

Computational Design and Fabrication of Portable MRI Systems

by

Patrick Christopher McDaniel

B.S., Massachusetts Institute of Technology (2013)
M.Eng., Massachusetts Institute of Technology (2015)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
August 7, 2020

Certified by
Lawrence L. Wald
Professor of Radiology, Harvard Medical School
Thesis Supervisor

Accepted by
Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Computational Design and Fabrication of Portable MRI Systems

by

Patrick Christopher McDaniel

Submitted to the Department of Electrical Engineering and Computer Science
on August 7, 2020, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering

Abstract

In this work, I have developed techniques for designing portable MRI scanners and have applied them to three portable systems for brain imaging. I first describe the procedure for designing a portable, low-field MRI scanner - in particular, how the constraints of compactness and portability affect the design of all system components (magnets, coils, sequences, RF pulses, and reconstruction schemes). I then describe the design of the principal hardware components of a portable MRI system: the B_0 magnet, the magnet shim array, the gradient coils, and the RF coils. This work makes novel use of numerical and computational tools for both sub-system design and physical construction. I next apply this paradigm to the design of gradient coils and a shim magnet array for portable whole-brain MRI scanner and demonstrate *in vivo* adult brain images. Finally, I describe two novel MRI scanners designed *ab ovo* using the approach described herein. The former is the "MR Cap", a single-sided MRI device designed for imaging over a reduced $8 \times 8 \times 3\text{cm}^3$ region of the adult brain; the latter is the "Helmet MRI", a whole-brain scanner optimized specifically for the head geometry.

Thesis Supervisor: Lawrence L. Wald

Title: Professor of Radiology, Harvard Medical School

Acknowledgments

I first must acknowledge my thesis committee: Larry, for giving me the latitude to try unproven things in lab, and yet still keeping my work oriented towards a clinical goal; Elfar, for first teaching me about MRI many years back and supporting my Master's research, and for serving as an advisor in countless other capacities since; Denny, for diving into the world of MRI and making me think about how my works connects with the outside world; and Clarissa, for sharing your technical insights and your own experiences in shoes similar to mine.

Three other individuals deserve mentioning for their specific contributions to this work: Thomas Witzel, for his endless knowledge of all things MRI and his willingness to share it; Matt Rosen, for teaching me much about low-field systems, experimental protocols, and mechanical craft; and Jason Stockmann, for insights into sequences, systems, and electronics. Magnets and pulses aside, I have also enjoyed our conversations on all sorts of topics.

I will also thank several other members of the Martinos community: Krista, for keeping the lab running; Erica, Eli, Melissa, and Monika for our many hours shared in lab; Bragi Sveinsson and Bo Zhao, for many hours of language instruction; and Choukri Mekkaoui, for always encouraging me to think beyond the walls of building 75. I am also grateful for the NTP fellowship for supporting much of my thesis work. Beyond this, there are too many more colleagues to list: I will simply say that getting to know them all as individual human beings has – more than anything else – made this a fulfilling place to work these past years.

I will acknowledge also my colleagues from MIT and Boston Children's Hospital. I have enjoyed our social outings (thanks, Elfar!); our journeys to far-off lands; and coffee hours at MIT. Among them, I would like to thank two people in particular: Prof. Audrey Fan, for mentoring me from the day I came across the acronym "PROM" until today; and to Jeff Stout, for being a friend and mentor through the entirety of my Ph.D (also to Bekah, for asking thought-provoking questions about my research, like: "But what's the point?").

Thank you to my acquaintances outside of MRI: barkeep Mike G and the Muddy crew; able seamen Joe and John; the Chinese department of MIT; my sheep-clear friends; roommates; and friends from my undergraduate days.

Last and most important of all: thank you to my mother, father, and sister for being so patient with me over the years.

List of Figures

1-1	Portable and high-end US, PET/SPECT, and CT systems	23
1-2	Early proposal for single-sided <i>in vivo</i> MRI	24
1-3	Range of MR system geometries	27
1-4	Single-sided devices for geophysical applications	30
1-5	Early mobile single-sided devices	31
1-6	Early Earth's field single-sided <i>in vivo</i> NMR	35
2-1	Defining an ROI for MRI system design	41
2-2	Head-Optimized target ROI and initial system geometry	42
2-3	Bore ordering for a conventional scanner and for the Head-Optimized MRI	45
3-1	Comparison of permanent and resistive magnet designs with a toy model	53
3-2	Permanent magnet torque, mechanical tolerance, and position error	54
3-3	Rectangular magnet block multipole field terms up to L=5	62
3-4	Exploiting magnet symmetry (MR Cap example)	63
3-5	Reducing ROI size by using only a boundary and by exploiting symmetry	64
3-6	Avoiding magnet intersections during magnet array design	66
3-7	Combining multiple stock magnets together to create custom magnet shapes and sizes	67
3-8	MR Cap mean B_0 vs. ROI field range plotted for 200 designs	71
3-9	Head-Optimized simulated B_0 maps for 3 different optimized shim designs	72
3-10	Test-print part used for magnet slot size calibration	72
3-11	Test-print part used for calibrating mechanical tolerances	74
3-12	Comparison of snap-off and dissolvable 3D printer support materials	75
3-13	Common overextrusions and burrs in SLA and FDM-printed parts	76

3-14	Design of the initial magnet former geometry and CAD model	77
3-15	Automatically-generated STL models of the optimized magnet design	78
3-16	Magnet former CAD model + STL binary operations	79
3-17	Splitting a magnet former into sections for 3D printing	80
3-18	Assembling printed sections of a magnet former into a single structure	81
3-19	Optimal order for adding magnet blocks to a Halbach-type magnet	82
4-1	Head-sized volume RF coil at $f_{RF} = 3.385MHz$ simulated with full-wave and magnetostatic solvers	85
4-2	Setup for a target field / surface stream function optimization problem	86
4-3	Wire packing density in a truncated gradient coil design	88
4-4	Scaling relations for solenoid coils with different numbers of winding layers	89
4-5	Test-print part for calibrating CAD model wire groove width	90
4-6	Example parametrized surface meshes for coil design	91
4-7	Combining a mesh structure with different surface geometry to create an enhanced mesh surface for coil design	92
4-8	Numerical artifact arising from a poorly-constructed target ROI/mesh geometry	93
4-9	Computation of discrete coil windings from a surface stream function by intermediate projection onto a flat surface	94
4-10	Converting numerical 1D wire winding paths into a volumetric model of the physical wire grooves	95
4-11	"Single-step" coil formers incorporating both electromagnetic and mechanical design elements	96
4-12	"Two-step" coil formers with separate parts for electromagnetic and mechanical interfacing elements	97
4-13	Design of a gradient coil former starting with a simple parametric design	98
4-14	Binary operations are performed between the coil former base and numerically-generated STL models of the wire windings	99
4-15	Wire-crossing locations for stepping from one coil winding to the next	100
4-16	Gradient coil construction from CAD model, to 3D-printed former, to wound coil	100

4-17	Practical tips for winding gradient coils using the method described herein . . .	100
5-1	Optimized Halbach magnet and internal geometry	104
5-2	Gx gradient coil designs created on two different surface geometries	105
5-3	Comparison of efficiency and bijective range for two Gx coil designs with different geometry	106
5-4	Gz gradient coil design surface	107
5-5	Gradient coil former mechanical interfacing elements	107
5-6	Optimized Gx and Gz coil designs	108
5-7	Simulated Gx and Gz coil B_z field maps	109
5-8	CAD model generation pipeline for the optimized Gx coil design	110
5-9	Sparse Halbach gradient coil former CAD model	111
5-10	Coil former that was 2mm too large for the magnet	112
5-11	Field-mapping the gradient coil in a 1.5 T MRI scanner	113
5-12	Measured Gx and Gz coil field maps	114
5-13	Sparse Halbach magnet shim tray locations and geometry	115
5-14	Sparse Halbach shim tray maximum allowable magnet size	115
5-15	Comparison of two shim array designs using different allowable magnet de- grees of freedom	116
5-16	Shim array target field pattern	117
5-17	Depiction of the optimized shim magnet array	117
5-18	Shim tray CAD model generation and physical assembly	118
5-19	Shim tray mechanical tolerance and angular position error	118
5-20	Shimmed and unshimmed B_0 field map comparison	119
5-21	RARE-type sequence used for Sparse Halbach imaging experiments	120
5-22	RF Tx/Rx coil used for initial imaging experiments	121
5-23	Phantom images acquired on the shimmed and unshimmed magnet	122
5-24	Initial test <i>in vivo</i> scans from the shimmed and unshimmed magnet	123
5-25	<i>In vivo</i> images reflecting subsequent RF coil, sequence, and reconstruction development	124
6-1	Depiction of the MR Cap system concept and imaging ROI	128
6-2	The MR Cap with principal system components labeled	129

6-3	MR Cap magnet design procedure	130
6-4	MR Cap magnet design available degrees of freedom	133
6-5	MR Cap optimized magnet design	134
6-6	MR Cap constructed magnet	135
6-7	MR Cap raw and filtered measured B_0 field maps	136
6-8	Comparison of MR Cap simulated and measured B_0 field maps	137
6-9	MR Cap gradient coils designed with 1st and 2nd-order target fields	138
6-10	MR Cap Gy and Gz gradient coil designs	138
6-11	MR Cap gradient coil CAD model generation pipeline	139
6-12	MR Cap constructed Gy and Gz coils	140
6-13	Field-mapping the MR Cap Gy coil	140
6-14	MR Cap gradient coil field maps with polynomial fitting	141
6-15	MR Cap RF coil design surface	141
6-16	MR Cap constructed RF Tx/Rx coil	142
6-17	MR Cap imaging setup with system elements labeled	143
6-18	Phantoms used for depth profiling and imaging experiments	144
6-19	RARE-type sequence used for MR Cap experiments	145
6-20	Narrow 1D phantom depth profiling results	147
6-21	Volumetric phantom depth profiles	147
6-22	Multislice images acquired with the volumetric imaging phantom	148
6-23	Artifacts observed in MR Cap images	150
7-1	Head-Optimized MRI concept drawing and constructed system photograph	156
7-2	Target ROI definition using a high-field brain MRI scan	156
7-3	Views of the Brain-Optimized MRI with spatial dimensions	157
7-4	Illustration of the many components and mechanical design elements in the Head-Optimized MRI	159
7-5	Vector and block models of the starting B_0 magnet design shown with the target ROI	160
7-6	The optimized B_0 magnet design	161
7-7	Plot of the distribution of magnet block sizes in the optimized B_0 magnet design	161

7-8	Comparison of the simulated B_0 fields from the continuous- and discrete-space magnet designs	162
7-9	Simulated B_0 maps for magnet designs with perturbed B_r , μ_r , and geometry	163
7-10	Comparison of B_0 field maps for the ideal design and a Monte-Carlo-perturbed design	164
7-11	Assembly of the 3D-printed magnet former	165
7-12	Adding the "booster" magnets to the former	166
7-13	Populating the magnet former with the magnet blocks	167
7-14	The fully-constructed B_0 magnet	168
7-15	Measured B_0 field maps compared with simulation	169
7-16	Head-Optimized MRI shim array geometry	170
7-17	Measured field maps of the shimmed magnet compared with simulation and the unshimmed magnet	171
7-18	Shim tray CAD model and 3D-printed parts	172
7-19	Shim tray with spray-on RF shield	173
7-20	Shim trays assembled inside the B_0 magnet with completed RF shield	173
7-21	Head-Optimized MRI gradient coil design surface	174
7-22	Comparison of Gx coil designs using different surface geometries	174
7-23	Optimized Gx, Gy, and Gz gradient coil designs	175
7-24	Simulated Gx, Gy, and Gz coil field maps	176
7-25	Gradient coil former CAD model	176
7-26	Constructed and wound gradient coils in the B_0 magnet	177
7-27	Measured gradient coil field maps	178
7-28	Head-Optimized MRI RF coil design surface and ROI	179
7-29	Comparison of x- and y-polarized RF coil designs	180
7-30	Optimized y-polarized RF coil design	181
7-31	CAD model, printed parts, and completed coil for the RF helmet coil	182
7-32	Measured RF coil $B_{1,xy}$ field maps in and out of the bore compared with simulation	183
7-33	RF coil transmission (S_{21}) plot vs. frequency	183
7-34	RF shield noise tests	184
7-35	RARE-type sequence used for imaging experiments	185

7-36	Imaging phantoms and corresponding high-field MRI scans	187
7-37	Phantom images acquired with different readout directions	188
7-38	Phantom images acquired with different readout strengths	188
7-39	High-resolution phantom images acquired with two readout directions	189
7-40	Observed gradient coil heating and center frequency drift	189

List of Tables

5.1	Sparse Halbach Gx and Gz gradient coil performance metrics	109
5.2	Sparse Halbach shim design allowable magnet options	125
5.3	Sparse Halbach initial imaging experiment sequence parameters	126
5.4	Sparse Halbach <i>in vivo</i> imaging experiment sequence parameters	126
6.1	MR Cap block available block design options	131
6.2	Manufacturer-supplied tolerances for magnet block linear dimensions and re- manent flux density	132
6.3	MR Cap B_0 magnet performance metrics	132
6.4	MR Cap imaging sequence parameters	146
6.5	MR Cap estimated total system cost	153
7.1	Head-Optimized MRI gradient coil performance metrics	171
7.2	Head-Optimized MRI RARE-type sequence parameters	186

0.1 Thesis Outline

This section lists the chapters of the thesis and briefly states what each one talks about.

Chapter 1: Introduction

I begin by arguing that there exists a clinical need for low-cost, portable, MRI systems. I then summarize the reasons why MRI scanners are expensive and immobile and propose how one might design low-cost, compact systems. Next, I summarize the existing literature on purpose-built, portable, and low-field MRI devices. Finally, I provide an introduction to the field of "Single-Sided NMR", in which low-cost, portable NMR devices have been developed for a variety of applications.

Chapter 2: System-level Design Considerations

Here I describe the general system-level design process for a portable MRI system. One first defines the geometry of both the scanner and of the imaging volume - practical considerations relating to scanner operation affect both. Next, I describe how one goes from a system imaging performance goal to specification for each of the system components. I further look at how the different parts of the scanner can be arranged, with an emphasis on opportunities and tradeoffs that exist in low-field, portable systems.

Chapter 3: Design and Construction of Permanent Magnet Arrays

I describe my approach to design and construct B_0 magnets and shim arrays for portable MRI scanners. I discuss the advantages and challenges of using permanent magnet material, methods for designing and simulating permanent magnet arrays, and practical manufacturing considerations. I then show methods for constructing physical prototype permanent magnet arrays.

Chapter 4: Design and Construction of Gradient and Low-frequency RF Coils

Here, I introduce the target field/current stream function approach for designing quasi-static coils. I present how I design gradient and RF coils with this approach, with an emphasis on the practical tradeoffs involved with short-bore or single-sided coils. I then show methods

for constructing the numerically-optimized coil designs using computational manufacturing methods such as 3D printing.

Chapter 5: Gradient and Shimming Hardware for a Portable Brain Scanner

I present the design and construction of two system components for the "Sparse Halbach", an existing portable brain MRI system: a pair of gradient coils for two axes of blipped phase encoding and a permanent magnet shim array for improving the built-in B_0 encoding field. The system performance is assessed, and phantom and *in vivo* images are presented.

Chapter 6: The MR Cap: A Single-sided MRI Device for Point-of-Care Applications

I introduce the "MR Cap" - a bicycle helmet-sized MRI system for imaging a region of the brain. I describe the design of this system from the ground up, including: the B_0 magnet, the gradient coils, the RF coil, and the imaging sequences/encoding schemes. Phantom images are shown, demonstrating the volumetric imaging capabilities of a single-sided MRI device.

Chapter 7: An Optimized Head-shaped MRI system for Whole-brain Imaging

I describe the "Head-Optimized MRI", a system optimized for performance over a head-shaped ROI and conforming a geometry that optimally fits over the human head. I show the design and construction of the B_0 magnet, gradient coils, RF coil, B_0 shim array, and mechanical apparatus. Phantom images are presented, demonstrating that the system can image a head-sized object in reasonable time.

Chapter 1

Introduction

1.1 Abstract

This chapter describes the existing clinical need for portable, low-cost, and point-of-care MR imaging; and reviews work in this area to date. I first cover clinical applications that could benefit from MRI, but for which existing scanners are not adequate because of cost, siting, or workflow reasons. I then give a brief overview of how conventional MRI scanners are built and operated, with an emphasis on how they could be made smaller and more portable. I next summarize prior work in the field of portable MRI scanners. Finally, I provide an introduction to the field of the "single-sided NMR", as that area has much to teach us about building portable, compact MR devices for unusual applications and settings.

1.2 Other Resources

1.2.1 Literature Reviews

The main motivation of this work is the development of portable, low-cost MRI scanners for situations ill-served by traditional MRI. Several recent reviews have looked at this problem from a physics perspective [1]; from a global economic perspective [2]; and from a system engineering perspective [3]. Taken together, these reviews provide a comprehensive introduction to the state of portable MRI as a field.

1.2.2 Technical Resources

I will not describe the basic physics of electromagnetism or MRI here, as many existing texts provide useful introductions to both. I have personally found [4] to be useful reference for EM; and have found [5, 6] to provide a good introduction to MRI. Esoteric or unusual technical principles (eg the magnetic pole density formalism, current stream functions) will be described where needed.

1.3 Clinical Need for Low-cost, Portable MRI systems

MRI is widely used for clinical diagnosis and has become a routine part of modern medicine. It is in particular used for brain imaging, as it enables a variety of soft-tissue contrasts and can easily image inside the skull. However, many clinical scenarios are poorly-served by conventional MRI systems and might benefit from a portable, low-cost, or low-field MRI scanner. Here, I describe the limitations of conventional MRI scanners and list some potential uses for a portable brain MRI system.

1.3.1 Detection of Cerebral Hemorrhage

Portable, low-field MRI systems may be useful for diagnosing cerebral hemorrhage (eg epidural vs. subdural) in cases of trauma [7] and in post-operative settings [8]. These are often either time-sensitive settings or involve critical patients. As such, they would benefit from point-of-care scanner or a system with reducing screening requirements. A portable MRI system could also enable continuous monitoring in these settings, which is not practical with conventional MRI machines.

1.3.2 Improving Access to MRI

Many developing-world settings suffer from a lack of access to diagnostic MRI [2], and could benefit from an MRI scanner that was less expensive and easier to site and maintain than a typical system. For example, inadequate imaging accessibility impedes the diagnosis and treatment of pediatric hydrocephalus in developing-world settings [9].

1.3.3 Acute Stroke Diagnosis

Portable, low-field MRI scanners could improve the diagnosis of acute ischemic stroke (AIS). MRI is better than CT for detecting and characterizing the severity of AIS, but routine use of MRI is not cost-effective at present with conventional scanners [10]. Reducing patient transport and pre-MRI screening time is also of critical importance during cases of acute stroke. Portable, easy-to-site, scanners might improve access to MRI and reduce transport times; an inexpensive scanner might reduce procedure cost; and lower-field, compact devices might have reduced pre-scan screening requirements.

1.3.4 Imaging of Critical Patients

A scanner that minimized the transport requirements for critical patients would be of value [11]. Namely, bringing the scanner to the patient could be both less risky for the patient and utilize ICU staff for efficiently. This might be possible with a portable scanner that could be wheeled to the patient's bedside. An example of such a case could be a post-operative patient being monitored for a bleed, as described above.

1.4 Building a Smaller MRI Scanner

Here I broach the paired questions: Why are normal MRI scanners so large and expensive? and: How can we make MRI scanners smaller and cheaper? Scanners' size and cost ultimately follow from their technical performance requirements. Making MRI systems smaller, more portable, and less expensive requires rethinking just how sophisticated that technical performance needs to be. In answering these questions, it is also illustrative to examine both the historical course of development of the modern "conventional" MRI scanner and examples of "hand-held", "point-of-care" devices in other imaging modalities.

1.4.1 Technical Perspective: Why are MRI Scanners Large and Expensive?

Scanner cost and size ultimately result from typical MRI systems being whole-body imaging systems that produce high-resolution images in a variety of contrasts. This follows from the historical development of clinical MRI: the first *in vivo* scans were "body" scans (as opposed to "head-only" scans) and employed resistive magnets big enough to accommodate a human

torso. There was not initially interest in targeted-organ (eg head-only) devices. Subsequent improvements in scanner performance accrued from increasing magnet field strength via more-sophisticated resistive magnets and, eventually, superconducting magnet designs. It was generally found that the installation and operational costs of superconducting magnets were offset by their faster acquisition and improved imaging performance in the case of whole-body imaging [12]. Thus, we are at our present state of affairs, where a typical MRI scanner contains a superconducting magnet large enough to accommodate a human torso.

Furthermore, conventional MRI scanners can produce images using an immense variety of pulse sequences. Some sequences require fast scanning (eg fMRI), chemically-selective saturation pulses, compensation for blood flow and fast diffusion encoding. These sequences require strong, fast-switching gradient coils and strong, homogeneous magnetic fields, resulting in expensive, immobile scanners.

1.4.2 Components of a Conventional MRI Scanner

The B_0 magnet is the main determinant of system performance in an MRI scanner. Its field strength determines available SNR. Its homogeneity determines image distortion, acceptable encoding schemes, the available SNR/scan time/resolution tradeoff, and RF system requirements. These considerations, in turn, further constrain the design of the gradient system, RF system, and pulse sequences. The resolution (typically $\approx 1mm^3$), scan time (typically ≈ 5 min), and variety of contrasts (T_2^* , T_2 , DWI, T_1 , perfusion, etc.) and encoding schemes (spiral, EPI, GRE, etc.) require a strong main field and high homogeneity over a large ROI. For example, the commercially-available Siemens Avanto and Prisma systems have B_0 fields of 1.5T and 3T with RMS homogeneities of 0.75ppm and 0.2ppm over a $D = 40cm$ spherical volume [13, 14]. Attaining these requirements requires a large superconductor-based magnet. Conventional superconducting magnets require liquid helium to operate, along with all the cryogenic infrastructure that entails (insulation, cold heads, quench pipe). The Prisma, for instance, is 1.98m long, weighs 10 tons, and requires 10's of kW of power during operation [14].

Other performance considerations influence the system installation footprint: a large, high-field magnet needs a large footprint for a 5-Gauss exclusion area (for patient safety); RF shielding (to shield both external signal interference from coupling in, and transmit pulse emissions from radiating out) requires a shielded room large enough to contain the magnet,

patient table, and scan subject.

The wide variety of imaging contrasts and encoding schemes also requires a large, high-power gradient system. The Prisma gradient coil requires up to 900 Amperes of current and 2250 Volts for max single-channel gradient amplitudes of $80 \frac{\text{mT}}{\text{m}}$ and slew rate of $200 \frac{\text{T}}{\text{m}\cdot\text{s}}$ [14].

In summary, the size, cost, and complexity of MRI scanners follow directly from their performance capabilities. A smaller, more-portable magnet will need to eschew superconductors and cryogenics, in turn requiring it to be smaller, lower-field, and/or less homogeneous. Reducing the system power consumption further requires eliminating the high-power gradient coils, reducing the available toolbox of imaging sequences. These are inherent tradeoffs in the design of portable MRI scanners.

1.4.3 Portable, Low-cost Devices in Other Medical Imaging Modalities

The tradeoff between system performance and system cost/complexity also exists in other medical imaging modalities. In each of ultrasound (US), positron emission tomography/single-photon emission computed tomography (PET/SPECT), and computed tomography (CT), there exists a range of devices along this tradeoff axis. That is, there exist several devices that exploit the same physics, but vary from high-performance, expensive systems to targeted-use, low-cost devices.

Ultrasound (US)

Review papers have examined the design of portable, low-cost ultrasound systems [15]; its use in developing-world settings [16]; its use by physicians (as opposed to technicians) [17]; and its use by EMS personnel [18]. These devices have been found to perform generally well when compared with high-end devices [19]. At present, portable ultrasound devices from many vendors are available on the market [20, 21, 22, 23, 24, 25].

The physics of US result in US probeheads naturally being compact compared with an MRI magnet, even for high-end systems. However, the size of complete US systems ranges from something the size of a cart (for high-end systems) to something that fits in a handbag (for point-of-care systems). Figure 1-1a shows the form factor difference among ultrasound systems.

PET/SPECT

PET and SPECT imaging work by detecting gamma ray photons that ultimately originate from a radiotracer injected into the patient. A typical PET/SPECT system is a whole-body imaging device similar in size to a conventional whole-body MRI system. However, the same principle of gamma photon detection has been exploited in a variety of portable "gamma camera" devices. There exist several reviews of these devices [26, 27, 28]. Portable gamma cameras have thus far mainly been developed for intraoperative applications, including testicular cancer sentinel node biopsy [29] and tumor margin assessment in breast cancer surgery [30, 31, 32, 33]. Figure 1-1b shows a whole-body PET scanner and an intraoperative gamma camera probe for reference.

Computed Tomography CT

CT utilizes the variation in x-ray photon absorption by human tissues and/or contrast agents to form images. A series of 2D projection images is acquired, from which a 3D image is then reconstructed. Several reviews have compared traditional (ie whole-body) and portable CT scanners [34]; have examined portable head-only CT systems [35, 36]; and have assessed the value of portable CT in cases of thoracic disease [37]. At present, several portable CT scanners are commercially available [38, 39, 40]. Figure 1-1c shows a whole-body CT scanner and a portable, head-only CT scanner for comparison.

Additionally, it is illustrative to compare a projection X-ray radiograph with a 3D CT scan. Both types of image exploit the same physical principles (namely, that of X-rays and their interaction with human tissues). However, distinct, purpose-built devices exist for producing either projection (2D) or 3D radiographs. Each family of devices is tailored to particular applications, and both types of device are widely used clinically in different settings. This paradigm represents a template for how purpose-built, portable, or low-cost MRI devices might coexist with established, whole-body scanners.

1.5 Existing Low-cost, Compact, and Portable MRI Systems

There exists a sizable body of work on the development of portable, low-cost, or purpose-built MR systems for clinical applications. This area continues to expand, with several such devices having been developed in the past few years.

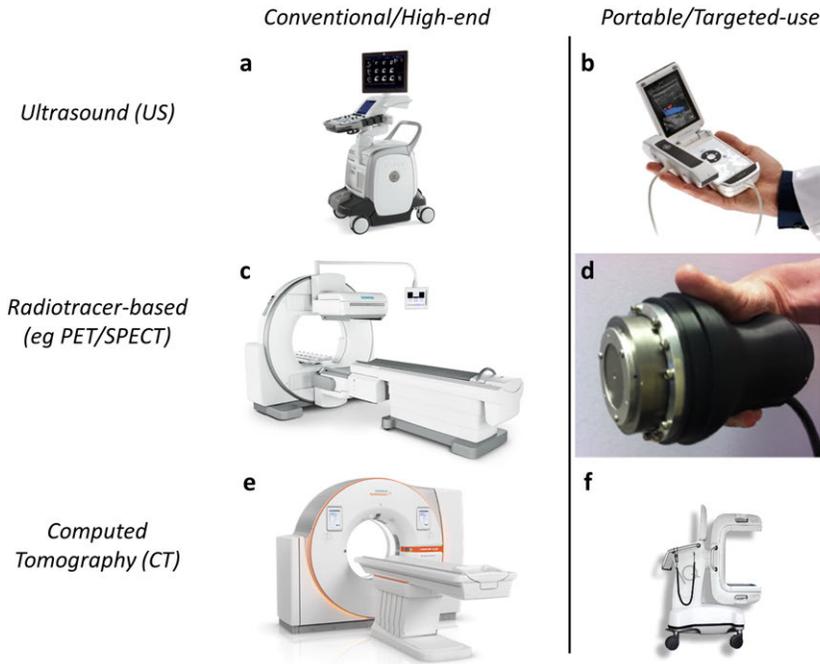


Figure 1-1: a. High-end cardiac Doppler US scanner [41]. b. Handheld point-of-care US scanner [22]. c. Whole-body SPECT scanner [42]. d. Handheld intraoperative “gamma camera” probe [27]. e. Whole-body CT scanner [38]. f. Portable head-only CT scanner [40].

1.5.1 Early interest

There already existed some inklings of awareness of the possibility of portable, low-cost MRI devices in the early 1980s. The Los Alamos well-logging NMR group noted that rastering a "sweet spot" magnet over a patient could potentially produce an *in vivo* image "in which the subject/patient need not enter the magnet" (Figure 1-2) [43]. Georges Béné in the late 1970s and early 1980s measured the relaxation times of various human tissues *in vivo* in the Earth's-field range [44, 45, 46], and argued that a low-field system "might solve the problems posed by the realization [of an MRI apparatus] for the use in medium size hospitals under reasonable technical and financial conditions" [47]. However, neither line of inquiry resulted in a practical clinical MRI device.

1.5.2 Open MRI systems

Open MRI systems are often designed for body imaging and are not typically considered portable or compact devices. However, like with portable MRI devices, these systems incorporate tradeoffs between imaging performance and other aspects of scanner operation (eg. patient comfort, bore access, or scanner siting requirements). Some open MRI systems are

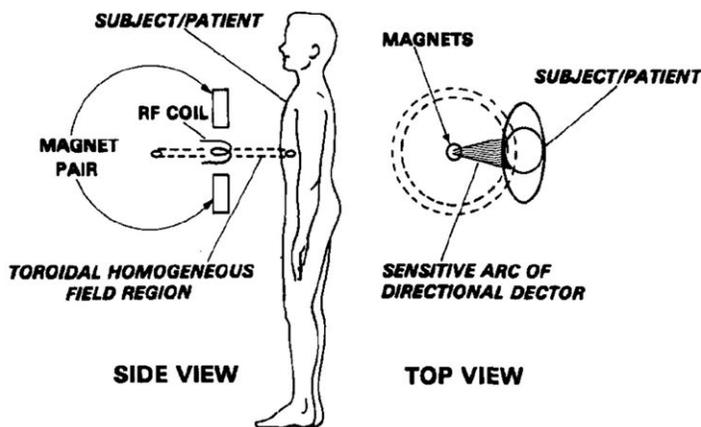


FIG. 9. New NMR imaging technique.

Figure 1-2: Proposal for *in vivo* MRI with an “inside-out”, rastered NMR device from 1980 [43].

further designed as targeted-use devices, especially for intraoperative use.

One early example from 1995 is a superconducting magnet-based MRI system for intraoperative MR scanning based on a pair of planar superconducting coils [48]. Resistive and permanent magnet-based open MRI scanners have also been developed [49]. Open MRI scanners may improve patient comfort, and were until recently produced by major MRI vendors at up to 1.0T. However, they were still bulky and difficult to site, and their benefits in many cases did not offset the reduction in image quality [12].

Open MRIs have also been used for neurosurgery [50, 51], with the open geometry providing greater patient access. Relatedly, MR-LINACs have been developed for image-guided radiotherapy. These scanners require more bore access than a typical scanner for the LINAC elements, and as a result have employed biplanar dipole magnets [52] and gapped cylindrical magnets [53, 54]. In general, all these systems have increased bore access, but reduced field strength when compared with a conventional MRI scanner.

1.5.3 Brain imaging systems

There has recently been significant interest in head-only MRI scanners, which has resulted in the development of a range of devices. At the high-performance end, these include high-field, superconducting systems with reduced footprints and cryogen requirements developed by Siemens (the Allegra system), GE [55] and Synaptiv [56]. Additionally, a research collaboration in the USA has been developing a compact, high-temperature superconductor-based

brain MRI scanner [57]. These systems are smaller and easier to site than typical MRIs and have field strengths of 1.5T to 3T. However, they are too large to be considered "portable" devices. On the other hand, portable scanners employing permanent magnet designs have been developed for point-of-care or developing-world use [58, 59, 60, 61]. These devices have field strengths between 50mT and 100mT and are true "portable" devices.

1.5.4 Extremity systems

Dedicated extremity MRI systems have also seen development over the years. These applications might seem a natural fit for compact MRI designs, as the arms and legs are significantly smaller in diameter than a whole human torso. However, clinical extremity musculoskeletal MRI often uses images with high in-plane resolution, which can be a challenge. Such devices have included: the MagneVu system [62]; the PoleStar iMRI (also used intraoperatively) [63]; the GE/ONI Ortho MRI [64]; the Esaote Oscan [65]; and a device developed at the University of Tsukuba [66, 67].

1.5.5 Prostate

A prostate MRI system has recently been underdevelopment at Promaxo [68]. This device utilizes a single-sided design optimized for imaging performance over an anatomical region overlaying the prostate.

1.5.6 Non-imaging systems

There has been further interest in novel non-imaging MR systems for clinical applications. This has included a single-sided MR device for monitoring tissue hydration state via T_2 -relaxometry [69, 70] and a single-sided system for assessing the placenta *in vivo* [71]. There has also been preclinical use of single-sided NMR devices for intraoperative intestine diffusion assessment [72].

1.5.7 Summary and Observations

There has been a proliferation of portable and low-cost MR devices over the past two decades, and especially within the past five years. The range of devices includes systems designed for the head, extremities, prostate, and placenta; it includes imaging and non-imaging systems;

and includes both commercially-available and research systems. These systems all make different tradeoffs between performance, portability, and cost. For example, the head-only 3T scanner was designed to be "movable in a standard freight elevator" prior to installation [55]; while other low-field (50-100 mT) head-only scanners were designed to be wheeled through a typical doorway [58, 59].

1.6 Single-sided MR Systems

The field of Single-Sided NMR devices has produced many unusual MR devices for chemical, analytical, industrial, agricultural, and other applications. Many of these devices are designed for specific use cases or are portable, compact, or inexpensive. As such, this area of human endeavor has much to teach the designer of portable MRI scanners. It has developed largely independently of traditional MRI scanners, and thus merits its own expository section.

In this section, I will attempt to define "single-sided" in this context and provide a comparison with conventional MR devices. I will then describe the existing space of "single-sided" MR devices and their use, with an emphasis on existing biomedical applications.

1.6.1 Other Resources

There have been several reviews and an entire book written about single-sided MR. In terms of comprehensiveness, none can top the 2008 review by Blumich, Perlo, and Casanova with its 689 references [73]. A 2011 book by the same authors also provides a good introduction to the field [74] (this book has fewer references than the aforementioned review paper).

1.6.2 "Single-sided" MR vs. Traditional MR

What distinguishes a "single-sided" device from a conventional MR device? Informally, one can state that a traditional "enclosed" device surrounds the sample under test, while a single-sided device has an open design and does not physically surround the sample. This geometric property of "enclosedness" represents a continuum on which we can place any MR system and relates to many aspects of scanner performance. Figure 1-3 illustrates this concept of "enclosedness". Scanners range in a continuous manner from long-bore whole-body MRI machines; through short-bore MRI's, "open plan" MRI's, "unilateral" NMR-MOUSE-type

devices; and ultimately to "inside-out" borehole prospecting devices. I typically think of "single-sided" as referring to systems where a contiguous half of the space overlaying the sensitive volume is free from the field-generating parts of the scanner (highlighted in Figure 1-3).

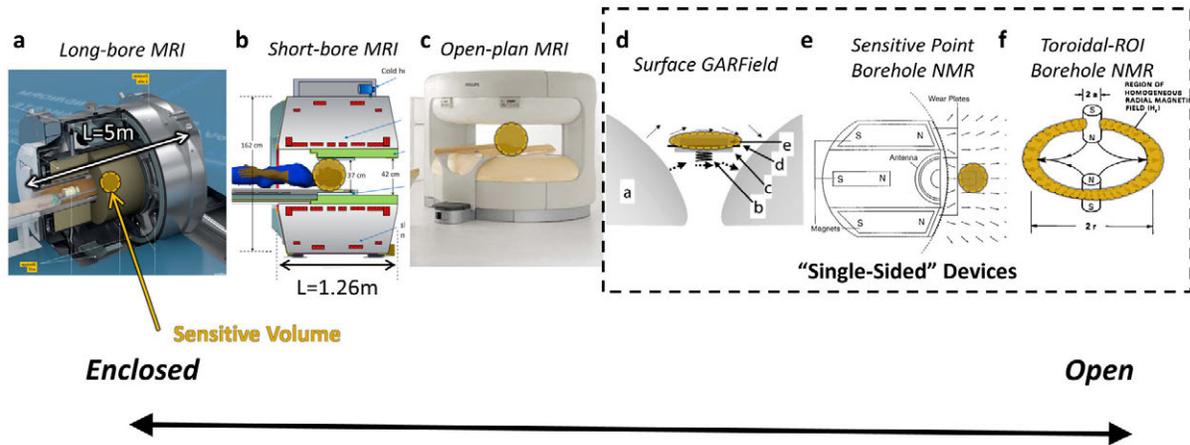


Figure 1-3: NMR and MRI device geometries differ greatly in terms of how much the field-generating components surround the sensitive volume. This can be visualized by plotting MR systems on an “Enclosed-Open” axis. a. A long-bore MRI scanner nearly completely surrounds the sensitive volume [75]. b. A truncated “short-bore” head-only MRI may be shorter, surrounding the imaging volume less [55]. c. Open MRI scanners allow more access to the patient/imaging volume than “closed-bore” scanners [12]. d. The Surface GARfield device only partly surrounds the sensitive volume [76]. e. The “sensitive point” downbore NMR system developed by Schlumberger has a convex geometry, and doesn’t surround the sensitive ROI at all [77]. f. “Inside-out” NMR systems have a sensitive volume that surrounds the NMR probe [43].

There are a couple important points relating to this "space" of magnet/system geometries. First, the term "single-sided" describes not just from the apparatus itself, but the combined geometry of the apparatus and the sensitive region. For example, a conventional solenoidal NMR-spectroscopy magnet is not "single-sided" because the sensitive region lies at the center of the solenoid. However, if the sample is placed at the opening of the solenoid (for instance, to obtain a higher field gradient as in [78]), then one might say the same machine is being operated in "single-sided mode" because the sensitive ROI has been moved. Relatedly, designing an NMR or MRI system requires one to jointly define both the sensitive volume and the device volume. This concept is expanded upon in Section 2.3.

Secondly, the degree of "openness" of a system relates to many aspects of its performance, especially with regard to field efficiency and uniformity (see also Section 2.3). The more

"single-sided" a device with given size or weight is, the lower its field efficiency and/or field uniformity will be when compared with a more-enclosed design. Prior work has explored the geometry vs. field homogeneity/efficiency tradeoff in the design of MRI magnets [79, 80] and gradient coils [81], whereby a longer (ie "more-enclosed") design is associated with improved performance. This trend continuous past short-bore cylindrical designs into the realm of "single-sided" geometries.

1.6.3 Review of Existing Single-Sided NMR Devices

In this section I summarize the existing range of single-sided NMR devices. These systems vary greatly and represent over 60 years of instrumentation development. Beyond this, the following book chapter contains a fine overview of single-sided NMR, including a summary of the historical literature [73].

Geophysical Prospecting

The ancient human lust to strike upon riches in the heart of the earth has not spared the field of NMR. The earliest single-sided, remote, and inside-out MR devices were developed for geophysical prospecting, particularly to look for subsurface water and oil. The earliest such proposed device is a patent filed in 1956 by Russell Varian [82] who had previously with Martin Packard demonstrated the first prepolarization + Earth's field NMR measurement in 1953 [83]. This proposed device consisted entirely of a large Tx/Rx "radio-frequency" coil (actually at $f_c \approx 2000\text{Hz}$ to match the Earth's magnetic field) that was placed on the ground. To operate the device, one would play an RF pulse to excite the proton spins in underground liquid that had achieved thermal polarization with the Earth's magnetic field. One would then switch the device to "receive mode" and would measure the voltage induced in the RF coil by the precessing spins. It does not appear that such a device was ever actually built at the time; the first instance of such a physical realization seems to have come at the hand of one A G Semenov in 1982 [84]. I confess that I have personally been unable to find this publication and that it apparently only exists in Russian. A slightly later account of subsurface NMR measurements obtained with a surface device can be found in the 1996 work of Oleg A Shushakov [85]. Nowadays, there exist commercially available devices that are conceptually identical to Varian's 1956 design [86]. The Numis Poly (Iris Instruments), for example, consists of a $150\text{m} \times 150\text{m}$ square RF coil that is laid on the

surface of the ground and can measure water to a specified depth of 150m (see Figure 1-4a). Reviews of the use of surface NMR for subsurface geophysical measurement can be found in Section 4.7 of [87] and in [88].

I will also mention a 1965 patent for an airplane-mounted groundwater prospecting NMR device [89], but only for its novelty. It strikes me as wholly impractical based on both NMR and aviation considerations.

Our unquenchable thirst for oil also fueled the development of "downbore" or "inside-out" NMR devices. These apparatus are lowered down boreholes regularly drilled during prospecting, and measure signals from liquids in the rock just outside the hole. Patents for prepolarization + Earth's field downbore systems were filed in 1956 by Varian Associates [90]; in 1960 by Mobil [91]; and in 1965 by Chevron [92]. The first downbore NMR experiments were performed in 1960 by Brown and Gamson (of the Chevron patent) and demonstrated the ability to discern water and oil *in situ* via relaxation curves [93]. Later in 1977, John C Cook reviewed the state of downbore NMR, looking at technical considerations, oil/water prospecting applications, and ESR-based mineral prospecting applications [94]. Cook further proposed a downbore NMR system consisting of a pair of longitudinally-aligned, antiparallel magnets that produce a toroidal "uniform field" region outside the borehole and an RF coil sensitive to that ROI (Figure 1-4b). This design would be adopted and refined by researchers at Los Alamos a few years later [95, 96, 43] and subsequently Schlumberger [77]. Note that initial iterations of the Schlumberger design used an axially-symmetric magnet with a toroidal ROI [43], while further improvements resulted in a design with a more-compact, simply-connected ROI as used by the NMR-MOUSE and other prototypical single-sided devices [77].

In recent years, borehole NMR has also been used as a tool in hydrogeology [97]. This has spurred the development of new downbore NMR devices tailored for the needs of scientific hydrology, such as the 2"-diameter (51mm-diameter) instrument in [98] (shown in Figure 1-4c).

Finally, I must mention a recent single-sided NMR device designed to detect oil spills under Arctic sea ice [99] (shown in Figure 1-4d). This is not strictly speaking a "prospecting" device. However, its intended use is for the oil industry, and the system itself has much in common with "subsurface NMR" devices. This device consists of two principal components: a Tx/Rx "RF" coil, and a prepolarization field coil. These coils (diameter \approx 6m are placed

atop sea ice and make NMR measurements to assess whether or not oil lies beneath the ice. The system was designed to detect oil beneath 2 meters of sea ice.

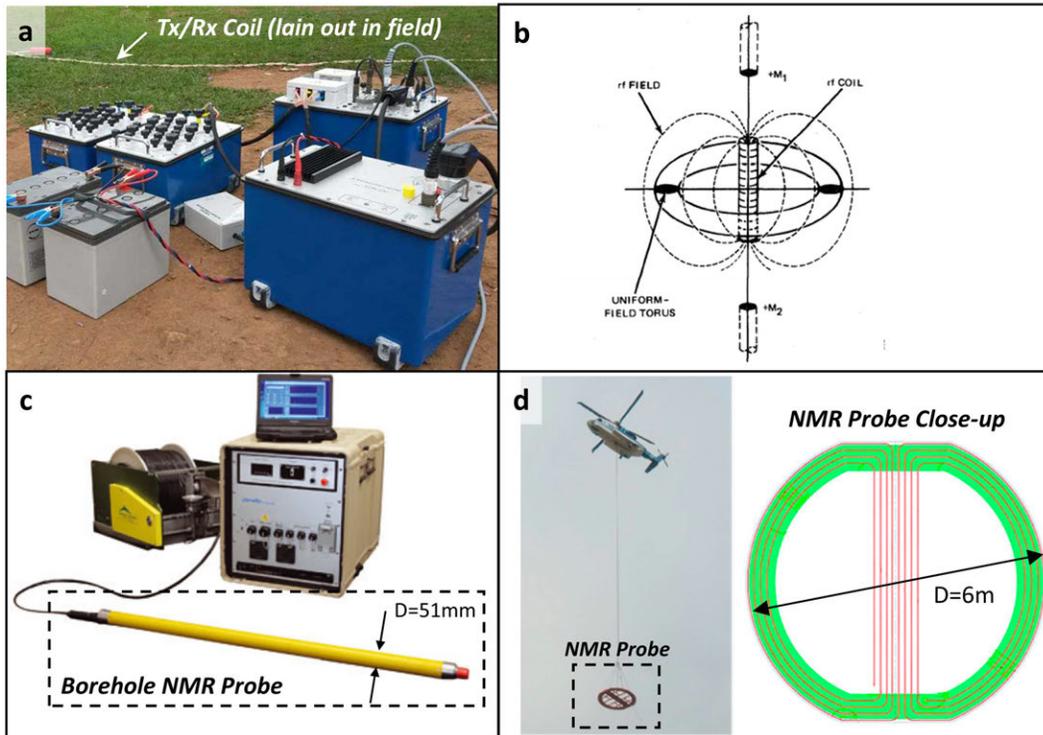


Figure 1-4: Geophysical and geological NMR systems vary greatly in size, shape, and mode of operation. a. The Numis-Poly is an Earth's field system consisting of an $L = 150\text{m}$ square Tx/Rx coil that's placed on the ground to detect subsurface water [86]. b. The axially-antisymmetric downbore magnet design with a toroidal homogeneous-field ROI was initially described by Cook et al in 1977 [94]. c. Downbore NMR probes have recently seen use in scientific hydrology, including with the $D = 51\text{mm}$ probe shown [98]. d. A prepolarization + Earth's field system was recently developed to detect oil spills under Arctic sea ice, and consists of a $D = 6\text{m}$ NMR probe that can be carried by helicopter [99].

Early Single-Sided NMR Systems

Early "single-sided" NMR devices (the first described as such) were published in the 1980s for several applications. These included: a device for measuring soil water content with a tractor-pulled NMR device [100]; a device for determining moisture levels in structural concrete [101, 102], particularly in bridges; devices for detecting explosives [103] and testing soild rocket motors [104]; and devices for measuring grain kernel moisture content as it moves through processing [105, 106]. This work arose from the desire for non-destructive testing (as with explosive detection); for in-situ testing (as with soil moisture assessment); or both (as with bridge concrete moisture measurement). These constraints precluded the use of

conventional NMR instruments and necessitated the development of "single-sided" and/or portable NMR instruments with unusual geometry. This led to the first mobile, single-sided devices. Figure 1-5a illustrates the subsurface water-measuring probe; Figure 1-5b shows the device for measuring moisture in bridge concrete.

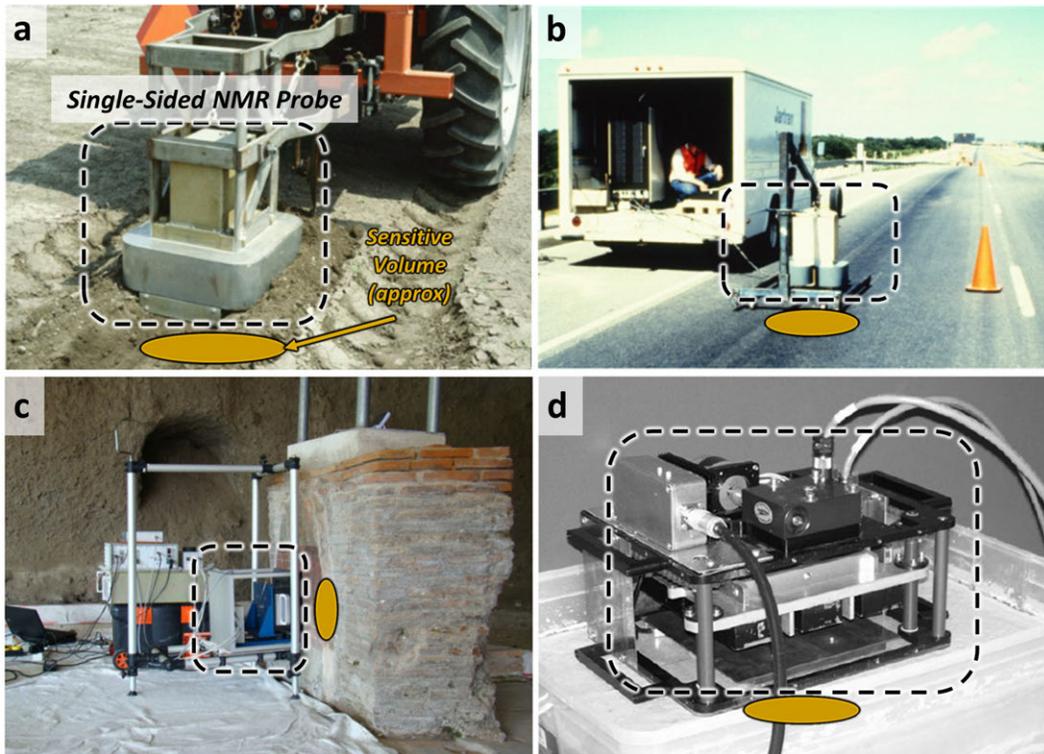


Figure 1-5: a. Tractor-pulled single-sided NMR probe for subsurface soil moisture measurements [100, 73]. b. Mobile single-sided NMR scanner for detection moisture in bridge deck concrete *in situ* [101, 102, 73]. c. The NMR-MOUSE single-sided device, shown being used to profile ancient Roman frescoes *in situ* [107]. d. The Surface GARfield single-sided NMR probe [108].

A particularly interesting example is the soil hydration system developed in [100]. This work used the inherently-inhomogeneous B_0 field for depth encoding. The requirement that it be "mobile" specifically meant that it could be lifted by people and driven to a site in a vehicle. The manget design used a yoked electromagnet with pole piece extensions [100]. Other systems also used yoked electromagnet designs, and produced B_0 fields between 50mT and 80mT. In these systems, the flat, single-sided shape arose from the geometry of the problems of interest. For example, fields and road beds are large and flat, requiring an NMR system with a sensitive ROI offset from a flat surface: i.e. the archetypal "single-sided NMR" scanner.

Additional single-sided NMR work used the fringe fields of conventional "enclosing"

magnet designs to allow for single-sided operation [78]. This emphasizes the point that "single-sided" describes not just a system, but the operational mode of the system + sample together.

The NMR-MOUSE and Similar Devices

The NMR-MOUSE (Mobile Universal Surface Explorer) [109] is a single-sided MR device conceptually similar to the earlier devices depicted in Figure 1-5. However, it represents a significant advance in the field for several reasons:

- The NMR-MOUSE probe (including B_0 magnet and RF coil) are compact and lightweight (weight = 2.5kg; diameter = 5.5cm). This allows the NMR-MOUSE to be easily positioned by hand, carried to remote sites, and handled in a laboratory. This distinguishes it from earlier, car- or tractor-toted systems.
- It was successfully commercialized [110], facilitating its use in a wide variety of applications. It was estimated that 40-50 NMR-MOUSE devices were in use as of 2008 [111], and the NMR-MOUSE continues to see widespread use.
- The NMR-MOUSE was the first device with its form factor. Since its initial publication, numerous single-sided NMR devices have been developed that build on and refine the original NMR-MOUSE design.
- The NMR-MOUSE also represents an early use of an NdFeB rare-earth material-based B_0 magnet. An NdFeB-based human MRI magnet had previously been designed by Philips [112], but I can find no indication that it was ever built. However, I am unable to verify whether or not the NMR-MOUSE was *the* first NMR system to use an NdFeB B_0 magnet.

Similar devices have include a system for assessing fruit ripeness [113]; systems for analytical applications with highly-homogeneous [114] and constant-gradient sweet spots [108, 115]; and a recent system for hydration assessment *in vivo* via T_2 relaxometry [116]. Applications of these devices are described in Section 1.6.4.

Single-Sided Elements of Conventional Systems

Several "single-sided" hardware components have been developed for use in conventional cylindrical-bore MRI scanners. These are not "single-sided" imaging devices, but nonetheless illustrate situations where a single-sided component is advantageous. They also represent potential opportunities for a complete single-sided imaging scanner. Note that these devices are often described as "local", rather than "single-sided".

First, the surface receive coil arrays for conventional clinical MRI scanners provide a precedent for single-sided MR devices [117, 118]. An NMR device requires, at a minimum, a magnetic field for polarization/precession, a means of spin excitation, and a way to detect Larmor precession. These surface coil arrays demonstrate that a single-sided receive coil can perform Larmor detection for *in vivo* 3D imaging; that their inhomogeneous B^1 fields can be managed; and that such a coil array can achieve greater SNR than a volume receive coil. In conventional scanner operation, each individual coil in the array is only sensitive to a small region of the imaging volume - a fact exploited via parallel imaging techniques to shorten image encoding time. SNR considerations generally preclude such tricks at low field, but the fact remains that a single-sided receive coil inherently contains spatial information. This has recently been exploited for spatial encoding in a nonbijective, low-field scanner [119].

Single-sided RF and field gradient coils have also been used to for spin-labeling brain angiography by tagging blood flowing in the carotids, which are located superficially in the neck. In one case, a long saturation pulse is played on a local RF coil placed over a carotid while the z-axis scanner gradient is played. Blood flowing up the artery (along the gradient axis) experiences adiabatic inversion and is thus spin-labeled [120]. This same experiment has been performed using a local gradient coil (as opposed to the scanner gradient coil), which allows for greater sequence flexibility [121].

"Surface-spoiling" single-sided devices have also been used to eliminate signal from unwanted anatomical regions. This was shown early on in a ^{31}P phantom with a surface sheet containing ferromagnetic material that produces locally-inhomogeneous fields [122]. The same concept has been demonstrated with a switched electromagnetic coil to eliminate signal from the back during cardiac imaging [123]. Both of these devices exploited the nature of single-sided devices to efficiently produce locally inhomogeneous fields, enabling spin

dephasing (spoiling) in the anatomy located close to the single-sided coil.

These examples demonstrate a couple of things. First: that people in MRI are aware of the uses of "single-sided" field-generating devices; thus, there is precedent to build on in designing a single-sided MRI system. Second, "field inhomogeneity" is usually a bad thing in MRI. However, it can either be managed (as with RF receive fields) or it can be harnessed (as with spoiling and local spin labeling/saturation). Thus, there may be other unconventional opportunities with as-yet unbuilt single-sided MR devices.

1.6.4 Applications of Single-Sided MR Systems

Single-sided NMR devices have been utilized in a variety of different settings. This section - while not a complete literature review - aims to convey the scope of this range of applications.

Non-biomedical Applications of Single-sided MR

A range of industrial and materials analysis applications has been explored with single-sided NMR [124]. This includes: the analysis of polyethylene pipe wall integrity [125]; the detection of UV-degraded rubber [126]; and the analysis of structural cement *in situ* [127]. It has been used for non-destructive testing of objects with cultural or historical significance, as reviewed in [128]. Applications have included: the *in situ* analysis of ancient Roman walls and wall paintings [129, 107]; laminar analysis of painting layers [130]; and the analysis of paper in historical codices [131, 132]. Agricultural and food science applications are reviewed in [133, 134], and include: analyzing the fat content of foods [135]; measuring kiwifruit ripeness [136]; detecting spoilage of tomatoes in sealed drums [137]; and detecting the adulteration of olive oil [138]. Single-sided NMR devices generally have large field gradients, enabling diffusion encoding. This has been applied to the study of sedimentary rock porosity [139]. Single-sided NMR is applicable in a wide range of settings, especially where nondestructive or *in situ* analysis is desired. Both properties are often beneficial in a clinical diagnostic tool, suggesting that single-sided MR may translate well to clinical scenarios.

Biomedical Applications of Single-sided MR

Single-sided MR has been used in an increasingly diverse range of settings in biology in medicine. These use cases range from the analysis of biological samples to the development

of systems for clinical uses.

The first *in vivo* single-sided MR measurements were published by Georges Béné in 1980 [46, 45, 47]. The apparatus consisted of a single solenoid coil which performed both prepolarization and RF Tx/Rx functions. Figure 1-6 illustrates the human experimentation setup. He performed *in vivo* measurements with the detector placed over the heart, stomach, and bladder, and reported multiple signal components with different T_2 values; the long T_2 values were interpreted as the fluid (blood, water or urine); short T_2 components were interpreted as arising in anatomical tissues. Also reported was a pre/post-micturation experiment where signal was measured over the bladder using a spin-echo train sequence. In this case, it was found that the long- T_2 component was reduced in the post-micturation case, suggesting that it indeed corresponded to urine [47]. This early research does not appear to have inspired much subsequent work, but it bears similarities to recent attempts to measure hydration with single-sided NMR via multiexponential T_2 relaxometry [70].



Fig. 12. Measurement of T_2 of the blood in cardiac region

Figure 1-6: The first *in vivo* single-sided NMR experiments used a single prepolarization + Rx coil with the Earth's field for Larmor precession [47].

Later, a 2013 review [124] found described two biomedical applications of single-sided NMR: the measurement of diffusion and/or relaxation anisotropies in tendon [140] and cartilage [141]; and the profiling of human skin [76], including for measuring the absorption

of moisturizer [142]. These studies used either an NMR-MOUSE-type system [109] or a "Surface GARField"-type system [108].

More-recent studies have used single sided NMR for skin profiling [143, 144] and burn depth assessment [145]. The NMR-MOUSE has also been used for preclinical (porcine) intraoperative measurements of intestinal diffusion [146, 72]. Relaxometry-based applications of single-sided MR have included: mammographic density assessment via T_1 [147] and T_2 [148] measurements; measurement of placenta T_2 relaxation *in vivo* [71]; and the measurement of hydration state in dialysis patients via T_2 relaxometry in the calf muscle [149, 150, 151, 116, 70]. This work reflects the growing interest in the potential of single-sided MR and its potential in clinical cases not normally served by MRI.

Imaging with Single-sided MR systems

Prior work has also looked at imaging with single-sided MR devices. This includes 1D profiling using the intrinsic B_0 field gradient for readout encoding [152], and 3D imaging using the built-in B_0 gradient in conjunction with single-sided phase-encoding gradient coils [153, 154]. However, the *in vivo* applications and performance of this work have not been evaluated, and no single-sided systems have been designed specifically for *in vivo* volumetric imaging.

1.7 Main Points and Overview of This Work

Brain imaging with MRI is diagnostically useful in a variety of clinical settings, but its applicability is presently limited by the cost and operational limitations of conventional scanners. Portable, low-cost, or low-field MRI scanners could open up MRI to an array of currently-underserved settings. Much work has explored possibility, and the field of portable, compact MRI is presently growing rapidly. The field of "single-sided MR" has also produced portable NMR devices for a wide variety of applications and usage settings. In this thesis, I build on this work to develop techniques for the design and fabrication of portable MRI scanners. I then apply this approach to three systems for brain imaging that span the space of possible MRI scanner designs. Those three systems are: a short-bore cylindrical system for whole-brain imaging; a "single-sided" device for imaging a reduced volume within the brain; and a hybrid cylinder+sphere geometry system for brain imaging.

Chapter 2

System-level Design Considerations

2.1 Abstract

This section deals with the high-level design of a portable MRI scanner. Basically, how does one go from an idea for an MRI system down to concrete parameters and performance goals for the essential system components (B_0 magnet, gradient coils, and RF coil)? In this chapter, I first discuss the design goals for portable MRI scanners and the design tradeoffs that these entail. I next describe my general design approach: First, I decide on a specific anatomical region of interest (ROI) and system use case, along with operational constraints for the system. Next, I create the basic geometries of the system and design ROI: these system-level geometric designs dictate the geometry of the lower-level system components. I conclude by describing some of the design choices I made for the systems in this work and their implications for system performance.

2.2 System Design Objectives

The goal of this work is to design portable MRI systems that can form images of a specified region of human anatomy and potentially be used in unconventional settings. Specifically, I am concerned with neuroimaging and systems that can image all or part of the brain. These goals entail both desired performance specifications and constraints on system design.

2.2.1 Imaging Performance Objectives

The most important requirement of any MRI system is that it produce images of diagnostic quality in a reasonable acquisition time. This means that we want the highest SNR/CNR possible. The definition of a "reasonable" acquisition time depends on context. In this work, I generally use 5-10 minutes as a "reasonable" acquisition length. Images should be "good enough" for the targeted clinical application, and also be free from artifacts and distortion to the extent that they affect their diagnostic quality.

2.2.2 System Operation Requirements

Additional system requirements come into play when designing portable MRI scanners. The system should not require any special electrical power infrastructure, and should run either on batteries or a standard wall outlet. Portable magnet designs also generally eschew cryogenics and active liquid cooling. The system should not require a shielded room for EMI shielding, and should be light and compact enough to be maneuvered down hallways and through doors.

2.2.3 Tradeoffs in Portable MRI Design

The requirements for high SNR and CNR mean we want as strong a B_0 field as possible. However, the impracticality of using cryogenics precludes the use of superconducting magnets, requiring the use of a permanent magnet or resistive magnet design. This, in turn, greatly limits the maximum allowable B_0 to typically below $0.2T$ for human head- or body-sized scanner. Safety considerations and scanner footprint size also favor a lower-field magnet. The system mobility requirement limits the maximum size of the magnet. In some cases, we may want a head-only design (ie a design that doesn't reach below the patient's shoulders), further constraining the magnet's size. The combination of resistive and/or permanent magnet design and limited system size results in a crucial homogeneity-vs-field strength trade off that must be carefully managed [79].

The need for distortion-free images impacts the design of the spatial encoding system and image reconstruction scheme, in that the encoding and reconstruction schemes should together result in well-conditioned inverse problem. One is in principle able to choose a non-isotropic or non-cartesian reconstruction image resolution, but we will assume the use of

isotropic resolution in this work. For the textbook case of k-space encoding with perfectly linear gradient fields, this reconstruction problem is simply the inverse DFT, which has unity condition number. Encoding problems with non-linear encoding fields will often have higher condition number, but can often still be solved with minor approximations [155]. In general, it is fine to have non-linear encoding fields as long as they are suitably monotonic along some spatial axis; the "suitability" can be assessed by reconstructing simulated data acquired in the encoding field of interest with added noise [119]. For readout encoding, the encoding field is either the readout gradient coil field combined with the B_0 field (for switched gradient readout encoding) or just the B_0 field (for built-in B_0 encoding). Thus, image encoding concerns play a major role in B_0 magnet design.

The limited electrical supply available from a typical wall outlet and thermal management considerations constrain the design of resistive magnets, gradient coils, and RF coils. In this work, I addressed this by using low-power gradient coils, permanent B_0 magnets, and close-fitting RF coils. I typically aimed to keep peak RF power during transit to below $1kW$ and gradient coils to below $R_{dc} = 2\Omega$ with peak current below $10A$. The operation of the switched RF and gradient coils in turn, depends on the homogeneity of the B_0 field. A less homogeneous field requires a higher RF Tx bandwidth to produce the same excitation; thus, reduced B_0 homogeneity adversely affects RF Tx performance. Similarly, in the case of a switched readout gradient, a less-homogeneous B_0 field requires a stronger readout gradient to overcome the inherent field inhomogeneities. Thus, worse magnet homogeneity results in both greater gradient power requirements and a reduction in voxel SNR.

Finally, all of these tradeoffs have implications for the pulse sequence and encoding used. First, SNR considerations practically require the use of RARE-type sequences given the gross B_0 inhomogeneities in these systems (>10000 ppm). Additionally, a given peak RF power limit and a desired pulse bandwidth (in turn influenced by the spatial B_0 variation) specify a hard constraint on minimum pulse length. Magnet homogeneity determines linewidth and T_2^* , and thus SNR available during a readout window. Gradient peak current, coil efficiency, and desired image resolution determine the shortest allowable gradient blip duration. Maximizing SNR requires packing as many data acquisition windows as possible within one repetition time. However, the available degree of temporal compression of the sequence is limited by gradient and RF considerations, which are ultimately influenced by field strength, available RF and gradient power, and - perhaps most significantly - B_0 field

homogeneity.

2.3 High-Level System Design

This section describes the process of going from an idea concept to concrete design problems for system components. Note that even though this section is ostensibly about "high-level design", that does not mean we can ignore the details of the individual sub-systems. An understanding of how the magnets and coils will be built is essential for developing a plausible system-level design. This means in practice that it is valuable to have experience building magnets and coils and a realistic idea of performance/design tradeoffs involved. Furthermore, it means that there is a lot of iteration involved in this process. Even though this section is written going from "high-level-design" on to "component-level-design", the designer should not unwaveringly follow this trajectory. One should be willing to prematurely dip into the design of a system component if it's unclear how that component impacts the high-level design. Likewise, once a big-picture design has been decided on, it can and often should be changed in light of insights gleaned from more detailed design later on. And lastly, remember: it is much easier to make design changes by tweaking a parameter in CAD than by using a Dremel tool in our physical world.

