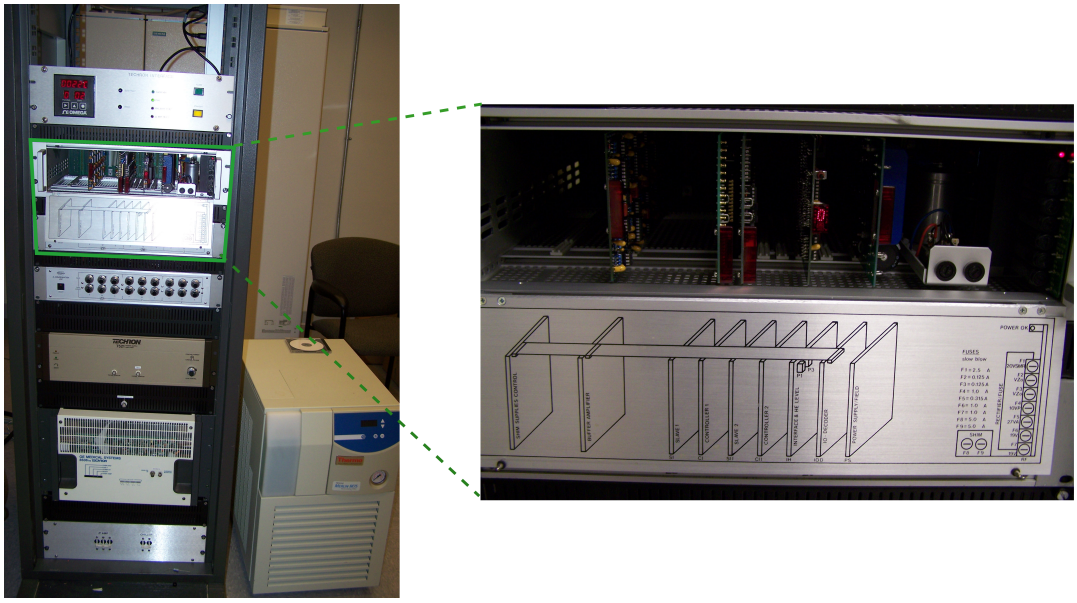

Figure 3.3: DSU rack with inset image of custom-engineered dynamic shimming controller boards. The rack includes, from bottom to top, a switch box, Techron amplifier for the Z^2 field, a smaller amplifier for the Z_0 field, pre-emphasis controls for eddy current compensation, dynamic-shimming control circuits (also inset), and display/gauge box. The water coolant pump is at lower right.

The pre-emphasis controls for each field component include six knobs for shaping the exponential waveform corrections: three for the amplitude and three for the time constant. Using a superposition of three exponential terms, most eddy current-induced distortions in an SEM pulse waveform can be compensated, restoring the desired trapezoidal waveform.

3.3 Custom RF electronics

To perform accelerated parallel imaging, localized RF coil sensitivities are needed (42). The more coil channels used, the more independent measurements are available for image reconstruction from undersampled data, minimizing noise amplification and/or resolution

loss. For O-Space imaging, we hypothesize that the ideal coil array uses localized coils encircling the object. Arrays of small, localized coils provide SNR gains over uniform coils (48), with the improvements increasingly confined to the FOV periphery as the elements grow more numerous (39). In choosing the number of elements, one must consider the available space as well as the mutual inductive coupling (4, 82) that will occur between coils laid out in a given pattern. The challenge for the present work was to fit as many encircling elements as possible into a 12 cm diameter while maximize the remaining volume for imaging. Since the Siemens scanner's built-in RF body coil can not penetrate the shielding and windings of the Z^2 insert coil, the chosen RF coil(s) must be capable of both transmit and receive mode.

Two RF coil designs were considered for this application. The first approach was to use a conventional birdcage coil to transmit a homogeneous field along with a separate multi-channel receive array of loop elements (40). The alternative is to use a single set of elements for transmission and reception (8, 83). After careful consideration of the physical and electrical requirements, the latter configuration was chosen. Using the same elements for transmission and reception obviates the need for decoupling circuits in the separate transmit and receive coils to disable them during RF reception and transmission, respectively. Also, a single coil array could be easily fit within a 1 cm wall thickness, while arranging two independent coils in such a limited annulus would have been challenging. While the transmit-receive design requires extra electronics (TR modules and a power splitter) outside the bore of the Z^2 coil, it conserves precious space within the bore, maximizing the remaining diameter available for imaging.

For the coil elements, microstrip transmission lines were selected for their (a) uniform B_1 distribution along the z-direction and (b) relatively good decoupling properties on resonance. While circular arrays do not enjoy the full intrinsic decoupling of planar arrays (7), they still exhibit adequate decoupling when elements are connected by capacitors to cancel mutual inductance (8). For microstrip arrays, there is a trade-off between penetra-

tion depth on the one hand and decoupling performance on the other. The parameters involved in this trade-off are the substrate thickness, element width, and inter-element spacing (84). As a general rule, to prevent excessive coupling, inter-element spacing (between the midpoints of each line) should be greater than twice the substrate thickness (85)[1]. For the teflon low-loss dielectric substrate used in this study (permittivity ≈ 2.1), the substrate thickness and element width were both chosen to be 5 mm. For the available circumference, eight elements were chosen as the optimal array size.

Capacitive shortening is required due to the fact that a half-wave resonator over teflon at 123.2 MHz is over 90 cm while useful region of the Z^2 insert coil is only 12 cm. Shunt capacitors are used at either end of the microstrip lines for capacitive shortening, achieving a half-wave resonance with a line only 12 cm in length[2]. The resonant frequency is tuned using a non-magnetic trimmer capacitor (Voltronics, Denville, NJ), C_t in Fig. 3.4, actuated via an acrylic rod with a thin copper tooth that mates with the capacitor screw head. The impedance match between the 50 ohm coaxial cable feeding an element and the microstrip transmission line is similarly tuned with a series non-magnetic trimmer capacitor, C_m. For optimal performance, the coil needs to be tuned every time the sample is changed, due to an unequal ratio of capacitive to inductive coupling between coil elements and the sample[3]. A single decoupling capacitor is used between the inputs of each element to cancel mutual inductance. Eddy currents are measured on a network analyzer and verified to be less than -15 dB for adjacent elements, and less than -19 dB for next-nearest neighbors. Due to the presence of TR modules in the design, pre-amplifier decoupling (4) was not used, but coil coupling levels were acceptable even without use of this technique.

[1] Equations exist for modeling coupling between microstrip lines on a shared planar surface (7), but not for the more complicated case of microstrip lines distributed around a cylinder. In this case, the most accurate way to model coil element coupling is to use 3-D electromagnetic modeling software.

[2] Care must be take to choose the capacitance on each end properly so that the current flowing in the microstrip is at or near the peak of the half-sine wave distribution for the resonant half-wave mode. If the currents flowing in the microstrip are not at the peak of the sine curve, then B_1 will be reduced and the coil SNR will suffer.

[3] In practice, coil performance was found to be adequate using a single tuning/matching setting for a variety of imaging phantoms, saving users considerable time and effort during experiments.

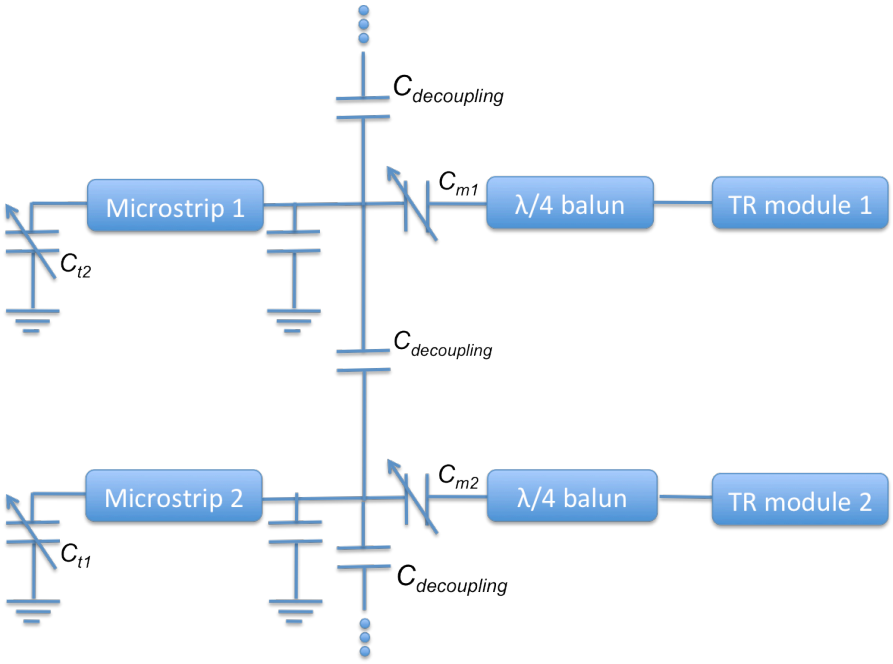

Figure 3.4: Circuit diagram for two of the eight capacitively-shortened microstrip coil elements. Capacitor C_t is used for tuning the resonant frequency while . A single capacitor is used to cancel the mutual inductance between neighboring elements.

Eddy currents are suppressed in the shield through the use of a thin, segmented copper layer (5 micron) comparable to the skin depth at 123.2 MHz. Neighboring elements are decoupled using single capacitors bridging their inputs. A circularly-polarized birdcage mode is approximated by introducing a 2π progressive phase using lumped element phase shifters in each transmit path at multiples of $45\,^{\circ}$. The design is optimized for compactness at the expense of B_1 homogeneity near the periphery, but as the experiment is a proof of concept for quadratic field encoding and not for RF coil design, this is considered an acceptable compromise.

"Bazooka" baluns are used to suppress common-mode currents on the coaxial shield. The desired transverse electromagnetic mode in coaxial cable consists of currents flowing in one direction on the center wire and the opposite direction on the inner surface of the coaxial shield. However, common-mode currents can also be excited in which current flows in the same direction on the center wire and on the outer surface of the shield. Common-mode

currents are often excited at transitions between a "balanced" circuit (such as the coaxial cable) to an "unbalanced" circuit like a receive coil element. The common-mode currents produce fields around the coaxial cable, leading to losses and reductions in SNR. The fields also interact with nearby objects, disrupting the impedance match of the coil circuit. [1]

3.3.1 Transmit-receive modules

In the transmit path, an 8-to-1 power divider (Taylor Microwave, Clifton, NJ) is used to provide transmit power to each element. A circuit known as a transmit-receive module, or TR module, is required to isolate the transmit and receive paths in each channel. TR modules were built in-house according to the schematic in Fig. 3.6, achieving transmit-to-receive path isolation of better than -54 dB in all eight channels. The circuit takes advantage of positive bias provided by the Siemens scanner that is intended to bias PIN diodes in detuning circuits. Each Siemens RF channel is equipped with a 100 mA current source during forward bias and a -30 V reverse bias for switching diodes off. The available 100 mA are divided among a high-power, transmit-path PIN diode ($D1$) and two receive-path PIN diodes ($D2$ and $D3$ [not shown]). RF is cut off from the DC supply and ground wires using large inductors, which function as open circuits at very high frequencies.

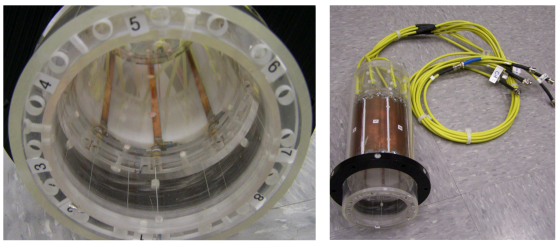

Figure 3.5: Eight-channel microstrip array for use inside the 12 cm Z^2 insert coil. Copper microstrip elements are plated onto 5 mm teflon layers shielded on the outside by a 5 micron copper layer. The thinness of the shield, comparable to the skin depth at 123.2 MHz, helped suppress eddy currents. Hole in the black nylon collar allow the array to the bolted into place using bolt circle "B" shown in Fig. 3.2

The two arms of the receive path consist of $\lambda_l/4$ phase delays followed immediately

[1] When the scanner body coil is used for RF transmission, large common-mode currents can be induced in RF receive cables, threatening to burn the patient.

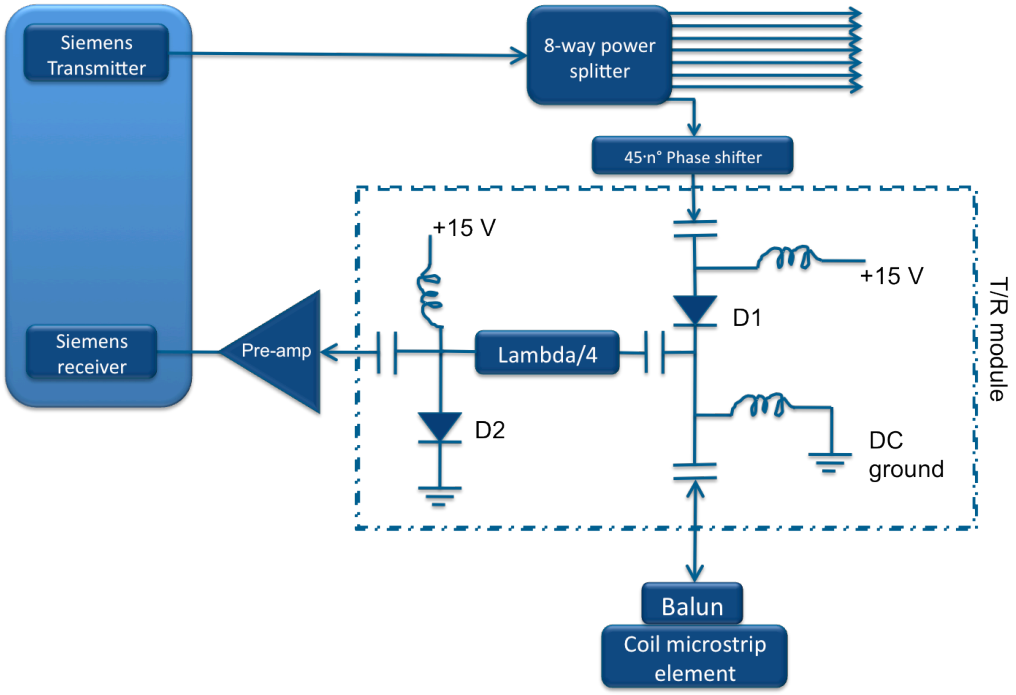

Figure 3.6: RF front end design used in combination with the 8-channel array. Transmit-receive modules consisted of a transmit path and a receive path with two $\lambda_l/4$ segments ending in shunt PIN diodes (only one segment shown here). Isolation of better than -54 dB was achieved between the transmit and receive paths.

by a PIN diode, where λ_l is the RF wavelength in the microstrip transmission line. The first phase delay is realized using a quarter-wave segment of coaxial cable. The second delay takes the form of an LCL pi-network made using discrete components. The pi-network is designed for a 50 ohm impedance match and 90° phase delay. During forward bias (receive mode), each of the two diodes forms a short circuit which is impedance transformed through the quarter-wave delay to an open circuit. The open circuit prevents currents flowing down the transmit path from leaking into the receive path and burning out the sensitive pre-amplifiers at the other end of the line. Conversely, during RF reception, the diodes are reversed bias, permitting the receive signal to reach the pre-amp but blocking it from traveling up the transmit path to the transmitter.

3.3.2 Pre-amplifiers

Low-noise amplifiers (pre-amps) are a critical part of the RF receiver front-end. The pre-amp scales both signal and noise by an amount specified by the gain of the amplifier, 27 dB for Siemens MRI pre-amps. The pre-amp also injects additional noise, degrading the SNR. The change in the ratio of signal-to-noise is described by the "noise figure". Pre-amplifiers are designed to achieve the desired gain with minimum noise figure (typically 0.5 dB for Siemens pre-amps).

After the pre-amp in the receive chain, additional sources of noise have little impact on the final system SNR, since new noise sources are typically very small compared to the amplified signal. For this reason, pre-amps should be placed as close to the receive coil/antenna as possible. Long cable runs invariably introduce some noise due to dissipative losses in the cable, hurting the SNR. For the configuration used in this study, the pre-amp/TR-module assembly was placed 1.5 meters from RF coil, as shown in 3.1. While this placement is not optimized for SNR, it was motivated by (a.) the lack of space for pre-amp modules in the

small confines of the Z^2 insert bore and (b.) the desire to minimize eddy currents induces in the pre-amp/TR-module housing by the switching Z^2 and linear SEM fields.

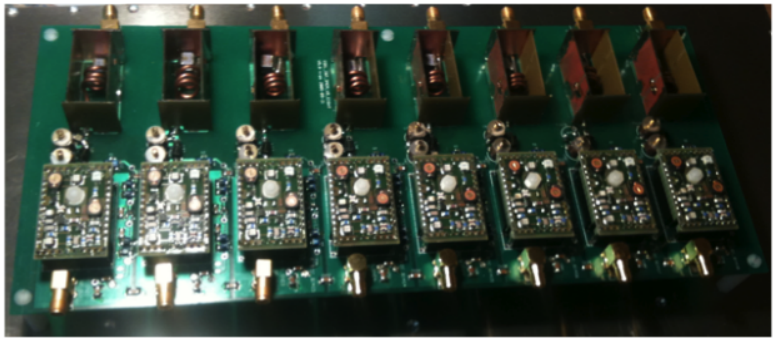

Figure 3.7: A bank of eight pre-amp modules mounted on a dedicated pre-amp board. Parallel LC circuit RF traps are visible near the top of the board. Input impedance tuning capacitors are visible between the RF traps and pre-amp modules.

Siemens pre-amplifiers and matching boards (Fig. 3.7) were provided to us by Larry Wald of the Martinos Center, Massachusetts General Hospital. The boards include RF traps at the inputs to suppress common mode currents. In the traps, the RF is routed through semi-rigid coaxial cable that is wound through several turns to form an inductor. A capacitor is soldered in parallel with the inductor windings to form a parallel LC circuit. The L and C values are chosen to form an open circuit on resonance, blocking RF energy from propagating on the outer surface of the cable shield. The pre-amplifier gain was measured using a network analyzer and verified to be at least 25 dB. Pre-amplifier noise performance was then measured using a noise figure meter along with a calibrated noise source on the input. The input impedance was adjusted to minimize the noise figure, as measured on a noise figure meter. Coil elements were tuned using an Agilent network analyzer (Palo Alto, CA).

Capacitive or inductive decoupling is generally used to cancel mutual inductance between neighboring coil elements, but is insufficient to decouple other pairs of elements. For this purpose, a method known as pre-amp decoupling is used (4). The input impedance of the pre-amp is deliberately tuned to a small value (usually only a few ohms), approximating

a short circuit and creating an impedance mismatch. A half-wave segment of cable between the pre-amp and the coil element then transforms this impedance to a short circuit at the input terminals of the coil element, where there is typically a parallel LC circuit along with a shunt diode that serves to detune the receive element during RF transmission. For the case of pre-amp decoupling, the diode is not forward biased, but the inductor in the LC circuit is shunted to ground by the low transformed impedance from the pre-amp input. This has the effect of forming a high-impedance LC resonator, blocking current form being induced in the coil element by other elements. However, voltages induces in the coil by the magnetization signal source are still able to couple through the LC circuit and into the pre-amp.

The entire assembly consisting of the TR module and pre-amp housings, the 8-way power divider, and the cable breakout box is shown in Fig. 3.8.

3.3.3 Coil array performance

Coil profiles are estimated from gradient echo reference images using the adaptive algorithm described in (41). Coil sensitivities for the RF transmit-receive array are shown in Fig. 3.9. Significant flip angle variation occurs near the periphery due to the close proximity of the sample to the microstrip elements, whose combined fields only approximate a birdcage mode within the centermost 6.5 cm diameter of the coil. The drawback to the coil design is that for samples approaching the available diameter of 10 cm, bright spots occur adjacent to the microstrip lines and signal voids are observed in the region between elements where the B_1 fields of neighboring elements are oriented in opposite directions, causing cancelation.

The performance of the coil for conventional parallel imaging is shown in Fig. 3.10. The *g*-factor is calculated using the coil maps shown in 3.9 for the case of regular undersampling in k-space along the phase encode (vertical) direction. The *g*-factor is essentially a voxel-wise map of the condition number for the small matrix inversions performed during

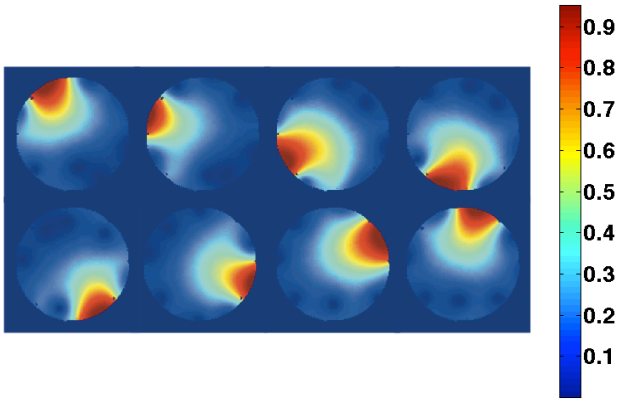

Figure 3.9: Coil profiles as measured on a uniform water phantom using a gradient echo sequence and the adaptive methods described in (41). The plots are normalized so as to lie on the scale 0-to-1. The maps are proportional to the coils' B_1 fields (as measured in gauss), but without specialized measurements, it is difficult to separate the multiple contributions to the image intensity (proton spin density, T_1, T_2, etc.) from the effect of the B_1 sensitivity on the image intensity.

SENSE unwrapping. As expected, the array performs satisfactorily for R=2 SENSE parallel reconstructions, introducing little noise amplification, because the 8 coils are able to easily unwrap the 2 signals aliased onto each voxel in the reduced FOV. For any pair of voxels in the reduced and full FOV, there are always a pair of coil profiles whose sensitivities are high in one voxel and low in the other.

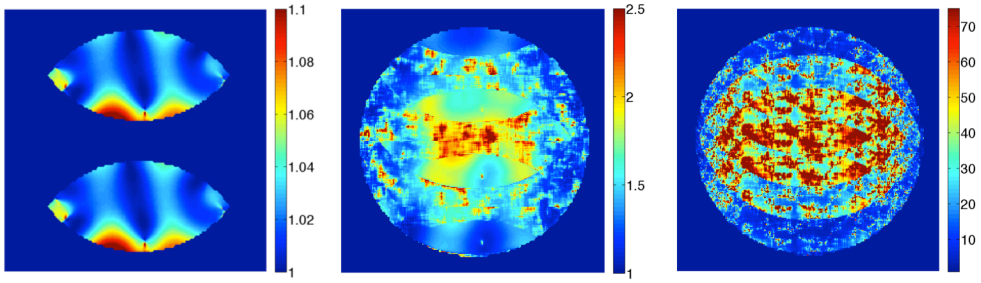

Figure 3.10: *G*-factor describing acceleration performance for experimental 8-channel array performing SENSE parallel imaging with undersampling along the phase encode (vertical) direction. The *g*-factor maps the factor of noise amplification introduced at each voxel by SENSE parallel image reconstruction. From left to right, the acceleration factors are R=2, R=4, and R=8. At R=8 the problem is badly ill-conditioned and noise grows uncontrollably in the absence of regularization.

At R=4, it is difficult to unwrap four aliased signals from each voxel, since the coils form an encircling array rather than a linearly array along the phase encode direction. This

limitation is reflect in the growth of the g-factor. Also, at R=4, the g-factor is higher in the central regions, where the coil profiles are less distinct along the phase encode direction and therefore less able to unwrap the aliased signals. At R=8 the matrix problem for unwrapping each aliased voxel is highly ill-conditioned throughout the FOV, and the g-factor grows by orders of magnitude. SENSE parallel imaging at R=8 with 8 coils is therefore impossible.

By contrast, undersampled projections methods such as radial and O-Space degrade more gracefully in the presence of undersampling, as will be shown below. The three primary reasons for this are (a) the improved complementarity between the SEMs and the coil sensitivities, (b) the intrinsic regularization of the Kaczmarz iterative reconstruction algorithm, and (c) the fact that undersampling these trajectories introduces both noise amplification as well as blurring, while in Cartesian SENSE the undersampling causes only noise amplification.

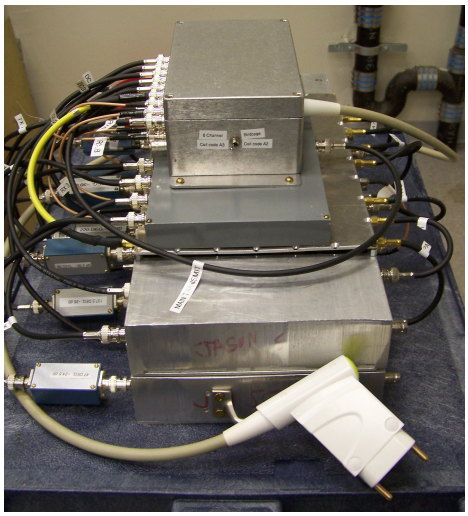

Figure 3.8: RF front end including, from top to bottom, breakout box for interfacing with Siemens RF and DC cables; 8-way power divider for distributing transmit pulse among all 8 coil elements; pre-amp board embedded in aluminum housing; transmit-receive modules positioned below pre-amp board within the housing. Lumped element phase shifters in each transmit path are visible at the far left. Modular design allows for easy troubleshooting and for using the electronics with different RF coils.

3.4 SEM calibration: A minute to learn, a lifetime to master

The most important experimental task is to calibrate the DSU strength setting, which ranges from 0% to 100%, so that it corresponds to a field strength in Hz/cm^2. To accomplish this, calibration maps are acquired on a uniform cylindrical water phantom doped with manganese to reduce T_1. Field mapping of SEMs is typically performed by acquiring multiple Cartesian images, each with a different amount of phase preparation applied by an SEM pulse prior to readout. Frequency maps are calculated using the voxel-wise derivative of the phase through this image series. However, in nonlinear projection imaging approaches, while field maps provide a useful estimate of field strengths and higher-order impurities, they have not provided sufficient accuracy to ensure high quality image reconstruction with nonlinear SEMs, as will be shown below.

