Trajectory Optimization and Machine Learning Radiofrequency Pulses for Enhanced

Magnetic Resonance Imaging


By

Julianna Denise Ianni


Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Biomedical Engineering

December 16, 2017

Nashville, Tennessee


Approved:

William A. Grissom, Ph.D.

Adam W. Anderson, Ph.D.

Bennett A. Landman, Ph.D.

David S. Smith, Ph.D.

E. Brian Welch, Ph.D.

# ACKNOWLEDGMENTS

I am fortunate to have the problem of having almost too many people to thank for their help in reaching this point. I've had the good fortune of being advised by Will Grissom. It is rare to be advised by someone both easy-going and brilliant, and I don't take this for granted. I'm extremely grateful for his vote of confidence in bringing me on board, and that he's entrusted me to bring some wild ideas to fruition. I want to also thank Mark Does, who first piqued my interest in MRI while I was an undergraduate. When I later asked for advice on where I should continue my MR studies, he simply chuckled–to him, Vanderbilt was the only "right" option. I would also like to thank Brian Welch, who has served on my committee and offered helpful feedback both in my dissertation work and my career. I respect Brian for his intellect and especially for his ability to make anyone feel an equal.

My deepest thanks is to my Grandmother, Pat, who has been rooting for me to get my Ph.D. at least since I was five – though I don't think she envisioned it so far from home! I also owe thanks to my Dad, Emile, who is my confidence and from whom I have inherited the perseverance and drive which brought me to this point.

My friends have often been closer to family, and I'm so incredibly grateful for them all. I have had no lack of people to look to for support at Vanderbilt, some even long after their careers have led them elsewhere – Melonie Sexton, Aliya Gifford, Pooja Gaur, Alex Smith, and Megan Poorman in particular. Melonie gave me the extra nudge to go to graduate school, and somehow didn't talk me out of it! Aliya has given her time freely as an informal career coach. Pooja's advice, encouragement, and example have been invaluable. Alex has been a cheerful comrade, and so generously fielded my many questions over the years. Megan has helped keep my sanity and brought laughter into that basement office. Alex and my other first graduate school officemates - Meghan Bowler, Mark Baglia, Oscar Ayala - are a constant source of support and provided the antics to make that first year bearable. There exists an even longer list of those who have helped along the way, one I don't have the room to write here, and I hope I have had opportunity to thank them individually.

There are a few more that I must single out. Jackie has been a most dependable sounding board and always kept me amused. I can't express enough gratitude to Robert Furr, who has helped me begin to take my life back from endometriosis. I owe a lot to Lisa and Michael, who have practically adopted me throughout my time here - their support means so much. Finally, I could not have made it to this point without John, my fiancé, who's supported me through all of life's hurdles. He has made so many sacrifices for me, and brings balance to my life. Most importantly, he has made sure that I take breaks from my work to admire Cousteau, our dog, who is Joy.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Introduction

## 1.1 Objective

Magnetic Resonance Imaging (MRI) enables imaging the human body with a wide variety of manipulatable contrasts and views, and produces images without the use of harmful ionizing radiation. Its ability to acquire high-resolution images in a short time makes it ideally suited for many challenges in medicine today. While most clinical scanners today are 3 Tesla (3T) or 1.5T systems, there is motivation to move to higher field strengths given the many benefits, including increased spectral resolution, better contrast due to longer $T_1$ relaxation times, higher signal-to-noise ratio (SNR), and better parallel imaging performance. The latter two properties can be traded for increased speed and resolution. These strong advantages are the basis for much research occurring at field strengths of 7T and above. Despite these advantages, high field MRI presents several challenges to be addressed before it can see widespread adoption for clinical use. In practice, many imaging techniques require fast imaging readouts and strong flip-angle uniformity to eliminate image artifacts and non-uniform decreases in SNR.

One obstacle in rapid acquisitions is that many fast techniques use non-Cartesian k-space readouts, which are particularly sensitive to trajectory errors caused by gradient eddy currents, delays, and non-ideal gradient amplifier characteristics, which can result in severe image distortions. Another commonly-used trajectory for fast readouts is echo-planar imaging (EPI), which also is very susceptible to trajectory errors because of the fast switching of the gradients. Existing techniques to correct for trajectory errors typically require measurement probes or extra calibration scans. In Chapter 2, we present a calibrationless method called Trajectory Auto-Corrected image Reconstruction (TrACR) to reconstruct images free of trajectory errors. This method uses a flexible gradient-based trajectory op-

timization approach. It jointly estimates images and k-space errors, can be adapted to multiple trajectories, and can be used with multiple existing parallel imaging reconstruction techniques. In Chapter 3, the TrACR algorithm is extended to incorporate the unique trajectory and phase errors encountered in fast EPI acquisitions.

Another challenge encountered in imaging at high field is that $B_1$ field inhomogeneities often critically affect scan results. At field strengths of 7T and above, the wavelength of the transmit radio-frequency (RF) pulse is on the same order of magnitude as the size of the imaged object; this causes spatially varying flip angle and hence spatially varying SNR and changes in tissue contrast. In practice, this inhomogeneity could cause pathology to be mistaken for normal tissue or vice versa.

Several methods exist to address the problem of flip-angle homogeneity, but the most successful approach for general purposes has been the design of patient-tailored pulses for parallel transmission. However, this requires measuring transmit sensitivities with the patient already in the scanner and then optimizing a set of shim weights to excite the desired pattern. To do this requires a design time that is prohibitively long for clinical use, since optimization of these RF shims requires the solution of a non-convex optimization problem, and is therefore computationally complex. Chapter 4 addresses this problem by implementing machine learning techniques to discover relationships between optimized RF shims and the transmit RF field maps for which they are tailored, with minimal calibration data. The algorithm presented can quickly predict patient-tailored RF shims to produce homogeneous excitations, without explicitly simulating the underlying MR physics.

Overall, this work significantly reduces image artifacts due to gradient eddy currents and decreases RF inhomogeneities in high field MR. This is accomplished by employing functions that 1) optimize and exploit redundancies inherent in parallel imaging and 2) exploit redundant information in multi-subject data to learn characteristic relationships between RF pulses and patient-specific parameters.

## 1.2  MR Theory

This chapter gives a brief review of the origin of the MR signal, before introducing some MR excitation and image reconstruction principles, techniques, and shortcomings that will be expanded in later chapters.

An MR scanner works by manipulating nuclear spins. Hydrogen nuclei –protons– are typically targeted in MRI because of their prevalence due to the abundance of water and fat within biological tissue. Spin is a property of atomic nuclei, analogous to the nucleus having an angular momentum. This spin gives rise to a small magnetic moment defined by:

$$\mu = \gamma \mathbf{S}, \tag{1.1}$$

where $\gamma$ is the Larmor frequency of the spins (42.58 MHz/T for hydrogen nuclei). At resting state, nuclear spins within an external magnetic field are oriented according to the Boltzmann distribution:

$$\frac{N^+}{N^-} = e^{-\Delta E/kT}, \tag{1.2}$$

where $N^-$ and $N^+$ represent the number of nuclear spins in lower and upper energy states, respectively, $\Delta E$ is the difference in energy between these states, $k$ is Boltzmann's constant $(1.3805 \times 10^{-23} J/K)$, and $T$ is temperature. The energy difference $\Delta E$ is dependent on the Larmor frequency of the spins $\gamma$ (42.58 MHz/T for protons) and the strength of the main magnetic field $B_0$:

$$\Delta E = \hbar \gamma B_0 \tag{1.3}$$

where $\hbar$ is Planck's constant. At equilibrium, slightly more spins are aligned with the main magnetic field (parallel) than against it (antiparallel). The bulk magnetization of the tissue $M$ (an integral over the magnetic moments $\mu$ of the targeted nuclei within a region of tissue) tends toward a static value $M_0$:

$$M_0 \approx \frac{B_0 \gamma^2 \hbar^2 P}{4kT}, \tag{1.4}$$

where $P$ is the proton density. The bulk magnetization $\mathbf{M} = (M_x, M_y, M_z)$ bears a relationship to a bulk angular momentum $\mathbf{J}$ that mirrors Eq. 1.1:

$$\mathbf{M} = \gamma \mathbf{J}. \tag{1.5}$$

The presence of a dynamically varying external magnetic field generates a torque, according to

$$\frac{d\mathbf{J}}{dt} = \mathbf{M} \times \mathbf{B}, \tag{1.6}$$

and this translates to an expression for the rate of change of the magnetization $\mathbf{M}$:

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} \tag{1.7}$$

Therefore, the time course of the magnetization that is not initially aligned with the magnetic field is a precession about an axis parallel to it with frequency $\omega_0 = \gamma B_0$. If we define the direction of the main magnetic field as $\hat{z}$, then the component of magnetization that is actually measured in MRI is that which lies in the transverse ($xy$) plane. Due to precession, this magnetization rotating in the radiofrequency (RF) range yields a signal that can be detected as a voltage induced in an RF coil placed near the spin system. But all this requires that some spins are initially not aligned with the main magnetic field; otherwise, the RF coil will not detect any signal. RF excitation is the means by which spins are rotated into the transverse plane.