2.3.1 Conceptualizing the System

We first need to have an idea of how the system will interact with the patient: how will the patient be positioned in the device (or the device positioned on the patient)? How much space should there be between the patient and the most proximal part of the device? What subset of patients will this be designed for? (What head size percentile? What about bed-ridden patients? Patients with external monitors? etc.) These considerations dictate some of the most fundamental aspects of system design: spatial dimensions, access holes/openings, key mechanical considerations (eg does the system need a hinge to open up?) and so forth. This step need not be quantitative at all - figure out what the scanner should be like, then figure out the quantitative dimensions. Furthermore, we should think about the ultimate imaging performance of the system at this stage (given the desired clinical applications) and develop reasonable metrics for the component-level design process. For instance, if we want a system that can be positioned on the head of a bed-ridden patient,

it may need to be a single-sided device as opposed to a traditional "patient-in-bore" MRI scanner design (for example, see Ch 6). A single-sided device will have a much less favorable B_0 strength/homogeneity tradeoff compared with an "enveloping" design of comparable size, and we must adjust our expectations of magnet and imaging performance accordingly (always keeping in mind the ultimate clinical application). There will always be room to adjust design dimensions and parameters to play with performance tradeoffs, but some things will already be baked in even at this early stage.

2.3.2 Selecting System and ROI Geometry

The next step involves going from our concept to an initial design with numerical dimensions (this is a key step - adjusting and tweaking dimensions is comparatively easy once you have a starting design to play with!) At this stage, we want to create a volumetric model of the system with dimensions, along with a corresponding coordinate system. First, we should pick a design ROI based on our anatomical imaging region of interest. Two examples in this work were the MR Cap, which used a $\approx 8\text{cm}$ lenticular ROI; and the Head-Optimized MRI, with used a brain-shaped ROI. In both cases, a high-field anatomical MR image was used to help select the ROI (Fig 2-1).

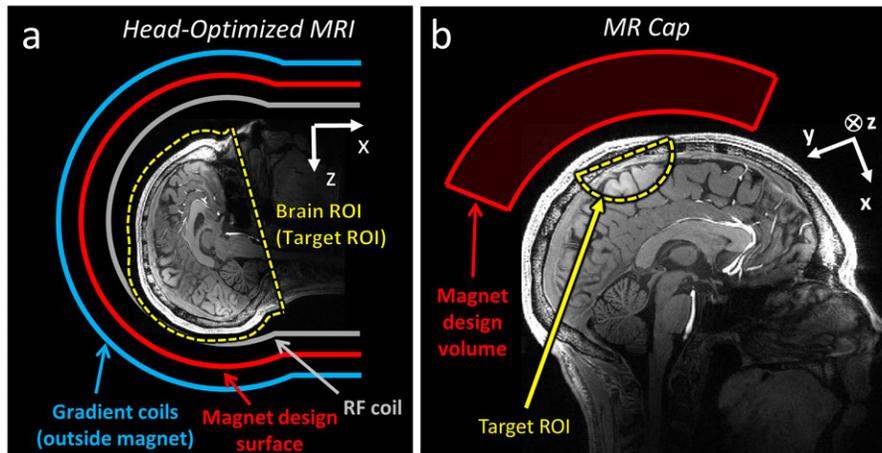


Figure 2-1: a. ROI definition for the Head-Optimized MRI scanner. The ROI was chosen to include all of the brain anatomy as identified on a high-field structural scan. b. ROI definition for the MR Cap. This ROI was $\approx 8\text{cm}$ in size and lenticular in shape.

The design ROI need not precisely correspond to the anatomical region of interest, and there may be more than one ROI for different aspect of system design. For example, the design of the Head-Optimized RF coil (Sec 7.6) used two target ROIs: one where I wanted

uniform, nonzero B_1 coverage; and one where I wanted zero B_1 field to prevent signal aliasing. The design ROI is the region in which we will evaluate system performance during design and testing, be that B_0 field homogeneity, gradient field linearity, or some other performance metric.

I next determine quantitative values for the geometric parameters of the system. At this stage, I personally find it easiest to work in a CAD program - this allows you to see an accurate rendering of your quantitative model in real time; allows you to faithfully see un-physical model intersections; lets you see how a human body model fits; and lets you arrange the system together with the ROI. It also helps to use a high-field MRI scan (if available) to precisely measure the relevant anatomical distances. This can help quantify the dimensions of your system design (Fig 2-2). It is fine to just make up dimensions as needed: they can always be edited later.

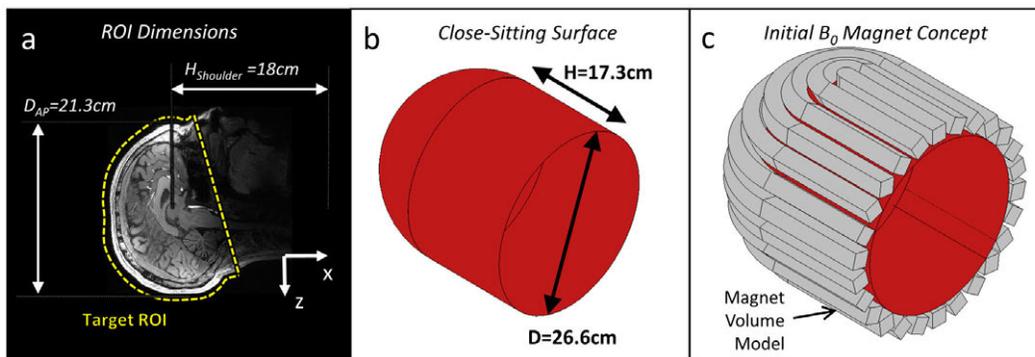


Figure 2-2: a. The Head-Optimized MRI target ROI dimensions were measured from a high-field MRI scan; the max A-P size was 21.3cm. The distance from isocenter to the shoulders was measured on an adult person to be 18cm. This specified the maximum extent of the system geometry below isocenter. b. These dimensions were used to create a close-fitting surface as the first step in specifying the system geometry. The surface diameter is greater than the ROI A-P distance; its length is less than the head-shoulder distance. c. A preliminary magnet design was created around the close-fitting surface. At this point, there is a design to play around with and optimize.

2.3.3 The "Bore Hierarchy"

We now have a volumetric model of the system geometry, and we can think about how the B_0 magnet, gradient coils, shims, RF coils, and other components fill that volume. They will usually occupy concentric shells in the bore in a sequence I call the "Bore Hierarchy". There exists a performance tradeoff in determining where something should go in the bore. In general, putting a component closer to the patient at a smaller radius provides greater field efficiency and less field uniformity; being away from the patient (at a larger radius) results in less field

efficiency and greater field uniformity. Other component-specific considerations also come into play. These constraints are different and generally more restrictive for a conventional superconducting high-field MRI scanner as compared with a low-field, permanent magnet-based system.

Efficiency/Uniformity Tradeoff

Greater efficiency is attained by being at an inner shell near the ROI: RF Tx power efficiency ($\frac{\mu T}{\sqrt{W}}$), B_0 strength (for a given size or mass magnet), and gradient efficiency ($\frac{mT}{m \cdot A}$) will all be higher than if positioned closer to the patient. However, increased spatial uniformity is attained when further out in the bore: the $B1+$, $B1-$, and B_0 fields will all be more uniform, and gradient coil fields will be more linear. Though the same tradeoff exists for all these field-generating components, the different components operate at different points on that tradeoff curve. This helps determine a "bore hierarchy" in a systematic way.

B_0 Magnets and B_0 Shimming

This tradeoff is illustrated with the relative placement of the B_0 magnet and the fixed shim magnets in a permanent magnet system. We want the B_0 magnet to produce a field of low spatial order (either a uniform field or a linearly-varying field). Thus, the B_0 magnet should be placed further from the patient based on field uniformity considerations. The shims, however, should accurately cancel out the residual non-uniformities or non-linearities of the B_0 field; these residual field terms are of higher spatial order. We therefore actually want the shim magnets to produce non-uniform, high-order fields, and should naturally place the shims closer to the design ROI. From this one can argue that shim magnets should be placed inside of the main B_0 magnets.

However, we also want a strong B_0 field, which suggests that we should place the B_0 magnet at an inner bore shell. This strength/homogeneity tradeoff must be determined on a system-by-system basis, and is bounded by the available strength (ie remanent flux density B_r) of available permanent magnet materials.

Component-specific Considerations

RF Tx and Rx Coils

Operating an RF coil in Rx mode has different considerations from when in Tx mode. During Rx, the primary concern is not directly field efficiency, but SNR. A smaller coil fitting closer to the signal-generating patient anatomy provides greater SNR than a larger coil. This provides a compelling reason to keep the RF Rx coil at the innermost part of the bore fitting closest to the patient.

Another consideration is the effect of an RF shield placed outside the RF coil in the bore. This structure may filter unwanted environmental and gradient-coil-fed interference, and also helps keep coil loading (and thus tuning and matching) consistent. However, a grounded RF shield also reduces Tx power efficiency; the closer the shield is to the coil, the greater the effect. In our experience, we have found it helps to have $\approx 3\text{cm}$ between the RF coil and the RF shield. This space requirement in turn influences the geometric system design.

Gradient Coils

The impacts of gradient coil switching and Ohmic heating also come into play. These are of great concern for high-field scanners and their rapid imaging sequences (see below), and also have effects in permanent magnet-based scanners. Eddy currents induced in permanent magnet material and thermal contact between the warm gradient coil windings and the magnets will both tend to warm up the magnet. This will tend to decrease the B_0 field and possibly change the spatial B_0 field map. The importance of this effect depends greatly on the acquisition, encoding, and reconstruction scheme used, and it should be considered during system design.

In this work, I created gradient coils with adequate field efficiency on a 30cm to 40cm-diameter geometry. In those systems, there was no compelling field efficiency argument to put the gradient coils any closer to that patient than 30cm to 40cm. Gradient coils could thus be placed at a relatively large bore diameter - in some cases, even outside the B_0 magnet.

Comparison with Traditional Scanners

The "bore hierarchy" for a commercial superconducting MRI scanner is shown in Fig 2-3. The order goes (from outermost to innermost): B_0 magnet, shim coils, permanent magnet B_0 shims, gradient coils, RF Tx coil, and RF Rx coil. Not shown are various RF and eddy current shields. There are several similarities with the observations described above: The RF coils are placed close to the patient; and the shim magnets are placed inside the main B_0 magnet.

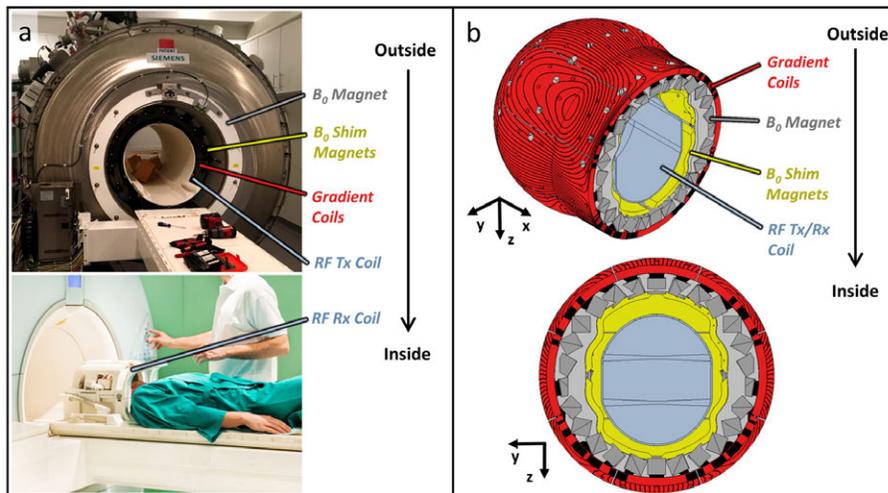


Figure 2-3: a. The "Bore Hierarchy" of a conventional superconducting magnet places the superconducting B_0 magnet on the outside, followed by the B_0 shim magnets, gradient coils, RF Tx, and lastly RF Rx coil. b. The bore ordering for the Head-Optimized magnet placed the gradient coils outside the B_0 magnet. The shims and RF coil were placed inside the B_0 magnet.

That said, the bore hierarchy for a normal scanner is dominated by concerns about cryogenic magnet operation. The necessity of putting the B_0 magnet in a cryostat all but requires it to be on the outside of the magnet. The imperative of keeping it cold requires a substantial cryostat between the gradient coils and the superconducting windings. Minimizing eddy currents in this cryoshield precludes putting the gradient coils outside the magnet. Furthermore, these gradient coils are always self-shielded: the contain a counter-wound set of windings that cancels the field outside the coil at the cost of reduced field efficiency inside. Similar concerns about shim iron heating and stability motivate the decision to keep the shim material outside of the gradient coils or between the primary/shielding windings. (If the shims were placed inside the gradient coils, the eddy currents would adversely impact imaging performance, and the heating of the iron would have deleterious effects on field stability and homogeneity.) The lack of many of these considerations in

designing a resistive- or permanent magnet-based system gives the designer greater flexibility in ordering the different system components in the bore.

Portable System Example: The Head-Optimized MRI

Fig 2-3b shows the bore ordering for the Head-Optimized MRI system. This system had several similarities with a conventional scanner - the B_0 shim magnets sit inside the main B_0 magnet and the RF coil is at the innermost shell. However, the gradient coils were placed outside the B_0 magnet and B_0 shim magnets; this was only possible because a superconducting magnet design was not used. The NdFeB magnet material can still heat up via eddy currents or thermal conduction from the ohmically-heated gradient coils: these concerns are less critical than in superconducting systems and are looked at in Section 7.8.2.

2.3.4 Conceptualizing the System Components

At this stage, one will have developed an initial design specified by a set of geometric parameters (eg overall radius, overall length, ROI size), along with an idea for ordering the different parts of the bore. It is now useful to simulate a non-optimized design for certain system components and think about how the MRI system will be manufactured.

Preliminary Field Simulations

I find it illustrative to run simulations on standard, non-optimized designs with the initial dimensions I've selected. For example, it may be worth simulating a standard dipole magnet, Halbach magnet, or Golay-type gradient coil with the working system dimensions to get a handle on expected B_0 magnet or gradient performance. Even though I'll be optimizing all these components, some aspects of performance are determined more by raw size and scale than by how clever the design is. Mean B_0 strength, for instance, will not change much as you optimize your design.

Thus, if you simulate a boilerplate magnet design at this stage and find a mean B_0 value that's significantly lower than needed, then you may need to iterate on the overall system design and dimensions. Running these simulations for different initial dimensions will also give you a quantitative idea about the performance tradeoffs of your system. Additionally, simulating a "vanilla" design gives you a baseline with which to compare your optimized magnets and coils later on in the design process.

Manufacturability Considerations

It is important to consider how the system and its constituent parts will actually be built before delving deep into the design of magnets and coils. Subsequent chapters describe in detail the techniques I use for permanent magnet (Ch 3) and coil (Ch 4) design. However, some universal notions of designing for manufacturability still hold. First, there obviously can't be self-intersections between magnets or copper wire windings, no matter how easy it is for CAD programs to handle them. Also, when mating pieces together – whether it's two machined ABS pieces or bonding a magnet in a magnet former – there must always be some mechanical tolerance for "wobble-room". This tolerance propagates back into the design as an effective mismatch or error source between the ideal design and the practical result to be expected. Finally, "manufacturability" describes not just the final constructed system, but also every intermediate step along the way. For instance, one must have an idea not just about how an array of permanent magnets will be held together in the ultimate prototype, but how the magnets will, one at a time, be put together.

2.4 Implications for Component-Level Design

This section describes several paradigms and rules-of-thumb used in this work for the design of portable MRI system components. These decisions are ultimately motivated by the requirements of designing a compact, portable MRI scanner.

2.4.1 B_0 Magnet Design Principles

As previously discussed, portability precludes the use of a superconducting magnet. In this work, I use exclusively permanent magnet-based B_0 magnets. This work aims for truncated head-only or reduced FoV magnet designs, and for a main B_0 field greater than 60 mT. The magnets are built of either stock NdFeB magnet sizes (eg cubes) or easily-manufactured custom magnet sizes (eg a rectangular block with only one non-stock dimension). As a result of these choices, these magnets are highly inhomogeneous (1-6% homogeneity peak-to-peak). This B_0 inhomogeneity impacts gradient, RF, and sequence design.

2.4.2 Gradient and Encoding System Design

Gradient coils are designed to be as linear as possible. It is usually possible to design a gradient coil that produces a bijective - but not perfectly linear - encoding field. A non-linear gradient coil will result in imaging distortions that can in principle be fixed by using a model-based reconstruction. The gradient coils do not need to be self-shielded and can be placed outside the magnet. The former design choice improves coil field efficiency; the latter opens up new options for system design.

Built-in B_0 encoding is possible in addition to switched gradient coils. In some cases (eg with a single-sided magnet design) it is difficult to avoid a having a built-in B_0 inhomogeneity. In such cases, one may optimize the shape of this spatially-encoding inhomogeneity during the B_0 magnet design process.

Blipped, balanced phase encoding can be used for spatial encoding during RARE-type sequences in inhomogeneous fields. The spin-echo train refocuses all field inhomogeneities except for that incurred by the phase-encode blips. Blipped phase encoding often produces less distortion than readout or slice-selection encoding. The former depends only on the gradient coil field pattern, while the latter two both depend on the inhomogeneous B_0 field.

The properties of low-field scanners contain several synergies. Notably, the reduced SNR requires running lower-resolution scans. This in turn reduces the extent of k-space that must be sampled, in turn reducing the maximum required gradient coil current. Thus, low-power, compact gradient coils are a natural fit for low-field MRI systems. Low-field and compact scanners also have smaller fringe fields, making them inherently safer, easier to site, and more portable than high-field body systems. The lower Larmor frequency leads to significantly lower SAR levels, which is beneficial when using a spin-echo train-based sequence in an inhomogeneous magnetic field. One should be aware of these overlaps and exploit them during system design.

2.4.3 RF System Design

Single-channel Tx/Rx coils are used in this work, though encoding benefits can be achieved by using multiple coils [119]. Coils in this work needed to resonate at $f_{RF} \in [2.5, 3.5MHz]$. Covering the entire RF pulse or signal bandwidth requires a bandwidth between $50kHz$ and $200kHz$; coils were therefore designed with quality factor Q between 10 and 30. Achieving

such a low Q was sometimes done with a series 1Ω resistor in the coil, though improved methods are theoretically possible [156].

For a fixed pulse length and flip angle, covering a wider bandwidth requires greater RF power. In this work pulses used up to 500 W RF peak power (for 180° refocusing pulses). This in turn requires a substantial RFPA. Keeping the Tx coil close to the target ROI also helps manage peak power requirements.

2.4.4 Sequences and Encoding Schemes

RARE-type sequences are exclusively used in this work due to the highly inhomogeneous B_0 fields of the systems developed. Field inhomogeneity results in short T_2^* that precludes the use of gradient-echo sequences. Using a multi-echo train improves acquisition efficiency (tSNR). Many common MRI contrasts can still be acquired with RARE-type sequences, including: T_2 , T_1 , FLAIR- T_2 , diffusion, and proton density. Susceptibility-weighted contrasts are generally not practical without gradient-echo sequences, and are physically much weaker at low B_0 fields.

Significant pulse sequence time is taken up by high-BW frequency-swept ("chirped") pulses - typically 1-3ms per pulse. Phase encoding gradients are typically 1ms long, and are played in a balance fashion about each acquisition window. Finally, the use of swept pulses means that on even echoes, different frequency components of the signal are refocused at different times (a "spectral echo"). Acquiring this signal thus requires an acquisition window as long as the RF pulse (1-3 ms). All of these considerations reduce the number of available echoes per shot, thus reducing available SNR for a fixed imaging time. One often needs to average to obtain acceptable SNR.

Chapter 3

Design and Construction of Permanent Magnet Arrays

3.1 Abstract

"Hard" permanent magnet materials are a mass- and volume-efficient source of magnetic fields, and are thus suitable for use in portable MRI systems. Magnet designs based on arrays of blocks of permanent magnet material enable designs with large numbers of degrees of freedom. Here I describe my approach to design and construct permanent magnet arrays for MRI main field magnets and for static magnet shimming. I discuss the formulation of the design optimization problem, the conversion of the resulting numerically-optimized design into a geometric CAD model, and the construction of physical prototypes.

3.2 Advantages and Drawbacks of Permanent Magnets for MRI

In this work, I use rare-earth NdFeB permanent magnet material to build B_0 magnets and for static B_0 shimming. Permanent magnet materials provide several advantages for these system components over resistive coil designs. Namely: permanent magnets require no quiescent power source; they are more a more efficient field source (more dipole moment per volume) than resistive magnets over length scales under 1m; and, relatedly, they can more efficiently create high-spatial-order field patterns. That said, permanent magnet designs also

come with drawbacks that must be managed: the properties of many materials drift with temperature; blocks of magnetic material have significant inter- and intra- batch variation (compared with the ability to precisely control current in a resistive magnet); and the handling and assembly of permanent magnets can be challenging.

3.2.1 Advantages

A typical N52 NdFeB permanent magnet with $B_r = 1.42\text{T}$ is a more space-efficient source of magnet dipole moment than a resistive current loop operating at $J = 5 \frac{\text{A}}{\text{mm}^2}$ at small spatial scales. This ultimately derives from the fact that a moving charge produces a field proportional to $\frac{1}{r^2}$ while a magnetic moment (such as an electron spin) produces a field proportional to $\frac{1}{r^3}$ where r is the distance from the particle. This leads to the familiar expressions for the magnetic moment of a current loop and a sample of magnetized material:

$$\text{Current loop: } \mu_{loop} = I \cdot A \quad \text{Magnetized volume: } \mu_{vol} = M \cdot V$$

Where I and A are the loop current and area, and M and V are the magnetization (ie magnetic dipole moment per volume) and the sample volume. These differing dependencies on magnet geometry (one depends on area, the other on volume) are of consequence for magnet engineering. This is illustrated with the example in Fig 3-1. Shown therein are 3 dipole magnet designs with their spatial dimensions: (1) a magnet made of a disc of NdFeB ($B_r = 1.42\text{T}$, $M = 1.13 \cdot 10^6 \text{Am}^2$); (2) a current loop made of two copper rings operating at $J = 5 \frac{\text{A}}{\text{mm}^2}$; (3) a "current disc" made of concentric current-carrying rings, each one carrying azimuthal current density $J = 5 \frac{\text{A}}{\text{mm}^2}$.

Fig 3-1b lists the mass, volume, and magnetic dipole moment for these three magnet designs. Magnetic dipole moment per volume is constant for a permanent magnet, but equals $J \cdot \frac{R}{3}$ for a solid current disc electromagnet. These expressions are equal when $R = 3 \cdot M/J = 0.68\text{m}$. Thus, at length scales smaller than this, permanent magnets are a more efficient source of magnetic dipole moment.

The suitability of using permanent magnet material for permanent B_0 shimming (improving the B_0 spatial field pattern) derives from this observation. Shimming typically requires producing magnetic fields with high-spatial-order components, which are more efficiently produced by small magnetic sources (in this work, shim magnets are 3-13 mm in size, 1-2 orders of magnitude smaller than the linear dimension of the main magnet). Permanent

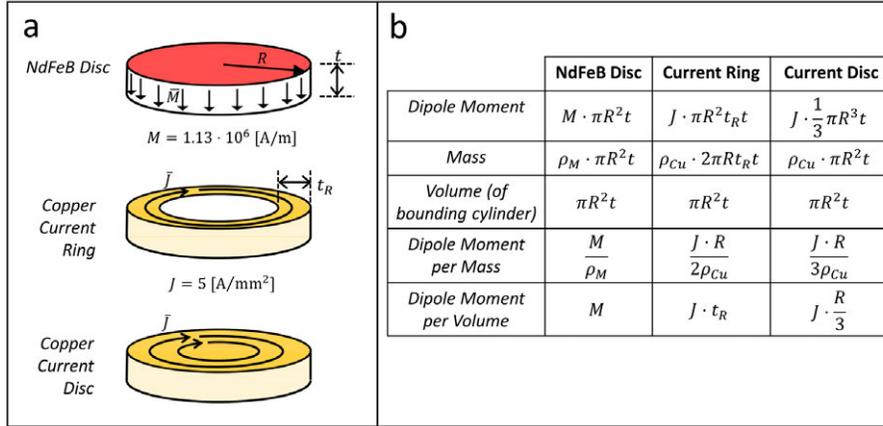


Figure 3-1: a. Three magnet designs with the same thickness and radius are compared - a disc with volume magnetization M ; a copper ring of given radial thickness t_R carrying current density J ; and a copper disc carrying tangential current density J . b. Computed physical properties for each of the three designs.

magnet material is a good choice for such operations given this spatial scale.

3.2.2 Drawbacks

Many permanent magnet materials have magnetic properties that vary with temperature, the most consequential being the remanent flux density B_r . The B_r of NdFeB magnets is typically around $-0.001 \frac{1}{^\circ C}$, meaning a $1^\circ C$ change in temperature leads to a 0.1% drop in B_r . For an NdFeB MRI magnet with a Larmor frequency of $3MHz$, this corresponds to a shift of $-3kHz$. There exist magnet materials with significantly lower or negative T_c (such as SmCo), but they tend to be more expensive and to have lower B_r . NMR magnets have also been developed using multiple materials with different T_c to create a low-field drift scanner [157]. In this work, temperature-induced field drift in NdFeB magnets is accounted for with phase navigation when needed.

Variances and tolerances in the manufacture of permanent-magnet structures can adversely affect MRI magnet performance. When using permanent magnet material, these production errors are literally cast in stone, and can effect temporal field drift (eg due to temperature) and spatial field homogeneity. In a resistive electromagnet, temporal field drift can be minimized with a field measurement - electromagnet current feedback system. A temperature-control system could in principle be used with an NdFeB permanent magnet, but would be limited by the thermal time constant of the system. Other approaches have been developed to manage field errors arising from NdFeB magnet manufacturing variations,

such as:

- Exploiting batch randomness when designing the magnet [158]
- Rejecting bad magnets that don't meet the desired specifications
- Additional magnet shimming steps

Handling and assembling permanent magnets can be both challenging and hazardous. In particular, the tendency of permanent magnets to exert forces and torques on other magnets can make it difficult to position a magnet precisely. There must always be some mechanical tolerance when designing an assembly to hold permanent magnets, and the internal torques and forces will tend to shift the magnets around within this "tolerance space". This can lead to errors between the designed and realized spatial positions of the magnet blocks, and must be contended with during magnet design and construction. Fig 3-2 shows an example of Sparse Halbach shim tray angular position error resulting from these internal torques.

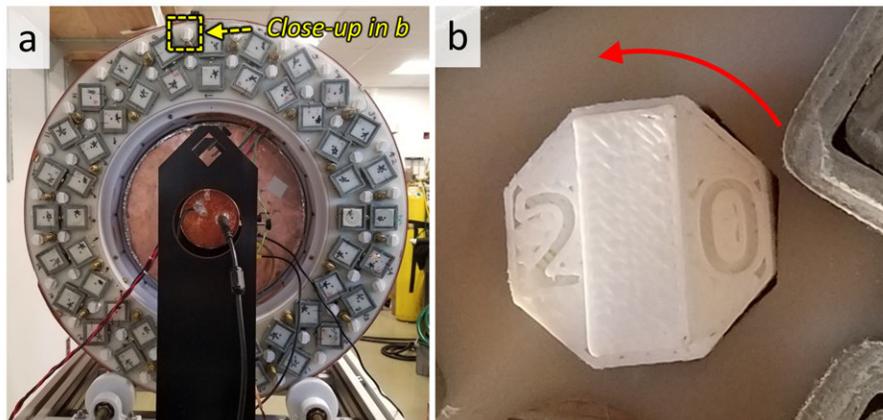


Figure 3-2: a. The optimized Sparse Halbach magnet contained 48 shim trays with permanent magnet material. b. The mechanical tolerance for shim tray insertion allowed the trays to rotate due to the torque exerted on them by the B_0 magnet blocks.

3.3 Permanent Magnets and Block Arrays for MR Systems

The potential for NdFeB materials in the design of MRI/NMR magnets was understood only shortly after its initial appearance in 1984 [159]. Early system designs using this material included a whole-body MRI concept from 1989 [112] and the NMR-MOUSE [109], and these materials have since seen widespread use. Book chapters on the design of permanent

magnets for compact and single-sided NMR systems have been published [160, 161]. One design (often used in this work) is the "Halbach" cylindrical magnet design [162]. This design produces a homogeneous field inside a cylindrical shell and zero field outside the shell, making it practical for compact MR systems.

One method to design and build a permanent magnet is to make it out of multiple smaller magnetic blocks, ie: to build a "block array". This has been explored for building "Halbach-type" NMR magnets [163, 164]. Using this approach, one can build magnets using a set of a small number of distinct block sizes, which simplifies and lowers the cost of the magnet. All magnets designed in this work were based on NdFeB block arrays.

3.4 Magnet Block Field Computation

Designing a magnet requires that we be able to accurately compute the magnetic field. The iterative optimization procedure used here further requires fast field computation. Here, I describe several ways to compute the spatial B-fields produced by magnet blocks: finite-element model (FEM) techniques; block precomputation; and the multipole model.

3.4.1 Need for Fast, Accurate Field Computation

In this work, we assess magnet design performance based on the field it produces in the target ROI. The performance metric could be: the mean field; encoding properties of field; or field homogeneity. It is not possible to compute the field analytically except for the simplest magnet designs, so we need a numerical way to compute the field. Field computation also needs to be fast for the iterative solvers used here. A typical iterative optimization solver works as follows:

1. We start with a design or set of designs at each iteration.
2. Several "perturbations" to that set of designs are created. For example, this could be in order to compute gradient of the cost function (as in gradient-descent), or it could involve quasi-random variation of the current design (as in genetic optimization).
3. The field is then computed and the cost function evaluated for each of these "perturbed" designs.

4. A new design set is created for the next iteration, based on the performance of the "perturbed" designs and the specifics of the optimization procedure.

Step (3) involves the field computation of a number of designs. Since that step is part of one iteration of the solver, that set of field computations must be repeated many times. In practice, run time often limits the size of the search space. Thus, it is important to quickly compute the field produced by a magnet design without sacrificing accuracy.

3.4.2 Dipole Model

The simplest way to compute the field produced by a permanent magnet block is to model the block as a point magnetic dipole. This is a superposition-based approach: the fields from all magnet blocks are computed separately and then added together. Like all superposition-based approaches, it does not model block-block field interactions. In this case, we take the value of a block's magnetic dipole moment to be $\boldsymbol{\mu}_{Blk} = \int_{Vol} dV \mathbf{M}(\mathbf{r})$. The magnetic field produced by such a dipole located at \mathbf{r}_0 is given by:

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \left(\frac{3(\mathbf{r} - \mathbf{r}_0)(\boldsymbol{\mu} \cdot (\mathbf{r} - \mathbf{r}_0))}{|\mathbf{r} - \mathbf{r}_0|^5} - \frac{\boldsymbol{\mu}}{|\mathbf{r} - \mathbf{r}_0|^3} \right) \quad (3.1)$$

This is a linear operation in $\boldsymbol{\mu}$ and can be written:

$$\mathbf{B}(\mathbf{r}) = D_{1,r,r_0} \boldsymbol{\mu} \quad (3.2)$$

where D_{1,r,r_0} is the 3x3 field computation matrix (D for dipole, 1 for 1-st order multipole model). If, in our optimization, the magnet is allowed to rotate but is translationally fixed, then D_{1,r,r_0} will not change. Thus, we only need to compute the dipole matrix for each magnet and each field computation point once. Field computation at each iteration of the optimization then only involves the following operation:

$$[B] = [D_1][\boldsymbol{\mu}] \quad (3.3)$$

B is a $3N_{ROI} \times 1$ vector containing all field components at all ROI points; $\boldsymbol{\mu}$ is a $3N_{Mag} \times 1$ vector containing all dipole moment components of all magnet blocks; D_1 is a $3N_{ROI} \times 3N_{mag}$ matrix containing the appropriate concatenation all all the 1st-order dipole matrices.

This approach results in a simple and fast field computation. However, it implicitly

models all magnet blocks as uniformly-magnetized spheres. It glosses over all internal details of the magnet blocks and may prove less accurate in certain situations, such as with rectangular blocks.

3.4.3 Finite-Element Model (FEM) Computation

FEM approaches first discretize all of space – including the magnet blocks – into a graph or mesh. Associated with this mesh is a set of weighting functions with which one can cast the original continuous problem as a discrete, computationally tractable problem. The solution to this problem will give the magnetic field at each node in the mesh. Many commercial FEM software packages exist, including COMSOL, Ansys Maxwell, and others.

The essential operation of these solvers requires that:

- The magnetic properties of the magnet blocks (BH curve, permeability, B_r , B_r orientation) are all expressed as boundary conditions on the problem.
- Adjusting the magnet array geometry (ie rotating a block, shifting a block, changing a block’s size) all require either recomputation of the mesh or adjusting the node coordinates of the mesh.
- The computation time required for these two overhead steps is typically longer than the actual time required to compute the field from the assembled problem.

This method is often too slow for use in iterative solvers given the computational overhead required. However, I use FEM field simulation to validate magnet designs in this work. I also assume that FEM solvers are the most accurate of all methods listed here. This is because each block is modeled as a distribution of magnetic material and block-to-block field interactions are accounted for.

3.4.4 Precomputing the Field

The fastest way to compute the field produced by a block is to have already computed it, and to just read it in from memory. One application of this approach is when performing superposition-based field computation with a finite number of possible states for each magnet block. This has been performed in the design of a genetically-optimized Halbach magnet, where the position, size, and orientation of each block was fixed, and each block could take

one of three possible material grades [165]. This approach has also been used in the design of a magnet consisting of a stack of concentric, stacked rings of differing radii. The fields produced by the rings of blocks (not the individual blocks) were precomputed for different ring diameters, and the diameters of all the rings were then optimized using those precomputed fields [166].

This approach has the benefit of being fast and generally accurate. Block-to-block (or ring-to-ring) interactions are still not modeled, but any method of field computation can be used (eg FEM) since it need not run many times during optimization. However, this technique becomes impractical when the design space is continuous (eg for continuous block translation, rotation, or size variation).

3.4.5 Multipole Models

A spatial distribution of magnetization can be modeled as a series of multipole field terms localized at a single point. An accessible introduction to this approach with applications to geomagnetism is given in [167]. This section is based on that treatment.

The Multipole Expansion

We can perform a multipole expansion of the field produced by a bounded region of magnetization. We start with the magnetic pole density describing this magnetization distribution:

$$\rho_m(\mathbf{r}') = -\nabla \cdot \mathbf{M}(\mathbf{r}') \quad (3.4)$$

Outside the magnetization region, we can define a magnetic scalar potential $\phi(\mathbf{r})$. This potential can be approximated by an N_ℓ -order multipole expansion:

$$\phi_{N_\ell}(\mathbf{r}) = \frac{1}{4\pi} \sum_{\ell=0}^{N_\ell} \sum_{m=-\ell}^{+\ell} A_{\ell,m} \frac{1}{r^{\ell+1}} S_{\ell,m}(\theta, \varphi) \quad (3.5)$$

where $S_{\ell,m}$ is the real spherical harmonic of degree ℓ and order m . The $A_{\ell,m}$ are the multipole terms of the magnetization region, and are given by an integral over the magnetization volume:

$$A_{\ell,m} = \int_{Vol} dV' \rho_m(\mathbf{r}') r'^{\ell} S_{\ell,m}(\theta', \varphi') \quad (3.6)$$

Computing the magnetic field requires taking the gradient of the scalar potential. From this, we find:

$$\mathbf{B}(\mathbf{r}) = -\nabla\phi(\mathbf{r}) \quad (3.7)$$

$$= -\frac{1}{4\pi} \sum_{\ell,m} \left[A_{\ell,m} \nabla \left(\frac{1}{r^{\ell+1}} S_{\ell,m}(\theta, \varphi) \right) \right] \quad (3.8)$$

This can be cast as a matrix multiplication of a field computation matrix with the vector of multipole terms:

$$[\mathbf{B}] = [D_{N_\ell}] [A_{\{\ell\},\{m\}}] \quad (3.9)$$

\mathbf{B} is the 3×1 vector of field components; $A_{\{\ell\},\{m\}}$ is the $(N_\ell + 1)^2 \times 1$ vector of multipole coefficients up to degree N_ℓ ; and D_{N_ℓ} is the $3 \times (N_\ell + 1)^2$ field computation matrix. Note that the *0th*-degree multipole term (a magnetic monopole) is nonphysical. For magnetization distributions with symmetry/antisymmetry, one also finds that the multipole terms of odd/even degree (respectively) are zero. Thus, one can often use fewer than all $(N_\ell + 1)^2$ terms and obtain a correspondingly smaller matrix multiplication. Note also that the dipole model described previously is precisely the $N_\ell = 1$ multipole model (ignoring the $\ell = 0$ term).

The matrices D_{N_ℓ} will not change for design problems where the magnet blocks do not undergo translation. Thus, they can be precomputed from the magnet/ROI geometry and don't need to be recalculated during the iterative optimization.

Computation of Multipole Terms

For simple magnetization geometries (such as a uniformly magnetized rectangular block), the integrals Eq 3.6 can be computed analytically. However, this is not possible when the geometry is complicated, or there are internal permeability/coercivity effects. To accurately compute the $A_{\ell,m}$, I first use COMSOL to compute ρ_m for a given region of magnetization and then numerically compute the integrals (3.6). I have not attempted to develop a quicker way to compute the multipole coefficients of arbitrary magnetization distributions; rather I have focused design-specific ways to speed up these calculations. Specifically, I have analyzed situations where the blocks are rectangular, are translationally fixed, and may undergo

rotations.

Multipole Terms of a Rotated Object

First, we consider applying a rotation matrix $R = R(\alpha_1, \alpha_2, \alpha_3)$ parametrized by three angles to a magnetization distribution with known $\{A_{\ell,m}\}$ and computing the rotated multipole terms of the transformed magnetization source $\{A_{\ell,m}^{Tx}\}$. We can compute the integral Eq 3.6 in coordinates that rotate along with the magnetization region. In this case, the integral is unchanged except for the $S_{\ell,m}$, as they are the spherical harmonics in unrotated coordinates:

$$A_{\ell,m}^{Tx} = \int_{Vol_R} dV'_R \cdot \rho_m(\mathbf{r}'_R) r'_R{}^\ell S_{\ell,m}(\theta_R, \varphi_R) \quad (3.10)$$

Note that $A_{\ell,m}^{Tx}$ describes the multipole terms of the rotated block in the unrotated multipole basis. Now, we exploit a useful property of spherical harmonics: under a coordinate rotation, the basis of spherical harmonics of degree ℓ transforms by a unitary transformation matrix parametrized by the 3 spatial rotation angles [168]:

$$[S_{\ell,\{m\}}^R] = [U_{\alpha_1, \alpha_2, \alpha_3}] [S_{\ell,\{m\}}] \quad (3.11)$$

The inverse transformation is likewise given by the matrix U^T . Substituting the inverse form of Eq 3.11 into Eq 3.10, we find:

$$A_{\ell,m}^{Tx} = \int_{Vol_R} dV'_R \cdot \rho_m(\mathbf{r}'_R) r'_R{}^\ell [U_{m,\{m'\}}^T] [S_{\ell,\{m'\}}^R(\theta_R, \varphi_R)] \quad (3.12)$$

$$= [U_{m,\{m'\}}^T] \left[\int_{Vol_R} dV'_R \cdot \rho_m(\mathbf{r}'_R) r'_R{}^\ell S_{\ell,\{m'\}}^R(\theta_R, \varphi_R) \right] \quad (3.13)$$

The expression in the right brackets is the same as Eq 3.6 (except written with rotated coordinates, a rotated magnetization distribution, and rotated spherical harmonics). Thus, it is equal to the appropriate multipole term of the unrotated block, ie $A_{\ell,m}$. In conclusion, rotation of magnetization source simply transforms the unrotated multipole terms of degree ℓ by a rotation matrix of size $2\ell + 1$:

$$[A_{\ell,m}^{Tx}] = [U_{-\alpha_3, -\alpha_2, -\alpha_1}^T] [A_{\ell,m}] \quad (3.14)$$

A technique for computing the matrix U for real spherical harmonics for arbitrary ℓ and rotation is found in [168]. As an illustrative case, the matrix $U_{\alpha_z, \ell=2}$ (a rotation about \hat{z} by α_z) is given by:

$$U_{\alpha_z, \ell=2} = \begin{bmatrix} \cos(2\alpha_z) & 0 & 0 & 0 & \sin(2\alpha_z) \\ 0 & \cos(\alpha_z) & 0 & \sin(\alpha_z) & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & -\sin(\alpha_z) & 0 & \cos(\alpha_z) & 0 \\ -\sin(2\alpha_z) & 0 & 0 & 0 & \cos(2\alpha_z) \end{bmatrix} \quad (3.15)$$

This approach allows for fast, in-optimization computation of the multipole terms of a rotated object. The terms themselves are precomputed; only the rotation matrices and rotation operations must be calculated at each iteration.

Multipole Terms of a Rectangular Block

Here I will consider the special case of a rectangular magnetized block of dimensions $A \times y \times A$, where A is fixed during magnet optimization and y is a variable/degree of freedom. For this example (taken from the design in Chapter 7), I will assume: $0 \leq y \leq A$, $A = 1'' = 25.4mm$, and $N_\ell = 5$. I assume also that the block is centered at the origin and is aligned with the coordinate axes (Fig 3-3a). From symmetry arguments, we find that there are only 6 nonzero multipole terms, given by: $(\ell, m) = \{(1, 1), (3, 0), (3, 2), (5, 0), (5, 2), (5, 4)\}$. These six multipole terms are then computed for several blocks with y -dimension varying between 0 and A (Fig 3-3). The multipole terms for a block of arbitrary y -dimension are then calculated by interpolating along these six curves. This can be done quickly at each optimization iteration.

This approach was only developed for rectangular blocks of magnetization that are parametrized by a single parameter. In principle, it could be possible to apply this approach to other single- or multi-parametric magnet block designs.

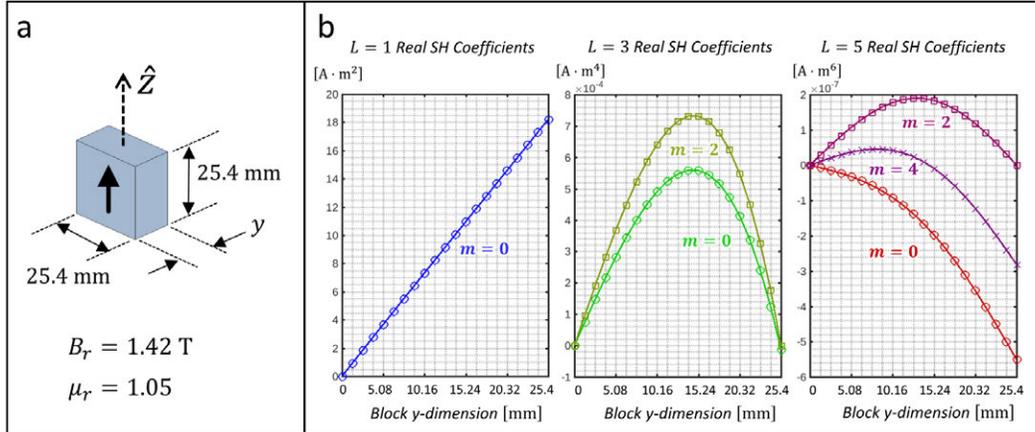


Figure 3-3: a. A magnetized block with size $25.4 \times y \times 25.4 \text{ mm}$ was aligned with the z-axis. b. This block has six nonzero multipole terms up to $L=5$; all are plotted against block dimension y .

3.4.6 Tricks to Speed Up Computation

Reducing field calculation time speeds up the optimization and allows for a larger design space or more accurate field computation. Various tricks can be used in many optimization problems to reduce the computational complexity of the magnetic field computation.

Exploiting Symmetry

We expect a symmetric field pattern over the ROI for a symmetric magnet and ROI geometry. In this case, computing any cost function only requires computing the field for part of the ROI, thereby reducing computation complexity. Furthermore, imposing symmetry on the magnet design reduces the number of degrees of freedom. This allows compressing the field computation problem to use only these independent DoFs, which reduces the necessary computation along with the size of the design search space. As an example, Fig 3-4 shows the MR Cap magnet design. This magnet was symmetric about the XY and XZ planes, which mean that the whole magnet design was fully specified by describing one quadrant of it (Fig 3-4d).

Using a Boundary as the Design ROI

We generally design our systems to provide adequate magnet performance inside a volumetric ROI. However, computing the cost function for that ROI may not require computing the field at every point in the volume. For example, many of the magnet optimization problems framed in this work use a "min-max" cost function to evaluate the B-field over the design

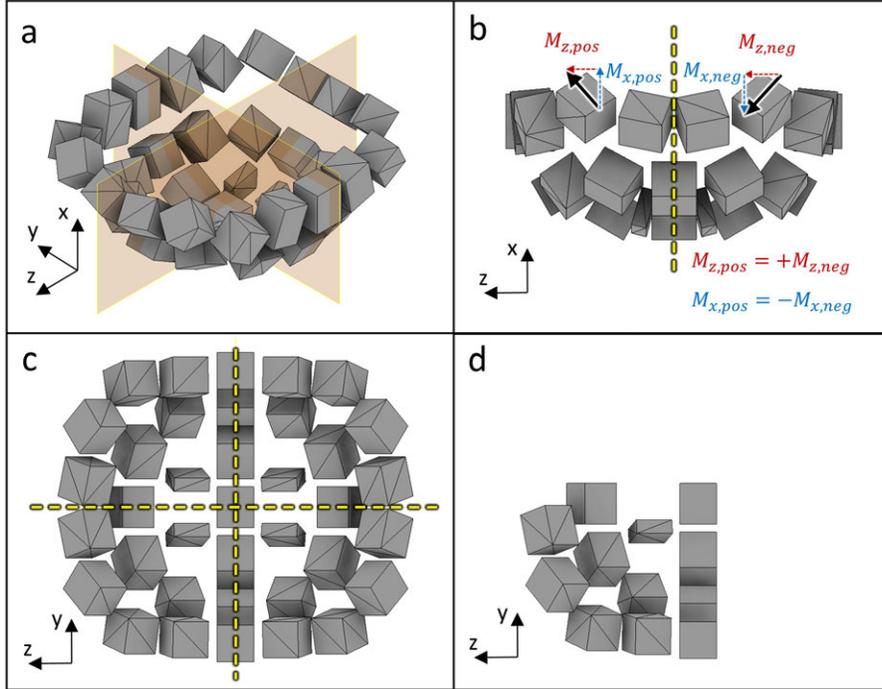


Figure 3-4: a. The MR Cap magnet was symmetric about the XY and XZ planes. b. Different degrees of freedom had positive and negative symmetry about different symmetry planes. c. The MR Cap design with symmetry planes shown in the YZ plane. d. Exploiting this symmetry allows one to specify the MR Cap design using only a subset of magnet blocks.

ROI. That is:

$$f_{cost} = \max_{\{ROI\}} (|\bar{B}|) - \min_{\{ROI\}} (|\bar{B}|)$$

It was found in this work that the extrema of $|\bar{B}|$ over the ROI were generally located on the ROI boundary. (Note that this is not always the case. For example, a "field-free point"-magnet has no field at the origin, but does have nonzero field everywhere on a spherical surface surrounding the origin.) Making this assumption (where justified) allows us to only compute the B-field on the boundary of the volume, greatly reducing the number of field evaluation points and speeding up computation. This approach was used for the helmet magnet design (Fig 3-5), which reduced the number of field computation points from 11617 (whole volume) to 3391 (boundary). Further exploiting magnet and ROI symmetry reduced the number of field computation points down to 1753.

Computing Only Necessary Components

In some optimization problems, one only needs to compute a single vector component of the B-field. This case arises when shimming an existing magnetic field under the assumption that

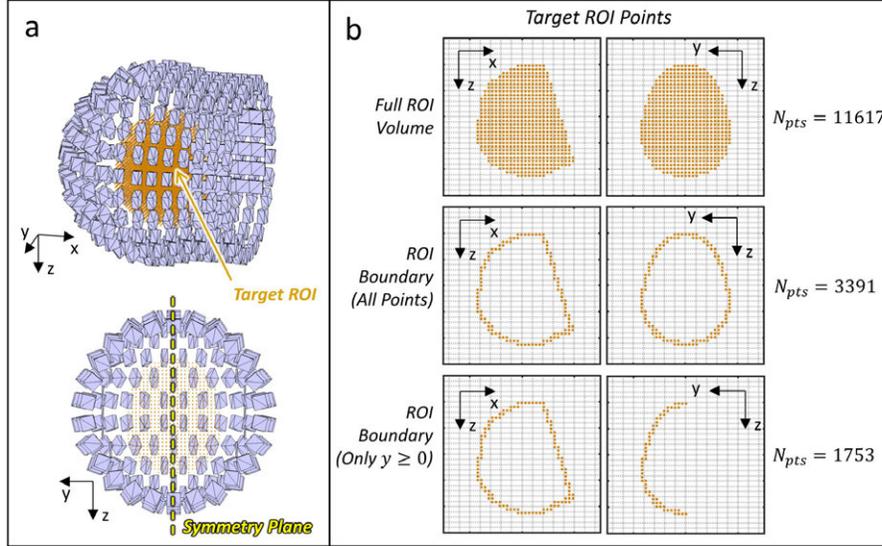


Figure 3-5: a. The Head-Optimized ROI used a volumetric head-shaped ROI and was symmetric about the XZ plane. b. Using only the boundary of the target ROI reduces the number of spatial locations at which the magnetic field must be calculated. Exploiting magnet and ROI symmetry reduces this number of points still further.

the shim magnets are too weak to perturb the direction of the existing B_0 field. Exploiting this assumption reduces the number of computations required by a factor of three - one only needs to compute the component of the shim field aligned with the known B_0 field direction.

3.5 Designing a Magnet Array

The goal of this step is to fully describe the magnetization distribution of our magnet. This means we must specify the shapes, sizes, material grades, positions and angular orientations of all blocks in our array. We must first consider the fabrication process and manufacturability concerns so that the final magnet be practical to construct. We next must cast the design problem in a form amenable to numerical optimization. This requires defining: a cost function (to assess design performance); a set of degrees of freedom/design variables and ranges of values they may take (to determine the space of possible designs); and constraints on allowable solutions. Finally, we assess the performance of the optimal design, with particular attention given to expected manufacturing tolerances and errors.

3.5.1 Manufacturability Considerations

Thinking about manufacturability is important at the outset of the design process. The way in which the magnet will be built determines the set of design parameters, the optimization constraints, and the magnet geometry.

Avoiding Block Intesections

Perhaps the most important practical consideration is that magnet blocks in a magnet array may not physically intersect. In this work, this is never explicitly enforced as an optimization constraint. Rather, it is implicitly "baked into the problem" so that the solution contains no intersecting blocks.

Generally speaking, the most direct way to do this is by leaving space between individual blocks in the geometric design of the magnet. This helps in design problems where blocks may move (whether rotationally or translationally) by giving them wiggle room free of other interlopers. However, we also want to pack the magnets as close together as possible in order to maximize performance. Thus, there exists a tradeoff between design simplicity and ease of manufacture, and magnet performance.

Another way to enforce this is by judicious choice of the DoFs. For example, the magnet array design problem framed by Cooley et al. [165] did not allow individual blocks to translate, rotate, or vary in size (ie the geometry was fixed). Thus, as long as the initial magnet geometry avoided self-intersections, the solution was guaranteed to as well. Another example is the permanent magnet shim array described in Ch 5. In this case, all magnet blocks were translationally fixed and only allowed to rotate along one axis (Fig 3-6). This geometric design also made self-intersections impossible.

One can also avoid intersections by limiting the range of the DoFs. In the "MR Cap" magnet described in Ch 6, magnets were allowed to translate linearly along some axis by up to $\pm 1cm$. This ensured that the solution had no self-intersections. If the solution had contained such intersections, it would have been possible to iterate by adjusting the constraints on those specific problematic DoFs and re-running to optimization until an acceptable solution was produced.

Finally, one can correct for self-intersections by playing with other geometric degrees of freedom besides those used by the optimization. An example of this can be seen in

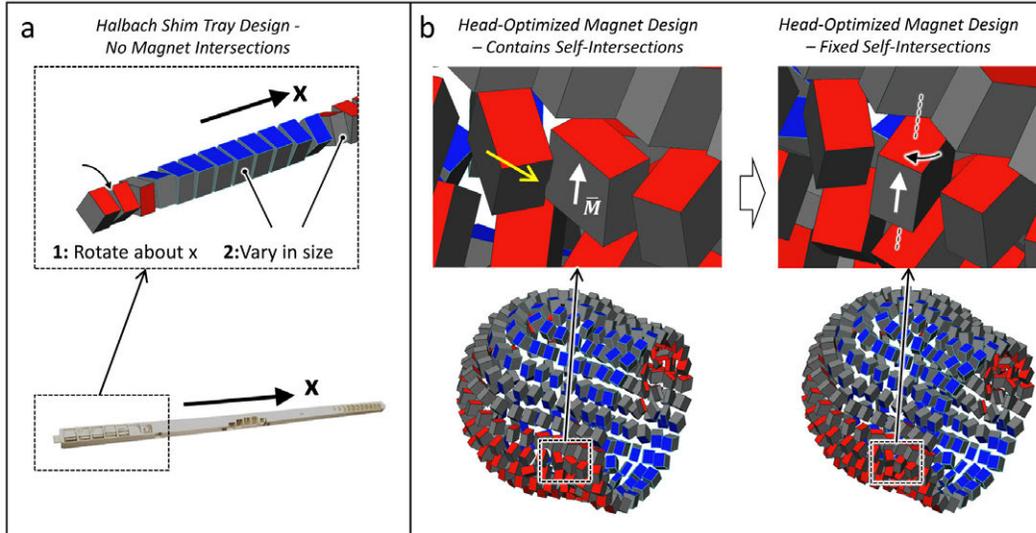


Figure 3-6: a. The Sparse Halbach shim trays were guaranteed to have no magnet intersections. The magnets could only rotate along one axis and could not translate, meaning intersections were impossible. b. The Head-Optimized magnet design allowed for the possibility of magnet intersections. These could be fixed by rotating the offending magnets about their own principle axes.

the head magnet designed in Ch 7. In this case, the optimized magnet design contained intersecting pairs of blocks as shown in Fig 3-6b. This was addressed post-optimization by applying a manually-specified rotations to the blocks about their own principal axes (Fig 3-6)b. This angular degree of freedom was not included in the numerical optimization. Furthermore, these rotations weren't expected to significantly change the magnet's field, as they didn't alter the magnet blocks' magnetic dipole moments. This assumption was verified by simulating the adjusted, intersection-free magnet.

Choosing the Magnet Blocks

Another key design choice is what sort of magnet blocks will be used and how those blocks will be fabricated. In general, it is easier and less expensive to use blocks of stock size/material combinations; or blocks in easy-to-manufacture shapes and orientations (eg a rectangular block magnetized along one of its main axes).

The choice of the magnet blocks also determines some of the DoF ranges and/or constraints for the optimization problem. Specifically, the choice of acceptable block material grades and block sizes constrains the max block magnetization magnitude.

It is possible to join together individual magnet blocks to create a sub-assembly of blocks with a new shape or size. This technique can be used to create a large set of available

magnet blocks out of a smaller number of stock sizes. This approach was used in designing the MR Cap (Ch 6) and the Sparse Halbach shim arrays (Sec 5.4). In the MR Cap case this technique allowed for the use of non-stock block sizes, such as the $1'' \times 1'' \times 1\frac{3}{8}''$ block shown in Fig 3-7b. For the Sparse Halbach shim trays, this allowed for more efficient packing of magnetization inside of a circumference of given diameter than could be obtained with stock sizes (Fig 3-7d). This technique requires individual magnets to be stuck together along their axis of magnetization.

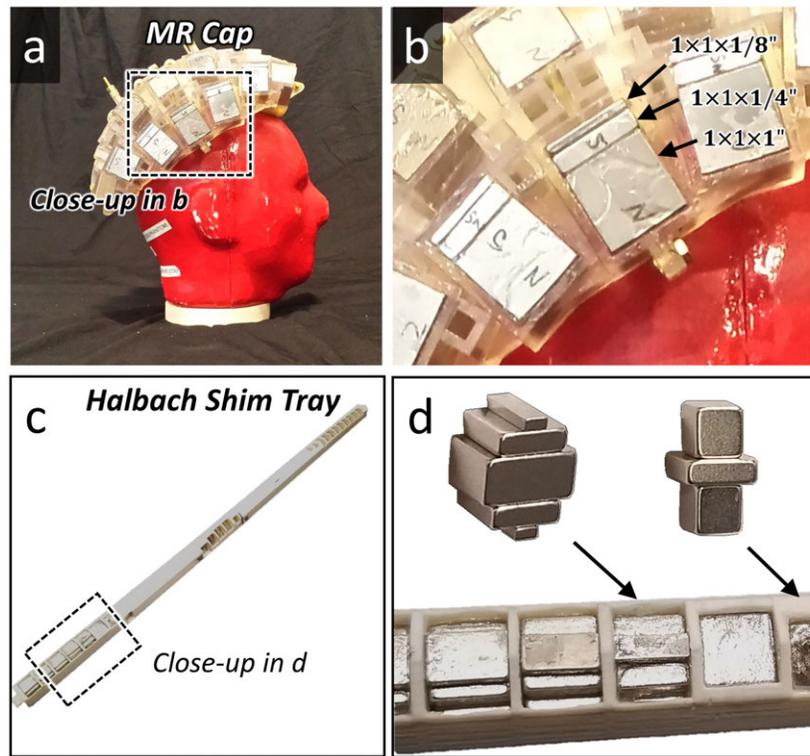


Figure 3-7: a. $1 \times 1 \times 1\frac{3}{8}''$ blocks were used in the MR Cap, and consisted of three separate blocks with $1'' \times 1''$ cross section and $1''$, $\frac{1}{4}''$, and $\frac{1}{8}''$ thickness stuck together. b. The Sparse Halbach shim trays used a variety of magnet configurations containing several small stock magnet sizes that were stuck together.

3.5.2 Setting up and Solving the Optimization Problem

Setting up the optimization problem involves picking a cost function to minimize, choosing the set of design variables to optimize over, and applying the appropriate constraints. One must also decide how to compute the cost functions and constraints, and how to solve the optimization problem.

Cost Function

The cost function is some scalar metric of magnet performance within the design ROI. In this work I use the min-max range of $|\mathbf{B}_0|$ over the ROI when designing a B_0 magnet (see Ch 6, 7). When shimming, I assume that the shim magnet field is much smaller than the \mathbf{B}_0 field, so that $|\mathbf{B}| \approx |\mathbf{B}_0 + \hat{B}_0 \cdot \mathbf{B}_{shim}|$. Under this approximation, the min-max uniformity needs only to be computed on one field component (see Sec 7.4). The cost function may also be the L^2 -norm of the residual from a linearly-varying target field (Sec 5.4), or the min-max range of the residual from a homogeneous target field (Sec 7.4).

Other work has explicitly included field gradient term in the cost function [165]. One could conceivably use a direct metric of image encoding performance as the cost function (eg the condition number of the reconstruction problem for a given imaging sequence and reconstruction resolution).

Here, the min-max metric of homogeneity is often used because that metric determines the necessary RF coil bandwidth to cover the whole ROI. Coil/pulse bandwidth is an important consideration when imaging in inhomogeneous B_0 fields. A lower required bandwidth leads to shorter pulses, reduced RF power requirements, reduced B_{1+} (Tx) and B_{1-} (Rx) artifacts, and a lower minimum readout bandwidth. These lead to faster acquisitions, fewer image artifacts, and higher image SNR.

Design Variables

Design variables may include the magnet material grades (Ch 6), magnet block sizes (Sec 5.4, Ch 7), magnet positions (Ch 6), or magnet orientations (Sec 5.4, Ch 7) and will be determined when deciding how to manufacture the magnet array. This choice has implications for the specific optimization techniques available. The number of degrees of freedom for magnets designed in this work varied from 26 for the MR Cap to 4032 for the Sparse Halbach shim array.

Constraints

Constraints on degrees of freedom often serve to impose physical reality on the numerically-computed solution. Examples include a maximum allowable magnet block size and maximum allowable magnet B_r (determined by material grade availability). These constraints

are used for all magnet designs in this work. Additional constraints may be used to impose performance requirements on the magnet array. For example, the degenerate $B_0(\mathbf{r}) = 0$ magnet is perfectly uniform, but utterly useless. This solution and ones like it are avoided by constraining the mean B_0 to be greater than some minimum value ($mean(B_0) \geq B_{min}$).

It was found for some optimization problems with continuous degrees of freedom, such as in Ch 7, that $mean(B_0) \rightarrow B_{min}$ as the solver converges to a solution. In these cases, imposing a minimum B_0 constraint effectively sets the mean B_0 for the magnet.

How to Compute

We want to compute the cost function and constraints quickly and accurately enough. The most computationally-intensive part of this is usually computing the B-field in the target ROI. There is a range of ways to do this, described in Sec 3.4. In this work, the dipole model was used for shimming problems (Sec 5.4, 7.4), the multipole model was used for continuous B_0 magnet design problems (Ch 7), superposition of precomputed fields was used for discrete problems (Ch 6).

Optimization Solvers

The main consideration in choosing a solver is whether the problem is discrete or continuous. Discrete problems can be solved with a genetic optimization. In this work, the Matlab `ga` tool is used for some discrete problems (Ch 6, Sec 7.4) following the work in [165]. Discrete or hybrid problems may also be solved by treating the discrete variables as continuous, using a continuous solver, and then approximating the solution using the discrete basis (eg Sec 5.4). It should be verified that this discretization post-processing step does not significantly impair magnet performance. In this work, the Matlab `fmincon` tool using the interior-point method is used for continuous magnet optimizations.

Certain tricks were found to improve the optimized design found by the optimizer.

- For `ga`, running multiple iterations (up to several hundred). For example, Fig 3-8 shows the results of 200 iterations of the MR Cap genetic optimization. The optimized designs had ROI field ranges between 5.3mT and 11.9mT, implying that one needs to run multiple iterations to be sure of adequately searching the design space. (See also [165].)

- Also with `ga`, it was found that introducing "dummy" options for discrete variables sometimes resulted in a better solution. For example, the HM shim array (Sec 7.4) used only two discrete options for magnet blocks: a N45 $\frac{1}{4}$ " \times $\frac{1}{4}$ " \times $\frac{1}{8}$ " block, and a null block (ie no magnet). Running the optimization problem with only these two options available resulted in the "2 Block Options" design shown in Fig 3-9. This design had an ROI field range (ie cost function) of 1.741 mT. The same problem was also run with a third intermediate option available: an N45 $\frac{1}{4}$ " \times $\frac{1}{8}$ " \times $\frac{1}{8}$ " block, shown as "3 Block Options" in Fig 3-9. This design only saw 4 out of 889 blocks use the intermediate option and had min-max uniformity of 0.721 mT. A third design was generated by converting these 4 intermediate blocks to "empty" magnet blocks. This design is a valid solution to the "2 Block Options" optimization problem, but has an ROI field range of 0.727 mT (over 50% lower). This effect was not analyzed in detail, but is something to try for the lazy and numerically-averse designer.
- The continuous problems solved in this work with the `fmincon` solver were found to converge very slowly. Often, when running an iterative solver, one sets a minimum step size for either the cost function or the distance in design space. When running `fmincon` with the min-max cost function, I set these step sizes to be as small as possible (down to the level of numerical precision) and set the stopping criterion to be 100 iterations with no change below that level. Using a higher minimum step size stops the solver prematurely, when one could still see significant improvement in design performance.

3.6 Physical Construction of the Magnet Array

Building a functional, physical prototype of the optimized design involves several steps. They are listed in order, but often one will iterate through these steps multiple times (eg the allowable performance deviation will inform how you plan to build the magnet, but the results of the manufacturing tests may conversely determine how much tolerance - and thus, performance deviation - you should expect).