A nonlinear encoding strategy known as 4D-RIO has been proposed (15) in which pairs of linear and second-order SEMs play out reciprocating, interleaved radial k-space trajectories. 4D-RIO images were initially degraded due to errors in the field maps for the second-order (PatLoc) fields. The poor image quality is thought to result from the effects of eddy currents and other defects in the field mapping sequence. This obstacle is overcome by decomposing 4D-RIO into its two component radial trajectories. Separate images are then acquired using only the linear SEMs and only the second-order SEMs. The xy and $x^2 - y^2$ fields are then scaled and rotated as necessary to align the nonlinear image with the linear image, providing improved field calibration for reconstructing the full 4D-RIO dataset.

Another confounding factor for O-Space imaging with the Z^2 field is through-plane dephasing from the three-dimensional variation of the nonlinear SEM (86). In multi-slice axial field mapping, in-plane frequency map bias has been reported due to the through-

plane evolution imparted by the z^2 field variation during readout. The observed bias grows with increasing slice thickness and with distance from isocenter.

The aforementioned problems motivate our use of a phase mapping, an extension of chemical shift imaging that uses phase encoding to map all spin dynamics in each voxel of the FOV. Phase maps can be used to directly populate the encoding matrix, greatly improving the fidelity of O-Space image reconstruction.

3.4.1 Limits of conventional field mapping

Conventional field maps are acquired using a Cartesian pulse sequence modified to include a test SEM pulse whose moment (amplitude $\times$ duration) is varied between TRs (3). For O-Space calibration, the test Z^2 strength is fixed and the duration is varied along with a reciprocal delay between the Z^2 pulse and the readout, such that TE remains constant. Phase images are computed for each SEM moment and the frequency is calculated as the slope of the phase through each voxel as a function of the pulse moment. Typically field maps are acquired over multiple slices in the imageable FOV and then approximated with best-fit 3-D spherical harmonics. This approach provides sufficient accuracy for measuring B_0 inhomogeneities, assessing the success of shimming, and even measuring nonlinearities in SEM fields (3).

Unfortunately, when voxel phase evolution is calculated using conventional frequency maps, very poor image quality results (Fig. 3.11). The principal cause of the artifacts was discovered to be a severe bias in the frequency maps due to the z^2 field variation of the $z^2 - \frac{1}{2}(x^2 + y^2)$ field. When slices evolve under a linear through-plane SEM, the slice phase dispersion is symmetric about the mid-point, preserving a net measured phase equal to that at the slice center. Thus, while through-plane SEM pulses can cause signal loss, they will not change the net phase of the slice, thereby leaving in-plane frequency maps unaf-

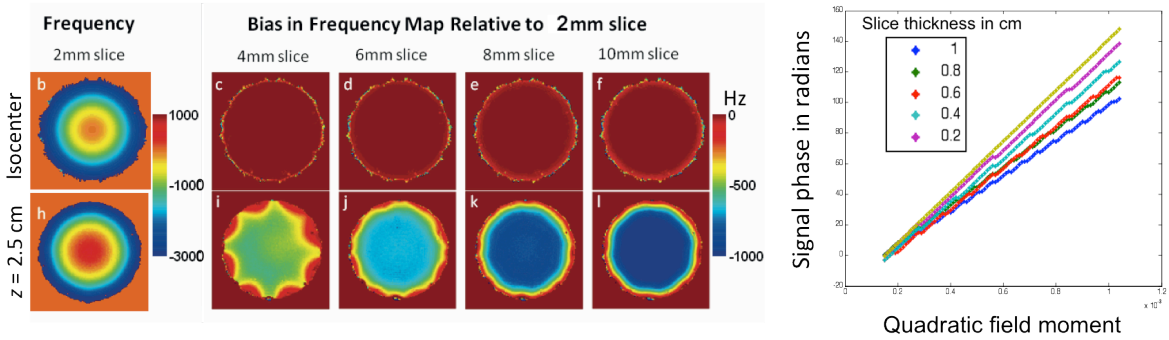

Figure 3.12: Frequency map bias due to through-plane phase evolution from a nonlinear SEM as a function of slice thickness and offset from isocenter. Slice phases evolving under a linear SEM do not show this bias. Reproduced from (87).

fected. By contrast, detailed investigation (87) has revealed that when a slice evolves under a quadratic field, the phase evolves asymmetrically about the midpoint, with a nonlinear component that can be described as precession about an ellipse. The resulting frequency map bias (Fig. 3.12) grows both with slice thickness and with offset from isocenter, since the z^2 grows rapidly.

3.4.2 Phase mapping

Phase mapping is an approach used in (88) to measure the k-space trajectories traversed by time-varying SEMs played during parallel RF excitation. The approach is a simple extension of chemical shift imaging (89) and point spread function mapping (90, 91). The sequence under test is modified through the addition of in-plane phase encoding. For the case of O-Space imaging, phase encoding steps are introduced along X and Y prior to readout. The O-Space sequence is then repeated $N_{PE_x} N_{PE_y}$ times, where N_{PE_x} and N_{PE_y} are equal to the number of phase encoding steps along X and Y, respectively. The number of phase encoding steps sets the resolution of the resulting phase map in each direction. As in conventional Cartesian pulse sequences, the size of each X phase encoding increment is chosen to be $1/FOV_x$ (likewise for the Y phase encodes). Acquired phase maps have four

dimensions: time, coil index, X phase encode, and Y phase encode. The data are Fourier transformed along both phase encoding directions, producing images of the spin phase evolution for each time point in the O-Space sequence (for all CPs and coils).

Care must be taken to sample the phase maps more finely than the final reconstructed image. Otherwise, intravoxel dephasing causes a loss of signal at points in the trajectory where a high degree of phase modulation is applied. Intravoxel dephasing occurs when a voxel is large enough so that the spread of spin phase within a voxel approaches π radians. This causes the vector sum of signal contributions from different spins in the voxel to cancel each other out, creating regions of signal loss in the image. In phase mapping during O-Space sequence, dephasing is of greatest concern for voxels furthest from the center placement, so the map voxel size should be chosen small enough to prevent substantial dephasing from occuring in voxels experiencing the strongest SEM field. Initial phase mapping of the O-Space SEM was performed at 64×64, but near the beginning and end of each acquired echo, many spins in the FOV were dephased, leaving only noise in these regions. Reacquiring the maps at 256×256 produced much clearer maps.

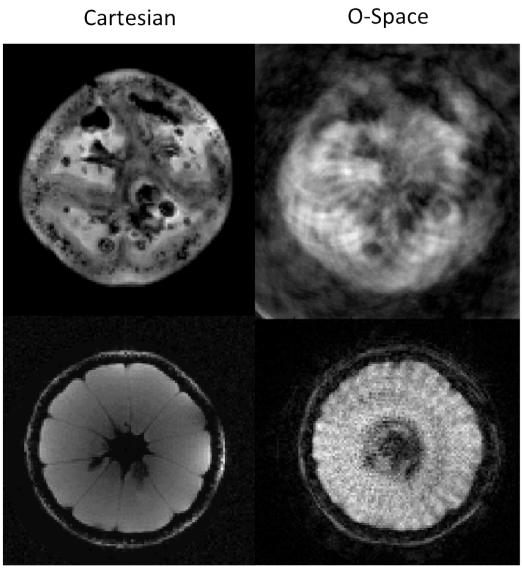

Figure 3.11: Cartesian and O-Space images reconstructed from experimental data acquired on tomato (top row) and orange phantoms. When conventional field maps are used to populate the encoding matrix for O-Space reconstruction, severe image artifacts occur. (FOV=10 cm, 256 readouts, Z^2 strength is $680 Hz/cm^2$). In these images, artifacts result because the encoding matrix, E, is populated with the wrong amplitude of each field component generated by the Z^2 insert coil. One of the key reasons for this systematic error is the bias introduced into Z^2 field maps by through-plane dephasing. O-Space reconstruction depends upon accurate knowledge of the amplitude of the desired Z^2 SEM as well as other impurity SEMs (uniform, linear, and higher-order).

Phase mapping, while tedious, enables proof-of-concept imaging on a dedicated high-power Z^2 field insert coil. Estimation of the spin dynamics from field probes (92, 93, 94) is more time-efficient but requires specialized apparatus. Accurate mechanical positioning of the field probes in the FOV is also crucial. Furthermore, field probes must be carefully designed to avoid inhomogeneous broadening. By contrast, phase mapping requires no special hardware or preparation other than a uniform water phantom and a simple modification of the existing O-Space imaging sequence. If the phase mapping acquisition time grows too long or the data load becomes too large, phase encodes may be skipped along both X and Y, accelerating the phase map by a factor of $R = R_x R_y$. After Fourier transformation along each phase encode direction, the resulting phase maps will be aliased. Unwrapping of the aliased maps can be performed using the multi-coil data and the SENSE algorithm, just as for conventional undersampled images.

3.4.2.1 Phase mapping experiments

First, a conventional field mapping sequence is run to measure any residual inhomogeneity in the object after it had been shimmed. The shim map is acquired by running ten gradient echo scans of size 128×128, with TE displaced by 400 μs between successive acquisition (3). The relatively short delay of 400 μs is chosen to minimize the possibility of phase wrap. The frequency is calculated from the slope of the phase in each voxel across all ten phase images.

Phase maps are then acquired at 256×256 resolution for each point in a Z^2 readout with 256 samples, producing roughly 1 GB of data across all 8 RF channels. For a TR of 40 ms, phase mapping for one center placement (or for the Z^2 SEM alone) lasts 45 minutes. Maps are acquired for several Z^2 strengths ranging up to 20%. The phase mapping sequence, shown in Fig. 3.13, is identical to the O-Space imaging sequence, except for two

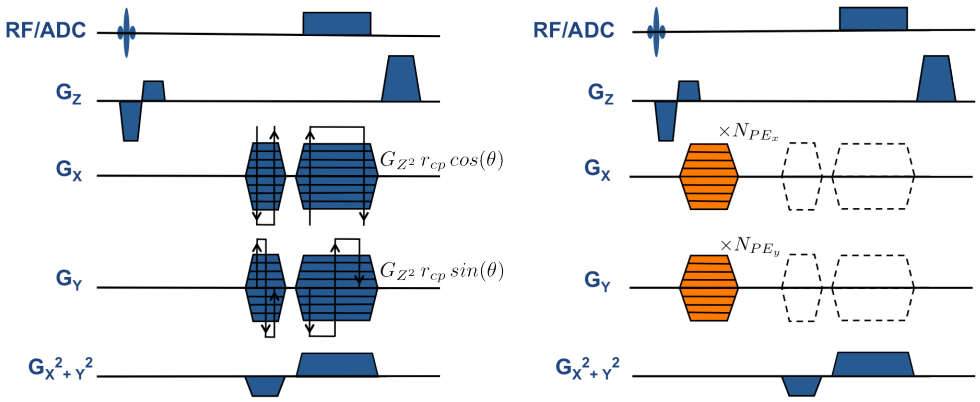

Figure 3.13: The phase mapping sequence (right) is identical to the O-Space sequence (left) except for the removal of the linear SEMs radial trajectory pulses (dotted lines), the addition of phase encoding pulses (shown in orange) along both X and Y, and the use of a shorter TR. The phase mapping sequence is repeated $N_{PE_x} N_{PE_y}$=256^2 times to provide transverse-plane phase maps with a resolution of 256×256. The grid must be chosen densely enough to prevent intravoxel dephasing at the periphery of the FOV. If desired, center placement encoding fields may be mapped by reintroducing the linear SEM pulses of the radial trajectory. See text for full description.

differences: (a.) the addition of phase encoding pulses in X and Y prior to readout and (b.) disabling of the X and Y dephasing and readout pulses in order to isolate the Z^2 field. In this way, it is assured that the phase maps accurately reflect the actual spin phase evolution produced during the O-Space imaging sequence. Because all images presented in this paper are taken at $z = 0$, only this single axial slice is used for phase mapping. Because of this, through-plane field evolution due to the z^2 variation of the SEM were kept to a minimum (86). But in the general case, phase mapping provides an accurate calibration of the Z^2 SEM for any slice, including oblique slices.

Along with the phase maps, coil profiles are acquired on the same uniform phantom using a standard Cartesian pulse sequence (41). This permits the phase maps from all 8 coil channels to be combined using a SENSE reconstruction with R=1 (1), removing the amplitude and phase information imparted by the coil profiles. The resulting phase maps that display only the phase evolution of the spins in the object. Combining the 8 coil channels also improves the SNR of the resulting phase maps, since each coil is only sensitive to one localized ROI. In the combined phase image, the slope through time in each voxel is used to calculate the local frequency, as shown in Fig. 3.14. A weighted least squares fit is used to determine the slope, with each point weighted according to the signal in the associated magnitude image for each readout point. Frequency maps are then decomposed into polynomials up to 6th order using a least squares fit weighted by the magnitude image for each voxel at each time point.

Phase maps taken at different Z^2 strengths yielded extremely similar normalized field maps and polynomial decompositions (Fig. 3.15), showing that the SEM scales linearly with the applied current. Linearity was further verified with oscilloscope measurements of the current supplied by the Techron amplifier.

Image reconstruction must account for all field components generated by the Z^2 in-

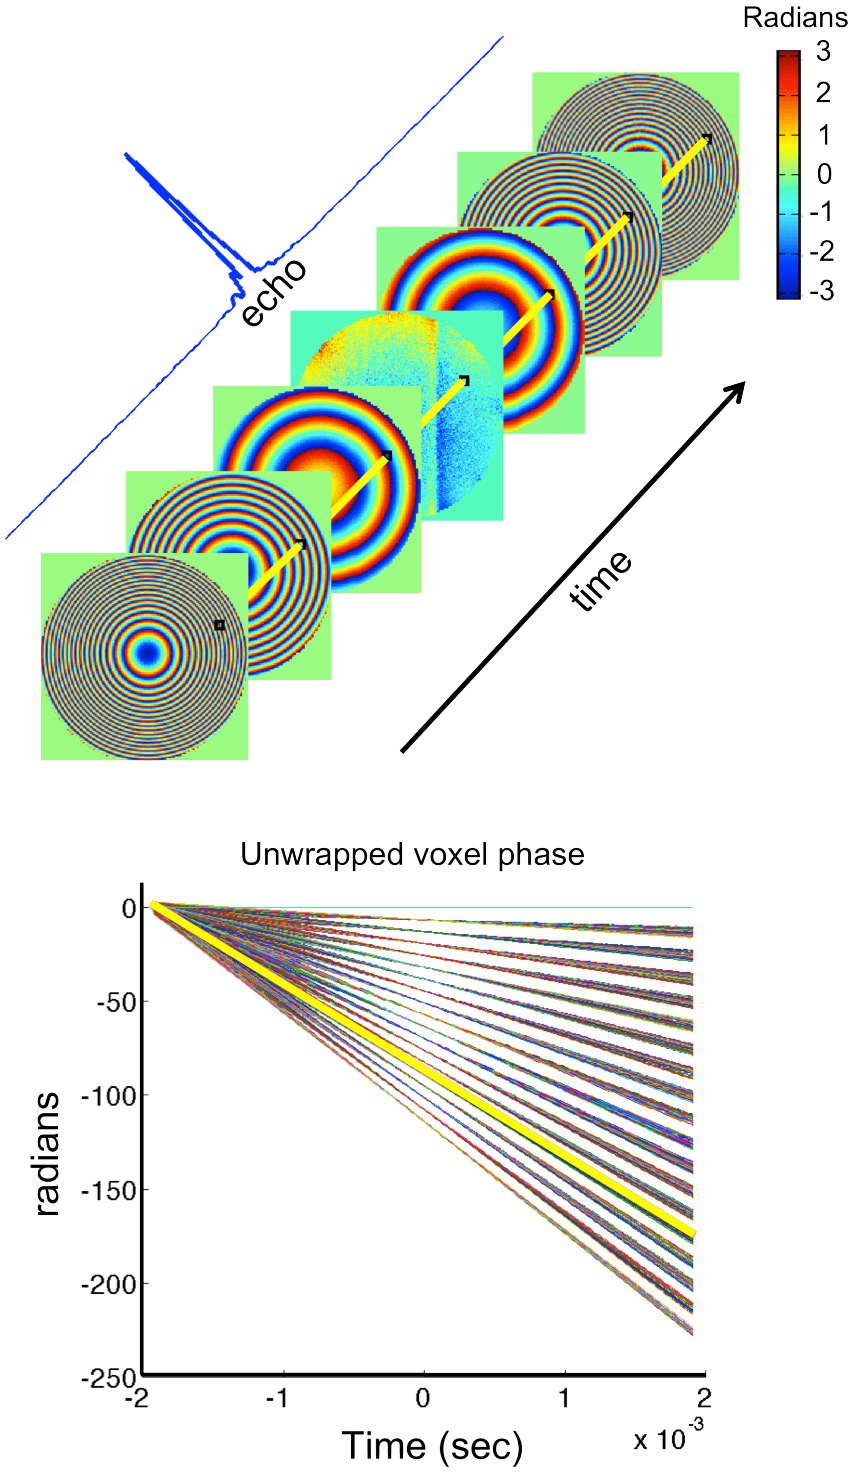

Figure 3.14: Phase maps of the evolving quadratic field are taken at each point in the readout, corresponding to different parts of the echo ranging from the flat phase at the echo peak to the tightly-wound phase at the edges. The unwrapped phase for a representative voxel is plotted over time in yellow. The slope of the plot gives the angular frequency at that voxel, while the intercept of the line provides timing information. Eddy currents would manifest as deviations from linearity.

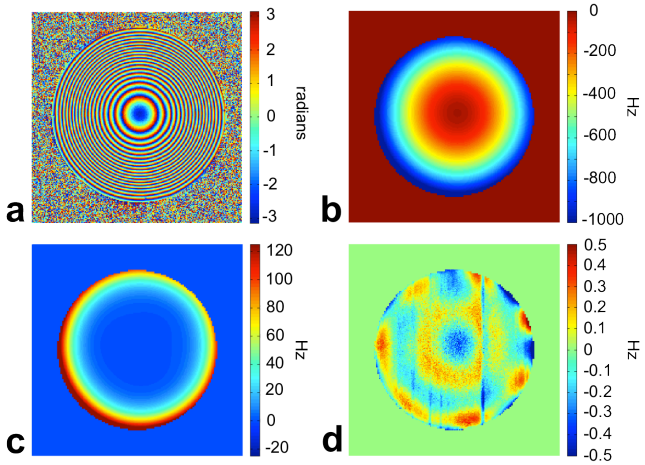

Figure 3.15: Top left: raw phase image at the beginning of a readout plotted in radians at the edge of a readout window, with maximum phase winding. Z^2 strength=680 Hz/cm^2. Top right: corresponding polynomial fit up to 2nd order plotted in Hz. Bottom left: Fit components up to 6th order in Hz. Bottom right: Residual field after fitting. The field components are normalized to describe the fields produced when the Z^2 is driven at 1% of its peak current level. Accurate knowledge of all field components is crucial for calculating the rows of the encoding matrix.

sert coil, shown in Table 3.2. As expected, the field is dominated by the quadratic $x^2 + y^2$ component. Linear components along X and Y also exist, shifting the quiescent center of the Z^2 by roughly +1 cm in the X and Y direction relative to the scanner isocenter. The slight difference between the x^2 and y^2 terms reflects the existence of a small "PatLoc" multipolar $x^2 - y^2$ field, counterpart to the small xy field also produced by the coil. A static field offset is also generated and must be carefully incorporated into image reconstruction to avoid a radial offset of each ring-like frequency isocontour form its true value. Minor higher-order components up to 6^{th} order were also detected. After the polynomial fit, the remaining residual was less than +/-1 Hz.

Best fit field components from phase maps acquired on different days agreed with each other to well within $\pm 1\%$. This result indicates that any mechanical displacement of the insert coil between experiments is small, and its effect on the magnetic fields is on the order of the noise floor of the phase mapping calibration. If the coil were to move in the z-direction between experiments, the resulting axial-slice phase map would experience a static offset of approximately 138 $Hz/cm^2/\%$ due to the translation of the z^2 hyperboloid pair relative to the scanner coordinate system.

To perform image reconstruction, the field components in the table are combined and scaled according to the strength of the applied Z^2 SEM. The resulting composite field is used, along with the coil profiles, to populate the encoding matrix, as shown in Fig. 2.9. To compute different rows of the encoding matrix, the phase winding at different time points is explicitly calculated by scaling the time variable in the exponent of phasor containing the composite SEM field.

P	C	P	C	P	C	P	C
00	-35.32	21	0.06	04	-0.05	60	0.04
10	20.73	12	-0.01	50	-0.05	51	0.00
01	23.67	03	0.10	41	-0.06	42	0.14
20	-68.50	40	-0.04	32	-0.10	33	0.00
11	-0.53	31	0.04	23	-0.11	24	0.14
02	-68.10	22	-0.12	14	-0.04	15	0.00
30	0.16	13	0.01	05	-0.05	06	0.05

Table 3.2: Best-fit components (C) of each polynomial (P) up to 6th order for the field produced by the Z^2 insert coil, as estimated from the phase maps. "20" denotes x^2, "11" denotes "xy", and so on. Coefficients are expressed in $Hz/cm^n/\%$, where n is the sum of the two subscript indices. The normalized field decomposed in this table corresponds to the Z^2 insert coil running at "1%" of peak current. To populate the encoding matrix, the field components in the table must be combined, scaled appropriately, and used to calculate the encoding function for each time point and coil channel (see Fig. 2.9).

3.4.3 Eddy currents

To test for the presence of eddy currents, the method presented in (30) is generalized to the case of nonlinear slices. The Z^2 SEM is used to select a hyperboloidal slice at a known frequency offset. Two readouts are acquired: a free induction decay and a readout during which the Z^2 SEM is switched on. The frequency of the Z^2 SEM is calculated as the derivative of the phase difference between the two readouts, removing any effects from coil phase and residual B_0 inhomogeneity. A variety of Z^2 strengths between -6% and +6% are investigated. These strengths are lower than the 10-20% range more typically used in imaging experiments, but stronger field strengths can not be used for the slice-phase readout mea-

surements, or else the test field causes too much through-plane dephasing, resulting in poor SNR.

Measurements of Z^2 field evolution based on a single slice show no evidence of significant eddy currents. The Z^2 field has a small overshoot caused by the Techron amplifier[1], so a settling time was added between the Z^2 pulse and the onset of analog-to-digital conversion. The plot also shows a roughly 65 microsecond delay between the rising edge of the linear SEMs and the Z^2 SEM; this delay is introduced by the DSU electronics and the filter plate between the amplifier and the scanner room. All timing changes in the Z^2 readout pulse were compensated by corresponding changes in the dephasing lobe to preserve the echo peak at $t = TE$.

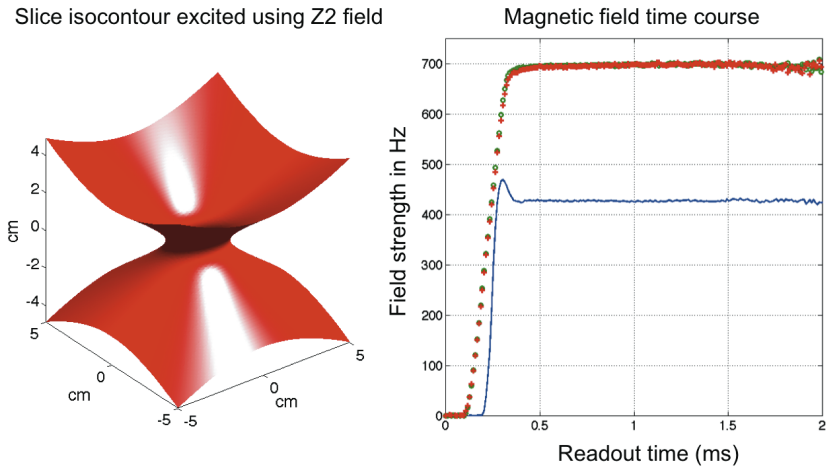

Figure 3.16: At left, the shape of the isocontour excited by playing a Z^2 field instead of a Z field during the RF pulse. At right, the field strength measured using for the Z^2 (blue), X (red), and Y (green) pulses.

In addition, the experiments described in (95) were reproduced to test for eddy currents. The Z^2 field was switched on, left on long enough for eddy currents to die down (1000 ms) and then switched off just prior to RF excitation in a Cartesian gradient echo

[1]The overshoot is likely due to the fact that the amplifier is typically operating at only 10-30% of its peak output in O-Space imaging during signal readout. The amplifier provides cleaner waveforms when it is operating closer to its maximum capacity of 120 amps.

sequence with a short readout (less than 1 ms) and short TE (less than 5 ms). The delay between the falling edge of the Z^2 pulse and the RF pulse was varied to sample eddy currents at different points in their time evolution. The phase of the resulting images was used to compute the local frequency evolution in each voxel. Eddy currents in the Z^2 and Z_0 components of the field were found to be extremely small, essentially insignificant for the imaging applications discussed in this paper.

Further evidence that eddy currents were minimal was provide by the slope of phase over time in each voxel of the phase maps, as shown in 3.14. Phase evolution was found to be highly linear in voxels from all locations in the FOV, betraying little influence of eddy currents, which would be expected to damp the phase evolution following the rising edge of the Z^2 pulse at the beginning of the readout. In summary, although the DSU system is equipped with pre-emphasis eddy current compensation for the Z^2 field at three separate time constants, it was deemed unnecessary for the present work.

3.4.4 Timing errors in the linear SEM pulses

For the radial sequence, two sets of projections are acquired such that they traverse k-space in opposing directions. Overlays of the spokes are compared using the method in (96) to remove any effects from coil sensitivities and B_0 inhomogeneity, permitting timing errors for the linear SEMs to be quantified and corrected.

Timing errors in the radial sequence were found to be quite modest for the readout bandwidth chosen, producing at most a 1/4-point offset from the true center of k-space. Based on the small size of this offset, and the quality of the radial images, it was assumed that the linear field trajectories were known with sufficient accuracy and required no further calibration.