## 1.3   RF Excitation

In order to generate a signal from these spins, we must disturb the system from its equilibrium state– by putting energy into the system. A magnetic field oscillating at the Larmor frequency of the spins can "tip" them such that they have a component perpendicular to the direction of the main magnetic field. This is known as RF excitation. For the following, we

will consider the simplifying assumption that the coordinate frame of reference is rotating about $\hat{z}$ at a frequency of $\omega_0$. When the RF field or transmit $B_1$ field, oriented perpendicular to $B_0$, is applied at this frequency, the spins experience another torque and tip away from the $z$ axis at an angle $\alpha$ determined by the strength of this applied magnetic field and the time $T$ over which it is applied:

$$\alpha(r) = \gamma s(r) \int_0^T B_1(t)dt, \tag{1.8}$$

where $s(r)$ is the transmit coil sensitivity ($B_1^+$) at position $r$, and $B_1(t)$ is an applied RF waveform known as the RF pulse. The signal measured by the receive RF coils is proportional to the transverse (xy) component of the magnetization vector, (which when tipped at angle $\alpha$, begins precessing). This signal will relax back to its original orientation along the z-axis at a rate $R_1 = 1/T_1$, where $T_1$ is the longitudinal relaxation time constant; meanwhile, the signal is also decaying in the $xy$ plane at a rate $R_2 = 1/T_2$, where $T_2$ is the transverse relaxation time constant. In this dissertation, we will neglect these relaxation effects. Ideally, the flip angle– and thus the excited signal profile– is uniform across the spatial volume of interest; however, in reality, the applied field interacts with and is distorted by the imaged object. This distortion manifests as a spatially varying, inhomogeneous flip angle $\alpha$.

### 1.3.1 Transmit Field Inhomogeneity

Transmit field inhomogeneity increases with field strength, and, at field strengths of 7T and above, often critically affects scan results. At 7T, the wavelength of the transmit RF pulse is on the same order of magnitude as the size of the imaged object (e.g. the brain); the resulting interaction of the RF field and tissue deforms the transmit RF field [2], causing spatially varying flip angles and hence spatially varying SNR and tissue contrast. [3–5] Often at high field, this results in a center-brightening effect in the brain, as shown in Fig. 1.1. The effects of RF inhomogeneity are pronounced in many scans, particularly in
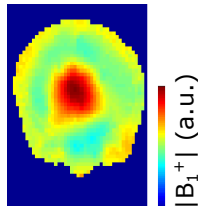
Figure 1.1: At field strengths of 7T and above, the transmit RF ($B_1^+$) field is inhomogeneous and contains a center brightening artifact.

$T_2^*$-weighted blood oxygenation level dependent (BOLD) functional magnetic resonance imaging (fMRI) [6], diffusion tensor imaging (DTI) [7], scans of structures that are located in regions of typically high inhomogeneity, such as the cerebellum and temporal lobes of the brain [8], as well abdomen and whole-body imaging [9].

Several methods exist to mitigate the effects of transmit RF inhomogeneity. Adiabatic pulses, which apply a frequency sweep during an RF excitation envelope, have been employed to combat the inhomogeneity problem [10]. However, these pulses are long and require high peak power, and therefore have relatively high specific absorption rate (SAR). (Limits are placed on SAR to ensure patient safety, as will be discussed in detail in Section 1.3.2.2.) These factors severely limit the use of adiabatic pulses for inhomogeneity correction. Another proposed solution to the inhomogeneity problem at high field is the use of dielectric pads [11, 12]. These pads have a high dielectric constant and are placed around the patient's head in the scanner. They have the effect of altering the field distribution such that it is flatter, especially close to the dielectric pads themselves. These dielectric pads improve field homogeneity in superficial regions of the brain, but the effect is not significant deep within the tissue. It is also simply impractical to make additions to the necessary equipment in an already crowded scanner bore.

### 1.3.2 Parallel Transmission

Parallel transmit (pTx) is a promising alternative to these other methods for mitigating inhomogeneity. Instead of a single coil being used for excitation, multiple coils are used,
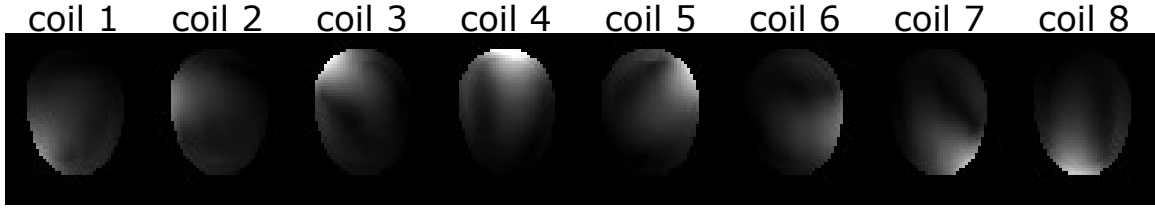
Figure 1.2: Example brain $B_1^+$ maps for a set of 8 transmit channels in a single axial slice.

each with a unique spatial sensitivity profile. Parallel transmission has many benefits, including the ability to drive each coil separately in order to "steer" the field dynamically to achieve different desired excitation patterns. In the past, MR transmit coils have been built with the capability to drive a single element and adjust its amplitude on a patient-by-patient basis based on a calibration of the 90° flip angle. However, having multiple transmit coils provides a number of degrees of freedom that can be used to achieve a more homogeneous excitation.

#### 1.3.2.1  RF Shimming

RF shimming is one method of using pTx for inhomogeneity correction. This is done by separately adjusting the transmit gain and phase and for each coil. This changes the spatial excitation profile of the coils' combined fields, and based on where the coils interfere constructively and destructively, careful choice of complex weights (shims) with which to drive each of the channels will result in a more uniform transmit field [4, 13]. Design of patient-tailored RF shims requires as input an individual's $B_1^+$ maps–showing the spatial sensitivity profile of each coil– acquired through a separate acquisition. Fig. 1.2 shows example brain $B_1^+$ maps for a set of 8 coils. However, the need to acquire these maps on an individual and scan-specific basis means that the patient and scanner must lie idle while the appropriate tailored shims are computed. Fast methods of $B_1^+$ field mapping and patient-tailored pulse design have therefore been a topic of much research (see, e.g. [14–18]).

For the RF shimming problem, we must consider the combined excitation pattern pro-

duced by all coils:

$$m(x) = \sum_{c=1}^{N_c} B_{1,c}^+(x)b_c,$$ (1.9)

where $m(r)$ is the total realized excitation pattern, which is a function of spatial location ($x$), $c$ indexes transmit coils, $N_c$ is the total number of coils, $B_{1,c}^+(x)$ is the measured spatial transmit sensitivity profile ($B_1$ map) corresponding to coil $c$, and $b_c$ is the complex weight (magnitude and phase) applied to coil $c$. We can concatentate the individual coil $B_1^+$ maps and RF shim weights and rewrite Eq. 1.9 as:

$$\mathbf{m} = \begin{bmatrix} \mathbf{B}_{1,1}^+ & \mathbf{B}_{1,2}^+ & \cdots & \mathbf{B}_{1,N_c}^+ \end{bmatrix} \begin{bmatrix} b_1 \\ \vdots \\ b_{N_c} \end{bmatrix} = \mathbf{Ab}$$ (1.10)

From here on, we will refer to only the full multicoil matrices above, and drop the subscripts for succinctness. The RF shimming problem can then be posed to minimize the difference between a desired excitation $\mathbf{m}$ (typically a vector of ones for a uniform profile) and the weighted $B_1^+$ maps:

$$\underset{\mathbf{b}}{\operatorname{argmin}} \|\mathbf{m} - \mathbf{Ab}\|_W^2$$ (1.11)

where $W$ are spatially dependent weights that may select only samples within the brain. RF shimming attempts to minimize the cost in Eq. 1.11 such that transmit sensitivity is as uniform as possible in a region of interest.

While shims can be obtained by directly minimizing Eqn. 1.11, in many cases, the phase of the desired excitation is not important to the design problem. This is true for cases in which only magnitude images are required. Therefore, it is possible to gain additional degrees of freedom and achieve a more homogeneous excitation by instead minimizing the difference between a desired excitation and the *magnitude* of the excited profile, as follows: [19, 20]

$$\underset{\mathbf{b}}{\operatorname{argmin}} \|\mathbf{m} - |\mathbf{Ab}|\|_W^2.$$ (1.12)

Eqn. 1.12 is known as the magnitude least squares (MLS) formulation of the RF shim problem. Although this formulation trades variation in phase for a more homogeneous global optimum, it introduces another problem in that it is non-convex, and therefore it can be difficult to solve in a reasonably efficient manner or to ensure a global optimum is found.

### 1.3.2.2  SAR and RF Power

Additionally, Eq. 1.12 only captures part of the problem; it is also necessary that designed RF shims meet constraints on specific absorption rate (SAR) in order to prevent excessive heat deposition in the tissue. SAR is a measure of the rate of RF energy absorbed in a given mass of tissue, and typically has units of W/kg . High SAR can cause tissue damage from RF heating, a serious concern for RF design, but it is difficult to explicitly design pulses to meet SAR constraints. It is desirable to directly constrain the temperature reached by the body tissue in the context of the MR scan, as tissue damage from RF fields is directly dependent on temperature. To date, however, it has been a challenge to efficiently and accurately measure local temperature and tissue electrical properties needed to simulate the thermodynamics involved; many factors affect local body temperature, including the ambient temperature, metabolism, and dissipation of heat that is dependent on chararacteristics of the tissue itself, as well as blood flow and the body's own thermoregulatory response. Therefore, SAR is used as an indirect safety measure. RF pulses are designed to stay well within a SAR safety factor set by the U.S. Food and Drug Administration (FDA) at 3 W/kg over 10 minutes, globally averaged over the head, and 4 W/kg over 15 minutes over the body [21].

Other requirements must also be satisfied in regards to the capabilities of the amplifier; the amplitude and average power for each channel must be limited. Regularization is an alternative to implementing strict constraints for SAR and RF power. This can be incorporated into the tailored shim design problem in Eq. 1.12 with a regularization term, as in the

following:

$$\underset{\mathbf{b}}{\text{argmin}} \ \|\mathbf{m} - |\mathbf{Ab}|\|_W^2 + \lambda \, \mathrm{R}(\mathbf{b}), \tag{1.13}$$

where $\lambda$ controls the strength of the regularization $\mathrm{R}(\mathbf{b})$.