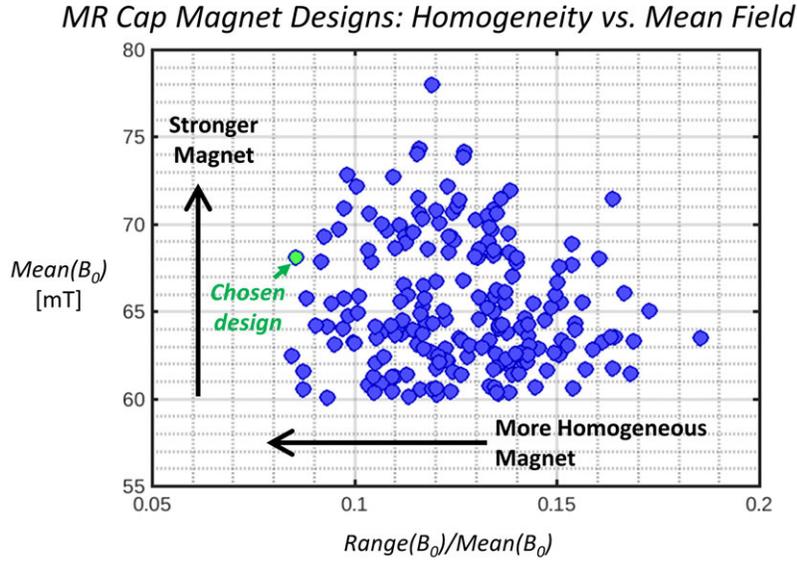


Figure 3-8: Magnet design mean field plotted against relative field range for 200 iterations of the MR Cap genetic optimization. The chosen design is highlighted in green.

3.6.1 Design Validation and Numerical Testing

First, I do any post-processing steps needed (eg discretize design variables if needed). Next, I compute the magnet's field and measure its performance using an FEM solver - this should be more accurate than the method used in-optimization (as it accounts for block-block interactions), and gives a "second opinion" to validate that the design works. I then consider sources of error that might arise when building the magnet. Their effects are estimated by adding the errors to the final design and resimulating the field. Such sources of error may include:

- Position error from mating the former pieces together, leading to position error in a large number of blocks. This is analyzed for the Helmet Magnet in Ch 7.
- Magnet B_r variation - NdFeB magnets are usually have B_r specified to a range of $\pm 1\%$.
- Random variations in B_r , μ_r , and block size.
- Systematic or random error in magnet block linear dimensions. In this work, I assumed independent uniform random distributions for all block y -dimensions and performed Monte-Carlo simulations based on that (Ch 7).

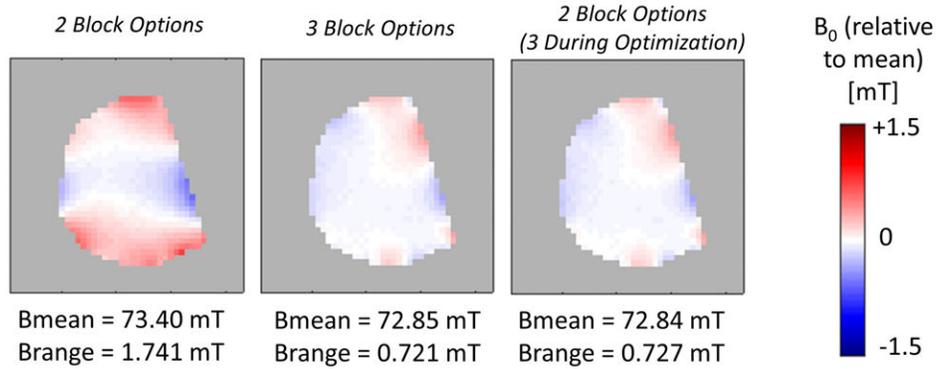


Figure 3-9: Simulated B_0 field maps (XZ plane) for the Head-Optimized magnet obtained using three different optimized shim designs. Using only two block options in the optimization produced significantly worse performance than using three designs, and eliminating the middle option as a post-processing step.

3.6.2 Process Tests

The exact procedure for assembling pieces of the magnet former (if necessary) and for bonding the magnets into the former should be tested beforehand. This is done by printing a mini-former to hold a magnet (Fig 3-10). This helps determine: strong, non-corrosive adhesive compounds for the 3D-printer material being used; tolerances needed for mating parts together and for mating magnets to the parts; and the proper orientation of the STL model relative to the 3D printer build platform.

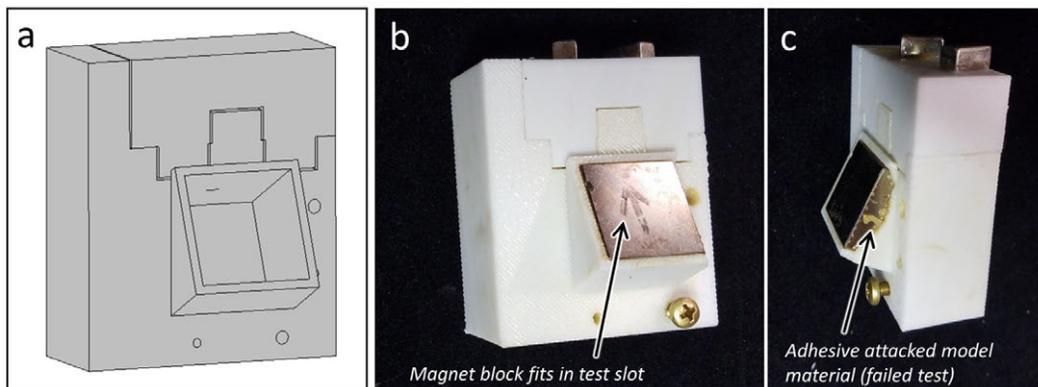


Figure 3-10: a. CAD model for a test part used to calibrate magnet slot size. b. The printed test part with a magnet satisfactorily bonded in place. c. Some adhesives were observed to attack the 3D printer model material during process testing.

Adhesive/Epoxy Tests

One should test adhesives that one might want to use. Many compatibility guides are available online, including guides supplied by manufacturers (both from the 3D printer and adhesive manufacturers - eg see: [169, 170]). A good adhesive should hold the magnet in place, but should not corrode the magnet or magnet former. For example, I found: that Loctite A330 and/or its activator attacked Stratasys PC-10 polycarbonate material (Fig 3-10c); that cyanoacrylate (Loctite 414) and 2-part 5-minute epoxy (Loctite EA608) worked well with Stratasys PC-10 and ABS-P430 materials; that Loctite 414 worked well with Stratasys ABS; and that Loctite EA608 worked well with Formlabs Acrylic FLPPCL04 resin.

The adhesive's work/curing time is also of importance. It was found in this work that fast-setting adhesives (eg Loctite 414, Loctit EA608) are useful when assembling arbitrarily-oriented magnets into a former. Often, one must use one's own two hands as clamps (given the absence of proper parallel clamping surfaces) - a situation where a long cure time is undesirable. Conversely, a long work time epoxy (eg Loctite EA-E60) is useful for complicated, lengthy operations, such as when joining two quadrants of the Helmet Magnet former with dowels, magnets, and screws (Fig 3-18).

Mating Tolerances

Any mate requires a built-in tolerance. This should be tested for both mating magnets to the former and mating multiple pieces of the former together.

In this work, magnet fit tolerance was tested by printing several test pieces with magnet slots of slightly different size (Fig 3-10a). The size of each magnet slot was the nominal magnet size, plus a different tolerance offset. For the example in Fig 3-10a, the tolerance was $0.2mm$ at each face (the total slot size thus had twice that amount added to it). Then, a test magnet was fit into the each slots. The slot where the magnet fit in snugly, but also fit all the way to the bottom of the slot without significant force, was chosen. This slot was further tested by ensuring the magnet could be bonded into it with the chosen adhesive.

The same approach was taken to choose an appropriate tolerance for mating two 3D-printed pieces together, as shown for the HM former (Fig 3-11a). A smaller "test" piece was printed with the same mating spline in order to determine the tolerance needed (Fig 3-11b).

In this step, the choice of tolerance was validated by measuring the nominal combined size of 2 mated pieces, compared with the nominal CAD model size (Fig 3-11b).

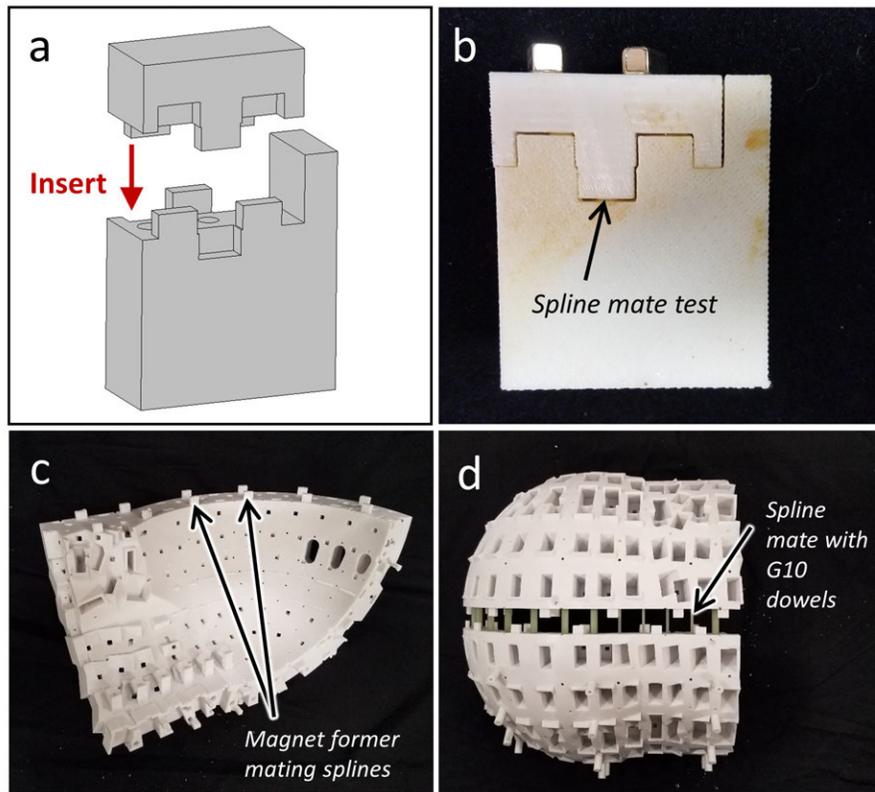


Figure 3-11: a. CAD model for a part for testing the mating structure for bonding together magnet former pieces. b. The printed test part allowed for the proper calibration of the model tolerance. c. Head-Optimized magnet quadrant showing the spline-type mate as used on the final part. d. Two magnet quadrants joined using the spline mating structure.

Configuring the 3D Printer

3D printers work by building up a part layer-by-layer. The layers are conventionally built in the XY plane, and the layers build up along Z. A 3D printer head can typically move with high precision in the XY plane, giving it a high XY resolution. This in-plane resolution varies between $20\mu\text{m}$ and $50\mu\text{m}$ for different printers. However, z-resolution is determined by layer thickness, which is often larger than in-plane resolution. For example, the Fortus 360mc used here was configured with 0.2mm layer thickness (other printers with finer layer thickness and comparable build volume are available). Many dimensions of the magnet former should be produced with sub-mm accuracy. Thus, one should take care that crucial dimensions do not perfectly align with the printer \hat{z} -axis.

The placement of support material during printing is another important consideration: it must be possible to access and remove it from the model. It may be prudent to use a dissolvable support material (Fig 3-12) for parts that must be printed with difficult-to-access support structures. This was done for the Sparse Halbach shim trays (Sec 5.4). Finally, some materials/printers (such as SLA printers) use the model material for building support structures, often resulting in the part having "nubs" where the supports attached (Fig 3-13a). It should be checked beforehand that these can be smoothed or otherwise pose no issue to assembling the magnet.

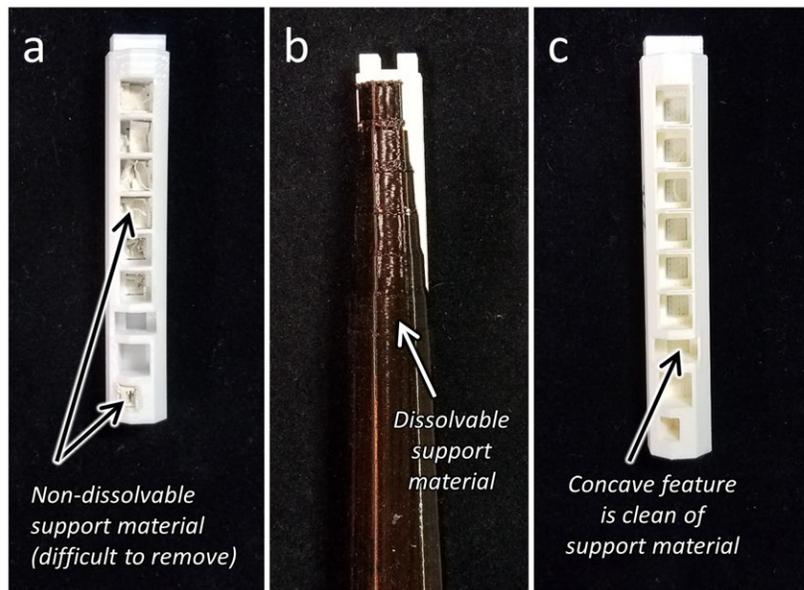


Figure 3-12: a. A test-print of a Sparse Halbach shim tray created with standard support material. The support material could not easily be removed from several small, concave regions of the model. b. Test Sparse Halbach shim tray printed with water-soluble support material. c. The same test shim tray after dissolving the support material.

3.6.3 Magnet Former CAD and Printing

The goal of this step is to go from the final, post-processed magnet design to a printable STL. This requires the synthesis of the final design, manufacturing/tolerance tests, and the system-level design. The basic process involves manually creating a magnet former in CAD, using a numerically-generated CAD model of the magnet blocks to create slots for the magnets to go, and preparing the model for 3D printing.

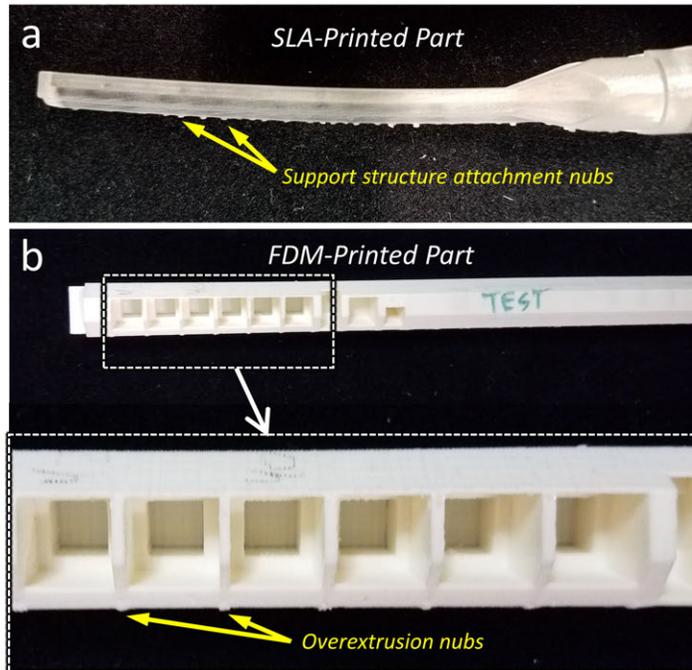


Figure 3-13: a. SLA printers use the model material to build support structures. Undesired "nubs" may be left on the model after removing these support structures. b. FDM printers may overextrude material when building models. These "burrs" may need to be manually cut or filed off.

Designing the Former Base

The magnet former base must have the proper geometry to fit in with the rest of the planned system (Sec 2.3) and to hold the magnets, given their positions, sizes, and other degrees of freedom (Sec 3.5.1). I find it best to start with a basic shape that fits the system-level design; shown in Fig 3-14a is the basic shape for the Head-Optimized MRI magnet. The main dimensions of this former is determined during the system-level design steps. To this rudimentary formers are then added all the holes, standoffs, bosses, and other features needed to allow mechanical interfacing with the B_0 magnet. Shown are standoffs and holes in the Helmet Magnet former for attaching gradient coils, shim trays, RF coil mechanical components, etc. (Fig 3-14c+d). At this stage, almost is everthing included in the CAD model besides the slots for the permanent magnets themselves. Note that some of these features (eg hole positions) may rely on knowing the final positions of the magnets blacks. It may be necessary to either add some features after including slots for the magnets, or to reiterate this step.

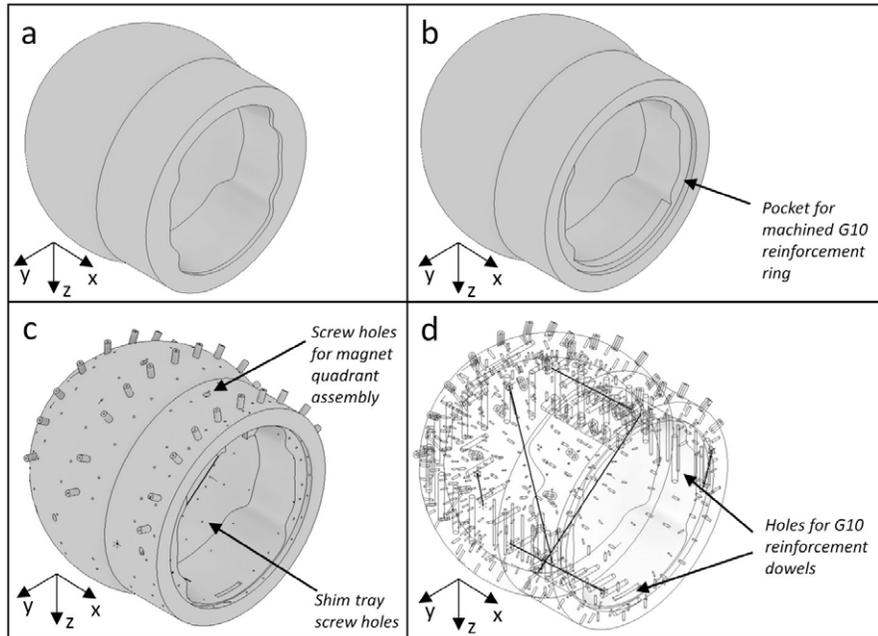


Figure 3-14: a. A basic magnet former shape is created to match the requirements of the system-level design. b. Mechanical provisions, such as regions for structural reinforcements are then added. c. Further features for attaching other system components are added to the model. d. Transparent-model view showing internal holes for further structural reinforcements.

Numerically-Generated Magnet Block Models

Incorporating the magnet design with a large number of degrees of freedom is done by automatically generating an STL, importing that into the CAD program, and performing a sequence of binary operations between the former base and the STL model. Multiple STL models need to be generated. These include:

- *Magnet Block STL* (Fig 3-15a): A model of the magnets themselves, plus the slot tolerance described in Sec 3.6.2. This can be used to etch away the model, leaving empty space where the magnet can be placed.
- *Extended Block STL* (Fig 3-15b): A model of the magnets, plus the slot tolerance, that has been swept along one axis. This axis is the direction along which the magnet will be inserted into its box. The extended STL model is used to etch the model, leaving a slot and clearing the space needed to insert the magnet.
- *Magnet Holder STL* (Fig 3-15c): A model of the magnets with all faces thickened by a specified dimension (usually $3mm$ in this work). This is used to create boxes

that completely surrounds the magnets, and is eventually etched by another STL (eg *Insertion STL*) to leave slots for the magnets.

- *Access Hole STL* (Fig 3-15d): A set of blocks with smaller cross section, and swept the opposite direction relative to *Insertion STL*. This is used to make a hole in the magnet box, which provides mechanical access to the bottom of the magnet and provides a pressure relief opening for epoxy when forcing the glue-covered magnets into their slots.

The essential binary operations are: union of the former base with the *Block STL*, shown in Fig 3-16a; and the subtraction of *Insertion STL* from the result (Fig 3-16b). The Access Hole STL may be subtracted from the magnet former model as a final step (Fig 3-16c).

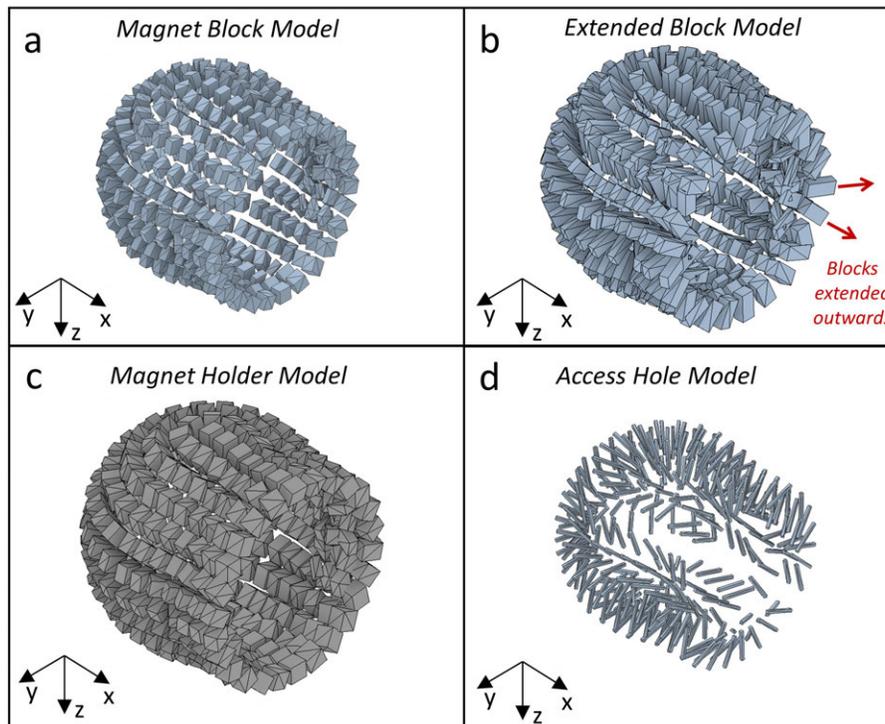


Figure 3-15: a. An automatically-generated STL of the optimized magnet block set. b. Auto-generated STL of the optimized magnet blocks, with each block swept outwards along its insertion direction. c. Auto-generated STL of the thickened optimized magnet blocks. This is used to create "coffers" that hold the magnet blocks. d. STL model for creating access holes in the bottom of the magnet block "coffers".

These operations may affect some of the holes and other features added to the former during earlier steps, and it may be necessary to re-build certain features (eg. re-bore holes). Features that weren't previously included in the former base because they depend on the

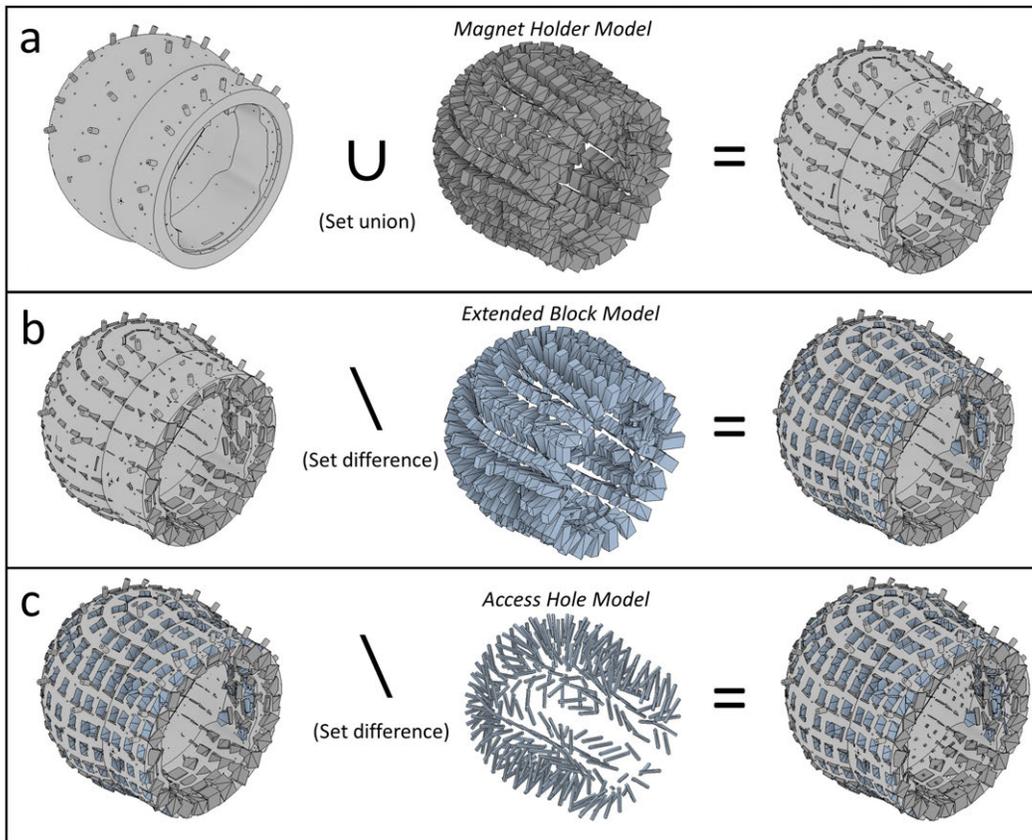


Figure 3-16: a. The manually-designed magnet former and the magnet holder STL model are joined together in CAD. b. The resulting union is then "etched" by taking the set difference with the magnet block insertion STL model. c. Another set difference operation is performed between the result and the access hole STL model.

positions of the magnet blocks should also be included at this stage.

Effects on the Other System Components

The binary operation of the former with *Block STL* generally adds volume to the former base; this augmented magnet former may intersect other system components (eg. gradient coil formers). It is important to keep track of changes like this. This problem can often be addressed by creating another STL that can be used to etch other system components, leaving space for it to mate with the modified B_0 magnet former. In generating this STL, it is recommended to include another tolerance/offset to allow for a free mate between the system parts.

Printing the Model

The B_0 magnet former model will often need to be printed in multiple pieces in order to fit in the printer's build volume. One must consider about how structural the mating interfaces will be when splitting up the magnet former model. For the MR Cap, this step involved face mate with a rib, epoxy (Loctite EA608) holding the faces together, and screws/nuts squeezing the former quadrants at flanges. The Helmet Magnet mating surfaces used a system of teeth included in the CAD model; G10 dowels (bonded in holes); epoxy (Loctite EA-E60) on the faces; and screws engaging with tapped holes in the former (Fig 3-17). Sparse Halbach shim trays were printed in two pieces and glued together in the middle with Loctite 414. These shim trays did not see significant axial or bending loads at this joint.

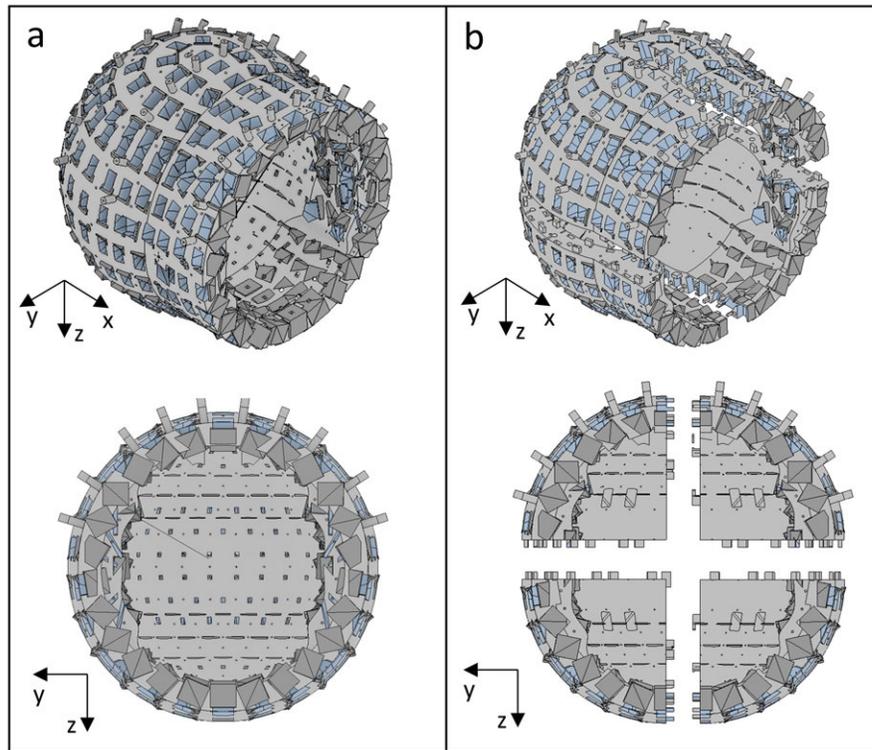


Figure 3-17: a. The Head-Optimized MRI magnet former shown as a monolithic part. b. The same magnet former split into four pieces for 3D printing. In addition to just splitting the model, mating splines and holes have been added where the pieces join together.

3.6.4 Assembling the Magnet Array

It's finally time to assemble the magnet. First, do any necessary mechanical steps, such as: removing support material/cure the printed parts; deburring/cutting off overextrusions;

widening and/or tapping holes; or other things particular to your design. Next bond former pieces together (if needed). In some cases, I bonded some magnets in during this step - they can serve as alignment fiducials and spacers/gauge blocks for the printed former pieces (Fig 3-18). Be careful during this step: errors here will propagate to errors in relative block positions.

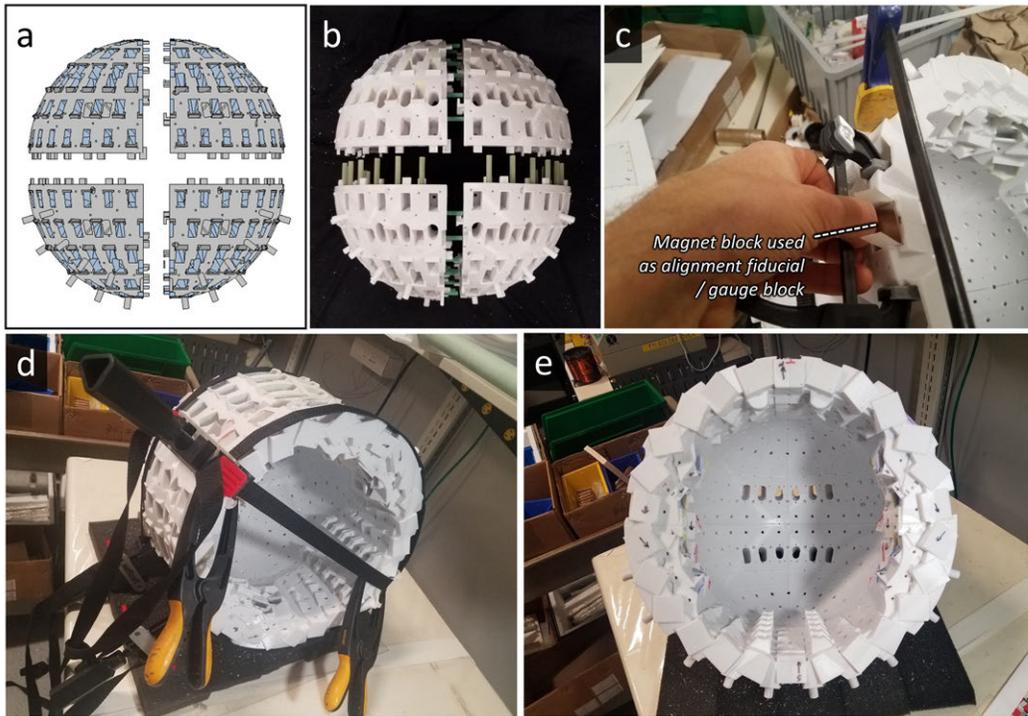


Figure 3-18: a. CAD rendering of the split magnet former model, ready for 3D printing. b. Partially-assembled, printed 4-part magnet former. c. Magnet blocks sometimes straddled the mating surface between magnet former pieces. These blocks served as alignment fiducials and spacers/gauge blocks when assembling the former. d. The magnet former pieces were set together with the proper reinforcement dowels, screws, and epoxy, and were clamped together until cured. e. Front view of the final, assembled Head-Optimized MRI magnet former.

Next, the magnets are bonded into place. I found it best to use clamps if possible (Fig 3-18). Be sure to clamp/hold the magnets in for the full adhesive setting time. Sometimes one can add magnets in an order where the magnetostatic force on a working magnet tends to pull it into its slot (thereby clamping it); this is preferable to the partially-assembled magnet array tending to push out the working magnet block. Consider the example of a ring of blocks in a Halbach orientation (Fig 3-19). The equatorial magnets should first be glued in place. Next, the magnets immediately above and below that magnet should be added. The orientation of the blocks means that these magnets will be pulled into their

boxes. This procedure can be continued several rungs up the ring.

When clamping with your hands, take breaks as needed to rest your thumbs.

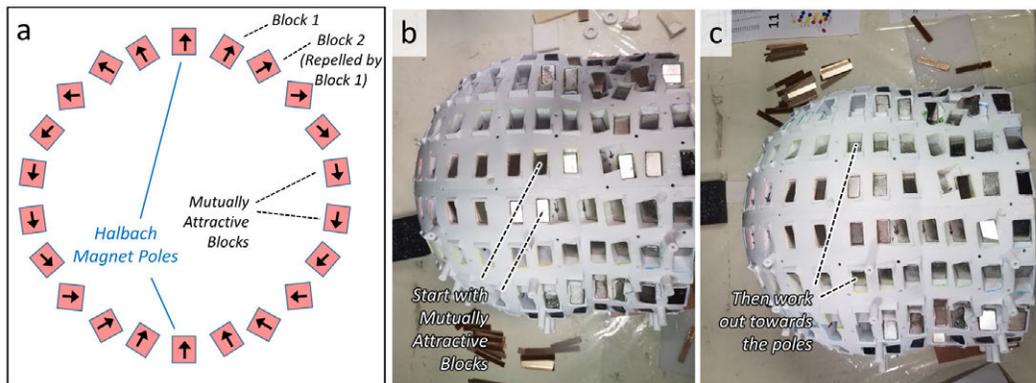


Figure 3-19: a. Transverse section schematic of a discretized Halbach magnet. Blocks near the equator are mutually attractive. Blocks near the poles may tend to be repelled by other blocks and pushed out of the magnet former. b. When populating a Halbach-type magnet, it is easiest to start with the mutually-attractive blocks near the equator. c. Blocks can then be added moving out towards the poles. This ensures that blocks will always be pulled into the magnet as they're added in.

Chapter 4

Design and Construction of Gradient and Low-Frequency RF Coils

4.1 Abstract

RF Coils are used by all MRI scanners for signal excitation and detection. Most scanners, including those in this work, also use gradient coils for image encoding. Both RF coils and gradient coils are operated quasi-statically in low-field ($B_0 \leq 0.2T$) MRI scanners. This simplifies coil field computation and design, and streamlines the design of numerically optimized gradient and RF coils. Here I describe my approach to target-field Boundary Element Method (BEM) computation quasi-static coil design for portable MRI systems. I then describe techniques for constructing physical versions of these numerical designs.

4.2 Prior MRI Coil Design Work

This work builds on prior work designing surface current patterns using current stream functions [171]. Much of the technical implementation is based on code from the toolbox in [172]. Target field and current stream function approaches have been used to design resistive MRI magnets [173], high-field system gradient coils [81, 174], and shim coils on non-cylindrical geometries [175]. It can also be used to design RF coils at sufficiently low frequency (such as for head-sized coils at 3MHz).

4.3 Quasi-static Coil Overview

The quasi-static method is a way to solve for E- and B- fields in slowly-varying systems where EM wave behavior is minimized. This is the case for low-field ($B_0 \leq 0.2T$) RF coils and for gradient coils. This assumption allows us to use Ampere's law to compute B-fields for electric current distributions, simplifying the design of coils by enabling Boundary Element Method (BEM) field computation in a stream function target field approach.

4.3.1 The Quasi-static Approximation

Maxwell's equations describe the behavior of E- and B-fields given charge density, current density, magnetic permeability and dielectric permittivity. Countless texts describe this physics [1]. In the quasi-static case, we neglect some of the time derivative terms in Maxwell's equations. When designing coils for MRI, we are generally interested in computing the B-field, so we start with the following equation:

$$\oint_C \mathbf{B} \cdot d\boldsymbol{\ell} = \iint_S \mu \mathbf{J} \cdot d\mathbf{S} + \iint_S \mu \varepsilon \frac{d}{dt} \mathbf{E} \cdot d\mathbf{S} \quad (4.1)$$

Ignoring the second term in Eq 4.1 gives:

$$\oint_C \mathbf{B} \cdot d\boldsymbol{\ell} = \iint_S \mu \mathbf{J} \cdot d\mathbf{S} \quad (4.2)$$

which is just the Ampere's Law / the Biot-Savart Law. This approximation incurs a fractional error in the computed magnetic field proportional to $\varepsilon \mu \omega^2 L^2 = \left(\frac{L}{\lambda}\right)^2$ where ω is the field angular frequency, λ is the wavelength in a medium with permittivity ε and permeability μ , and L is the length scale over which we integrate (ie the characteristic length of our system). This approximation is therefore valid when $L \ll \lambda$.

We are designing coils that generate and detect fields inside the human body. In the anatomical medium, $\varepsilon_r \approx 50$ and $\mu_r \approx 1$. Assume for head imaging that $L = 30cm$. The equality $L = \lambda$ then occurs when $f = 141.4MHz$. This is the Larmor frequency when $B_0 = 3.32T$ and is well above typical gradient switching frequencies. Thus, for low-field ($B_0 \leq 0.2T$) scanners, we are safely in the quasi-static regime for both gradient and RF coils. In this work, we also ignore eddy current effects when optimizing coils. (Note that the quasi-static approximation does not itself ignore eddy current effects; these perturb the

magnetic field via Lenz's law ($\nabla \times \mathbf{E} = -\frac{d\mathbf{B}}{dt}$) and Ohm's law ($\mathbf{J} = \sigma\mathbf{E}$) and not through time-varying electric fields.)

This approximation was numerically verified for an RF transmit/receive coil operating in a narrow band about $f = 3.385\text{MHz}$. The coil field was simulated using the Biot-Savart law (using the Biot-Savart software package) and using a full-wave field solver that made no approximations (HFSS). The simulated quasi-static and full-wave fields are shown in Fig 4-1. This approximation encapsulates another simplifying assumption: unlike in high-field RF coils, the coil spatial field patterns at this frequency are subject-independent. The effect of a human in the scanner is primarily to change the dielectric permittivity distribution $\sigma(\mathbf{r})$, which plays no role in quasi-static B-field computation.

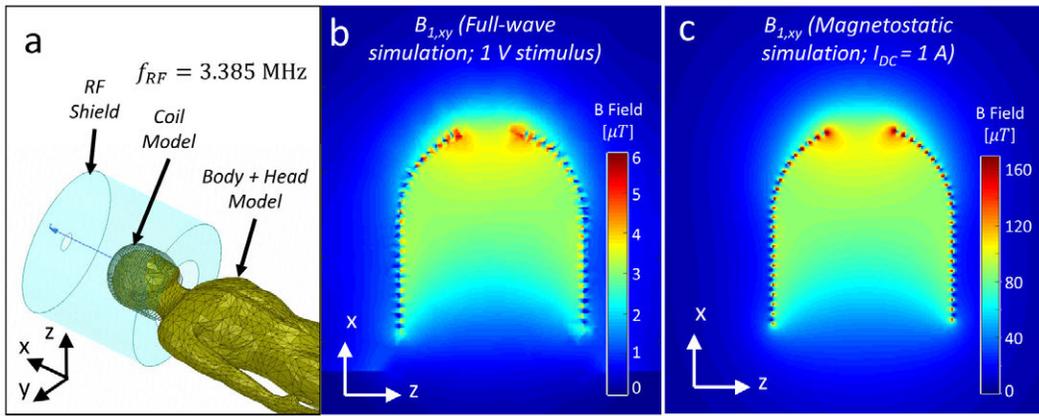


Figure 4-1: a. A full-wave field simulation was performed using a tuned 21-turn RF coil ($f = 3.385\text{MHz}$), a human body model, and an RF shield using HFSS. b. The simulated $B_{1,xy}$ field map in the XZ plane. c. The simulated magnetostatic/DC field produced by the same RF coil model.

4.3.2 The Current Stream Function

The differential form of Eq 4.2 is given by a form of the Biot-Savart Law:

$$\nabla \times \mathbf{B} = \mu\mathbf{J} \quad (4.3)$$

Under this quasi-static approximation, we can ignore E-fields, and therefore charge density (which only plays a role in the now-unimportant equation $\nabla \cdot \mathbf{E} = \frac{\rho_e}{\epsilon}$). The charge conservation equation is given by:

$$\nabla \cdot \mathbf{J} + \frac{d\rho_e}{dt} = 0 \quad (4.4)$$

Since we can assume $\rho_e(t) = 0$, this means $\nabla \cdot \mathbf{J} = 0$, implying we can specify \mathbf{J} as the curl of a vector potential.

Furthermore, when designing coils, we typically solve for an optimal surface current density localized to a 2D surface. In this case, vector potential must always point perpendicular to the 2D surface, and thus has only one non-zero component. The surface current density is then fully described by a scalar field ψ on the 2D surface, where ψ is the Surface Current Stream Function [171].

4.3.3 Target Field BEM Optimization

In this work, gradient and RF coils are designed using a target field method [172], which aims to produce a specified field (the "Target Field") at a chosen set of points (the "Target ROI"). Fig 4-2 illustrates such an ROI and target field pattern. The method uses a cost function of the form: $f_{cost} = norm(B_{coil}(\psi(\{\mathbf{r}_i\}) - B_{target}(\{\mathbf{r}_i\}))$. We then solve for the stream function ψ at the points $\{\mathbf{r}_i\}$ located on a surface, from which we can then get the surface current density (Fig 4-2). The surface is represented as a mesh; $\{\mathbf{r}_i\}$ are the mesh nodes. Associated with each node is a basis function (as in BEM/FEM problems) and a normal vector; these allow us to quickly compute the field produced by a given stream function. This computation is simply a matrix multiplication given by the stream function $[\Psi]$ left-multiplied by a precomputed geometry-dependent matrix $[D_{BEM}]$.

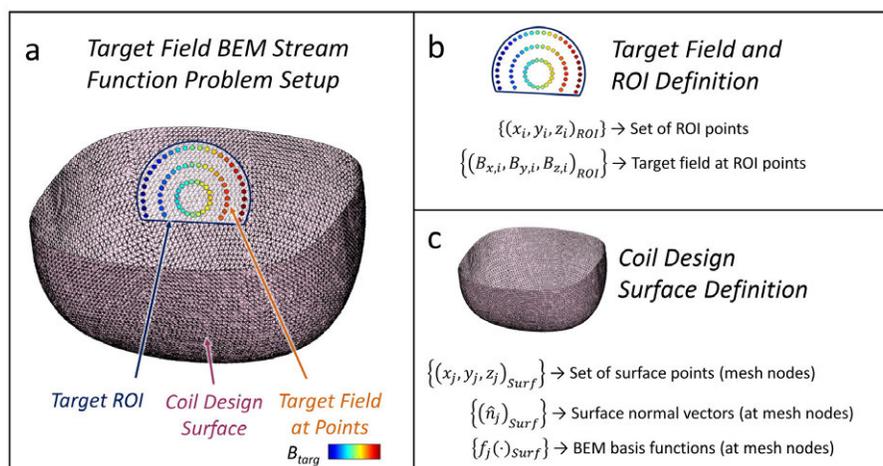


Figure 4-2: a. The geometry of a target field stream function optimization problem. Shown are the target ROI (with target field) and the coil design surface mesh. b. The target field and ROI are represented by a set of ROI points, and a set of target field values at those points. c. The coil design surface is described by a set of points, a normal vector at each point, and at BEM basis function at each point for field computation.

Other objective functions (eg coil inductance, coil resistance, stream function regularization) can be added to the optimization problem []. In this work, coil inductance and coil resistance were not included in the optimization. It was found by experience that inductance was not a major concern for unshielded coils of this size. Coil resistance was a design consideration, but was included in the design process in a different way (see Sec 4.4.1).

This work use target field BEM stream function optimization code based on the published toolbox [172]. All coil optimization problems incorporated L^2 regularization of the stream function and the boundary condition $\psi(\mathbf{r}_b) = 0$ for $\mathbf{r}_b \in \partial\Omega$ (with $\partial\Omega$ as the boundary of the mesh).

4.4 Coil Design

In this work, gradient coils are made by winding enameled magnet wire onto a former; RF coils use Litz wire wound on a former. The details of how the coils are built affect how they may be used in sequences and what amplifiers may be used to control them. In designing a coil, these constraints need to properly incorporated into the numerical design process. After this, the numerical optimization problem can be solved to generate an optimized stream function. The stream function is then post-processed (incorporating some of the aforementioned physical constraints) to generate a set of wire windings that fully describes the coil.

4.4.1 Physical Constraints and Process Tests

Gradient Coils

A major physical constraint when designing truncated coils (as in this work) is packing wires near the truncated end. This is illustrated in the single-sided gradient coils in Fig 4-3, where adjacent windings have clearance everywhere in the coil except near the top (truncation) edge. This impacts several design choices, including the number of layers of windings and the thickness of the wire used. Another consideration is DC resistance. In this work, we want coils to have $R_{DC} \leq 2\Omega$. This allows them to play arbitrary current waveforms with high fidelity using commercial GPAs that plug into standard wall outlets. It also enables the use of low-cost, compact "tabletop" gradient amplifiers [], as used for the MR Cap (Sec 6.5.2). This constrains coil length and wire diameter.

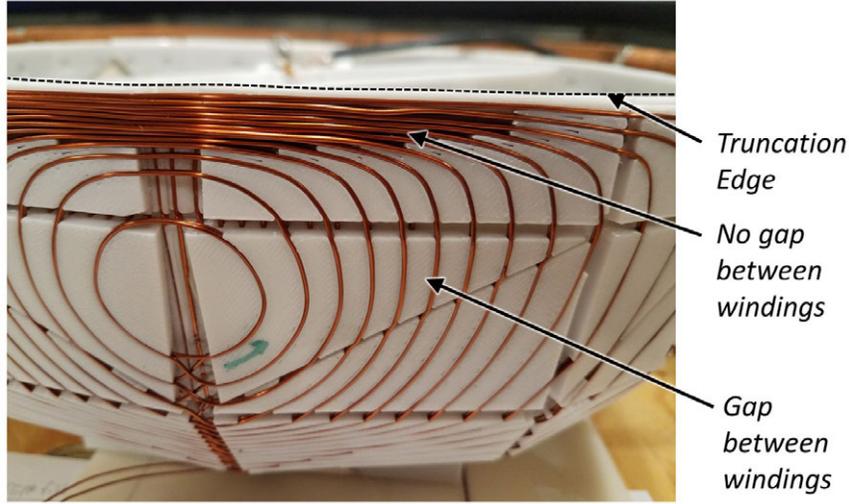


Figure 4-3: Close-up of the MR Cap gradient coils showing the truncation boundary of the surface. The windings are densest near the truncation edge, and have ample space away from it.

In this work, in-plane phase encode and readout encode coils were designed to produce at least $5 \frac{mT}{m}$ at $I = 10A$. This constrains coil field efficiency such that $Eff \geq 0.5 \frac{mT}{m \cdot A}$. A coil with this efficiency enables $2mm$ PE resolution with a rectangular $1ms$ gradient blip at a peak current of $3.74A$. Also, each gradient coil described here contains two layers of windings (directly on top of each other). For $D \approx 1mm$ wire, the coil is still quite thin. This doubles efficiency and DC resistance and quadruples inductance. Peak power is halved compared with keeping the same windings but doubling the current. This was found by experience to produce coils with acceptable DC resistance, efficiency over the design ROI, and inductance, all in an adequate geometrical form factor.

Adding winding layers also allows for design flexibility, especially in cases where inductance is not a major concern. For example, consider the case in Fig 4-4: adding layers of windings and using thicker wire increases efficiency and decreases DC resistance, at the cost of increased bore space (greater thickness) and increased inductance. These design choices can be imposed during post-processing steps. Wire spacing (and thus minimum permissible wire size) can be increased by smoothing the computed stream function, and the addition of a wire layer can be implemented by doubling the set of windings derived from it.

RF Coils

In this work, RF coils have a higher minimum wire spacing than gradient coils to prevent arcing. These coils are generally tuned to resonate a specific frequency with quality factor

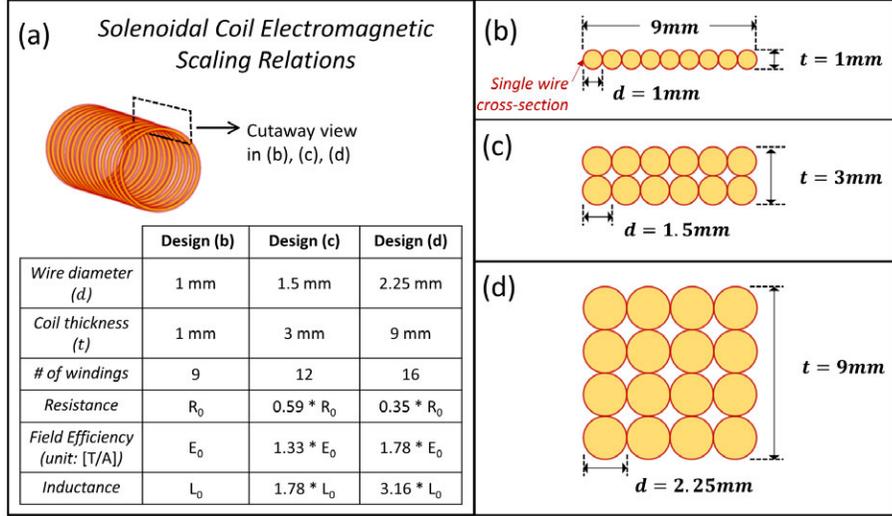


Figure 4-4: a. Scaling relations for electromagnetic properties of a solenoid coil with varying wire diameter and number of winding layers. Values of thickness, resistance, field efficiency, and inductance are calculated for: b. a single-layer coil with wire diameter of 1mm; c. a two-layer coil with wire diameter of 1.5mm; and d. a four-layer coil with wire diameter of 2.25mm.

$Q \approx 10$ (such as for the Head-Optimized MRI coil). The Head-Optimized MRI coil consisted of 4 distributed tuning capacitors. Thus, during RF transmission at $f_{RF} = 3.09MHz$ and peak power ($P_{max} = 423W$), the coil must store energy $E_{coil} = \frac{P_{pulse}Q}{2\pi f_{RF}} \approx 0.22mJ$. This is distributed over the 4 tuning capacitors, and corresponds to a peak capacitor voltage given by $V_{C,pk} = \sqrt{\frac{2E_{coil}}{4C}}$. With $C = 350pF$, this gives $V_{C,pk} = 560V$. This internal voltage and the dielectric breakdown field of dry air of $E_{max} \approx 3 \cdot 10^6 \frac{V}{m}$ together give a critical "breakdown gap" of $0.2mm$. In practice, it is wise to increase this limit by a safety factor (for the Helmet Magnet RF coil, a minimum gap of $1mm$ was used). Arcing considerations also mean multi-layer RF coils are not used in this work.

Process Testing

The coils are next constructed by press-fitting wire into grooves in a 3D printed former. It is important to get the wire fit to be correct – not too tight, but not too loose. This calibrated groove width is required for CAD model generation (Sec 4.4.4). The proper fit can be determined by test-fitting a wire into a set of grooves with slightly different widths. Fig 4-5 shows this being done for AWG18 wire (nominal diameter = $1.024mm$). Here, 11 grooves were printed with widths from $0.824mm$ to $1.324mm$ in $0.05mm$ increments, and the $1.074mm$ -groove worked best. This dimension was used to generate the CAD model to

make the wire grooves. Note that near the ends of the truncated gradient coils, the wires are so close together that there will generally not be grooves to hold each winding in place (Fig 4-3).

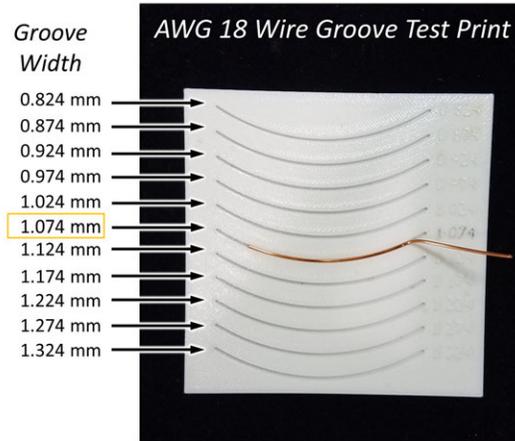


Figure 4-5: Test-print piece for calibrating the optimal groove width for press-fitting AWG18 enameled wire. Grooves were printed between 0.824mm and 1.324mm in 0.05mm increments. The nominal copper diameter for AWG18 wire is 1.024mm. The 1.074mm groove was found to give the best fit.

4.4.2 Generating a Surface Mesh

A 2D stream function surface is needed for BEM optimization. Volumetric aspects of coil (thickness, using actual wire, etc) come about during post-processing. Ideally, the design surface should be at the mid-level of where the current will be flowing in the physical, volumetric coil. The surface size, position, etc are determined by the system-level design.

We next need to turn the surface into a mesh. If the surface is parametrized, then the mesh can be created directly from the parametrization. This was done here for piecewise cylindrical, spherical, and planar coils (Fig 4-6). It is also possible to export surfaces to STLs in many CAD programs, such as Inventor (as used here). However, the meshes may contain facets with vastly different aspect ratio and size, and nodes connected to different numbers of other nodes. This effectively preconditions the optimization – in particular the regularization term – in an undesired way and can lead to numerical artifacts in the solution. This can be avoided by combining the proper surface geometry from CAD with a parametric mesh for a different surface to produce a well-conditioned mesh of an arbitrary surface. In the MR Cap example (Fig 4-7), this is done by fitting the spherical angular coordinates for the parametric mesh onto the CAD-generated surface mesh by interpolating the radial

coordinate values.

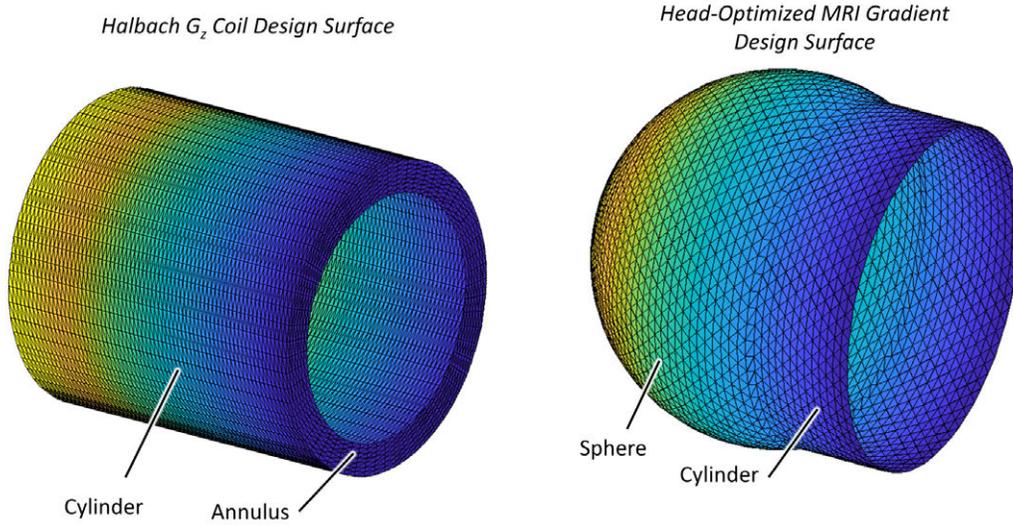


Figure 4-6: Parametrized meshes used to design the Sparse Halbach G_z gradient coil and the three Head-Optimized MRI gradient coils. The mesh coordinates were generated directly from the cylinder, annulus, and sphere surface parametrizations.

4.4.3 Setting Up and Running the Optimization

The details of the optimization implementation are described in [172]. In this work, I used the objective function:

$$f(\Psi) = \|D_{BEM}\Psi - B_{target}\|_2^2 + \lambda_r \|\Psi\|_2^2 \quad (4.5)$$

The regularization parameter λ_r was often just set to $\lambda_r = 4 \cdot 10^{-6}$. In some cases however, the optimization was solved with multiple values of λ_r . The solutions were then analyzed for field performance (ie the field target cost function) and for ease of construction (eg wire spacing, wire bend radii). A single solution was then manually chosen from this set of possibilities.

The field target varied depending on the system and the type of coil being designed. For gradient coils, the standard target is for the \mathbf{B}_0 -aligned field component to vary linearly with a spatial coordinate. This was used for the Sparse Halbach and Helmet Magnet coils. The MR Cap gradient coils incorporated additional 2nd-order spatial field terms in the field target. This had the effect of countering the tendency of these single-sided gradient coils to naturally vary in field strength in the x -direction (See Fig 6-9). The MR Cap RF coil used a

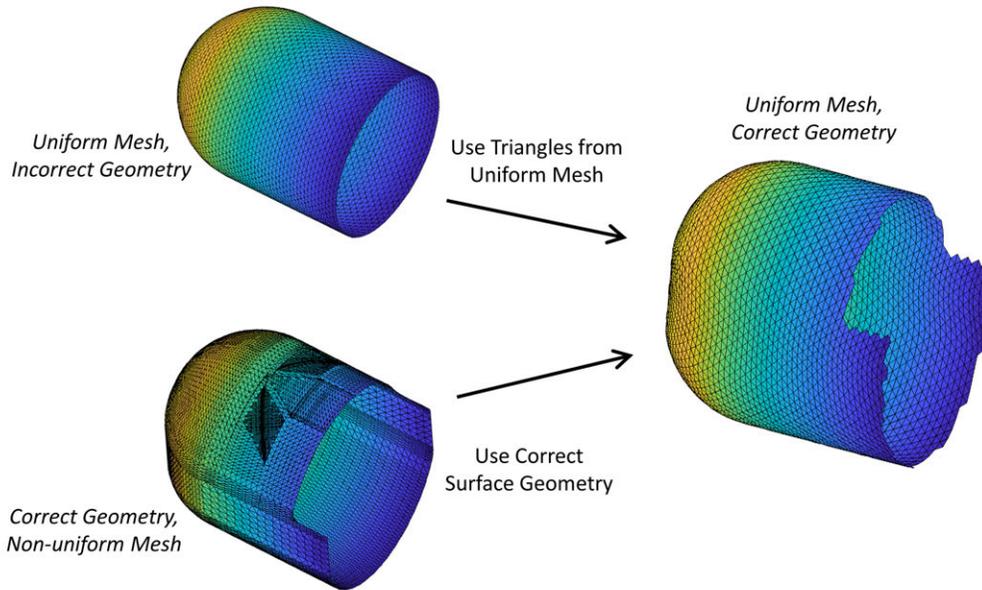


Figure 4-7: A uniform mesh can be generated for a non-parametrized (ie arbitrary) surface. A uniform, parametrized mesh for a different surface provides the triangles, but not the node coordinates. A non-uniform mesh with the correct node coordinates (such as that produced automatically by a CAD program) provides the surface geometry. The uniform mesh nodes are interpolated onto that surface, producing a uniform mesh with the correct surface geometry.

uniform field target for one field component oriented transverse to \mathbf{B}_0 . The Helmet Magnet RF coil optimization used a uniform field target within the imaging ROI, along with a null field target (ie no field produced) just inferior to the neck ROI. The aim was to localize coil sensitivity to the region of interest, and to explicitly not be sensitive outside that ROI (See Sec 7.6.1).

The set of target field points can also be freely-chosen. The natural first choice is the system-level design ROI. However, RF coils may fit so close to the target ROI that the discretization of the surface mesh and the target ROI can cause unwanted numerical artifacts in the resulting optimized stream function (Fig 4-8). In this situation, one can use a smaller target ROI with greater clearance from the design surface.

4.4.4 Creating a Coil CAD Model

After solving for the optimized stream function, we convert it into a set of 1D paths (for magnetostatic field simulation) and a space-filling 3D model (for use in coil former CAD design).

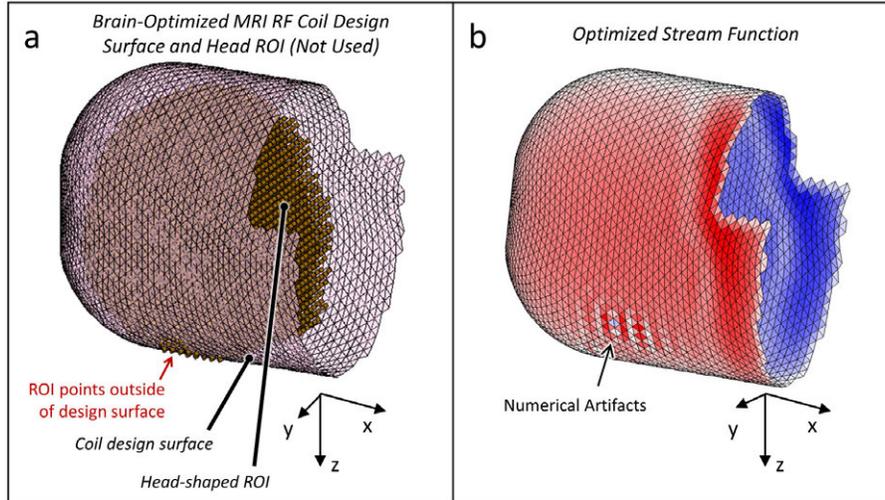


Figure 4-8: A target ROI that is too close to the stream function design surface can result in numerical artifacts. a. The close-fitting Brain-Optimized MRI RF coil design surface intersects the brain-shaped target ROI. b. This results in numerical artifacts in the optimized stream function.

Generating 1D Winding Paths

A surface stream function is approximated by a discrete set of current paths by taking equally-spaced level sets of the stream function. Each level set/isocontour defines a 1D path representing a wire carrying current, and all paths carry the same current. In this work, this is done by projecting the 2D stream function from the coil design surface onto a flat 2D surface (Fig 4-9b). This is shown for a cylindrical example (the Sparse Halbach Gx coil) and a spherical example (the Head-Optimized MRI Gy coil). Level sets/winding paths are then computed on the flat surface (Fig 4-9c) and projected back onto the 3D surface (Fig 4-9d). In going between surface geometries, the distance metric between points in the mesh is altered. This effect was not accounted for in the level set calculation or smoothing operations.

In this work, the stream function is smoothed with a Gaussian filter prior to discretizing it as a set of isocontours. This smooths over meshing artifacts and softens out sharp edges in the stream function. The latter effect is beneficial if one needs to increase the distance between adjacent windings (See Sec 4.4.1). The number of level sets chosen determines the number of windings and indirectly affects wire spacing. In this work, the number of level sets used ranged from 20-40 for gradient coils to 4-20 for RF coils.

After designing the winding paths, one may want to smooth them to increase bending radii and reduce kinks. Decimating and/or interpolating the winding paths was found to

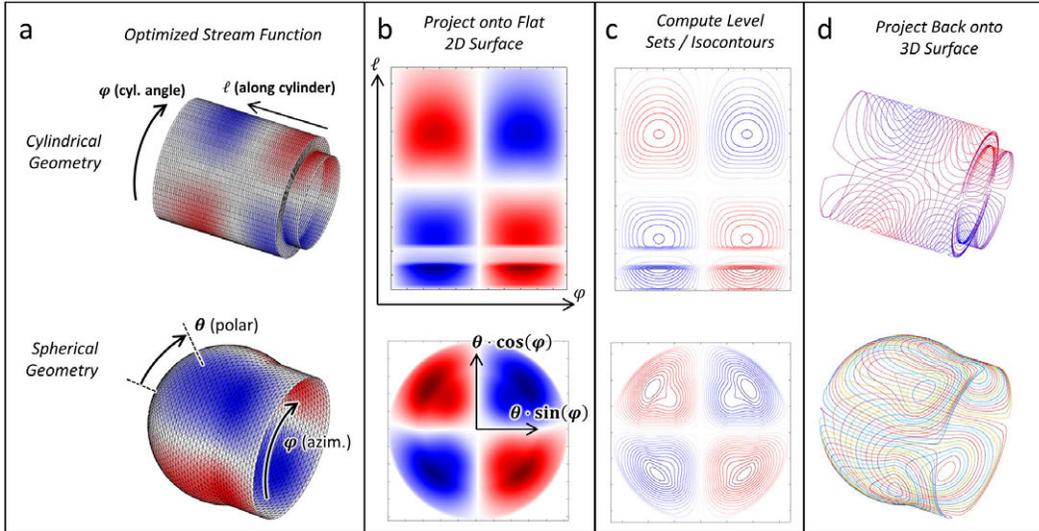


Figure 4-9: a. Stream functions may be designed on surfaces with cylindrical or spherical geometry. In both cases, the surface can be parametrized by two variables. b. The stream functions are next projected onto flat surfaces parametrized by these two variables. c. Isocontours/winding paths are computed on the flattened stream function. d. These windings paths are then projected back onto the original 3D surfaces.

sometimes help with generating a 3D volume model of the wires. In some cases, one may wish to build the coil on a different surface from what was used for design. For example, the Sparse Halbach gradient coils were designed on a largely cylindrical surface, but the actual coils were built on a 16-sided prism (approximating a cylinder). Implementing this requires fitting the winding path coordinates to the desired surface. The re-fit set of windings should be simulated again to ensure that performance has not suffered.