3.5 Experimental results

Radial and O-Space imaging are performed using the same pulse sequence, but with the Z^2 coil enabled during O-Space acquisitions. For all acquisitions, the TE is 9 ms, TR is 1.5 s, and the readout bandwidth per pixel is 260 Hz (corresponding to 67 KHz for a 256-point readout). In order to avoid aliasing in the radial direction, the sampling bandwidth and Z^2 SEM strength $G_Z^2 2$ were chosen so that the outermost frequency ring did not fall within the object for a particular CP. The linear SEMs are set to the same strengths used for a conventional radial sequence, with a factor of $\sqrt{2}$ additional coverage along the readout direction as compared to a Cartesian trajectory. The Z^2 strength is chosen to position the ring of center placements near the edge of the FOV, following the simulations in (59). Subsequently, acquisitions are performed with all SEM strengths doubled to explore the effect of dense SEM encoding on image resolution.

Institutional Review Board (IRB) approval was obtained from the Yale University School of Medicine Human Investigations Committee. Board approval permitted the O-Space experiments to include human anatomy, including axial-plane images of the hands of healthy volunteers.

Image reconstruction is performed using the Kaczmarz iterative algorithm (69), Eq. (2.5.1), with single-precision arithmetic to ease the computational demands. Although the conjugate gradients and Kaczmarz algorithms have similar convergence properties, the latter is useful for tuning reconstruction parameters because the image estimate can be observed on a time scale as fine as one data point backprojection at a time. By contrast, conjugate gradients requires a large matrix-vector product to be completed before an updated estimator is available.

Phase mapping information is incorporated into the Kaczmarz reconstruction in two

different ways. In the first case, the phase maps are weighted by the coil profiles and then used directly to populate the encoding matrix, E, for Kaczmarz iterations. In the latter case, the rows of the encoding matrix are calculated on the fly using normalized X, Y, and calibrated Z^2 field shapes along with a "k-space trajectory" describing the timing and strength of each SEM readout pulse[1]. Because close agreement is observed between the two sets of images, the best-fit approach is chosen for the reconstructions performed in this chapter for several reasons. First, the best-fit coefficients grant the investigator much greater flexibility, permitting the synthesis of fields with any amplitude. This is useful if, for example, the Z^2 pulse amplitude changes between acquisitions. This flexibility also allows the fields to be recalculated directly on any desired image grid, while the phase maps need to be interpolated if the grid size does not match their dimensions. A second advantage of the synthetic fields is that they are free from the noise present in the phase maps (though this noise still impacts the best-fit coefficients that are used to calculate the synthetic fields).

Image quality is compared across Cartesian gradient echo, O-Space, and radial images for different acceleration factors. Radial images of size 256x256 with 256 spokes are actually undersampled by a factor of $\pi/2$ according to the Nyquist criterion, but for purposes of comparison with O-Space and gradient echo images, they are referred to here as "R=1" datasets. For fairness of comparison, the radial images are reconstructed with the Kaczmarz algorithm rather than one of the conventional multi-coil re-gridding approaches (34, 35).

Cartesian gradient echo, radial, and O-Space images are compared in Fig. 3.17 for the case of double-strength SEMs, with the center placement unmoved. The center placement lies at r=4 cm and the FOV is 10 cm. The reference gradient echo is a fully-sampled 384×384 image. The radial and O-Space data are reconstructed to 512×512 to permit small object features to be visualized. The radial and O-Space images approach the level

[1]For the rectangular pulses used here, with minimal eddy current distortion, the k-space trajectory for each field component is simply a line.

of detail showed in the gradient echo image, which uses 1.5 times more data. In undersampled reconstructions, noise quickly overwhelms the Cartesian SENSE image, while the radial and O-Space images qualitatively seem to exhibit combination of noise amplification and loss of resolution. The O-Space images are comparable to the radial images, a sign of progress given the maturity of radial projection imaging and the novelty of the O-Space approach. While this illustrates the success of the careful O-Space calibration, it is expected that as O-Space methodology matures and any remaining sources of error are eliminated, the images will improve relative to radial images, and will more clearly display the spatially-variable resolution inherent to nonlinear SEMs.

3.5.1 *In-vivo* O-Space images

An *in-vivo* image comparison using a clenched human hand is shown in Fig. 3.18. The center placement remains at r=4 cm in a 10 cm FOV, as in the previous figure, but with the gradient strengths halved and the reconstruction grid set to 256×256. As in the phantom images, undersampling results in elevated noise levels and blurring, but most of the fine anatomical features remain visible at R=4, especially near the periphery, where the quadratic SEM variation is greatest. The periphery is where receive coil arrays have the best intrinsic ability to resolve, since receiver B_1 sensitivities are higher near the elements. This complementarity between quadratic SEMs and the SNR of coil arrays is an advantage that O-Space imaging shares with schemes such as PatLoc.

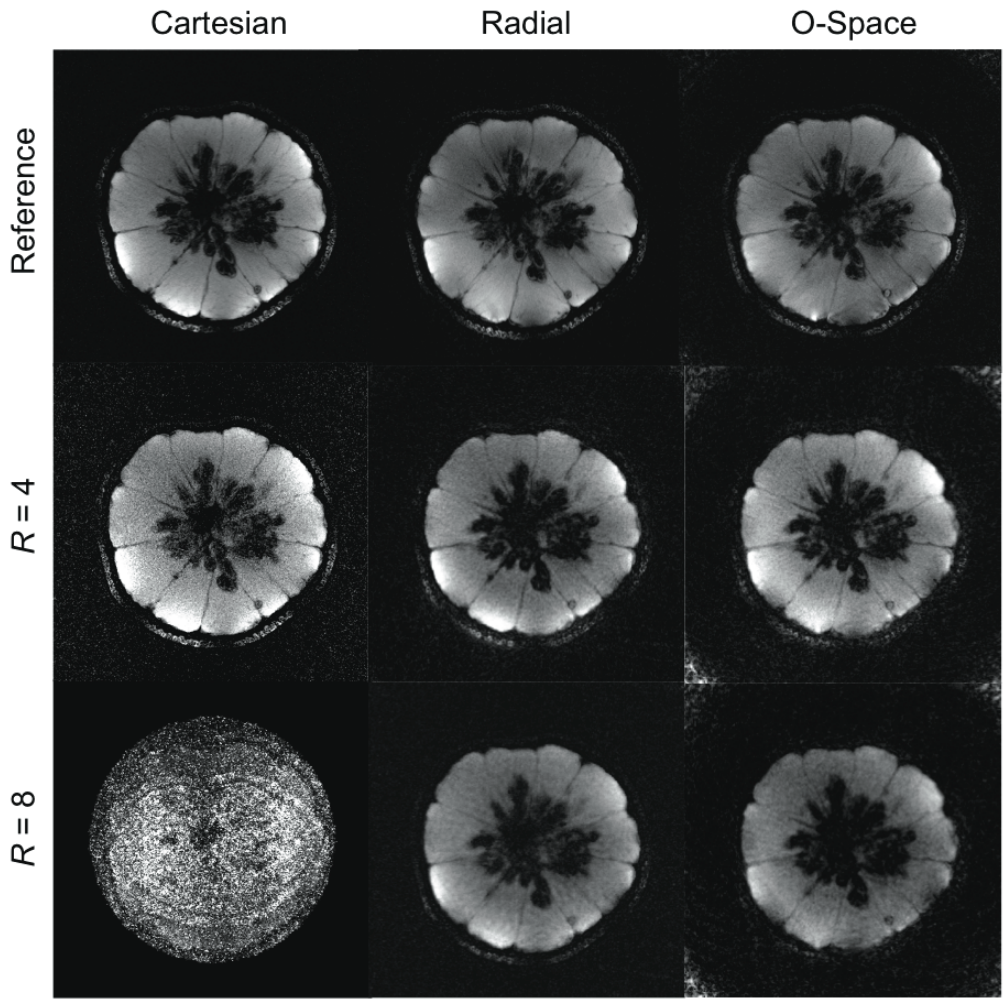

Figure 3.17: Comparison of axial-plane (z=0) gradient echo, radial, and O-Space images of an orange phantom. The top row compares a fully-sampled Cartesian image (384×384) with 256-spoke radial and O-Space images each reconstructed to 512×512. In subsequent rows, the data are undersampled by 4-fold and 8-fold and reconstructed to the same image grid. FOV is 10 cm.

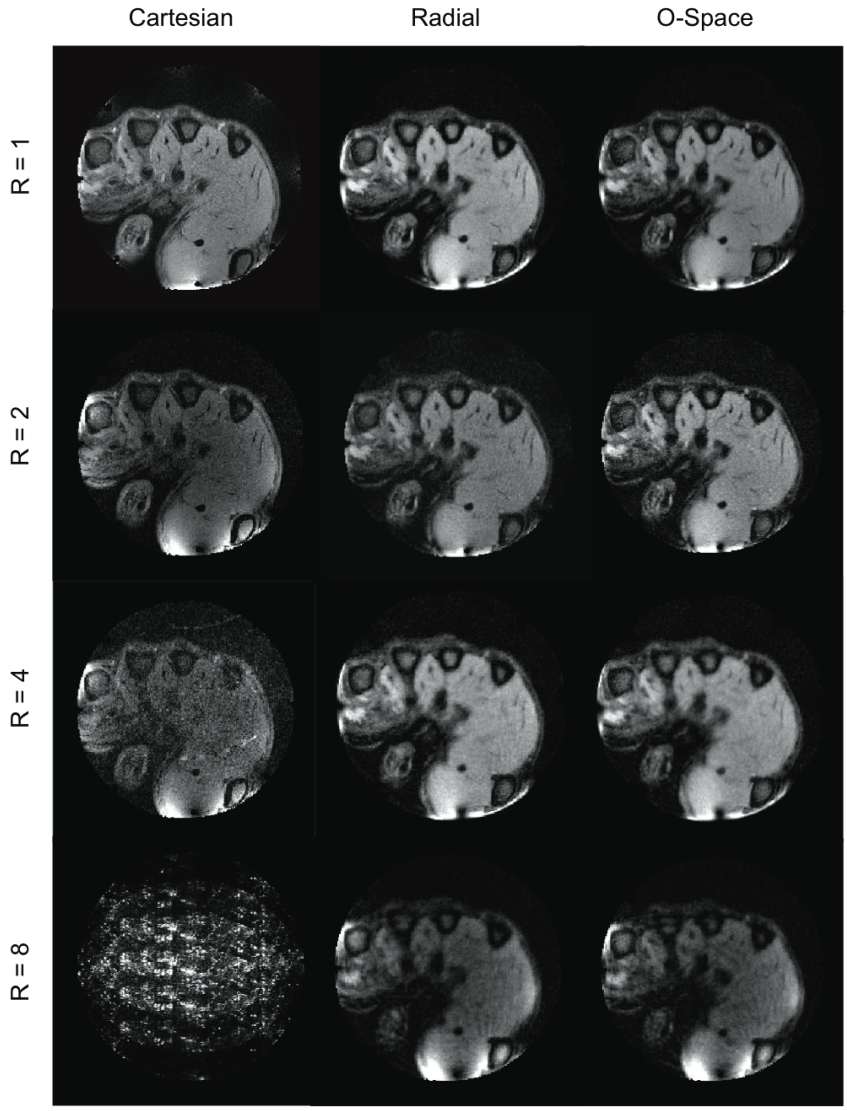

Figure 3.18: Comparison of axial-plane (z=0) Cartesian, radial, and O-Space 256×256 reconstructions of the hand of a healthy volunteer acquired using the 8-channel RF array inside the bore of the Z^2 insert coil. Subjects were instructed to lean in from the rear of the scanner bore and hold their hand inside the 12 cm Z^2 insert assembly (shown in Fig. 3.1). The patient table was positioned so as to be within arm's reach of the rear opening of the scanner bore. For each acceleration factor, R, equivalent amounts of data are used in each of the three images shown. A small B_1 artifact is visible at the bottom of the imaged anatomy. FOV is 10 cm.

3.6 Artifacts

3.6.1 Density of the encoding matrix

Fig. 3.19 shows the importance of choosing an appropriate reconstruction grid for O-Space encoding schemes. In the lower-resolution O-Space image, a banding artifact arises from intravoxel dephasing because the voxels described by the encoding matrix are too large.

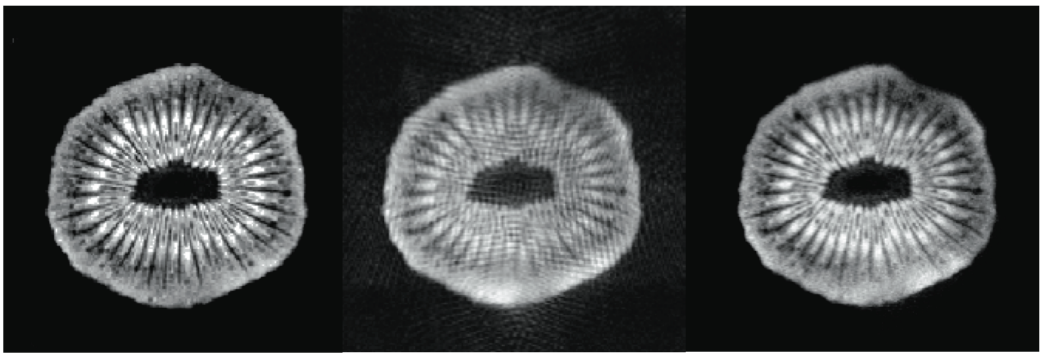

Figure 3.19: From left to right, 6-cm FOV gradient echo reference image at 128×128, 256-spoke O-Space dataset reconstructed to 128×128, and the same O-Space data reconstructed to 256×256. The strengths of the Z^2 and linear SEMs are $1360 Hz/cm^2$ and $13,196 Hz/cm$, respectively, causing over 2π radians of phase evolution between neighboring voxels on the 128×128 grid.

The problem is solved by reconstructing the same data onto a higher-resolution image grid. The disadvantage of this approach is that it requires longer reconstruction times due to the enlarged encoding matrix needed to describe the encoding fields on a finer grid of voxels[1]. The voxels in the reconstruction grid must be small enough to prevent the phase spread within a voxel from approaching π radians. After reconstruction, the high resolution image may be downsampled if a coarser resolution is desired for the final image.

[1] The time required to complete a Kaczmarz iteration grows linearly with the size of the encoding matrix

3.6.2 Effect of field strength and timing errors

To explore the effect of timing errors between the actual and assumed Z^2 pulse timing, images were deliberately reconstructed with a shift in the Z^2 field's "k-space" trajectory. This was achieved by setting the time variable, t, to the wrong value in computing the encoding matrix used for image reconstruction (see Fig. 2.9). Fig. 3.20 shows the effect of introducing a 10-point time shift, corresponding to 150 μs for the readout BW used (67 KHz). This models an unusually high degree of imbalance between the dephasing and rephasing pulses of the Z^2 SEM, larger than one would typically encounter in practice, so it can be considered a worst-case scenario.

The initial expectation was that timing errors would cause serious image degradation, as they do in conventional radial k-space trajectories. On the contrary, Z^2 timing errors were found only to introduce a quadratic phase on the reconstructed object. Because the timing and strength of the Z^2 pulse is the same for each TR, errors in the Z^2 pulse timing introduce a "k-space" shift vector (in the quadratic analogue of k-space). The reconstruction algorithm treats this additional quadratic phase as belonging to the object since it is not described by the encoding matrix. The situation would be more complicated if O-Space imaging did not play the same Z^2 pulse during every readout. If the amplitude were to change between TRs, the image quality would be much more sensitive to pulse timing errors. But for the present O-Space experiments, timing errors in the Z^2 pulse were ruled out as a source of harmful artifacts.

When the Z^2 SEM field components are deliberately scaled by the incorrect amount, such that mismatch exists between the applied field and the field modeled in the encoding matrix, severe "double-ghost" artifacts result (also encountered in section 4.5.2). This simulation models the effects of an inaccurate calibration, underscoring the need for accurate calibrations such as the phase mapping approach described in section 3.4.2.1. The ghost appears to be a new species of MRI artifact not previously described in the literature. It is

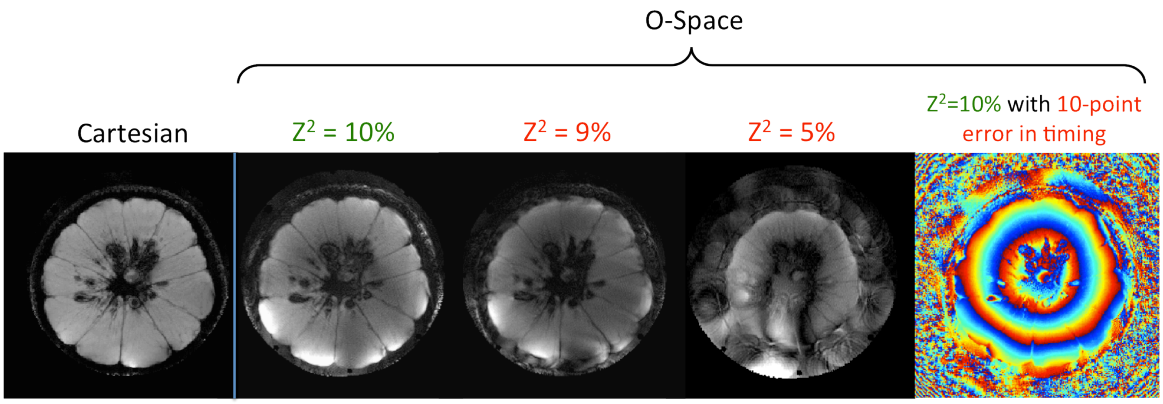

Figure 3.20: "Ghosting" and distortion occur when the wrong Z^2 amplitude is deliberately used during O-Space reconstruction. No ghosting occurs when the correct Z^2 field amplitude is used (10%), but even a 1% systematic error in the amplitude causes pronounced ghosting and signal loss. By contrast, the effects of pulse timing errors are benign. Since the Z^2 readout pulse is the same for every TR, timing errors manifest only as a quadratic phase in the image, leaving the magnitude image unaffected. The time shift modeled here is equivalent to 10 readout points (150 μs).

characterized by two distorted copies of the object, one of them slightly contracted and the other greatly enlarged (this can be seen more clearly in the images shown in Fig. 4.13). The ghosts seem to arise from improper assignment of signals to frequency isocontours in the reconstructed FOV. The sharper the Z^2 variation in a given region of the FOV, the larger the spatial displacement of the signals and the more expanded the ghost. This explains why the inner ghost is only moderately smaller than the true object while the outer ghost is substantially enlarged. Fortunately, the phase mapping calibration used in this study was sufficiently accurate to avoid visible ghosting.

3.6.3 Problems at high Z^2 strengths

In O-Space imaging, the center of the quadratic bowl moves through a sequence of radii, $r_l = \frac{|G_{Z^2}|}{\sqrt{G_{X_l}^2 + G_{Y_l}^2}}$, for center placements enumerated by $l = 1...L$. When the Z^2 SEM pulse is paired with a radial trajectory playing on the linear SEMs, as it is for all O-Space images shown in this thesis, the center placements simple lie on a ring of constant radius. When the

Z^2 SEM is strong relative to the linear SEMs, the center placements are brought well within the boundaries of the object. Although this induces no discernible artifacts in simulations, it causes a pronounced radially-oriented banding artifact in experimental images. For the images shown in Fig. 3.21, O-Space images with center placements just outside the object show no artifacts, while those with center placements inside the object show severe banding.

Because this artifact does not appear in simulations, it seems unlikely that it is attributable to a lack of encoding at the center of the quadratic bowl. Although it is true that no encoding occurs at the center of a given individual center placement, all other center placements will provide encoding at this voxel.

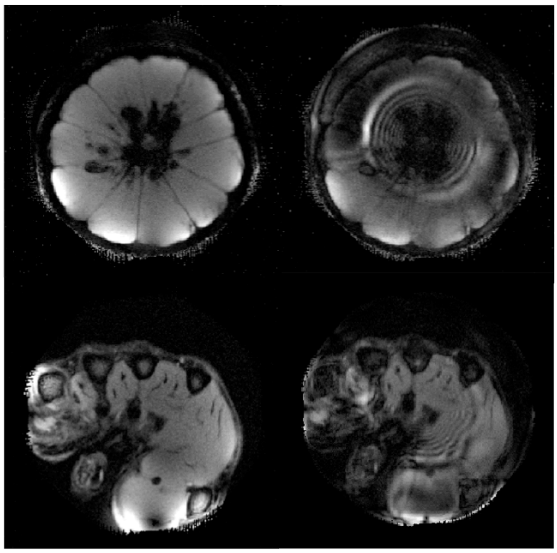

Figure 3.21: O-Space reconstructions in a 10 cm FOV with 10% (left) and 20% Z^2 strengths using identical linear SEM trajectories. This corresponds to center placement radii of 4.5 cm and 2.25 cm, respectively. The latter CP produces banding artifacts in experimental images that do not appear in simulations of the same encoding fields.

Efforts to eliminate the banding artifact by tweaking the SEM calibration have so far been unsuccessful. Attempts have been made to reduce the banding by tweaking the Z_0 as well as Z^2 components of the phase map calibration for the quadratic field, to no avail. The artifact also persists when the resolution is increased (as in Fig. 3.19), ruling out intravoxel dephasing as a possible source. The cause of this artifact therefore remains an open problem. Until a solution is found, center placement location will remain a constraint on practical implementation of O-Space imaging, with the best performance provided by CPs located outside the object.

3.7 Convergence properties of the Kaczmarz algorithm

One of the central questions in any iterative reconstruction problem is when to stop the iterations and accept the image as having "converged". For the Kaczmarz method, one must also choose a value between 0 and 2 for the relaxation parameter, λ[1]. Iterative regularization such as the Kaczmarz method have some "intrinsic" regularization due to the fact that with increasing numbers of iterations, the estimator tends to become less biased and more noisy. It is not easy to define the iteration for which the image has "converged", since extra iterations can always be computed, producing some minute change in the image. Most iterative algorithms are characterized by a classic trade-off between bias and noise, described by Fig. 3.22 below. In medical imaging, the question of convergence arises in multiple different modalities for a broad class of algorithms including conjugate gradients (97) and iterative expectation maximization (98) in addition to the Kaczmarz method used here.

To study the convergence properties of experimental *in vivo* O-Space images, the data shown in Fig. 3.18 was reconstructed using a variety of λ values for between 1 and 15 iterations. The plot in Fig. 3.22 shows the percentage change in mean voxel energy from one iteration to the next, with the first iteration omitted because the change is relative to an initial image of zeros. Convergence is quantified using the fractional change in mean voxel energy,

$$\delta_{n,\lambda} = E_x \frac{\| \hat{m}_n^2 \| - \| \hat{m}_{n-1}^2 \|}{\| \hat{m}_{n-1}^2 \|} \tag{3.7.1}$$

where n indexes the iteration and E_x is the expectation operator, calculated as a simple average over all voxels. The initial estimator, m_0, is taken to be an image of all zeros. Not surprisingly, for large values of λ, the algorithm approaches steady state more quickly than

[1]The conjugate gradients algorithm does not typically require a relaxation parameter, but one must still choose how many iterations through the encoding matrix to compute.

for small values. Inset against the plot are representative images for a few points on the plot. For the case of a small λ value of 0.01, $\hat{m}$ is qualitatively quite blurry on the 2^{nd} iteration and becomes gradually sharper and noisier as the iterations progress. By contrast, after only 2 iterations, the λ=.5 reconstructions are already quite noisy.

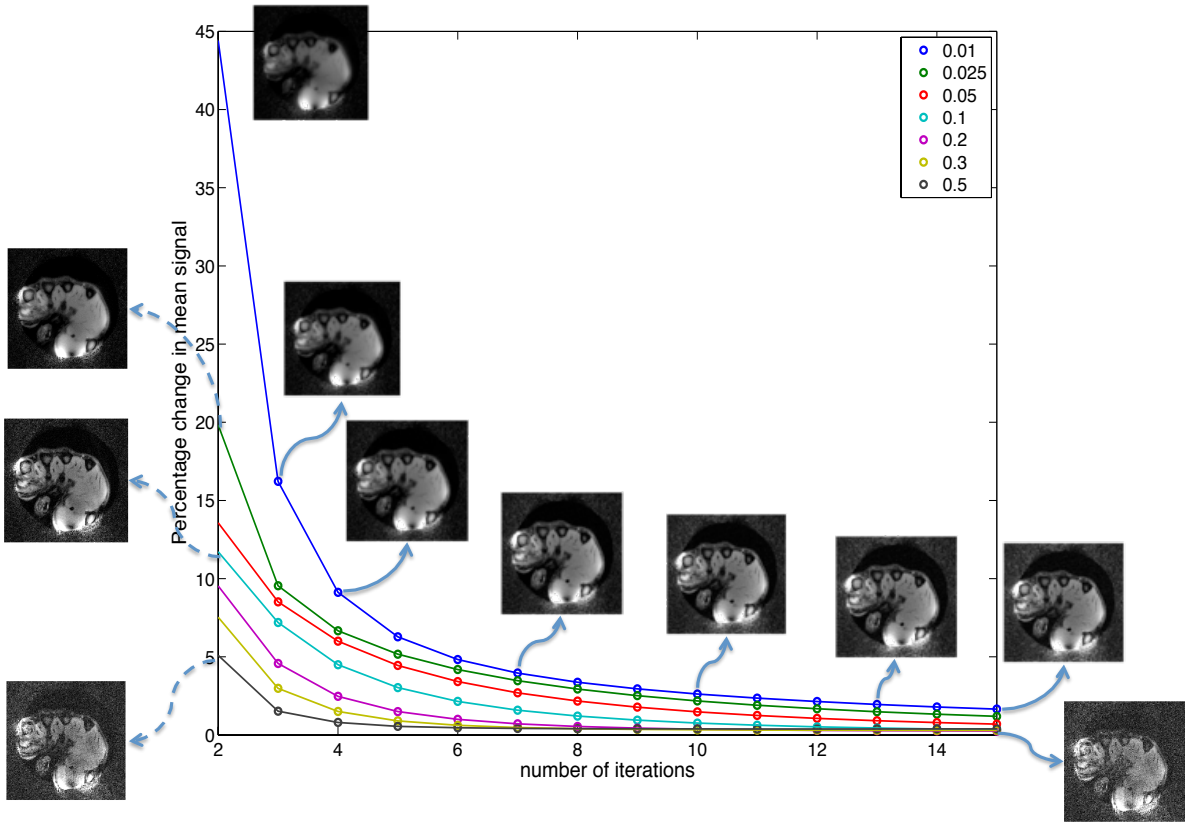

Figure 3.22: Convergence of the Kaczmarz algorithm based on the metric in Eq. (3.7.1) as a function of relaxation parameter λ. Representative images show a clear tendency for noise and bias to increase and decrease, respectively, as more iterations are used. Images appear blurry in the first few iterations when small values of λ are used. But for larger λ values, iterations beyond a certain point contribute primarily noise to the estimator.