Solving Eqn. 1.13 is time-consuming, and because it is non-convex and sensitive to initialization, there is no guarantee that the arrived-at solution will be globally optimal. Chapter 4 will introduce an efficient alternative – a machine learning approach to the problem that sidesteps the computational and scan burdens required for conventional RF shimming.

## 1.4   Image Acquisition

Once spins have been excited, or tipped into the transverse plane–ideally uniformly and in a SAR-efficient manner– they can be measured. However, if the transmit RF ($B_1$) field and main magnetic field ($B_0$) are the only fields present, the signal from all spatial locations (for the region over which the coil is sensitive) will be superimposed, and there will be no image to retrieve; there must be a way to localize spins in order to acquire an image.

### 1.4.1   Spatial Localization

Gradient fields are typically introduced in order to localize nuclear spins so that an image can be reconstructed. These are fields of the form $\mathbf{G} = (G_x, G_y, G_z)$ that are linearly proportional to position in each dimension, such that:

$$B_0 = B_0 + \mathbf{G}(t) \cdot \mathbf{r}, \tag{1.14}$$

where $\mathbf{r} = (x, y, z)$ is the position vector. The signal that is acquired becomes:

$$S(t) \propto \int \int \int m(x, y, z) e^{\left[ -i\gamma \int_0^T \mathbf{G}(t) \cdot \mathbf{r} \right] dt} \, dx \, dy \, dz, \tag{1.15}$$

10

where $m(x,y,z)$ is the spatially distributed magnetization (from the 3-dimensional target imaged object), and $t = 0$ corresponds to the beginning of the acquisition window. The received signal is the Fourier transform of the imaged object.

### 1.4.2   Trajectories and Transforms

Now, with a signal that has been excited, and a way to localize spins, the signal must be read out or acquired. The signal is read out in the Fourier domain, i.e. it is acquired at "locations" in k-space— or frequency space— along some trajectory. In Eqn. 1.15, we can substitute:

$$
\begin{aligned}
k_x(t) &= \frac{\gamma}{2\pi} \int_0^T G_x(t)dt \\
k_y(t) &= \frac{\gamma}{2\pi} \int_0^T G_y(t)dt \\
k_z(t) &= \frac{\gamma}{2\pi} \int_0^T G_z(t)dt,
\end{aligned}
\tag{1.16}
$$

where $\mathbf{k} = (k_x, k_y, k_z)$ is the 3-dimensional k-space trajectory in units of cycles per unit distance (typically $cm^{-1}$). Trajectories are used both for excitation and acquisition; this dissertation will discuss their application in acquisition only.

Readout trajectories for MRI have typically followed a Cartesian pattern through k-space, with sample points acquired at equispaced points in 2 and 3 dimensional k-space. Fig. 1.3a shows an example Cartesian trajectory. This allows image reconstruction with a simple inverse Fourier transform of the data, which can be performed efficiently using fast Fourier transforms (FFTs). However, sampling on a Cartesian grid is slow and does not allow for undersampling in an optimal manner in terms of the resulting signal-to-noise ratio (SNR).
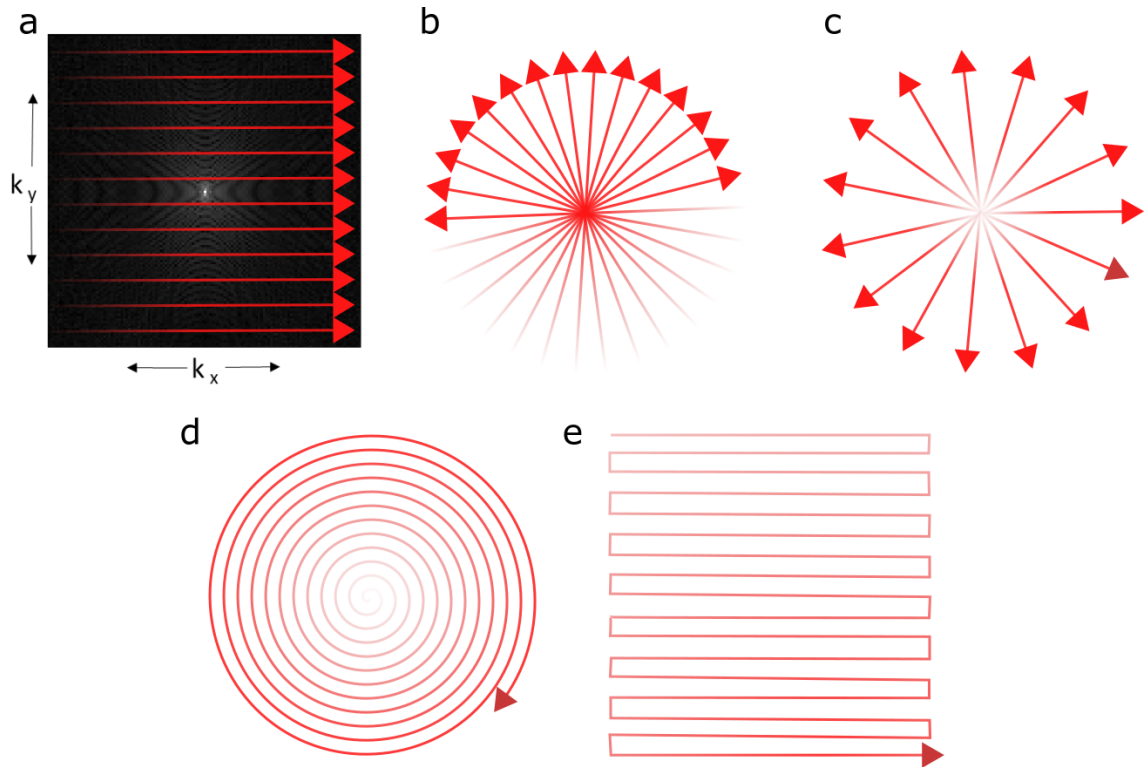
Figure 1.3: Different k-space trajectories used for data acquistion are shown. a) Cartesian acquisitions are the most common way to sample k-space. Non-Cartesian trajectories (b-d) such as (b) radial, (c) center-out radial and (d) spiral, are typically used for fast acquisitions. (e) Echo-planar imaging (EPI) trajectories are also used for fast readouts, but are designed to still acquire samples on a Cartesian grid.

#### 1.4.2.1 Non-Cartesian Readout Trajectories

Non-Cartesian trajectories include radial, center-out radial, and spiral, (examples shown in Fig. 1.3b-d) among many others. These were introduced as methods of increasing the speed of traversal of k-space while maintaining high SNR-efficiency. They are designed generally to oversample the center of k-space where signal is highest. They are often used in dynamic MRI, including functional brain imaging [22], cardiac imaging [23, 24], and in applications where the MR signal is short-lived, such as sodium imaging [25] and ultrashort echo time (UTE) imaging [26].

#### 1.4.2.2 NUFFTs

Non-Cartesian acquisitions require a more complicated image reconstruction than Cartesian trajectories; since the data are not acquired on a Cartesian grid, direct inverse FFTs cannot be applied to arrive at an image. The exact inverse operation to go from acquired non-Cartesian data to image is the inverse non-uniform discrete Fourier transform or NUDFT, which computes an image $f$ from the data $d$ in the Fourier domain:

$$f_j = \sum_{i=0}^{N-1} c_i d_i e^{-\iota 2\pi x_j \omega_i},$$

$$0 \leq j \leq N-1,$$

(1.17)

where $i$ indexes k-space frequencies, $j$ indexes spatial locations, $c_i$ are sampling-density compensation weights, $f_j$ is the complex image, $x$ are spatial positions, and $\omega_i$ are the k-space frequencies. However, this operation is too computationally complex for most applications; therefore, the inverse non-uniform FFT (NUFFT) is used instead. This involves gridding the non-Cartesian data to an oversampled Cartesian grid by convolving it with a kernel, then performing an inverse FFT and multiplying with a deapodization function to compensate for the effects of the kernel convolution in k-space. Applying density compensation is also necessary when performing an NUFFT, but in iterative image reconstructions

the need for this obviated. In Chapter 2, NUFFTs are employed for non-Cartesian image reconstructions; additionally, Chapter 3 will introduce an alternative to the NUFFT for EPI reconstructions, compensating for trajectory delays.

### 1.4.2.3 Gradient Eddy Currents

Compared to Cartesian k-space readouts, non-Cartesian trajectories are sensitive to trajectory errors caused by gradient eddy currents, delays, and non-ideal gradient amplifier characteristics, which can result in severe image distortions. This is due to the fact that non-Cartesian acquisitions typically use fast readouts and require rapid switching of the gradients. Gradient eddy currents produced on conducting structures in the scanner (e.g. gradient and RF coils, the imaged object) are the result of Faraday induction due to the fast switching of the gradients. These gradient eddy currents can be modeled by the derivative of the nominal gradient convolved with exponential decay functions of varying time constants $\tau$ [27] :

$$g_e(t,\tau) = \frac{dG(t)}{dt} * \left\{ H(t)e^{-t/\tau} \right\}, \tag{1.18}$$