It is also good to measure the total length of the coil during this step. This number was in the 10's of meters for the coils here, and tells you the minimum continuous length of wire that you need to build the coil.

Generating a 3D Volume Model

Finally, we need to turn the 1D windings into 3D volumes that represent the grooves to be etched into a CAD model for holding wire. This procedure involves creating 4 sets of paths with the appropriate offsets from the 1D level sets. These points on these paths will be the vertices of the coil winding STL volume. (Fig 4-10)

Depending on procedure used, one may need to smooth the 1D paths again to prevent mesh errors (eg intersecting faces). The desired wire groove width (including fit tolerance)

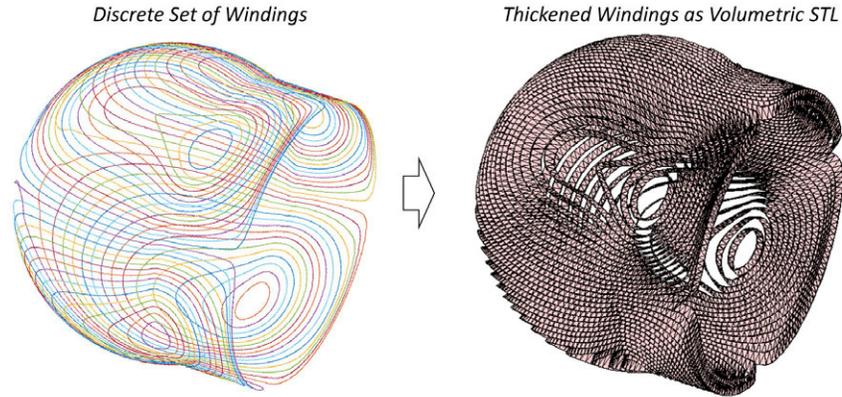


Figure 4-10: A set of computed coil winding paths is a set of 1D curves. These curves are "thickened" into 3D volume models that correspond to the desired wire press-fit grooves.

must be accounted for at this stage (Sec 4.4.1). Note that even though the wire typically has a circular cross-section, there is no problem making a flat bottomed groove. Doing so simplifies this mesh creation step.

One issue that may arise pertains to the closely-spaced wires near the truncated end of the coil. Some of these windings may only be spaced by the critical separation, and their winding groove volumes may therefore intersect. This is not a physical problem as long as the critical separation rule is still obeyed. However, it can cause numerical issues if the mesh volumes intersect. One workaround involves creating separate meshes for each winding, saving them as separate STL files, and merging them as solid bodies in the CAD program used (Inventor). In any case, the finished mesh should be saved as one or more STL files (or another mesh file format) for the next step.

4.5 Coil Construction

4.5.1 Coil Former Base

The first step is to make a CAD model for the coil former. This process has much in common with designing a permanent magnet array former (Sec 3.6.3). In this work, I used both a "1-piece" or "2-piece" approach:

- "1-piece": The wires go directly in the model you design. This paradigm has fewer parts, but the same piece must be structural and hold the wires. Fig 4-11 shows the MR Cap RF coil and the Sparse Halbach gradient coils - two examples of this type of

design.

- "2-piece": The former consists of 2 parts - a set of tiles that actually hold the wires, and a scaffold that holds the tiles in place. This decouples the electromagnetic design considerations from mechanical mating/structural considerations. (Fig 4-12) shows two examples of this approach: the Brain-Optimized MRI RF coil and the MR Cap gradient coils.

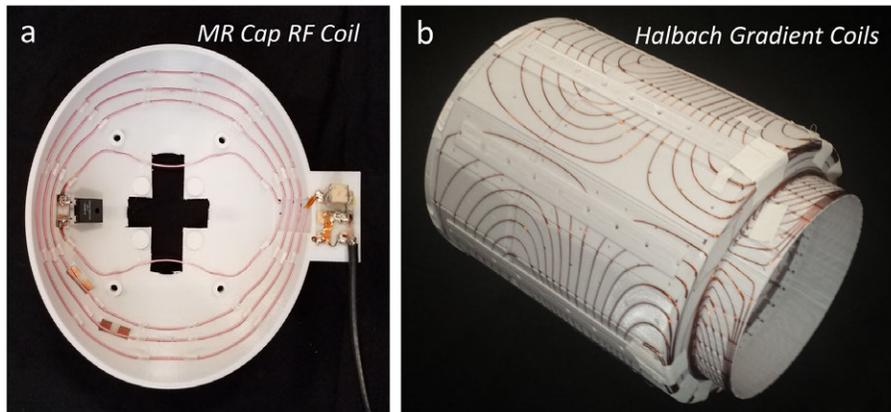


Figure 4-11: a. The MR Cap RF coil was printed in a single piece, which contained both wire winding grooves and mechanical interfacing elements. b. The Sparse Halbach gradient coils were printed in several pieces, which were then assembled into a self-supporting coil former. Each printed piece included both wire winding grooves and holes for mechanical interconnections.

In designing a former, I start with the basic geometry of what is needed based on the volumetric model from the system-level design process (Sec 2.3). Fig 4-13a shows the example of the Helmet Magnet gradient coils. The coil former base is more than thick enough to accommodate all layers of wire windings. Next, I add in any provisions needed for mating with the rest of the system, such as screw holes (Fig 4-13b). If needed, I then apply binary operations with magnet block STLs or other system parts to ensure a non-intersecting mate (see also Sec 3.6.3). This last step is shown for the Helmet Magnet in Fig 4-13c.

Gradient coils generally cannot be printed in one piece (with the exception of planar or almost-planar coils). It is important that the wire grooves be very precise in width so the wire can be press-fit in, which requires putting the grooves in the printer's XY plane (and eliminating slice resolution effects). Additionally, it would be nasty to have to pick support material out of a ton of $\approx 1mm$ grooves, as would be necessary if the grooves were oriented vertically. Printing in-plane allows all precise features to self-support, obviating the need for support material. Splitting the coil into multiple sections for 3D printing accomplishes this.

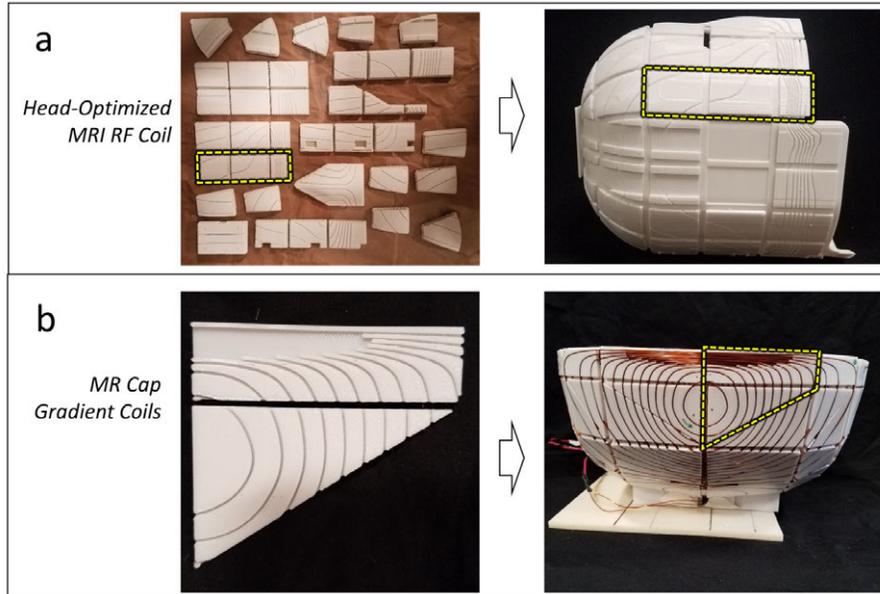


Figure 4-12: a. The Head-Optimized MRI RF coil consisted of several printed "tiles", each of which contained wire winding grooves. These were then bonded to a separate coil former scaffold, which dealt with all mechanical interconnections. b. The MR Cap gradient coils were printed as sets of tiles with winding grooves, which were then glued onto a separate structural former.

This will have already been done if using a "2-step" tile design (eg the MR Cap gradient coils) or a self-supporting split design (eg the Sparse Halbach Gradient Coils). Otherwise, it will need to be done now, as with the Helmet Magnet gradient coils (Fig 4-14a).

4.5.2 Etching the Coil Wire Grooves

The coil former base can now be etched with the imported 3D volume model of the winding grooves (Sec 4.4.4). This is done analogously to the magnet slot etching step for permanent magnet arrays (Sec 3.6.3). The process involves set subtraction of the wire groove model from the coil former base (Fig 4-14d). CAD program binary operation issues during this step, in my experience, mean that your STL has errors (often a self-intersection somewhere). This operation can be extremely memory- and CPU cycle-intensive; be patient.

One also needs to think about how the wires will be wound and will cross from one winding to the next. I find it easier to include a gap in the former design (Fig 4-15) than to modify the 3D winding mesh models. Gaps between coil former sections can also serve as conduits for crossing wires between windings. After this, the different etched sections of the coil can be separated and oriented into the proper configuration for 3D printing.

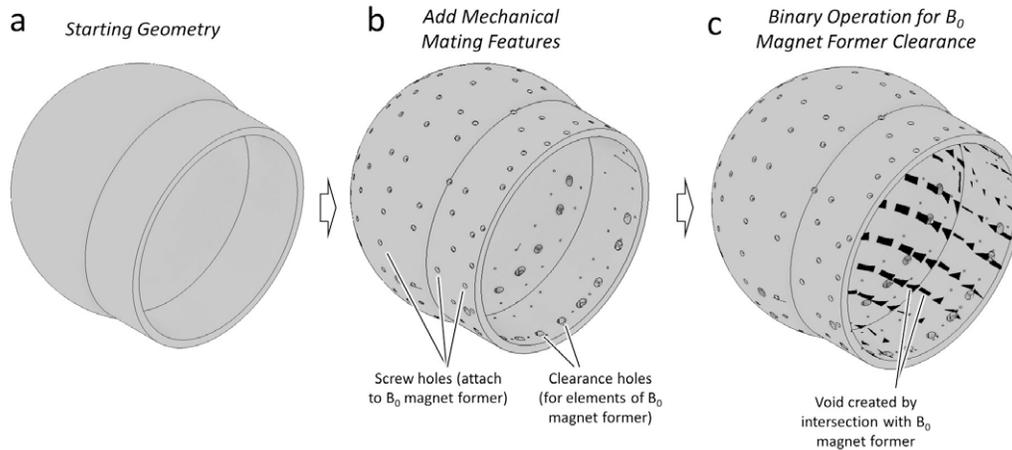


Figure 4-13: a. The Head-Optimized MRI gradient coil former began as a simple parametrized volume that met the system-level geometric constraints. b. To this former template were then added screw holes and other features for mechanical interfacing with the B_0 magnet former. c. A binary operation was then performed to subtract away parts of the gradient coil former that intersected the B_0 magnet former.

4.5.3 Assembling the Coil

The printed coil former sections are first assembled into the proper configuration, and the wires are pressed into the grooves (Fig 4-16). During this process, make sure the winding direction is correct – no one wants their x -gradient coil to come out as a 2nd-order xz coil. Also, make sure your copper wire is long enough for the whole coil. Use a blunt tool (eg a sanded down piece of G10 PCB stock as in Fig 4-17) to press the wires all the way into the grooves. Be careful not to tear into the insulation on the magnet wire: discovering a short after winding an entire coil would not be fun. Tape can be used while winding to keep troublesome wires in place, such as wires that aren't staying in their grooves (Fig 4-17).

Once the coil is wound, check its inductance and resistance on an LCR meter. Knowing these numbers helps with GPA compensation and RF coil tuning (Sec 4.5.4). Also, a significant mismatch between the simulated and measured L or R is an indication of a short somewhere in the coil.

After finishing the coil, it may be necessary to epoxy down some of the windings. This is often needed near the truncation end of a gradient coil where wires are bunched together and not held nicely in grooves. In this work, Loctite EA608 was used to hold down windings on the Sparse Halbach and Helmet Magnet gradient coils.

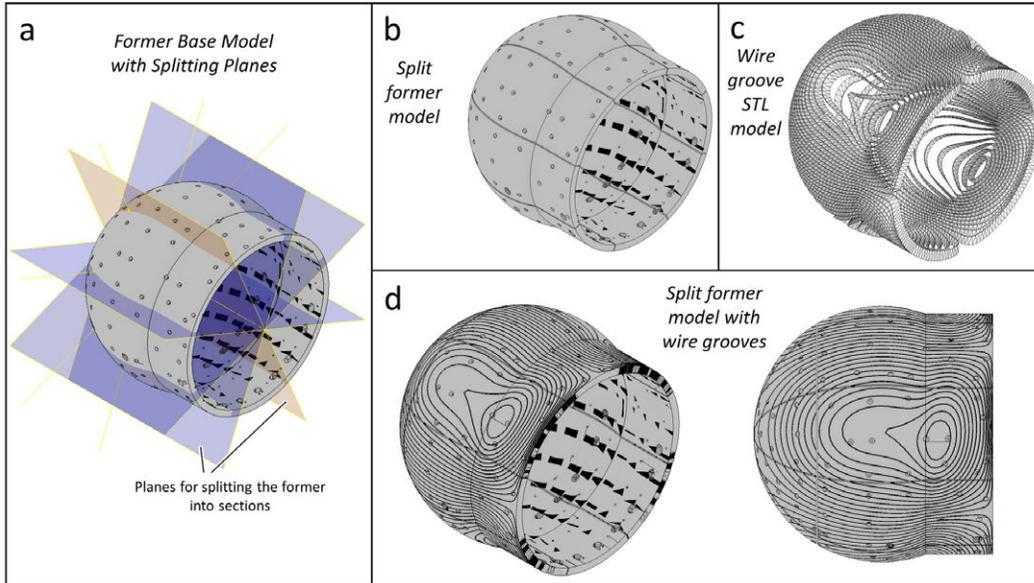


Figure 4-14: a. The gradient coil former is split into pieces in preparation for 3D printing. b. The split gradient coil former. c. The wire groove STL model for the Head-Optimized Gy coil. d. The split former model and the wire groove model are combined using the "set difference" binary operation to create a coil former model with the numerically-derived wire winding grooves.

4.5.4 Coil and Amplifier Tuning

Gradient coils are usually operated in current control (CC) mode, which requires properly compensating the GPA for the reactive gradient coil load. RF coils need to be tuned and matched to the correct frequency. Commercial amplifiers may come with a guide to compensation for a user-provided load and bandwidth.

RF coils are tuned by splitting the coil winding at a number of tuning points (between 1 for the MR Cap coil to 4 for the Helmet Magnet RF coil) and inserting a tuning capacitor. The number of tuning points chosen affects the peak capacitor voltage in the coil (see Sec 4.4.1). More tuning capacitors tends to increase the required capacitance and decrease the peak voltage. When tuning a large coil with multiple windings, one will see many resonant modes of the coil on the VNA. These generally have different spatial field profiles. When tuning the coil to the Larmor frequency, care must be taken to tune the correct spatial mode. This can be tested by using a sniffer loop and measuring the S_{21} through that loop at different points in the coil. For example, in the common case of a mode with uniform B_1 , there should be no spatial nodes where the field drops to zero and the phase changes across it. This can be verified using the sniffer probe.

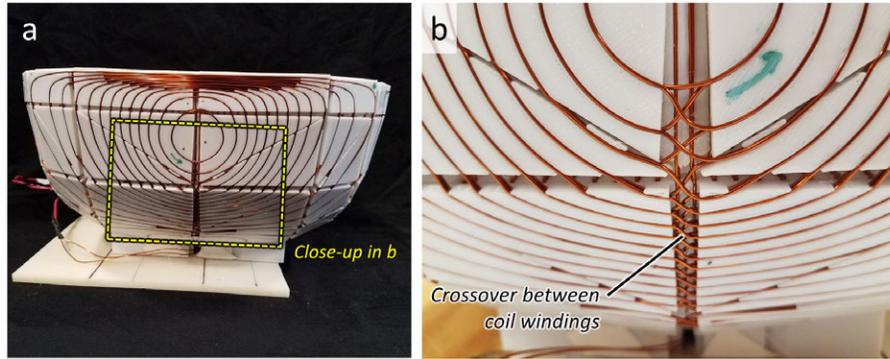


Figure 4-15: a. The MR Cap gradient coils (Gz coil visible). b. Close-up of the Gz coil showing the wire crossing from one winding to the next in a gap between the gradient former tiles.

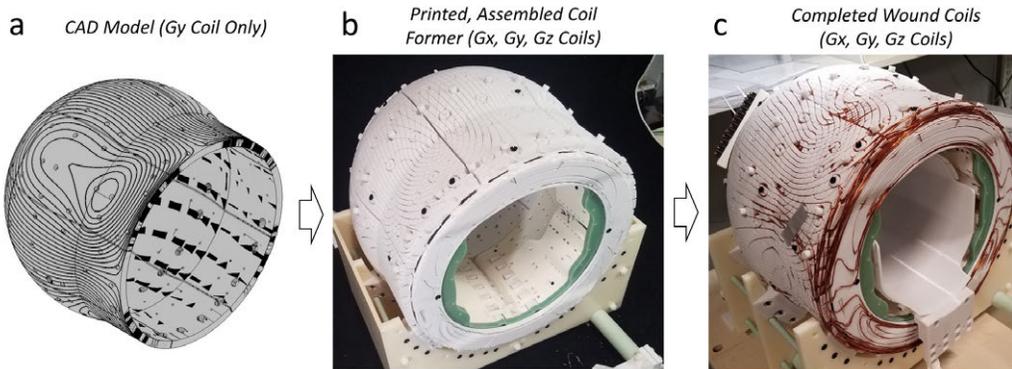


Figure 4-16: a. The CAD model of the Head-Optimized Gy coil. b. 3D-printed and assembled gradient coil former for the Head-Optimized Gx, Gy, and Gz coils. c. The completed, wound Head-Optimized Gx, Gy, and Gz coils.

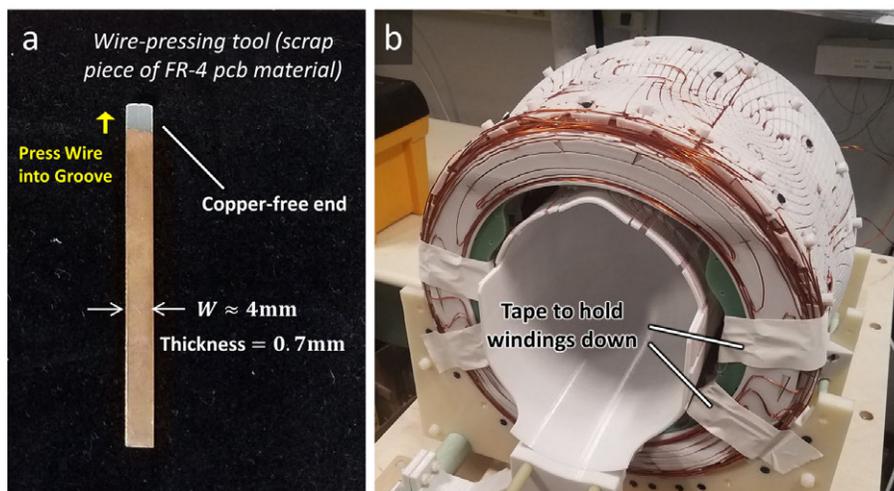


Figure 4-17: a. A precision-engineered wire-pressing tool was created from a sanded-down piece of scrap FR-4 PCB substrate. b. Tape was sometimes used to hold down gradient coil windings before the final epoxying step.

Chapter 5

Gradient and Shimming Hardware for a Portable Brain Scanner

5.1 Abstract

In this section, I describe the design and construction of gradient coils and a permanent magnet shim array to improve the imaging capabilities of a previously-published portable brain MRI scanner [165]. This scanner uses a permanent magnet-based B_0 magnet, has a mean field of 80 mT, and contains a readout encoding gradient built into the B_0 field. I designed and demonstrated a pair of gradient coils for partition and in-plane phase encoding along the other two spatial axis. This enabled Fourier-type encoding schemes and reduced image distortion. I then designed and constructed a permanent magnet shim set to improve the linearity of the built-in B_0 encoding field. This work has lead to *in vivo* imaging with this portable brain scanner.

5.2 Introduction

This work builds on an existing portable, head-only scanner with a permanent readout encoding gradient built into the B_0 field [165]. The magnet for this system was designed to use the inhomogeneous B_0 field as a built-in encoding field. To acquire an image, the magnet would be mechanically rotated in steps; a spin-echo acquisition would be acquired at each step (using the B_0 field as the readout-encoding field); and these "generalized projections" would be reconstructed as an image using a generalized (non-Fourier) reconstruction

scheme [119]. The combination of rotation + generalized projection provides two dimensions of spatial encoding. The 2D imaging capabilities of this approach had previously been demonstrated [165].

This system still required a way to perform partition encoding along the third spatial axis. In addition, the generalized projection/reconstruction approach was found to result in artifacts when imaging a large, head-sized object. Overcoming this would require improvements to the encoding scheme, the magnet (B_0 spatial field), and/or the reconstruction.

In this work, I designed hardware to enable a new imaging approach with this system. In this approach, the B_0 magnet would not need to be mechanically rotated - its inhomogeneous field would instead serve as a fixed readout encoding field along one spatial axis. I designed a pair of gradient coils for blipped phase encoding along the other two spatial axes, giving 3D-imaging capabilities. I also designed a set of shim trays to improve the encoding properties of the built-in encoding field. This work has ultimately enabled the acquisition of *in vivo* brain images using this scanner.

5.3 Gradient Coils

A pair of gradient coils was designed for two dimensions of phase encoding, with the magnet itself providing readout encoding along the third dimension. The coils were designed using the BEM target field approach (Sec 4.4) and were built by pressing wire into a 3D-printed former (Sec 4.5). The coils were field-mapped and found to be adequate for 3D imaging.

5.3.1 Motivation and Introduction

Phase encoding with gradient coils a good choice for imaging in a portable scanner with an inhomogeneous B_0 field. Such gradient coils can have:

- Reasonable power consumption
- Reasonable weight
- Ease of design and construction compared with RF encoding approaches, such as TRASE [176]
- The ability to work in inhomogeneous fields using balanced phase encode gradient moments

- Unshielded designs, given that eddy currents in the magnet are not a deal-breaker as they are with superconducting systems
- The property of enabling Fourier-type (k-space-ish) encoding

The last point in particular provides robustness to absolute field offsets and shifts in a way that projection-type imaging does not. Artifacts in Fourier imaging are also well-understood and generally localised to regions with non-bijective or flat encoding fields. Distortions are similarly localized and easy to diagnose. Gradient coils therefore enable new ways of acquiring images in this portable scanner, while not precluding the use of rotating-magnet or generalized-projection schemes [119].

5.3.2 System-Level and Gradient Coil Geometry Considerations

The gradient coils were designed to fit inside of an existing magnet [165], as shown in Fig 5-1. The first design task was to choose a design surface and an allowable amount of bore thickness for the coils to occupy. This was done by first making test G_x coil designs on different surfaces and looking at their expected encoding performance. In this geometry, the G_x coil is the hardest to make, as it must generate its maximum field at a point along its center axis near the truncated end (Fig 5-1). However, the coil physically ends at this point and is thus naturally not very field-efficient there.

The test coils were designed on two surfaces:

1. A cylinder that fit along the boundary of region A in Fig 5-1
2. The same cylinder with an extra "top-hat" surface that fit in region B (Fig 5-1)

Fig 5-2 shows the two surfaces and designed coils. The design optimization (described in Sec 5.3.3) used target fields which were linear over a $D = 21cm$ sphere. Fig 5-3 shows the field produced by the two coils with $I = 1A$ along two different axes. Labeled are the mean gradient efficiency (in $\frac{mT}{m \cdot A}$) and the extent over which the field is bijective. The latter metric is used as a measure of encoding performance. The cylinder + annulus/entry cylinder coil had a higher efficiency and a larger bijective region than the cylinder coil. That surface was chosen as the G_x design surface, even though it was more complex to build and required space in the already-tight "B" region of the magnet (Fig 5-1d).

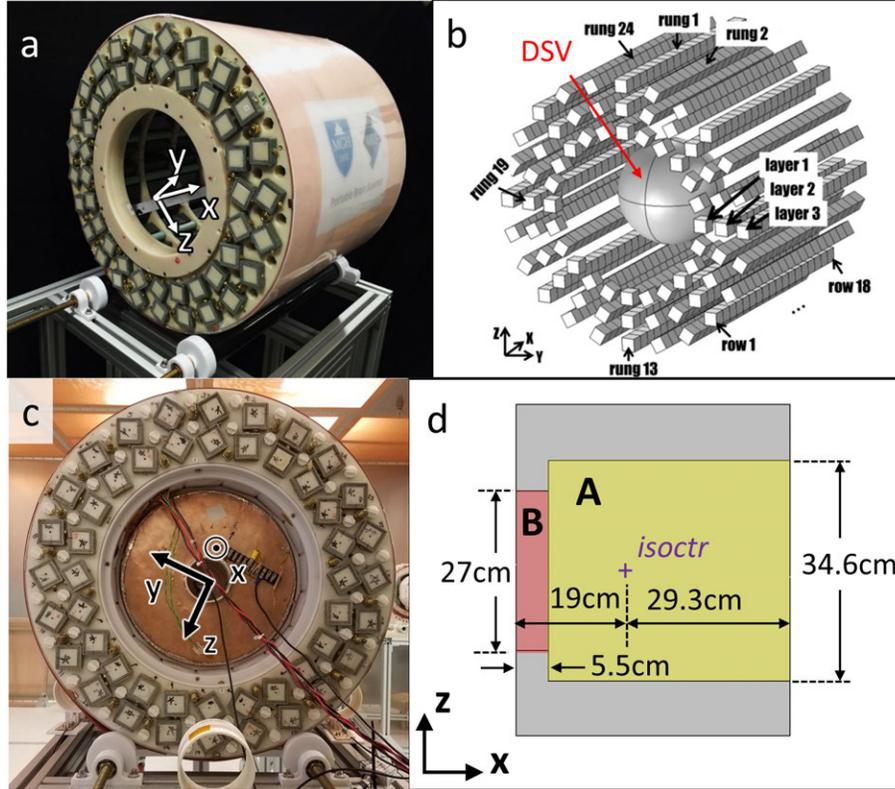


Figure 5-1: a. The constructed optimized Halbach magnet from [165]. Shown is the patient end (ie the end that the patient's head goes in). b. The magnet contained 24 rungs of magnets in 3 concentric layers and was optimized for performance in a $D=20\text{cm}$ DSV. c. Back view of the magnet showing dimeters of different parts of the bore. d. Cutaway (XY plane) view of the magnet showing the internal bore dimensions. Labeled are (A) the large-diameter main bore and (B) the narrower "entry" region of the bore.

A similar analysis found that using the "top-hat" region for the G_z coil did not justify the increased bore space required. A "cylinder+annulus" design (Fig 5-4) was thus used for G_z coil design.

The back ($-\hat{x}$) end of the magnet bore had an opening with the full large cylinder diameter. The gradient coils were inserted in this opening, and gradient feed lines entered through it.

The gradient coils were also designed for an RF coil and RF shield to slide inside. This was incorporated in the gradient coil design as a trio of 3D-printed scaffolding rings with built-in "slots" (Fig 5-5). The RF shield had "rails" attached to it that allowed it to slide in the gradient coil along these "slots" (Fig 5-5c). The G_x and G_z coils were both constructed on the same former using 2 layers of AWG18 wire windings. The total thickness of the gradient coil set was 7mm in region A (Fig 5-1) where both coils were wound, and 3.5mm

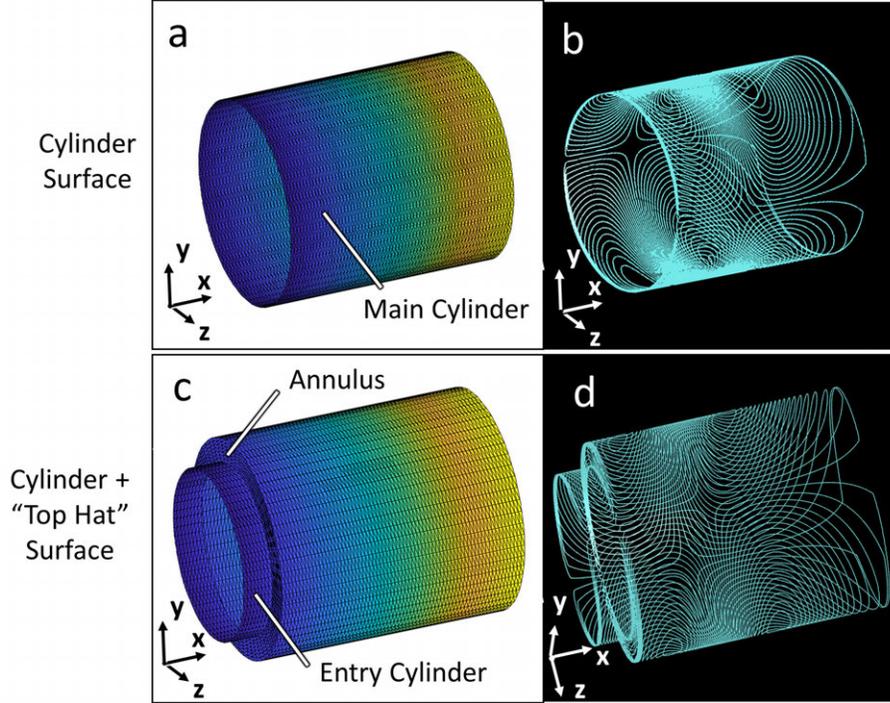


Figure 5-2: Two G_x coil designs were created for comparison on two different surfaces. a. The cylindrical surface. b. The cylindrical surface coil design. c. The cylindrical + “top hat” surface. d. The cylindrical + top hat surface coil design.

along region B, where only the G_x coil was wound. This was an acceptable amount of bore space to take up.

5.3.3 Coil Design

The coil design surfaces were piecewise cylinder/annulus structures and were easily parametrized. The meshes for BEM field computation were generated directly from this parametrization (see Sec 4.4.2). The coils were then designed using the target field stream function optimization procedure described in Sec 4.4, which incorporated the numerical optimization procedure in [172]. The target ROI was a $D = 21\text{cm}$ sphere for both the G_x and G_z coils. The field target was for the \hat{z} -aligned field component to vary linearly along x or z for the G_x or G_z coil, respectively. The \hat{y} - and \hat{x} -aligned components did not factor in. The optimization used L^2 regularization with $\lambda = 4 \cdot 10^{-6}$ (see Sec 4.4.3).

The optimized stream functions were projected and up-sampled onto a planar grid. The projected stream function was smoothed and where windings were computed from $N = 30$ level sets (Sec 4.4.4). The smoothing parameter was adjusted so that windings were spaced

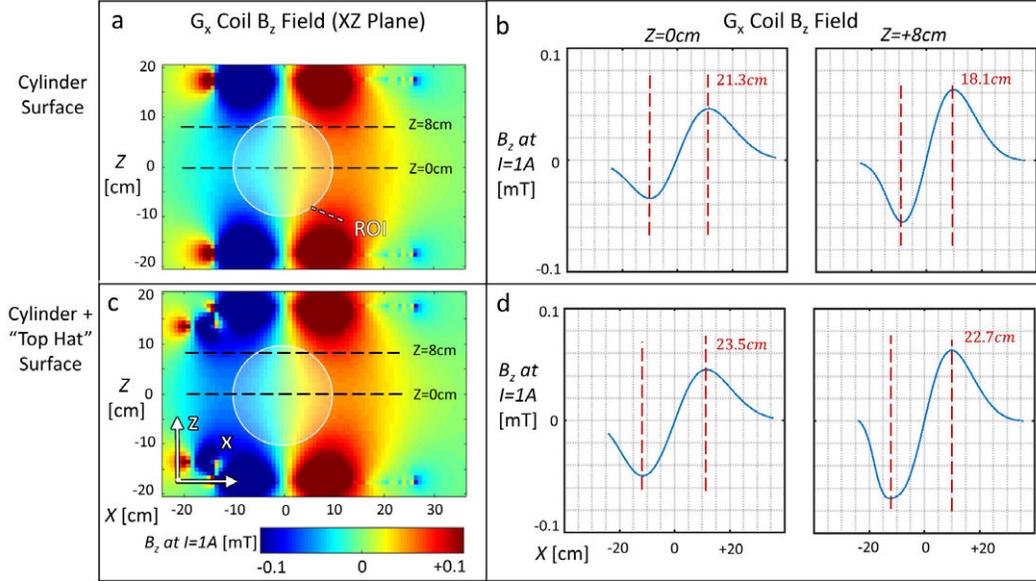


Figure 5-3: a. Simulated field map at $I = 1A$ for the cylindrical surface design. b. Cylindrical surface coil field plotted along $Z=0$ and $Z=8cm$ showing bijective range. c. Simulated field map at $I = 1A$ for the cylindrical + top hat coil. d. Field plots along $Z=0$ and $Z=8cm$ showing bijective range.

by at least the minimum spacing (see Sec 4.4.1). The optimized coils are shown in Fig 5-6 along with simulated field maps in Fig 5-7. The simulations show mean efficiencies of $0.577 \frac{mT}{m \cdot A}$ and $1.103 \frac{mT}{m \cdot A}$ for the G_x and G_z coils.

These windings were unfolded onto the constituent subsurfaces of the piecewise-smooth design surface. Fig 5-8 shows the unfolding of the G_x coil. The "top-hat" cylinder section has been unfolded into a planar surface; the annular windings have been separated; and the main cylinder windings have been unfolded.

The coil former base consisted of a set of 16 planar panels for the main cylinder section, a monolithic annular ring, and a set of 32 panels for the top-hat cylinder section (Fig 5-9). An additional structural ring mounted to the $+\hat{x}$ end of the 16 main cylinder panels. The pieces were designed to be connected with screws and/or epoxy (Loctite EA608). Next, the CAD models of the former base pieces were etched with the G_x and G_z wire grooves as described in Sec 4.5.2. The G_x coil fit was etched on the outside of the former; the G_z coil was etched on the inside. The etched coil former pieces were finally printed in PC-10 material on a Fortus 360mc printer.

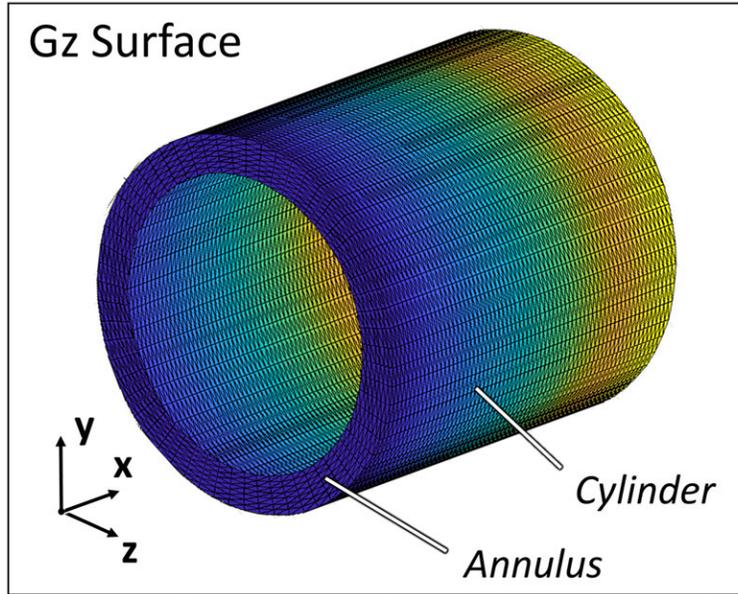


Figure 5-4: The G_z coil design surface consisted of a cylinder joined with an annulus.

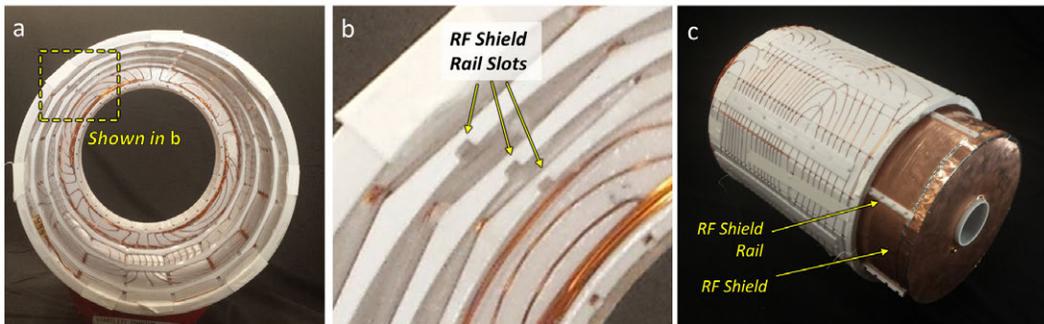


Figure 5-5: a. The gradient coil former incorporated several slots for inserting an RF shield. b. Close-up view of the RF shield rail slots. c. The RF shield used rails that allowed it to slide cleanly into the gradient coils.

5.3.4 Coil Construction and Testing

The coil former pieces were next assembled according to the design. The coils were then wound with AWG18 wire pressed into the grooves and subsequently glued down with Loctite EA608 epoxy. The G_x and G_z coils used $L = 87.2m$ and $L = 57.1m$ of wire, respectively.

It was then verified that the coil actually fit in the magnet. The first coil fit in most of the way, but could not be removed without destructive methods as the fit tolerance was not loose enough (Fig 5-10). This was remedied by making a second coil pair with all linear dimensions scaled down by 1%. (This operation was performed directly on the 3D printer STL files.) This was assumed to not affect wire groove or screw clearance hole tolerances

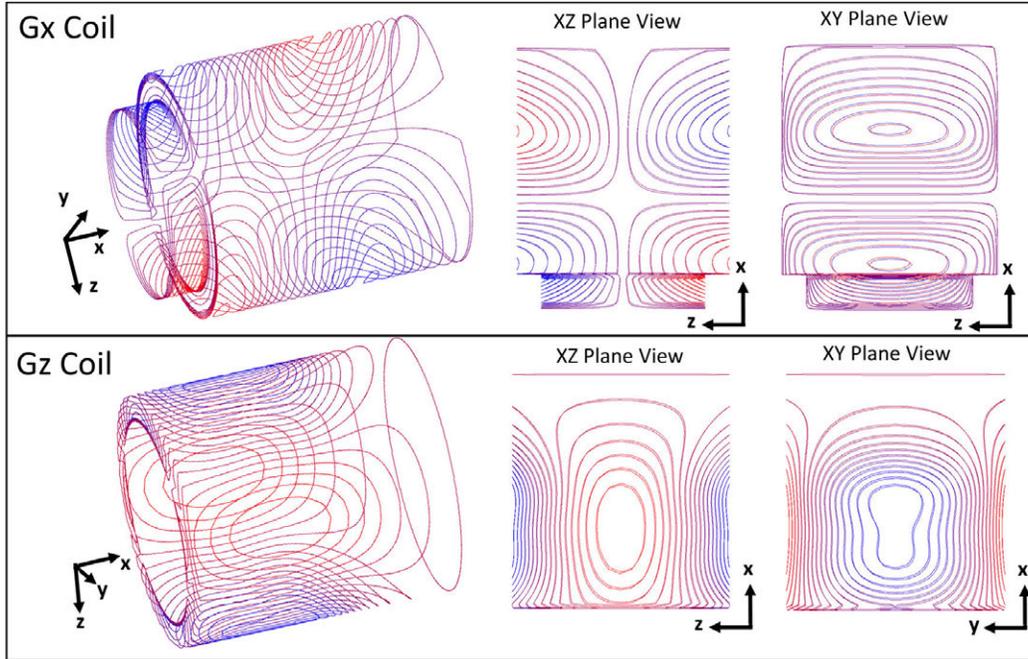


Figure 5-6: Winding patterns for the G_x and G_z coils obtained from the optimized stream functions.

(eg a 1% change was less than the precision of the wire tolerance tests in Fig 4-5). This second coil fit in the magnet.

Table 5.1 shows the measured resistance and inductance for each coil. These values enabled using a commercial GPA (Techron 7224) in current-mode operation with the built-in compensation network and $I_{max} > 10A$. The coils were next quickly field-mapped by scanning a $D = 24cm$ spherical phantom in a $1.5T$ superconducting MRI system (Fig 5-11). The phantom was placed inside the coils, which were energized with $I = \pm 100mA$. These experiments verified that the coils had approximately the expected spatial patterns and field efficiencies.

The coils were then mapped with greater accuracy and over a larger ROI using a 3-axis Hall probe (THM1176, Metrolab, Geneva, Switzerland) stepped by a 3-axis robot. Fig 5-12 shows gradient field maps at $I = 1A$ measured with the Hall probe. The mean efficiencies were $0.575 \frac{mT}{m \cdot A}$ and $0.815 \frac{mT}{m \cdot A}$ for the G_x and G_z coils.

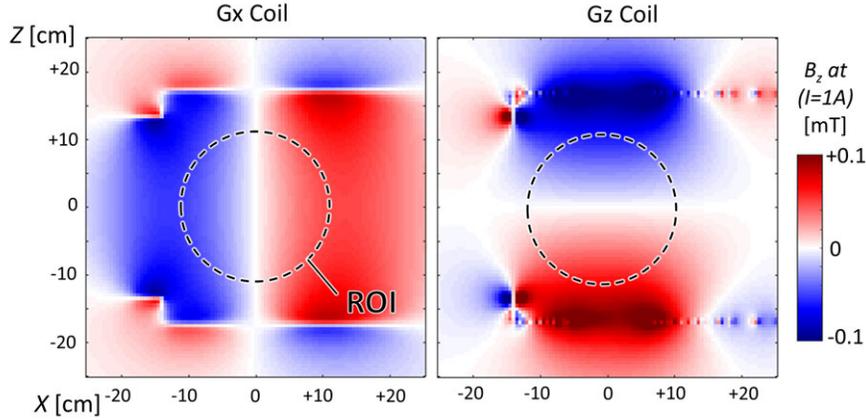


Figure 5-7: Simulated field maps for the Gx and Gz coils.

Table 5.1: Gx and Gz coil parameters and performance metrics

	Gx Coil	Gz Coil
<i>Wire Length</i>	87.2 m	57.1 m
<i>Inductance</i>	510 μ H	335 μ H
<i>Resistance</i>	1.85 Ω	1.21 Ω
<i>Isocenter Efficiency</i>	0.575 $\frac{mT}{m \cdot A}$	0.815 $\frac{mT}{m \cdot A}$

5.4 Target Field B_0 Shimming

5.4.1 System-Level and Geometric Considerations

An array of shim magnets was designed to fit inside shim tray slots that had been included in the initial construction of the B_0 magnet (Fig 5-13a). The magnet contained $N = 48$ shim tray slots arranged in two rings at $R_1 = cm$ and $R_2 = cm$. The slots had length $L = 48.3cm$, extended 19cm towards the patient end from isocenter, and had octagonal cross-section with incircle diameter $D = 17.78mm$ (Fig 5-13b). This volume specified the maximum extent of magnetic material in the shim trays. These slots were used primarily for convenience, as placing the shim trays at approximately the same radii as the main B_0 magnets was not expected to be optimal from a magnet performance perspective (see Sec 2.3.3).

Each shim tray was discretized into $\frac{3}{8}'' = 9.525mm$ -deep cells with a 2mm gap in between. Each cell represented a single magnet in the design. This resulted in $N = 42$ magnets in each shim tray for a total of 2016 in the whole array. Each magnet was allowed to vary

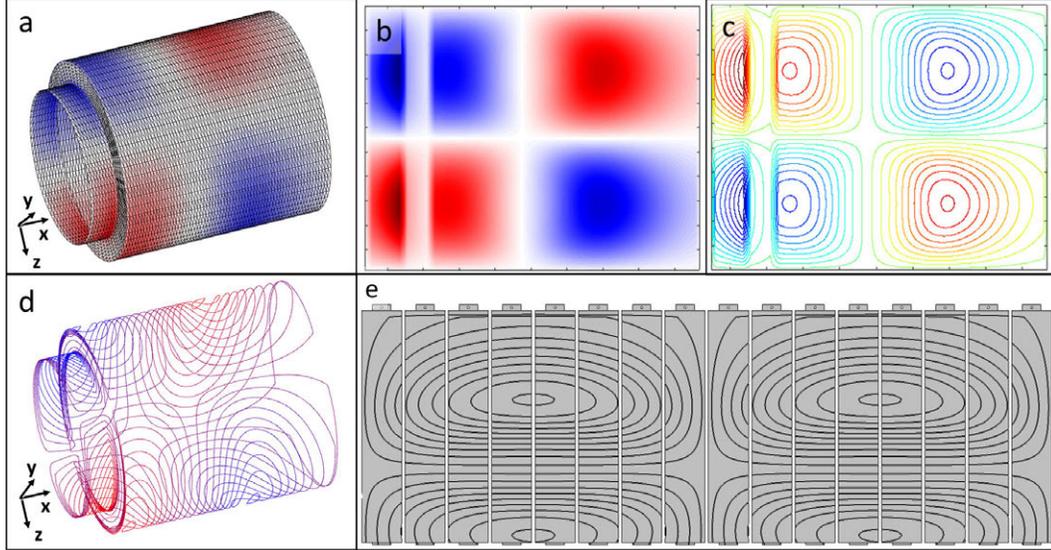


Figure 5-8: a. The optimization produces a stream function on the design surface. b. The stream function is projected onto a flat 2D surface. c. The flattened stream function is then discretized into a set of level sets/isocontours. d. The level sets are projected back onto the original design surface for simulation. e. The flattened 2D isocontours are used to generate the flattened CAD models for the 3D-printed coil pieces.

in size and to rotate about the longitudinal axis of the shim tray. This gave 2 DoFs for each magnet ($N = 4032$ DoFs total).

Individual shim magnets were constructed out of between 1 and 5 stock magnet blocks stuck together. The maximum-size magnet that could fit within one cell contained: one $\frac{1}{2}'' \times \frac{3}{8}'' \times \frac{3}{8}''$ N42 magnet, two $\frac{3}{8}'' \times \frac{3}{8}'' \times \frac{1}{16}''$ N42 magnets, and two $\frac{1}{4}'' \times \frac{3}{8}'' \times \frac{1}{16}''$ N42 magnets (Fig 5-14) and had a dipole moment of $\mu_{max} = 1.662A \cdot m^2$. Table 5.2 lists the 27 magnet configurations were used, ranging from an empty slot ($\mu_m = 0$) to this max value. All configurations were symmetric so that they contained no $\ell = 2$ multipole terms (Sec 3.4.5). The shim optimization assumed each component of each magnetization vector could vary continuously, even though there were 27 discrete block options. The resulting continuous-space design was discretized as a post-processing step.

The shim magnets were modeled as point dipole sources, as they were small ($\leq 13mm$ in size), far from the target ROI ($\geq 10cm$ away, and contained no $L = 2$ multipole terms).

The choice of 2 degrees of freedom per shim was compared with the hypothetical 3-DoF case where each magnet could vary in size and rotate with 2 spherical angles (Fig 5-15a). The 3-DoF case would be impractical to build, but was used as a best-case scenario for

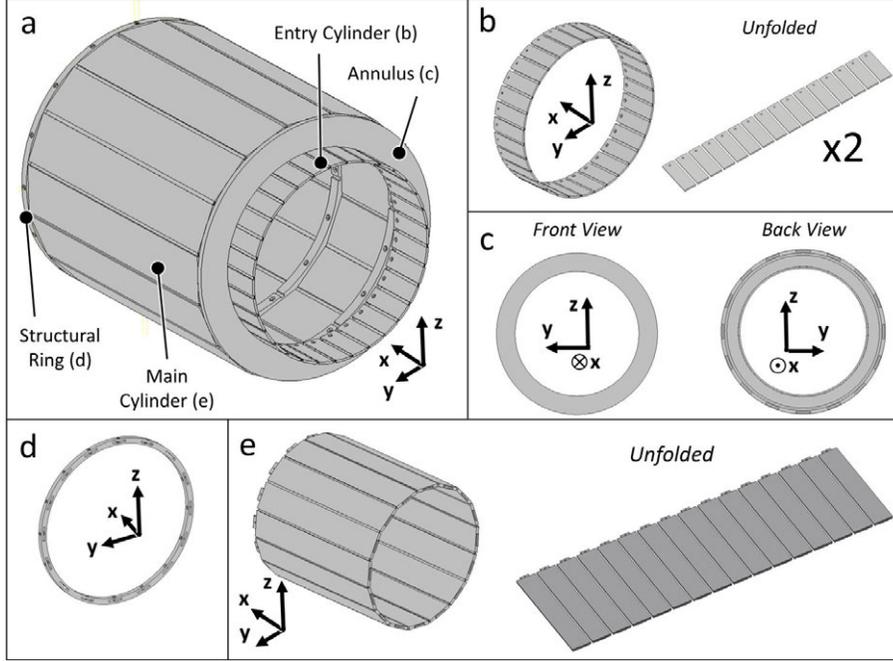


Figure 5-9: a. The combined G_x/G_z coil former base consisted of: b. 32 rectangular panels for the bore opening section; c. A one-piece annular section; d. A structural ring was also used at the back of the former; and e. 16 rectangular panels that approximated the main cylinder section. Sections were screwed and/or bonded together.

comparison. The shim configurations were optimized for both cases using the optimization procedure described here. The same number of magnets with the same max magnetization constraint was used for each optimization. Fig 5-15 shows field map residuals for the two cases. The SSE was $0.66 \cdot 10^{-2}$ for the 3-DoF case and $0.80 \cdot 10^{-2}$ for the 2-DoF case.

5.4.2 Shim Array Design

The shim optimization used a target-field optimization with an L^2 -norm cost function:

$$f_{cost} = \|D_z M - B_{z,targ}\|_2^2 \quad (5.1)$$

Here, D_z is the field computation matrix for the \hat{z} -component assuming the magnetic dipole model; M is the vector of magnetic dipole moment vector components for all shim blocks; and $B_{z,targ}$ is the field target. Fig 5-16 shows the field target $B_{z,targ}$, which is the difference between the measured B_0 map and the ideal linear encoding field. The field was optimized over a $D = 20cm$ sphere. No symmetry was imposed on allowable shim array designs, as the target field was not symmetric. The cost function was computed over a $D = 20cm$ sphere.



Figure 5-10: Coils that do not fit in the magnet may require drastic measures for removal.

The magnitude of each shim block's magnetic dipole moment was constrained to be $\leq \mu_{max}$ (Sec 5.4). The optimization was performed using an interior-point optimization with the Matlab `fmincon` tool.

The initial guess was supplied by the solution to the "softened" L2-regularized problem:

$$\{M_i\} = \underset{\{M\}}{\operatorname{argmin}} \left(\|D_z M - B_{z,targ}\|_2^2 + \lambda \|M\|_2^2 \right) \quad (5.2)$$

with the regularization parameter λ tuned such that the solution just met the hard constraints on max magnet size.

Fig 5-17 shows the optimized shim design as dipole moment vectors and as a set of magnet blocks. Each block in this optimized solution was rounded to the closest of the 27 allowable magnet configurations.



Figure 5-11: A coil can be quickly field-mapped using a large phantom and a high-field scanner. Using a field mapping sequence with low readout bandwidth reduces coupling to the gradient coil under test, but increases the distortion of the acquired field map.

5.4.3 Shim Tray Construction and Testing

Process tests were performed to figure out what printer, material, tolerance, and part orientation to use (see also Sec 3.6.2). It was found that vertically-oriented shim trays with dissolvable support material in ABS on Stratasys Dimension printer worked best. With this process, the angular position of the block was specified by high-precision XY-plane printer head movements. Loctite 414 cyanoacrylate adhesive worked well for bonding magnets in slots. Its 5-second work time was also a practical benefit, as $N = 2016$ magnet configurations would need to be glued in sequence, and using even 5-minute epoxy would have taken a very, very long time.

Next, STL models for shim tray slots were generated for CAD modeling of the shim trays themselves (see Sec 3.6.3). Fig 5-18 shows the STL model for one set of shim tray slots. The shim tray formers were designed to be printed in two parts in order to fit in the printer build volume. Identification numbers 01 through 48 were printed in the ends of the shim trays so they could be distinguished. The models were then etched with the shim tray slot STL models (see Sec 3.6.3). All shim tray halves were then printed.

The 2016 shim magnet configurations were then preassembled from $N = 6536$ individual stock magnets (suppliers: K+J magnets; Applied Magnets). They were then glued into the shim trays (Fig 5-18b). Once finished, the shim trays were inserted into the magnet. This was done carefully by supporting both ends of the shim tray as it was slid in (lest the

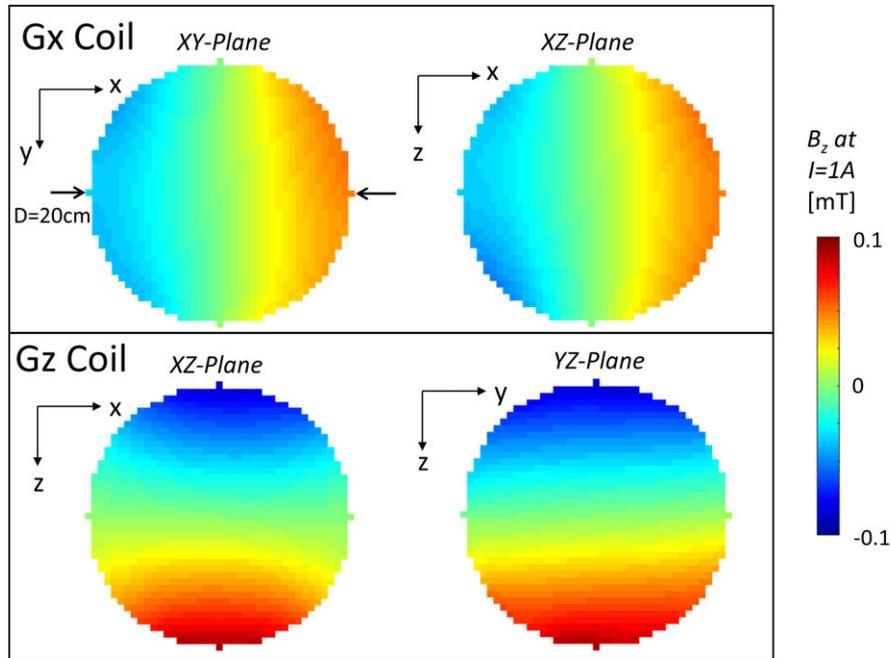


Figure 5-12: Gx and Gz coil field maps as measured with a 3-axis Hall probe stepped by a 3-axis positioning robot. Measurements were taken on the surface of a D=20cm sphere and fit to at 12th-order spherical harmonic expansion.

shim tray break inside the magnet and necessitate surgical removal – we learned our lesson). Fig 5-19 shows the magnet with all 48 shim trays inserted. Also illustrated is the angular misalignment of the shim trays due to the fit tolerance and the torques exerted by the main B_0 magnets.

The magnet was then re-mapped with the 3-axis Hall probe. The measured shimmed field is shown in Fig 5-20 along with the unshimmed B_0 field. The shimmed field looks more linear, and the residual non-linear RMSE was reduced by over 50%.

5.5 Imaging Experiments

5.5.1 Sequences and Experimental Protocols

All experiments used RARE-type sequences (Fig 5-21) with broadband swept WURST pulses [177] to cover the whole spin bandwidth. Sequence parameters for different experiments are listed in Table 5.3. Partition encoding was performed down the spin echo train; in-plane encoding was performed shot-to-shot. Imaging experiments all used:

- A spiral volume helmet coil (Fig 5-22)

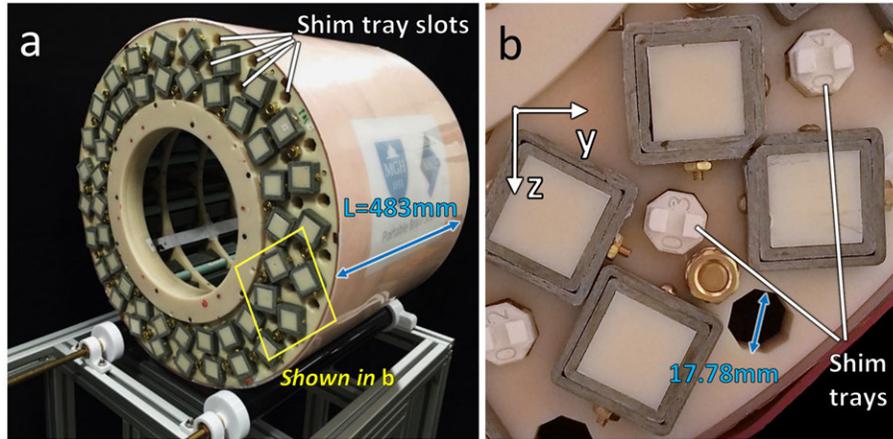


Figure 5-13: a. The magnet contained 48 shim tray slots distributed in between the main B_0 magnet rungs. b. Each shim tray slot had an octagonal cross-section with inner dimension 17.78mm.

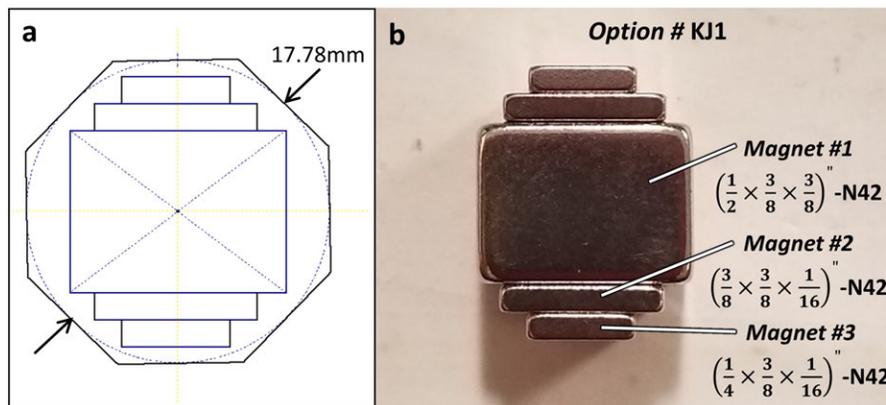


Figure 5-14: a. The maximum-sized magnet that could fit in a shim location was found by stacking stock magnet sizes inside the incircle of the octagonal shim tray cross-section. b. The maximal magnet contained: one $\frac{1}{2}J \times \frac{3}{8}J \times \frac{3}{8}J$ N42 block; two $\frac{3}{8}J \times \frac{3}{8}J \times \frac{1}{16}J$ N42 blocks; and two $\frac{1}{8}J \times \frac{3}{8}J \times \frac{1}{16}J$ N42 blocks.

- A passive lumped-element T/R switch
- A MITEQ AU-1583 preamp
- A Tecmag Apollo (phantom experiments) or Bluestone console (*in vivo* experiments) for waveform generation and signal detection
- A Tomco model # BTO2000-AlphaS-3MHz RFPA
- Techron 7224 GPAs for the Gx and Gz gradient coils

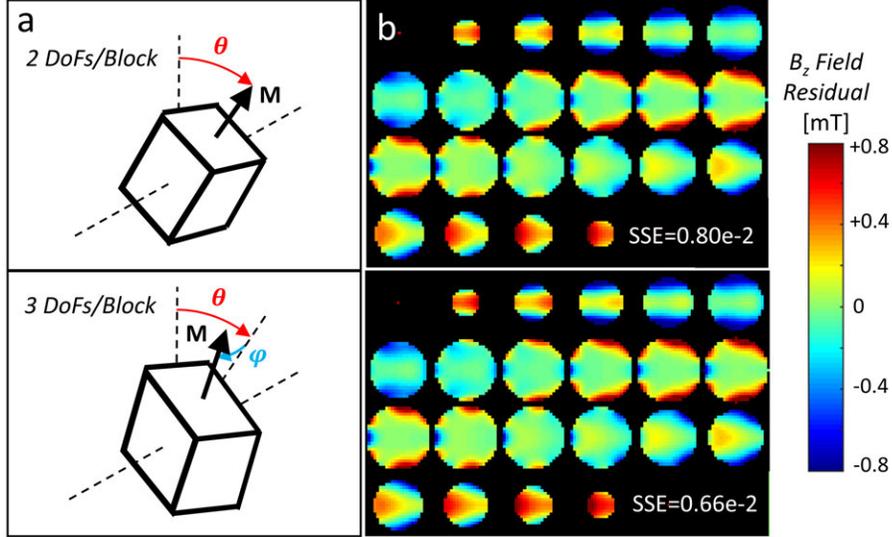


Figure 5-15: a. Shim arrays were designed allowing for in-plane “circular” magnet rotation and 2-angle “spherical” magnet rotation. b. Simulated residual error maps for the “circular” and “spherical” optimized designs. The summed squared error (SSE) was $0.8 \cdot 10^{-2}$ for the “circular” design and $0.663 \cdot 10^{-2}$ for the “spherical” case.

Phantom Imaging

Phantom experiments were performed using an anthropomorphic "red head" phantom filled with agarose gel and spheres containing solutions with different T_1 and T_2 relaxation times. Acquisitions used the built-in G_y encoding field as a readout gradient and the switched G_x and G_z coils for partition and in-plane phase encoding. Images were acquired in the unshimmed magnet and the shimmed magnet, and all images were acquired in an unshielded environment.

in vivo Imaging

in vivo images were acquired of a healthy female adult subject. These *in vivo* acquisitions were performed on the shimmed and unshimmed magnet in an unshielded environment as a preliminary assessment of image distortion.

5.5.2 Reconstruction Methods

Reconstructions were performed using Fourier reconstruction and using a generalized projection reconstruction approach [119]. The generalized projection reconstruction approach incorporates the nonlinear readout and in-plane phase encoding fields in a forward encoding

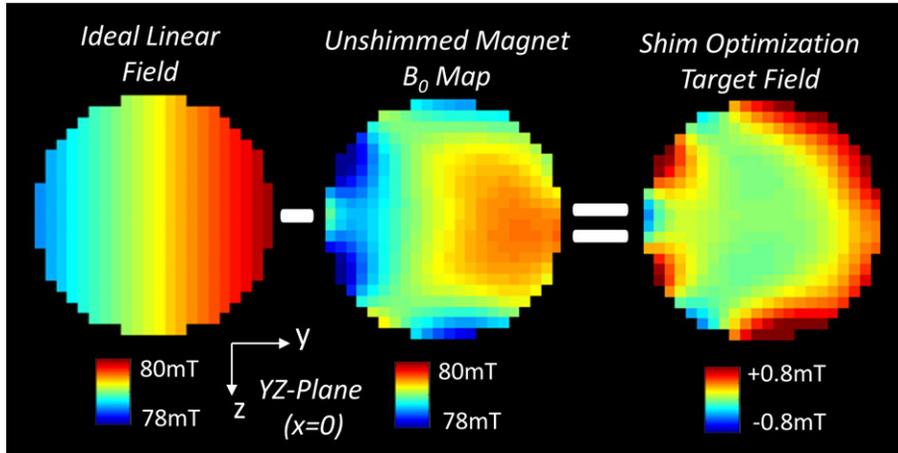


Figure 5-16: The shim optimization target was the different between the ideal linear encoding field and the measured magnet B_0 field.