To further quantify the performance of the Kaczmarz algorithm, regions are defined used for computing the signal (yellow) and noise (red) levels as shown in Fig. 3.23. Because coil profiles are only available within the circular ROI defined by a radius of 4 cm, the reconstruction problem outside of this circle is ill-defined and characterized by high noise levels. This outlying region is therefore ignored for purposes of the analysis and in practice

should be excluded from images with a mask.

The following metrics are used to assess the properties of O-Space image estimators:

- **The average intensity**, μ, within the signal ROI (yellow), normalized relative to the mean intensity over the entire circular FOV;
- **The average standard deviation**, σ, of the noise in the background noise ROI (red), also normalized relative to the mean intensity in circular FOV;
- **The signal-to-noise ratio**, computed as $SNR = \frac{\mu}{\sigma}$.
- L_2**-norm**, the total intensity of the image, defined using the complex modulus, $\| \hat{m} \|_2 = \sqrt{\sum_{\rho=1}^{N^2} |m_\rho|^2}$.

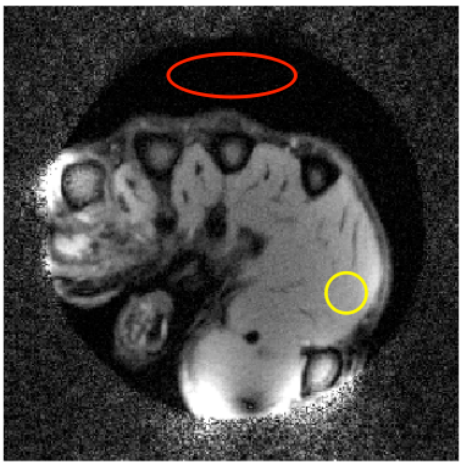

Figure 3.23: The ROIs defined for computing signal (yellow) and noise (red) statistics. Note that signal and noise are both computed in the circular region for which coil profiles are available. Outside this region, coil profiles can not be acquired, so the reconstructed voxels are meaningless.

These metrics are shown in Fig. 3.24. The signal intensity of the yellow ROI, μ, is normalized relative to the mean intensity in the circular FOV for ease of display. The noise standard deviation is also normalized relative to the mean signal level so that its plotted values can be referenced against a mean signal intensity of unity. For all values of λ, the signal levels off[1] after a few iterations while the noise continues to grow steadily. This causes the

[1] While some variation exists between steady-state signal levels as a function of λ, the differences are small compared to the variation among the corresponding noise levels at each iteration.

SNR plot to fall off consistently as the algorithm iterates. Comparing the convergence and SNR plots, it is clear that a trade-off is required between using too few iterations, which may be inadequate to resolve all image features, and too many iterations, which reduce the SNR by introducing unnecessary noise.

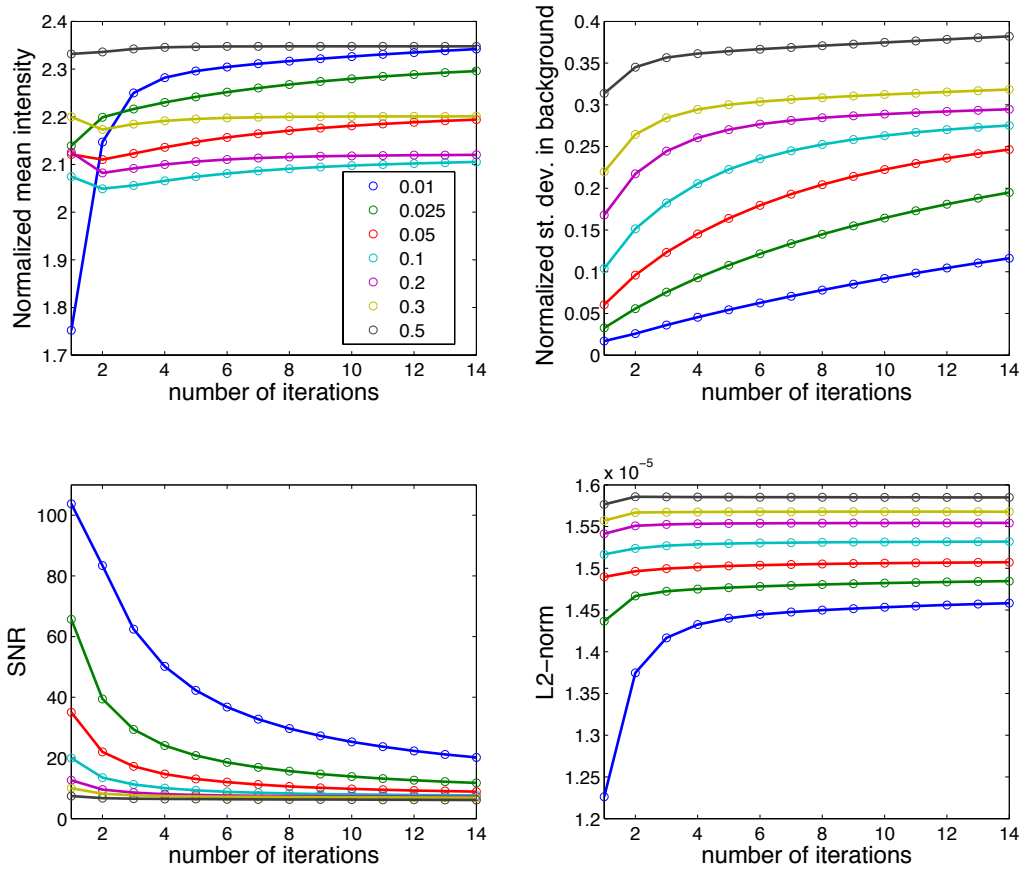

Figure 3.24: Variation in image statistics as a function of relaxation parameter λ and number of Kaczmarz iterations. Use of a small relaxation parameters provides the minimum-norm least squares estimator, providing superior SNR according to the metric used here. See text for more details.

Initially, the chosen stopping rule was to iterate until the mean change in intensity was less than 5%. However, a better approach is to use the signal vs. noise curve in Fig. 3.25. For small values of λ, there is a knee in the curve beyond which further iterations contribute mostly noise without increasing signal levels (at least for the chosen signal and noise regions used in this study). These curves also help compare the convergence achieved

by different values of λ. For the knee points indicated in the red circle, nearly the same signal/noise point is reached in 5 iterations when λ=0.01 but only 2 iterations when λ=0.025. The limitation of this plot is that it only considers one signal ROI and one noise ROI in the FOV. A more detailed study of the (possibly spatially-varying) convergence properties of O-Space images is reserved for future work.

In addition to defining a stopping criterion, it is necessary to choose a value to use in practice for the relaxation parameter. For very small λ – a condition known as strong underrelaxation – the algorithm requires a few more iterations for the signal to approach a steady state value, as compared with larger values of λ. For λ=0.01, the small change in mean signal and the qualitative high image quality above ~6 iterations shows that convergence has been achieved without elevated noise levels. This regime corresponds to the minimum norm least squares estimator, consistent with the plot of the L_2-norm in the figure. The Kaczmarz method has previously been shown to converge to the minimum-norm estimator in the strongly underrelaxed limit, with a proof provided in (72). In conclusion, while the Kaczmarz method is in principle stable for $0 < \lambda < 2$, in practice when using noisy data and an inconsistent matrix problem, the strongly underrelaxed reconstructions provide the best trade-off between noise and bias, at the expense of longer computation time. Use of a small

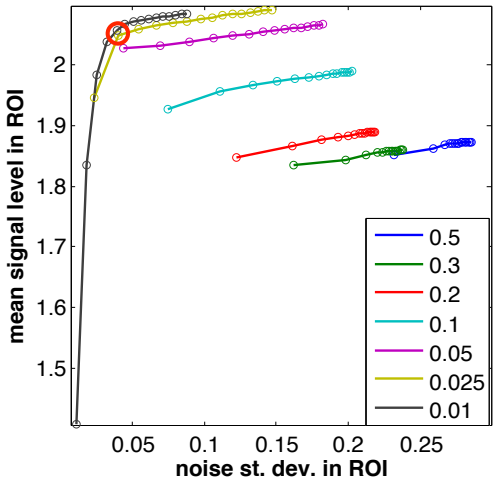

Figure 3.25: Mean signal level in the object ROI is plotted versus mean noise standard deviation in the background ROI from Fig. 3.23. Consistent with the SNR plot in Fig. 3.24, the images reconstructed with smaller values of λ achieve equivalent (or greater) signal levels with less noise. The knee of the curves is used to provide a stopping point for the reconstruction, since further iterations introduce significant amounts of noise without appreciable enhancement of the signal. For the stopping points shown in the red circle, a slightly larger λ value (0.025) achieves the same noise statistics in 2 iterations that require 5 iterations when λ=0.01.

λ also allows the user to fine tune the trade-off between noise and bias, reducing the likelihood that the image will overshoot the desired stopping point in the iterations.

3.7.1 Computational burden

Strongly underrelaxed Kaczmarz reconstructions take longer to converge, requiring additional processor time. Such reconstructions can not be performed in real-time on a clinical MRI scanner using a conventional processor. Fortunately, the Kaczmarz method can be easily implemented on parallel processors using a method called "string averaging" (99). In this approach, the encoding matrix is broken up into P equal sized segments, each of which is used by a separate processor to perform backprojection of the corresponding data. After each processor has cycled through all rows residing in its segment, all P resulting estimators for $\hat{m}$ are averaged. The resulting estimator is used for the next iteration through the encoding matrix. The process repeats on each processor, providing a convergent estimator in as little as $1/P$ the time required for computation on a single processor. This approach is particularly powerful when paired with Graphics Processing Units, or GPUs, which have recently attracted interest for computationally demanding MRI reconstructions (100, 101). GPUs typically have dozens or even hundreds of independent cores that are designed to perform fast floating point operations. By partitioning the encoding matrix into a large number of strings, Kaczmarz reconstruction times could be reduced by two orders of magnitude. However, GPU implementation of the Kaczmarz method lies beyond the scope of the current study.

3.7.2 Moving beyond the L_2-norm in O-Space imaging

Regularization in this work can be performed in two ways. The first is to use the "intrinsic regularization" of the iterative Kaczmarz method by choosing the number of iterations and the relaxation parameter so as to reach a desired point in the bias-noise convergence curve.

The second is to modify the Kaczmarz method to incorporate an explicit penalty on the norm of the image estimator $\hat{m}$. This latter approach was used unsuccessfully in a previous section to suppress artifacts caused by B_0 inhomogeneity. The regularization prior was able to limit artifact propagation, but only at the expense of introducing an unacceptable blur. Both of these methods, which can be used in combination, both have the effect of trading off between noise and bias, in the second case by explicitly imposing a penalty on the L_2-norm of the image estimate.

Many other priors have been widely used in the image reconstruction community. The most immediately relevant to O-Space image reconstruction, and MRI more generally, is the L_1-norm, which penalizes the sum of the absolute value of the coefficients in the image estimate: $\| \hat{m} \|_1 = \sum_\rho |m_\rho|$. In contrast to the L_2-norm, which tends to work as a smoothing operator, the L_1 norm has the potential to reduce noise while preserving edges. A simple illustration of this is to consider two neighboring voxels with intensities $[\,2 \;\; 0\,]$. Let Gaussian noise be added to each voxel's representation in encoding space. Both the L_1 and L_2 norms will tend to suppress the noise. But minimizing the L_2-norm during reconstruction will also have the effect of smoothing the discontinuity between the two voxels in an attempt to reduce the L_2-norm from 4 toward 2. The L_1 norm, however, has no "incentive" to smooth the discontinuity from $[\,2 \;\; 0\,]$ toward $[\,1 \;\; 1\,]$ because both have the same L_1 norm.

The L_1 norm has been successfully used in reconstructing undersampled radial data (35), which benefits from having artifacts that are less coherent that those produced in undersampled Cartesian images. Radial images were found to benefit both from denoising and from streak artifact suppression. The artifacts in O-Space imaging are arguably even less coherent than those in radial imaging, so O-Space benefits stand to reap even greater benefit from use of the L_1 norm. Another application of this norm is in compressed sensing, where it is used to enforce sparsity in a domain where the object is naturally sparse, such as the wavelet domain. The reconstruction then favors large coefficients in this sparse domain

while truncating small coefficients that primarily represent image noise, permitting high-SNR, highly-accelerated image reconstruction (102, 103). To explore these interesting possibilities, future work in O-Space image reconstruction will include use the L_1-norm as a way to reduce image noise while preserving object edges. A version of the Kaczmarz method that incorporates the L_1-norm is presently being sought.

4

Phase Scrambling and the Fractional Fourier Transform

4.1 Generalizing the Fourier transform

Chapters 2 and 3 have introduced the concept and experimental realization of O-Space imaging with quadratic fields. The present chapter changes gears to explore other uses of the Z^2 insert coil for localized and/or accelerated spatial encoding, as well as an alternative framework for reconstructing O-Space data:

- Quadratic field "lensing" for scalable FOVs with no aliasing;
- Gradient localization (GradLoc) for imaging a desired region of interest within the FOV;
- Reconstruction of O-Space data using the variable-order Fractional Fourier transform (FrFT).

Two of these applications take advantage of a generalization of the Fourier transform known as the Linear Canonical Transform, or LCT (104). The first order of business is to intro-

duce this operator and explain its connection to MRI acquisitions that use nonlinear SEMs.

4.1.1 The Linear Canonical Transform

The Linear Canonical Transform is a unitary integral operator based on a quadratic exponential kernel:

$$K_{a,b,c,d}(u,\rho) = \sqrt{\frac{1}{ib}} exp\left(i\pi\left(\frac{d}{b}\rho^2 - \frac{2}{b}\rho u + \frac{a}{b}u^2\right)\right) \tag{4.1.1}$$

where u is a unitless normalized spatial coordinate, $u = x\frac{\sqrt{N}}{FOV}$ (given N voxels in the FOV), and ρ is a normalized frequency-domain coordinate. For a normalized object $f(u)$, the explicit form of the LCT integral is

$$LCT_{a,b,c,d}(\rho) = \sqrt{\frac{1}{ib}} \int exp\left(i\pi\left(\frac{d}{b}\rho^2 - \frac{2}{b}\rho u + \frac{a}{b}u^2\right)\right) f(u)du \tag{4.1.2}$$

The LCT operator is more succinctly described by a 2×2 *ABCD* matrix,

$$M_{LCT} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \tag{4.1.3}$$

The ABCD matrix causes affine transformations of a function's time-frequency representation, or Wigner distribution (104). This representation is a spectrogram of the object in which the horizontal axis represents time or space (typically) and the vertical axis represents frequency. The magnitude or phase of an object can then be plotted in the third dimension or using a color scale. The following three instances of the LCT are important to the present discussion:

Rotation ⇒ The Fractional Fourier transform (FrFT) is defined as a rotation angleα on $[-\pi, \pi]$. At α = 0, the FrFT reduces to the identity operator, while at $\alpha = \pi/2$ it becomes the Fourier transform. At $\alpha = 2\pi$, the FrFT reproduces the object inverted through the origin, as shown in Fig. 4.1.

$$\begin{aligned} FrFT_\alpha\{f\}(\rho) &= \sqrt{1 - i\cot(\alpha)}e^{i\pi\cot(\alpha)\rho^2}\int e^{-i2\pi\csc(\alpha)\rho u}e^{i\pi\cot(\alpha)u^2}f(u)du \\ M_{FrFT_\alpha} &= \begin{bmatrix} \cos(\alpha) & \sin(\alpha) \\ -\sin(\alpha) & \cos(\alpha) \end{bmatrix} \end{aligned} \tag{4.1.4}$$

Shear along spatial axis ⇒ The Fresnel transform, which describes optical free-space diffraction. The Fresnel transform can also be described as a convolution of the source with a quadratic "chirp" kernel.

$$\begin{aligned} Fresnel_r\{f\}(\rho) &= \sqrt{\frac{1}{ir}}\int e^{i\pi(\rho-u)^2/r}f(u)du = f(\rho) * \sqrt{\frac{1}{ir}}e^{i\pi\rho^2/r} \\ M_{Fresnel_r} &= \begin{bmatrix} 1 & r \\ 0 & 1 \end{bmatrix} \end{aligned} \tag{4.1.5}$$

Shear along frequency axis ⇒ Corresponds to multiplication by a chirp function. This operation describes passage through a thin lens in an optical system.

$$\begin{aligned} Chirp_q\{f\}(\rho) &= f(\rho)e^{-i\pi q\rho^2} \\ M_{chirp_q} &= \begin{bmatrix} 1 & 0 \\ -q & 1 \end{bmatrix} \end{aligned} \tag{4.1.6}$$

For the FrFT, the rotation angle varies between $-\pi$ and π, with values outside this range repeating the behavior in this interval modulo 2π. For $\alpha = 0$, the FrFT reduces to

the identity operator. The Fourier transform corresponds to the special case of complete rotation into the frequency plane, with $\alpha = +/-\pi/2$,

$$\begin{aligned} FT\{f\}(\rho) &= e^{-i\pi/4}\int f(u)e^{-i2\pi u\rho}du \\ M_{FT} &= \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \end{aligned} \tag{4.1.7}$$

Optical wavefronts propagating in free space can be described using quadratic phase functions (105). Using the Fresnel transform, the complex-valued wave at a given distance from the source is calculated by convolving the source with a quadratic kernel describing the diffracting spherical wavefront. To a first-order approximation, the effect of a thin lens is to impart a quadratic phase to a wavefront that causes it to either diverge away from or converge toward a focal point. This is equivalent to a chirp multiplication such as that shown above. Together, lenses and free space propagation are illustrated by the optical system in Fig. 4.1, which can be modeled using LCT operators.

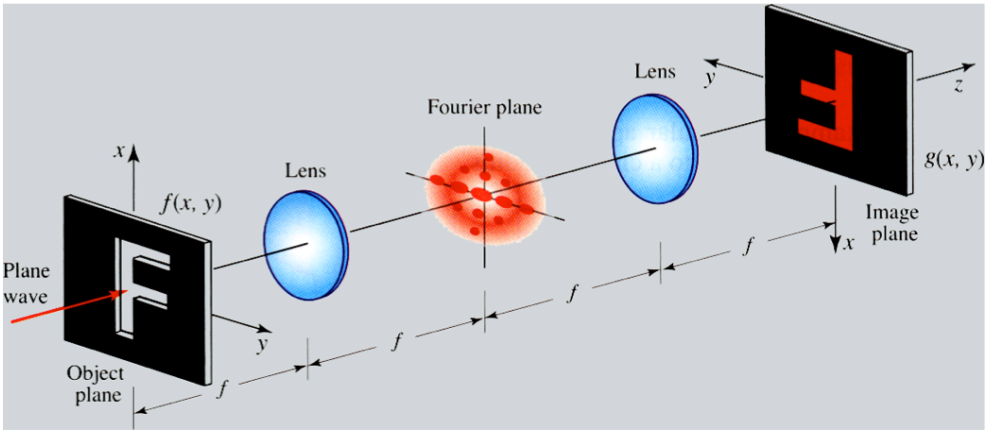

Figure 4.1: Optical system consisting of a light source, free space propagation, thin lenses, and an image plane. The effect of each part of the system on the wavefront can be modeled using the linear canonical transform and its daughter functions. Reproduced from (106).

Using multiple LCT matrices to model the effects of free space propagation and passage through a thin lens, the optical wavefront can be calculated at any plane in the system (104).

Like the Fourier transform, the LCT is a unitary operator. The unitarity of the LCT can be proved by showing that inner products are preserved by the transform. Consider two arbitrary functions f and g:

$$\begin{aligned}\int f(u)g^*(u)du &= \int \left[\int F_{a,b,c,d}(\rho)K_{d,-b,-c,a}(\rho,u)d\rho\right] g^*(u)du \\ &= \int\int F_{a,b,c,d}(\rho)g^*(u)K_{d,-b,-c,a}(\rho,a)dud\rho \\ &= \int\int F_{a,b,c,d}(\rho)g^*(u)K_{a,-b,-c,d}(u,\rho)dud\rho \\ &= \int\int F_{a,b,c,d}(\rho)g^*(u)K^*_{a,b,c,d}(u,\rho)dud\rho \\ &= \int\int F(\rho)G^*(\rho).\end{aligned} \tag{4.1.8}$$

This result is a generalization of Parseval's relation, familiar from Fourier theory, to the LCT operator. Unitarity implies that the inverse of an LCT operator is simply its complex conjugate transpose, a property which holds both for the LCT's integral operator as well as for its matrix representation.

4.2 Lensing with quadratic encoding fields

Just as the LCT and its daughter transforms can be used to model optical system, so too can they describe the effect of quadratic phase modulations in MRI. Scalable reconstructions have been achieved through the addition of a "phase scrambling" quadratic SEM pulse to a conventional Cartesian pulse sequence. Phase scrambling has previously been used for dynamic range compression (107) since it changes the echo time on a voxel-by-voxel basis, spreading signal away from the peak of the echo. But the applied quadratic phase also mimics the effect of an optical system consisting of thin lenses and free space propagation, such as that shown in Fig. 4.1. The quadratic phase enables reconstruction of alias-free

images from undersampled data using the Fresnel transform (108, 109) or the chirp-Z transform (110). The resolution of the final image is still determined by the coverage of the acquired k-space lines in the frequency domain, but the technique enables flexible reconstruction using a single dataset, potentially benefitting localizer scans in MRI protocols.

Phase scrambled MRI data can be cast as a Fresnel transform (free-space propagation) in combination with multiplication by chirp functions (thin lens). This permits fast reconstruction either via convolution with a chirp kernel (109) or via inverse filtering in the Fourier domain (108). The image is then multiplied by the same chirp function, evaluated in spatial coordinates, to remove any residual quadratic phase.

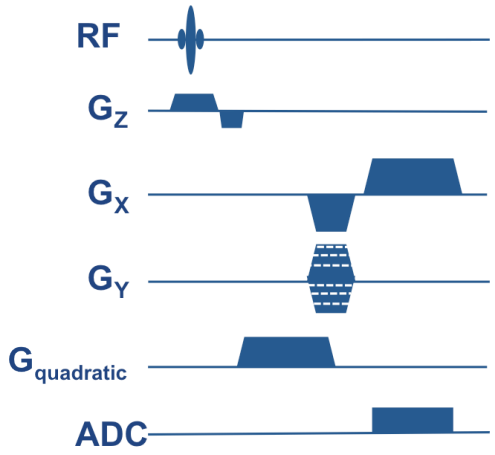

Figure 4.2: Pulse sequence used for scalable-FOV lensing using phase scrambling quadratic pulse.

In the present work (111), we show that phase-scrambled MRI data can be described as the Fractional Fourier transform of the object. The relationship between the FrFT and Fresnel transforms is demonstrated using *ABCD* matrices. We then present scalable MR images acquired using the Z^2 insert coil and reconstructed via the FrFT. In previous work on phase scrambled reconstruction, the quadratic SEM moment was applied using low-power shim coils, necessitating the use of impractically long TE times (108, 109).

4.2.1 Phase scrambled MRI signal

Assume that an object $f(u)$ is imaged using a Cartesian sequence (Fig. 4.2) with a quadratic trapezoidal pulse of strength $-\gamma H$ applied for τ seconds prior to readout, giving a pulse

moment $-\ \gamma H\tau$. Let u be a unitless spatial coordinate as defined in Eq. (4.1.1) and k be the conventional k-space coordinate. In the discrete case used for practical computation, u and k are both normalized, such that $k = u = [-N/2, N/2-1]/N$, for a FOV with N voxels.

For the 1-D case, the MR signal in the presence of a quadratic SEM moment is

$$s\left(k(t)\right) = \int \exp\left(-i2\pi\left(k(t)u + \gamma H\tau u^2\right)\right) f(u)du \tag{4.2.1}$$

where $\gamma H\tau$ is quadratic SEM moment. The signal equation can be recast in the form of the FrFT by setting

$$\begin{aligned} k_u &= \rho \csc\left(\alpha\right) \\ \alpha &= \cot^{-1}\left(-2\gamma H\tau\right). \end{aligned} \tag{4.2.2}$$

The following chirp multiplier is also required to complete the square,

$$f_\alpha(\rho) = \sqrt{1 - i\cot(\alpha)}\exp\left(i\pi\cot(\alpha)\rho^2\right) s\left(k(t)\right) \tag{4.2.3}$$

With these substitutions and post-multipliers, the signal takes the form of the FrFT,

$$f_\alpha(\rho) = \sqrt{1 - i\cot(\alpha)}\exp\left(i\pi\cot(\alpha)\rho^2\right)\int \exp\left(-i\pi\left(2\rho u\csc(\alpha) + u^2\cot(\alpha)\right)\right) f(u)du \tag{4.2.4}$$

where ρ varies over time during data acquisition while linear SEMs are being applied. Care must be taken to evaluate the chirp function at the values of k where data is taken. The modified signal now represents the α^{th} order FrFT. As such, the object can be reconstructed

without quadratic phase via the inverse operation: the FrFT of order $-\alpha$.