where $G$ is the time-varying gradient waveform, $H(t)$ is the Heaviside step function, and $*$ denotes the convolution operation. These eddy currents have the effect of applying a low-pass filter to the requested gradient waveforms. The resulting errors are pronounced in non-Cartesian trajectories, particularly because they often involve sampling on the ramps of the gradients. Figure 1.4a shows an example gradient used for a single line of a center-out radial trajectory. The desired gradient waveform is shown in dashed green. Without compensation, gradient eddy currents distort the waveform so that it looks like the dashed black waveform. Modern scanners typically apply gradient pre-emphasis, which attempts to correct for some of these gradient errors. A pre-emphasized waveform is essentially a high-pass filtered version of the nominal or desired gradient waveform (blue line in Fig. 1.4a), such that it compensates in advance for the low-pass filtering effect of gradient eddy currents. However, the gradient pre-emphasis methods that scanners typically use to com-
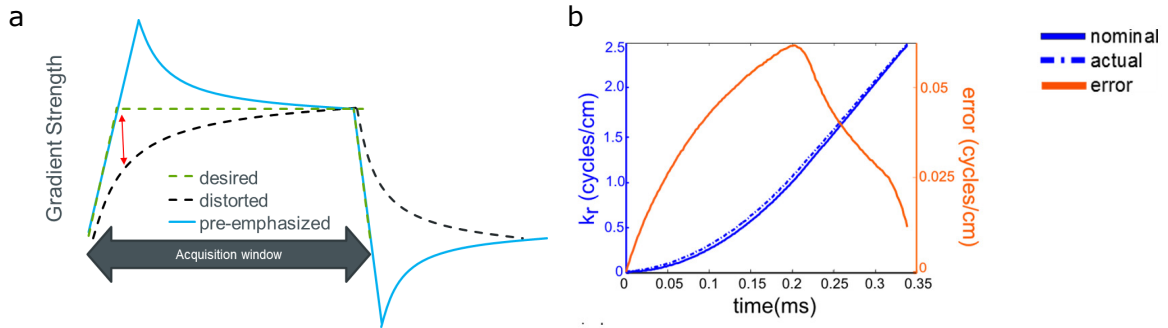
Figure 1.4: In non-Cartesian acquisitions, gradient eddy currents can cause significant trajectory errors, as shown in this example for a center-out radial trajectory. a) Shown are the nominal (desired) gradient waveform (dashed green), the waveform distorted by eddy currents (dashed black) and the pre-emphasized waveform. b) The corresponding k-space trajectory (solid blue) is shown, as well as the (actual) trajectory distorted by eddy currents (dashed blue), and the error between the two (orange).

pensate for these types of eddy current errors are often targeted to optimize gradient trapezoids used for Cartesian imaging only, and are limited in the scope of errors for which they can compensate. Fig. 1.4b shows an example of center-out radial trajectory errors for a single line of the readout; the solid and dashed blue lines correspond to the nominal and distorted (actual) k-space trajectory, and the orange line shows the error between the two. Residual uncompensated trajectory errors such as these can result in severe image distortions.

### 1.4.2.4  Echo-Planar Imaging

Echo-planar imaging (EPI) is another fast-imaging technique which can yield significant image distortions due to the fast switching of the gradients. Samples of adjacent lines in k-space are acquired directly by reversing the direction of the gradient, instead of rewinding after each line, as shown in Fig. 1.3e. The errors resulting from effective delays applied to alternating lines in k-space produce replicas of the imaged object, known as Nyquist ghosts. Furthermore, the small magnitude of the phase-encoding blips relative to the gradient-switching in the fast-readout dimension results in much larger EPI trajectory errors in the readout dimension than in the phase-encoding dimension. In addition

15

to trajectory errors, EPI acquisitions are susceptible to global phase differences between odd and even lines or multiple shots. A variety of sources contribute to these errors, including patient respiration and motion, heating of the magnet, gradient and shim coils, and eddy currents in the main magnet and shim coils [28]. These errors are normally corrected by estimating a k-space shift and phase difference that can be applied to all odd or even lines, based on scanner calibration measurements. Jesmanowicz et al. demonstrated that EPI phase errors could be estimated by playing out several lines of an EPI trajectory without the phase blips and calibrating out the difference in the repeatedly acquired line [29]. Similarly, one can also integrate a calibration into the existing EPI sequence, in which the center of k-space is acquired twice, once with an odd echo and once with an even one. Hu and Le presented a method to calibrate out the differences of two acquisitions, one that is shifted by a line such that odd echoes in the first acquisition correspond to even in the second [30]. All of these methods, however, are limited in the fact that they require the acquisition of reference data. A new method for correcting trajectory and phase errors in EPI trajectories without calibration data is discussed in detail in Chapter 3.

## 1.4.3   Parallel Imaging Reconstructions

Another common technique for efficient signal acquisition in MRI is parallel imaging, in which several coils placed around the imaged volume are used to sample the signal simultaneously. Since the coils are sensitive to different regions in the volume, they each collect a uniquely weighted version of the underlying signal. This redundant sampling enables image reconstruction that exploits this fact and allows higher fidelity recovery of the underlying signal, providing a signal-to-noise ratio (SNR) benefit. Further, it enables undersampling of the signal in k-space with recovery of the signal via the spatial encoding provided by the separate coils. In turn, this allows increased temporal resolution and provides a basis for reduced-artifact image reconstructions that take advantage of redundantly sampled data over multiple coils.

There are several methods of exploiting the redundancies in multicoil datasets. The two methods specifically used in this work are SENSE and SPIRiT reconstructions, and these are discussed in detail in the following sections.

### 1.4.3.1 SENSE

A brief summary of the SENSE (sensitivity encoding) technique used for the parallel imaging reconstructions in this work follows. A more detailed account can be found in the work of Pruessmann et al.[31, 32]. When we measure signal along an arbitrary trajectory in k-space, the corresponding signal model is:

$$y_c(\vec{k}) = \int_{VOI} p(\vec{r}) e_c(\vec{k}, \vec{r}) d\vec{r} \tag{1.19}$$

where $y_c(\vec{k})$ is the k-space signal measured by coil $c$ as a function of k-space position, $p(\vec{r})$ is a function of the imaged object and imaging parameters at spatial coordinates $\vec{r}$, and $e_c(\vec{k}, \vec{r})$ are encoding functions. To reconstruct an image, these encoding functions are discretized and form the elements of an encoding matrix $E$ defined as follows:

$$E_{(ci),j} = e^{i\vec{k}_i \cdot \vec{r}_j} s_{cj} \tag{1.20}$$

where $\vec{k}_i$ is the $i^{th}$ sampled k-space coordinate, and $s_{cj}$ is coil $c$'s sensitivity at $\vec{r}_j$. Image reconstruction then follows a linear model:

$$\mathbf{f} = \mathbf{G}\mathbf{y} \tag{1.21}$$

where $\mathbf{f}$ represents the image object, and the reconstruction matrix $\mathbf{G}$ is the pseudoinverse:

$$\mathbf{G} = (\mathbf{E}^H \Psi^{-1} \mathbf{E})^{-1} \mathbf{E}^H \Psi^{-1} \tag{1.22}$$

where $\Psi$ is the sample noise covariance matrix, which here is used to optimize SNR in the reconstruction. However, direct solution of this equation is extremely computationally expensive, and therefore, iterative methods are typically employed. Since the noise is complex additive white Gaussian, the maximum-likelihood solution to the non-Cartesian image reconstruction problem is the minimizer of a cost $J$ as follows:

$$J(\boldsymbol{f}) = \frac{1}{2} \sum_{c=1}^{N_c} \sum_{i=1}^{N_k} d_{ci} \left| y_{ci} - \sum_{j=1}^{N_s} e^{-\imath 2\pi(\vec{k}_i)\cdot\vec{r}_j} s_{cj} f_j \right|^2 \tag{1.23}$$

where $N_c$ is the number of receive coils, the $d_{ci}$ are optional coil- and k-space location-dependent weights, which compensate for sampling-density, $\vec{r}_j$ is the $j$th spatial coordinate in the image, and $s_{cj}$ is coil $c$'s receive sensitivity at $\vec{r}_j$. This minimization problem can be solved using a conjugate gradient (CG) routine. The reconstruction requires accurate measurement of coil sensitivities. The resulting image $\hat{f}_j$ will be a *body-coil* or *ground truth* image, on top of which the known or measured coil sensitivities are multiplied to get back the full-coil image set. The technique enables corrections of undersampled data, which can lead to pixel aliasing. While the aliased spatial pattern is the same for all coils, the aliasing weights are different for each coil, so coil sensitivity information can be exploited to account for undersampling errors.

### 1.4.3.2 SPIRiT

Iterative self-consistent parallel imaging reconstruction (SPIRiT) is a method that utilizes the redundancies in a multicoil MR dataset to minimize errors, particularly in non-Cartesian and undersampled image reconstructions [33]. Like the earlier GRAPPA method [34], SPIRiT introduces a calibration kernel that enforces consistent relationships between each acquired k-space sample and its surrounding neighborhood. The SPIRiT kernel is iteratively applied to the reconstructed, coil-dependent, fully-sampled k-space data arrays to ensure that data can be synthesized from its neighbors in a manner consistent with the

parallel imaging signal model. The SPIRiT kernel is calibrated on multi-coil data at the center of k-space; it develops a set of weights that minimize the difference between the acquired data points and the *same* points synthesized by the weighted sum of neighboring data in k-space (including that from other coils). Again, the data are represented by a linear model:

$$\mathbf{y} = \mathbf{Gf} \tag{1.24}$$

where $\mathbf{G}$ is a reconstruction system matrix that relates the reconstructed image $\mathbf{f}$ to the collected k-space data $\mathbf{y}$. This is a requirement that reconstruction be consistent with the acquired data. Additionally, SPIRiT requires consistency between acquired and synthesized k-space data:

$$\mathbf{x} = \mathbf{Cx} \tag{1.25}$$

where $\mathbf{x}$ represents acquired and non-acquired k-space data, and $\mathbf{C}$ performs a series of convolution operations based on the derived set of SPIRiT calibration weights. Figure 1.5 demonstrates, for Cartesian and non-Cartesian trajectories, which k-space points may influence the data consistency and kernel consistency constraints for a given sample.