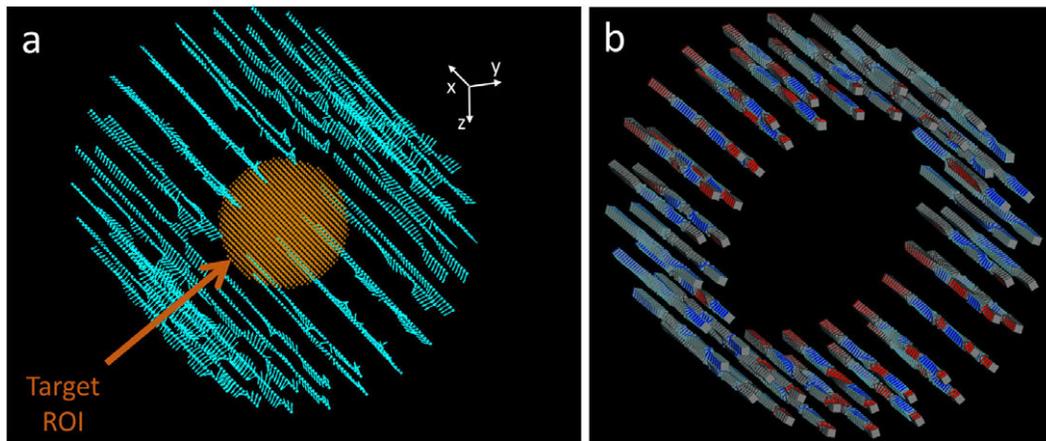


Figure 5-17: a. Shim optimization solution show as a vector plot (each vector represents a shim block magnetic dipole moment); the design ROI is also shown. b. Shim optimization solution shown as a block model. Blue and red faces signify the south and north poles of a magnet, respectively.

model. This encoding model better approximates this scanner and acquisition than ideal Fourier encoding.

5.5.3 3D Imaging Results

Phantom Experiments

Fig 5-23 shows the red head images acquired using: the unshimmed magnet and Fourier reconstruction; the shimmed magnet and Fourier reconstruction; and the shimmed magnet and generalized reconstruction. The Fourier reconstructed, shimmed magnet images have less distortion than the unshimmed magnet images. The shimmed magnet generalized re-

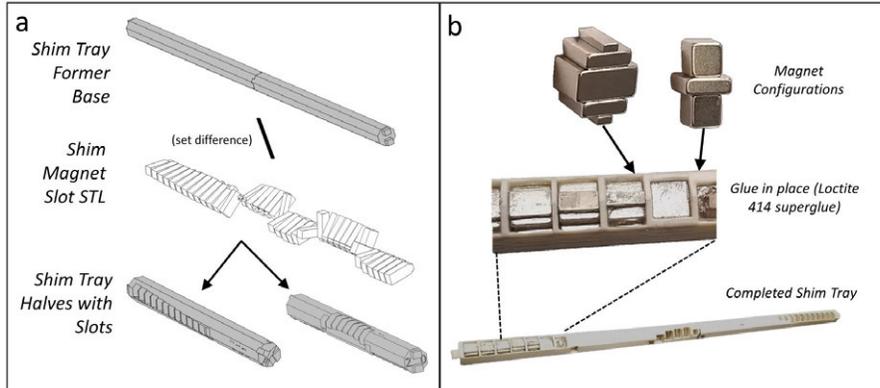


Figure 5-18: a. Shim tray CAD design used automatically-generated STLs for the shim magnet slots and 2-piece shim tray former base. b. Shim magnet configurations were pressed and glued into the 3D-printed shim tray formers.

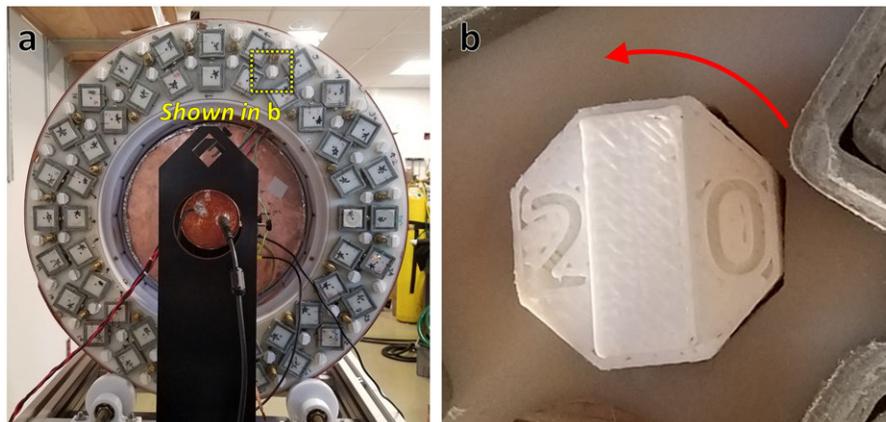


Figure 5-19: a. All 48 shim trays were inserted in the back end of the magnet. b. Some shim trays were angularly misaligned due to the mechanical tolerance and the torques exerted by the main B_0 magnets.

construction images have less distortion than the shimmed magnet Fourier reconstruction images. All images nonetheless contain distortions and hyperintense regions.

in vivo Experiments

Fig 5-24 shows initial *in vivo* images acquired using the shimmed and unshimmed magnets and reconstructed with FFT. Distortion is reduced in the shimmed magnet images relative to the unshimmed magnet images.

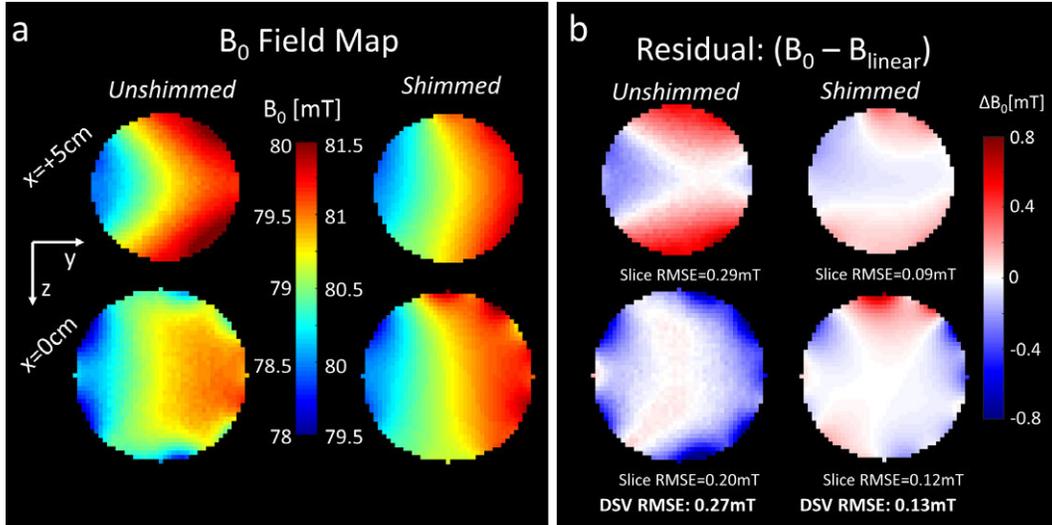


Figure 5-20: a. The measured shimmed field maps appeared more linear and monotonic than the unshimmed field maps. b. Error maps of the unshimmed and shimmed magnetic fields (relative to an ideal linear field). The shim array reduced the residual error by more than 50%.

5.5.4 Subsequent *in vivo* Imaging

Fig 5-25 shows later *in vivo* images acquired with this scanner. This work reflects subsequent improvement of the imaging sequences, RF coils, and reconstruction procedure [59]. These images were acquired of a healthy adult subject (male; age=63y). Shown are both T_1 - and T_2 -weighted images. The T_1 sequence used an inversion pulse at $TI = 400ms$. Total acquisition times for these images were 19m24s and 11m46s, respectively. Table 5.4 gives the sequence and pulse parameters for these acquisitions.

5.6 Discussion and Future Work

Here, I have demonstrated improved spatial encoding and image reconstruction using a portable, 80mT scanner by using target field shimming and close-fitting blipped gradient coils. This work, in conjunction with subsequent developments, has demonstrating the *in vivo* imaging capabilities of this scanner.s The chief remaining challenges for *in vivo* imaging are reducing RF interference and improving Rx coil performance. Reducing interference requires improved patient shielding and/or interference cancellation, while improving Rx coil performance requires non-resistive Q-lowering strategies.

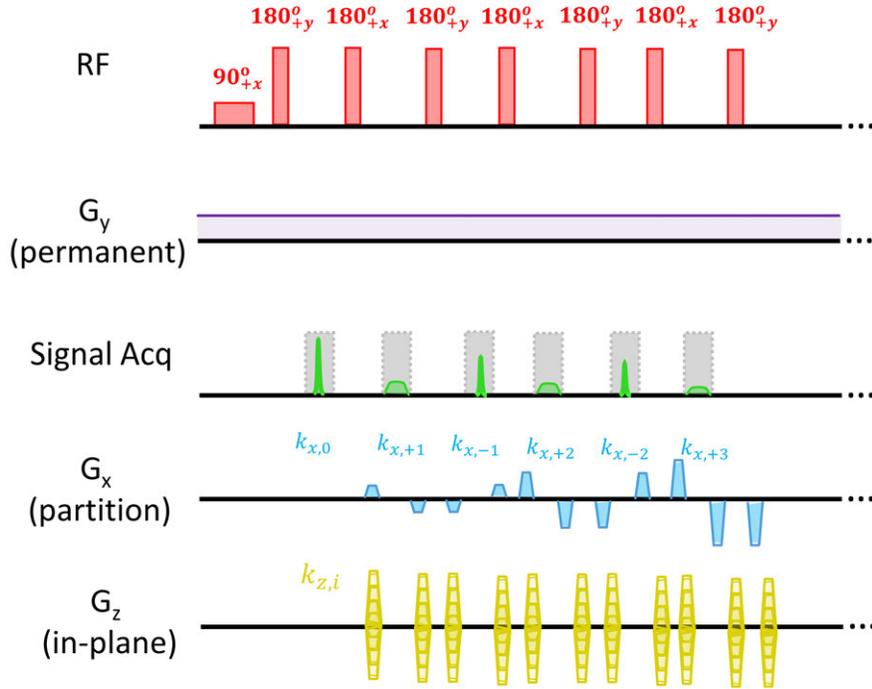


Figure 5-21: The RARE-type sequence for imaging experiments used spin-echo train with phase cycling; the built-in G_x gradient for readout encoding; blipped partition (x) phase encoding down the echo train; and in-plane (z) phase encoding from shot-to-shot.

5.6.1 System Performance

The shimmed Sparse Halbach B_0 field was 80mT, compared with 1.5 – 3T for a conventional scanner (a factor of 188 to 375 greater). This difference leads to inferior SNR with the Sparse Halbach as compared with a high-field MRI. A direct comparison of field homogeneity is not particularly useful, as the Sparse Halbach uses an intentionally-inhomogeneous field for readout encoding. A more useful comparison is between the homogeneity of a standard magnet and the nonlinear residual component of the shimmed Sparse Halbach B_0 field. For the Sparse Halbach, this residual is $130\mu\text{T}$ (RMS) for a corresponding Larmor frequency variation of 5.1kHz; a conventional magnet has approximately $1\mu\text{T}$ [14, 13] (over a $D = 40\text{cm}$ sphere), giving $\approx 40\text{Hz}$. These metrics suggest that greater image distortions will be seen when using Fourier encoding from the Sparse Halbach than from a high-field superconducting system. For example, the Sparse Halbach acquisitions use a readout bandwidth of approximately $500\frac{\text{Hz}}{\text{mm}}$. This bandwidth would be expected to lead to spatial distortions on the order of 1cm in the reconstructed images, while they would generate distortions under 0.1mm on a conventional system. This emphasizes the importance of using a model-based



Figure 5-22: A spiral helmet RF coil was used for transmit and receive for all acquisitions. The coil B_1 field was linearly polarized along x.

image reconstruction approach that accounts for these geometric artifacts.

Additionally, the Sparse Halbach performance was evaluated over a $D = 20\text{cm}$ sphere (for head imaging). The homogeneity of a conventional magnet is better still over an ROI of this size.

Field homogeneity also determines the RF pulse and spin excitation requirements. The Sparse Halbach has a total ROI field range (peak-peak) of $\approx 2.5\text{mT}$, which covers a total spin bandwidth of $\approx 106\text{kHz}$. Thus, 100kHz -pulses were used for RF excitation. Exciting all spins within this range while still meeting peak RF power limits (500W peak) required the use of chirped pulses of 1.5ms to 3ms duration. A conventional scanner can utilize reduced-bandwidth pulses because of the improved field homogeneity, helping to control RF power consumption. For example, the reduced spin bandwidth in highly-homogeneous high-field scanners enables the use of short ($\approx 100\mu\text{s}$) "hard" pulses. This is an example of the reduced sequence flexibility generally available in portable, low-field MRI systems.

Gradient power consumption during Sparse Halbach sequences was at most 121W on the Gz channel ($R_{dc} = 1.21\Omega$, $I_{max} = 10\text{A}$). This produced a field gradient of $8.15\frac{\text{mT}}{\text{m}}$ (efficiency = $0.815\frac{\text{mT}}{\text{m}\cdot\text{A}}$). For comparison, the Prisma gradient coils produce $80\frac{\text{mT}}{\text{m}}$ field gradients with 900A of current, giving an efficiency of $0.089\frac{\text{mT}}{\text{m}\cdot\text{A}}$ - a factor of 10 lower. This improvement in efficiency stems from two sources: the unshielded design used for the Sparse Halbach



Figure 5-23: Phantom images acquired with the shimmed and unshimmed magnet, and reconstructed using FFT and the generalized reconstruction.

gradient coils, and their smaller bore diameter.

The Sparse Halbach system is under 60cm in linear size; fits on a cart that can be wheeled through a building; and weights 123 kg (unshimmed magnet weight). This contrasts with a $\approx 2\text{m}$, 10-ton conventional scanner located in a dedicated suite. Note that the Sparse Halbach images shown here were acquired in a shield room, effectively increasing the size of the total system. However, work continues on ways to eliminate the need for a shielded room (see Section 5.6.2). The estimated material cost for the Sparse Halbach magnet is \$5000 (USD), with an estimated total system parts cost of \$20000 (USD). A new high-field body scanner typically costs over \$1 million (USD). These figures are not directly comparable (one is a parts cost, one is a sale price). However, the two estimates differ by a factor of 50, suggesting that the Sparse Halbach would be significantly cheaper than a typical MRI scanner.

5.6.2 Future Work

The *in vivo* images present reflect the subsequent work by others. This work continues, and specifically includes: improvements to the sequences and reconstructions; scanning

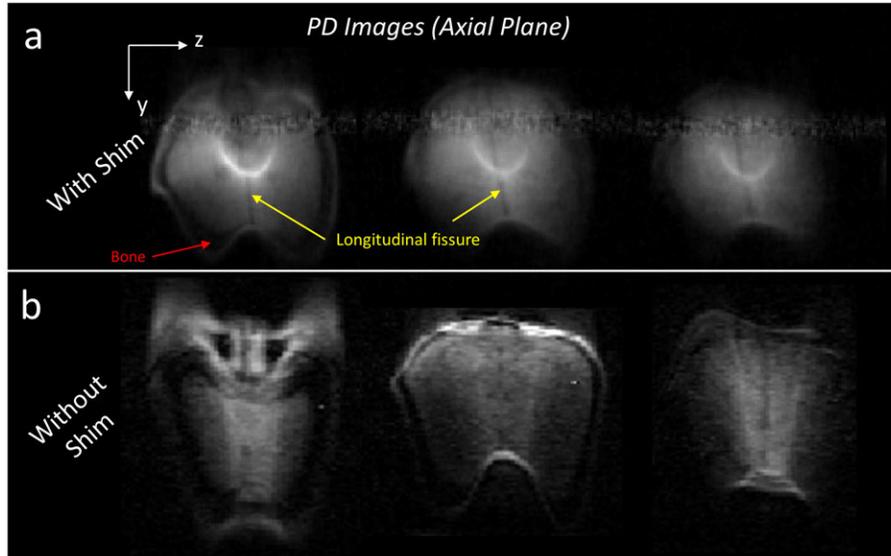


Figure 5-24: a. Early *in vivo* images acquired on the shimmed magnet and reconstructed using FFT. b. *In vivo* test image acquired on the unshimmed magnet and reconstructed using FFT.

additional human subjects, including non-healthy subjects; and RF interference cancellation [178]. RF interference cancellation for NMR/MRI has previously been explored [179, 180], and in this case will enable *in vivo* scanning outside a shielded environment.

Reducing geometric distortion in the reconstructed images could involve reshimming the magnet to further improve the spatially-encoding B_0 field. This could be done by re-optimizing and replacing the shim trays designed in this work. However, it may be preferable to use a closer-fitting shim set (ie, to place shims closer to the volume of interest), as described in Section 2.3.3. Another option is to use a "shim array insert" (such as with [181, 61]). This may offer complementary performance gains by being located at an inner part of the bore, and would not require removing the already-inserted shim trays.

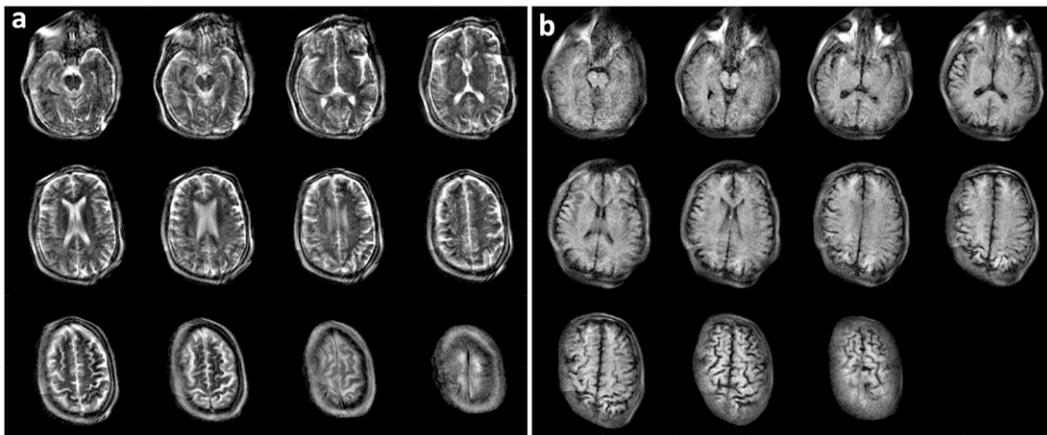


Figure 5-25: In vivo images acquired on the shimmed magnet after subsequent improvements to the RF coil, acquisition, and reconstruction. a. T_2 -weighted images. b. T_1 -weighted images incorporating an inversion-prep module in the sequence.

Table 5.2: Allowable shim options included 27 different magnet configurations. Configurations used between zero (the “empty” option) and five magnet blocks. All configurations fit in the incircle depicted in Fig 5-14.

Option #	Magnet #1	Magnet #2	Magnet #3	$\mu_m (A \cdot m^2)$
1	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{3}{8})''$ - N42	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{16})''$ - N42	$2 \times (\frac{1}{4} \times \frac{3}{8} \times \frac{1}{16})''$ - N42	1.662
2	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{1}{4})''$ - N42	$2 \times (\frac{3}{4} \times \frac{3}{8} \times \frac{1}{8})''$ - N52	$2 \times (\frac{1}{4} \times \frac{3}{8} \times \frac{1}{16})''$ - N42	1.633
3	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{1}{4})''$ - N42	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{4} \times \frac{3}{8} \times \frac{1}{16})''$ - N42	1.583
4	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{1}{4})''$ - N42	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{3}{16} \times \frac{3}{8} \times \frac{1}{16})''$ - N42	1.534
5	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{1}{4})''$ - N42	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{8} \times \frac{3}{8} \times \frac{1}{16})''$ - N42	1.485
6	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{1}{4})''$ - N42	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{3}{16} \times \frac{1}{8} \times \frac{1}{16})''$ - N42	1.436
7	$1 \times (\frac{1}{2} \times \frac{3}{8} \times \frac{1}{4})''$ - N42	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45		1.387
8	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{3}{8})''$ - N52	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N52		1.311
9	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{3}{8})''$ - N52	$2 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8})''$ - N45		1.252
10	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{3}{8})''$ - N52	$2 \times (\frac{1}{4} \times \frac{3}{8} \times \frac{1}{16})''$ - N42		1.178
11	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{3}{8})''$ - N52	$2 \times (\frac{1}{4} \times \frac{1}{8} \times \frac{1}{8})''$ - N42		1.113
12	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{3}{8})''$ - N52	$2 \times (\frac{1}{8} \times \frac{1}{8} \times \frac{1}{8})''$ - N48		1.053
13	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{3}{8})''$ - N52			0.983
14	$1 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{4})''$ - N52		0.885
15	$1 \times (\frac{1}{2} \times \frac{1}{4} \times \frac{1}{4})''$ - N42	$2 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8})''$ - N45		0.790
16	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{4} \times \frac{1}{8} \times \frac{1}{8})''$ - N42		0.735
17	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{8} \times \frac{1}{8} \times \frac{1}{8})''$ - N48		0.675
18	$2 \times (\frac{3}{8} \times \frac{3}{8} \times \frac{1}{8})''$ - N45			0.605
19	$1 \times (\frac{1}{2} \times \frac{1}{4} \times \frac{1}{4})''$ - N42			0.521
20	$3 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{8} \times \frac{1}{8} \times \frac{1}{8})''$ - N48		0.473
21	$3 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8})''$ - N45			0.403
22	$1 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{4})''$ - N52	$2 \times (\frac{1}{8} \times \frac{1}{8} \times \frac{1}{8})''$ - N48		0.361
23	$1 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{4})''$ - N52			0.291
24	$1 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8})''$ - N45	$2 \times (\frac{1}{8} \times \frac{1}{8} \times \frac{1}{8})''$ - N48		0.204
25	$1 \times (\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8})''$ - N45			0.134
26	$1 \times (\frac{1}{4} \times \frac{1}{8} \times \frac{1}{8})''$ - N42			0.065
27				0

Table 5.3: Imaging sequence parameters used for the phantom PD and in vivo PD acquisitions

	Phantom PD	<i>In vivo</i> PD
<i>TE</i>	10 ms	10 ms
<i>FOV</i>	$410 \times 26.4 \times 39\text{cm}$ cm	$410 \times 26.4 \times 39\text{cm}$ cm
<i>Resolution</i>	$0.8 \times 2.4 \times 7.5\text{mm}$	$0.8 \times 2.4 \times 7.5\text{mm}$
<i>Matrix Size</i>	$512 \times 101 \times 52$	$512 \times 101 \times 52$
<i>Averages</i>	4	8
<i>TA</i>	7m	14m
<i>TR</i>	1.0s	1.0s

Table 5.4: Imaging sequence parameters used to acquire the images shown in Fig 5-25

	Phantom PD	<i>In vivo</i> PD
<i>TE</i>	14 ms	167 ms
<i>TI</i>	400 ms	N/A
<i>FOV</i>	$18 \times 20 \times 16$ cm	$18 \times 20 \times 16$ cm
<i>Resolution</i>	$2 \times 2 \times 7\text{mm}$	$2 \times 2 \times 7\text{mm}$
<i>Averages</i>	2	4
<i>TA</i>	11m46s	19m24s
<i>TR</i>	1.8s	3.0s
<i>Peak Gx/Gz Current</i>	4.5/9.0 A	4.5/9.0 A

Chapter 6

The MR Cap: A Single-sided MRI Device for Point-of-Care Applications

6.1 Abstract

Conventional MRI systems are typically heavy, costly, and immobile, and are usually sited in a dedicated suite within a hospital. This operational paradigm is amenable to general-purpose, whole-body imaging where the patient is transported to the MR imaging suite, but precludes the use of MRI in a monitoring or point-of-care (PoC) setting. To address this, I developed the "MR Cap" – a lightweight ($<10\text{kg}$), portable MRI system for imaging over a reduced field-of-view (FOV) in the adult brain. This scanner is the size of bicycle helmet and can be positioned around the head for imaging over $\sim 8 \times 8 \times 3\text{cm}^3$ of brain.

The MR Cap used a numerically-optimized magnet design that produced a 64mT B_0 field with 4mT variation across the FOV. The gradient and RF hardware was also numerically optimized for imaging performance with the design ROI. System performance was demonstrated with 1D object projections and with 3D imaging in the ROI.

This work demonstrates the feasibility of a low-cost, lightweight MRI system potentially amenable to PoC applications. Potential applications, limitations, and technical challenges of this work are discussed. Much of this work has previously been published as a conference paper [182] and as a journal paper [183].

6.2 Introduction

6.2.1 Design Goals

Fig 6-1 illustrates the concept of the MR cap. We wanted system that could be supported by an articulating arm and positioned by hand, and that could image with 2mm in-plane resolution over an $\sim 8 \times 8 \times 3\text{cm}^3$ region of an adult brain. The system should a mean B_0 field between 50 and 100 mT, should be homogeneous to within 10% (ideally, 5%) over the sensitive ROI, and should have low material cost.

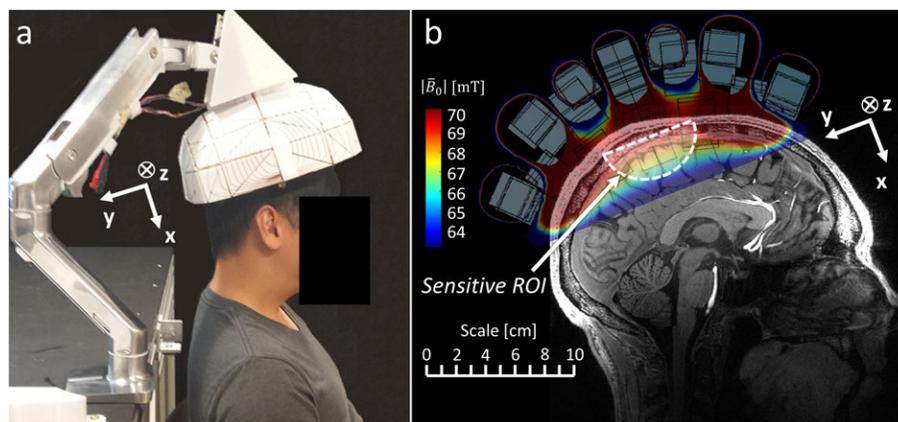


Figure 6-1: a. The concept for an MRI scanner than can be hand-positioned to image an arbitrary region of the brain. b. The scanner's sensitive ROI should overlap with the patient's skull, meninges, and cerebral cortex.

6.2.2 System description

The system is shown in Fig 6-2. The principal hardware components are the B_0 magnet (with built-in G_x field gradient), gradient coils along the y - and z -axes, and the RF Tx/Rx coil. Such single-sided magnets generally have highly inhomogeneous B_0 fields, meaning that we'd be limited to using spin echo train-type sequences. The system was designed for imaging in an $\sim 8 \times 8 \times 3\text{cm}^3$ lenticular ROI. This region corresponds to the volume bounded by 2 B_1 isochromats. This was chosen because the relation between signal available during a TSE-type sequence and B_1 strength has a sharp transition [184]. Thus, it was expected that the spatial B_1 profile would be the main determinant of imaging FoV.

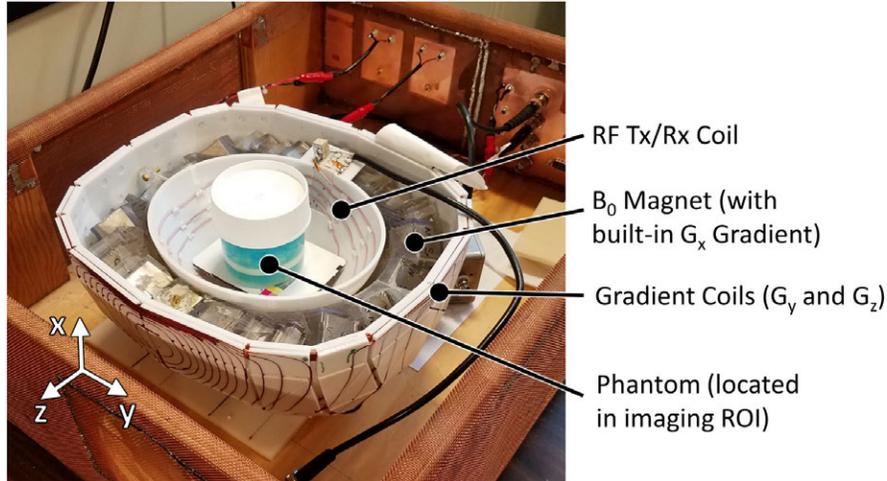


Figure 6-2: The MR Cap system with principal components labeled

6.3 B_0 Magnet

The magnet was designed by optimizing the positions and compositions of a set of magnet blocks arranged on a "cap-shaped" surface and was constructed using the procedures described in detail in Sec 3.6. It was then measured and assessed using several metrics of performance.

6.3.1 Magnet Design

The magnet was designed from NdFeB magnet block array with the magnets arranged in a cap-shaped configuration to fit about an ROI in the head. This design was expected to inherently contain a strong x-gradient. It was designed for a transverse B_0 (oriented along \hat{z}). The design began as a section of a Halbach Sphere magnet [185]. It was then discretized into blocks (for ease of manufacture), which were optimized by varying their sizes, compositions, and positions (Fig 6-3). The angular orientation of each block was fixed based on the analytical Halbach Sphere design and the block's initial position.

Optimization of the block array used a discrete set of possible blocks, given in Table 6.1. Blocks were purchased from Applied Magnets (Plano, Tx, USA) and had tolerances specified in Table 6.2. The optimization also allowed all blocks to move along x , and for six blocks were allowed to move along z as well (Fig 6-4). Blocks could move up to $\pm 1\text{cm}$ in 1mm increments. The use of a discrete design space (given the finite sets of block options and block positions) made it natural to use a genetic optimization (here, Matlab's `ga` tool).

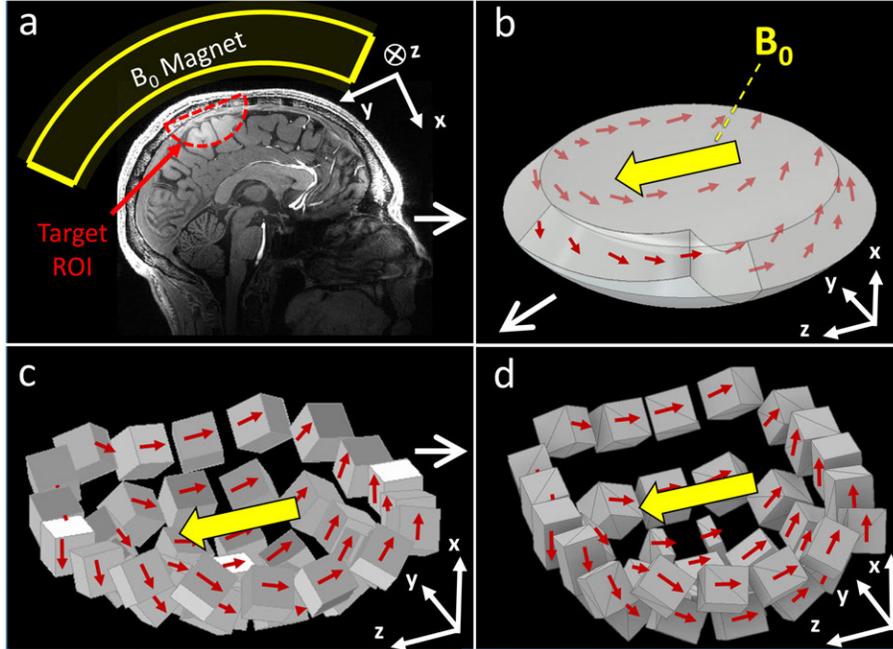


Figure 6-3: a: Magnet volume geometry and target ROI in relation to an adult head. b: Preliminary magnet design showing transverse-aligned B_0 field (yellow) and magnetization field (red). c: Preliminary magnet design discretized into 37 blocks. d: Optimized magnet design. Blocks have different positions and sizes relative to the unoptimized design.

The fields for all possible block options were precomputed in Comsol, and they were superimposed during optimization to compute the block array field. This was used because of the discrete design space. Fields were precomputed on a 1mm isotropic grid, matching the allowable motion increment of the blocks. Thus, computing the field of a shifted block only required shifting the entries of the precomputed field maps.

The optimization used the cost function $f_c = \frac{\max(B_0) - \min(B_0)}{\text{mean}(B_0)}$. This corresponds to $\frac{1}{Q}$ for an RF coil with enough bandwidth to cover the whole ROI (see Sec 3.5.2). Minimum mean B_0 was constrained so that $\text{mean}(B_0) \geq 50\text{mT}$. The initial guess for the optimization initial was the unperturbed, discretized Halbach Sphere configuration. Symmetry about the XZ and XY planes was required on allowable magnet designs. This reduced the number of optimization degrees of freedom to 24, and the appropriate choice of symmetry/antisymmetry ensured that the B_0 field pointed along \hat{z} . 20 iterations of the genetic optimization were run and the design with the lowest cost function was chosen.

The optimized design is shown in Fig 6-5 and contained four different block options. The optimized design was pinched together near poles compared with the starting guess (Fig 6-5). Simulated field maps (Fig 6-8) show a mean field of 67.5mT with an expected range of

Table 6.1: Block design options available to the optimization

Option Number	NdFeB Material Grade	Block Dimensions		Effective B_r if 1" Cube (T)
		SI Units (mm^3)	Mfgr. Spec.	
1	Empty	Empty	Empty	0
2	N45	$25.4 \times 25.4 \times 9.525mm^3$	$1 \times 1 \times 0.375in^3$	0.51
3	N42	$25.4 \times 25.4 \times 19.05mm^3$	$1 \times 1 \times 0.75in^3$	0.97
4	N42	$25.4 \times 25.4 \times 22.225mm^3$	$1 \times 1 \times 0.875in^3$	1.13
5	N42	$25.4 \times 25.4 \times 25.4mm^3$	$1 \times 1 \times 1in^3$	1.29
6	N52	$25.4 \times 25.4 \times 25.4mm^3$	$1 \times 1 \times 1in^3$	1.46
7	N52	$25.4 \times 25.4 \times 28.575mm^3$	$1 \times 1 \times 1.125in^3$	1.64
8	N52	$25.4 \times 25.4 \times 31.75mm^3$	$1 \times 1 \times 1.25in^3$	1.83
9	N52	$25.4 \times 25.4 \times 34.925mm^3$	$1 \times 1 \times 1.375in^3$	2.01

4.77 mT over the target ROI. Covering this bandwidth would require a coil with $Q = 14.2$. The cost of the magnets needed to realize this design was under \$450 USD (in year 2017).

6.3.2 Magnet Construction

The magnet former was 3D printed in 4 segments out of an acrylic resin (model # FLGPCL04) on a Form 2 SLA printer (Formlabs, Somerville, MA). The former contained several holes and standoffs for attaching RF coils, gradient coils, and mounting hardware. After printing the former was glued together (Loctite EA608 epoxy) and fastened together with screws. The magnets were next glued in with Loctite EA608 5-minute epoxy. They were first added at the center of the array, and subsequently by working out towards the pole ends (Fig 6-6).

6.3.3 Magnet Measurement and Performance Assessment

The magnet B_0 field was mapped with a 3-axis Hall effect probe (model # THM1176; Metrolab, Geneva, Switzerland) stepped by a 3-axis CNC robot. The field was measured at a resolution of $2.5 \times 5 \times 5mm^3$. The x -resolution was higher because of the expected steeper field gradient along that axis. Fig 6-7 shows the raw measured field and the measured

Table 6.2: Manufacturer-supplied tolerances for magnet block linear dimensions and remanent flux density

		Specification/Tolerance (mfr. spec in parentheses)
Block Linear Dimension		$\pm 51\mu m$ ($\pm 2mil$)
Material B_r	N42	1.29-1.32 T
	N45	1.35-1.38 T
	N52	1.46-1.48 T

data fit to a 12th-order spatial polynomial. Fig 6-8 shows the polynomial-fit measured field map together with the simulated field maps. Table 6.3 shows several metrics of magnet performance. The mean field was measured to be 63.6 mT ($f_L = 2.71Hz$) with an ROI range of 4.4 mT ($BW = 187Hz$), requiring a coil with $Q = 14.5$. The built-in field gradient was oriented along \hat{x} and $117\frac{mT}{m}$ in strength. The constructed magnet weighed 6.3kg and had linear dimensions of $11.3 \times 25.5 \times 21.8cm$. Prior to designing other system components, 1D NMR depth profiles were acquired using the built-in G_x gradient for readout encoding. This is described in Sec 6.6.1.

Table 6.3: B_0 magnet size, weight, cost, and performance metrics

Magnet Parameters	
<i>Size (along x)</i>	11.3 cm
<i>Size (along y)</i>	25.5 cm
<i>Size (along z)</i>	21.8 cm
<i>Weight (magnets+former)</i>	6.3 kg
<i>Cost (magnets)</i>	\$450 USD
<i>Mean B_0 (at ctr. of ROI range)</i>	63.6 mT
<i>B_0 Range over ROI</i>	4.4 mT
<i>Built-in B_0 gradient</i>	$117\frac{mT}{m}$

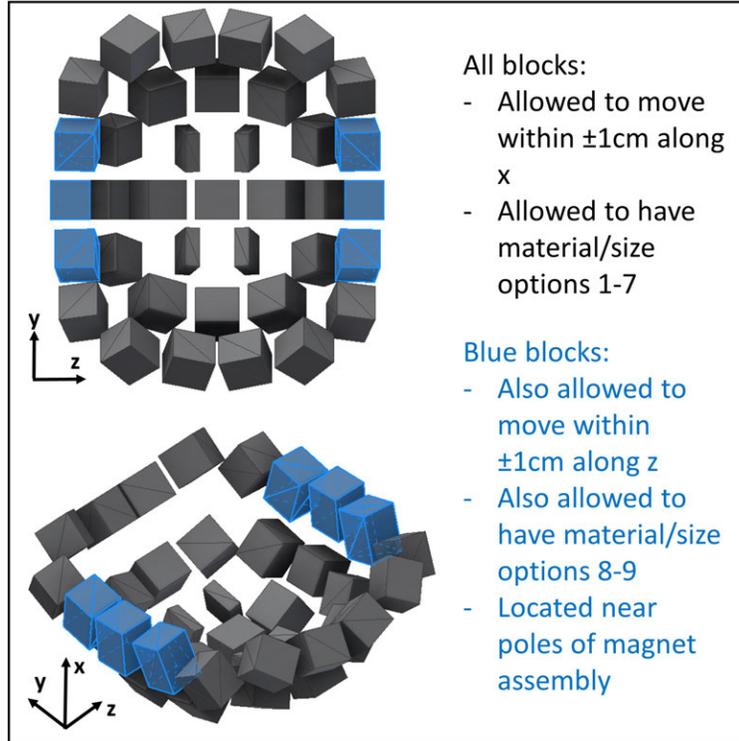


Figure 6-4: Degrees of freedom and allowable ranges thereof available to the magnet design optimizer. Different design options were available to different blocks.

6.4 Gradient Coils

Two gradient coils were constructed for blipped phase encoding along the y and z axes on a surface that fit outside the B_0 magnet (Fig 6-2). It was found that eddy currents and Ohmic heating did not appreciably affect the B_0 field or impair system performance. These coils were designed using the BEM stream function approach in Sec 4.4 and were built using the fabrication process in Sec 4.5.

6.4.1 Gradient Coil Design

The coils were designed on a bowl-shaped surface that fit outside the B_0 magnet (Fig 6-2). The surface extended 2.5cm beyond the top flat surface of the magnet array in order to better encircle the imaging ROI. This change was implemented after an earlier coil prototype – designed on a surface without this extra extension – did not have adequate linearity or imaging performance in the target ROI. It slightly increased the size of the scanner, but did not make it more likely to intersect the patient anatomy.

The coil optimization used the BEM stream function target field approach with reg-

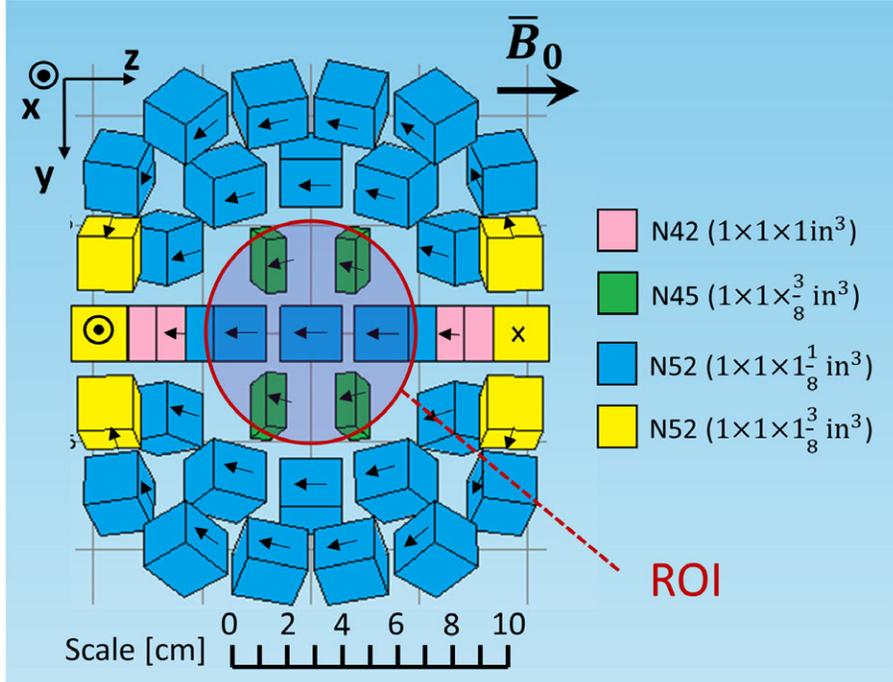


Figure 6-5: The optimized B_0 magnet design used 4 different block size/material options.

ularization. Each coil contained 2 layers of windings and was made with AWG20 wire ($OD = 0.812 \text{ mm}$). This diameter determined the minimum wire spacing. The stream function regularization parameter was fixed during optimization. Smoothing was performed with a Gaussian filter on the computed stream function (see Sec 4.4.4).

We ultimately wanted both coils to produce pure 1st-order spatial fields in the imaging ROI. However, the optimizer's target field pattern contained the desired 1st-order gradient term along with a 2nd-order term (XY for the y-gradient coil, XZ for the z-gradient coil). This was done to counter the tendency of a single-sided coil to naturally fall off in efficiency as one moves away from its surface (in this case, along \hat{x} . For example, two G_z coils designed using first-order (Z) and first+second-order (Z + XZ) field targets are shown in Fig 6-9, along with data summarizing the performance of the two designs. The 1st + 2nd-order coil shows much better uniformity of linearity within the target ROI compared to the pure 1st-order target field coil. The relative weighting of the 1st- and 2nd-order field terms was chosen from the result of the 1st-order target field optimization. The amplitude of the 2nd-order residual field term in this coil's field map was used as the 2nd-order target term for the hybrid 1st+2nd order target field optimization. The design ROI was the same as for the B_0 magnet.

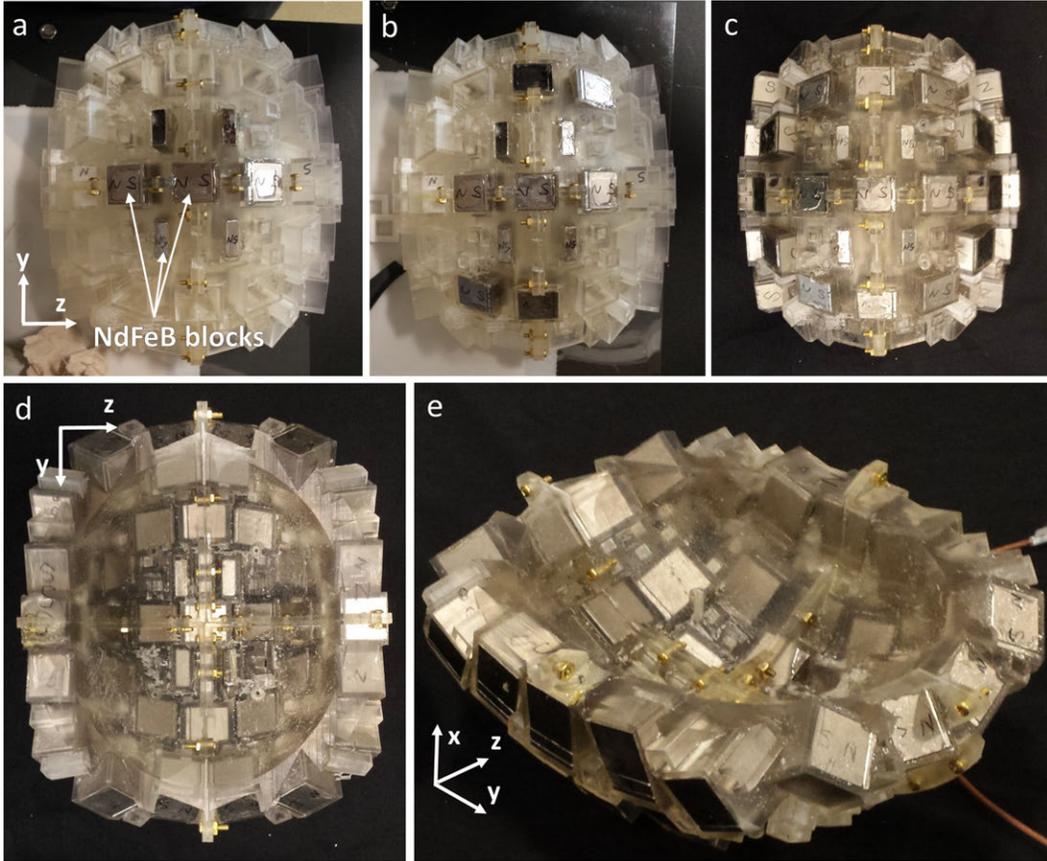


Figure 6-6: Construction of the B_0 magnet. a. Magnets were attached starting in the middle of the array at the bottom of the “bowl”. b. They were then added working outwards along y . c. The completed magnet shown from below. d. Completed magnet shown from above. e. Completed magnet show in an isometric view.

The optimized stream function was then smoothed and discretized into 1D paths (see Sec 4.4.4). This step used $N = 20$ level sets for the G_y and G_z coils. Fig 6-10 shows the resulting G_y and G_z coils. These coils were constructed by gluing a set of flat triangular and rectangular tiles onto a coil former. The tiles contained etched grooves for holding the windings. The surface and windings were unfolded onto a polyhedral net so that the set of tiles could be printed as a single flat part (Fig 6-11). Note that while this may satiate your inner geometry nerd, it is a more cumbersome process than the CAD pipeline used for the Head-Optimized MRI gradient coils (Sec 7.5.3): figuring out the transformations to “unfold” the windings paths takes a fair amount of manual work. It is easier to not unfold the windings, to etch the coil former in its appropriate geometry, and to then reorient the etched coil former pieces for printing. The windings were projected onto this piecewise planar surface and the result was simulated (Biot-Savard software, Ripplon).

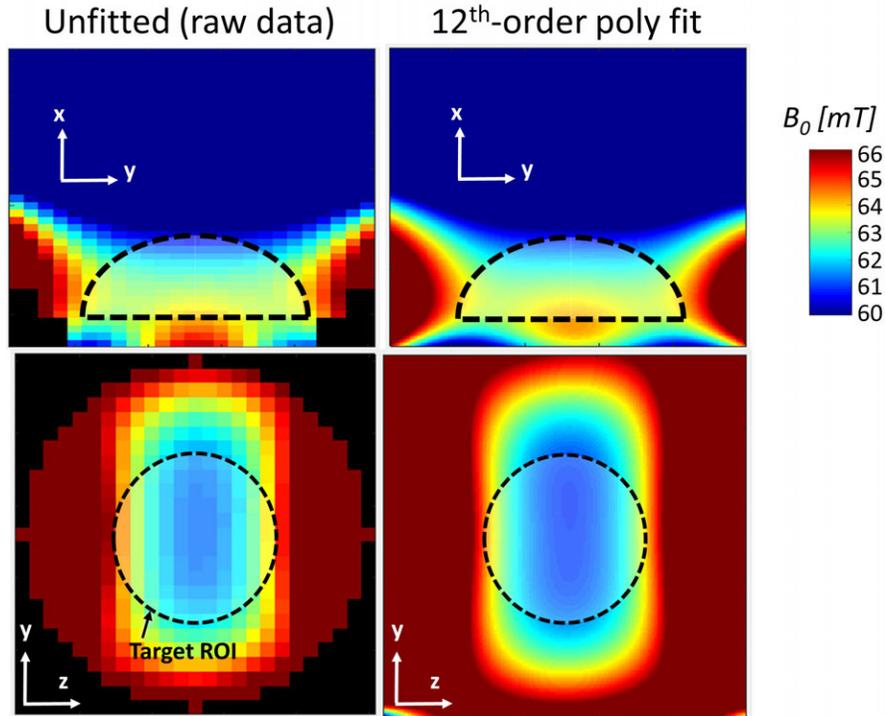


Figure 6-7: Measured, unfiltered B_0 maps and measured B_0 maps fit to a 12th-order 3D polynomial

6.4.2 Gradient Coil Construction

Each of the gradient coils was built from AWG20 enameled wire press-fit into a set of grooves. The grooves were etched into two sets of tiles (one for the G_y coil and one for the G_z coil) which were 3D printed in PC-10 polycarbonate material (Fortus 360mc printer; Stratasys, Eden Prairie, MN). The G_y tiles were first epoxied (Loctite EA608) onto a 2mm-thick former. The wire was then pressed into the grooves to wind the coil. Next, the G_z were epoxied on top of the completed G_y coil, and the G_z coil was wound. The total thickness of both coils was 7mm. The coils weighed under 2kg. Fig 6-12 shows the wound coils.

6.4.3 Coil Measurement and Assessment

The gradient coil fields were measured with the THM1176 3-axis hall probe stepped by a 3-axis robot (Fig 6-13), measured at 5mm isotropic resolution. The field was measured with the coil at $I = 0A$ and at $I = 2A$, taking the difference between them gave the field with background offsets removed. The raw data were noisy, and were fit to 5th-order polynomials with the appropriate symmetry constraints in order to estimate efficiency and linearity. Fig 6-14 shows these polynomial-fit field maps. The G_y and G_z coils were measured to have

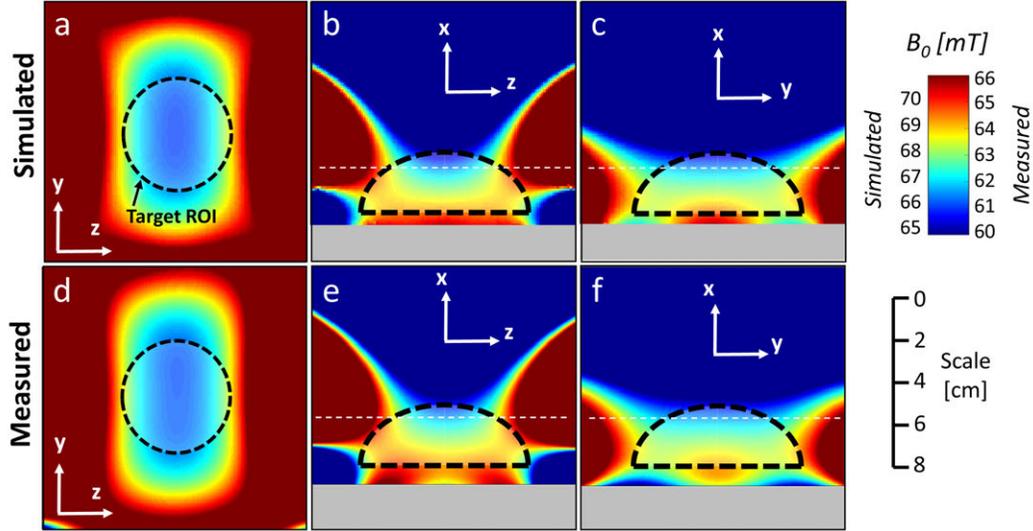


Figure 6-8: Simulated and measured B_0 field maps; note that the simulated and measured field maps have different color scaling. a. Simulated - YZ (transverse) plane. b. Simulated, XZ plane. c. Simulated, XY plane. d. Measured, YZ plane. e. Measured, XZ plane. f. Measured, XY plane.

mean efficiencies/linearities of $1.117 \frac{mT}{m \cdot A} / 33.3\%$ (along \hat{y}) and $1.019 \frac{mT}{m \cdot A} / 27.8\%$ (along \hat{z}). The efficiencies varied from 0.979 to $1.348 \frac{mT}{m \cdot A}$ and 0.845 to $1.03 \frac{mT}{m \cdot A}$ over the design ROI, for G_y and G_z , respectively.

6.5 Radiofrequency (RF) Coils and Imaging Setup

6.5.1 RF Coils

Two RF coils were used in experiments: a "volumetric" RF coil for imaging over the whole sensitive ROI, and a "narrow" coil for initial experiments on a narrow phantom.

Volume Imaging RF Coil

The volumetric RF coil was designed on a saucer-shaped and -sized surface fitting insides the B_0 magnet (Fig 6-15) using the stream function target field approach (Sec 4.4). This coil was designed for imaging over the entire sensitive volume, and coil design used the same ROI as the B_0 magnet and gradient coils. The field target was for a uniform x -aligned B_1 field, with the y -component equal to zero. The optimized coil stream function was discretized into 4 windings (Fig 6-16). It was then constructed by press-fitting Litz wire into a 3D-printed polycarbonate former (PC-10 material; Fortus 360mc printer). The coil was next tuned to $f = 2.67 MHz$ and matched to 50Ω . Its Q-factor was decreased by adding a 1Ω -resistor.

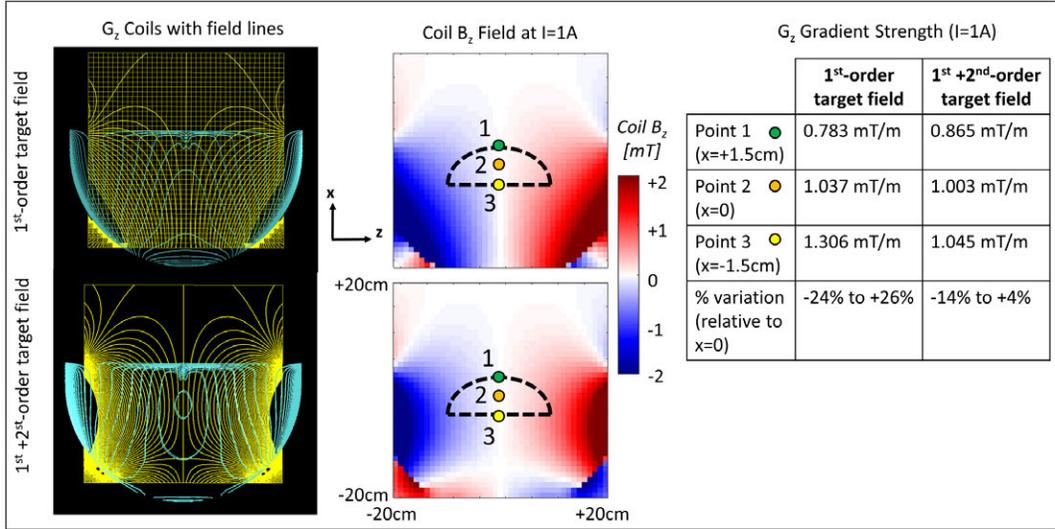


Figure 6-9: G_z coils designed with and without a 2nd-order spatial target field term, along with simulated field maps. The table shows the measured gradient strength at three points in the imaging ROI. The simulated field was more linear when using the additional 2nd-order target field term and the efficiency was only reduced by 3%.

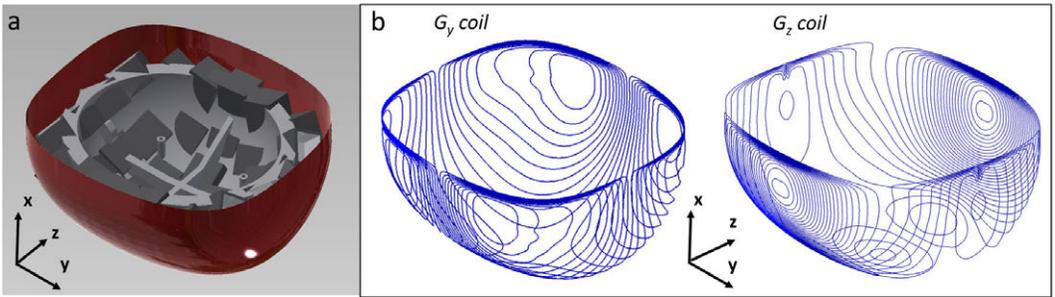


Figure 6-10: Optimized gradient coil designs. a. The gradient design surface (red) shown outside the B₀ magnet. b. G_y coil design. c. G_z coil design

This raised its 3dB-BW from 72kHz to 157kHz.

Narrow RF Coil for 1D Profiling

Initial 1D profiling experiments used a smaller coil that fit closely around a 1D test phantom (Fig 6-16). This coil contained 5 turns of AWG22 wire, was $L = 3\text{cm}$ long and $D = 1.8\text{cm}$ wide. The coil was tuned to $f = 2.67\text{MHz}$, was matched to 50Ω and had a bandwidth of 138kHz.

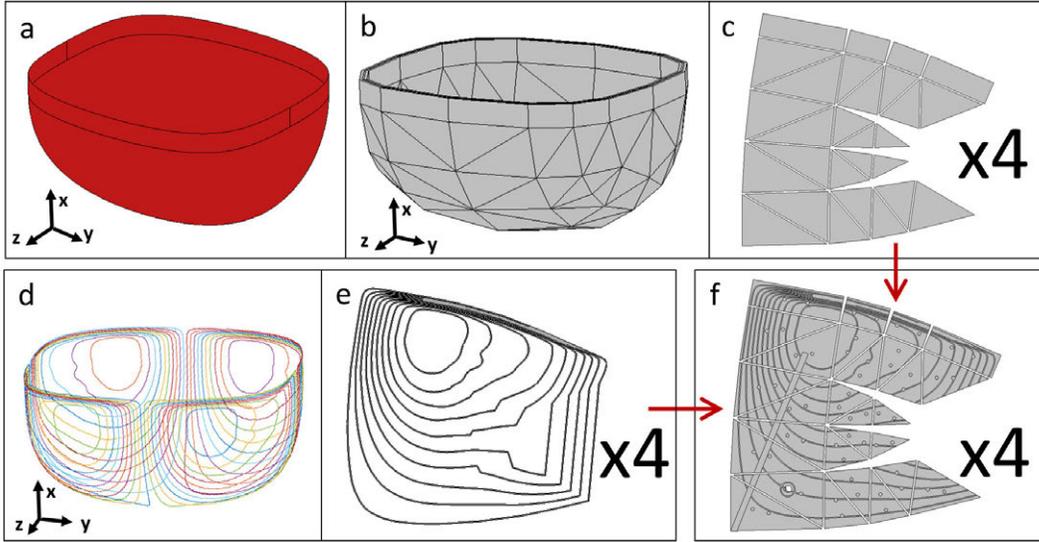


Figure 6-11: a. Smooth gradient design surface. b. Discretized piecewise-linear gradient design surface. c. Unfolded quadrant of the piece-wise linear surface. d. Optimized Gy coil windings designed on the smooth design surface. e. One quadrant of Gy windings unfolded onto the unfolded, discretized surface. f. Final design of one Gy coil former quadrant. Shown are the etched wire paths and features for interfacing with other parts of the system.

6.5.2 Imaging Console and Amplifiers

All experiments used an Apollo console (Tecmag, Houston, TX, USA) for gradient and RF pulse generation and signal recording. The RF system used a passive crossed-diode T/R switch with 2 quarter wave elements and two sets of series crossed diodes in the transmit path. A wideband 50Ω -input MITEQ AU-1583 (MITEQ, Hauppauge, NY, USA) preamplifier was used for receive. RF pulse transmit used a $2.5kW$, $2.5MHz$ - $4.5MHz$ RFPA (Tomco model BTO2000-AlphaS-3MHz, Stepney, SA, Australia). The gradient coils used home-built current-mode amplifiers incorporating a push-pull pair of linear op-amps (OPA549, Texas Instruments, Dallas, TX, USA) [186]. Each amplifier was capable of driving $\pm 10A$ at a voltage of $\pm 15V$ at a bandwidth of $20kHz$. All experiments were performed in a shielded box (Fig 6-17).

6.5.3 Imaging Phantoms

All phantoms used in this work are shown in Fig 6-18. All were filled with an aqueous phantom solution consisting of 0.9% NaCl and 0.2% Gd-DPTA, and dyed with blue food coloring. The phantoms were constructed out of various plastics, including PC-10 polycarbonate, ABS, acrylic, and Formlabs acrylic FLGPCL04 resin.

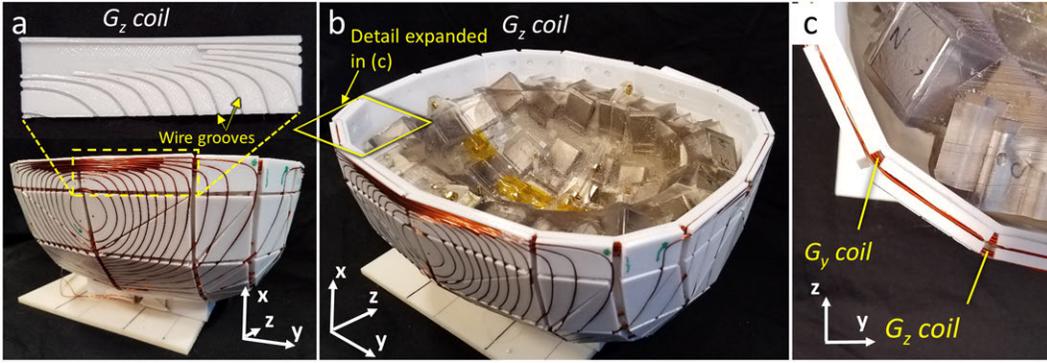


Figure 6-12: Constructed gradient coil set. a. The coil was built from AWG20 wire press-fit into 3D-printed tiles. b. Isometric view of the coils fitting outside the B_0 magnet. c. The G_y coil was built first, and the G_z coil was then built atop it.

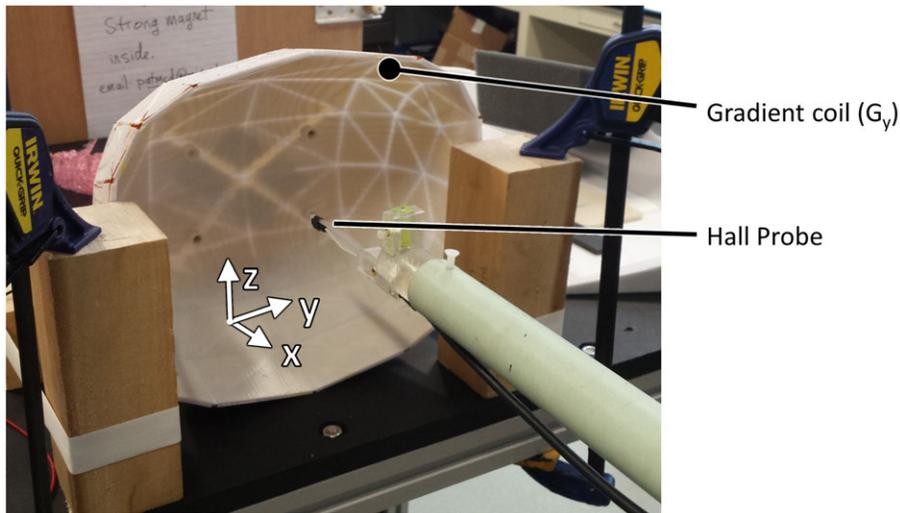


Figure 6-13: Field-mapping the G_y gradient coil

Narrow 1D Profiling Phantom

Fig 6-18a shows the 1D phantom used for initial depth profiling experiments. The phantom had an $ID = 10\text{mm}$ circular cross-section and contained three 5mm -thick regions of phantom solution spaced 5mm apart. This phantom was 3D printed in an acrylic resin on a Form 2 printer. The narrow RF coil was wound tightly around the phantom (see Sec 6.5.1).