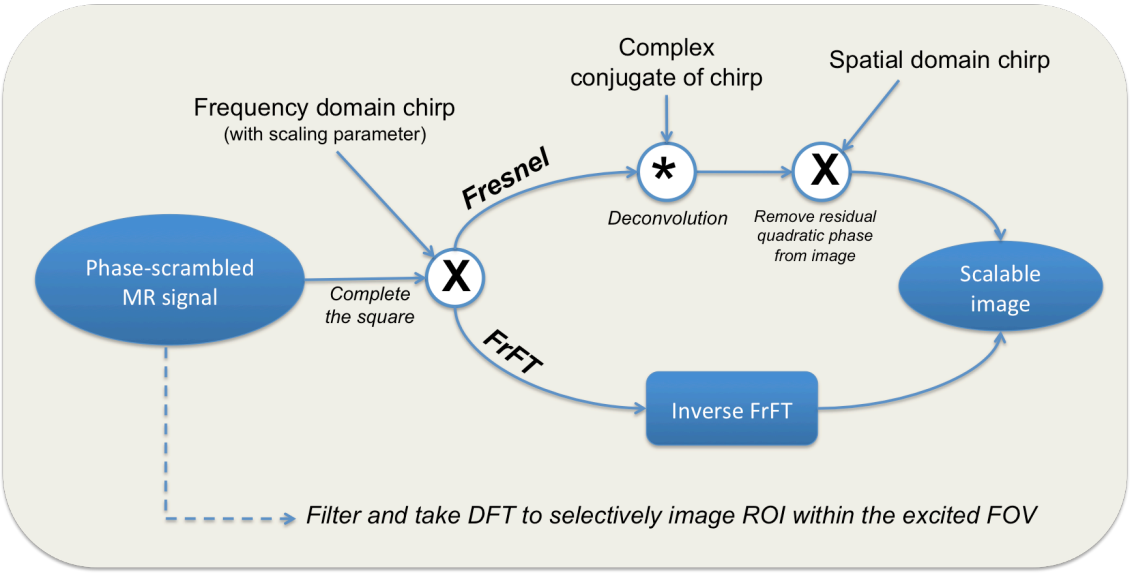

Figure 4.3: The Fresnel transform and FrFT can both be used for scalable reconstruction. The Fresnel transform is equivalent to the FrFT when combined with chirp multiplication of suitable strength.

The equivalent steps for FrFT and Fresnel image scaling are shown in Fig. 4.3. If the acquired signal is multiplied by the appropriate chirp function, the signal equation is recast in the form of a FrFT or a convolution with a Fresnel chirp kernel. In the Fresnel approach, image reconstruction proceeds via a deconvolution with the Fresnel kernel, realized either as convolution with the complex conjugate of the kernel (109) or via the inverse filtering method in the Fourier domain (108).

In the present work, we perform the same operations using the FrFT, exploiting the fact that LCT matrices can be decomposed into multiple elementary LCT operations. The FrFT of order α can be decomposed into three LCT matrices (104): a chirp multiplication, a Fresnel transform, and a second chirp multiplication,

$$\begin{bmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \cot\alpha - \csc\alpha & 1 \end{bmatrix} \begin{bmatrix} 1 & \sin\alpha \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \cot\alpha - \csc\alpha & 1 \end{bmatrix} \tag{4.2.5}$$

In the optical analogy, the phase scrambling of MRI data is equivalent to passage through a thin lens followed by free space propagation followed by passage through a second thin lens. This combination of is described more succinctly in a single operation using the FrFT. To summarize, the phase-scrambled signal in Eq. (4.2.1) can be put into the form of the FrFT through chirp multiplication, shown in Eq. (4.2.3), where α and ρ are set (Eq. (4.2.2)) to match the definition of the FrFT in Eq. (4.1.4).

Care must be taken to evaluate the chirp function at values of ρ which correspond to the values of k in the signal $s\left(k\left(t\right)\right)$, i.e., the chirp must be evaluated at $k = \csc(\alpha)\rho$. The image is then reconstructed using the FrFT of angle α. To scale the image by factor ζ, the chirp multiplier in Eq. (4.2.3) is changed to $\exp\left(i\pi\zeta\cot(\alpha)\rho^2\right)$. To reconstruct a properly "focused" image, the FrFT with the new angle

$$\alpha_s = \cot^{-1}(\zeta\cot(\alpha)) \tag{4.2.6}$$

must be performed. Setting $\zeta > 1$ results in less quadratic phase across $f(\rho)$, essentially "tricking" the FrFT into reconstructing the object within a larger FOV. This model is readily generalized to two dimensions.

4.2.2 Experiments

Data were acquired on a Siemens Trio 3T scanner equipped with a high-power 12 cm SEM insert coil capable of generating the Z^2 spherical harmonic, $G\left(x,y,z\right) = z^2\frac{1}{2}\left(x^2+y^2\right)$, providing the desired quadratic variation in the transverse plane. Images of size 64×64 were acquired with a uniform-sensitivity "birdcage" coil (112) using a gradient echo sequence (TE = 6 ms; aliased FOV = 3 cm). For the scaled images shown below, a quadratic field of 3.5 KHz/cm^2 was pulsed for 550 μs prior to readout, corresponding to the FrFT

order of $\alpha = 0.34\pi$. The FrFT was performed using a fast Matlab code made available by (113).

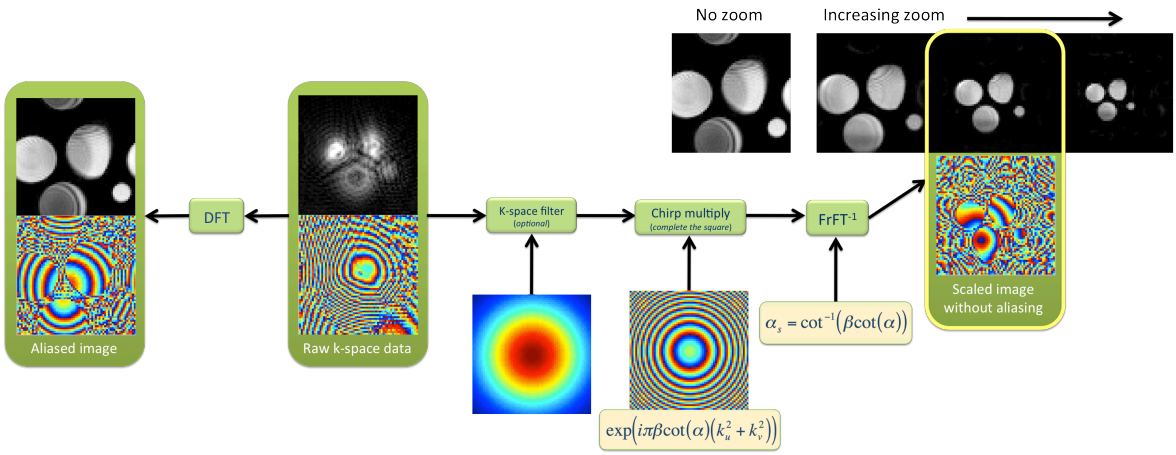

Figure 4.4: Illustration of FrFT reconstruction from phase scrambled data. Without lensing, two of the water tubes extend beyond the FOV and alias back into view on the other side of the image. When quadratic phase modulation is applied during the experiment, the image may be scaled using the FrFT to bring all objects within the FOV. The amount of scaling depends on the order of the FrFT used. The scaling range is limited by the amount of applied phase scrambling.

This work demonstrates the equivalence of Fresnel deconvolution and FrFT processing for obtaining scalable unaliased images from phase-scrambled data. The concept is proved with scalable reconstructions of short-TE data obtained using a dedicated high-power quadratic gradient insert.

Interestingly, if the phase scrambled data are reconstructed with the FT instead of the FrFT, a localized image results. This effect was first demonstrated by (114) and represents and alternative but potentially more useful application of quadratic phase scrambling SEMs. Gradient localization, known as "GradLoc", is discussed further in the next section.

4.3 Phase scrambled FOV localization (GradLoc)

Spatial encoding magnetic fields (SEMs) with quadratic or hyperbolic shapes have been proposed for imaging using multipolar pairs of fields (10), localizing NMR spectra (115), and exciting curvilinear regions using RF pulses (116, 117). Nonlinear SEMs have also been used to traverse k-space in a voxel-specific way (88) for parallel RF transmission schemes (55), which are useful for shape selection excitation and for improving B_1 homogeneity. Quadratic phase RF pulses have also been used to selectively refocus thin layers of a slice (118). More recently, phase scrambling has been shown to enable gradient localization (GradLoc) of signals from a target ROI within the larger FOV (114). Localization is achieved by applying quadratic modulation prior to readout in a conventional Cartesian pulse sequence. This is the same approach used in the preceding section on scalable-FOV imaging, except in this case the aim is to spoil signals residing outside the target ROI. The SEM pulse moments are correspondingly stronger in this case. The quadratic SEM pulse is governed by the relation $\gamma H\tau = k_{max}/\Delta_{ROI}$, where Δ_{ROI} is the size of the target ROI, τ is the pulse length, and γH is the SEM field strength.

The SEM pulse in GradLoc has the effect of convolving the object's spectrum with a Gaussian kernel, spreading out the spectral energy and causing the spectrum to resemble the object in a 1-to-1 mapping (114). As pointed out in (114), quadratic fields are well suited to GradLoc because the k-space shifts are linear and proportional to the applied SEM pulse moment. As shown in Eq. (2.3.3), nonlinear SEMs apply a spatially-variable k-space shift vector to each voxel's k-space trajectory. If the applied nonlinear pulse moment is large enough, a given voxel's k-space trajectory may be translated outside the acquired k-space region. This shifts the voxel's echo so that it occurs outside the readout window so that refocusing does not occur, spoiling most of the voxel's contribution to the acquired signal (114). Seen another way, the phase scrambling SEM shape is scanned across the desired ROI by the linear SEMs as the pulse sequence plays, dephasing spins that reside outside this

region. One drawback of the method is that at least part of a spoiled voxel's echo still resides in the readout window, contributing high spatial frequency information for the voxel that can cause artifacts.

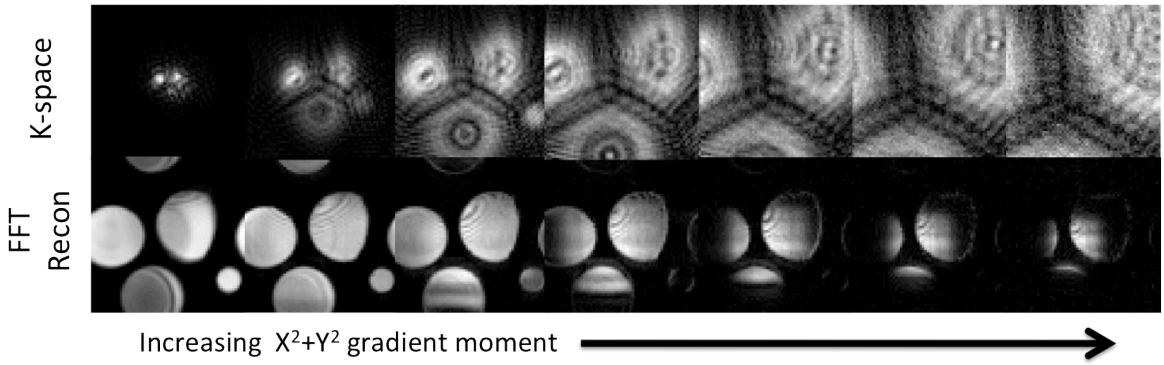

Figure 4.5: As the phase-scrambling SEM moment grows, signal from progressively smaller regions can be acquired without aliasing. With phase-scrambling, the k-space data take on a coarse resemblance to the object.

GradLoc permits reduced FOV imaging, thus enabling (a) faster acquisition of the desired ROI at equivalent resolution or (b) higher resolution acquisition of the ROI in equivalent time. The approach has been demonstrated experimentally (114) using second-order PatLoc fields that vary only in two dimensions, i.e., hyperbolic fields that vary as $2xy$ and $x^2 - y^2$. The primary aim of the present work is to extend the utility of GradLoc to 3-D SEMs such as the Z^2 spherical harmonic, which varies as $G_{Z^2}(x, y, z) = z^2 - \frac{1}{2}(x^2 + y^2)$. The Z^2 has advantageous properties from a coil designer's perspective and has already found promising applications for spatial encoding, including those presented previously in this thesis as well as emerging methods from other research groups (119). The quadratic axial-plane variation of this SEM is ideally suited to performing GradLoc. This is clearly illustrated by the images in Fig. 4.5, which are based on the same data used to reconstruct the scaled images in the previous section. The only difference is that the images in this case are reconstructed via FFT instead of FrFT[1].

[1]For the GradLoc images in Fig. 4.5 the data are also multiplied in k-space by a Gaussian window to limit Gibbs ringing arising due to the k-space truncation cased by the quadratic SEM.

The drawback to using the Z^2 field for GradLoc is that the SEM varies twice as strongly in the z-direction as it does within the axial-plane. Because the phase scrambling pulse is not refocused, this through-slice field causes extensive dephasing, similar to that shown in Fig. 3.12 in relation to field mapping of the Z^2 SEM. This dephasing seriously degrades the acquired signal even for thin slices at small z-offsets or typical slice thicknesses with no offset (Fig. 4.6).

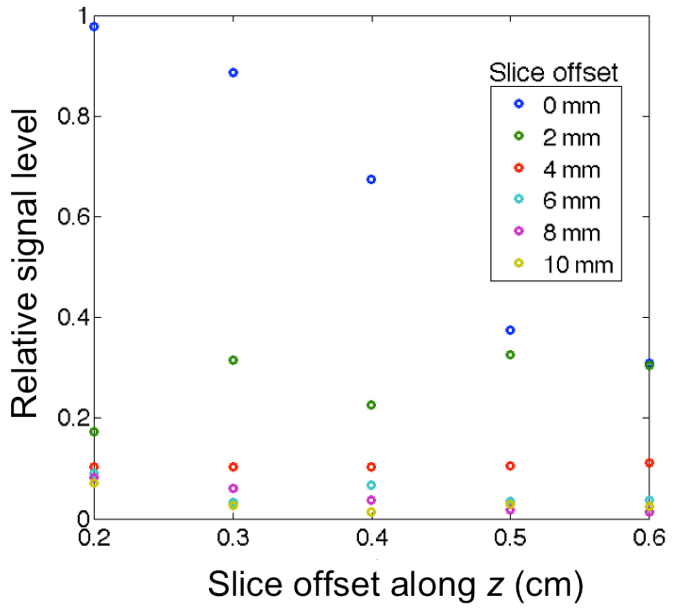

Figure 4.6: Signal level as a function of slice offset and thickness. The Z^2 causes dephasing in the through-slice direction, severely reducing the signal level at offset of only z=2 mm. The calculation shown assumes a phase scrambling SEM sufficient to acquire signal from an ROI measuring $\frac{1}{2}$ of a 10 cm FOV in each dimension.

The primary goal of the present study is to investigate the use of a quadratic-phase RF pulse in combination with a conventional slice select SEM to precompensate for the through-slice phase applied by the Z^2 SEM. As a secondary goal, we investigate the use of SENSE parallel imaging in combination with GradLoc to image target ROIs at high acceleration factors.

A quadratic phase RF pulse is designed using the low flip angle approximation. The target slice profile is assumed to be a rectangular slab with quadratic phase equivalent to that applied by the phase-scrambling SEM, but with opposite sign (Fig. 4.7). The Fourier

transform is used to calculate an RF pulse that will produce approximately this desired slice profile. GradLoc images of target ROIs are acquired in the axial-plane using the Z^2 SEM coil. The 8-channel transmit-receive array (Fig. 3.5) is nested within the bore of the SEM insert coil and used for RF excitation and reception. Coil profiles in the target ROI are obtained from fully-sampled GradLoc images based on the adaptive method in (41) (Fig. 4.8).

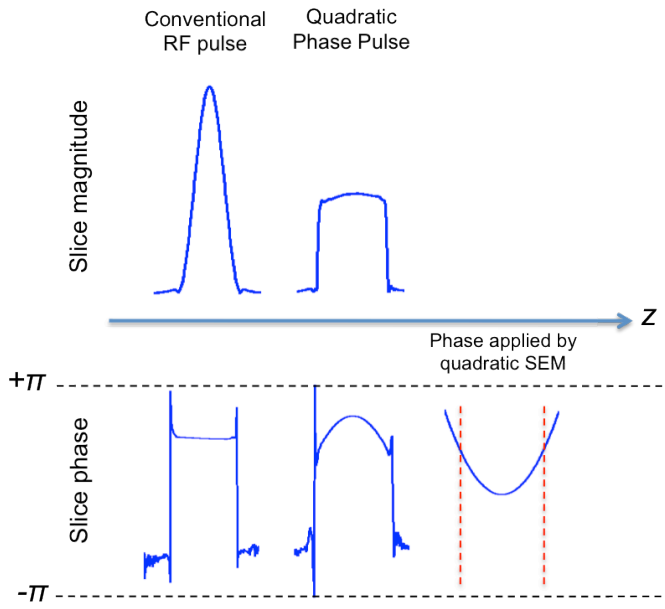

Figure 4.7: Comparison of measured slice profiles acquired using a conventional RF pulse (left) and quadratic-phase RF pulse (center). The quadratic-phase pulse applies a through-slice phase (simulated at right) to precompensate the phase that will be applied by the Z2 SEM prior to readout. The edges of the slice are indicated by the red dotted line. Example shown is a 3 mm slice at z=0.

The conventional GradLoc pulse sequence from Fig. 4.2 is modified to include the quadratic phase RF pulse (Fig. 4.9). For offset slices, the amplitude of the slice select SEM is increased in order to shift the center of the quadratic phase profile produced by the RF pulse. Similarly, linear SEMs pulsed prior to readout may be used to shift the GradLoc ROI within the FOV. To test the efficacy of SENSE reconstruction, GradLoc data are undersampled in k-space and the resulting images are unaliased using SENSE (1).

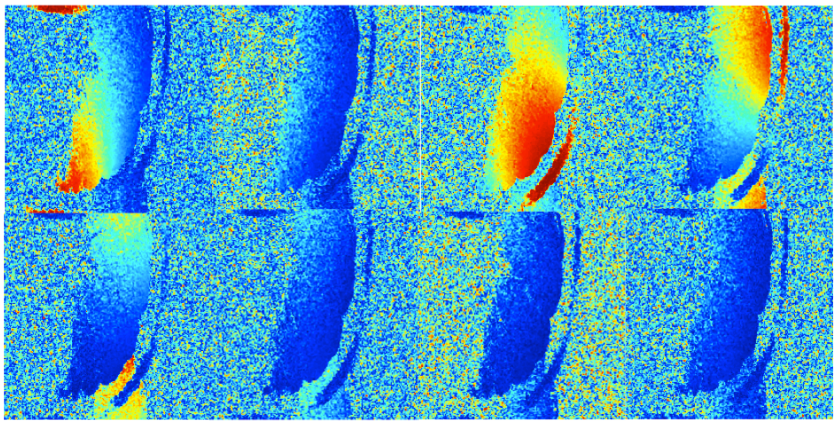

Figure 4.8: Coil profiles are obtained from fully-sampled GradLoc images in the 5 cm ROI for an orange phantom. Some residual aliasing is present due to the highly variable flip angle of the transmit-receive RF array used.

GradLoc images with 4 mm slice thickness are successfully acquired (Fig. 4.10) using the Z2 SEM along with the quadratic-phase RF pulse for slice-phase precompensation. SENSE reconstructions are performed in a GradLoc ROI (Fig. 4.11) in one quarter of the original FOV, yielding a net acceleration factor of $2R$, where R is the k-space undersampling factor. SENSE image quality degrades more quickly than would be expected based on the number of available receive RF coils. This is likely due to the fact that only 4 out of the 8 RF coils have significant spatial variation and B1 sensitivity within the GradLoc ROI.

As already demonstrated previous to this work, GradLoc is a promising alternative to Transmit SENSE for achieving localized acquisitions, circumventing the complexity of Transmit SENSE and the safety limits on RF power deposition that arise when using shape-selective RF pulses. The present work shows that parallel imaging may be combined with GradLoc to image local ROIs with high net acceleration factors as compared with imaging the full FOV. We also show that GradLoc may be performed using 3-D SEMs if care is taken to compensate for through-plane dephasing using a tailored RF pulse. Future efforts will seek to refine the RF pulse to more closely match the phase imparted by the Z^2 SEM. Auto-calibrated reconstructions such as GRAPPA (43) will also be explored as alternatives to the use of fully-sampled GradLoc images for obtaining SENSE coil profiles. The Z^2

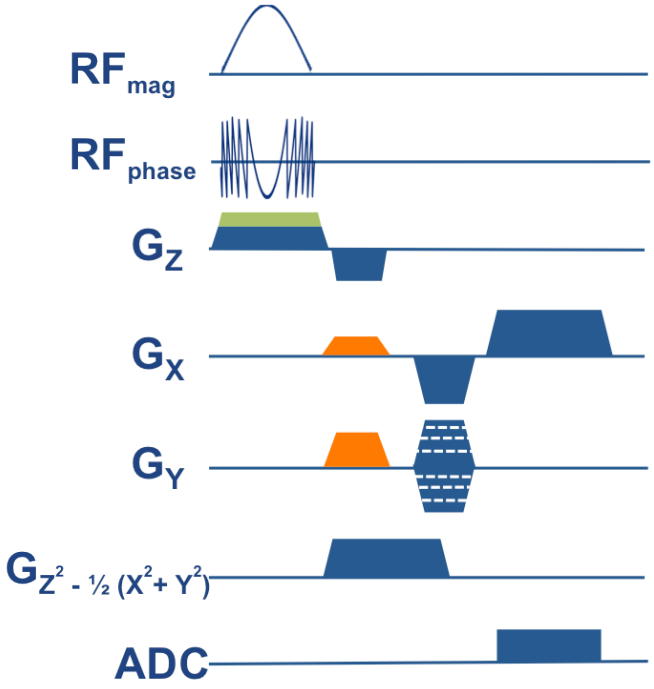

Figure 4.9: GradLoc pulse sequence based on a Cartesian gradient echo sequence. A quadratic phase RF pulse is used to precompensate the through-slice phase applied by the Z^2 phase scrambling SEM. The GradLoc window size is determine by the quadratic SEM pulse moment, while the offset depends on the orange lobes of the linear SEMs, which serve to shift the center of the quadratic bowl. The additional green lobe of the slice select SEM shifts the quadratic phase along Z if a slice offset is used (similar to the illustration in Fig. 2.4).

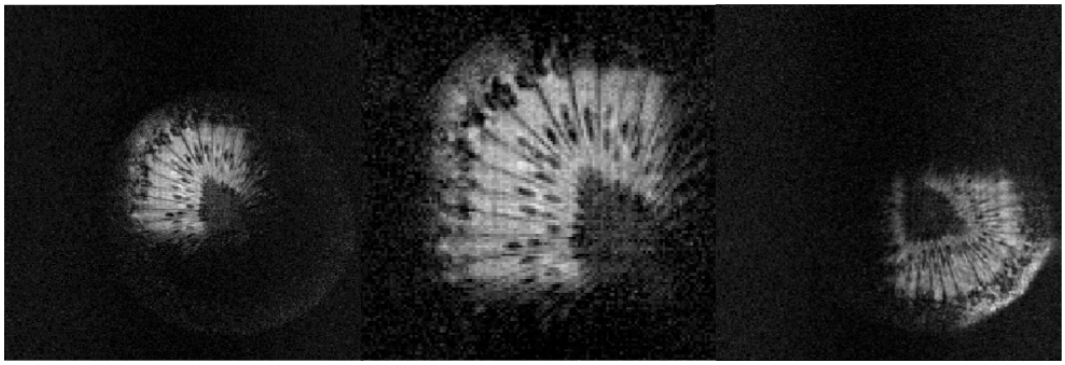

Figure 4.10: GradLoc images of a kiwi with a 7 cm FOV (left), a 3.5 cm FOV with the same voxel size (center), and an ROI offset using extra linear SEM pulses prior to readout.

SEM will also be tested for GradLoc imaging in the sagittal and coronal planes, albeit with rectangular instead of square ROIs.

4.4 Compressed sensing with quadratic modulation

Because quadratic phase modulation in the image domain becomes a convolution of the object's spectrum with a Gaussian kernel, phase scrambling has the effect of making the spectrum less sparse. This has been shown to improve image reconstruction from randomly-undersampled k-space data using a sparsity constraint (120). By spreading out the object's energy in k-space, phase scrambling reduces the mutual coherence between the sensing basis (k-space basis) and a given sparse basis (often a wavelet basis). In principle, this reduces the number of k-space coefficients required for accurate recovery of the image. The Yale MRRC Z^2 insert coil provides an excellent platform for exploring this approach to compressed sensing. Previously, only single-coil experiments have been shown. In future work, we hope to explore this topic using multi-coil data acquired using the Z^2 insert coil.

4.5 Variable-order FrFT

Previous sections in this chapter have discussed the utility of adding a quadratic phase scrambling pulse to a Cartesian sequence prior to readout. This phase modulation was found to enable either scalable-FOV imaging or selective acquisition of a target ROI, depending on the amount of applied quadratic phase and the chosen reconstruction method. In this section, the Fractional Fourier transform (FrFT) is discussed in relation to quadratic pulses played during readout - in other words, O-Space imaging. Iterative O-Space reconstructions based on the full encoding matrix have been shown to be sensitive to systematic errors in the encoding matrix arising from B_0 drift, local B_0 inhomogeneity, and field calibration errors. An alternative reconstruction based on the variable-order FrFT VO-FrFT) holds the potential to be less sensitive to these systematic errors, translating the errors into geometric distortions instead of phase cancelations. This robustness arises from the fact that the VO-FrFT operates as a transform instead of a matrix solver. The promise as well as the limitations of the VO-FrFT will be demonstrated on empirical data.

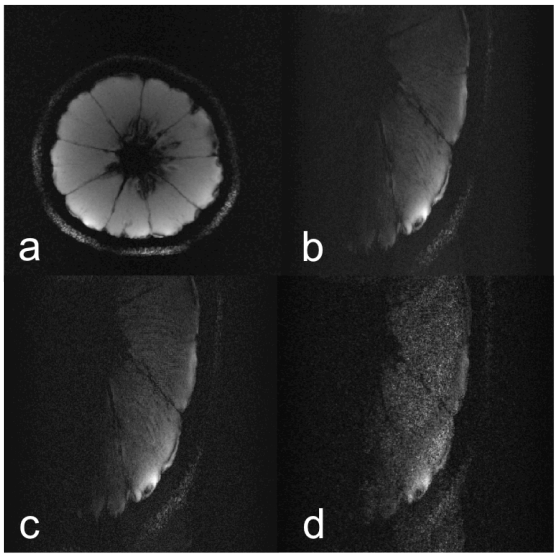

Figure 4.11: (a) Reference image of an orange (256×256, 10 cm FOV) and (b) fully sampled GradLoc image (128×128, 5 cm FOV). SENSE parallel reconstructions of undersampled GradLoc images are shown with R=2 (c) and R=4 images (d), effecting net acceleration factors of 4 and 8, respectively.