Image reconstruction is then formulated as a minimization problem that enforces both data consistency and self-consistency with the SPIRiT calibration kernel:

$$J(\mathbf{f}) = \frac{1}{2} \sum_{c=1}^{N_c} \sum_{i=1}^{N_k} d_{ci} \left| y_{ci} - \sum_{j=1}^{N_s} e^{-\imath 2\pi (\vec{k}_i) \cdot \vec{r}_j} f_{cj} \right|^2 + \frac{\lambda}{2} ||\mathbf{Sf}||^2 \tag{1.26}$$

where $\mathbf{S}$ is the SPIRiT calibration kernel, and $\lambda$ controls the strength of the SPIRiT regularization (i.e. the relative importance of consistency between synthesized and acquired k-space).

Both the SPIRiT and SENSE reconstruction methods lend themselves to modifications further exploiting data redundancy. In that vein, Chapter 2 introduces a similar cost function to that employed here, in order to allow the estimation of trajectory errors typically caused
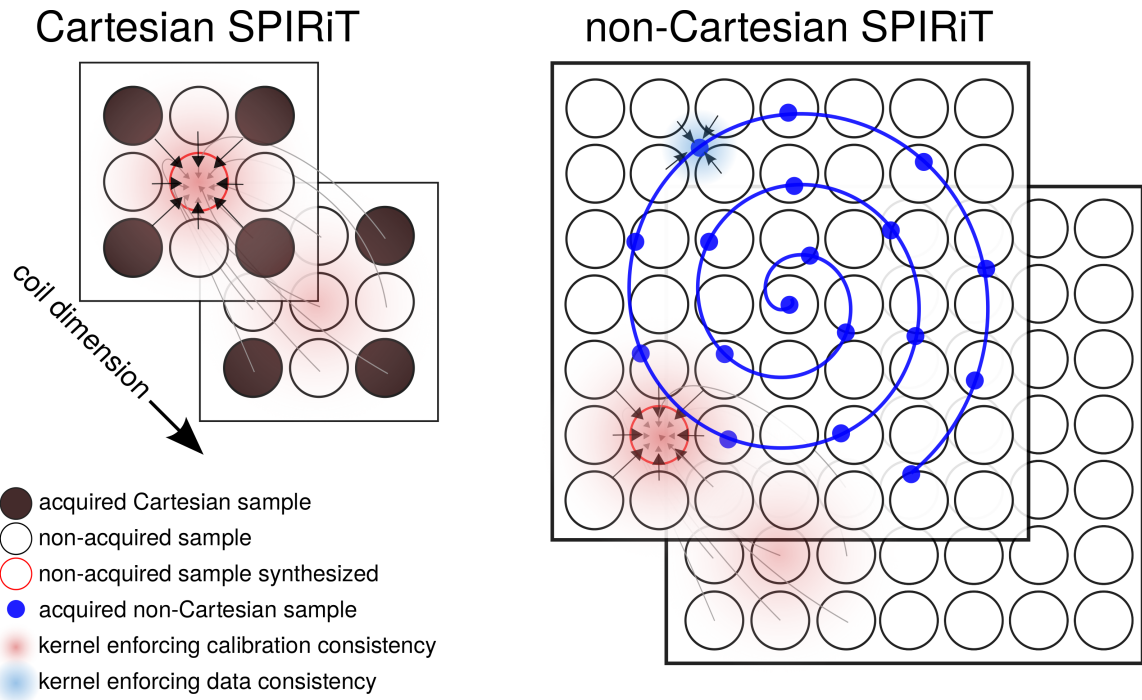
Figure 1.5: In Cartesian SPIRiT, a k-space sample not acquired (red circle) is synthesized from a weighted kernel applied to those acquired (solid black) in the surrounding neighborhood, including other coils. Arrows indicate the samples that contribute to this point. For non-Cartesian SPIRiT reconstruction, data consistency is enforced between sampled non-Cartesian points and the synthesized points on the Cartesian grid, and calibration consistency is enforced with the Cartesian data synthesized from the surrounding Cartesian k-space points.

by gradient eddy currents, and Chapter 3 extends this method to allow the estimation of additional trajectory and phase errors in EPI trajectories.

Chapter 2

Automatic Correction of Non-Cartesian Trajectory Errors

## 2.1 Introduction

This chapter proposes a new method for automatic correction of trajectory errors for non-Cartesian acquisitions, without the need for additional measurements or hardware. Non-Cartesian k-space readout trajectories are used in several MRI applications, including functional brain imaging [22], cardiac imaging [23, 24], sodium imaging [25], and ultra-short echo time (UTE) imaging [26]. However, compared to Cartesian k-space readouts they are particularly sensitive to trajectory errors caused by gradient eddy currents, delays, and non-ideal gradient amplifier characteristics, which can result in severe image distortions. Modern scanners use gradient coil shielding and waveform pre-emphasis to prospectively avoid significant trajectory deviations in typical MR acquisitions. However, those methods are limited in terms of the magnitude and temporal dynamics of the errors for which they can compensate, and non-Cartesian trajectory gradient waveforms can easily push past those limits. Consequently, much research has focused on developing methods to compensate for non-Cartesian trajectory errors retrospectively.

One approach to retrospective trajectory error correction is to measure or predict the erroneous trajectory waveforms and use them in place of the nominal trajectory for image reconstruction. Mason et al. [35] proposed an early gradient waveform measurement technique using a point phantom to localize and measure the phase progression of a set of spins. To eliminate the need for precise placement of a physical phantom, Duyn et al. [36] and Zhang et al. [37] proposed measuring the phase accrual of spins due to a gradient waveform of interest by performing slice selection along the same axis as the encoding gradient. Gurney et al. [38] later introduced a modification to Duyn's method to additionally allow

the measurement of $B_0$ eddy currents. Magnetic field monitoring is another measurement approach that uses susceptibility-matched NMR probes placed around the imaging volume in the scanner and has the advantage of flexibility, in that it can be performed concurrently with any scan protocol [39]. The primary disadvantage of this approach is that it requires specialized hardware to be situated inside the already-crowded magnet bore. All gradient measurement approaches share the disadvantage that they cannot be performed retroactively as a post-processing step, for example after attempts at image reconstruction without corrections reveal the presence of artifacts in previously-acquired data. Predictive methods have also been proposed based on the calibration of a gradient system model; subsequently this model may be applied to predict errors for new input waveforms that might differ in terms of the orientation of the imaging plane [40] or the trajectory itself [41–43]. All of these techniques require calibration scans that can lengthen overall examination time. Some require only one-time or periodic calibration, but cannot predict transient gradient errors such as those caused by variations in the gradient system response with gradient coil temperature increases. Predictive methods are also fundamentally limited by the models on which they are based. For example, linear time-invariant models cannot predict errors due to gradient amplifier nonlinearity, and most do not account for concomitant gradient terms due to the difficulty in measuring them.

Several recently-proposed methods for trajectory error correction do not require additional measurements, calibration scans or hardware [44–46], and focus on correcting errors in radial scans. These are iterative methods that estimate trajectory errors from the k-space data themselves, and work by exploiting a) data redundancy resulting from oversampling in the center of k-space, which is a universal characteristic of non-Cartesian trajectories in use today, and b) data redundancy provided by parallel imaging. They are members of a broader class of methods that aim to jointly estimate both images and other quantities such as receive sensitivity maps [47] and off-resonance maps [48] from k-space data. Deshmane et al. [44] proposed a method that iteratively shifts data in k-space with the goal of

23

finding the set of shifts that produces the highest sum-of-squares (SOS) signal; the set of best shifts is then used to update the k-space trajectory used for reconstruction. Wech et al. [46] proposed a method that iteratively shifts radial projections in k-space, choosing the direction of those shifts based on the concordance of the resulting k-space data with the remainder of the dataset. These methods have the advantage that transient gradient errors can be captured retrospectively without the need for additional measurements. However, they are limited in their range of potential applications because the need to select specified shift directions and magnitudes would make for a large and potentially intractable combinatorial solution space when applying them to trajectories other than radial.

In the following, we propose a more general method to reconstruct images free of trajectory errors, called TRajectory Auto-Corrected image Reconstruction (TrACR), that is based on the same basic idea as the aforementioned measurement-free methods but uses a flexible gradient-based trajectory optimization approach. The method jointly estimates images and k-space errors, can be adapted to multiple trajectories, and can be used with multiple existing non-Cartesian parallel imaging reconstruction techniques. The method is evaluated with in vivo 7 Tesla brain data from radial, center-out radial, and spiral acquisitions in five human subjects. Performance of the method is investigated as a function of k-space acceleration factor and the number of receive coils. Center-out radial and spiral trajectory error estimates are validated against trajectory measurements.