Full-FoV 1D Profiling Phantom

Fig 6-18b shows the 1D phantom used for initial depth profiling experiments within a 3D volume. The phantom had a $ID = 63\text{mm}$ circular cross-section and contained three 5mm -thick regions of phantom solution spaced 5mm apart. This phantom was built with an ABS

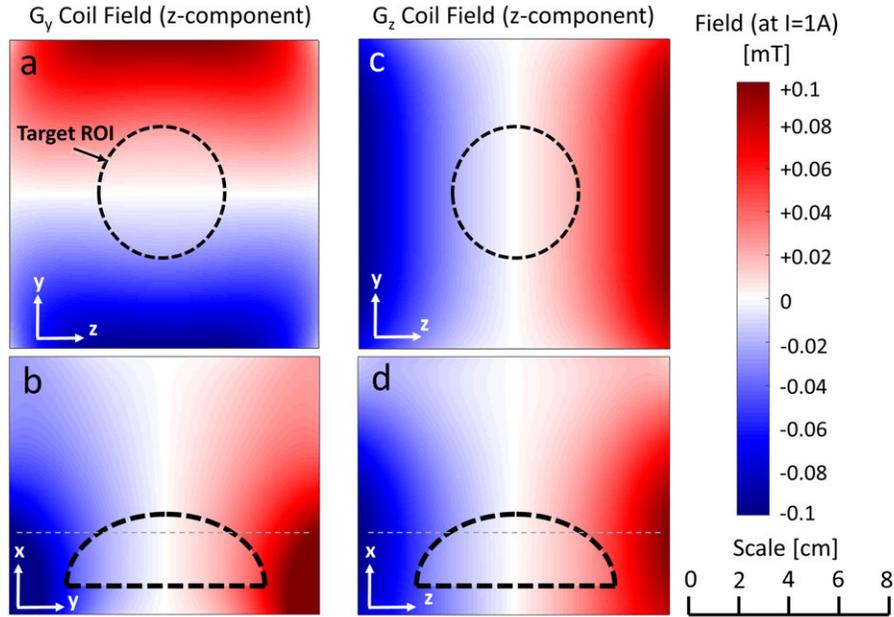


Figure 6-14: Field maps of the gradient coils at $I=1A$. a. G_y coil field, YZ plane. b. G_y coil field, XY plane. c. G_z coil field, YZ plane. d. G_z coil field, XZ plane.

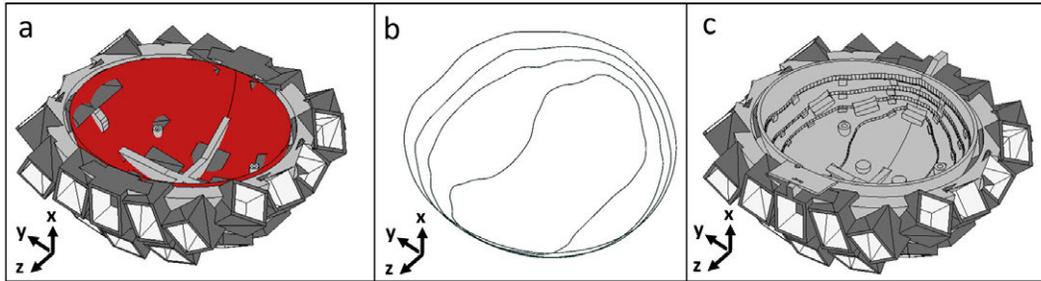


Figure 6-15: a. RF coil design surface (red) shown inside the B_0 magnet former. b. Numerically-optimized 4-winding RF coil. c. Model of the finished RF coil shown inside the B_0 magnet former.

structure placed in a plastic bottle.

3D Volume Imaging Phantom

Fig 6-18c shows the 3D phantom used for multislice imaging. The phantom had $ID = 63mm$ circular cross-section and contained four layers between $6mm$ and $10mm$ thick. The different layers contained different structure, shown in Fig 6-18c. The phantom was built from a machined ABS/acrylic object placed in a plastic bottle that was filled with phantom solution.

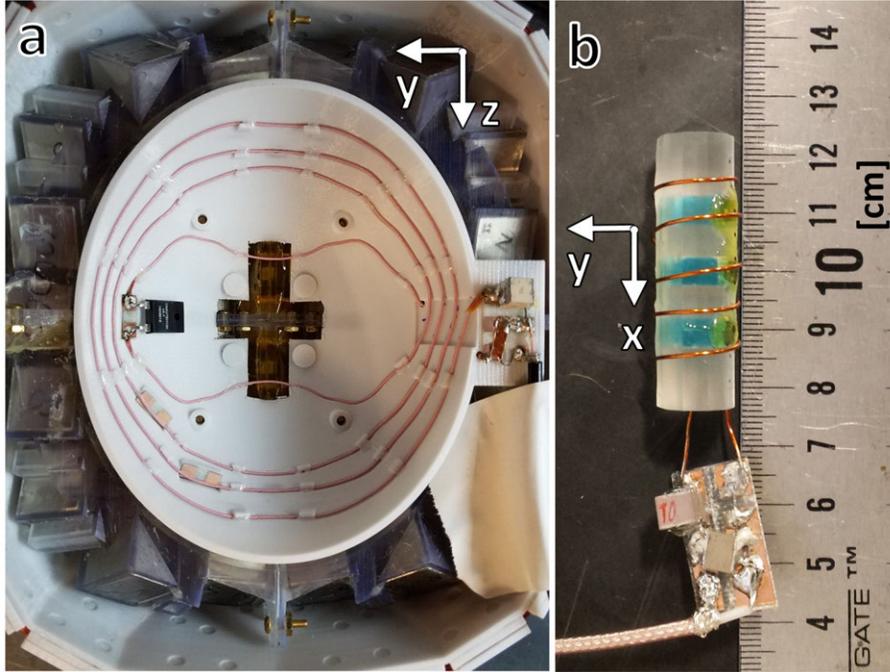


Figure 6-16: a. Volumetric RF coil for imaging over the whole ROI. b. Narrow RF coil for initial phantom experiments.

6.6 Sequences and Imaging

Both high-resolution 1D "Depth Profiling" acquisitions and lower-resolution multislice 2D imaging acquisitions were performed on phantoms with the MR Cap. Fig 6-19 shows the basic CPMG sequence used for all acquisitions, and Table 6.4 shows the sequence parameters used for the different experiments.

6.6.1 Depth Profiling

1D depth profiling used a RARE-type sequence with data acquired at each echo with the built-in B_0 gradient providing readout encoding. Data from different echoes in the spin echo train were averaged together. Different experiments were performing using different coils/phantoms, and with/without YZ-plane phase encoding.

Narrow 1D Phantom/Coil Experiments

Hard pulses were used for experiments with the narrow 1D phantom (Sec 6.5.3) and narrow RF coil (Sec 6.5.1). No phase encoding was used. The 90° and 180° pulses were of durations $\tau = 2\mu s$ and $\tau = 4\mu s$ and powers $125W$ and $500W$, respectively. 64 averages were used

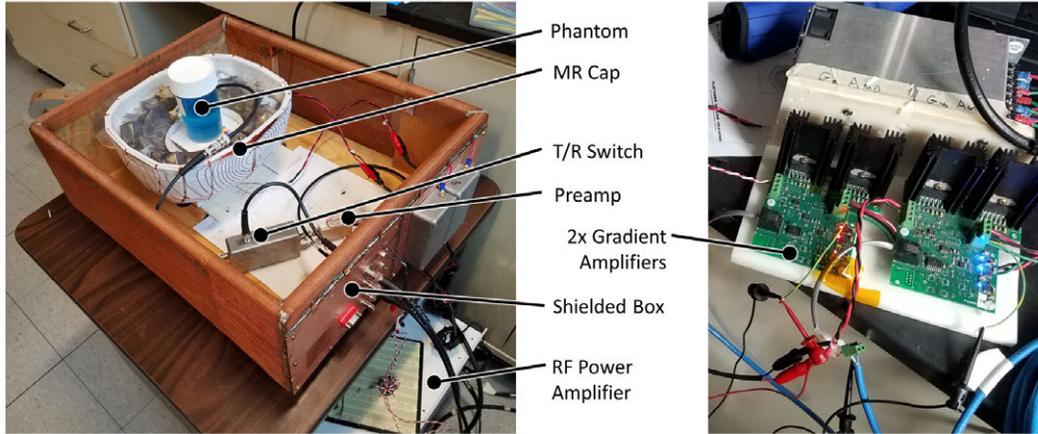


Figure 6-17: System configuration for imaging experiments showing all components except the console and interfacing PC

and the total acquisition time 60s. The readout window was truncated from 512 points to 128 points during reconstruction, giving an effective resolution of 0.25 mm and readout bandwidth of $1221 \frac{Hz}{Px}$.

Fig 6-20 shows the inverse-DFT reconstructed frequency spectrum/1D depth profile. The reconstructed profile matches the geometry of the phantom. However, the spectrum is geometrically distorted near the ends of the phantom, corresponding to regions of steeper G_x field gradient.

Volume Phantom/Coil 1D Experiments

Volume coil experiments used frequency-swept WURST pulses [177] with $\tau_{90/180} = 3ms$ and $BW = 200kHz$. Peak pulse power was $32W$. A hard pulse would have required an impractically-high power level to cover the spin bandwidth of the sample. This choice of pulses resulted in all echoes being "spectral"-type echoes with different frequencies being refocused at different times. These spectral echoes were retrospectively converted in to FID-type echoes by subtracting out the quadratic phase profile of the RF pulse [177]. The data were then truncated in the time-domain to produce an effective resolution of $0.9mm$ and bandwidth of $4464 \frac{Hz}{Px}$. 64 averages were used, and the acquisition time was 30s.

Fig 6-21 (purple trace) shows the 1D depth profile obtained from this experiment. The profile shows three peaks, but they overlap significantly. This was postulated to be due to the layers of the phantom being perfectly flat, while the isocontours of the B_0 field curve away from the YZ plane (Fig 6-8).

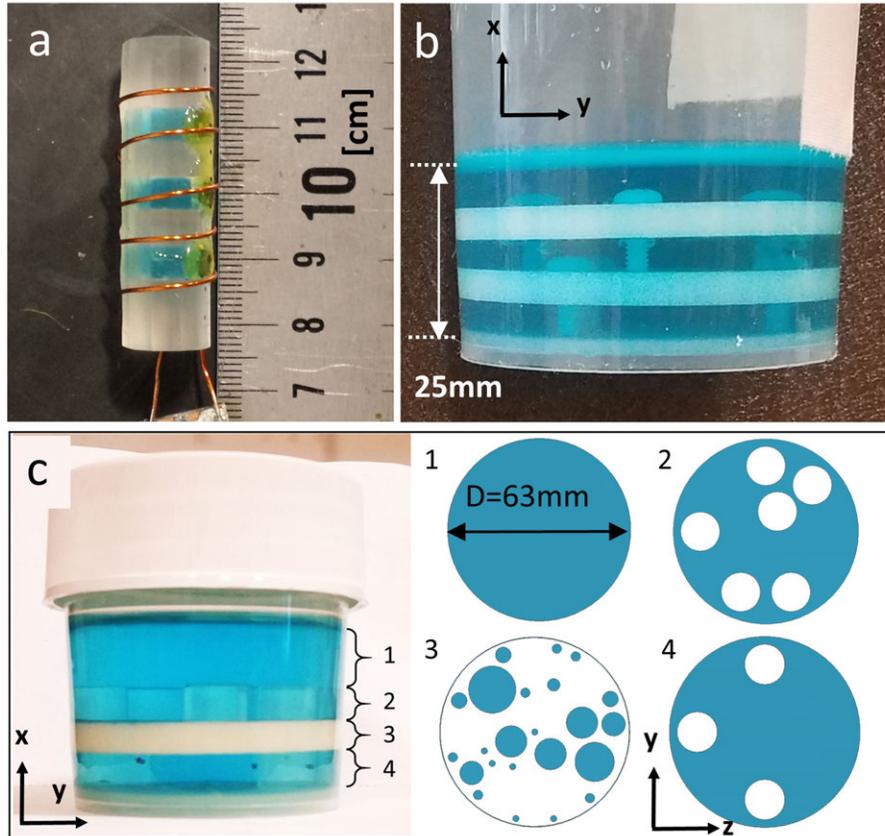


Figure 6-18: a. Narrow 3-disc phantom used for initial depth profiling experiments. b. 3-layer phantom for volumetric depth profiling experiments. c. Volumetric phantom with internal structure used for 3D imaging experiments.

Volume Phantom/Coil 3D Experiments

These experiments used the same sequence and pulses as the 1D volume coil/phantom experiments with the addition of in-plane (YZ) phase encoding. The phase encoding was applied in a balanced manner about each readout window (Fig 6-19). Each shot used the same phase encoding for all echoes; the phase encode was varied from shot-to-shot. Phase encoding used trapezoidal gradient blips with $\tau = 1ms$ and max current of $I = 0.62A$. In-plane resolution was $14 \times 14mm^2$ and a $N_y \times N_z = 5 \times 5$ grid was acquired. 64 averages were acquired, giving an acquisition time of 11m.

Fig 6-21 (blue traces) show the 3D depth profiles acquired at three phase encode pixels. These depth profiles show three peaks with better definition than the 1D depth profile.

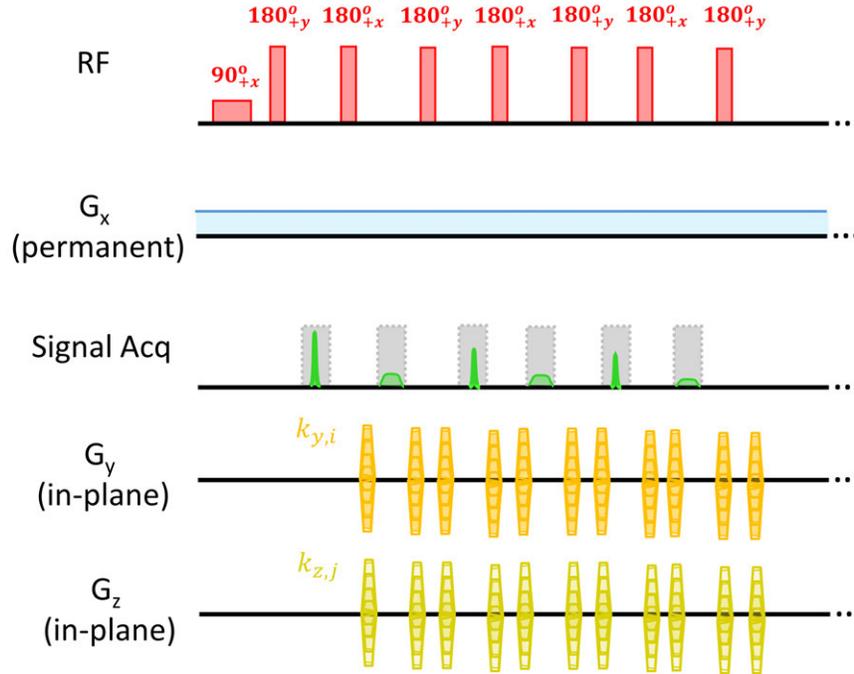


Figure 6-19: Pulse sequence used for MR Cap imaging experiments. G_y and G_z – when used – were iterated over from shot to shot. Pulse, gradient moments, and sequence timing were different for different experiments.

6.6.2 2D Multislice Imaging

Multislice imaging experiments used the 3D volume imaging phantom (Sec 6.5.3) and the volume RF coil. Y-axis phase encoding was performed down the echo train; z-axis encoding is performed shot-to-shot. Peak gradient current was $I = 4.62A$ with trapezoidal $\tau = 1ms$ pulses. The in-plane matrix size was 31×31 and the nominal resolution was $2 \times 2mm$. The RF pulses were Tukey-windowed sinc pulses with $BW = 30kHz$, $\tau = 110\mu s$, and peak power $315W$. Six interleaved slices were acquired with $f_{RF} = \{2.60, 2.63, 2.66, 2.69, 2.72, 2.75MHz\}$. The built-in B_0 gradient was used as a slice-select gradient. The gradient strength gave slices with thickness between $4mm$ and $8mm$. $N = 24$ averages were acquired, giving $TA = 12m40s$. Only the center one or three points of each spin echo were used; ie there was no readout encoding. This was done to increase image SNR. Reconstruction was performed with Fourier reconstruction and exponential-window apodization along the z-encoding dimension (the shot-to-shot dimension).

Figure 6-22 shows the acquired images. The different slices of the phantom can be seen, along with different structures within each slice. However, several types of imaging artifact

Table 6.4: Sequence parameters and details for the three imaging experiments performed.
 * - Effective parameter after retrospective down-sampling.

	Narrow Profiling	1D	Volume Profiling	Depth	Volume Imaging	3D
N_{TE}	6		2		32	
N_{samp}	128*		56*		64	
N_{avg}	64		64		24	
<i>Readout BW</i>	$1211 \frac{Hz}{Px}^*$		$4464 \frac{Hz}{Px}^*$		$1739 \frac{Hz}{Px}$	
<i>TR</i>	923 ms		500 ms		1.0 s	
<i>TE/ΔTE</i>	3.5/3.5 ms		10/10 ms		3/3 ms	
<i>TA</i>	60s		30s (1D)/11m (3D)		12m40s	
<i>Pulse Type</i>	Hard Pulse		Chirped Pulse		Windowed Sinc	
<i>90/180 Pulse BW</i>	500/250 kHz		200 kHz		30 kHz	
<i>90/180 Pulse τ</i>	2/4 μs		3/3 ms		110/110 μs	
<i>RF Peak Pow.</i>	500 W		32 W		315 W	
<i>PE Matrix</i>	N/A		5×5		31×31	
<i>PE Resolution</i>	N/A		$14 \times 14 mm^2$		$2 \times 2 mm^2$	
<i>Gradient I_{max}</i>	N/A		0.62 A		4.62 A	

are present.

6.7 Discussion and Next Steps

The system demonstrated the ability of a point-of-care system to acquire 1D and 3D images over a reduced FOV sensitive region. As desired, the system is lightweight (under 10kg) and can be attached to a posable arm. The cost of system components is low (magnet material under \$450 USD and gradient driver parts cost \$100 USD per channel.) In short, the prototype MR Cap satisfies our design objective to create a 3D brain imaging device that is lightweight, portable, and low-cost. However, several limitations exist in the presented system.

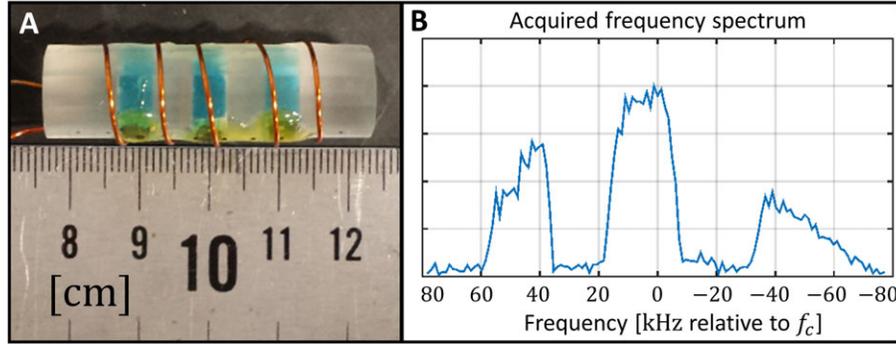


Figure 6-20: Narrow 1D imaging phantom and measured NMR frequency spectrum.

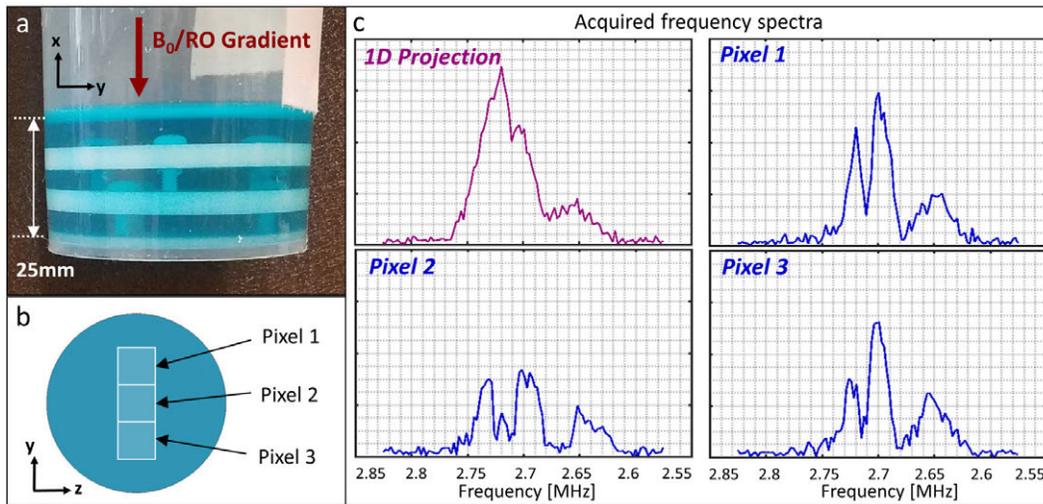


Figure 6-21: Volumetric depth profiling phantom and depth profiles. Depth profiles were measured with no in-plane encoding (purple) and at three phase-encoded pixels in the YZ plane (blue).

6.7.1 Depth Profiling

The measured depth profiles in Fig 6-21 do not precisely match the structure of the phantom. This is especially apparent for the depth profile measured without YZ-plane phase encoding: in this measurement, three peaks are visible, but they overlap and their boundaries are blurred. This is likely a consequence of the misalignment between the imaging phantom structure and the isochromats of the B_0 /readout encoding field whereby phantom structures are spread out over a large number of readout frequencies. This is in contrast to traditional NMR depth profiling with highly-planar isochromats matching a sample with planar, laminar structure. Precise alignment of the magnetic field and sample enables the acquisition of high-resolution depth profiles [124]. This approach has been used to measure artwork composition [130, 129], to assess surface degradation of materials [126], and for biomedical applications

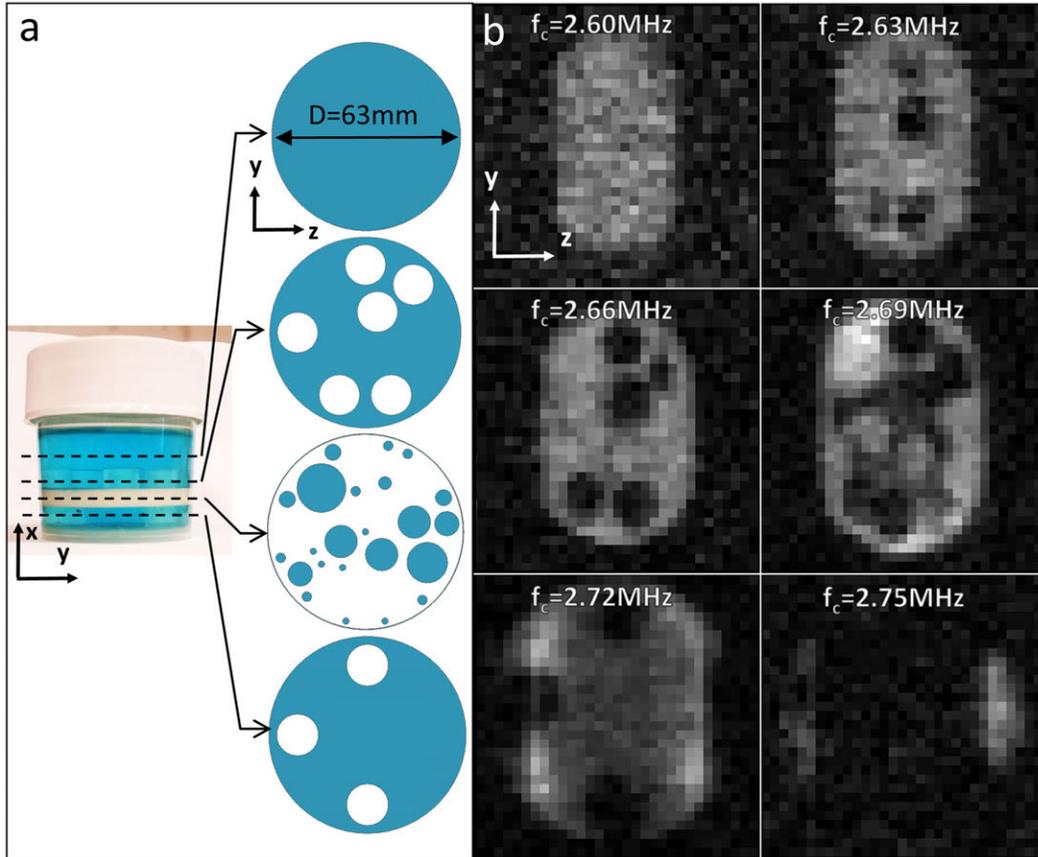


Figure 6-22: Volumetric imaging phantom and reconstructed 2D multislice images.

[124] such as skin profiling [144]. Successful extension to depth profiles 1-2cm into the human brain could potentially to assess pathologies of the dura and detect subdural hematomas. However, the acquisition of high-resolution in vivo cerebral depth profiles will be challenging given human anatomical variation and the difficulty of aligning the anatomy of interest with the magnetic field isochromats. In this work, we found that a 3D acquisition with low-resolution in-plane phase encoding produced higher-quality depth profiles of a phantom than a true 1D depth acquisition. This suggests that such an approach may also help in vivo.

6.7.2 Image Artifacts

The z -dimension of the sensitive ROI was truncated by approximately 4cm less than the design target in several imaging slices (Fig 6-23a. This likely results from the highly-inhomogeneous B_0 field in that region – i.e. the static gradient in that location is steeper. One effect of this is a reduction in slice thickness and an associated reduction in signal

level. Additional issues that may arise are T_2 signal losses due to diffusion effects and a shortening of the local T_2^* . T_2^* decay is largely refocused at the center of each spin echo, but we averaged together the center three points of each spin echo during image reconstruction, introducing a dependence on local T_2^* . Mitigation of this effect (restoration of this lost FOV) would require either the application of shimming capability to the existing magnet or a new design where the optimization more explicitly penalized peak ("min-max") inhomogeneities. Furthermore, until more experience is gained with such systems, the clinical implication of the reduced FOV is not clear. In fact, reducing the FOV is sometimes desired in MRI to focus on an anatomical region of interest or to reduce the image encoding burden required to avoid aliasing [187, 188]. The MR Cap, by design, is such a "zoomed" acquisition device.

Fig 6-23b shows another image artifact: hyperintense regions in the $fc = 2.72$ MHz slice of the 3D acquisition. These likely result from the non-uniform slice thickness arising from the B_0 field and the fixed BW RF pulses. The RF pulses were designed to have uniform frequency content in a 30 kHz bandwidth (in this case, between 2.705 MHz and 2.735 MHz). However, given the curving shape and changing gradient within the B_0 map, the excited slice varies in thickness throughout the imaging volume (Fig 6-23b). The flattened region of B_0 variation results in the two peaks in the slice thickness at the edge of the z FOV, as shown in the computed plot of slice thickness as a function of z location (Fig 6-23b). This artifact could be superficially addressed with an image normalization procedure.

Because of the Gy gradient strength non-linearity, the $fc = 2.72$ MHz and $fc = 2.60$ MHz slices were observed to have different fields-of-view along the y -axis (FOVy). The phantom ($D=63\text{mm}$) extended 31 pixels and 25 pixels in the two slices, corresponding to an FOVy of 63mm ($fc = 2.72$ MHz slice) and 78mm ($fc = 2.60$ MHz slice) (Fig 6-23c). Using the measured Gy coil field at $z=0$ in these two slices in combination with the known acquisition parameters, the expected FOVys are 59mm ($fc = 2.72$ MHz slice) and 75mm ($fc = 2.60$ MHz slice). These are 6.3% and 3.9% lower than the respective image-derived FOVys, but explain the observed trend. These gradient non-linearity effects could be mitigated by refining the coil design or compensated for in the sequence by adjusting the encoded FOV for each slice. Other options include post-processing approaches which apply a gradient nonlinearity correction [189, 190] or generalized image reconstruction [155] approach.

The image SNR in Fig 6-23 varies from slice to slice. There are likely three reasons for this: (1) a progressive reduction in slice thickness resulting from the steeper B_0 gradient

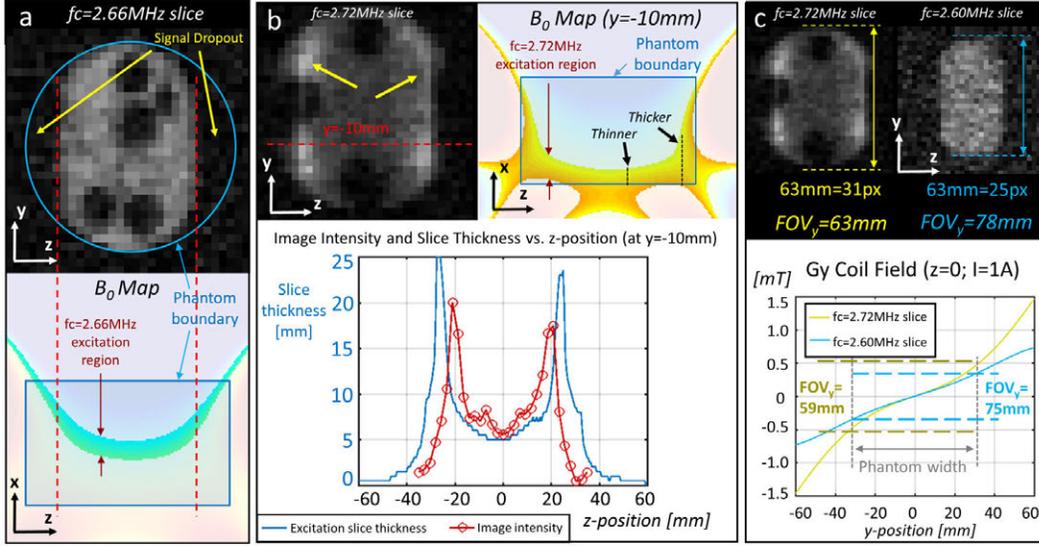


Figure 6-23: a. Signal dropout in multislice images corresponds to regions of high field inhomogeneity and thin slice thickness. b. Hyperintense regions in some images align with thicker regions of the excitation slab. c. Image distortion artifacts correspond to non-uniform gradient strength throughout the ROI.

(i.e., slice-select gradient) at deeper locations; (2) the RF coil provides progressively reduced Rx and Tx capabilities for pixels regions more distant from the RF coil/magnet assembly; and (3) the resonance frequency for the deeper slices begin to exceed the RF coil bandwidth.

As described above, control of the built-in B_0 gradient is critical to providing equal amounts of signal per unit bandwidth at different positions in the ROI. An improved magnet design with higher linearity (but limited gradient strength) would mitigate artifacts. Additionally an RF coil with increased spatial uniformity and coverage would improve images. This would likely involve either a physically larger RF coil or an RF coil with more windings. Adding windings of increasing size boosts the inductance of the coil more quickly than the resistance. Unfortunately, this in turn would increase the Q of the coil and decrease its bandwidth, exacerbating the coil BW issue. Resolving issues stemming from narrow coil bandwidth can be approached by shaping the spectral resonance response of the coil. Several approaches have been proposed for effectively creating a coil with a more uniform frequency response than the typical simple LC circuit without using a series resistor, as we did. These include quasi-transmission line coils [191], coupled resonant structures [192], use of a low-impedance preamplifier [117], and inductively coupled negative feedback mechanisms [193]. However, it must be verified that any employed Q-broadening technique not only creates a wide-bandwidth receive chain, but actually does so in a loss-less way and thus improves the

SNR across that bandwidth.

While spatially non-uniform sensitivity can be problematic, it is routinely encountered and managed in modern high-field MRI scanners using multi-channel Rx arrays [118]. Practitioners have come to appreciate that it is more important to achieve acceptable image SNR over the desired field of view. In summary, this could be done either with improved system hardware or with more general image reconstruction or post-processing schemes.

6.7.3 SNR

The MR-Cap, like other low field MR approaches, has intrinsic SNR limitations. A stronger B_0 magnet in the same form factor is likely achievable by allowing for a higher density of magnetic material or perhaps by adding a second layer or otherwise increasing the cap thickness. A more uniform magnet could likely be realized by more thoroughly exploring the parameter space of possible magnet designs or the use of B_0 shim coils or shim material.

An improved RF coil would also increase SNR. Adding the 1Ω resistor increased the coil 3dB BW from 72kHz to 157kHz, corresponding to an increase in the series resistance of the coil from 0.85Ω to 1.85Ω . This causes a 48% increase in noise standard deviation and thus a 32% drop in SNR. Removing this resistor and increasing the BW with more sophisticated, lossless approaches would improve performance. This coil has intrinsic resistance of 0.85Ω (measured from the unloaded Q) and a measured inductance of $3.5\mu H$. Theoretically, the Fano theory of matching networks suggests that a lossless matching network could achieve a flat, -5.9dB match across 180kHz BW for this coil [156]. It is not clear how to obtain this in practice and the noise figure of the practical network would need to be assessed.

Acquisition improvements could also increase image SNR. Weighting the sampling density to the center of k-space or utilizing sparsity priors such as compressed-sensing type acquisitions or denoising approaches have been used to boost SNR in other low-field systems [194]. A flipback pulse after each RARE train may assist with longitudinal Mz recovery and increase available signal [195]. The spin echoes in the 3D acquisition are temporally very narrow (due to field inhomogeneity), and the acquisition window length could be significantly shortened. This would free up sequence time for more echoes, allowing averaging down the echo train. Such an approach may be impractical at high field due to safety limits on the RF specific absorption rate (SAR), but at 64mT the SAR for such an acquisition is negligible.

6.7.4 Diffusion Effects

The MR Cap B_0 field contains a large built-in Gx gradient (Gx=117 mT/m), which leads to diffusion-induced signal attenuation for the first echo the RARE characterized by the b value: $b_{first} = (\frac{1}{12}\gamma^2 G^2 D \cdot (TE)^3)$. For the first echo of the multislice sequence (TE=3ms), $b_{first} = 2.2 \frac{s}{mm^2}$ for the TE = 3ms used resulting in a negligible (0.5%) signal loss to diffusion for the center of k-space for this sequence. For the nth echo of a RARE sequence, $b = b_{first} * n$ (58) and larger diffusion weightings can be obtained by either increasing the echospacing or re-ordering kspace to sample k=0 at a later echo. Thus, using a RARE readout train allows one to vary diffusion weighting independently of TE [196], and has been explored as a way to obtain diffusion-weighted images in inhomogeneous fields [197].

6.7.5 Practical Considerations

There are several practical requirements to overcome before the MR Cap could be deployed for point-of-care clinical care. First, the multislice phantom images were acquired in a shielded box and thus did not demonstrate in vivo imaging in unshielded environment. Practical point-of-care use would require either a passive shielding approach such as draped conductive cloth, or an active interference cancellation system [179, 180]. Permanent magnet systems are notorious for temperature induced drift in B_0 . While this was not controlled here, standard approaches include either a feedback system controlling a heater to stabilize the temperature, use of a combination of rare-earth materials with differing temperature coefficients [157], or the use of a field probe to measure B_0 drift for incorporation into a model-based image reconstruction algorithm (13). Finally, performance of the MR Cap for scanning a head-sized object that is larger than the target FOV must be assessed.

Additionally, while the MR Cap itself, T/R switch, GPAs, and preamplifier are all highly compact and cost effective, several high-cost, laboratory grade instruments were used. For example, the RF power amplifier (cost \$15000 USD) and console (\$35000 USD) used in these experiments were not optimized for cost or portability. The 2 kW RF amplifier used was larger than needed and could also likely be reduced in size by using a narrow-band and lower power design. For example, it could potentially be replaced by a previously published low-cost LDMOS amplifier capable of achieving 1.4kW at a cost of \$1500 USD [198]. Likewise, there has been recent effort towards compact, low-cost MRI consoles [199, 200, 201, 202, 203]

that could be a great cost-savings, such as the \$500 FPGA based open-source system [200]. A total bill of materials cost using these two low-cost options is provided in Table 6.5, detailing a potential total system parts cost of \$3350 USD.

Table 6.5: Estimated system cost (References go with MR Cap paper citations)

Component	Cost
NdFeB Material	\$450
RFPA (63)	\$1500
GPA (39)	\$300
Console (65)	\$500
Gradient Coils	\$200
RF Coil	\$100
Preamp	\$300
TOTAL:	\$3350

The MR Cap was designed to be highly portable and lightweight capable of use with minimal patient movement or transport. Realizing such a light-weight device required a reduced imaging FoV and a highly-inhomogeneous main magnetic field. In contrast, a head-sized cylindrical magnet could achieve a larger imaging FoV and greater field homogeneity, and would potentially have greater imaging capabilities. However, our previous work (13,34) suggests that this approach results in a heavier magnet (roughly 10x). Additionally, like conventional systems, it would require moving the patient’s head into the bore. We do not attempt to break this trade-off, but simply demonstrate a system at the extreme edge of this spectrum. Ultimately, if the full spectrum of scanner tradeoffs is available, each usage setting and clinical application will decide the appropriate scanner for the job.

6.7.6 Potential Applications

Possible clinical applications for the MR Cap include point-of-care assessment of cerebral hemorrhage such as subdural hematomas in trauma [7] or post-surgery [8]. In the latter, we envision potentially using the device as a monitoring instrument during post-surgical recovery in a neuro-ICU to provide early warning of hemorrhage. Also it would be nice if a PoC system could rule out hemorrhage in potential stroke, allowing earlier rTPA administration (i.e. by EMT prior to transport to a hospital). This latter application will require a clinical

trial to make sure we can see hemorrhages, especially deep pathology without T2* contrast (which is difficult at low field.) Other areas where a PoC MR device might be useful is in the neonatal ICU, where transport out of the facility is undesired. While small-footprint scanners that can be sited in the neonatal ICU are on the market, it might be beneficial to have a device that can provide diagnostic information such cerebral hemorrhage assessment without removing the neonate from the isolette.

In addition to extending MRI applications to PoC sites, significantly reducing the equipment and siting costs of an MRI scanner might extend the impact of MRI diagnostic techniques, even if the scanner had a narrow clinical focus. As discussed in the Introduction, there is precedent for this in other imaging modalities, and there are large sections of the Global South and Global North with relatively sparse access to MRI (6).

From a physics point of view, many of these clinical applications would require further refinement, including a PoC usable, flexible, conductive fabric shielding or active RF interference cancellation systems. Sequence requirements include the ability to obtain T2, FLAIR T2, proton density and possibly diffusion contrast, all of which are accessible to a RARE sequence. Certain types of sequences are not well suited to the built-in gradient of the system, including most of the gradient-echo class of sequences (further handicapping T2* imaging). These "RARE sequence accessible" contrasts can also, of course, be obtained with reduced image encoding, such as depth profile imaging. Here, reduced image information is traded for PoC simplicity and speed to answer a simple "yes/no" clinical question in the field, such as: "does the accident victim have a bleed?"

Finally, the built-in B_0 gradient field of 117mT/m is larger than the gradient amplitudes achievable on existing clinical scanners. Nonetheless, the RARE echo train can be arranged to reduce diffusion weighting, and the porous media community has developed pulse sequences to probe diffusion contrast even in the presence of strong static field gradients by modulating the RARE acquisition [204]. This method was recently evaluated for applicability to in vivo brain imaging [197].

Chapter 7

The "Head-Optimized MRI" System for Whole-brain Imaging

7.1 Abstract

This chapter describes the design of a "Head-Optimized" MRI scanner - a portable MRI system for imaging the whole head. The system was designed on a "head-shaped" hybrid cylinder/sphere geometry to best contour to the shape of the human head. System components (magnets, coils) were optimized for performance within an explicitly head-shaped ROI. It included a magnet optimized for field homogeneity, three gradient coils (one for readout encoding, two for phase encoding), and a transverse-polarized RF Tx/Rx coil. This allowed for increased flexibility in sequence and image encoding scheme design. The constructed Head-Optimized MRI prototype is under 40cm in size along linear dimensions (including magnet, gradient coils, shims, and RF coil), weighs under 50 kg, has a mean B_0 of 72.5mT. Phantom images are presented which demonstrate the ability of the Head-Optimized MRI to image a head-sized and -shaped object in reasonable time.

7.2 System-Level Design

The Head-Optimized MRI is a portable MRI system for whole-brain imaging of adults. During a scan, the patient lies on a bed with their head resting in the RF coil. Then, the rest of the system - including the B_0 magnet, RF shield, and gradient coils - is slid over the coil and patient. The patient's head is now aligned with the imaging region of the scanner,

and anatomical images can be acquired. Fig 7-1 shows the evolution of this system from initial concept to completed, working prototype.

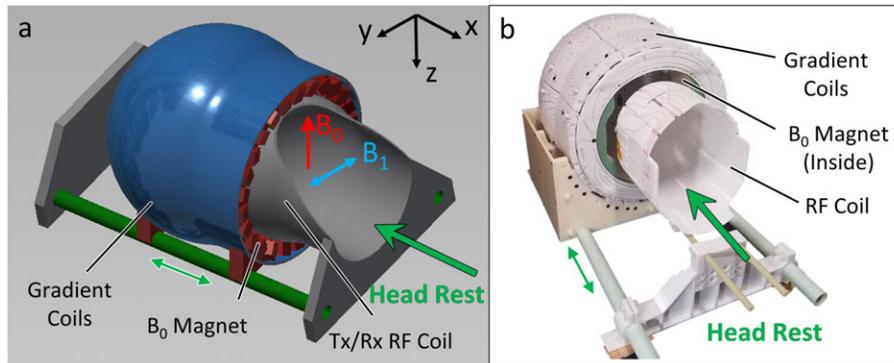


Figure 7-1: a. The concept for the Head-Optimized MRI, showing the essential system components. The patient lies with their head in the RF coil, which the scanner then slides over (analogously to close-fitting high-field transmit RF coils). b. The constructed prototype, showing the essential components.

7.2.1 Imaging ROI

The Head-Optimized MRI was designed for imaging performance specifically within a head-shaped ROI, as determined from a high-field scan of an adult head (Fig 7-2). The ROI was first demarcated by a line running from the eyeballs to the caudal end of the medulla. It was then padded in the foot ($+\hat{x}$) direction by 14mm to account for patients with shorter necks.

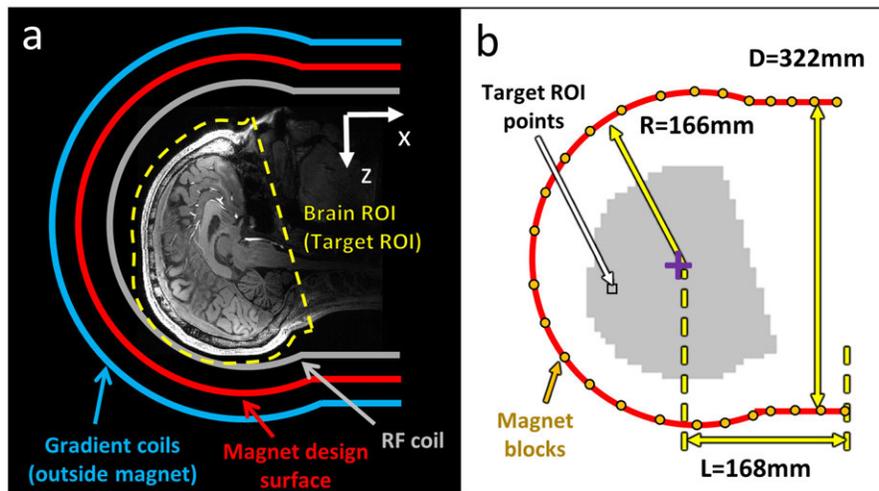


Figure 7-2: a. The target ROI was determined from a high-field scan on an adult brain, and overlaid all anatomy above a plane running approximately from the eyeballs to the caudal medulla. b. The magnet was designed on a spliced sphere+cylinder surface.

7.2.2 System Geometry

The distance from the foot-directed ($+\hat{x}$) limit of the magnet to isocenter was chosen to be 18cm based on the high-field anatomical scan. This determined the maximum extent of the magnet arrays and coils in the system. The system geometry was based on a symmetric hybrid sphere+cylinder form (Fig 7-3) [205] to match the bulbar shape the human head and neck. The system extended inward near the $+\hat{y}$ and $-\hat{y}$ limits of the bore to allow more space for "booster" B_0 magnets; the RF coil similarly was narrower along its y -axis than along its z -axis (Fig 7-3). This choice was possible because the human head is generally narrower along the LR axis than the AP axis (aligned with the scanner \hat{y} and \hat{z} axes, respectively). Linear scanner dimensions are also shown in Fig 7-3; the bore opening measures 22.4cm by 27.2cm ; the outer cylindrical and spherical diameters are 37.6 and 38.6cm ; the length overall (LOA) is 38.7cm .

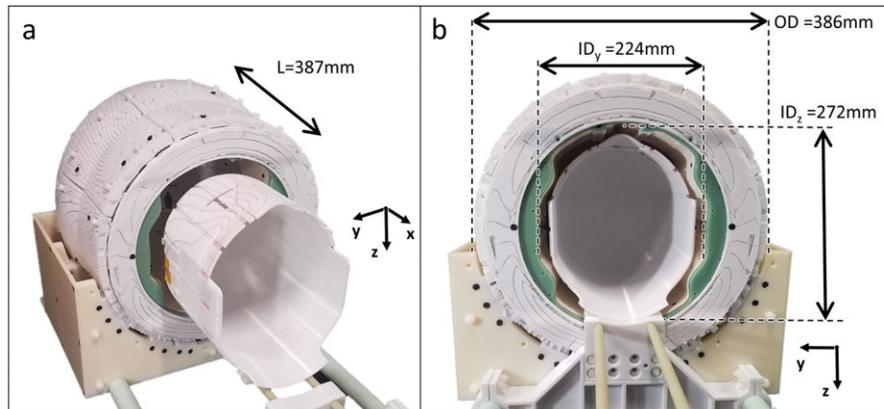


Figure 7-3: a. Isometric view of the scanner with overall length labeled. b. Front view of the scanner showing the major (z) and minor (y) bore diameters and the overall system outer diameter.

The order of system components in the bore (see also Sec 2.3.3) went, from innermost to outermost:

- (innermost) RF Tx/Rx coil
- B_0 shim magnets
- RF shield
- The main B_0 magnet
- (outermost) gradient coils

The system had 3 gradient coils for encoding along the \hat{x} , \hat{y} , and \hat{z} axes. During imaging acquisitions, one would be used as a readout gradient while the other two would perform phase encoding. The gradient coils were placed outside the B_0 magnet, as with the MR Cap (Ch 6). The use of three switched gradients provides greater image encoding flexibility than using a fixed readout gradient.

7.2.3 Performance Goals

The Helmet MRI was designed to keep the total spin bandwidth across the imaging ROI as low as possible. Reducing total spin bandwidth allows for: shorter or lower-power RF pulses; less image distortion/better-conditioned reconstruction problems; and greater flexibility in image encoding (see also Sec 2.2.3). It also enables the use of higher-Q RF coils. Based on prior experience, a min-max field homogeneity of 1% was the design target (allowing for a $Q = 100$ RF coil).

7.2.4 Mechanical Design

Fig 7-4 shows the mechanical design for the Helmet MRI. Emphasized are aspects of the design that deal with the interfacing/mating of different system components. The system contains a G10 fiberglass ring for structural reinforcement that screws directly into the B_0 former. The gradient coil former attaches directly to the outside of the B_0 former with screws. Shim trays with the RF shield on the outside are similarly are screwed into the inside of the B_0 former. The whole system is mounted on a "cradle" that allows it to slide of a pair of G10 tubes which serve as rails. The G10 rails are mounted to a 3D-printed fixed base, which also holds the RF coil. The magnet+gradient+shim assembly slides underneath the RF coil, as a bearing assembling using G10 tubes as rails supports the RF coil. The RF coil can itself slide on these G10 rails, allowing the exact position of the RF coil to be adjusted independent of the base or the magnet/gradient assembly.

Wire feeds are needed for: the RF coil, the RF shield (for grounding), and the gradient coils. The gradient coil lines feed the coils at the back of the magnet, while the RF and RF shield lines run through a hole in the gradient coils, magnet, and shim trays to access the coil/shield. A cart was designed to hold the whole system and to make with an existing patient table [59], enabling human subjects to be scanned.

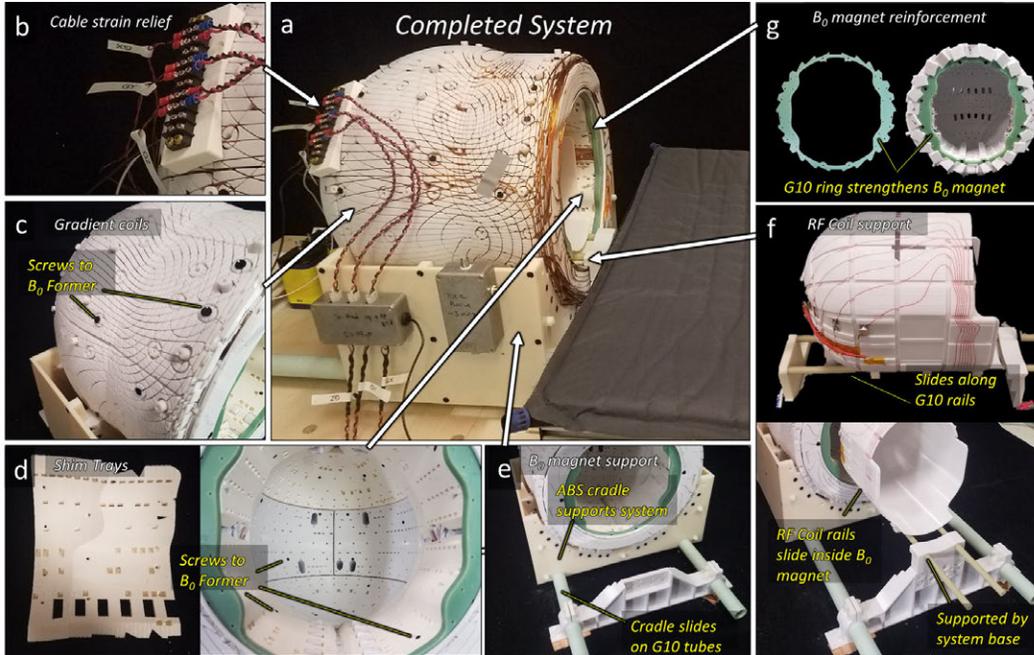


Figure 7-4: a. (Top center) Overall view of the constructed Brain-Optimized MRI scanner with mechanical elements labeled. b. A terminal block provided strain relief for gradient and RF feed lines. c. The gradient coil former screwed into the B_0 former with M4-0.7 screws. d. The shim trays fit inside the B_0 magnet and attached with M3-0.5 screws. e. A machined ABS cradle supported the B_0 magnet assembly and allowed it to slide on a pair of OD=1.25" G10 tubes. f. The RF coil slides on two 7/16" G10 tubes, which are attached to the main system base. These G10 tubes also slide atop the inside of the B_0 former, allowing the magnet and RF coil to be independently positioned. g. A machined $t=3/8$ " G10 ring provides structural reinforcement at the bore opening of the B_0 former.

7.3 B_0 Magnet

7.3.1 Magnet Design

The B_0 magnet was designed for homogeneity over the head-shaped ROI. The design was based on the "test-tube" hybrid Halbach cylinder/sphere design and prior work on close-fitting helmet-shaped RF and shim coils [175], consisted of $N = 326$ discrete blocks. Of these, the "main" blocks numbered 296 and could have a max size of $1'' \times 1'' \times 1'' = 25.4 \times 25.4 \times 25.4\text{mm}$, while the other 30 "booster" blocks were placed near the patient opening and could have a max size of $1'' \times \frac{1}{2}'' \times 1''$ (Fig 7-5). The initial magnet configuration followed the "test-tube" design, whereby all blocks were pointed in the "Halbach" orientation and the cylinder magnets had $\frac{4}{3}$ times the size of the sphere section magnets [205]. This number of magnets was a chosen to pack magnetization densely in the available volume while still making it possible to avoid block self-intersections (see Sec 3.5.1).

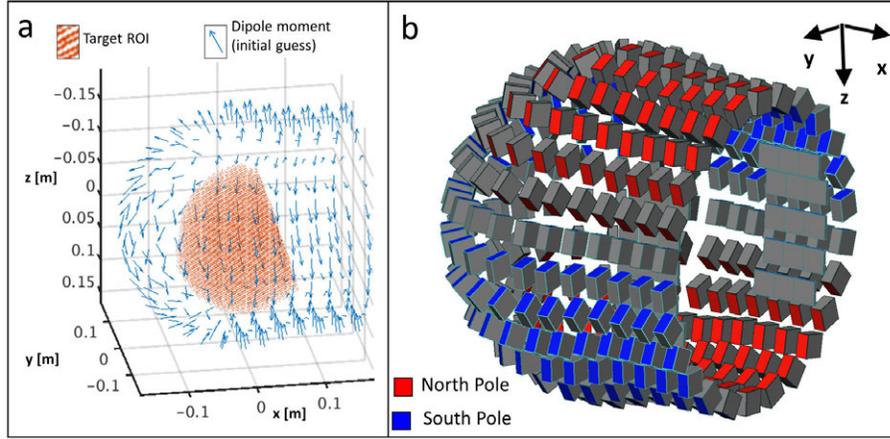


Figure 7-5: a. Dipole vector model of the starting magnet design shown with the target ROI. b. Block model of the starting magnet design.

The design approach allowed all 326 blocks vary in size and angular position. All magnets were N52-NdFeB rectangular blocks magnetized along \hat{z} with size parametrized by their y -dimension (the x and z dimensions were both always equal to $1''$). Angular position was described with two rotation angles. In total, this allowed the dipole moment of each block to vary freely in three dimensions (subject to max block size constraints) but fixed the third Euler angle of each block to be zero. This choice of magnet block design and parametrization was chosen to simplify their manufacturing. The magnet optimization assumed block size and position angle could vary continuously; block sizes would be rounded to a discrete set as a post-processing step, with the choice of that set to be determined at that time. All block translational positions were fixed.

The continuous-design-space magnet optimization was performed using the interior-point method with the Matlab `fmincon` tool. The optimization used a min-max field homogeneity cost function and constrained $mean(B_0) \geq 70mT$. The magnetic field computation modeled the blocks as $\ell = 5$ multipole sources and used the field computation procedure described in Sec 3.4.5. The final design was simulated in Comsol and was found to match the in-optimization field computation. Fig 7-6 shows the optimized magnet design. The block sizes were then rounded to the closest of 13 different discrete block sizes to simplify purchasing the magnet material. Fig 7-7a shows the distribution of magnet blocks over these 13 discrete sizes. Fig 7-7b shows the cumulative sum of this distribution, along with the cumulative sum of the continuous-space magnet design. The discretization step incurred an RMS block size error of 0.30mm.

Fig 7-8 shows the field maps for the discrete and continuous-space magnet designs as simulated in Comsol. The discrete magnet had a mean B_0 of 70.0 mT with a range of 1.02 mT across the ROI (1.46% homogeneity). The continuous-space magnet had mean B_0 of 69.9 mT with an ROI range of 0.76mT (1.09%).

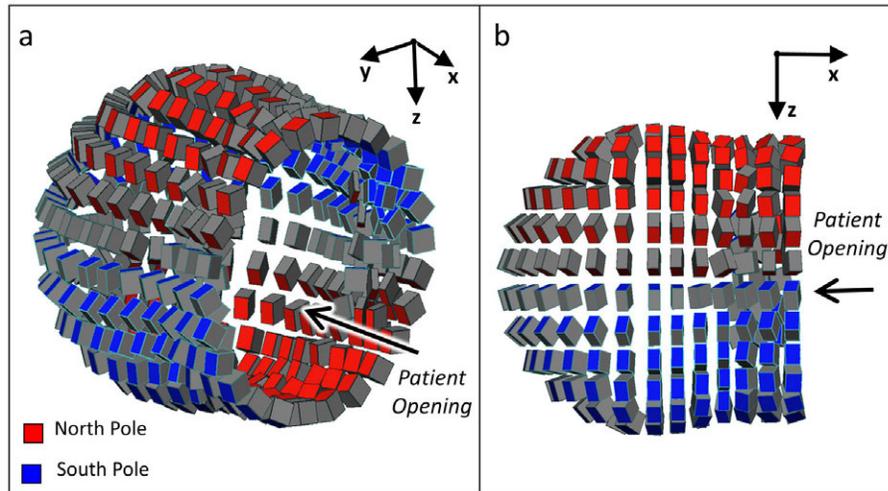


Figure 7-6: a. Isometric block model rendering of the numerically-optimized magnet design. b. Side view (XZ plane view) of the optimized magnet design.

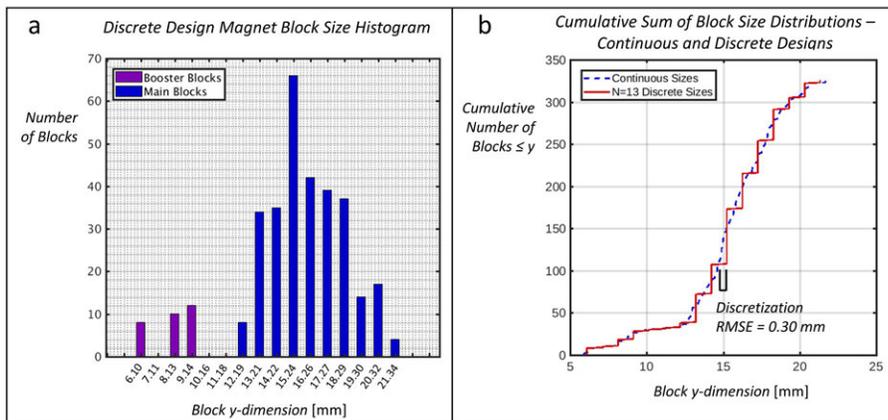


Figure 7-7: a. Bar plot showing the distribution of magnet blocks with all available y-dimensions. b. Cumulative sums of the distribution in a and the continuous-space design distribution. The RMS block size error between the two was 0.30mm.

7.3.2 Systematic and Random Design Error Simulations

This design was then analyzed with added perturbations to simulate the effects of systematic manufacturing errors and tolerances. Simulations were performed looking at global

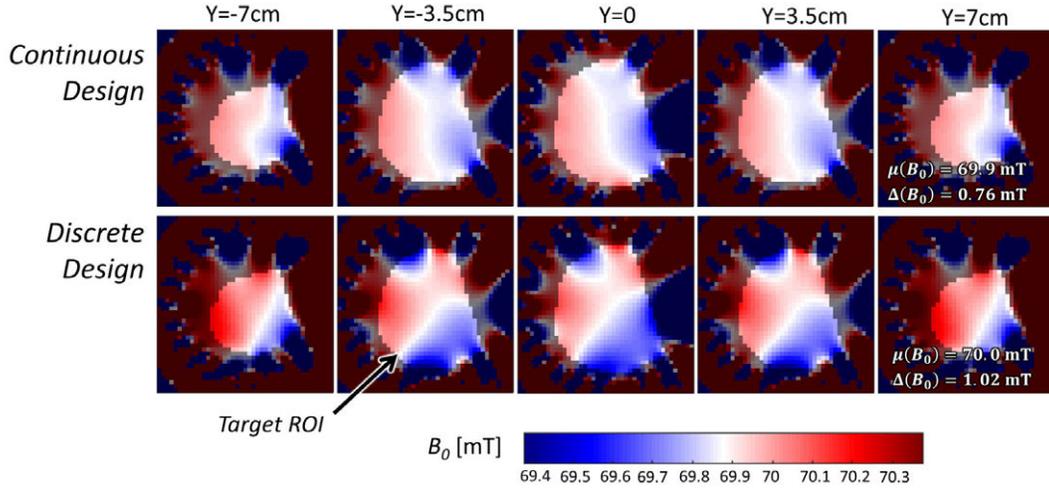


Figure 7-8: Simulated B_0 fields in the XZ plane for the continuous-space and discrete-space magnet designs.

increases and decreases in B_0 and μ_r , and with magnet position errors arising from tolerances in the magnet quadrant mates (see Sec 7.3.3). Simulations were performed with $B_r = \{1.40T, 1.44T\}$, with $\mu_r = \{1.04, 1.06\}$, and with mate offset errors of $\{-0.4mm, +0.4mm\}$. The unperturbed design used $B_r = 1.42T$, $\mu_r = 1.05$, and no mate offset error. Fig 7-9 shows simulated the simulated field maps for all cases along with the computed mean field and field range over the ROI. The mean B_0 varied between 69.0 mT and 71.0 mT across all perturbed designs; the ROI B_0 range varied between 1.04 mT and 1.13 mT (compared with 70.0 mT and 1.02 mT, respectively, for the unperturbed design).

Monte-Carlo simulations were also performed to measure the effect of random errors in block size on field performance. Note that block size errors, B_r errors, and μ_r errors all affect the block dipole moment, which is the primary determinant of the spatial field map. Thus, it was expected that this Monte-Carlo simulation could serve as proxy for any source of random error from block to block. To perform the tests, a random error for each block's y -dimension was chosen from a uniform distribution on $[-1mm, +1mm]$. Fig 7-10 shows the unperturbed field map along with the field map from a design with random block errors. The mean field in this case did not significantly change (70.0 mT vs. 70.1 mT), but the ROI range increased from 1.02 mT to 2.18 mT. This design was used as a test case to assess the performance of the B_0 shimming procedure prior to constructing the B_0 magnet (Sec 7.4.3).

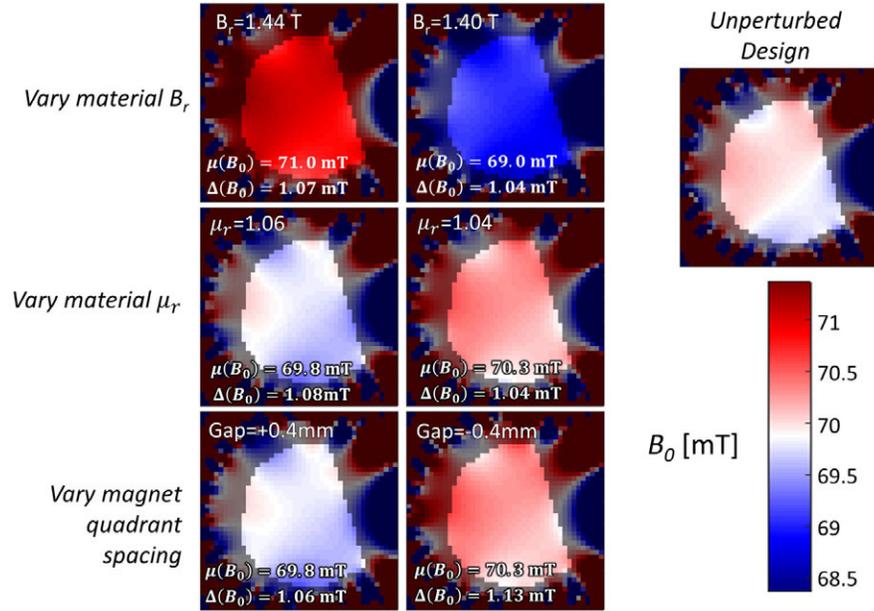


Figure 7-9: Simulated B_0 field maps (XZ plane) for perturbed magnet designs with different B_r , μ_r , and magnet quadrant gap size. On the right is the unperturbed design field map for comparison.

7.3.3 Magnet Construction

The magnet former was 3D-printed in four quadrants out of PC-10 polycarbonate material on a Fortus 360mc printer. The pieces were deburred and about 500 holes (for shim tray and gradient coil attachment) were widened and threaded. The joints between the quadrants had a spline-type mate, contained reinforcement elements, screwed together, and were bonded with Loctite EA-E60 epoxy. The reinforcement fiberglass front ring (see Fig 7-4) was screwed in place prior to populating the former with magnets (Fig 7-11).

The magnet former was then populated with magnets. The booster magnets were glued in first, as shown in Fig 7-12. Magnets were then glued into the main outer portion of the magnet former (Fig 7-13). Fig 7-14 shows the completed magnet.