4.5.1 Analytical model for O-Space signal equation

The FrFT has recently been used as an analytical approach to correcting B_0 inhomogeneity during image reconstruction (121). Often local inhomogeneities can be approximated as quadratic functions in two dimensions. The quadratic inhomogeneity causes spin phase

evolution that can be modeled as an FrFT whose order changes at each time point following RF excitation of the spins. The SEM used for spatial encoding in O-Space imaging can be modeled as a very strong inhomogeneity that is switched on and off during a pulse sequence. This invites exploration of the VO-FrFT as a way to perform O-Space image reconstruction (122).

In previous chapters, "brute force" iterative matrix solvers were used for O-Space reconstruction. Because they require explicit calculation of an encoding matrix of size $n_K n_c \times N^2$, iterative approaches such as conjugate gradients and the Kaczmarz method are computationally demanding. They are also highly sensitive to inconsistency in the matrix equations arising from off-resonance effects or systematic errors in calibration. The FrFT offers the prospect of a faster reconstruction with greater robustness to calibration errors.

Recall that the FrFT integral operator, $FrFT_\alpha$, is defined in one dimension as:

$$FrFT_\alpha \{f\}(\rho) = \sqrt{1 - i\cot(\alpha)}\, e^{i\pi \cot(\alpha)\rho^2} \int e^{-i2\pi \csc(\alpha)\rho u} e^{i\pi \cot(\alpha)u^2} f(u)du \quad (4.5.1)$$

To cast the acquired O-Space signal in the form of the FrFT expressed in 4.5.1, let $k(t)$ and $l(t)$ denote the accumulated SEM pulse moments of the linear, $\gamma G(t)$, and quadratic, $\gamma H(t)$, fields, respectively:

$$\begin{aligned} k(t) &= -\gamma \int_0^{t'} G(t')dt' \\ l(t) &= -\gamma \int_0^{t'} H(t')dt' \end{aligned} \quad (4.5.2)$$

If we then use the following mapping between gradient trajectories and FrFT parameters (78),

$$\begin{aligned} \alpha(t) &= \cot^{-1}\left(-2l(t)\right) \\ \rho(t) &= \frac{k(t)}{\csc(\alpha)} = \frac{k(t)}{\sqrt{1+4l^2(t)}} \end{aligned} \tag{4.5.3}$$

then the O-Space signal equation at time point t_n begins to resemble the FrFT integral:

$$\begin{aligned} s(t_n) &= \int \exp\left(i\pi\left(-2k(t_n)u - 2l(t_n)u^2\right)\right) f(u)du \\ &= \int \exp\left(i\pi\left(u^2\cot(\alpha_n) - 2u\rho\csc(\alpha_n)\right)\right) f(u)du \end{aligned} \tag{4.5.4}$$

Off-resonance effects can be described by an additional factor of $\exp(-i2\pi\Delta B_0 t)$ inside the integral. Multiplying $s(t)$ by a chirp function and constant then "completes the square" of the FrFT integral kernel[1]:

$$f_\alpha(\rho_n) = \sqrt{1-i\cot(\alpha_n)}\exp\left(i\pi\rho_n^2\cot(\alpha_n)\right)s(t_n) \tag{4.5.5}$$

Because the quadratic SEM plays during readout, α varies on a point-by-point basis, preventing reconstruction from being performed by a single FrFT of $f_\alpha\left(\rho_n\right)$. Instead, reconstruction is attempted using the "variable-order FrFT" (VO-FrFT), an approach used in (78) to reconstruct Cartesian images in the presence of a quadratically-shaped field inhomogeneity. This approach exploits the fact that the FrFT is a unitary operator by applying the conjugate of the forward integral kernel at each point in $s(t_n)$. We modify this formalism by adding a k-space density compensation function (DCF) to make it compatible with the radial trajectory traversed by the linear SEMs in O-Space imaging,

$$\hat{f}(u) = \sum_{n=1}^{N_s} DCF(\rho_n)s(t_n)|\csc(\alpha_n)|\exp\left(i\pi\left(u^2\cot(\alpha_n) - 2u\rho_n\cot(\alpha_n)\right)\right). \tag{4.5.6}$$

[1]The extra phase factor $\sqrt{1-i\cot(\alpha_n)}$ does not appear explicitly in the LCT precursor transform, but it is commonly used in the optics literature, so we adopt it here out of convention.

The DCF varies inversely with the density of the k-space traversed by the linear SEMs, weighting data less in the center of k-space and more at the edges, where there are larger gaps between points: $DCF(\rho_n) = k_n^2$. In this point-by-point variable-order implementation, the FrFT is essentially a discrete Fourier transform modified to include a quadratic integral kernel. The above formalism is easily generalized to two dimensions. Moreover, using the non-separable linear canonical transform as described in (123, 124), we expect that FrFT reconstruction can be applied to the case of PatLoc fields with hyperbolic rather than quadratic exponentials in the encoding integral kernel.

As observed in (78), this approach is related to the method of conjugate phase (125, 126), which has been used to correct for the effects of higher-order off-resonant fields in conventional MRI. When field maps of inhomogeneous regions are available, the off-resonant components can be compensated during reconstruction by modifying the usual discrete Fourier transform to also include a factor with the conjugate of the phase evolution caused by the off-resonant fields. The primary differences between this approach and the approaches are (a) the inclusion of the DCF in the FrFT formalism and (b) the amplitude of the higher-order components being "corrected". In off-resonance correction, the perturbing higher-order fields are much smaller than the SEMs, while in O-Space imaging the "perturbing" quadratic field is comparable in amplitude to the linear SEMs.

4.5.2 VO-FrFT O-Space Simulations

To investigate the utility of the VO-FrFT for O-Space reconstruction, simulations of O-Space encoded acquisitions are performed using a single uniform RF coil and a numerical reference phantom designed to highlight geometric distortions. Additive white noise is omitted in order to focus on the effects of calibration errors on the images. Reconstructions of size 256×256 are shown based on simulated radial and O-Space echo signals for the case of a single, uniform RF coil, 512 readout samples, and 256 readout "spokes". Off-

resonance effects and errors in the calibrated SEM field, $\mathcal{F}H(u)$, are compared for ART and FrFT reconstructions.

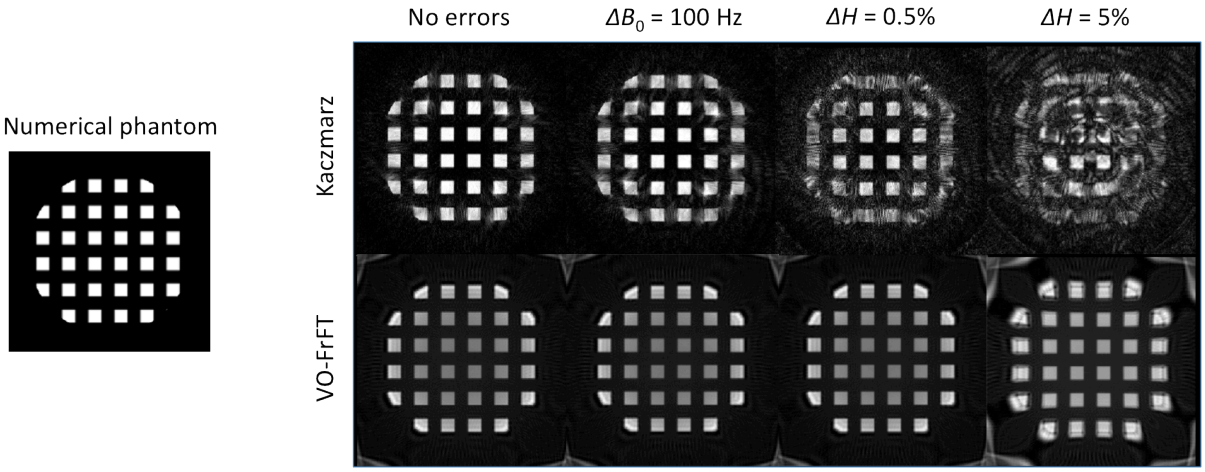

Figure 4.12: VO-FrFT reconstructions degrade much more gracefully than Kaczmarz reconstructions in the presence of global off-resonance ΔB_0 and quadratic field strength error ΔH. These systematic errors cause geometric distortion in the FrFT reconstructions, in contrast to the destructive phase cancelation of the Kaczmarz images.

In Fig. 4.12, FrFT reconstructions show markedly reduced sensitivity to off-resonance effects and errors in SEM calibration, each of which cause deleterious phase cancelation in the Kaczmarz reconstructions[1]. FrFT reconstructions degrade gracefully, showing geometric distortions at the periphery where the quadratic field is strongest. Preliminary comparisons show a roughly 2-fold improvement in reconstruction times using the FrFT as compared with the Kaczmarz algorithm, an advantage that is expected to grow as more efficient implementations of the variable-order FrFT are developed.

Images reconstructed from acquired data do not perform as well as the simulations (Fig. 4.13). Although the images exhibit geometric distortion instead of phase cancelation, the distortion could not be eliminated by tweaking the reconstruction parameters to account for a possible calibration error. The reconstructions show a distinct "double-ghost" pattern

[1]When there are errors in the amplitude of the quadratic SEM, Kaczmarz reconstructions degrade more severely near the periphery, where the quadratic field is strongest.

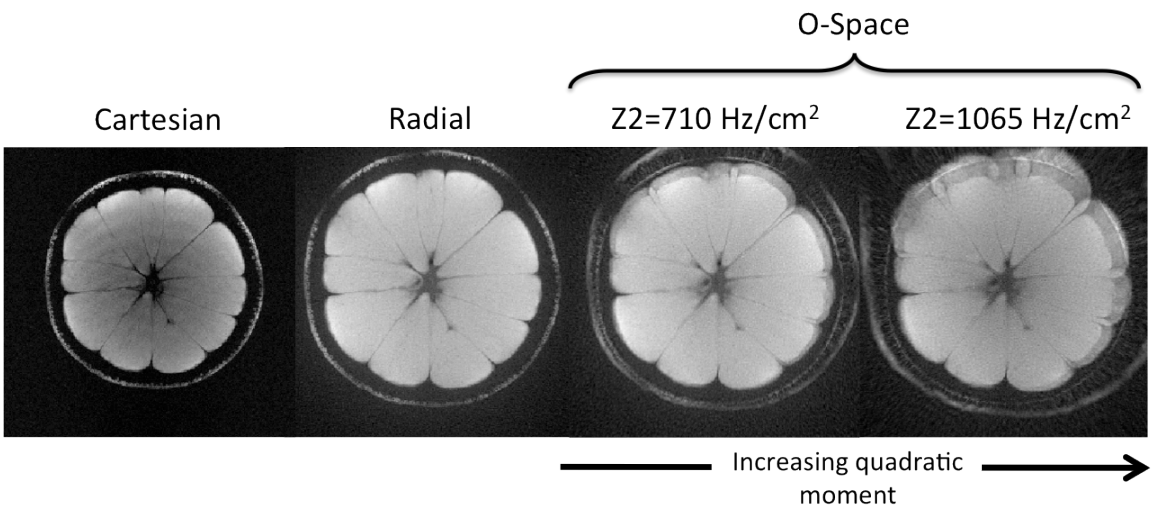

Figure 4.13: O-Space reconstructions using the VO-FrFT compared with a radial image reconstructed using the DFT and a Cartesian image reconstructed via the FFT. As the quadratic SEM strength increases, the VO-FrFT becomes more non-unitary, causing severe ghosting and distortion.

that grows more severe in regions where the quadratic field is strongest[2]. The ghosting also becomes more pronounced as the Z^2 SEM increases in amplitude. It is very similar to the ghosting observed in section 3.6.2 when incorrectly-scaled Z^2 SEMs are used to populate the encoding matrix during image reconstruction (modeling systematic error). Incorporating an additional factor into Eq. (4.5.6) to compensate for potential off-resonance effects does not diminish the appearance of the ghosts.

4.5.3 Non-unitarity of the VO-FrFT

At first glance, it seems possible that the ghosting arose from an error in the SEM calibration that had escaped detection. But on closer examination, fault is found to lie with the variable-order FrFT itself. Because the order of the operator changes at every data point, the VO-FrFT ceases to be a unitary linear operator. Strictly speaking, this implies that the inverse of the VO-FrFT is no longer equal to the complex conjugate of the forward transform,

[2]The double-ghost has been observed in simulated VO-FrFT reconstructions as well, suggesting that is caused by the non-unitarity of the VO-FrFT rather than by experimental calibration errors.

invalidating Eq. 4.5.6. The loss of unitary is illustrated in Fig. 4.14, where plots of $E^H E$ and FE clearly differ, with E and F defined as the encoding and reconstruction matrices for a one-dimensional O-Space test case[1]. The non-unitarity grows more pronounced for strong Z^2 SEMs, explaining why the ghosting in the Fig. 4.13 images grows more pronounced with increasing Z^2 encoding amplitude.

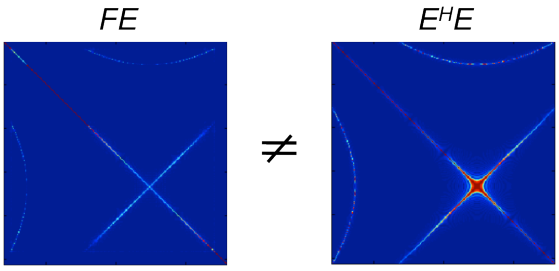

Figure 4.14: For strong Z^2 SEMs, the O-Space encoding matrix becomes increasingly non-unitary, shown here for the 1-D case. The SEM strength simulated is 6.8 KHz/cm^2, corresponding to an FrFT order of $\alpha = 0.29$ at the end of the readout. Note that FE does not equal $E^H E$, violating the definition of unitarity.

Although these results are not entirely promising, it remains an open question whether the VO-FrFT could be modified to permit ghost-free O-Space reconstructions. In particular, the density compensation function used in Eq. (4.5.6) invites closer scrutiny. It was chosen naively based only on the density of points in the underlying linear k-space domain, ignoring the extra complexity created by the evolution of spins under the quadratic field. It is not clear how the DCF could be generalized to account for the density of both the linear and nonlinear "k-space trajectories".

4.5.4 Future challenges

Should the problem of non-unitarity problem be surmounted, the next task would be to generalize the VO-FrFT to the multi-coil case. By analogy to GRAPPA (43), auto-calibrated multi-coil reconstruction might be possible using the frequency shift rule for the FrFT, Eq. (4.5.7), a direct generalization of the analogous rule for the Fourier transform,

[1] For a small one-dimensional test case, it is feasible to explicitly calculate the least-squares reconstruction matrix, F, via the pseudoinverse.

$$\begin{aligned} f_\alpha\left(\rho - \xi\sin\alpha\right) \quad &\Longleftrightarrow \quad \exp\left(i\pi\xi\sin\alpha\cos\alpha\right)\exp\left(i2\pi\xi u\left(1-\cos\alpha\right)\right)f(u) \\ &= \quad \eta_{\alpha,\xi}\exp\left(i2\pi\left(1-\cos\alpha\right)\right)f(u) \end{aligned} \tag{4.5.7}$$

This property may permit auto-calibrated O-Space reconstructions by using multi-coil data to fill in missing lines of spectral information, a natural extension of GRAPPA parallel imaging into the fractional Fourier domain. In addition to providing auto-calibration of coil profiles, this approach is expected to introduce additional time savings as compared with the Kaczmarz method, which must backproject each coil's data separately[1].

[1] As discussed in Chapter 3, this limitation can be overcome by implementing string-averaged parallel Kaczmarz on multiple processors or GPU cores.

5

Outlook

5.1 Summary

This thesis represents an initial investigation of the possibilities for spatial encoding offered by the Yale MRRC's high-powered 12 cm Z^2 insert coil, a unique resource in the emerging field of nonlinear spatial encoding. A new approach called O-Space imaging is introduced, iterative image reconstruction is described, experimental obstacles are overcome, and images are compared to radial projection imaging and standard Cartesian imaging. An alternative, analytical reconstruction based on the Fractional Fourier transform is then discussed, but its utility for O-Space imaging appears to be limited because the order of the transform changes continuously during signal readout. The thesis then turns its attention to two applications of phase scrambling with quadratic pulses. The first application, "lensing" of scalable-FOV MRI, has been demonstrated elsewhere via Fresnel deconvolution but is performed here for the first time with the Fractional Fourier transform. Phase scrambled data are then used to create GradLoc images of target ROIs without aliasing using standard Fourier reconstruction. GradLoc is extended here through (a) the use of a quadratic phase RF pulse to precompensate for through-slice dephasing caused by the 3-D SEM and (b) combination with SENSE parallel imaging for highly accelerated GradLoc imaging.

Several potential directions for future work have been proposed, including the following:

- Incorporation of the L_1 norm into the Kaczmarz algorithm for improved O-Space image quality, particularly for highly accelerated acquisitions;
- Exploration of highly accelerated, multi-coil, phase scrambled reconstructions, taking advantage of the quadratic field's ability to spread out k-space energy and reduce coherence between the data domain and the sparse (wavelet) domain;
- Refinements of the quadratic-phase RF pulse used for GradLoc imaging in order to handle stronger Z^2 pulse moments and to perform more accurate precompensation of the slice phase profile;
- More detailed exploration of the noise and resolution properties of O-Space imaging;
- Continued exploration of the variable-order FrFT as a reconstruction method for O-Space imaging. In particular, redefinition and refinement of the proposed density compensation function.

Most significantly, future experiments at Yale may involve imaging with a human head insert equipped with multiple second-order spherical harmonic SEMs, bringing O-Space imaging and other methods much closer to potential clinical and neuroscientific application.

5.2 Emerging methods for nonlinear encoding

Since the introduction of PatLoc and O-Space imaging, a number of other promising encoding strategies have been proposed that exploit the properties of nonlinear SEMs. In the space below I offer a brief overview of these methods along with speculation on their potential utility.

5.2.1 COGNAC

In keeping with a venerable naming tradition that includes the likes of GRAPPA and CAIPIRINHA, researchers from Würzburg have introduced an approach called "Combinations of arbitrary gradient encoding fields using SPACE-RIP", or COGNAC (119). This method employs linear SEMs along with the Z^2 SEM to perform axial imaging. Unlike O-Space imaging, however, one set of fields is played as the readout and the other as a phase encoding field. This approach allows fast reconstruction of the data in polar coordinates using a hybrid image-domain/k-space domain method called SPACE-RIP (44). The resulting image is then interpolated from polar to Cartesian coordinates to create the final image. This reconstruction enjoys the advantage of being (a) fast and (b) comparatively robust to systematic calibration errors. Several variants on COGNAC imaging have been proposed, including:

- A constant Z^2 SEM for readout with the Y SEM played as the phase encode;
- A constant Z^2 SEM for readout while the X and Y together play a radially-varying phase encode;
- A sinusoidal X and Y SEMs to trace a circle in k-space while the Z^2 is played as a phase encode to provide high-resolution encoding in the radial direction. Because the Z^2 pulse is not refocused, signal dephasing occurs in the central region of the FOV, suppressing unwanted signal from this region if outer-volume selection if desired.

The disadvantage that each of these COGNAC approaches provides no encoding at the center of the FOV. However, when high resolution is desired at the periphery, or signal from the central part of the FOV needs to be suppressed, COGNAC provides a useful and efficient approach.

5.2.2 4D-RIO

A new PatLoc encoding strategy, 4D-RIO, has recently been developed (15). This encoding scheme also uses quadrupolar encoding fields as in previous PatLoc experiments (61),but couples these fields with the linear SEMs, thus combining features from previous PatLoc encoding strategies and O-Space imaging. 4D-RIO is a projection-class method that plays reciprocating radial k-space trajectories on each pair of SEMs, linear and quadrupolar. Each pair traverses a radial "k-space" trajectory, but they are out of phase with one another, such that one trajectory reaches the center of k-space as the other reaches the outer extents. In this way, the spins do not all refocus at the same time. 4D-RIO holds the potential to outperform previous PatLoc and O-Space imaging methods in undersampled acquisitions.

While calibrating four SEMs is more challenging than calibrating three SEMs (one of which plays the same pulse for every readout), 4D-RIO calibration benefits from the ability to play only one pair of SEMs or the other and reconstruct a test image. Comparing the test image from the linear SEMs with that of the PatLoc SEMs, the PatLoc field map can be tweaked (rotated and scaled) until the two reconstructed images overlay one another. This sleight-of-hand has enabled high-quality 4D-RIO images on phantoms, with *in vivo* 4D-RIO expected to be achieved in the near future.

5.2.3 Null space imaging

O-Space imaging is just one of a multitude of potential nonlinear projection encoding schemes. Investigators have begun searching for ways to obtain and evaluate the optimal SEM shapes to complement a particular surface coil array. One approach (127) uses the singular value decomposition of the matrix holding the coil profiles to find the orthogonal modes spanning the space of the profiles. Pairs of frequency/phase encoding PatLoc SEMs are then chosen so as to approximate these modes. Similarly, in an approach known as Null Space Imaging (128), nonlinear SEMs are chosen to approximate functions residing in the null space of

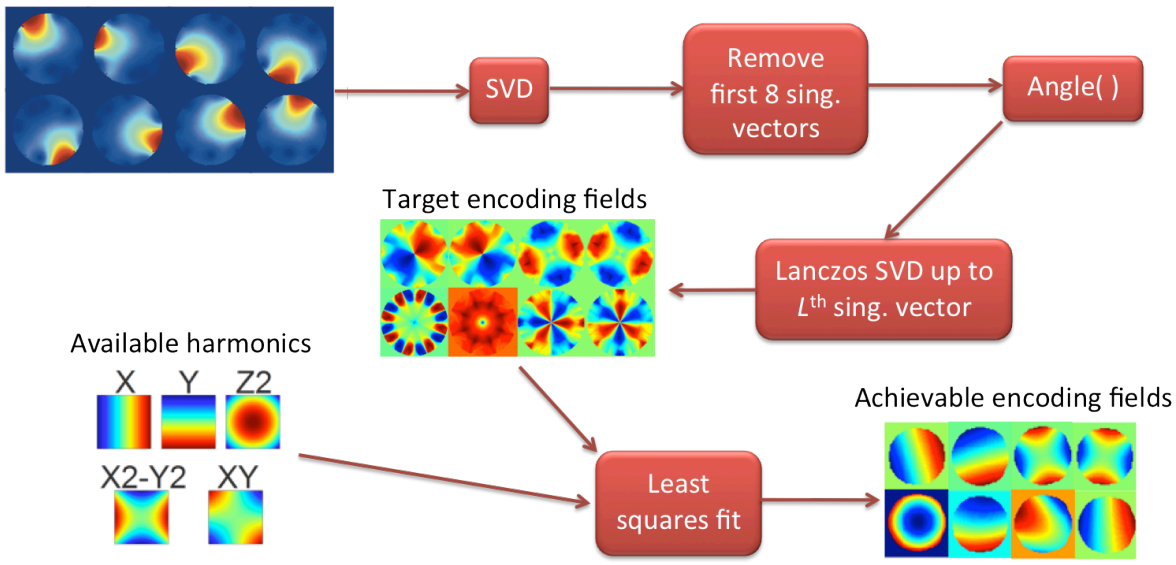

Figure 5.1: Target encoding fields are obtained from the null space of the coil profile projection matrix. The set of available harmonics are then used to generate achievable encoding fields that approximate the target fields in a least squares sense. Target and ideal field plots are reproduced from (129).

the coil profile matrix. Because they satisfy Laplace's equation, spherical harmonics form one appealing basis set with which to approximate desired SEM shapes. In principle, linear combinations of spherical harmonics can be used to create linearly independent projections of the object at each echo. As the object is projected onto each of these basis functions, successive echoes capture object features that were not resolved in the preceding echoes. Orthogonality between all pairs of spherical harmonics exists only over a spherical volume. This suggests that a suitable combination of first-order and higher-order spherical harmonics, an extension of the approach used in this paper, might be suitable for three-dimensional imaging.

Recent work at Yale and elsewhere (66, 128, 129, 130) shows that as the number of higher-order SEMs grows, reconstruction error decreases and resolution improves. Efforts are now underway to design and implement null space trajectories in the hopes of optimally harnessing the spatial encoding of the available RF coil sensitivities. Specific aims include highly-accelerated axial imaging, compressed sensing reconstruction with enforced spar-

sity, and 3-D volumetric reconstructions from highly undersampled data. The question of how to optimally encode in three dimensions using 2-D and 3-D nonlinear SEMs is an open question that invites the use of approaches like null space imaging, which is easily generalized from 2-D to 3-D.

5.2.4 Alternative coil arrays for generating nonlinear magnetic fields

Up until now, nonlinear encoding has only been performed using dedicated, high-power insert coils able to synthesize one or two nonlinear SEMs. As suggested by the null space approach, however, optimal spatial encoding may require use of numerous higher-order spherical harmonics, or even another basis set of field shapes. One intriguing possibility is to synthesize target fields from scalable arrays of individual, generic coils, as has been recently demonstrated in the context of multi-coil imaging (131) and shimming (132). Because they are placed closed to the object being imaged, arrays of shim coils offer the advantage of high efficiency, typically requiring less than 1 amp of current per coil to perform shimming or imaging. Also, with numerous degrees of freedom, the target field can be shaped with great flexibility, achieving third, fourth, and fifth-order spherical harmonics, in addition to other arbitrary shapes, depending on the number of available shim coils.

Alternatively, a design has recently been proposed for generating a set of three orthogonal PatLoc fields using a planar array of three coils inset in the scanner patient table (133). Such an SEM insert provides an extended FOV for PatLoc imaging of the body. As compared with previous PatLoc coils (134), the planar array also generalizes localization from two to three dimensions. A torque-balanced experimental coil is now under development for use in PatLoc and GradLoc imaging of the human body.