## 2.2   Theory

### 2.2.1   Problem Formulations

The TrACR method is formulated as a joint estimation of images and k-space trajectory errors, using extensions of the cost functions for SENSE [31, 32] and SPIRiT [33] non-Cartesian parallel imaging reconstruction to incorporate trajectory errors as additional

variables. The cost function used for SENSE reconstruction is:

$$\Psi(\boldsymbol{f}, \Delta \vec{\boldsymbol{k}}) = \frac{1}{2} \sum_{c=1}^{N_c} \sum_{i=1}^{N_k} d_{ci} \left| y_{ci} - \sum_{j=1}^{N_s} e^{-\iota 2\pi(\vec{k}_i + \Delta \vec{k}_i)\cdot \vec{r}_j} s_{cj} f_j \right|^2, \qquad (2.1)$$

where $\boldsymbol{f}$ is a length-$N_s$ vector of image samples to be reconstructed, $\Delta \vec{\boldsymbol{k}}$ is a length-$N_k$ vector of trajectory errors to be estimated, $N_c$ is the number of receive coils, the $d_{ci}$ are optional coil- and k-space location-dependent weights, $y_{ci}$ is coil $c$'s $i$th k-space data sample, $\iota = \sqrt{-1}$, $\vec{k}_i$ is the nominal $i$th k-space location, $\vec{r}_j$ is the $j$th spatial coordinate in the image, and $s_{cj}$ is coil $c$'s receive sensitivity at $\vec{r}_j$. In this work the $d_{ci}$ are used to apply k-space density compensation to accelerate algorithm convergence. The cost function used for SPIRiT reconstruction is:

$$\Psi(\boldsymbol{f}, \Delta \vec{\boldsymbol{k}}) = \frac{1}{2} \sum_{c=1}^{N_c} \sum_{i=1}^{N_k} d_{ci} \left| y_{ci} - \sum_{j=1}^{N_s} e^{-\iota 2\pi(\vec{k}_i + \Delta \vec{k}_i)\cdot \vec{r}_j} f_{cj} \right|^2 + \frac{\lambda}{2} \|\boldsymbol{S f}\|^2, \qquad (2.2)$$

where $\boldsymbol{f}$ is now a length-$N_s N_c$ vector of images for all coils, and $\frac{\lambda}{2} \|\boldsymbol{S f}\|^2$ is the SPIRiT regularization, where $\lambda$ is a user-specified regularization parameter and $\boldsymbol{S}$ is the SPIRiT operator. Equation 2.2 is an extension of Eq. 10 in Ref. [33]. The individual coil images can be combined after SPIRiT reconstruction using any coil combination method [33]. Sum-of-squares coil combination was used in this work. In both the SENSE and SPIRiT cases we model the k-space trajectory errors $\Delta \vec{k}_i$ as a sum of weighted error basis functions:

$$\Delta \vec{k}_i = \sum_{b=1}^{N_b} \vec{e}_{bi} w_b = \left\{ \vec{\boldsymbol{E}} \boldsymbol{w} \right\}_i, \qquad (2.3)$$

where $N_b$ is the number of error basis functions $\vec{e}_b$. In order to minimize the required number of error basis functions, we construct them in a trajectory-dependent manner. Useful error basis construction approaches for radial, center-out radial, and spiral trajectories are described further in the Methods.

### 2.2.2 Algorithm

TrACR is an iterative method based on an alternating minimization approach, in which one of the parameters $f$ or $\Delta\vec{k}$ is kept fixed while the other is updated. Accordingly, the algorithm comprises an outer loop which in each iteration invokes an $f$ update, followed by a $\Delta\vec{k}$ update. For fixed $\Delta\vec{k}$, the cost functions in Eqs. 2.1 and 2.2 reduce to the original non-Cartesian SENSE and SPIRiT reconstruction problems and are typically minimized with respect to $f$ using the Conjugate Gradient (CG) algorithm [49]. To update the k-space error weights $w$ in Eq. 2.3 for fixed $f$, a nonlinear Polak-Ribière CG algorithm is used [49]. Each iteration of that algorithm requires the derivatives of the cost function with respect to $w$, in order to calculate the next search direction. Since by Eq. 2.3 each error weight $w_b$ affects all k-space trajectory dimensions, the total derivative for each weight will comprise a sum over the k-space dimensions. For the SENSE reconstruction problem, the contribution to the derivative of $w_b$ from the $k_x$-dimension is:

$$\left\{ \frac{\partial \Psi}{\partial w_b} \right\}_x = \sum_{c=1}^{N_c} \sum_{i=1}^{N_k} \Re \left\{ \sum_{j=1}^{N_s} -\imath 2\pi x_j e_{bi}^x e^{\imath 2\pi(\vec{k}_i + \Delta\vec{k}_i)\cdot\vec{r}_j} d_{ci} s_{cj}^* f_j^* r_{ci} \right\}, \tag{2.4}$$

where $\Re$ denotes the real part of the complex number in the braces, $^*$ indicates a complex conjugate, and $r_{ci}$ is the residual:

$$r_{ci} = y_{ci} - \sum_{j=1}^{N_s} e^{-\imath 2\pi(\vec{k}_i + \Delta\vec{k}_i)\cdot\vec{r}_j} s_{cj} f_j. \tag{2.5}$$

Once the derivatives for each k-space dimension are computed, they are summed to obtain the total derivative for each weight, and the gradient vector of collected derivatives for all weights is returned to the CG algorithm. The derivatives of the SPIRiT cost function are obtained by replacing $s_{cj} f_j$ with $f_{cj}$ in Eqs. 2.4 and 2.5. The TrACR algorithm alternates between image and k-space error weight updates until a stopping criterion is met.

## 2.3    Methods

### 2.3.1    Algorithm Implementation

The TrACR algorithm was implemented in MATLAB 2014a (The Mathworks, Natick, MA, USA) on a desktop PC with an Intel Xeon E3-1240 3.4 GHz CPU (Intel Corporation, Santa Clara, CA, USA) and 16 GB of RAM. Image updates were initialized with zeros each outer iteration to avoid noise amplification. Except where otherwise noted, all images were reconstructed using MATLAB's `lsqr` function, with a fixed tolerance of $10^{-2}$, both inside and outside the TrACR algorithm. This allowed the number of CG image iterations in each image update step to vary as needed; typically 2 to 10 iterations were used. All non-uniform discrete Fourier transforms were computed using a non-uniform fast Fourier transform (NUFFT) algorithm [50]. Density compensation weights ($d_{ci}$ in Eqs. 2.1 and 2.2) were calculated using the method of Zwart et al. [51] using the nominal trajectories. For SPIRiT image reconstructions, the regularization parameter $\lambda$ (Eq. 2.2) was fixed to 10% of the median of the absolute value of the k-space data. To enable direct comparison of SENSE and SPIRiT reconstructions, the SPIRiT kernel was calibrated using images obtained by applying the receive sensitivities measured for SENSE to a sum-of-squares Cartesian reconstruction. The CG algorithm for the k-space updates used a maximum of 5 iterations and a backtracking line search ([52], p. 464) with a maximum allowed trajectory change in one CG iteration of 1/FOV, where FOV is the reconstructed field-of-view. The outer loop of the TrACR algorithm was stopped when the k-space backtracking line search returned a zero step size in its first iteration. MATLAB code to implement the algorithm and a demonstration with an in vivo radial dataset are available at https://bitbucket.org/wgrissom/tracr/downloads.

### 2.3.2 Experiments

In vivo experiments were performed at 7 Tesla (Philips Achieva, Philips Healthcare, Best, Netherlands) using a quadrature volume coil for excitation and a 32-channel head coil (Nova Medical, Wilmington, MA, USA) for reception. Scans were performed in 5 healthy volunteers with approval from the Institutional Review Board of Vanderbilt University. Data were collected using 3 non-Cartesian trajectories: golden-angle (GA) radial, center-out radial, and multi-shot spiral, detailed further below. Cartesian scans were also collected and used to synthesize a body coil image and an estimate of the sum-of-squares receive sensitivity using a polynomial fit; the receive sensitivity was divided out of the reconstructed non-Cartesian images. Coil sensitivity measurements were collected for SENSE reconstructions. All scans were gradient echo sequences with repetition time and echo time matched for all trajectories, at 200 ms and 7.9 ms, respectively, and with 2.5 mm slice thickness. Center-out radial and spiral trajectory measurements for validation were collected in a spherical phantom for each scan session using a modified Duyn method [36, 38]. Trajectory measurements began 1 ms prior to the expected start of the gradient waveforms in order to capture components generated by the scanner's gradient pre-emphasis.

The GA radial trajectory comprised 201 projections, each containing 256 sample points. Trajectory-specific acquisition parameters were: readout duration 0.46 ms, maximum gradient amplitude 16.1 mT/m and maximum gradient slew rate 7.9 T/m/s, water/fat shift 0.741 pixels. The center-out radial trajectory comprised 402 projections, each containing 170 sample points. Trajectory-specific acquisition parameters were: readout duration 0.34 ms, maximum gradient amplitude 25.5 mT/m, maximum gradient slew rate 114 T/m/s. The spiral trajectory comprised 16 shots of length 5.7 ms. Maximum gradient amplitude was 14.3 mT/m and the maximum gradient slew rate was 70 T/m/s. The trajectory was designed using Brian Hargreaves' spiral design toolbox [53]. The resolution of each trajectory matched that of the $128 \times 128$ reconstruction grid. The reconstructed field of view

was 25.6 cm and all trajectories were designed to sample k-space to a maximum frequency of $\pm 2.5$ cycles/cm. Where indicated, the data were coil-compressed prior to image reconstruction using singular value truncation [54].

### 2.3.2.1   Incorporating $B_0$ inhomogeneity correction

An additional experiment was performed to demonstrate the capability of using TrACR with off-resonance correction incorporated into the reconstruction system matrix as follows. A $B_0$ map was measured in the first subject. Reconstructions were then performed incorporating off-resonance correction using the time-segmentation method of Fessler et. al [1]. In this method, the system matrix is split up into a series of $L$ time segments The off-resonance term is approximated as:

$$e^{-\iota(\Delta\omega)_j t_i} \approx \sum_{l=1}^{L} b_{il} c_{lj}, \tag{2.6}$$

where $(\Delta\omega)_j$ is the measured $B_0$ offset at spatial location $j$, $i$ indexes time, $l$ indexes time segments, $b_{il}$ are temporal interpolation functions, and $c_{lj}$ are complex exponential functions evaluated at the middle of the interpolation windows, given by:

$$c_{lj} = e^{-\iota(\Delta\omega)_j t_l}, \tag{2.7}$$

where $t_l$ is time at the center of window $l$. This enables efficient computation using $L$ separate NUFFTs that are subsequently combined. The signal model is as follows:

$$y_i = \sum_{l=1}^{L} b_{il} \left[ \sum_{j=1}^{N} f_j c_{lj} e^{-\iota 2\pi \vec{k}_i \cdot \vec{r}_j} \right], \tag{2.8}$$

where $y_i$ is the data at time $t_i$, $f_j$ is the imaged object at spatial location $j$, $\vec{k}_i$ are the k-space locations, and $\vec{r}_j$ are the spatial coordinates.