7.3.4 Magnet Field Mapping

The assembled magnet was the mapped with a 3-axis Hall probe (Metrolab THM1176) stepped by a CNC robot. The field map was acquired with an isotropic 7mm resolution. Figure 7-15 shows the measured field maps. Note that the design of the robot arm did not allow for every point in the design ROI to be measured - in particular, the occipetal region of the head could not be reached. The measured field was therefore linearly extrapolated into

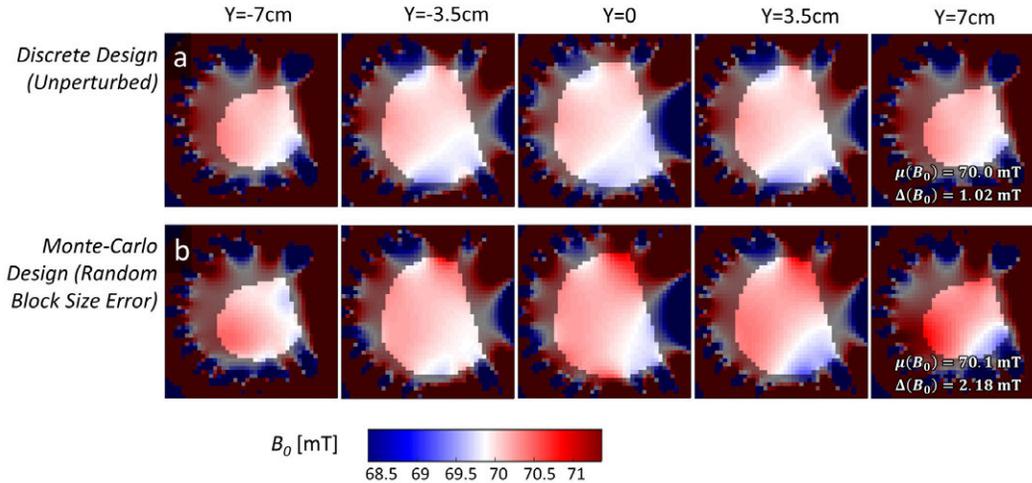


Figure 7-10: Simulated B_0 field maps (XZ plane) for the unperturbed design and a Monte-Carlo-perturbed design.

these points in order to enable comparison with simulation. The measured, extrapolated B_0 had a mean of 71.6 mT and ROI range of 3.08 mT, compared with 70.0 mT and 1.02 mT for the ideal simulated magnet. The constructed magnet field errors were positive near the anterior and posterior poles of the head ROI, and negative near the left and right ends of the ROI.

7.4 B_0 Shimming

7.4.1 Magnet Shim Array Design

A set of shim magnets was designed to fit inside the main B_0 magnet and to improve field homogeneity over the target ROI. The shim array design allowed for up to $N = 889$ individual magnets distributed over 14 horseshoe-shaped arcs situated inside the magnet (Fig 7-16). The distance between the centers of adjacent shim magnets was 10 mm - this allowed for a $\frac{1}{4}$ "-deep magnet at each position with a 3.65 mm gap on each side. The shim magnets were about 4cm closer to the imaging ROI than the main magnet blocks (Fig 7-16). The exact positions of the shims were chosen as a compromise between maximizing interior bore space while avoiding intersections with the B_0 magnets, and were only be determined with certainty once the B_0 magnet had been optimized and all its blocks fixed.

Once the positions were fixed, the shim blocks were not allowed to move either translationally or rotationally. The angular orientation of each block was specified by the Halbach

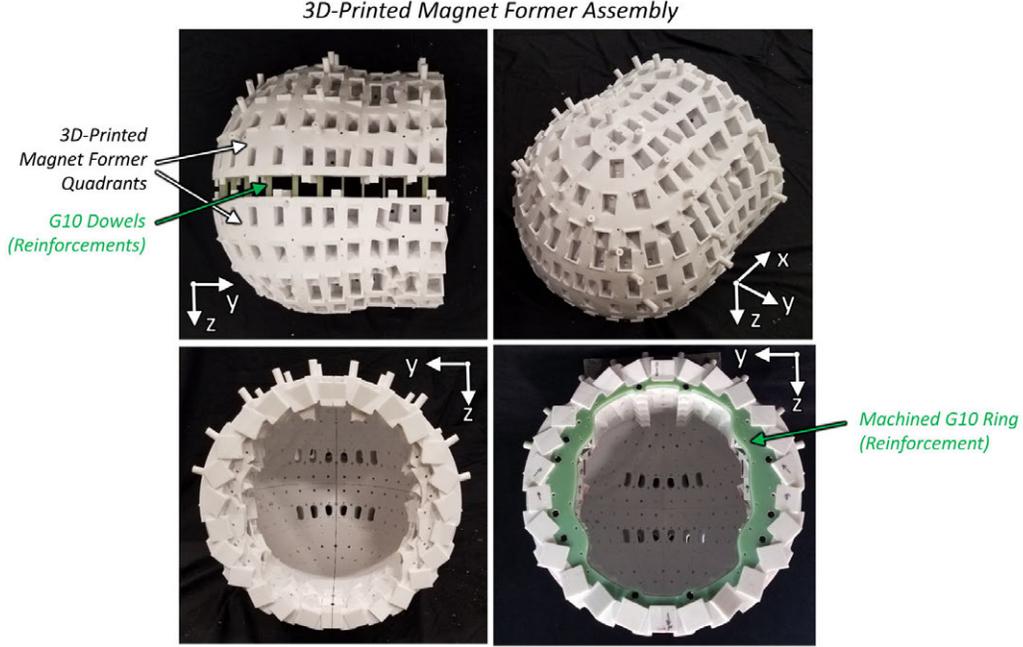


Figure 7-11: The magnet former consisted of four 3D-printed quadrants. These were bonded together with G10 dowels as structural reinforcements. The G10 reinforcement ring was screwed into place in the bore opening prior to populating the former with magnets.

magnet relation [162] $\theta_{rot} = 2 * \theta_{pos}$, where θ_{rot} is the polar angle describing block orientation and θ_{pos} is the polar angle describing block position. Each shim location could take one of two block options: it could be empty (no magnetization); or it could contain a $\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8}$ " N45-grade NdFeB block. This block option was chosen because our magnet supplier (Applied Magnets) had it in stock, and it met all geometric requirements.

7.4.2 Shim Magnet Optimization

The set of shim blocks was optimized using a genetic optimization (Matlab `ga` toolbox) on account of the discrete design space. The optimization allowed for three block options: empty/no magnetization; a $\frac{1}{4} \times \frac{1}{4} \times \frac{1}{8}$ " N45-grade NdFeB block; and an intermediate $\frac{1}{4} \times \frac{1}{8} \times \frac{1}{8}$ " N45-grade NdFeB block. The half-sized block option was not used in the final design, only as an optimization trick (see Sec 3.5.2). The optimization cost function was the absolute field homogeneity over the ROI for the sum of the shim magnet field and the measured, extrapolated B_0 field:

$$f_{cost} = \max(D_1 M_{sh} + B_{0,ext}) - \min(D_1 M_{sh} + B_{0,ext}) \quad (7.1)$$

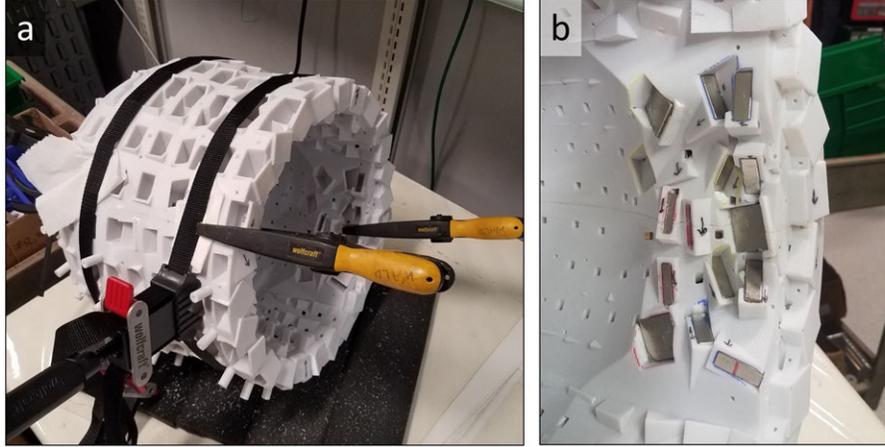


Figure 7-12: a. The booster magnets were the first ones glued into place. Clamps were used where possible. b. Set of 15 booster magnets (ie one side's worth) bonded into the former.

D_1 is the field computation matrix; M_{sh} is the set of shim magnetic dipole moments; and $B_{0,ext}$ is the measured, extrapolated B_0 field. The shim magnets were modeled as point dipole sources since they were small (linear dimension = $6.35mm$). The optimization used no constraints: the constraints on magnet size were implicit in the set of block options available to the genetic optimization, and it was assumed that the shims would not appreciably affect the mean B_0 field strength. 25 iterations of the ga optimization were performed and solution with the lowest cost function among them was chosen as the final design.

7.4.3 Shim Design Procedure Tests

The aforementioned design and optimization procedure was tested on simulated B_0 magnet designs prior to actually constructing the main B_0 magnet. A shim magnet array was designed for the simulated B_0 field produced by the optimized magnet with random Monte-Carlo block size errors (Sec 7.3.2). Fig 7-17 shows the shimmed and unshimmed Monte-carlo-magnet field maps along with the optimized shim array. The shim array reduced field variation over the ROI from 2.18 mT to 0.667 mT, while mean B_0 increased from 70.1 mT to 72.1 mT. This provided evidence that the proposed shimming paradigm would in fact work as designed.

7.4.4 The Optimized Shim Array

A shim array was designed for the actual B_0 magnet once it had been constructed and field mapped. In simulation, the shim array decreased ROI field inhomogeneity from 3.08 mT to



Figure 7-13: a. Magnets were populated into the outer surface of the former starting at the bore opening. Clamps were used when mechanical access allowed. b. Magnets were then added to the remainder of the surface, generally using fingers and thumbs as clamping elements.

0.72 mT and increased the mean field from 71.6 mT to 72.9 mT (Fig 7-17). The optimized design contained 527 empty slots and $363 \frac{1}{4} \times \frac{1}{4} \times \frac{1}{8}''$ magnets. The empty slot count includes 4 locations to which the optimization assigned the third "fictitious" block option, and which were converted to empty slots.

7.4.5 Shim Array Construction

The shim array was then constructed and inserted in the B_0 magnet. The shim array former consisted of 8 shim plates that contained the magnets and screwed into the main magnet former (Fig 7-18). During CAD design, an STL of the B_0 magnet block boxes was used to etch spaces in the shim trays that allowed them to mate with the B_0 magnet (see Fig 7-18). The shim trays were 3D-printed in ABS or PC-10 polycarbonate on a Dimension or Fortus 360mc printer. The magnets were then inserted and glued in with "instant" cyanoacrylate adhesive (Loctite 414).

The RF shield was added to the outside of the shim trays after magnet assembly. The RF shield used 6 coats of spray-on MG Chemicals silvered-copper-bead EMI shield. The spray-on substance itself reeks of blue cheese for about a day after being applied and should

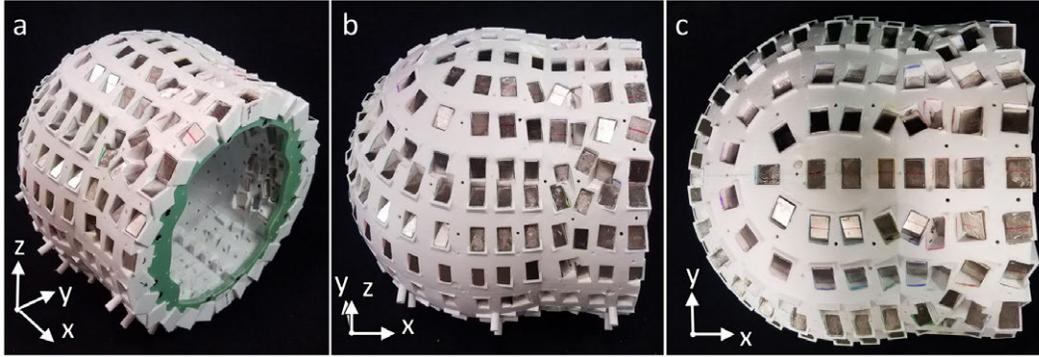


Figure 7-14: a. Isometric view of the completed magnet. b. Diagonal view. c. Top (XY plane) view.

only be used in a facility properly equipped for spraypainting. Six coats were also sprayed on a test piece of 3D-printed ABS plastic to test how well it attenuate the coupling between a pair of untuned loops. The test piece measured $20 \times 20\text{cm}$ and the untuned multi-turn loops had $D = 5\text{cm}$ and were placed 15cm apart. The grounded sprayed-on shield attenuated the coupling by 5.5dB . For comparison, a grounded copper mesh sheet of the same size gave 7.0dB attenuation. The spray was applied to the entire outer surface and to a 1cm -thick edge on the inside of all 8 shim trays (Fig 7-19).

The shim trays were then inserted into the magnet and screwed in place. Copper tape was applied on the conductive, 1cm -thick edges of adjacent shim trays to connect the segments of the RF shield (Fig 7-20). A wire was soldered to an accessible portion of the RF shield and fed through the cable feed port in the back of the magnet (Fig 7-4) to serve as a grounding connection.

7.4.6 Shimmed Magnet Field Mapping

The shimmed magnet was field mapped using the same procedure as used for the unshimmed magnet (Sec 7.3.4). Fig 7-17 shows measured the shimmed magnet field, the simulated shimmed field map, and the measured unshimmed magnet field. The shimmed magnet had an ROI field range of 1.02 mT and a mean of 72.5 mT , compared with 3.08 mT and 71.6 mT for the unshimmed magnet. The constructed shimmed magnet achieved a 67% improvement in field homogeneity, compared with the simulated 77% improvement.

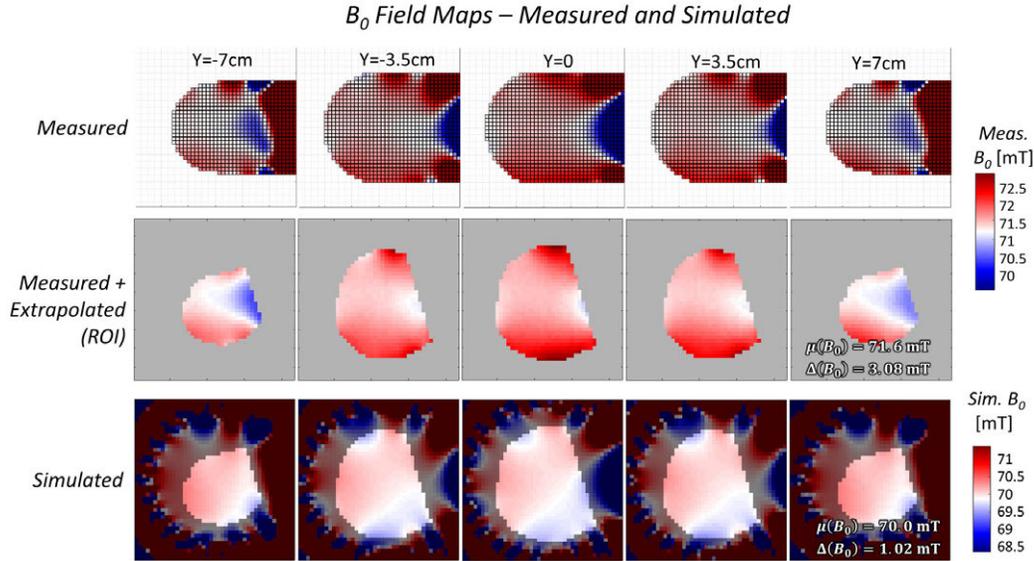


Figure 7-15: Measured B_0 field of the constructed magnet. The top row ("Measured") shows the field at all measured spatial location. The second row ("Measured + Extrapolated (ROI)") shows the field over the ROI with linear extrapolation used in the occipetal pole. The third row ("Simulated") shows the ideal simulated field.

7.5 Gradient Coils

7.5.1 Gradient Coil Geometry

Three gradient coils were designed to fit in a volume outside the B_0 magnet and to provide phase or readout encoding along x, y, and z. The three gradient coils were made with AWG18 enameled wire and sat on top of each other in the order (from in to out) x-y-z. The design surface for the y and z coils was the outer cylindrical+spherical surface of the B_0 magnet; the design surface for the x coil used same surface, plus an additional annulus on the patient-side face of the magnet (Fig 7-21). This design was compared with an x-coil designed without the extra annular flange (Fig 7-22), and was found to increase the field efficiency along the x-axis by 10% and provide the same bijective encoding range.

7.5.2 Gradient Coil Optimization

The three coils were optimized using the target field approach described in Sec 4.4. The target ROI for all optimizations was the imaging target ROI (Sec 7.2.1) and the target fields varied linearly in strength along the each coil's respective spatial encoding coordinate. The optimization used the regularization parameter of $\lambda = 1.0 \cdot 10^{-6}$. Smoothing was also

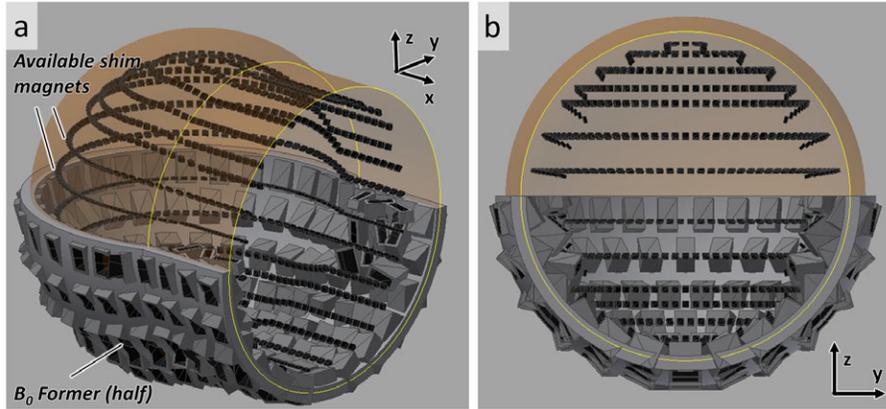


Figure 7-16: a. Isometric view of the available shim locations in relation to the B_0 magnet former. b. Front view (YZ plane).

performed on the stream function to obtain an acceptable wire spacing. The x, y, and z coils all used $N = 30$ level sets to generate windings from the stream functions. Fig 7-23 shows the optimized winding patterns and stream functions for the x, y, and z coils. Figure 7-24 shows the simulated coil field maps. The simulated efficiencies for the x, y, and z coils were: $0.488 \frac{mT}{m \cdot A}$, $0.946 \frac{mT}{m \cdot A}$, and $1.081 \frac{mT}{m \cdot A}$, respectively.

7.5.3 Gradient Coil Construction

A single coil former base was designed in CAD to hold all three coils (Fig 7-25). The former attached to the B_0 magnet with screws that mated with tapped holes in the magnet former. The former was split into 17 pieces in order to be printed: 16 pieces for the sphere+cylinder portion, and 1 for the front ring portion (Fig 7-25). The wire paths were converted into STLs, which were then etched into the gradient former in CAD (see Sec 4.5.2). The B_0 former was also etched into the gradient former to ensure a non-intersecting mate, as was done with the shim trays (Sec 7.4.5). The pieces were printed in PC-10 polycarbonate material on a Fortus 360mc printer. The wire was then press-fit into the grooves (Fig 7-26). Thinned-down adhesive (Loctite EA608 thinned with isopropanol) was painted over wires that weren't held fast in grooves, such as near the patient end of the coils (Fig 7-26). The x, y, and z coils required 84.7, 84.6, and 91.1 meters of wire, respectively.

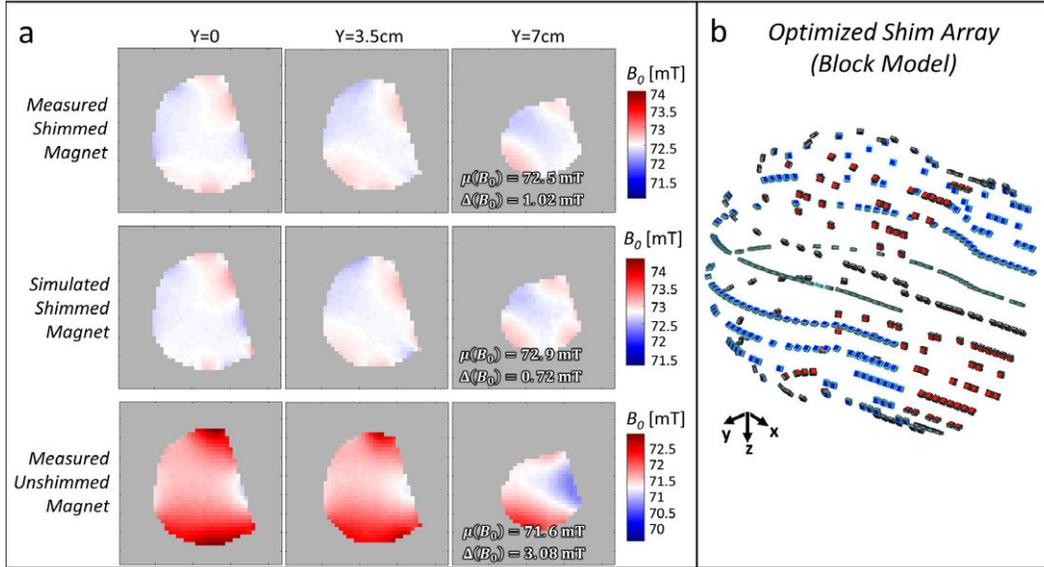


Figure 7-17: a. Simulated field maps (XZ plane) of the shimmed Monte-Carlo magnet design and the unshimmed Monte-Carlo magnet design. b. Block model rendering of the optimized set of shim magnets.

7.5.4 Gradient Coil Measurement and Testing

The coils were mapped using the same setup as the shimmed and unshimmed B_0 magnet. However, the coils were built directly on the B_0 magnet, which made mapping the coils alone impossible. Thus, the field was mapped with current in each coil ($I = 3A$ for the y and z coils; $I = 4A$ for the x coil) and with $I = 0A$ for all coils. The background ($I = 0A$) field was subtracted from the field maps with the different gradient coils energized to obtain field maps for the coils.

Fig 7-27 shows the measured field maps thus obtained. The inaccessible points of the ROI were extrapolated using the same procedure as for the B_0 magnet field maps. Table 7.1 shows the measured inductances, resistances, mean efficiencies, and efficiency ROI ranges for the three coils.

Table 7.1: Gradient coil parameters and performance metrics

	Gx Coil	Gy Coil	Gz Coil
<i>Wire Length</i>	84.7 m	84.6 m	91.1 m
<i>Inductance</i>	566 μH	837 μH	959 μH
<i>Resistance</i>	1.693 Ω	1.816 Ω	1.816 Ω
<i>Isocenter Efficiency</i>	0.275 $\frac{mT}{m \cdot A}$	0.939 $\frac{mT}{m \cdot A}$	1.078 $\frac{mT}{m \cdot A}$

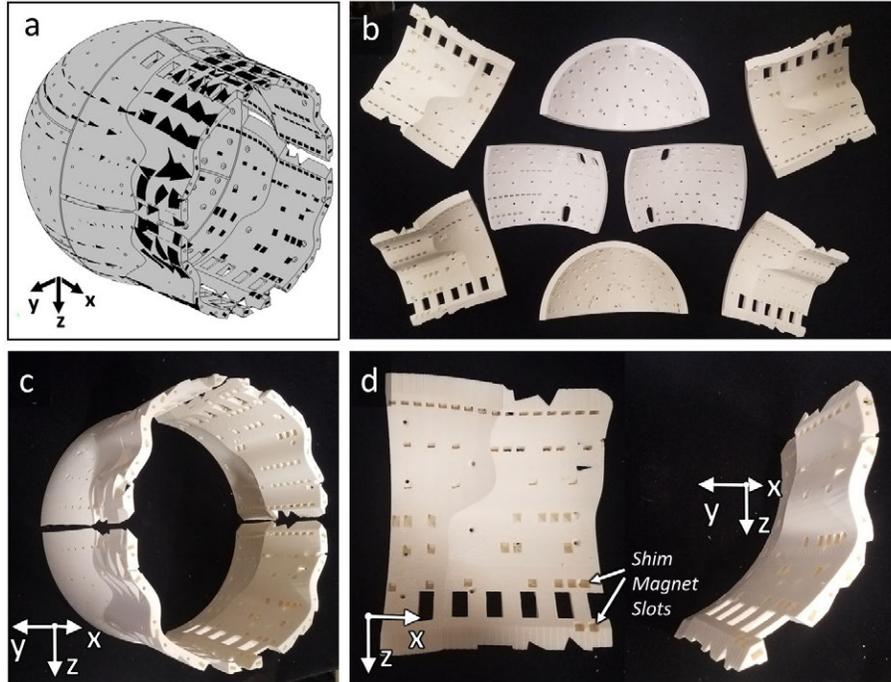


Figure 7-18: a. CAD rendering of the 8-piece shim tray former. b. The 8 3D-printed shim trays. c. The front 4 3D-printed shim trays arranged in their approximate in-bore configuration. d. Close-up of one shim tray showing the shim magnet slots.

7.6 RF Coil and Shielding

7.6.1 Coil Geometry

The RF coil was designed for transmit and receive on a helmet-shaped surface that fit close to the patient's head (Fig 7-28). The coil design optimization used a different target ROI from the magnet and gradient coils. Specifically, the ROI was reduced in size so as to avoid numerical artifacts resulting from the proximity of the target ROI and the stream function surface (see Sec 4.4.3). The ROI also extended further into the patient's neck and was split in to two sub-regions. A uniform B_1 field was desired in the first region (the "Uniform ROI"), while zero field was desired in the second region (the "Null ROI"). Fig 7-28b shows these two ROIs. This was done to reduce the coil sensitivity in the poorly-encoded neck region, thereby reducing the available signal that could alias into the image.

The coil had y-axis linear polarization, resulting in a design that was inherently balanced about the center feed point. The y-polarized design was compared with an x-polarized design produced with the same optimization, stream function surface, and target ROI/field. Fig 7-29 shows: the x- and y-polarized coils; their normalized simulated x-y component field maps

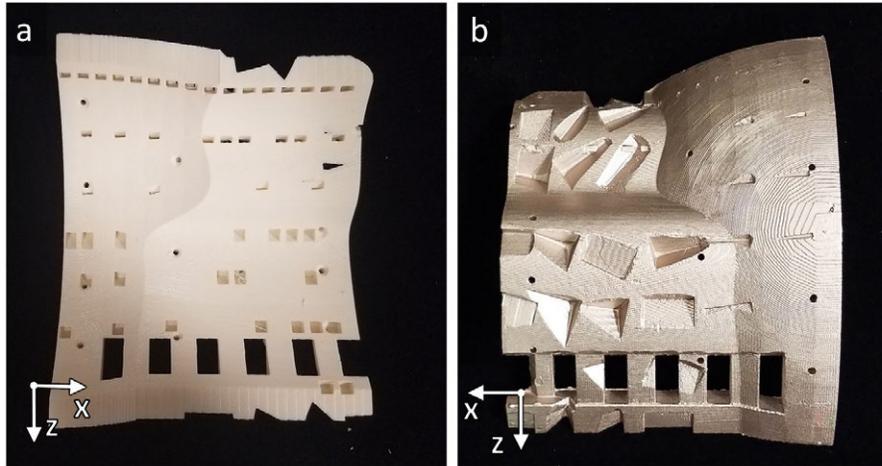


Figure 7-19: a. Inner surface of shim tray (no RF shield). b. Outer surface of shim tray with sprayed-on RF shield.

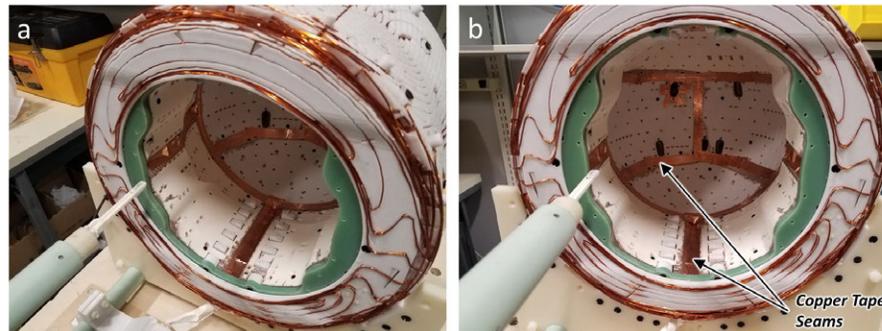


Figure 7-20: a. Isometric view of the bore interior showing the inserted shim trays. b. Front view of the bore interior showing the copper tape seams that connect the sections of the RF shield.

($B_{1,xy} = 1$ at isocenter for each coil); and histograms showing the field in the "uniform" and "null" ROIs. The y-polarized coil field had a mean efficiency of $49.4 \frac{\mu T}{A}$ and a standard deviation of $7.35 \frac{\mu T}{A}$ (14.9%) within the "uniform" ROI, and produced $8.20 \frac{\mu T}{A}$ (17%) in the null ROI. The x-polarized coil had a mean efficiency of $33.7 \frac{\mu T}{A}$ and a standard deviation of $12.8 \frac{\mu T}{A}$ (38%) within the "uniform" ROI, and produced $6.62 \frac{\mu T}{A}$ (20%) in the null ROI. The y-polarized design thus has a higher field efficiency, better uniform ROI field homogeneity, and a lower mean null ROI field.

7.6.2 Coil Design Optimization

Coil optimization used the procedure described in Sec 4.4. The regularization parameter for the stream function optimization was tuned such that the resulting stream function looked smooth, and had $> 1\text{mm}$ wire spacing when discretized with $N = 14$ level sets. Figure 7-30

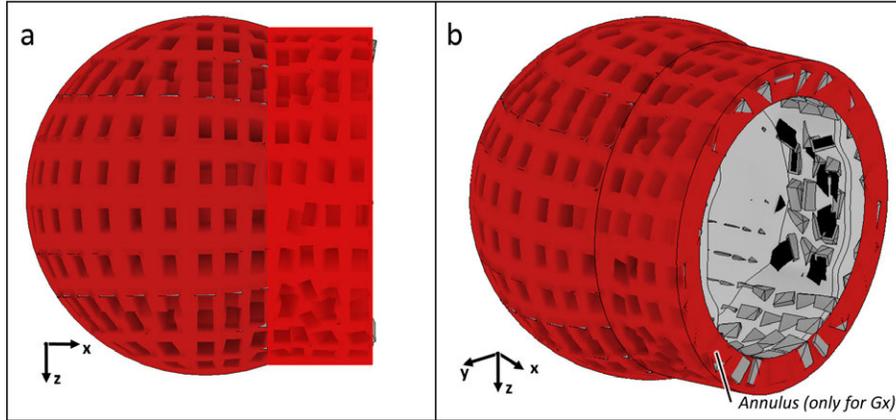


Figure 7-21: a. Side view (XZ plane) of the gradient design surface (red). b. Isometric view of the gradient design surface showing the annular portion used for Gx coil design.

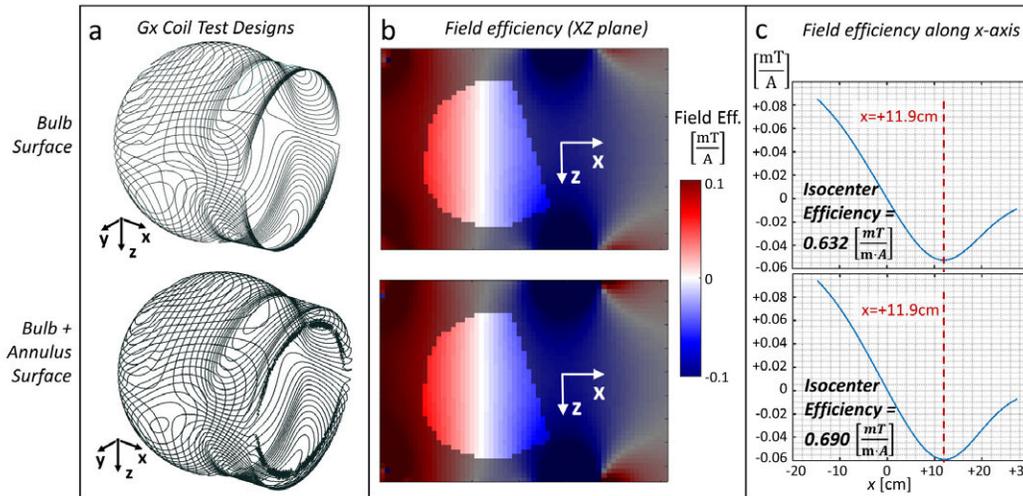


Figure 7-22: a. Gx coil designs created both with and without the extra "annular" section. b. Simulated coil field maps for the two designs. c. Field efficiency plotted along the x-axis for the two designs.

shows the optimized coil stream function and windings.

7.6.3 Coil Construction

The RF coil former used "tiles" with etched grooves as shown in Fig 7-31. The tiles were etched in CAD, 3D-printed in PC-10 (Fortus 360mc), and epoxied onto a helmet-shaped former. AWG36-strand Litz wire was then pressed into the grooves. The coil had a routed circuit board in back for a coaxial cable feed and for tuning/matching circuits (Fig 7-31f). The wound coil inductance was measured at $L = 31\mu H$ at $f = 1kHz$ prior to tuning the coil. The coil winding was then split with four distributed tuning capacitors $C \approx 360pF$ in a

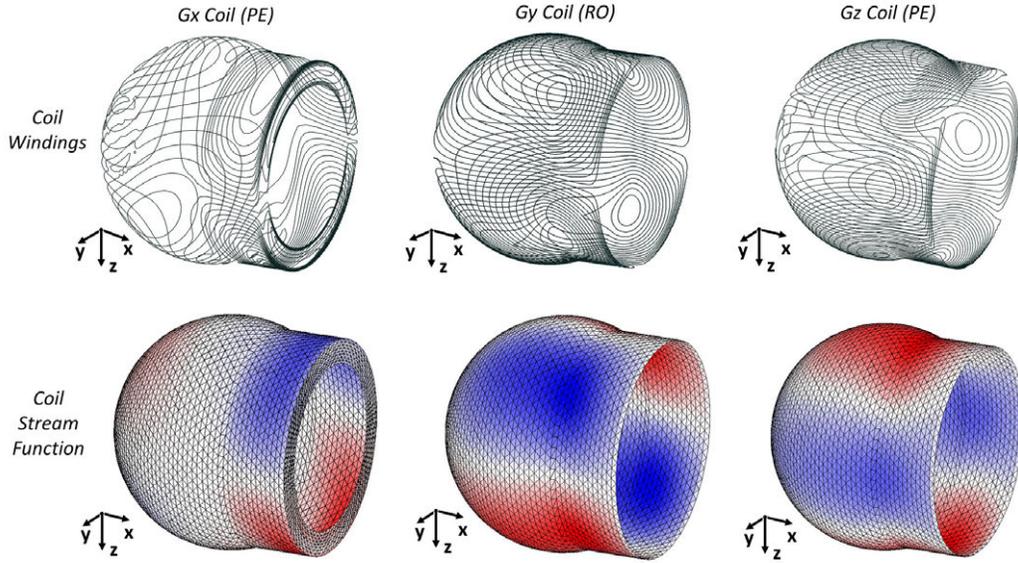


Figure 7-23: Optimized coil winding and stream function designs for the Gx (partion phase encode), Gy (readout), and Gz (in-plane phase encode) coils.

balanced fashion. The coil was then tuned and matched on the feed board to $f = 3.09MHz$ and 50Ω . The coil was initially tuned and matched in free space for field mapping, and was then tuned and matched inside the magnet+RF shield.

7.6.4 Coil Field Mapping

The x- and y-components of B_1 field was measured by using a pair of x- and y- polarized sniffer loops stepped through the imaging ROI inside the RF coil. A VNA measured S_{21} using the sniffer loop with a narrow-band stimulus applied to the tuned RF coil. Note that this was technically a $\frac{dB}{dt}$ measurement and not a B -field measurement; however, the narrow-band measurement makes this derivative measurement approximately proportional to the field itself. The measurement setup was not calibrated for absolute field measurements. Separate measurements were made for the x and y field components, from which was derived the total transverse field component $B_{1,xy}$.

Fig 7-32 shows the simulated, free space measured, and in-bore (shielded) measured $B_{1,xy}$ field maps. Note that the free space and in-bore measurements share the same color scale, even though neither is in physically-meaningful units. The field maps have the same spatial shape, but coil field efficiency drops by a factor of 4 in the bore compared with in free space.

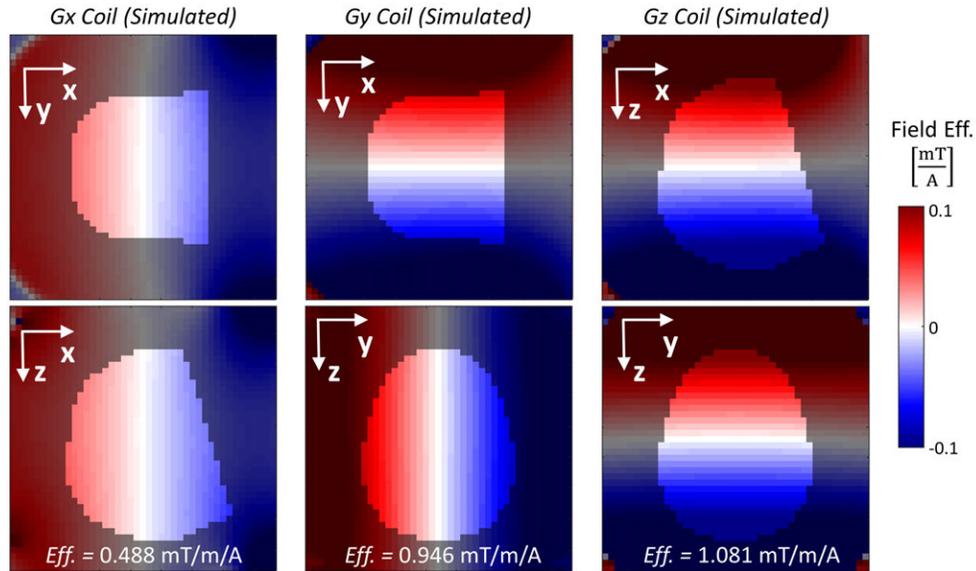


Figure 7-24: Simulated field efficiency (ie mT/A) maps for the Gx, Gy, and Gz coil designs.

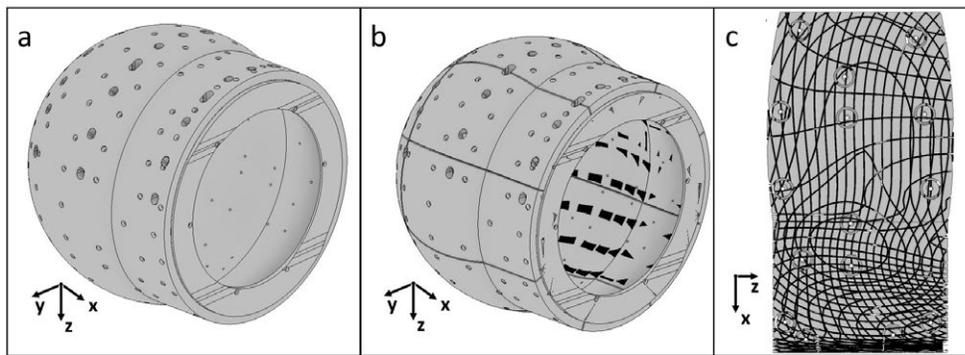


Figure 7-25: a. Gradient coil former CAD model. b. The gradient coil former, having been split into 17 separate pieces. c. The CAD model of one of the 17 pieces showing the wire winding paths that are "etched" into it.

Fig 7-33 shows the S_{21} parameters (magnitude and phase) for the tuned coil inside the shielded bore. The S_{21} parameter was measured by stimulating the coil and recording with an untuned sniffer loop normal to the y-axis at isocenter. The 3dB bandwidth was 433 kHz and the center frequency was 3.105 MHz. The measurement did not actually cover the whole 3dB-BW. The full-width BW was estimate by doubling the measured half-width 3dB-BW.

7.6.5 RF Shielding Tests

Noise spectra were measured as recorded from the RF coil connected through a MITEQ AU-1583 preamp. Measurements were performed outside the magnet on the bench; in the

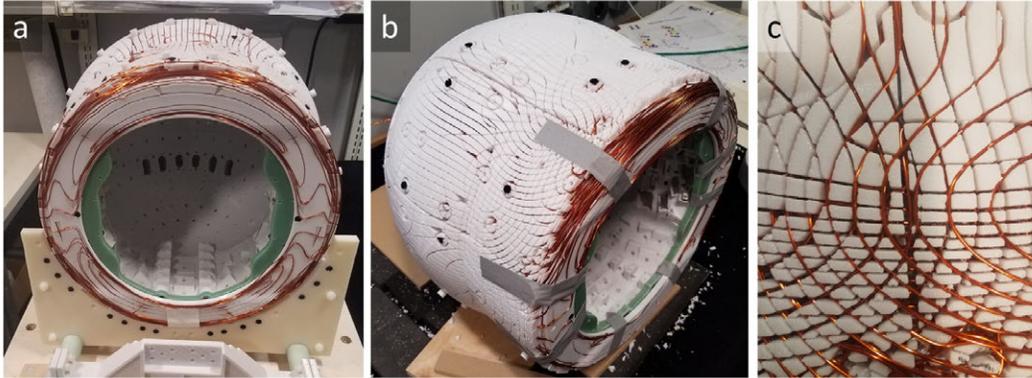


Figure 7-26: a. Front view of the wound gradient coils. b. Isometric view of the wound gradient coils. c. Close-up of the gradient coils showing the wire press-fit into the 3D-printed grooves.

magnet without shim trays or the RF shield; and in the magnet with shim trays and the grounded RF shield. The coil was re-tuned between measurements as the shield altered the tune/match of the coil, and all measurements were acquired at the coil resonant frequency. Fig 7-34 shows the acquired spectra under all three conditions, along with the 50 kHz bandwidth over which total noise was integrated. The noise levels were $-75.46dBm$ and $-63.98dBm$ for the in bore shielded and unshield cases, respectively. The shield thus lowered the noise floor by a factor of $11.48dB$ but decreased efficiency by $\approx 12dB$. This specifically looked at environmental interference, and did not assess how well it shielded noise coupling in on the gradient coils.

7.7 Imaging Sequences and Protocols

7.7.1 Sequences

Imaging tests used a RARE-type sequence with frequency-swept WURST pulses. Two gradient channels were used for phase encoding and the third was used for balanced readout encoding. The first two pairs of spin echos in the sequence did not have any gradient encoding – this allowed them to be used as navigators for phase correction down the echo train for each shot. Fig 7-35 shows this sequence with y selected as the readout direction, and x and z as phase encode axes.

Table 7.2 lists the sequence parameters used for this sequence during the various acquired acquisitions. All sequences used an even number of averages with the shot-to-shot phase cycling scheme shown in Fig 7-35. With this phase cycling scheme, the 180-degree RF

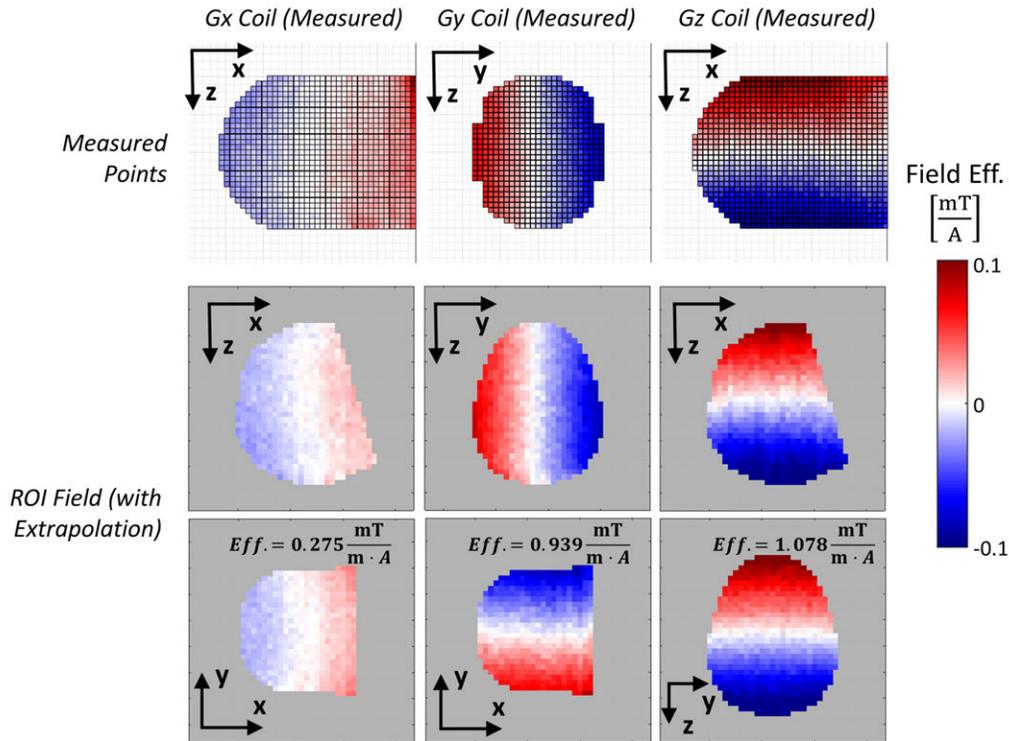


Figure 7-27: Measured gradient coil field efficiency maps (mT/A).

pulses and refocused signals are alternately in phase and out of phase on successive shots. Appropriate averaging thus enables the residual ringdown of the transmit pulse to be mostly cancelled out.

7.7.2 System Setup and Components

All experiments used: a passive, lumped-element T/R switch; two preamps – a MITEQ AU-1583 as the first stage, and a Mini-Circuits ZFL-500LN+ as the second stage; a 12V lantern battery for preamp power; a Tomco BTO2000-AlphaS-3MHz RFPA; a Tecmag Bluestone console for waveform generation and signal detection; two Techron 7224 and one Techron 7548 amplifiers as GPAs; and a shielded room with Schaffner FN 7513-16-M4 100nF capacitor feed throughs for the gradient coil lines.

7.7.3 Imaging Phantoms

In cryptæ images were acquired using an undead "mummy" phantom (aka "King Tut") and a smaller "Red Head" phantom (Fig 7-36a+c). Fig 7-36b+d show 3 T axial T_2 -weighted

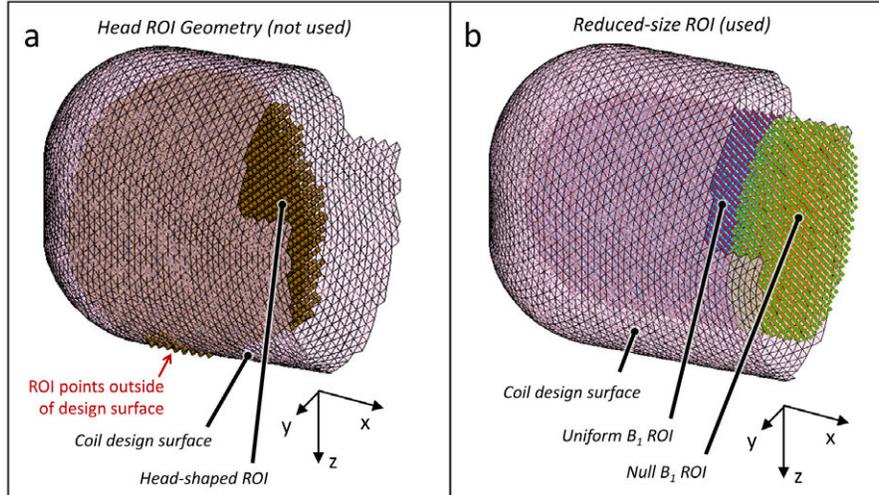


Figure 7-28: a. RF coil design surface shown in relation to the head-shaped ROI. Some ROI points sit outside of the RF coil design surface. b. The RF coil design surface shown with the modified "uniform" and "null" ROIs.

MRI images of both phantoms. The A-P phantom dimensions as measured from these 3T acquisitions were: 192mm (King Tut) and 171mm (Red Head).

7.7.4 Imaging Protocols

Phantom: Vary Readout Direction

Phantom images were acquired using this sequence with different readout directions. Acquisitions were performed using $+y$, $-y$, $+z$, and $-z$ as the readout direction and the parameters shown in Table 7.2. Acquisitions were reconstructed individually with FFT. k-space data were cropped along the readout direction to 256 points (RO_y) and 506 points (RO_z).

Phantom: Vary Readout Strength

Phantom images were acquired with $+y$ as the readout axis, and with readout gradient strengths of 8.9, 14.2, 17.8, and 21.4 $\frac{mT}{m}$. Sequence parameters are shown in Table 7.2. Acquisitions were reconstructed individually with FFT. k-space data were cropped along the readout axis to $N=128$ points. Partial k-space was used along the partition phase encode direction ($N_{PE} = 25$ acq/35 recon).

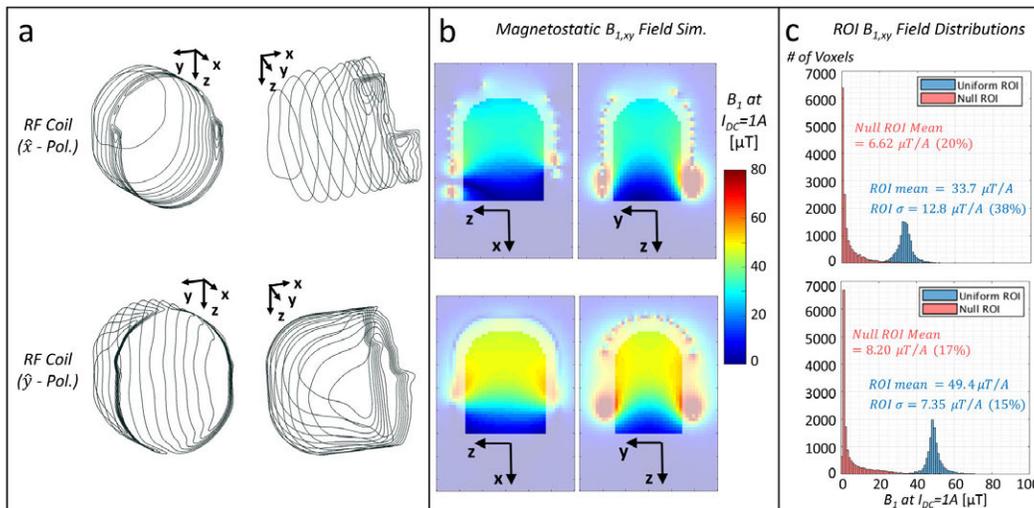


Figure 7-29: a. Optimized RF coil designs that produce x-polarized and y-polarized B_1 fields. b. Simulated spatial map of the transverse ($B_{1,xy}$) field component for the two coils. c. Field histograms for the "Uniform" and "Null" ROIs for the two coil designs.

Phantom: High-Resolution

High-resolution acquisitions used the sequence parameters shown in Table 7.2. Acquisitions were acquired with y as the readout axis using both positive and negative polarities. Images were reconstructed using FFT. k-space data were cropped along the readout direction to $N=128$ points. Partial k-space was used along both the partition ($N_{PE} = 25$ acq/35 recon) and in-plane ($N_{PE} = 81$ acq/101 recon) phase encode directions.

7.8 Imaging Results

7.8.1 Phantom Images

Fig 7-37 shows phantom images obtained using a different readout directions. Fig 7-38 shows phantom images obtained using a different readout strengths. Fig 7-39 shows high-resolution King Tut images obtained with two readout directions.

7.8.2 Other Observations

The gradient coils heated up after several hours of running them. Fig 7-40 shows the scanner after a hard day's work acquiring data. The gradients were observed to heat up from 22°C (ambient) to 41°C . The center frequency of the magnet as measured by NMR was $f = 3.053\text{MHz}$ when the magnet was cold, but dropped down to $f = 3.006\text{MHz}$ by

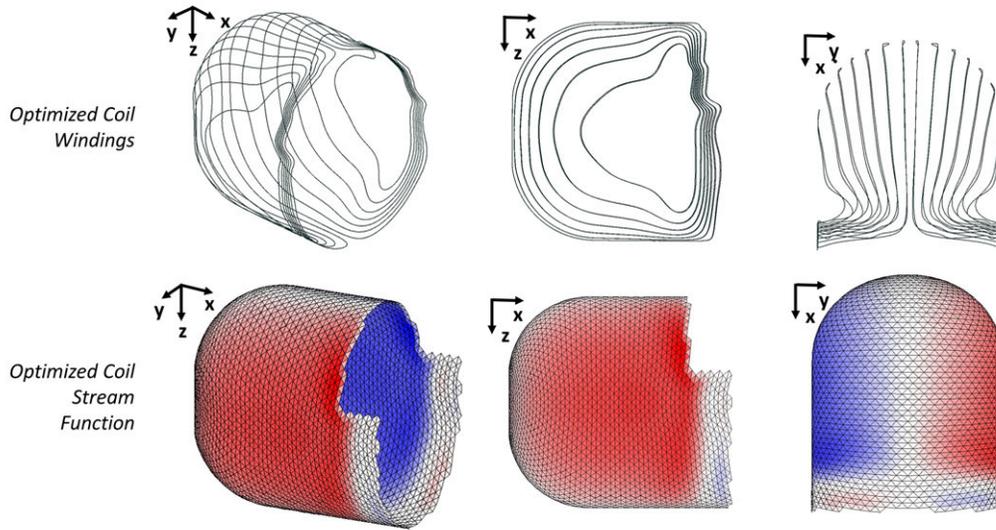


Figure 7-30: Optimized y-polarized coil windings and stream function.

the end of the day (Fig 7-40c). This represents a 1.5% drop. Given that the remanent flux density T_c of NdFeB magnets is approximately $-0.001 \frac{1}{\circ C}$, this frequency change corresponds to a $\Delta T = 15^\circ C$ temperature rise of the magnet. This is close to the observed temperature rise of the gradient coils.

7.9 Discussion and Future Work

- Presented the design of the Head-Optimized MRI including all system components.
- It is a 40cm, 50kg system (magnet, gradient, shims, RF) with 72.5mT, head-shaped and -sized ROI. It has 3 switched gradient coils for spatial encoding.
- Demonstrated its imaging capabilities with RARE-type sequences.
- Lower-SNR, lower-power, and smaller than a normal scanner. Comparable in many ways to the Sparse Halbach system described in Chapter 5.
- Future improvements are possible to the hardware, sequences and reconstruction.
- Applications and *in vivo* imaging are also a natural next step.

7.9.1 System Performance

The Head-Optimized MRI B_0 field was 72.5mT, comparable to the 64mT MR Cap and the 80mT Sparse Halbach, while a factor of 200 to 400 less than a 1.5–3T conventional scanner. Based on this, the Head-Optimized MRI should have a similar "SNR ceiling" as the Sparse Halbach system. Field homogeneity was 1.02mT (pk-pk)/99 μ T (RMS)/4.2kHz (RMS) for

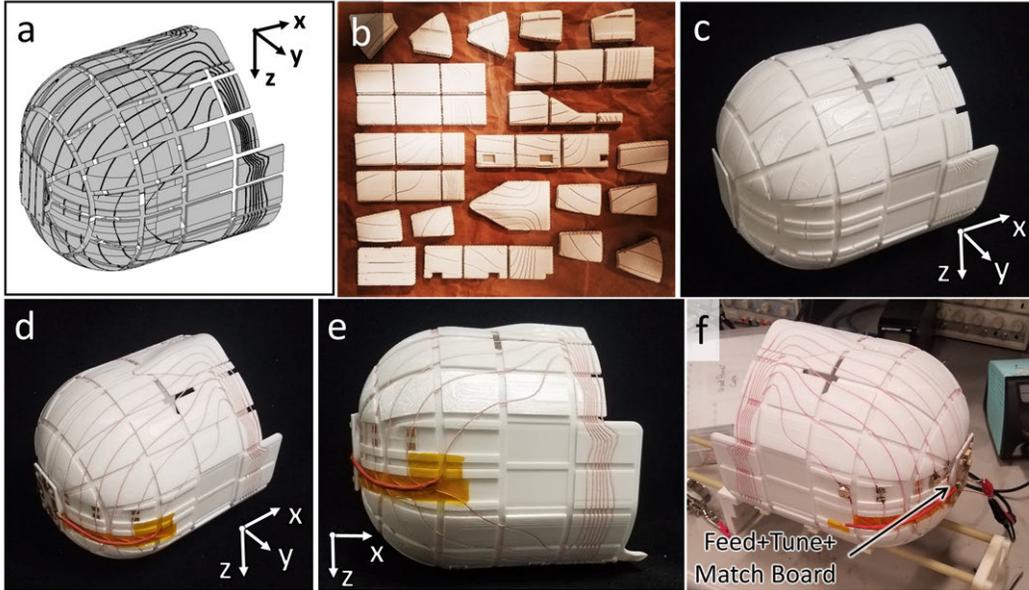


Figure 7-31: a. CAD rendering of the RF coil former tiles with optimized wire grooves. b. The 3D-printed coil former tiles. c. The assembled, unwound coil. d. Assembled, wound coil (isometric view). e. Assembled, wound coil (Size / XZ plane view). f. Assembled coil showing the tune/match/feed board.

the Head-Optimized scanner, compared with $130\mu\text{T}$ (RMS)/ 5.1kHz (RMS) for the Sparse Halbach nonlinear residual, and approximately $1\mu\text{T}$ (RMS)/ 40Hz (RMS) for a high-field system [14, 13]. The similar residuals for the Head-Optimized and Sparse Halbach systems suggests that for acquisitions with similar readout bandwidth, the resulting geometric distortions should be comparable. Note that these metrics are computed over different ROIs: the Head-Optimized scanner used a head-shaped ROI; the Sparse Halbach uses a $D = 20\text{cm}$ sphere, and the quoted high-field system metrics use a $D = 40\text{cm}$ sphere.

RF pulse and sequence constraints differ slightly for the Head-Optimized and Sparse Halbach systems. The Sparse Halbach has a total ROI field range (peak-peak) of $\approx 2.5\text{mT}$, which covers a total spin bandwidth of $\approx 106\text{kHz}$. The Head-Optimized system has no built-in gradient and a peak-peak ROI field range of $1.02\text{mT} / 43\text{kHz}$. This enables the use of lower-bandwidth RF pulses, which can be shorter and/or lower-power. The Head-Optimized scanner field is still quite inhomogeneous when compared with a high-field system, so many sequences and encoding schemes will likely not be practical. However, slice-selective excitation may be possible: Suppose we want to excite a 1cm -thick slab using a 4kHz pulse. From an RF engineering perspective, such a pulse would be within the capabilities of the RF coil and RFPA. From a gradient perspective, one would require a slice-select gradient

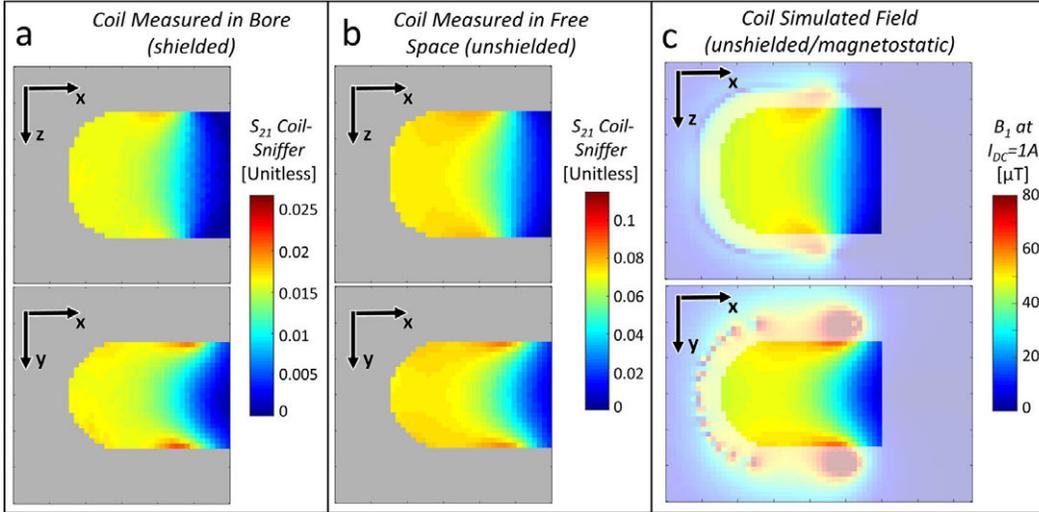


Figure 7-32: a. In-bore measured spatial B_1 field. b. Free-space (unshielded) measured coil B_1 field. c. Simulated magnetostatic B_1 field.

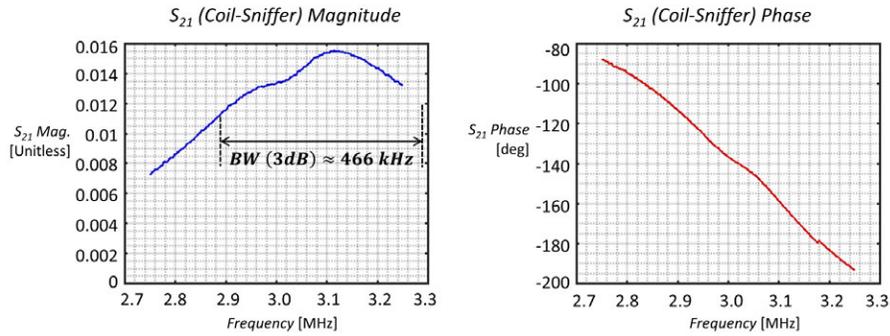


Figure 7-33: S_{21} as measured by transmitting into the coil and receiving with an untuned loop.

of: $G_{ss} = \frac{4\text{kHz}}{\gamma \cdot 1\text{cm}} \approx 10 \frac{\text{mT}}{\text{m}}$. This is feasible with the Head-Optimized gradient system, and was exceeded during several acquisitions. Thus, slice-selective encoding schemes should be explored with this device.

Max gradient power consumption during readout varied between 163W and 941W for the acquisitions listed in Table 7.2. The higher end of that range is still attainable from a wall outlet, but is higher than desired (for example, three gradient channels could not all draw that power simultaneously). However, the max phase encoding gradient power consumption was comparable to the Sparse Halbach acquisitions performed in Chapter 5. The gradient coil efficiencies ($0.275 \frac{\text{mT}}{\text{m}\cdot\text{A}}$ to $1.078 \frac{\text{mT}}{\text{m}\cdot\text{A}}$) were comparable to the MR Cap ($1.019 \frac{\text{mT}}{\text{m}\cdot\text{A}}$ to $1.117 \frac{\text{mT}}{\text{m}\cdot\text{A}}$) and Sparse Halbach gradient coil efficiencies ($0.58 \frac{\text{mT}}{\text{m}\cdot\text{A}}$ to $0.82 \frac{\text{mT}}{\text{m}\cdot\text{A}}$), and were significantly higher than those of a high-field scanner ($\approx 0.01 \frac{\text{mT}}{\text{m}\cdot\text{A}}$).

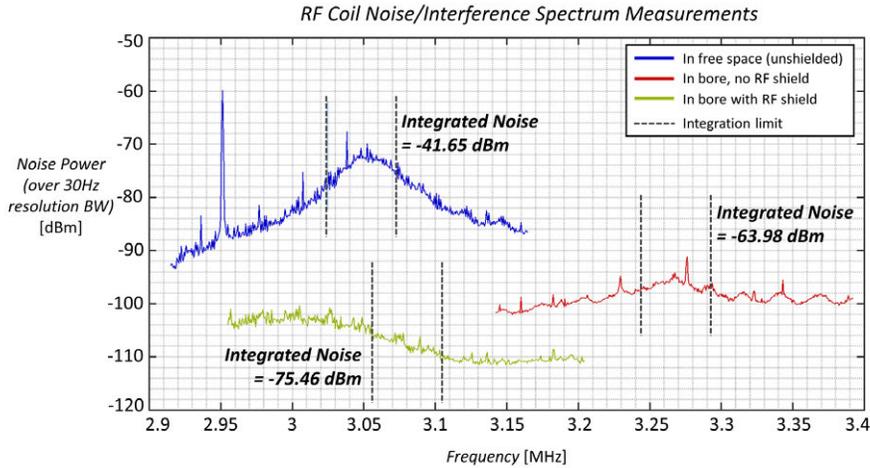


Figure 7-34: Measured noise spectra and integrated noise (over 50 kHz BW) for the coil in free space, in the bore without the spray-on RF shield, and in the bore with the spray-on shield.