5.2.5 Resolution and noise amplification metrics for nonlinear fields

Quantifying noise levels and resolution in O-Space imaging is challenging because both properties are spatially-varying. By contrast, when linear SEMs are used, the point spread function and resolution are approximately the same everywhere in the FOV. Recently, a new approach has been introduced to estimate the resolution of O-Space images using the frame matrix (135), which is calculated using the encoding matrix and the noise covariance matrix, $Q = E^H \tilde{\Psi} E$. Each row of the frame matrix provides an approximation of the point spread function (PSF) for a particular voxel. For a given voxel, resolution is quantified as the ratio between the PSF intensity in the voxel of interest and the sum of the PSF intensity over all voxels.

The most accurate estimate for the PSF would be provided by the matrix product FE, which models the mapping of each source voxel from image space into the data domain (encoding by E) and then back again (reconstruction by F). For a particular reconstruction matrix F, this matrix product shows explicitly how signals originating in a particular voxel spread throughout the FOV in the reconstructed image. The frame matrix $E^H E$ only equals FE for the case of a unitary encoding matrix E, for which the pseudoinverse is equal to the conjugate transpose. This is true for the discrete Fourier transform, whose rows and columns are orthonormal, but the assumption of unitarity can not be made for nonlinear encoding schemes in general (as described in the section 4.5.3). The resolution metric described above also does not take into account the spatial distribution of the PSF "sidelobes" relative to the source voxel. A more robust metric might take into account both the PSF intensities in each voxel and their distance from the source voxel.

The frame matrix can be used to compute the noise variance on a voxel-by-voxel basis (136), providing an O-Space performance metric analogous to the SENSE g-factor for linear SEMs. Direct calculation of the noise matrix (Eq. 1.3.2) is computationally intractable for O-Space imaging. However, it has been shown that the diagonal elements

of the noise matrix can be calculated using a Taylor series of the inverse frame matrix, Q^{-1}. Because the method does not require explicit inversion of the frame matrix, it provides the first feasible way to calculate the noise variance map of an approach like O-Space imaging. Excellent agreement has been shown between variance maps obtained through this method with those obtained from Monte Carlo simulations (136). In the future, the frame matrix promises to be a useful tool for evaluating and comparing a range of nonlinear encoding strategies.

5.2.6 RF encoding with nonlinear fields

In an approach known as Excitation Localization, or ExLoc, PatLoc SEMs are played during RF excitation in order to select curvilinear slices with desirable properties (117). Such slices may be used, for instance, to efficiently image the cerebral cortex, particularly when combined with PatLoc spatial encoding.

RF excitation with the Z^2 SEM has been proposed as a way to create "virtual coil profiles" for use in highly accelerated imaging (137). By playing the Z^2 field along with an appropriate frequency offset in the RF pulse, rings of the object can be excited. By interleaving several such excitations in the span of a single TR, signal can be acquired from the entire object, but partitioned in between 2 and 8 rings. When the rings are weighted by surface coil profiles, the effective number of coils is equal to the number of coils multiplied by the number of rings. The set of virtual coil profiles are highly localized from one another, providing efficient spatial encoding. This strategy has been employed in combination with conventional radial pulse sequences to perform highly undersampled reconstructions with only 8 receive coils. High image quality has been achieved at acceleration factors that would not be possible using the 8 coils alone without RF localization of rings.

5.2.6.1 Reconstruction in the Fourier transform domain

It has been shown that nonlinear SEM reconstructions can be accelerated by sparsifying the encoding matrix (74). One way to achieve sparsity is to Fourier transform the encoding matrix and the data and then truncate all entries in the encoding matrix falling below a certain threshold. Reconstruction can then be performed using methods that are optimized for sparse matrices, such as the LSQR algorithm(67). This can speed reconstruction by orders of magnitude while introducing only a very slight blurring of the image due to loss of some high frequency information.

More recently, efforts have begun at Yale to perform reconstruction of PatLoc, O-Space, and related datasets in the frequency domain using a precomputed reconstruction matrix. This approach computes the Fourier transform of the coil-weighted phasor used at each time point in the acquisition. This spectral representation of the encoding function, or "k-space stamp", provides a way to assemble the k-space data of the reconstructed image. First, the appropriate linear combinations of all available k-space stamps must be obtained. The required linear combination for a given point in k-space is the weighted sum of all stamps that best approximates a delta function at the specified k-space point. The set of all such linear combinations is used to populate the reconstruction matrix, which can be directly applied to acquired data to generate the spectrum of the image. Fourier transformation then provides the final image. For a given acquisition strategy, the reconstruction matrix only needs to be calculated once. Fast reconstruction via a single matrix-vector multiplication can then be performed. This holds the promise to greatly speed the reconstruction of data encoded with nonlinear SEMs. Furthermore, it may be compatible with k-space parallel imaging methods like GRAPPA (43), permitting auto-calibrated O-Space reconstructions for the first time.

6

Appendix A:

Angenäherte Auflösung von Systemen linearer Gleichungen

6.1 Approximate Solution for Systems of Linear Equations (English translation)

By Stefan Kaczmarz (1895-1939)

Unable to find an English translation of Kaczmarz's original article (69), I undertook the following translation from German.

Although the approximate solution of the equation $f(x) = 0$ with an unknown variable has been extensively treated in the literature, we nevertheless know little about the solutions of sets of equations, even when they are linear. There are only two methods [1,2] which permit solutions, but they are limited to equations whose diagonal coefficients, a_{ii},

are large in comparison with the other coefficients. For this reason, one finds in the work of Mises and Pollaczek-Geiringer the condition $a_{ii} >= (n-1)a_{ik}$, and so forth.

In general, however, one can completely solve the equations by means of iterative methods. We assume that a system is given,

$$a_{i1}x_1 + a_{i2}x_2 + ... + a_{in}x_n + b_i = 0, \qquad i = 1, 2, ...n \tag{6.1.1}$$

which possesses only one solution, $x_1, x_2, ... x_n$. We presuppose that equation (6.1.1) by multiplication with appropriate constants is transformed in such a way that

$$\sum_{k=1}^{n} a_{ik}^2 = 1 \tag{6.1.2}$$

is valid for all $i = 1, 2, ... n$.

Let there exist a first estimate, $x_1^0, x_2^0, ...x_n^0$, which can be completely arbitrary. We divide successive estimates into groups, each consisting of n (approximate) solutions,

$$x_1^{(p,1)}, x_2^{(p,1)}, ... x_n^{(p,1)},$$

$$. \quad . \quad . \quad .$$

$$x_1^{(p,n)}, x_2^{(p,n)}, ... x_n^{(p,n)},$$

the n solutions of the p^{th} group. We define the $(p+1)^{th}$ group as follows

$$
\begin{aligned}
x_i^{(p+1,1)} &= x_i^{(p,n)} - a_{1i}L_1^{(p,n)} \\
x_i^{(p+1,2)} &= x_i^{(p+1,1)} - a_{2i}L_2^{(p+1,1)} \\
&\cdot \quad \cdot \quad \cdot \quad \cdot \quad \cdot \quad \cdot \\
x_i^{(p+1,n)} &= x_i^{(p+1,n-1)} - a_{ni}L_n^{(p+1,n-1)} \\
&(i = 1, 2, \ldots n)
\end{aligned}
$$

In these equations, $L_i^{(r,s)} = \sum_{k=1}^{n} a_{ik}x_k^{(r,s)} + b_i$ defines the left side of the i^{th} equation in (6.1.1), wherein the s^{th} solution to the r^{th} group is used for the variables $x_1, x_2, \ldots x_n$.

Consequently, we submit that the solution

$$x_1^{(r,s)}, x_2^{(r,s)}, \ldots x_n^{(r,s)} \qquad s = 1, 2, \ldots n, r = 1, 2, \ldots \tag{6.1.3}$$

converges toward $x_1, x_2, \ldots x_n$. From (6.1.3) it follows that

$$x_i^{(r,s+1)} - x_i = x_i^{(r,s)} - x_i - a_{s+1,i}L_{s+1}^{(r,s)}, \qquad s = 0, 1, \ldots n-1, \tag{6.1.4}$$

(for $s = 0$, we have $x_i^{(r,1)} - x_i = x_i^{(r-1,n)} - x_i - a_{1i}L_i^{(r-1,n)}$). By squaring and summing from $i = 1$ to n one obtains [1]

$$\sum_{i=1}^{n}(x_i^{(r,s+1)} - x_i)^2 = \sum_{i=1}^{n}(x_i^{(r,s)})^2 + L_{s+1}^2 - 2L_{s+1}\sum_{i=1}^{n} a_{s+1,i}(x_i^{(r,s)} - x_i) \tag{6.1.5}$$

But because

$$\sum a_{s+1,i}(x_i^{(r,s)} - x_i) = L_{s+1} - b_{i+1} + b_{i+1} = L_{s+1}, \tag{6.1.6}$$

[1] In (6.1.5) the double subscript for L is omitted.

if one substitutes

$$y_{r,s} = \sum_{i=1}^{n} (x_i^{(r,s)} - x_i)^2 \tag{6.1.7}$$

for the sake of conciseness, one then obtains from (6.1.5),

$$y_{r,s+1} = y_{r,s} - L_{s+1}^2. \tag{6.1.8}$$

The sequence $y_{r,s}$ thus falls monotonically toward the boundary y. Since $y_{r,s}$ are limited, one can find a subsequence of $x_i^{(r,s)}$ which approaches the bound $y_1, y_2, \ldots y_n$. From (6.1.8), it follows that

$$\lim_{r\to\infty} L_{s+1}^{(r,s)} = L_{s+1}(y_1, y_2, \ldots y_n) = 0, \tag{6.1.9}$$

and also $\lim x_i^{(r,s)} = y_i = x_i$, the system in (6.1.1) possesses only one solution. Therefore $y_{r,s}$ must converge toward zero, which means that $x_i^{(r,s)}$ tends toward x_i.

When the system (6.1.1) has more solutions, then naturally $y_1, y_2, \ldots y_n$ is one of the possible solution.

For the practical execution of calculations it is important to note that the condition $\sum_1^n a_{ik}^2 = 1$ does not have to be strictly fulfilled. If one sets $\sum_1^n a_{ik}^2 = c_i$, $c_i \neq 1$, then relation (6.1.5) turns into

$$y_{r,s+1} = y_{r,s} + c_{s+1} L_{s+1}^2 - 2L_{s+1}^2. \tag{6.1.10}$$

We therefore see that $y_{r,s+1} < y_{r,s}$ (and thus converges) if $c_{s+1} < 2$. One can choose coefficients c_i such that $c_i < 2$ and the multiplication becomes very easy.

Geometrically the algorithm means the following:

The point $x_1^0, \dots x_n^0$ is projected orthogonally onto the first hyperplane $L_1 = 0$; this projection is just the point $x_1^{(1,1)}, \dots x_n^{(1,1)}$, which is now cast onto $L_2 = 0$, etc. The point $x_1^{(1,n)}, \dots x_n^{(1,n)}$ is again cast onto $L_1 = 0$, giving the point $x_1^{(2,1)}, \dots x_n^{(2,1)}$, etc. In this way, the convergence of the algorithm is easily plausible.

References:

[1] R. Mises, H. Pollaczek-Geiringer, Praktische Verfahren der Gleichungsauflösung, Zeit. ang. Math. u. Mech. 9 (1929), S. 58-77, 152-164.

[2] J. Morris, A successive approximative process for solving simultaneous linear equations, Aeronaut. Res. Com. Rep. Number 1711 (1936) p. 1-12.

Przybliżone rozwiązywanie układów równań liniowych. — Angenäherte Auflösung von Systemen linearer Gleichungen.

Note

de M. *S. KACZMARZ,*

présentée le 14 Juin 1937 par M. Th. Banachiewicz m. t.

Obgleich die angenäherte Auflösung der Gleichung $f(x)=0$ mit einer Unbekannten zahlreiche Bearbeitungen in der Literatur aufweist, wissen wir dennoch sehr wenig über die Auflösung von Gleichungssystemen, sogar wenn sie linear sind. Es gibt zwar einige Methoden [1]), welche die Auflösung gestatten, doch sind sie auf Gleichungen beschränkt, für welche die Diagonalkoeffizienten a_{ii} überwiegend groß sind im Vergleich mit den übrigen Koeffizienten; so findet man z. B. bei Mises und Pollaczek-Geiringer die Bedingung $a_{ii} \geqslant (n-1)\,a_{ik}$ und desgleichen.

Man kann jedoch die Aufgabe mittels eines Iterationsverfahrens ganz allgemein lösen. Nehmen wir an, daß ein System

$$a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n + b_i = 0, \qquad (i=1,2,\ldots n) \quad (1)$$

gegeben ist, welches eine einzige Lösung $x_1, x_2, \ldots x_n$ besitzt. Wir setzen voraus, daß die Gleichungen (1) durch Multiplikation mit entsprechenden Konstanten so umgeformt wurden, daß

$$\sum_{k=1}^{n} a_{ik}^2 = 1$$

für alle $i=1,2\ldots n$ gilt.

[1]) Vgl. R. Mises und H. Pollaczek-Geiringer, Praktische Verfahren der Gleichungsauflösung, Zeit. ang. Math. u. Mech. 9 (1929), S. 58—77, 152—164.

J. Morris, A successive approximative process for solving simultaneous linear equations, Aeronaut. Res. Com. Rep. Nr 1711 (1936) p. 1—12.

Es sei $x_1^0, x_2^0, \dots x_n^0$ eine erste Annäherung, die ganz beliebig sein kann. Die folgenden Annäherungen teilen wir in Gruppen, deren jede aus n (angenäherten) Lösungen besteht. Die Lösung $x_1^0, x_2^0, \dots x_n^0$ bildet die nullte Gruppe. Es seien

$$\begin{array}{l} x_1^{(p,1)}, x_2^{(p,1)}, \dots x_n^{(p,1)}, \\ \dots \quad \dots \quad \dots \\ x_1^{(p,n)}, x^{(p,n)} \dots x_n^{(p,n)}, \end{array}$$

die n Lösungen der p-ten Gruppe; wir definieren die $(p+1)$-te Gruppe folgendermaßen

$$\begin{array}{ll} x_i^{(p+1,1)} = x_i^{(p,n)} & - a_{1i} L_1^{(p,n)}, \\ x_i^{(p+1,2)} = x_i^{(p+1,1)} & - a_{2i} L_2^{(p+1,1)} \quad (i=1,2 \dots n) \quad (2) \\ \dots \quad \dots & \\ x_i^{(p+1,n)} = x_i^{(p+1,n-1)} & - a_{ni} L_n^{(p+1,n-1)}. \end{array}$$

In diesen Gleichungen bezeichnet $L_i^{(r,s)} = \sum_{k=1}^{n} a_{ik} x_k^{(r,s)} + b_i$ die linke Seite der i-ten Gleichung (1), wo für die Unbekannten $x_1, x_2, \dots x_n$ die s-te Lösung der r-ten Gruppe eingesetzt wurde.

Wir behaupten, daß die Folge der Lösungen

$$x_1^{(r,s)}, x_2^{(r,s}, \dots x_n^{(r,s)} \qquad s=1,2,\dots n, \; r=1,2,\dots$$

gegen $x_1, x_2 \dots x_n$ konvergiert. Aus (2) folgt nämlich

$$x_i^{(r,s+1)} - x_i = x_i^{(r,s)} - x_i - a_{s+1,i} L_{s+1}^{(r,s)}, \qquad s=0,1,\dots n-1,$$

(für $s=0$, ist $x_i^{(r,1)} - x_i = x_i^{(r-1,n)} - x_i - a_{1i} L_i^{(r-1,n)}$). Durch Quadrieren und Summieren von $i=1$ bis n erhält man [1])

$$\sum_{i=1}^{n} (x_i^{(r,s+1)} - x_i)^2 = \sum_{i=1}^{n} (x_i^{(r,s)} - x_i)^2 + L_{s+1}^2 - 2L_{s+1} \sum_{i=1}^{n} a_{s+1,i} (x_i^{(r,s)} - x_i). \quad (3)$$

Da aber

$$\sum a_{s+1,i} (x_i^{(r,s)} - x_i) = L_{s+1} - b_{i+1} + b_{i+1} = L_{s+1},$$

so bekommt man, wenn

$$y_{r,s} = \sum_{i=1}^{n} (x_i^{(r,s)} - x_i)^2$$

[1]) In (3) ist der obere Doppelindex bei L weggelassen.

der Kürze halber gesetzt wird, aus (3) die Gleichung

$$y_{r,s+1} = y_{r,s} - L_{s+1}^2. \tag{4}$$

Die Folge $\{y_{r,s}\}$ ist also monoton fallend und besitzt somit eine Grenze y; da die $y_{r,s}$ beschränkt sind, kann man eine Teilfolge in der Folge $\{x_i^{(r,s)}\}$ finden, welche gegen die Grenze $y_1, y_2, \dots y_n$ strebt. Aus der Gleichung (4) folgt, daß

$$\lim_{r\to\infty} L_{s+1}^{(r,s)} = L_{s+1}(y_1, y_2, \dots y_n) = 0,$$

also auch $\lim x_i^{(r,s)} = y_i = x_i$, da das System (1) nur eine einzige Lösung besitzt. Es muß also $y_{r,s}$ gegen Null konvergieren, was besagt, daß $x_i^{(r,s)}$ gegen x_i streben.

Wenn das System (1) mehrere Lösungen besitzt, ist natürlich $y_1, y_2, \dots y_n$ eine der möglichen Lösungen.

Für die praktische Ausführung der Rechnungen ist es wichtig, daß die Bedingung $\sum_1^n a_{ik}^2 = 1$ nicht streng erfüllt werden muß. Wenn man nämlich $\sum_1^n a_{ik}^2 = c_i$ setzt, $c_i \neq 1$, so geht die Relation (3) in

$$y_{r,s+1} = y_{r,s} + c_{s+1} L_{s+1}^2 - 2 L_{s+1}^2$$

über. Wir sehen also, daß $y_{r,s+1} < y_{r,s}$ (und somit konvergiert), wenn $c_{s+1} < 2$. Man kann also die c_i so wählen, daß $c_i < 2$ und die Multiplikation sehr einfach werde.

Geometrisch bedeutet das Verfahren folgendes:

Der Punkt $x_1^0, \dots x_n^0$ wird orthogonal auf die erste Hyperebene $L_1 = 0$ projiziert; die Projektion ist eben der Punkt $x_1^{(1,1)}, \dots x_n^{(1,1)}$, welcher nun auf $L_2 = 0$ geworfen wird usw. Der Punkt $x_1^{(1,n)}, \dots x_n^{(1,n)}$ wird wiederum auf $L_1 = 0$ geworfen und gibt den Punkt $x_1^{(2,1)}, \dots x_n^{(2,1)}$, usw. Die Konvergenz des Verfahrens ist geometrisch ohne weiteres einleuchtend.

References

[1] K. Pruessmann, M. Weiger, M. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Magnetic Resonance in medicine*, vol. 42, no. 5, pp. 952–962, 1999.

[2] M. Bernstein, K. King, and X. Zhou, "Handbook of MRI Pulse Sequences," 2004.

[3] R. De Graaf, *In vivo NMR spectroscopy*. John Wiley and Sons, 2007.

[4] P. Roemer, W. Edelstein, C. Hayes, S. Souza, and O. Mueller, "The nmr phased array," *Magnetic Resonance in Medicine*, vol. 16, no. 2, pp. 192–225, 1990.

[5] J. Vaughan, H. Hetherington, J. Otu, J. Pan, and G. Pohost, "High frequency volume coils for clinical NMR imaging and spectroscopy," *Magnetic Resonance in Medicine*, vol. 32, no. 2, pp. 206–218, 1994.

[6] L. Bolinger, M. Prammer, and J. Leigh, "A multiple-frequency coil with a highly uniform B1 field," *Journal of Magnetic Resonance*, vol. 81, no. 1, pp. 162–166, 1989.

[7] R. Lee, C. Westgate, R. Weiss, D. Newman, and P. Bottomley, "Planar strip array (PSA) for MRI," *Magnetic resonance in medicine*, vol. 45, no. 4, pp. 673–683, 2001.

[8] G. Adriany, P. Van de Moortele, F. Wiesinger, S. Moeller, J. Strupp, P. Andersen, C. Snyder, X. Zhang, W. Chen, K. Pruessmann, P. Boesiger, J. Vaughan, and

K. Ugurbil, "Transmit and receive transmission line arrays for 7 Tesla parallel imaging," *Magnetic Resonance in Medicine*, vol. 53, no. 2, pp. 434–445, 2005.

[9] P. Lauterbur *et al.*, "Image formation by induced local interactions: examples employing nuclear magnetic resonance," *Nature*, vol. 242, no. 5394, pp. 190–191, 1973.

[10] J. Hennig, A. Welz, G. Schultz, J. Korvink, Z. Liu, O. Speck, and M. Zaitsev, "Parallel imaging in non-bijective, curvilinear magnetic field gradients: a concept study," *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 21, no. 1, pp. 5–14, 2008.

[11] A. Kumar, D. Welti, and R. Ernst, "NMR Fourier zeugmatography," *Journal of Magnetic Resonance (1969)*, vol. 18, no. 1, pp. 69–83, 1975.

[12] S. Deans, *The Radon transform and some of its applications*. Wiley, New York, 1983.

[13] R. Bracewell, *The Fourier transform and its applications*. McGraw-Hill, New York, 2000.

[14] J. Stockmann, P. Ciris, G. Galiana, L. Tam, and R. Constable, "O-Space imaging: Highly efficient parallel imaging using second-order nonlinear fields as encoding gradients with no phase encoding," *Magnetic Resonance in Medicine*, vol. 64, no. 2, pp. 447–456, 2010.

[15] D. Gallichan, C. Cocosco, A. Dewdney, G. Schultz, A. Welz, J. Hennig, and Z. M., "Simultaneously driven linear and nonlinear spatial encoding fields in MRI," *Magnetic Resonance in Medicine*, vol. 65, no. 3, pp. 702–714, 2011.

[16] J. Earls, V. Ho, T. Foo, E. Castillo, and S. Flamm, "Cardiac mri: recent progress and continued challenges," *Journal of magnetic resonance imaging*, vol. 16, no. 2, pp. 111–127, 2002.

[17] F. Epstein, J. London, D. Peters, L. Goncalves, K. Agyeman, J. Taylor, R. Balaban, and A. Arai, "Multislice first-pass cardiac perfusion mri: Validation in a model of myocardial infarction†," *Magnetic resonance in medicine*, vol. 47, no. 3, pp. 482–491, 2002.

[18] S. Sampath, J. Derbyshire, E. Atalar, N. Osman, and J. Prince, "Real-time imaging of two-dimensional cardiac strain using a harmonic phase magnetic resonance imaging (harp-mri) pulse sequence," *Magnetic Resonance in Medicine*, vol. 50, no. 1, pp. 154–163, 2003.

[19] E. Merkle and B. Dale, "Abdominal mri at 3.0 t: the basics revisited," *American Journal of Roentgenology*, vol. 186, no. 6, pp. 1524–1532, 2006.

[20] F. Calamante, D. Gadian, and A. Connelly, "Quantification of perfusion using bolus tracking magnetic resonance imaging in stroke," *Stroke*, vol. 33, no. 4, pp. 1146–1151, 2002.

[21] S. Ogawa, T. Lee, A. Kay, and D. Tank, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation," *Proceedings of the National Academy of Sciences*, vol. 87, no. 24, p. 9868, 1990.

[22] K. Odegard, J. DiNardo, B. Tsai-Goodman, A. Powell, T. Geva, and P. Laussen, "Anaesthesia considerations for cardiac mri in infants and small children," *Pediatric Anesthesia*, vol. 14, no. 6, pp. 471–476, 2004.

[23] P. Jezzard and R. Balaban, "Correction for geometric distortion in echo planar images from b0 field variations," *Magnetic Resonance in Medicine*, vol. 34, no. 1, pp. 65–73, 1995.

[24] M. Stehling, R. Turner, and P. Mansfield, "Echo-planar imaging: magnetic resonance imaging in a fraction of a second," *Science*, vol. 254, no. 5028, p. 43, 1991.

[25] J. Hennig, A. Nauerth, and H. Friedburg, "Rare imaging: a fast imaging method for clinical mr," *Magnetic Resonance in Medicine*, vol. 3, no. 6, pp. 823–833, 1986.

[26] D. Noll, "Multishot rosette trajectories for spectrally selective MR imaging," *IEEE Transactions on Medical Imaging*, vol. 16, no. 4, pp. 372–377, 1997.

[27] "A primer on medical device interactions with magnetic resonance imaging systems," *www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm107721*. 2009.

[28] C. Ham, J. Engels, G. Van de Wiel, and A. Machielsen, "Peripheral nerve stimulation during MRI: effects of high gradient amplitudes and switching rates," *Journal of Magnetic Resonance Imaging*, vol. 7, no. 5, pp. 933–937, 1997.

[29] H. Tan and C. Meyer, "Estimation of k-space trajectories in spiral mri," *Magnetic Resonance in Medicine*, vol. 61, no. 6, pp. 1396–1404, 2009.

[30] J. Duyn, Y. Yang, J. Frank, and J. van der Veen, "Simple correction method for k-space trajectory deviations in MRI," *Journal of Magnetic Resonance*, vol. 132, no. 1, pp. 150–153, 1998.

[31] J. Jackson, C. Meyer, D. Nishimura, and A. Macovski, "Selection of a convolution function for fourier inversion using gridding," *Medical Imaging, IEEE Transactions on*, vol. 10, no. 3, pp. 473–478, 1991.

[32] J. Pipe and P. Menon, "Sampling density compensation in mri: rationale and an iterative numerical solution," *Magnetic Resonance in Medicine*, vol. 41, no. 1, pp. 179–186, 1999.

[33] J. Fessler, "On nufft-based gridding for non-cartesian mri," *Journal of Magnetic Resonance*, vol. 188, no. 2, pp. 191–195, 2007.

[34] K. Pruessmann, M. Weiger, P. Boernert, and P. Boesiger, "Advances in sensitivity encoding with arbitrary k-space trajectories," *Magnetic Resonance in Medicine*, vol. 46, no. 4, pp. 638–651, 2001.

[35] K. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1086–1098, 2007.

[36] J. Beck and K. Arnold, *Parameter estimation in engineering and science*. Wiley New York, 1977.

[37] T. Lawry, M. Weiner, and G. Matson, "Computer modeling of surface coil sensitivity," *Magnetic resonance in medicine*, vol. 16, no. 2, pp. 294–302, 1990.