In this work, 4 time segments were used. Off-resonance was incorporated in TrACR re-

constructions of the first subject's spiral scan in two schema: 1) Off-resonance incorporated throughout TrACR and in the final reconstruction system matrix also and 2) Off-resonance incorporated solely in the final image reconstruction system matrix.

### 2.3.2.2 Error Basis Generation

**2.3.2.2.1 Golden-Angle Radial** Data sampling occurs only during the flat parts of the trapezoids in conventional and golden-angle radial acquisitions, so the majority of trajectory errors can be captured by linear translations of the radial lines in k-space [44]. This leads to a straightforward trajectory error basis matrix construction $\vec{E} = (E_x, E_y)$, as:

$$E_x = \left[ \begin{array}{cc} I_{N_{proj}} \otimes 1_{N_{samp} \times 1} & 0_{N_{proj}N_{samp} \times N_{proj}} \end{array} \right], \tag{2.9}$$

$$E_y = \left[ \begin{array}{cc} 0_{N_{proj}N_{samp} \times N_{proj}} & I_{N_{proj}} \otimes 1_{N_{samp} \times 1} \end{array} \right], \tag{2.10}$$

where $I_{N_{proj}}$ is an $N_{proj} \times N_{proj}$ identity matrix in which $N_{proj}$ is the number of radial projections, $\otimes$ represents a Kronecker product, $1_{N_{samp} \times 1}$ is a length-$N_{samp}$ vector of ones with $N_{samp}$ being the number of sample points per projection, and $0_{N_{proj}N_{samp} \times N_{proj}}$ is a matrix of zeros.

**2.3.2.2.2 Center-Out Radial and Spiral** In ramp-sampled center-out radial and spiral acquisitions, data are acquired while the gradients change amplitudes. The majority of trajectory errors are the result of eddy currents generated on conducting structures in the scanner which produce gradient field errors. These gradient field errors are typically modeled as a weighted linear combination of terms of the form ([55], p. 320):

$$g_e(t, \tau) = -\frac{dG(t)}{dt} * \left\{ H(t)e^{-t/\tau} \right\}, \tag{2.11}$$

where each term has a different value of the time constant $\tau$, $t$ is time, $G(t)$ is the nominal gradient waveform, and $H(t)$ is the Heaviside step function. For a given center-out radial

or spiral readout gradient waveform $G(t)$, this model was used to calculate error basis vectors by generating functions $g_e(t, \tau)$ for 1000 time constants spaced linearly between 1 $\mu$s and 2 ms, and sampled with the same dwell time as the measured k-space data. The functions were integrated to arrive at a k-space error basis set, then compressed down to six linearly-independent waveforms by stacking them into a matrix $G_e$, calculating its singular value decomposition (SVD) $G_e = USV'$, and taking the first six columns of the matrix $U$ (corresponding to the six largest singular values) as the error basis for that input gradient waveform. For the center-out radial case, the compressed error basis matrix was calculated for a single trapezoid and was rotated for each projection, forming the final error basis matrix $\vec{E} = (E_x, E_y)$ as:

$$
E_x = \begin{bmatrix} \tilde{G} \\ \tilde{G} \cos \frac{2\pi}{N_{proj}} \\ \vdots \\ \tilde{G} \cos \frac{2\pi(N_{proj}-1)}{N_{proj}} \end{bmatrix} ,
\tag{2.12}
$$

$$
E_y = \begin{bmatrix} 0 \\ \tilde{G} \sin \frac{2\pi}{N_{proj}} \\ \vdots \\ \tilde{G} \sin \frac{2\pi(N_{proj}-1)}{N_{proj}} \end{bmatrix} ,
\tag{2.13}
$$

where $\tilde{G}$ is the SVD-compressed error basis matrix. For the spiral case, compressed error basis matrices were calculated for the $G_x(t)$ and $G_y(t)$ waveforms for one of the 16 shots.

These were then rotated to form the final error basis matrix $\vec{E} = (E_x, E_y)$ as:

$$
E_x = \begin{bmatrix} \tilde{G}_x & 0 \\ \tilde{G}_x \cos\frac{2\pi}{N_{shot}} & \tilde{G}_y \sin\frac{2\pi}{N_{shot}} \\ \vdots & \vdots \\ \tilde{G}_x \cos\frac{2\pi(N_{shot}-1)}{N_{shot}} & \tilde{G}_y \sin\frac{2\pi(N_{shot}-1)}{N_{shot}} \end{bmatrix}, \tag{2.14}
$$

$$
E_y = \begin{bmatrix} 0 & \tilde{G}_y \\ -\tilde{G}_x \sin\frac{2\pi}{N_{shot}} & \tilde{G}_y \cos\frac{2\pi}{N_{shot}} \\ \vdots & \vdots \\ -\tilde{G}_x \sin\frac{2\pi(N_{shot}-1)}{N_{shot}} & \tilde{G}_y \cos\frac{2\pi(N_{shot}-1)}{N_{shot}} \end{bmatrix}. \tag{2.15}
$$

In total, 6 error weights were fit in the center-out radial case and 12 in the spiral case. The definitions in Eqs. 2.12-2.15 are based on the empirical observation that the trajectory errors were very similar for the $x$ and $y$ gradient channels, so a single set of error coefficients can be estimated that applies to all shots/projections. For non-axial slice planes, it may be more accurate to estimate separate error coefficients for each gradient channel. Finally, on our scanner, the vendor's gradient pre-emphasis resulted in a temporal gradient delay that was found to be constant between scans, subjects, and trajectories; this shift was measured and applied to the nominal gradient waveforms provided to our algorithm. Alternatively, one could disable waveform pre-emphasis for such acquisitions.

## 2.4   Results

Figure 2.1 illustrates the accuracy with which the generated error basis functions can be fit to the measured trajectory errors in the center-out radial and spiral cases. This was investigated by directly fitting (by least-squares, without the TrACR algorithm) the error basis functions to the measured errors for the first projection/shot, while varying the size of the SVD-compressed error bases. Figure 2.1a plots the root-mean-square (RMS) error
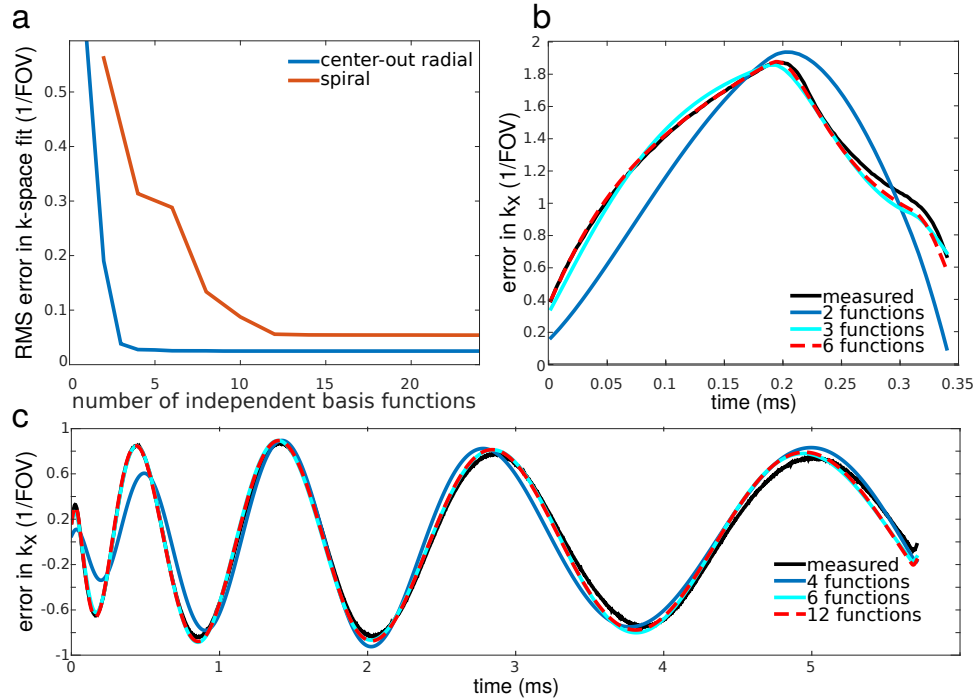
Figure 2.1: Investigation of the number of SVD-compressed error basis functions necessary to accurately model trajectory errors. (a) Residual error for direct least-squares fits of basis functions to the measured trajectory error for the center-out radial and spiral trajectories versus the number of independent basis functions used. (b) Direct least-squares fits of 2, 3, or 6 independent basis functions to the measured error for one projection of the center-out radial trajectory. (c) Direct least-squares fits of 4, 6, or 12 independent basis functions to the measured error for one shot of the spiral trajectory in the $k_x$ dimension.

in fitting the measured trajectory errors, as a function of the number of basis functions used. For both trajectories, as the size of the SVD-compressed basis set increased, the error monotonically decreased to a minimum value and then flattened out. Figure 2.1b plots least-squares fits of 2, 3 and 6 basis functions to the error measured for the first projection of the center-out radial trajectory. For 6 basis functions, the measured and fit curves nearly coincide.

Figure 2.1c plots least-squares fits of 4, 6 and 12 basis functions to the error measured for the first shot of the spiral trajectory (only the $k_x$ error waveform is shown). For 12 basis functions, the measured and fit curves nearly coincide. These results support the use of 6 basis functions for the center-out radial TrACR reconstructions that follow, and 12 basis functions for the spiral TrACR reconstructions.