The Head-Optimized field-producing assembly (including B_0 magnet, shims, gradient coils, and RF coil) is under 40cm along all linear dimensions. It is slightly smaller than the Sparse Halbach assembly (outer diameter \approx 56cm. Like the Sparse Halbach, it fits on a cart that can be wheeled through a building. The Head-Optimized field-producing assembly weighs 50 kg, roughly half that of the Sparse Halbach assembly (123 kg). Both are significantly smaller than a \approx 2m, 10-ton, immobile superconducting system. The Head-Optimized acquisitions were performed in a shielded room, and the same concerns/potential improvements exist as with the Sparse Halbach system (see Section 5.6.2).

7.9.2 Future Work

The first step would be to incorporate the generalized reconstruction pipeline currently being used with the Sparse Halbach system (Chapter 5). A preliminary step would involve re-measuring the spatial field maps of the gradient coils and the B_0 magnet, as the generalized reconstruction has been observed to be sensitive to mismatches between the physical system and the assumed encoding model. This would enable correction of the geometric distortion observed in the acquired phantom images.

The spatial B_0 field could be further improved by redoing the shimming step. The shim trays presented here are easy to design, build, insert, and remove. It would be relatively

RARE-Type Imaging Sequence with Phase Navigation

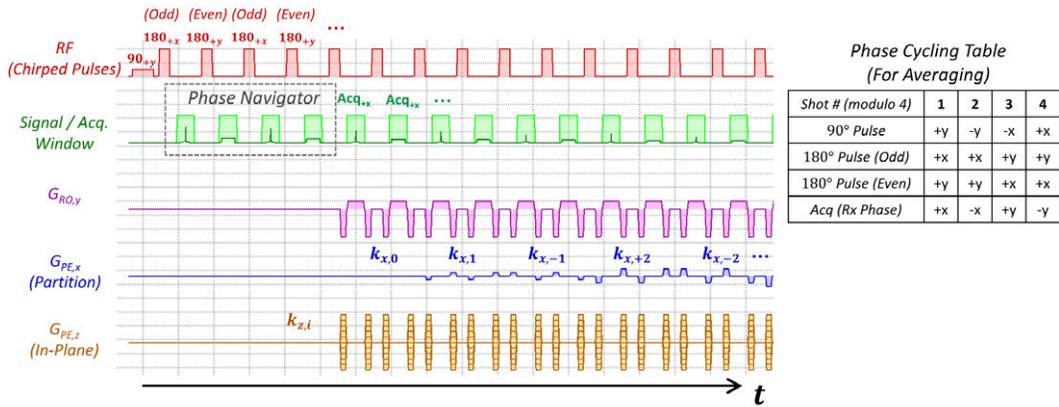


Figure 7-35: RARE-type imaging sequence used for phantom imaging experiments. Partition (x) phase encoding was performed down the echo train; in-plane (z) phase encoding was performed shot-to-shot. Phase cycling was used when averaging multiple shots together.

straightforward to design an improved set of shim trays. This improvement may be a different design that minimizes the same peak-to-peak field homogeneity cost function (as used here), or it could be a design that minimizes a different cost function (for example, the field RMS over the imaging volume, as in Chapter 5).

Addition work is needed to eliminate certain image artifacts that manifest as bands across the image (especially apparent in Figure 7-39). These are thought to result from interference coupling in on the readout gradient coil feeds, possibly being produced by the readout gradient GPA. Finally, healthy and clinical *in vivo* scans should be performed using the Head-Optimized MRI.

Table 7.2: RARE-type Imaging Sequence Parameters

	Vary RO Direction	Vary RO Strength	High-Resolution
<i>TE</i>	7.85 ms	7.85 ms	7.85 ms
ΔTE	7.85 ms	7.85 ms	7.85 ms
<i>Averages</i>	2	2	4
<i>TR</i>	1.3 s	1.3 s	1.3 s
<i>TA</i>	3m03s	3m03s	8m14s
<i>Matrix Size (x-y-z; Recon)</i>	$25 \times 256 \times 61 (RO_y) / 25 \times 61 \times 506 (RO_z)$	$35 \times 128 \times 61$	$35 \times 128 \times 101$
<i>Resolution (x-y-z; Recon)</i>	$12 \times 2.2 \times 4.0 \text{ mm } (RO_y) / 12 \times 4.0 \times 1.9 \text{ mm } (RO_z)$	$8 \times 4.4 \times 4.0 \text{ mm (RO min) to } 8 \times 1.8 \times 4.0 \text{ mm (RO max)}$	$8 \times 2.2 \times 2.4 \text{ mm}$
<i>FoV (x-y-z)</i>	$280 \times 562 \times 242 \text{ mm } (RO_y) / 280 \times 203 \times 939 \text{ mm } (RO_z)$	$280 \times 560 \times 242 \text{ mm (RO min) to } 280 \times 233 \times 242 \text{ mm (RO max)}$	$280 \times 281 \times 242 \text{ mm}$
<i>Matrix Size (x-y-z; Acq)</i>	$25 \times 512 \times 61 (RO_y) / 25 \times 61 \times 512 (RO_z)$	$25 \times 512 \times 61$	$25 \times 512 \times 81$
<i>Readout BW (Recon)</i>	Varies	Varies	$1667 \frac{Hz}{Px}$
<i>Readout Current</i>	$9.48 \text{ A } (RO_y) / 4.63 \text{ A } (RO_z)$	9.48 - 22.76 A	18.9 A
<i>Readout Gradient</i>	$8.9 \frac{mT}{m} (RO_y) / 5.0 \frac{mT}{m} (RO_z)$	8.9 - 21.37 $\frac{mT}{m}$	$17.8 \frac{mT}{m}$
<i>RF Power (Peak)</i>	423 W	423 W	423 W
<i>Pulse Length (t₉₀/t₁₈₀)</i>	3/1.1 ms	3/1.1 ms	3/1.1 ms
<i>Pulse BW</i>	50 kHz	50 kHz	50 kHz

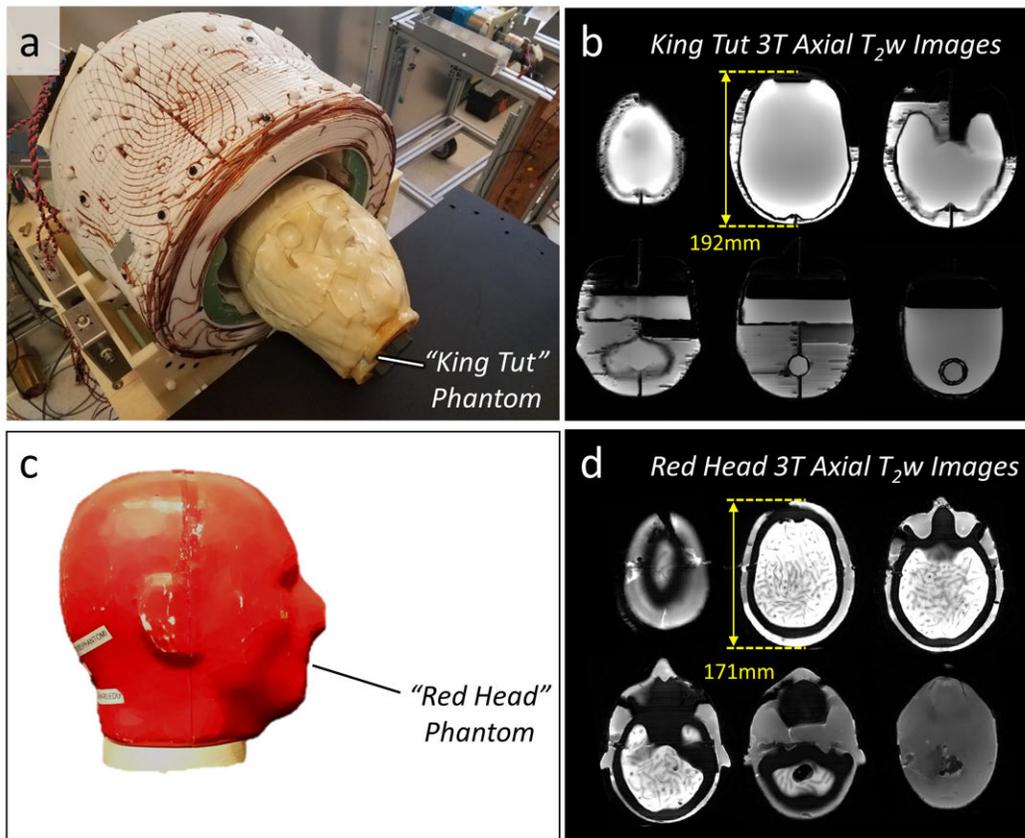


Figure 7-36: a. "King Tut" anthropomorphic head phantom. b. King Tut high-field axial MRI scans with anterior-posterior dimension labeled. c. "Red Head" anthropomorphic head phantom. d. Red Head high-field axial MRI scan with anterior-posterior dimension labeled.

Red Head Axial PD Images

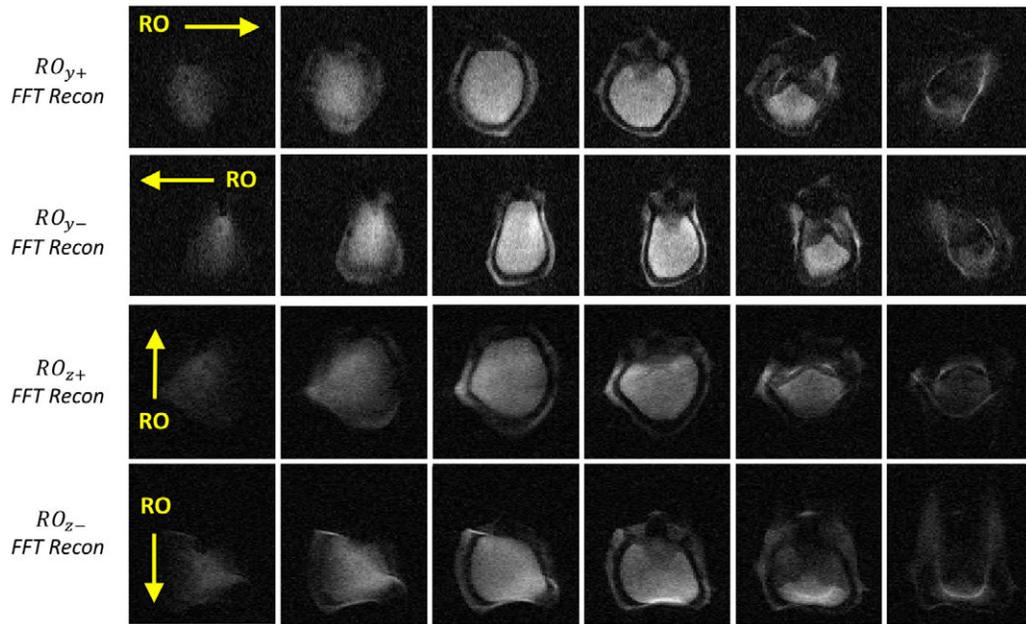


Figure 7-37: Red Head phantom images (axial plane) acquired with four different readout directions (+y, -y, +z, -z).

King Tut Axial PD Images

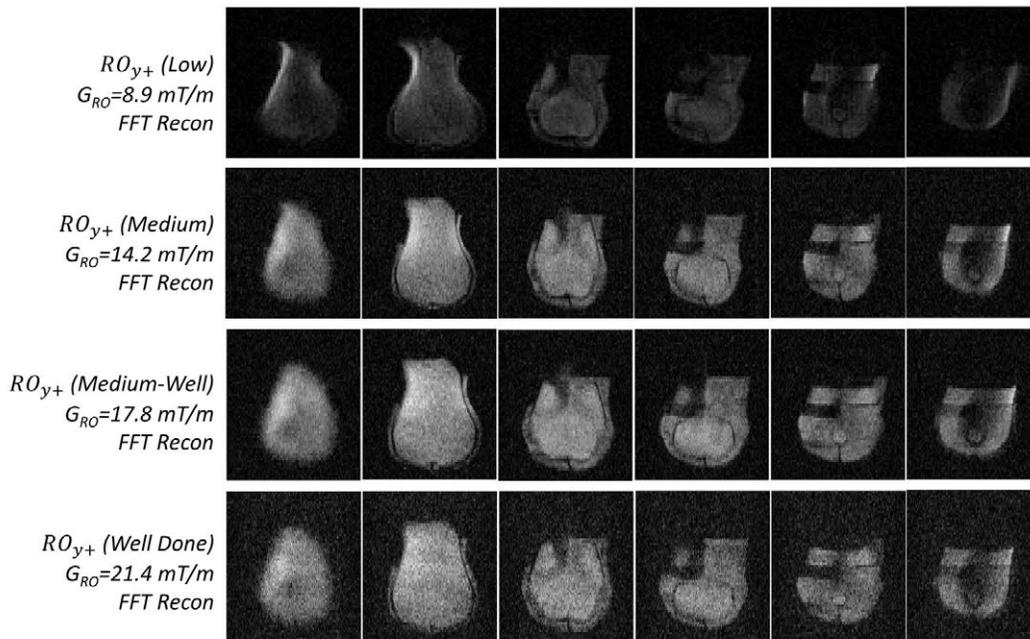


Figure 7-38: King Tut phantom images (axial plane) acquired with four different readout strengths along the +y direction.

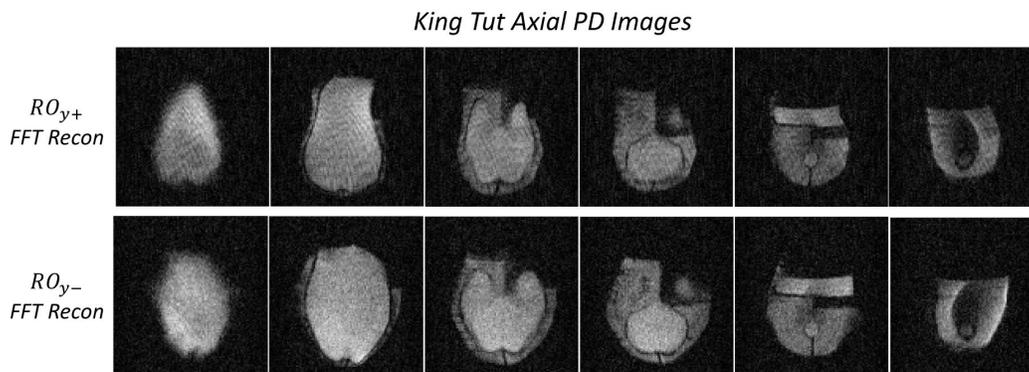


Figure 7-39: King Tut phantom images (axial plane) acquired with two different readout directions (+y and -y).

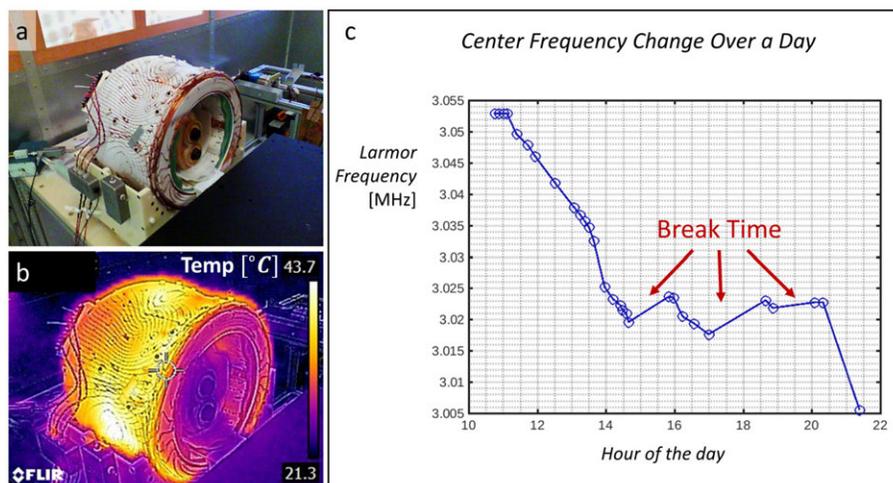


Figure 7-40: a. Visible-light image of the Head-Optimized scanner immediately post-scan. b. Simultaneous infrared (temperature) image of the Head-Optimized scanner. c. Change in NMR center frequency over the course of a day. Shown are break times when no scans were run.

Bibliography

- [1] José P. Marques, Frank F.J. Simonis, and Andrew G. Webb. Low-field MRI: An MR physics perspective. *Journal of Magnetic Resonance Imaging*, 49(6):1528–1542, June 2019.
- [2] Sairam Geethanath and John Thomas Vaughan. Accessible magnetic resonance imaging: A review. *Journal of Magnetic Resonance Imaging*, 49(7):e65–e77, June 2019.
- [3] Lawrence L. Wald, Patrick C. McDaniel, Thomas Witzel, Jason P. Stockmann, and Clarissa Zimmerman Cooley. Low-cost and portable MRI. *Journal of Magnetic Resonance Imaging*, October 2019.
- [4] Edward M. Purcell. *Electricity and Magnetism*. Cambridge University Press, Cambridge, third edition edition, 2013.
- [5] Richard Buxton. *Introduction to Functional Magnetic Resonance Imaging*. Cambridge University Press, second edition, 2009.
- [6] Dwight G Nishimura. *Principles of Magnetic Resonance Imaging*.
- [7] Bruce Lee and Andrew Newberg. Neuroimaging in Traumatic Brain Imaging. *NeuroRx*, 2:372–383, 2005.
- [8] Chris Zacko and Peter LeRoux. *Perioperative Neurosurgical Critical Care*, 2013.
- [9] Ryan T. Muir, Shelly Wang, and Benjamin C. Warf. Global surgery for pediatric hydrocephalus in the developing world: A review of the history, challenges, and future directions. *Neurosurgical Focus*, 41(5):E11, November 2016.
- [10] William J. Powers, Alejandro A. Rabinstein, Teri Ackerson, Opeolu M. Adeoye, Nicholas C. Bambakidis, Kyra Becker, José Biller, Michael Brown, Bart M. Demaer-

schalk, Brian Hoh, Edward C. Jauch, Chelsea S. Kidwell, Thabele M. Leslie-Mazwi, Bruce Ovbiagele, Phillip A. Scott, Kevin N. Sheth, Andrew M. Southerland, Deborah V. Summers, and David L. Tirschwell. 2018 Guidelines for the Early Management of Patients With Acute Ischemic Stroke: A Guideline for Healthcare Professionals From the American Heart Association/American Stroke Association. *Stroke*, 49(3), March 2018.

- [11] Patrick H Knight, Neelabh Maheshwari, Jafar Hussain, Michael Scholl, Michael Hughes, ThomasJ Papadimos, WeidunAlan Guo, James Cipolla, StanislawP Stawicki, and Nicholas Latchana. Complications during intrahospital transport of critically ill patients: Focus on risk identification and prevention. *International Journal of Critical Illness and Injury Science*, 5(4):256, 2015.
- [12] Thomas C. Cosmos and Michael Parizh. Advances in Whole-Body MRI Magnets. *IEEE Transactions on Applied Superconductivity*, 21(3):2104–2109, June 2011.
- [13] Mathias Blasche and Daniel Fischer. Magnet Homogeneity and Shimming. page 8, 2017.
- [14] Siemens MAGNETOM Prisma/Prisma fit System Owner Manual, 2015.
- [15] Jonathan M Baran and John G Webster. Design of low-cost portable ultrasound systems: Review. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 792–795, Minneapolis, MN, September 2009. IEEE.
- [16] Dawn M. Becker, Chelsea A. Tafoya, Sören L. Becker, Grant H. Kruger, Matthew J. Tafoya, and Torben K. Becker. The use of portable ultrasound devices in low- and middle-income countries: A systematic review of the literature. *Tropical Medicine & International Health*, 21(3):294–311, 2016.
- [17] Colin F. Royse, David J. Canty, John Faris, Darsim L. Haji, Michael Veltman, and Alistair Royse. Core Review: Physician-Performed Ultrasound. *Anesthesia & Analgesia*, 115(5):1007–1028, November 2012.
- [18] Bret P. Nelson and Kevin Chason. Use of ultrasound by emergency medical services: A review. *International Journal of Emergency Medicine*, 1(4):253–259, December 2008.

- [19] Rykkje, Carlsen, and Nielsen. Hand-Held Ultrasound Devices Compared with High-End Ultrasound Systems: A Systematic Review. *Diagnostics*, 9(2):61, June 2019.
- [20] Butterfly iQ. store.butterflynetwork.com/us/en/iq/, 2020.
- [21] Clarius. clarius.com, 2020.
- [22] GE Vscan. www.gehealthcare.com/products/ultrasound/point-of-care-ultrasound, 2020.
- [23] KOSMOS. kosmosplatform.com/ultrasound-lp, 2020.
- [24] Philips Lumify. www.usa.philips.com/healthcare/sites/lumify, 2020.
- [25] Sonosite iViz. secure.sonosite.com/iviz-handheld-ultrasound, 2020.
- [26] Sherman Heller and Pat Zanzonico. Nuclear Probes and Intraoperative Gamma Cameras. *Seminars in Nuclear Medicine*, 41(3):166–181, May 2011.
- [27] S.L. Bugby, J.E. Lees, B.S. Bhatia, and A.C. Perkins. Characterisation of a high resolution small field of view portable gamma camera. *Physica Medica*, 30(3):331–339, May 2014.
- [28] Makoto Tsuchimochi and Kazuhide Hayama. Intraoperative gamma cameras for radioguided surgery: Technical characteristics, performance parameters, and clinical applications. *Physica Medica*, 29(2):126–138, March 2013.
- [29] O. R. Brouwer, R. A. Valdes Olmos, L. Vermeeren, C. A. Hoefnagel, O. E. Nieweg, and S. Horenblas. SPECT/CT and a Portable γ -Camera for Image-Guided Laparoscopic Sentinel Node Biopsy in Testicular Cancer. *Journal of Nuclear Medicine*, 52(4):551–554, April 2011.
- [30] A. Bricou, M.-A. Duval, Y. Charon, and E. Barranger. Mobile gamma cameras in breast cancer care – A review. *European Journal of Surgical Oncology (EJSO)*, 39(5):409–416, May 2013.
- [31] I. Stoffels, T. Poeppel, C. Boy, S. Mueller, F. Wichmann, J. Dissemond, D. Schaden-dorf, S. Rosenbaum-Krumme, and J. Klode. Radio-guided surgery: Advantages of

- a new portable γ -camera (Sentinella®) for intraoperative real time imaging and detection of sentinel lymph nodes in cutaneous malignancies: Intraoperative real time imaging of SLN. *Journal of the European Academy of Dermatology and Venereology*, 26(3):308–313, March 2012.
- [32] Marc Thill. MarginProbe® : Intraoperative margin assessment during breast conserving surgery by using radiofrequency spectroscopy. *Expert Review of Medical Devices*, 10(3):301–315, May 2013.
- [33] Marc Thill and Kristin Baumann. New Technologies in Breast Cancer Surgery. *Breast Care*, 7(5):370–376, 2012.
- [34] Dominik Fleischmann and F. Edward Boas. Computed tomography—old ideas and new technology. *European Radiology*, 21(3):510–517, March 2011.
- [35] Kaitlin Peace, Eileen Maloney Wilensky, Suzanne Frangos, Eileen MacMurtrie, Elizabeth Shields, Marianne Hujcs, Joshua Levine, Andrew Kofke, Wei Yang, and Peter D. Le Roux. The Use of a Portable Head CT Scanner in the Intensive Care Unit:. *Journal of Neuroscience Nursing*, 42(2):109–116, April 2010.
- [36] Z. Rumboldt, W. Huda, and J.W. All. Review of Portable CT with Assessment of a Dedicated Head CT Scanner. *American Journal of Neuroradiology*, 30(9):1630–1636, October 2009.
- [37] C S White, C A Meyer, J Wu, and S E Mirvis. Portable CT: Assessing thoracic disease in the intensive care unit. *American Journal of Roentgenology*, 173(5):1351–1356, November 1999.
- [38] Siemens Somatom. www.siemens-healthineers.com/en-us/computed-tomography/mobile-head-ct, 2020.
- [39] Neurologica Ceretom. www.neurologica.com/portable-ct, 2020.
- [40] Xoran Xcat. xorantech.com/xcat-iq/, 2020.
- [41] GE Vivid-e95 Ultrasound. www.gehealthcare.com/products/ultrasound/vivid/vivid-e95, 2020.

- [42] Siemens Symbia Evo SPECT Scanner. www.siemens-healthineers.com/molecular-imaging/spect-and-spect-ct/symbia-evo, 2020.
- [43] Jasper A Jackson, Lowell J Burnett, and J.Frank Harmon. Remote (inside-out) NMR. III. Detection of nuclear magnetic resonance in a remotely produced region of homogeneous magnetic field. *Journal of Magnetic Resonance (1969)*, 41(3):411–421, December 1980.
- [44] G. J. Béné. Medical Diagnosis by Nuclear Magnetism. *Europhysics News*, 8(9):1–3, 1977.
- [45] Georges J. Béné. Nuclear magnetism of liquid systems in the earth field range. *Physics Reports*, 58(4):213–267, February 1980.
- [46] G J Bene, B Borcard, E Hiltbrand, and P Magnin. *In Situ* identification of human physiological fluids by nuclear magnetism in the Earth’s field. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 289(1037):501–502, June 1980.
- [47] R. Damadian, Peter Diehl, Ekkehard Fluck, and Robert Kosfeld, editors. *NMR in Medicine*, volume 19 of *NMR Basic Principles and Progress*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1981.
- [48] J F Schenck, F A Jolesz, P B Roemer, H E Cline, W E Lorensen, R Kikinis, S G Silverman, C J Hardy, W D Barber, E T Laskaris, B Dorri, R W Newman, C E Holley, B D Collick, D P Dietz, D C Mack, S P Souza, C L Dumoulin, R D Darrow, R L St. Peters, K W Rohling, R D Watkins, D R Eisner, S M Blumenfeld, and K G Vosburgh. Superconducting open-configuration MR imaging system for image-guided therapy. *Radiology*, 195(3):805–814, June 1995.
- [49] X. Jiang, G. Shen, Y. Lai, and J. Tian. Development of an Open 0.3 T NdFeB MRI Magnet. *IEEE Transactions on Applied Superconductivity*, 14(2):1621–1623, June 2004.
- [50] V. Seifert, M. Zimmermann, C. Trantakis, H.-E. Vitzthum, K. Kühnel, A. Raabe, F. Bootz, J.-P. Schneider, F. Schmidt, and J. Dietrich. Open MRI-Guided Neurosurgery. *Acta Neurochirurgica*, 141(5):455–464, May 1999.

- [51] Roberto T. Blanco, Risto Ojala, Juho Kariniemi, Jukka Perälä, Jaakko Niinimäki, and Osmo Tervonen. Interventional and intraoperative MRI at low field scanner – a review. *European Journal of Radiology*, 56:130–142, November 2005.
- [52] Paul J. Keall, Michael Barton, and Stuart Crozier. The Australian Magnetic Resonance Imaging–Linac Program. *Seminars in Radiation Oncology*, 24(3):203–206, July 2014.
- [53] Sebastian Klüter. Technical design and concept of a 0.35 T MR-Linac. *Clinical and Translational Radiation Oncology*, 18:98–101, September 2019.
- [54] ViewRay MRI-guided ROAR. viewray.com/mri-guided-roar/, 2020.
- [55] Thomas K. F. Foo, Evangelos Laskaris, Mark Vermilyea, Minfeng Xu, Paul Thompson, Gene Conte, Christopher Epps, Christopher Immer, Seung-Kyun Lee, Ek T. Tan, Dominic Graziani, Jean-Baptiste Mathieu, Christopher J. Hardy, John F. Schenck, Eric Fiveland, Wolfgang Stautner, Justin Ricci, Joseph Piel, Keith Park, Yihe Hua, Ye Bai, Alex Kagan, David Stanley, Paul T. Weavers, Erin Gray, Yunhong Shu, Matthew A. Frick, Norbert G. Campeau, Joshua Trzasko, John Huston, and Matt A. Bernstein. Lightweight, compact, and high-performance 3 T MR system for imaging the brain and extremities. *Magnetic Resonance in Medicine*, 80(5):2232–2245, November 2018.
- [56] Alex Panther, Gilbert Thevathasen, Ian R O Connell, Yi Yao, Curtis N Wiens, Andrew T Curtis, Geron A Bindseil, Chad T Harris, Philip J Beatty, Jeff A Stainsby, Charles H Cunningham, B A Chronik, and Cameron Prion. A Dedicated Head-Only MRI Scanner for Point-of-Care Imaging. In *Intl Symp Magn Reson Med*, page 3679, Montreal, Canada, 2019.
- [57] J Thomas Vaughan, Bert Wang, Djaudat Idiyatullin, Sung-min Sohn, Albert Jang, Lance DelaBarre, and Michael Garwood. Progress Towards a Portable MRI System for Human Brain Imaging. In *Intl Symp Magn Reson Med*, page 498, Singapore, 2016.
- [58] Hyperfine. www.hyperfine.io, 2020.
- [59] Clarissa Z Cooley, Patrick C McDaniel, Jason P Stockmann, Sai Abitha Srinivas, and Lawrence L Wald. A point-of-care MRI scanner for human brain imaging. In *Proc ISMRM*, 2020.

- [60] Yucheng He, Zheng Xu, Jiamin Wu, and Liang Tan. A movable MRI system for brain imaging and its pre-clinical experiments. In *Proceedings of the ICMRM*, page 51, Paris, France, 2019.
- [61] T. O'Reilly, W.M. Teeuwisse, and A.G. Webb. Three-dimensional MRI in a homogeneous 27 cm diameter bore Halbach array magnet. *Journal of Magnetic Resonance*, 307:106578, October 2019.
- [62] G Gold, D Theodorou, T Blair, G Garcia, C Crowley, F Rose, D Trudell, and D Resnick. MR Imaging of the Wrist with a Portable Extremity Scanner. In *Intl Symp Mag Reson Med*, page 2039, 1999.
- [63] Vasileios Ntoukas, Rene Krishnan, and Volker Seifert. The New Generation Polestar N20 for Conventional Neurosurgical Operating Rooms: A Preliminary Report. *Operative Neurosurgery* 1, 62:582–590, 2008.
- [64] Mei Chung, Issa J Dahabreh, Nira Hadar, Sara J Raticheck, James M Gaylor, Thomas A Trikalinos, and Joseph Lau. Emerging MRI Technologies for Imaging Musculoskeletal Disorders Under Loading Stress. Technical Report 11-EHC024-EF, Agency for Healthcare Research and Quality, 2011.
- [65] Esaote Oscan. www.esaote.com/en-US/dedicated-mri/mri-systems/p/o-scan/, 2020.
- [66] Katsumi Kose and Tomoyuki Haishi. High Resolution NMR Imaging using a High Field Yokeless Permanent Magnet. *Magnetic Resonance in Medical Sciences*, 10(3):159–167, 2011.
- [67] Mayu Nakagomi, Michiru Kajiwara, Jumpei Matsuzaki, Katsumasa Tanabe, Sodai Hoshiai, Yoshikazu Okamoto, and Yasuhiko Terada. Development of a small car-mounted magnetic resonance imaging system for human elbows using a 0.2 T permanent magnet. *Journal of Magnetic Resonance*, 304:1–6, July 2019.
- [68] Promaxo Prostate MRI. promaxo.com, 2020.
- [69] Ashvin Bashyam, Chris J. Frangieh, Matthew Li, and Michael J. Cima. Dehydration assessment via a portable, single sided magnetic resonance sensor. *Magnetic Resonance in Medicine*, 83(4):1390–1404, April 2020.

- [70] Lina A. Colucci, Kristin M. Corapi, Matthew Li, Xavier Vela Parada, Andrew S. Allegretti, Herbert Y. Lin, Dennis A. Ausiello, Matthew S. Rosen, and Michael J. Cima. Fluid assessment in dialysis patients by point-of-care magnetic relaxometry. *Science Translational Medicine*, 11(502):eaau1749, July 2019.
- [71] Robert H. Morris, Najlaa K. Almazrouei, Christophe L. Trabi, and Michael I. Newton. NMR CAPIBarA: Proof of Principle of a Low-Field Unilateral Magnetic Resonance System for Monitoring of the Placenta during Pregnancy. *Applied Sciences*, 10(1):162, December 2019.
- [72] Paula R. Keschenau, Hanna Klingel, Silke Reuter, Ann Christina Foldenauer, Jochen Vieß, Dennis Weidener, Julia Andruszkow, Bernhard Bluemich, René Tolba, Michael J. Jacobs, and Johannes Kalder. Evaluation of the NMR-MOUSE as a new method for continuous functional monitoring of the small intestine during different perfusion states in a porcine model. *PLOS ONE*, 13(11):e0206697, November 2018.
- [73] B. Blümich, J. Perlo, and F. Casanova. Mobile single-sided NMR. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 52(4):197–269, May 2008.
- [74] Federico Casanova, Juan Perlo, and Bernhard Blümich, editors. *Single-Sided NMR*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [75] Iseult Project: Travel of the Largest and Most Powerful MRI Magnet in the World, May 2017.
- [76] M Dias, J Hadgraft, P M Glover, and P J McDonald. Stray field magnetic resonance imaging: A preliminary study of skin hydration. *Journal of Physics D: Applied Physics*, 36(4):364–368, February 2003.
- [77] Robert L. Kleinberg. NMR well logging at Schlumberger. *Concepts in Magnetic Resonance*, 13(6):396–403, 2001.
- [78] A A Samoilenko, D Yu Artemov, and L A Sibel'dina. Formation of Sensitive Layer in Experiments on NMR Subsurface Imaging of Solids. *JETP Lett*, 47(7):417–419, 1988.
- [79] H Xu, S M Conolly, G C Scott, and A Macovski. Fundamental Scaling Relations for Homogeneous Magnets. In *Intl Symp Magn Reson Med*, page 475, 1999.

- [80] Beibei Zhang, Carl Gazdzinski, Blaine A. Chronik, Hao Xu, Steven M. Conolly, and Brian K. Rutt. Simple design guidelines for short MRI systems. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 25B(1):53–59, April 2005.
- [81] H. Sanchez, F. Liu, A. Trakic, and S. Crozier. A Simple Relationship for High Efficiency–Gradient Uniformity Tradeoff in Multilayer Asymmetric Gradient Coils for Magnetic Resonance Imaging. *IEEE Transactions on Magnetics*, 43(2):523–532, February 2007.
- [82] Russell Harrison Varian. Ground Liquid Prospecting Method and Apparatus, January 1962.
- [83] Martin Packard and Russel Varian. Free Nuclear Induction in the Earth’s Magnetic Field. In *Minutes of the Stanford Meeting*, page 941. American Physical Society, 1954.
- [84] A G Semenov, A Y Pusep, and M D Schirov. Hydroscope - an installation for prospecting without drilling. In *USSR Academy of Sciences*, pages 1–26, Novosibirsk, 1982.
- [85] Oleg A. Shushakov. Groundwater NMR in conductive water. *Geophysics*, 61(4):998–1006, July 1996.
- [86] Numis Poly. www.iris-instruments.com/numis-poly.html, 2020.
- [87] Klaus Knödel, Gerhard Lange, and Hans-Jürgen Voigt, editors. *Environmental Geology: Handbook of Field Methods and Case Studies*. Springer, Berlin ; New York, 2007.
- [88] Anatoly Legchenko, Jean-Michel Baltassat, Alain Beauce, and Jean Bernard. Nuclear magnetic resonance as a geophysical tool for hydrogeologists. *Journal of Applied Geophysics*, 50(1-2):21–46, May 2002.
- [89] Anthony Rene Barringer and Joseph F White. Groundwater Survey Method and Apparatus, 1968.
- [90] Arnold L Bloom and Dolan L Mansir. Gyromagnetic Resonance Methods and Apparatus, 1967.
- [91] Gustave L Hoehn, Jr, Donald E Woessner, and John R Zimmerman, Jr. Nuclear Magnetic Resonance Well Logging, 1968.

- [92] Robert J S Brown and Bernard W Gamson. Nuclear Magnetism Well Logging Method and Apparatus, 1968.
- [93] R.J.S. Brown and B.W. Gamson. Nuclear Magnetism Logging. *Transactions of the AIME*, 219(01):201–209, December 1960.
- [94] John C Cook. Electromagnetic Resonance Borehole Assay Logging. Technical Report BuMines OFR 69-78, United States Department of the Interior Bureau of Mines, November 1977.
- [95] Richard K Cooper and Jasper A Jackson. Remote (inside-out) NMR. I. Remote production of region of homogeneous magnetic field. *Journal of Magnetic Resonance (1969)*, 41(3):400–405, December 1980.
- [96] Lowell J Burnett and Jasper A Jackson. Remote (inside-out) NMR. II. Sensitivity of NMR detection for external samples. *Journal of Magnetic Resonance (1969)*, 41(3):406–410, December 1980.
- [97] John W. Lane, Frederick D. Day-Lewis, Judith L. Robinson, Cian B Dawson, Lee D. Slater, and Carole D. Johnson. Borehole Nuclear Magnetic Resonance (NMR): A valuable tool for environmental site management. 2017.
- [98] David Walsh, Peter Turner, Elliot Grunewald, Hong Zhang, James J. Butler, Ed Reboulet, Steve Knobbe, Tom Christy, John W. Lane, Carole D. Johnson, Tim Munday, and Andrew Fitzpatrick. A Small-Diameter NMR Logging Tool for Groundwater Investigations: D. Walsh Ground Water xx, no. xx: Xx-xx. *Groundwater*, 51(6):914–926, November 2013.
- [99] S.A. Altobelli, Mark S. Conradi, E. Fukushima, J. Hodgson, T.J. Nedwed, D.A. Pallandro, A. Peach, N.J. Sowko, and H. Thomann. Helicopter-borne NMR for detection of oil under sea-ice. *Marine Pollution Bulletin*, 144:160–166, July 2019.
- [100] R. F. Paetzold, G. A. Matzkanin, and A. De Los Santos. Surface Soil Water Content Measurement Using Pulsed Nuclear Magnetic Resonance Techniques. *Soil Science Society of America Journal*, 49(3):537–540, May 1985.

- [101] G A Matzkanin, A De Los Santos, and D A Whiting. Determination of Moisture Levels in Structural Concrete Using Pulsed NMR. Technical Report DOT-FH-11-9487, Southwest Research Institute, 1982.
- [102] G A Matzkanin and A De Los Santos. In-situ determination of moisture levels in structural concrete by a modified nuclear magnetic resonance (NMR) method. Technical Report FHWA/RD-85/101, Southwest Research Institute, 1985.
- [103] L. J. Burnett and M. A. Fineman. Evaluation of One-Sided Nuclear Magnetic Resonance for Remote Detection of Explosives. Technical report, Defense Technical Information Center, Fort Belvoir, VA, October 1987.
- [104] Lowell J. Burnett, Dale R. McKay, Erik M. Magnuson, and E. J. VanderHeiden. Solid Rocket Motor NDE Using Nuclear Magnetic Resonance. In Donald O. Thompson and Dale E. Chimenti, editors, *Review of Progress in Quantitative Nondestructive Evaluation*, pages 663–670. Springer US, Boston, MA, 1993.
- [105] R M Pearson and D I Wetzel. Rapid on-line moisture determination of whole kernel wheat by pulsed nuclear magnetic resonance. *Cereal Foods World*, 30:563–564, 1985.
- [106] R M Pearson, L R Ream, C Job, and J Adams. Use of NMR spectrometers for process control. *Cereal Foods World*, 32:822–826, 1987.
- [107] Agnes Haber, Bernhard Blümich, Daria Souvorova, and Eleonora Del Federico. Ancient Roman wall paintings mapped nondestructively by portable NMR. *Analytical and Bioanalytical Chemistry*, 401(4):1441–1452, September 2011.
- [108] P.M. Glover, P.S. Aptaker, J.R. Bowler, E. Ciampi, and P.J. McDonald. A Novel High-Gradient Permanent Magnet for the Profiling of Planar Films and Coatings. *Journal of Magnetic Resonance*, 139(1):90–97, July 1999.
- [109] G. Eidmann, R. Savelsberg, P. Blümmler, and B. Blümich. The NMR MOUSE, a Mobile Universal Surface Explorer. *Journal of Magnetic Resonance, Series A*, 122(1):104–109, September 1996.
- [110] Magritek NMR-MOUSE. www.magritek.com/products/nmr-mouse/, 2020.

- [111] Bernhard Blümich. The Incredible Shrinking Scanner. *Scientific American*, 299(5):92–98, November 2008.
- [112] W Baran, P R Locher, W Susse, and H Zijlstra. Permanent Magnet System for Magnetic Resonance Imaging. In *Concerted European Action on Magnets (CEAM)*. 1989.
- [113] Shin Utsuzawa and Eiichi Fukushima. Unilateral NMR with a barrel magnet. *Journal of Magnetic Resonance*, 282:104–113, September 2017.
- [114] Andrew E. Marble, Igor V. Mastikhin, Bruce G. Colpitts, and Bruce J. Balcom. A compact permanent magnet array with a remote homogeneous field. *Journal of Magnetic Resonance*, 186(1):100–104, May 2007.
- [115] Andrew E. Marble, Igor V. Mastikhin, Bruce G. Colpitts, and Bruce J. Balcom. A constant gradient unilateral magnet for near-surface MRI profiling. *Journal of Magnetic Resonance*, 183(2):228–234, December 2006.
- [116] Ashvin Bashyam, Matthew Li, and Michael J. Cima. Design and experimental validation of Unilateral Linear Halbach magnet arrays for single-sided magnetic resonance. *Journal of Magnetic Resonance*, 292:36–43, 2018.
- [117] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The NMR phased array. *Magnetic Resonance in Medicine*, 16(2):192–225, November 1990.
- [118] Cecil E. Hayes, Neil Hattes, and Peter B. Roemer. Volume imaging with MR phased arrays. *Magnetic Resonance in Medicine*, 18(2):309–319, April 1991.
- [119] Clarissa Zimmerman Cooley, Jason P. Stockmann, Brandon D. Armstrong, Mathieu Sarracanie, Michael H. Lev, Matthew S. Rosen, and Lawrence L. Wald. Two-dimensional imaging in a lightweight portable MRI scanner without gradient coils: Lightweight MRI Scanner without Gradient Coils. *Magnetic Resonance in Medicine*, 73(2):872–883, February 2015.
- [120] W. Thomas Dixon, Leila N. Du, David D. Faul, Mokhtar Gado, and Susan Rossnick. Projection angiograms of blood labeled by adiabatic fast passage. *Magnetic Resonance in Medicine*, 3(3):454–462, June 1986.

- [121] Robert Trampel, Toralf Mildner, Ute Goerke, Andreas Schaefer, Wolfgang Driesel, and David G. Norris. Continuous arterial spin labeling using a local magnetic field gradient coil. *Magnetic Resonance in Medicine*, 48(3):543–546, September 2002.
- [122] J Hennig, Ch Boesch, R Gruetter, and E Martin. Homogeneity Spoil Spectroscopy as a Tool of Spectrum Localization for in Vivo Spectroscopy. *Journal of Magnetic Resonance*, 75:179–183, 1987.
- [123] David G. Wiesler, Han Wen, Steven D. Wolff, and Robert S. Balaban. Reduction of field of view in MRI using a surface-spoiling local gradient insert. *Journal of Magnetic Resonance Imaging*, 8(4):981–988, July 1998.
- [124] Ernesto Danieli and Bernhard Blümich. Single-sided magnetic resonance profiling in biological and materials science. *Journal of Magnetic Resonance*, 229:142–154, April 2013.
- [125] B. Blümich, F. Casanova, A. Buda, K. Kremer, and T. Wegener. Mobile NMR for Analysis of Polyethylene Pipes. *Acta Physica Polonica A*, 108(1):13–23, 2005.
- [126] N.O. Goga, D.E. Demco, J. Kolz, R. Ferencz, A. Haber, F. Casanova, and B. Blümich. Surface UV aging of elastomers investigated with microscopic resolution by single-sided NMR. *Journal of Magnetic Resonance*, 192(1):1–7, May 2008.
- [127] P.J. McDonald, P.S. Aptaker, J. Mitchell, and M. Mulheron. A unilateral NMR magnet for sub-structure analysis in the built environment: The Surface GARField. *Journal of Magnetic Resonance*, 185(1):1–11, March 2007.
- [128] Maria Brai, Mara Camaiti, Cinzia Casieri, Francesco De Luca, and Paola Fantazzini. Nuclear magnetic resonance for cultural heritage. *Magnetic Resonance Imaging*, 25(4):461–465, May 2007.
- [129] Bernhard Blümich, Agnes Haber, Federico Casanova, Eleonora Del Federico, Victoria Boardman, Gerhard Wahl, Antonella Stilliano, and Licio Isolani. Noninvasive depth profiling of walls by portable nuclear magnetic resonance. *Analytical and Bioanalytical Chemistry*, 397:3117–3125, 2010.

- [130] Federica Presciutti, Juan Perlo, Federico Casanova, Stefan Glöggler, Costanza Miliani, Bernhard Blümich, Brunetto Giovanni Brunetti, and Antonio Sgamellotti. Noninvasive nuclear magnetic resonance profiling of painting layers. *Applied Physics Letters*, 93(3):033505, July 2008.
- [131] I. Viola, S. Bubici, C. Casieri, and F. De Luca. The Codex Major of the Collectio Altaempsiana: A non-invasive NMR study of paper. *Journal of Cultural Heritage*, 5(3):257–261, July 2004.
- [132] C Casieri, S Bubici, I Viola, and F De Luca. A low-resolution non-invasive NMR characterization of ancient paper. *Solid State Nuclear Magnetic Resonance*, 26(2):65–73, September 2004.
- [133] Donatella Capitani, Anatoly P. Sobolev, Valeria Di Tullio, Luisa Mannina, and Noemi Proietti. Portable NMR in food analysis. *Chemical and Biological Technologies in Agriculture*, 4(1):17, December 2017.
- [134] Gisela Guthausen, Harald Todt, Wolfgang Burk, Dieter Schmalbein, Andreas Guthausen, and Andreas Kamlowksi. Single-Sided NMR in Foods. In Graham A. Webb, editor, *Modern Magnetic Resonance*, pages 1895–1897. Springer Netherlands, Dordrecht, 2006.
- [135] Andreas Guthausen, Gisela Guthausen, Andreas Kamlowksi, Harald Todt, Wolfgang Burk, and Dieter Schmalbein. Measurement of fat content of food with single-sided NMR. *Journal of the American Oil Chemists' Society*, 81(8):727–731, August 2004.
- [136] D. Capitani, L. Mannina, N. Proietti, A.P. Sobolev, A. Tomassini, A. Miccheli, M.E. Di Cocco, G. Capuani, R. De Salvador, and M. Delfini. Monitoring of metabolic profiling and water status of Hayward kiwifruits by nuclear magnetic resonance. *Talanta*, 82(5):1826–1838, October 2010.
- [137] Michael Pinter, Tod Harter, Michael McCarthy, and Matthew Augustine. Towards Using NMR to Screen for Spoiled Tomatoes Stored in 1,000 L, Aseptically Sealed, Metal-Lined Totes. *Sensors*, 14(3):4167–4176, March 2014.

- [138] Zheng Xu, Robert Morris, Martin Bencsik, and Michael Newton. Detection of Virgin Olive Oil Adulteration Using Low Field Unilateral NMR. *Sensors*, 14(2):2028–2035, January 2014.
- [139] G. Maddinelli and E. Peron. Characterization of sedimentary rocks with a single-sided NMR instrument. *Applied Magnetic Resonance*, 29(4):549–559, December 2005.
- [140] R Haken and B Blümich. Anisotropy in Tendon Investigated in Vivo by a Portable NMR Scanner, the NMR-MOUSE. *Journal of Magnetic Resonance*, 144(2):195–199, June 2000.
- [141] Erik Rössler, Carlos Mattea, Ayret Mollova, and Siegfried Stapf. Low-field one-dimensional and direction-dependent relaxation imaging of bovine articular cartilage. *Journal of Magnetic Resonance*, 213(1):112–118, December 2011.
- [142] Elisabetta Ciampi, Michael van Ginkel, Peter J. McDonald, Simon Pitts, Eleanor Y. M. Bonnist, Scott Singleton, and Ann-Marie Williamson. Dynamic in vivo mapping of model moisturiser ingress into human skin by GARfield MRI. *NMR in Biomedicine*, 24(2):135–144, February 2011.
- [143] Patrick Judeinstein, Filippo Ferdeghini, Rodrigo Oliveira-Silva, Jean-Marc Zanotti, and Dimitrios Sakellariou. Low-field single-sided NMR for one-shot 1D-mapping: Application to membranes. *Journal of Magnetic Resonance*, 277:25–29, April 2017.
- [144] Maxime Van Landeghem, Ernesto Danieli, Juan Perlo, Bernhard Blümich, and Federico Casanova. Low-gradient single-sided NMR sensor for one-shot profiling of human skin. *Journal of Magnetic Resonance*, 215:74–84, February 2012.
- [145] Zhonghua He, Wei He, Jiamin Wu, and Zheng Xu. The Novel Design of a Single-Sided MRI Probe for Assessing Burn Depth. *Sensors*, 17(3):526, March 2017.
- [146] Paula R. Keschenau, Nina Simons, Hanna Klingel, Silke Reuter, Ann Christina Foldenauer, Jochen Vieß, Dennis Weidener, Julia Andruszkow, Bernhard Blümich, René Tolba, Michael J. Jacobs, and Johannes Kalder. Perfusion-related changes in intestinal diffusion detected by NMR-MOUSE® monitoring in minipigs. *Microvascular Research*, 125:103876, September 2019.

- [147] Monique C. Tourell, Tonima S. Ali, Honor J Hugo, Chris Pyke, Samuel Yang, Thomas Lloyd, Erik W. Thompson, and Konstantin I. Momot. T_1 -based sensing of mammographic density using single-sided portable NMR. *Magnetic Resonance in Medicine*, 80(3):1243–1251, September 2018.
- [148] Tonima S. Ali, Monique C. Tourell, Honor J. Hugo, Chris Pyke, Samuel Yang, Thomas Lloyd, Erik W. Thompson, and Konstantin I. Momot. Transverse relaxation-based assessment of mammographic density and breast tissue composition by single-sided portable NMR. *Magnetic Resonance in Medicine*, pages 1199–1213, April 2019.
- [149] Matthew Li, Christophoros C. Vassiliou, Lina A. Colucci, and Michael J. Cima. ^1H nuclear magnetic resonance (NMR) as a tool to measure dehydration in mice. *NMR in Biomedicine*, 28(8):1031–1039, August 2015.
- [150] Matthew Li. *Methods and Devices for Noninvasive Fluid Volume Assessment*. PhD thesis, MIT, 2016.
- [151] Michael J Cima, Matthew Li, Christophoros C Vassiliou, Negar Tavassolian, and Lina A. Colucci. NMR sensor and methods for rapid non-invasive determination of hydration state or vascular volume of a subject, May 2016.
- [152] S. Rahmatallah, Y. Li, H.C. Seton, I.S. Mackenzie, J.S. Gregory, and R.M. Aspden. NMR detection and one-dimensional imaging using the inhomogeneous magnetic field of a portable single-sided magnet. *Journal of Magnetic Resonance*, 173(1):23–28, March 2005.
- [153] F. Casanova, J. Perlo, B. Blümich, and K. Kremer. Multi-echo imaging in highly inhomogeneous magnetic fields. *Journal of Magnetic Resonance*, 166(1):76–81, January 2004.
- [154] J. Perlo, F. Casanova, and B. Blümich. 3D imaging with a single-sided sensor: An open tomograph. *Journal of Magnetic Resonance*, 166(2):228–235, February 2004.
- [155] Fa-Hsuan Lin, Thomas Witzel, Gerrit Schultz, Daniel Gallichan, Wen-Jui Kuo, Fu-Nien Wang, Juergen Hennig, Maxim Zaitsev, and John W. Belliveau. Reconstruction of MRI data encoded by multiple nonbijective curvilinear magnetic fields. *Magnetic Resonance in Medicine*, 68(4):1145–1156, October 2012.

- [156] R. M. Fano. Theoretical Limitations on the Broadband Matching of Arbitrary Impedances. *Journal of the Franklin Institute*, 249(1):57–83, 1950.
- [157] E. Danieli, J. Perlo, B. Blümich, and F. Casanova. Highly Stable and Finely Tuned Magnetic Fields Generated by Permanent Magnet Assemblies. *Physical Review Letters*, 110(18):180801, April 2013.
- [158] Anna J. Parker, Wasif Zia, Christian W.G. Rehorn, and Bernhard Blümich. Shimming Halbach magnets utilizing genetic algorithms to profit from material imperfections. *Journal of Magnetic Resonance*, 265:83–89, April 2016.
- [159] M. Sagawa, S. Fujimura, N. Togawa, H. Yamamoto, and Y. Matsuura. New material for permanent magnets on a base of Nd and Fe (invited). *Journal of Applied Physics*, 55(6):2083–2087, March 1984.
- [160] Peter Blümler and Federico Casanova. Chapter 5: Hardware Developments: Halbach Magnet Arrays. In Michael L Johns, Einar O Fridjonsson, Sarah J Vogt, and Agnes Haber, editors, *New Developments in NMR*, pages 133–157. Royal Society of Chemistry, Cambridge, 2015.
- [161] Peter Blümler and Federico Casanova. Chapter 4: Hardware Developments: Single-Sided Magnets. In *Mobile NMR and MRI: Developments and Applications*, New Developments in NMR No. 5, page 23. The Royal Society of Chemistry, 2016.
- [162] K. Halbach. Design of permanent multipole magnets with oriented rare earth cobalt material. *Nuclear Instruments and Methods*, 169(1):1–10, February 1980.
- [163] H. Raich and P. Blümler. Design and construction of a dipolar Halbach array with a homogeneous field from identical bar magnets: NMR Mandhalas: Design of a Dipolar Halbach Array. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 23B(1):16–25, October 2004.
- [164] Giorgio Moresi and Richard Magin. Miniature permanent magnet for table-top NMR. *Concepts in Magnetic Resonance*, 19B(1):35–43, 2003.
- [165] Clarissa Zimmerman Cooley, Melissa W. Haskell, Stephen F. Cauley, Charlotte Sappo, Cristen D. Lapierre, Christopher G. Ha, Jason P. Stockmann, and Lawrence L. Wald.

- Design of Sparse Halbach Magnet Arrays for Portable MRI Using a Genetic Algorithm. *IEEE Transactions on Magnetics*, 54(1):1–12, January 2018.
- [166] Thomas O’Reilly, Wouter Teeuwisse, Lukas Winter, and Andrew Webb. The Design of a Homogeneous Large-bore Halbach Array for Low Field MRI. In *Intl Symp Magn Reson Med*, page 272, Montreal, Canada, 2019.
- [167] F. J. Lowes and B. Duka. Magnetic multipole moments (Gauss coefficients) and vector potential given by an arbitrary current distribution. *Earth, Planets and Space*, 63(12):i–vi, December 2011.
- [168] Didier Pinchon and Philip E Hoggan. Rotation matrices for real spherical harmonics: General rotations of atomic orbitals in space-fixed axes. *Journal of Physics A: Mathematical and Theoretical*, 40(7):1597–1610, February 2007.
- [169] Stratasy. Comparison of Bonding Methods for FDM Materials, 2014.
- [170] Loctite Design Guide for Bonding Plastic, 2011.
- [171] G.N. Peeren. Stream function approach for determining optimal surface currents. *Journal of Computational Physics*, 191(1):305–321, October 2003.
- [172] Gael Bringout, Ksenija Grafe, and Thorsten M. Buzug. Performance of Shielded Electromagnet - Evaluation Under Low-Frequency Excitation. *IEEE Transactions on Magnetics*, 51(2):1–4, February 2015.
- [173] H. SanchezLopez, C. GarridoSalmon, C. CabalMirabal, and H. Saint-Jalmes. Designing an Efficient Resistive Magnet for Magnetic Resonance Imaging. *IEEE Transactions on Magnetics*, 40(5):3378–3381, September 2004.
- [174] Michael Poole and Richard Bowtell. Novel gradient coils designed using a boundary element method. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 31B(3):162–175, August 2007.
- [175] Michael Poole. *Improved Equipment and Techniques for Dynamic Shimming in High Field MRI*. PhD thesis, The University of Nottingham, 2007.
- [176] Jonathan C. Sharp and Scott B. King. MRI using radiofrequency magnetic field phase gradients. *Magnetic Resonance in Medicine*, pages NA–NA, 2009.

- [177] Jason P. Stockmann, Clarissa Z. Cooley, Bastien Guerin, Matthew S. Rosen, and Lawrence L. Wald. Transmit Array Spatial Encoding (TRASE) using broadband WURST pulses for RF spatial encoding in inhomogeneous B₀ fields. *Journal of Magnetic Resonance*, 268:36–48, July 2016.
- [178] Sai Abitha Srinivas, Clarissa Z Cooley, Jason P Stockmann, Patrick C McDaniel, and Lawrence L Wald. Retrospective Electromagnetic Interference Mitigation in a Portable Low-Field MRI System. In *Proceedings of the Annual Meeting of the ISMRM*, 2020.
- [179] David O Walsh. Noise cancelling in-situ NMR detection, December 2011.
- [180] Todd Rearick, Gergory L Charvat, Matthew Scot Rosen, and Jonathan M Rothberg. Noise suppression methods and apparatus, April 2017.
- [181] Y. Terada, S. Kono, K. Ishizawa, S. Inamura, T. Uchiumi, D. Tamada, and K. Kose. Magnetic field shimming of a permanent magnet using a combination of pieces of permanent magnets and a single-channel shim coil for skeletal age assessment of children. *Journal of Magnetic Resonance*, 230:125–133, May 2013.
- [182] Patrick C Mcdaniel, Clarissa Z Cooley, Jason P Stockmann, and Lawrence L Wald. A 6.3kg single-sided magnet for 3D, point-of-care brain imaging. In *Proc ISMRM*, page 943, Paris, France, 2018.
- [183] Patrick C. McDaniel, Clarissa Zimmerman Cooley, Jason P. Stockmann, and Lawrence L. Wald. The MR Cap: A single-sided MRI system designed for potential point-of-care limited field-of-view brain imaging. *Magnetic Resonance in Medicine*, 82(5):1946–1960, November 2019.
- [184] M.D Hürlimann and D.D Griffin. Spin Dynamics of Carr–Purcell–Meiboom–Gill-like Sequences in Grossly Inhomogeneous B₀ and B₁ Fields and Application to NMR Well Logging. *Journal of Magnetic Resonance*, 143(1):120–135, March 2000.
- [185] H. Leupold and E. Potenziani. Novel high-field permanent-magnet flux sources. *IEEE Transactions on Magnetics*, 23(5):3628–3629, September 1987.
- [186] Clarissa Zimmerman Cooley, Jason P. Stockmann, Thomas Witzel, Cris LaPierre, Azma Mareyam, Feng Jia, Maxim Zaitsev, Yang Wenhui, Wang Zheng, Pascal Stang,

- Greig Scott, Elfar Adalsteinsson, Jacob K. White, and Lawrence L. Wald. Design and implementation of a low-cost, tabletop MRI scanner for education and research prototyping. *Journal of Magnetic Resonance*, 310:106625, January 2020.
- [187] D A Feinberg, J C Hoenninger, L E Crooks, L Kaufman, J C Watts, and M Arakawa. Inner volume MR imaging: Technical concepts and their application. *Radiology*, 156(3):743–747, September 1985.
- [188] Y. Luo, R.A. de Graaf, L. DelaBarre, A. Tannús, and M. Garwood. BISTRO: An outer-volume suppression method that tolerates RF field inhomogeneity: B₁-Insensitive Outer-Volume Suppression. *Magnetic Resonance in Medicine*, 45(6):1095–1102, June 2001.
- [189] Thilaka Sumanaweera, Gary Glover, Samuel Song, John Adler, and Sandy Napel. Quantifying MRI geometric distortion in tissue. *Magnetic Resonance in Medicine*, 31(1):40–47, January 1994.
- [190] Serge Langlois, M Desvignes, J M Constans, and M Revenu. MRI geometric distortion: A simple approach to correcting the effects of non-linear gradient fields. *Journal of Magnetic Resonance Imaging*, 9:821–831, 1999.
- [191] Hermann Scharfetter, Andreas Petrovic, Heidi Eggenhofer, and Rudolf Stollberger. A no-tune no-match wideband probe for nuclear quadrupole resonance spectroscopy in the VHF range. *Measurement Science and Technology*, 25(12):125501, December 2014.
- [192] Ahmad Raad and Luc Darrasse. Optimization of NMR receiver bandwidth by inductive coupling. *Magnetic Resonance Imaging*, 10(1):55–65, January 1992.
- [193] S. Mandal, S. Utsuzawa, D.G. Cory, M. Hürlimann, M. Poitzsch, and Y.-Q. Song. An ultra-broadband low-frequency magnetic resonance system. *Journal of Magnetic Resonance*, 242:113–125, May 2014.
- [194] Mathieu Sarracanie, Brandon D. Armstrong, Jason Stockmann, and Matthew S. Rosen. High speed 3D overhauser-enhanced MRI using combined b-SSFP and compressed sensing: 3D OMRI Using b-SSFP and Compressed Sensing. *Magnetic Resonance in Medicine*, 71(2):735–745, February 2014.

- [195] Edwin D. Becker, James A. Feretti, and Thomas C. Farrar. Driven Equilibrium Fourier Transform Spectroscopy. A New Method for Nuclear Magnetic Resonance Signal Enhancement. *Journal of the American Chemical Society*, 91(27):7784–7785, December 1969.
- [196] H. Y. Carr and E. M. Purcell. Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Physical Review*, 94(3):630–638, May 1954.
- [197] Jason P Stockmann, Patrick C McDaniel, Christopher Vaughn, Clarissa Z Cooley, and Lawrence L Wald. Feasibility of Brain Pathology Assessment with Diffusion Imaging on a Portable Scanner using a Fixed Encoding Field. In *Intl Symp Magn Reson Med*, page 1196, Montreal, Canada, 2019.
- [198] Christian Bluecher, Haopeng Han, Werner Hoffmann, Reiner Seemann, Frank Seifert, Thoralf Niendorf, and Lukas Winter. COSI Transmit: Open Source Soft- and Hardware Transmission System for traditional and rotating MR. In *Intl Symp Magn Reson Med*, page 1055, Honolulu, HI, USA, 2017.
- [199] Christopher J. Hasselwander, Zhipeng Cao, and William A. Grissom. Gr-MRI: A software package for magnetic resonance imaging using software defined radios. *Journal of Magnetic Resonance*, 270:47–55, September 2016.
- [200] Suma Anand, Jason P Stockmann, Lawrence L Wald, and Thomas Witzel. A low-cost (< \$500 USD) FPGA-based console capable of real-time control. In *Intl Symp Magn Reson Med*, page 948, Paris, France, 2018.
- [201] J. Beau W. Webber and Pavel Demin. Credit-card sized field and benchtop NMR relaxometers using field programmable gate arrays. *Magnetic Resonance Imaging*, 56:45–51, February 2019.
- [202] Kumar Khandagle and S. S. Rathod. Implementation of MRI gradient generation system and controller on field programmable gate array(FPGA). In *2018 International Conference on Communication Information and Computing Technology (ICCICT)*, pages 1–4, Mumbai, February 2018. IEEE.

- [203] Hsueh-Ying Chen, Yaewon Kim, Pulak Nath, and Christian Hilty. An ultra-low cost NMR device with arbitrary pulse programming. *Journal of Magnetic Resonance*, 255:100–105, June 2015.
- [204] M.D. Hürlimann. Diffusion and Relaxation Effects in General Stray Field NMR Experiments. *Journal of Magnetic Resonance*, 148(2):367–378, February 2001.
- [205] Jizhong Chen and Chunyan Xu. Design and Analysis of the Novel Test Tube Magnet as a Device for Portable Nuclear Magnetic Resonance. *IEEE Transactions on Magnetics*, 43(9):3555–3557, September 2007.