[38] D. Gensanne, G. Josse, J. Lagarde, and D. Vincensini, "High spatial resolution quantitative mr images: an experimental study of dedicated surface coils," *Physics in medicine and biology*, vol. 51, p. 2843, 2006.

[39] S. King, S. Varosi, and G. Duensing, "Eigenmode analysis for understanding phased array coils and their limits," *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, vol. 29, no. 1, pp. 42–49, 2006.

[40] G. Wiggins, C. Triantafyllou, A. Potthast, A. Reykowski, M. Nittka, and L. Wald, "32-channel 3 Tesla receive-only phased-array head coil with soccer-ball element geometry," *Magnetic Resonance in Medicine*, vol. 56, no. 1, pp. 216–223, 2006.

[41] D. Walsh, A. Gmitro, and M. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magnetic Resonance in Medicine*, vol. 43, no. 5, pp. 682–690, 2000.

[42] D. Sodickson and W. Manning, "Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays," *Magnetic Resonance in Medicine*, vol. 38, no. 4, pp. 591–603, 1997.

[43] M. Griswold, P. Jakob, R. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (grappa)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.

[44] W. Kyriakos, L. Panych, D. Kacher, C. Westin, S. Bao, R. Mulkern, and F. Jolesz, "Sensitivity profiles from an array of coils for encoding and reconstruction in parallel (space rip)," *Magnetic Resonance in Medicine*, vol. 44, no. 2, pp. 301–308, 2000.

[45] M. Bydder, D. J. Larkman, and J. V. Hajnal, "Generalized SMASH imaging," *Magnetic Resonance in Medicine*, vol. 47, pp. 160–70, Jan 2002.

[46] B. Madore, G. Glover, and N. Pelc, "Unaliasing by fourier-encoding the overlaps using the temporal dimension (unfold), applied to cardiac imaging and fmri," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 813–828, 1999.

[47] F. Breuer, M. Blaimer, R. Heidemann, M. Mueller, M. Griswold, and P. Jakob, "Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging," *Magnetic Resonance in Medicine*, vol. 53, no. 3, pp. 684–691, 2005.

[48] G. Wiggins, J. Polimeni, A. Potthast, M. Schmitt, V. Alagappan, and L. Wald, "96-channel receive-only head coil for 3 tesla: design optimization and evaluation," *Magnetic Resonance in Medicine*, vol. 62, no. 3, p. 754, 2009.

[49] B. Liu, K. King, M. Steckner, J. Xie, J. Sheng, and L. Ying, "Regularized sensitivity encoding (sense) reconstruction using bregman iterations," *Magnetic Resonance in Medicine*, vol. 61, no. 1, pp. 145–152, 2009.

[50] P. Jakob, M. Grisowld, R. Edelman, and D. Sodickson, "Auto-smash: a self-calibrating technique for smash imaging," *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 7, no. 1, pp. 42–54, 1998.

[51] R. Heidemann, M. Griswold, A. Haase, and P. Jakob, "Vd-auto-smash imaging," *Magnetic Resonance in Medicine*, vol. 45, no. 6, pp. 1066–1074, 2001.

[52] F. Breuer, S. Kannengiesser, M. Blaimer, N. Seiberlich, P. Jakob, and M. Griswold, "General formulation for quantitative G-factor calculation in GRAPPA reconstructions," *Magnetic Resonance in Medicine*, vol. 62, no. 3, pp. 739–746, 2009.

[53] J. Shewchuk, *An introduction to the conjugate gradient method without the agonizing pain*. Carnegie-Mellon University, Pittsburgh, PA, 1994.

[54] N. Seiberlich, F. Breuer, M. Blaimer, K. Barkauskas, P. Jakob, and M. Griswold, "Non-cartesian data reconstruction using grappa operator gridding (grog)," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1257–1265, 2007.

[55] U. Katscher, P. Börnert, C. Leussler, and J. S. van den Brink, "Transmit SENSE," *Magn Reson Med*, vol. 49, pp. 144–50, Jan 2003.

[56] U. Katscher and P. Börnert, "Parallel RF transmission in MRI," *NMR Biomed*, vol. 19, pp. 393–400, May 2006.

[57] Y. Zhu, R. Watkins, R. Giaquinto, C. Hardy, G. Kenwood, S. Mathias, T. Valent, M. Denzin, J. Hopkins, W. Peterson, *et al.*, "Parallel excitation on an eight transmit-channel mri system," in *Proc. Int. Soc. Magn. Reson. Med*, vol. 13, p. 14, 2005.

[58] F. Wiesinger, P. Boesiger, and K. Pruessmann, "Electrodynamics and ultimate snr in parallel mr imaging," *Magnetic resonance in medicine*, vol. 52, no. 2, pp. 376–390, 2004.

[59] P. Ciris, J. Stockmann, and R. Constable, "O-space imaging: Tailoring encoding gradients to coil profiles for highly accelerated imaging," in *Proceedings 17th Scientific Meeting, ISMRM, Honolulu, HI*, p. 4556, 2009.

[60] F. Roméo and D. Hoult, "Magnet field profiling: analysis and correcting coil design," *Magnetic Resonance in Medicine*, vol. 1, no. 1, pp. 44–65, 1984.

[61] G. Schultz, P. Ullmann, H. Lehr, A. Welz, J. Hennig, and M. Zaitsev, "Reconstruction of MRI data encoded with arbitrarily shaped, curvilinear, nonbijective magnetic fields," *Magnetic Resonance in Medicine*, vol. 64, no. 5, pp. 1390–1403, 2010.

[62] B. Kusse and E. Westwig, *Mathematical physics: applied mathematics for scientists and engineers*. Wiley-VCH, 2006.

[63] G. Schultz, H. Weber, D. Gallichan, W. Witschey, A. Welz, C. Cocosco, J. Hennig, and M. Zaitsev, "Radial Imaging with Multipolar Magnetic Encoding Fields," *IEEE Trans. Med. Imag.*, vol. 16, p. 17, 2011.

[64] L. Sacolick, *Method development in in-vivo NMR spectroscopy and spectroscopic imaging*. PhD thesis, Yale University, 2008.

[65] R. Turner, "A target field approach to optimal coil design," *Journal of physics D: Applied physics*, vol. 19, p. L147, 1986.

[66] F. Lin, T. Witzel, A. Nummenmaa, P. Vesanen, R. Ilmoniemi, and J. Belliveau, "Multi-dimensional encoded (MDE) magnetic resonance imaging Multi-dimensional encoded (MDE) magnetic resonance imaging," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 480, 2011.

[67] C. Paige and M. Saunders, "Lsqr: An algorithm for sparse linear equations and sparse least squares," *ACM Transactions on Mathematical Software (TOMS)*, vol. 8, no. 1, pp. 43–71, 1982.

[68] W. Smythe, "Static and dynamic electricity," 1950.

[69] S. Kaczmarz, "Angenaeherte aufloesung von systemen linearer gleichungen," *Bull. Acad. Polon. Sci. Lett. A*, vol. 35, no. 355–357, 1937.

[70] G. Herman and A. Lent, "Iterative reconstruction algorithms," *Computers in biology and medicine*, vol. 6, no. 4, pp. 273–294, 1976.

[71] T. Strohmer and R. Vershynin, "A randomized kaczmarz algorithm with exponential convergence," *Journal of Fourier Analysis and Applications*, vol. 15, no. 2, pp. 262–278, 2009.

[72] Y. Censor, P. Eggermont, and D. Gordon, "Strong underrelaxation in kaczmarz's method for inconsistent systems," *Numerische Mathematik*, vol. 41, no. 1, pp. 83–92, 1983.

[73] F. Lin, K. Kwong, J. Belliveau, and L. Wald, "Parallel imaging reconstruction using automatic regularization," *Magnetic Resonance in Medicine*, vol. 51, no. 3, pp. 559–567, 2004.

[74] G. Schultz, D. Gallichan, M. Reisert, M. Zaitsev, and H. J., "Fast image reconstruction for generalized projection imaging," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 2868, 2011.

[75] D. Needell, "Randomized kaczmarz solver for noisy linear systems," *BIT Numerical Mathematics*, vol. 50, no. 2, pp. 395–403, 2010.

[76] K. Koch, *In Vivo Static Field Perturbations in Magnetic Resonance*. PhD thesis, Yale University, 2007.

[77] Y. Censor, "Row-action methods for huge and sparse systems and their applications," *SIAM review*, vol. 23, no. 4, pp. 444–466, 1981.

[78] V. Parot, C. Sing-Long, C. Lizama, C. Tejos, S. Uribe, and P. Irarrazaval, "Application of the fractional fourier transform to image reconstruction in MRI," *Magnetic Resonance in Medicine*, vol. in press, 2011.

[79] R. de Graaf, P. Brown, S. McIntyre, D. Rothman, and T. Nixon, "Dynamic shim updating (DSU) for multislice signal acquisition," *Magnetic resonance in medicine*, vol. 49, no. 3, pp. 409–416, 2003.

[80] G. Morrell and D. Spielman, "Dynamic shimming for multi-slice magnetic resonance imaging," *Magnetic resonance in medicine*, vol. 38, no. 3, pp. 477–483, 1997.

[81] K. Koch, L. Sacolick, T. Nixon, S. McIntyre, D. Rothman, and R. de Graaf, "Dynamically shimmed multivoxel 1h magnetic resonance spectroscopy and multislice magnetic resonance spectroscopic imaging of the human brain," *Magnetic Resonance in Medicine*, vol. 57, no. 3, pp. 587–591, 2007.

[82] R. Lee, R. Giaquinto, and C. Hardy, "Coupling and decoupling theory and its application to the MRI phased array," *Magnetic resonance in medicine*, vol. 48, no. 1, pp. 203–213, 2002.

[83] N. Avdievich, J. Pan, J. Baehring, D. Spencer, and H. Hetherington, "Short echo spectroscopic imaging of the human brain at 7T using transceiver arrays," *Magnetic Resonance in Medicine*, vol. 62, no. 1, pp. 17–25, 2009.

[84] D. Pozar, *Microwave engineering*. Wiley-India, 2009.

[85] X. Zhang, K. Ugurbil, and W. Chen, "A microstrip transmission line volume coil for human head mr imaging at 4 t," *Journal of Magnetic Resonance*, vol. 161, no. 2, pp. 242–251, 2003.

[86] G. Galiana, J. Stockmann, L. Tam, and R. Todd Constable, "Spin dephasing under nonlinear gradients: Implications for imaging and field mapping," *Magnetic Resonance in Medicine*, 2011.

[87] G. Galiana, J. Stockmann, L. Tam, and R. Constable, "The inner lives of voxels: Revisiting the basics for nonlinear gradient imaging," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 4621, 2011.

[88] J. Schneider, M. Haas, W. Ruhm, J. Hennig, and P. Ullmann, "Robust spatially selective excitation using radiofrequency pulses adapted to the effective spatially encoding magnetic fields," *Magnetic Resonance in Medicine*, vol. 65, no. 2, pp. 409–421, 2011.

[89] T. R. Brown, B. M. Kincaid, and K. Ugurbil, "NMR chemical shift imaging in three dimensions," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 79, no. 11, pp. 3523–3526, 1982.

[90] M. Robson, J. Gore, and R. Constable, "Measurement of the point spread function in MRI using constant time imaging," *Magnetic Resonance in Medicine*, vol. 38, no. 5, pp. 733–740, 1997.

[91] M. Zaitsev, J. Hennig, and O. Speck, "Point spread function mapping with parallel imaging techniques and high acceleration factors: Fast, robust, and flexible method for echo-planar imaging distortion correction," *Magnetic Resonance in Medicine*, vol. 52, no. 5, pp. 1156–1166, 2004.

[92] N. De Zanche, C. Barmet, J. Nordmeyer-Massner, and K. Pruessmann, "Nmr probes for measuring magnetic fields and field dynamics in mr systems," *Magnetic Resonance in Medicine*, vol. 60, no. 1, pp. 176–186, 2008.

[93] C. Barmet, N. Zanche, and K. Pruessmann, "Spatiotemporal magnetic field monitoring for mr," *Magnetic Resonance in Medicine*, vol. 60, no. 1, pp. 187–197, 2008.

[94] C. Barmet, N. De Zanche, B. Wilm, and K. Pruessmann, "A transmit/receive system for magnetic field monitoring of in vivo mri," *Magnetic Resonance in Medicine*, vol. 62, no. 1, pp. 269–276, 2009.

[95] A. Welz, D. Gallichan, C. Cocosco, R. Kumar, F. Jia, J. Snyder, A. Dewdney, J. Korvink, J. Hennig, and M. Zaitsev, "Characterization of a PatLoc gradient coil," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 1527, 2010.

[96] V. Rasche, D. Holz, and R. Proksa, "MR flouroscopy using projection reconstruction multi-gradient-echo (prMGE) MRI," *Magnetic Resonance in Medicine*, vol. 42, pp. 324–334, 1999.

[97] M. Hestenes and E. Stiefel, "Methods of conjugate gradients for solving linear systems," *Journal cf Research of the National Bureau of Standards Vol*, vol. 49, no. 6, 1952.

[98] K. Lange and R. Carson, "Em reconstruction algorithms for emission and transmission tomography.," *Journal of computer assisted tomography*, vol. 8, no. 2, p. 306, 1984.

[99] Y. Censor, "On sequential and parallel projection algorithms for feasibility and optimization," in *Proceedings of SPIE*, vol. 4553, pp. 1–9, 2001.

[100] S. Stone, J. Haldar, S. Tsao, W. Hwu, B. Sutton, and Z. Liang, "Accelerating advanced MRI reconstructions on GPUs," *Journal of Parallel and Distributed Computing*, vol. 68, no. 10, pp. 1307–1318, 2008.

[101] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, "Second order total generalized variation (TGV) for MRI," *Magnetic Resonance in Medicine*, 2010.

[102] M. Lustig, D. Donoho, and J. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[103] E. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *Information Theory, IEEE Transactions on*, vol. 52, no. 12, pp. 5406–5425, 2006.

[104] H. Ozaktas, M. Kutay, and Z. Zalevsky, *The fractional Fourier transform with applications in optics and signal processing*. Wiley, 2001.

[105] J. Goodman, *Introduction to Fourier optics*. Roberts & Company Publishers, 2005.

[106] B. Saleh and M. Teich, *Fundamentals of photonics*. John Wiley & Sons New York, 2007.

[107] V. Wedeen, Y. Chao, and J. Ackerman, "Dynamic range compression in MRI by means of a nonlinear gradient pulse," *Magnetic Resonance in Medicine*, vol. 6, no. 3, pp. 287–295, 1988.

[108] S. Ito and Y. Yamada, "Alias-free image reconstruction using fresnel transform in the phase-scrambling fourier imaging technique," *Magnetic Resonance in Medicine*, vol. 60, no. 2, pp. 422–430, 2008.

[109] M. Zaitsev, G. Schultz, and J. Hennig, "Extended anti-aliasing reconstruction for phase-scrambled MRI with quadratic phase modulation," in *Proceedings 17th Scientific Meeting, ISMRM, Honolulu, HI*, p. 2859, 2009.

[110] J. Kaffanke, T. Dierkes, S. Romanzetti, M. Halse, J. Rioux, M. Leach, B. Balcom, and N. Shah, "Application of the chirp z-transform to MRI data," *Journal of Magnetic Resonance*, vol. 178, no. 1, pp. 121–128, 2006.

[111] J. Stockmann, G. Galiana, V. Parot, L. Tam, and R. Constable, "Scalable anti-aliasing image reconstruction in the presence of a quadratic "phase-scrambling" gradient using the Fractional Fourier Transform," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 2867, 2011.

[112] C. Hayes, W. Edelstein, J. Schenck, O. Mueller, and M. Eash, "An efficient, highly homogeneous radiofrequency coil for whole-body NMR imaging at 1.5 T," *J Magn Reson*, vol. 63, no. 3, pp. 622–628, 1985.

[113] A. Bultheel and H. Martinez-Sulbaran, "http://nalag.cs.kuleuven.be/research/software/FRFT/," 2011.

[114] W. Witschey, C. Cocosco, D. Gallichan, G. Schultz, H. Weber, A. Welz, J. Hennig, and M. Zaitsev, "Localization by nonlinear phase preparation and k-space trajectory design," *Magnetic Resonance in Medicine*, vol. in press, 2011.

[115] R. Pohmann, E. Rommel, and M. von Kienlin, "Beyond k-space: Spectral localization using higher order gradients," *Journal of Magnetic Resonance*, vol. 141, no. 2, pp. 197–206, 1999.

[116] R. de Graaf, D. Rothman, and T. Nixon, "Spatial Localization with Pulsed Second-Order Shims," in *Proceedings 15th Scientific Meeting, ISMRM, Berlin, Germany*, 2007.

[117] H. Weber, D. Gallichan, G. Schultz, W. Witschey, A. Welz, C. Cocosco, J. Hennig, and M. Zaitsev, "ExLoc: Excitation and encoding of curved slices," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, 2011.

[118] J. Pipe, "Spatial encoding and reconstruction in MRI with quadratic phase profiles," *Magnetic Resonance in Medicine*, vol. 33, no. 1, pp. 24–33, 1995.

[119] J. Asslaender, M. Balimer, F. Breuer, M. Zaitsev, and P. Jakob, "Combination of arbitrary gradient encoding fields using SPACE-RIP for reconstruction (COGNAC)," *Proceedings of the ISMRM 19th Annual Meeting*, p. 2870, 2011.

[120] G. Puy, J. Marques, R. Gruetter, J. Thiran, D. Van De Ville, P. Vandergheynst, and Y. Wiaux, "Spread spectrum magnetic resonance imaging," *Medical Imaging, IEEE Transactions on*, vol. PP, no. 99, p. 1, 2011.

[121] P. Irarrazaval, C. Lizama, V. Parot, C. Sing-Long, and C. Tejos, "The fractional Fourier transform and quadratic field magnetic resonance imaging," *Computers & Mathematics with Applications*, 2011.

[122] J. Stockmann, G. Galiana, V. Parot, L. Tam, and R. Constable, "The variable-order fractional Fourier transform: A new tool for efficient reconstruction of images encoded by linear and quadratic gradients with reduced sensitivity to calibration errors.," *Proceedings of the ISMRM 19th Annual Meeting*, p. 744, 2011.

[123] A. Sahin, M. Kutay, and H. Ozaktas, "Nonseparable two-dimensional fractional fourier transform," *Applied Optics*, vol. 37, no. 23, pp. 5444–5453, 1998.

[124] A. Koc, H. Ozaktas, and L. Hesselink, "Fast and accurate computation of two-dimensional non-separable quadratic-phase integrals," *JOSA A*, vol. 27, no. 6, pp. 1288–1302, 2010.

[125] D. Noll, C. Meyer, J. Pauly, D. Nishimura, and A. Macovski, "A homogeneity correction method for magnetic resonance imaging with time-varying gradients," *Medical Imaging, IEEE Transactions on*, vol. 10, no. 4, pp. 629–637, 1991.

[126] H. Schomberg, "Off-resonance correction of MR images," *Medical Imaging, IEEE Transactions on*, vol. 18, no. 6, pp. 481–495, 1999.

[127] F. Lin, T. Witzel, J. Polimeni, J. Hennig, G. Schultz, J. Belliveau, and L. Wald, "Parallel imaging technique using localized gradients (PatLoc) reconstruction using orthogonal mode decomposition," in *Proceedings 17th Scientific Meeting, ISMRM, Honolulu, HIth Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 4557, 2009.

[128] L. Tam, J. Stockmann, G. Galiana, and R. Constable, "Null Space Imaging: Nonlinear Magnetic Encoding Fields Designed Complementary to Receiver Coil Sensitivities for Improved Acceleration in Parallel Imaging," *Magnetic Resonance in Medicine*, vol. in press, 2011.

[129] L. Tam, J. Stockmann, and R. Constable, "Null Space Imaging: a Novel Gradient Encoding Strategy for Highly Efficient Parallel Imaging," in *Proceedings 18th Scientific Meeting, ISMRM, Stockholm*, p. 2868, 2010.

[130] L. Tam, J. Stockmann, G. Galiana, and R. Constable, "Magnetic gradient shape optimization for highly accelerated null space imaging," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 721, 2011.

[131] C. Juchem, T. Nixon, S. McIntyre, D. Rothman, and R. Graaf, "Magnetic field modeling with a set of individual localized coils," *Journal of Magnetic Resonance*, vol. 204, no. 2, pp. 281–289, 2010.

[132] C. Juchem, T. Nixon, S. McIntyre, V. O. Boer, D. Rothman, and R. de Graaf, "Dynamic multi-coil shimming of the human brain at 7T," *J Magn Reson*, vol. 212, pp. 280–8, Oct 2011.

[133] S. Littin, A. Welz, D. Gallichan, G. Schultz, C. Cocosco, J. Hennig, W. de Boer, and M. Zaitsev, "Planar gradient system for imaging with non-linear gradients," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, p. 1837, 2011.

[134] A. Welz, M. Zaitsev, F. Jia, M. Semmler, J. Korvink, D. Gallichan, H. Lehr, G. Schultz, J. Hennig, and J. Hennig, "Development of a non-shielded patloc gradient insert for human head imaging," in *Proc Joint Annual Meeting ISMRMESMRMB*, vol. 453, pp. 9208–9208, 2009.

[135] K. Layton, M. Morelande, P. Farrell, B. Moran, and L. Johnston, "A performance measure for MRI with nonlinear encoding fields," in *Proceedings of the ISMRM 19th Annual Meeting*, p. 482, 2011.

[136] K. Layton, M. Morelande, P. Farrell, B. Moran, and L. Johnston, "Performance analysis for magnetic resonance imaging with nonlinear encoding fields," *Medical Imaging, IEEE Transactions on*, vol. PP, no. 99, p. 1, 2011.

[137] G. Galiana, J. Stockmann, L. Tam, and R. Constable, "Using RF to create nonlinear virtual coil profiles," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, 2011.

Declaration

I herewith declare that I have produced this paper without the prohibited assistance of third parties and without making use of aids other than those specified; notions taken over directly or indirectly from other sources have been identified as such. This paper has not previously been presented in identical or similar form to any other American or foreign examination committee.

The thesis work was conducted from September, 2005 to September, 2011 under the supervision of R. Todd Constable at the Yale University Magnetic Resonance Research Center. It would not have been possible without the contributions of many individuals whose help was indispensable at every stage of the work. The division of labor is described below.

Todd Constable conceived the idea of using nonlinear SEMs to provide spatial encoding that optimally complements the information provided by RF receive coils. Together with the candidate, he developed the projection imaging approach described in this thesis. He oversaw and guided all subsequent stages of the project.

Gigi Galiana provided tremendous support throughout this project. She wrote the first field mapping, phase mapping, eddy current mapping, and O-Space imaging sequences used in this project. The candidate subsequently wrote his own phase mapping and O-Space imaging sequences to acquire the majority of the images shown in Chapter 3, but Dr. Galiana's work served as a crucial

template to build upon. She also helped debug the candidate's imaging experiments both at the level of pulse sequence programming and the underlying spin physics. In addition, Dr. Galiana performed a detailed study of through-plane phase evolution, published in (86), that helped clarify the reasons the candidate's field mapping approaches were not working for O-Space imaging. This contributed to the candidate's suggestion that our team instead use phase mapping to capture all spin phase evolution in-plane as well as through-plane. Dr. Galiana also helped fabricate the RF hardware described in Chapter 3.

Pelin Aksit Ciris performed the Center Placement comparative analysis discussed in Chapter 2, as reported in (59). The candidate suggested a matrix formalism to use for image reconstruction, but the decision to use the LSQR algorithm in Matlab was made by Ms. Ciris. She also helped develop the use of u-substitution to express the polar-coordinates form of the O-Space imaging equation as a Fourier transform in u. She also christened the new technique as "O-Space imaging" in honor of its ring-like frequency isocontours.

Leo Tam performed all work related to null space imaging referred to in this thesis; notably, the target and ideal field shapes shown in Chapter 5. Dr. Tam also assisted in fabrication of the RF hardware presented in Chapter 3.

Dana Peters checked the timing of the radial pulse sequence that underlies the O-Space sequence, as reported in Chapter 3. Using the method in (96), Dr. Peters verified the timing errors in the radial pulses were less than 1/4-point in k-space, too small to cause artifacts in the reconstructions.

Terence Nixon developed the Dynamic Shim Updating hardware and associated interfacing software used in this thesis. The hardware had been previ-

ously used for multivoxel spectroscopy and other applications (79). Mr. Nixon modified the hardware for use with the Z^2 insert coil described in this thesis.

Piotr Starewicz and William Punchard (Resonance Research, Inc.) designed and manufactured the Z^2 insert coil and delivered it to Yale as part of an NIH-funded business research partnership.

Peter Brown provided valuable advice on the fabrication of the transmit-receive RF coil, including training on making coaxial cables, baluns, and resonant RF circuits. Mr. Brown also helped orient the candidate for safe and efficient use of the MRRC machine shop.

Christoph Juchem helped the authors develop a field mapping sequence during early pilot experiments in O-Space imaging on a Bruker animal scanner. Dr. Juchem also shared his field map analysis code, FMAP, for use in processing the Bruker data. This assistance helped the candidate later develop field mapping methods for the Siemens scanner. Additional training on programming the Bruker scanner was provided by Robin de Graaf. Dr. de Graaf also provided helpful advice on the topics of shimming, field mapping, gradient calibration, and RF pulse design.

Hemant Tagare was an invaluable resource on the topic of mathematical image reconstruction. He suggested use of the Kaczmarz iterative algorithm for image reconstruction, an approach that has been invaluable to the candidate at every stage of the research presented here.

Vicente Parot and Pablo Irarrazaval (Catholic University, Santiago, Chile) developed the variable-order FrFT formalism presented in Chapter 4. They advised the candidate on how to extend this formalism for use in reconstructing

O-Space data. The candidate built upon their work fur further suggesting the use of a density compensation function, which improved the resulting image quality.

Lawrence Wald (Massachusetts General Hospital) provided the candidate with pre-amplifiers and matching circuit boards for use in the 8-channel RF transmit-receive array, speeding the processing of hardware development for the candidate's experiments.

YALE UNIVERSITY
NEW HAVEN, CONN.