Figures 2.2-2.4 show golden-angle radial, center-out radial, and spiral images reconstructed using the nominal trajectories, the measured trajectories (center-out radial and spiral only) and the trajectories estimated using TrACR with SENSE and SPIRiT, in the same subject and slice. For each case, the displayed image was formed as a sum-of-squares combination of the individual coil images reconstructed by CG using the final trajectory. Before running TrACR, the 32-channel coil data was compressed to 15 channels. In all cases, the uncorrected image contains considerable intensity modulations and blurring across the brain, which are removed in both SENSE and SPIRiT TrACR reconstructions. The difference images are similar in all cases, indicating that both TrACR-SENSE and -SPIRiT were effective in estimating the corrected trajectories, and (in the center-out radial and spiral cases) yielded similar image reconstructions as the measured trajectories. Across subjects, the mean number of TrACR iterations was: 27 (golden-angle radial), 779 (center-out radial), and 205 (spiral). The mean compute time was: 8.7 minutes (golden-angle radial), 2.7 hours (center-out radial), and 1.1 hours (spiral). The TrACR-SPIRiT reconstructions required between 10-40% more iterations/longer compute times.

Figure 2.5 illustrates the estimated trajectories and errors for the subject shown in Figs.

34
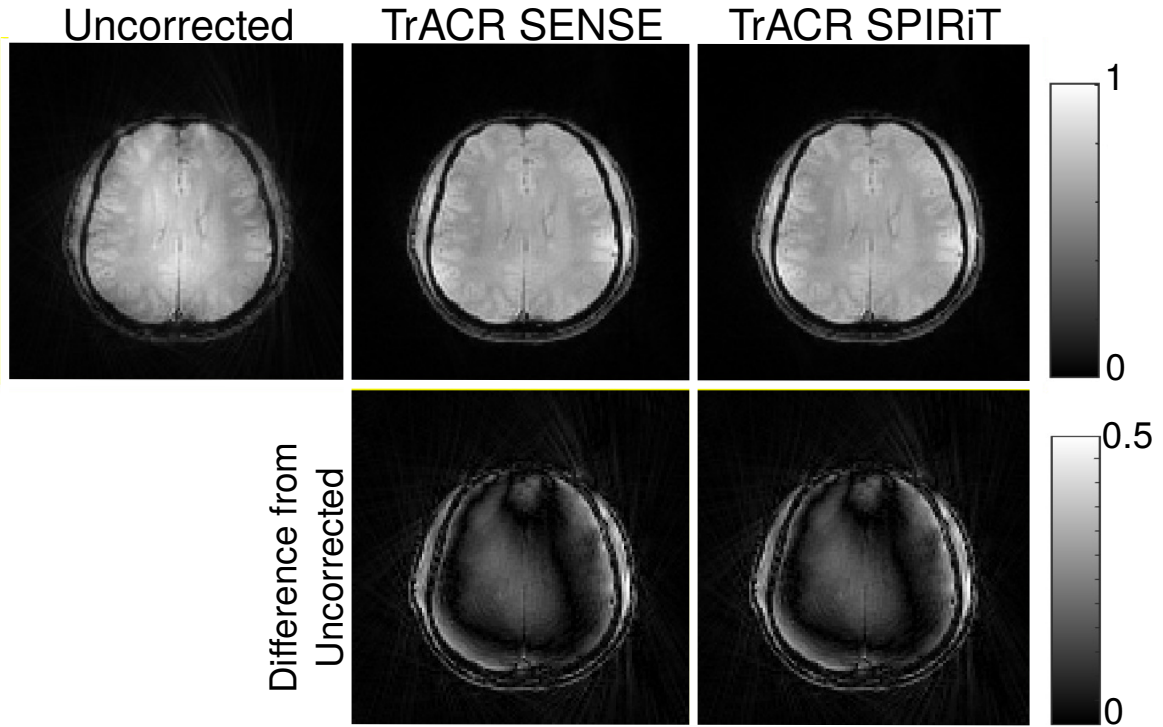
Figure 2.2: Final CG image reconstructions on nominal (uncorrected), TrACR-SENSE, and TrACR-SPIRiT trajectories for the golden-angle radial dataset in one subject. The second row shows intensity differences between the TrACR reconstructions and the uncorrected image.
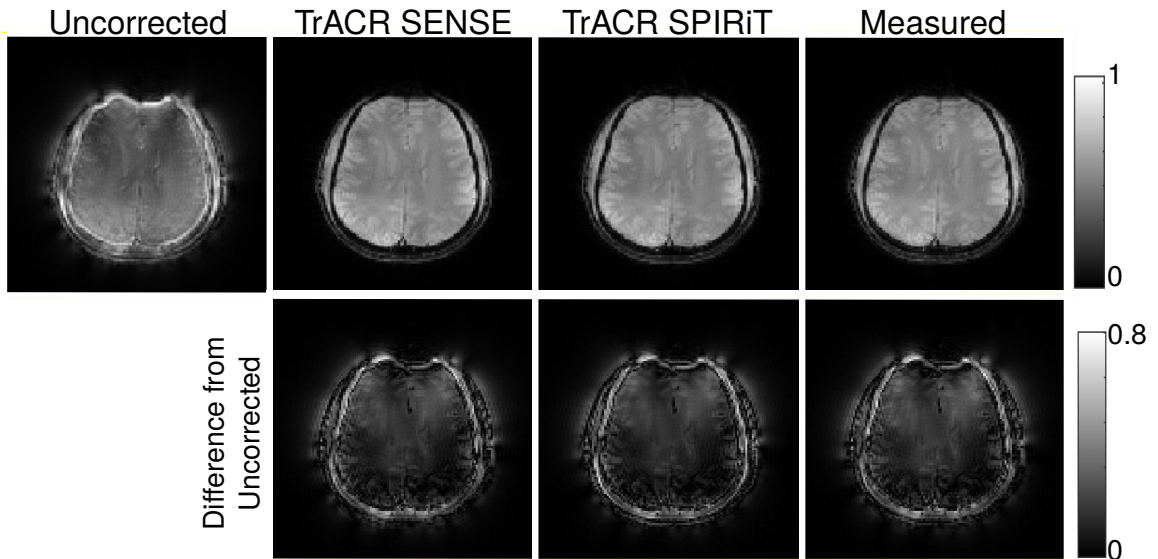


Figure 2.3: Final CG image reconstructions on nominal, TrACR-SENSE, TrACR-SPIRiT, and measured k-space trajectories for the center-out radial dataset in one subject. The second row shows intensity differences between the corrected reconstructions and the uncorrected image.
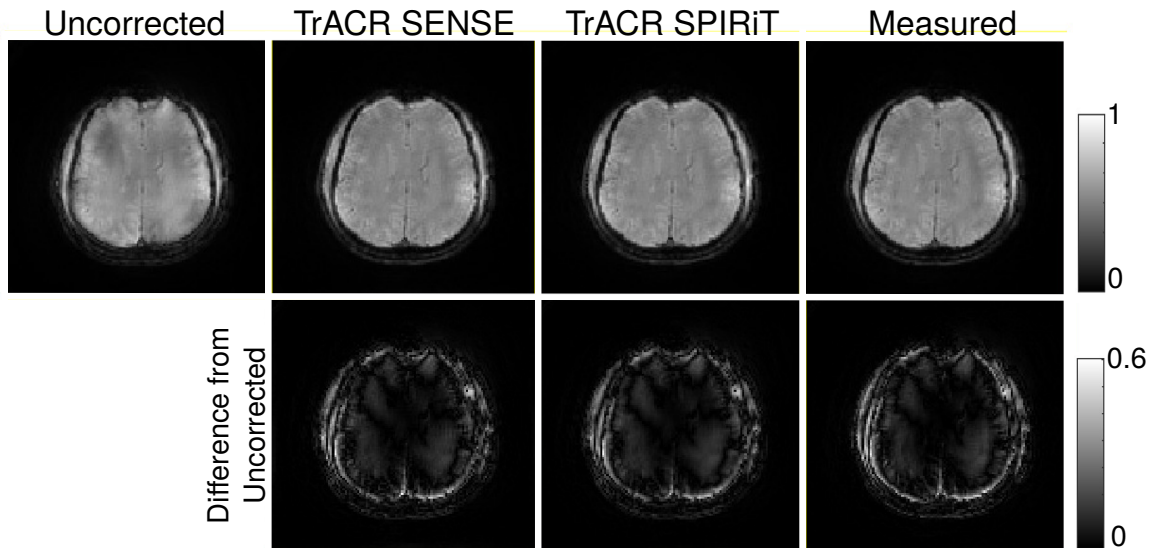
Figure 2.4: Final CG image reconstructions on nominal, TrACR-SENSE, TrACR-SPIRiT, and measured k-space trajectories for the spiral dataset in one subject. The second row shows intensity differences between the corrected reconstructions and the uncorrected image.

2.2-2.4.

Figure 2.5a shows the center of the nominal golden-angle radial trajectory and the trajectories estimated by TrACR-SENSE and -SPIRiT. The TrACR-SENSE and -SPIRiT trajectories are indistinguishable on this plot, and differ considerably from the nominal trajectory. Figure 2.5b plots the measured and estimated k-space errors for a single projection of the center-out radial trajectory, as a function of time. While both trajectory estimates fit the measured error well at the beginning of the projection in the center of k-space, at higher spatial frequencies near the end of the projection, the TrACR-SENSE estimate is closer to the measured error. The higher accuracy achieved at low spatial frequencies reflects the fact that the MR signal amplitude is much higher in the center of k-space, so the algorithm favors minimizing trajectory errors there. Figure 2.5c shows an analogous result for the spiral case: the error estimates are very close to the measured trajectory near the center of k-space, and diverge somewhat with increasing time/spatial frequency, and the TrACR-SENSE estimate comes closer to the measured trajectory than does TrACR-
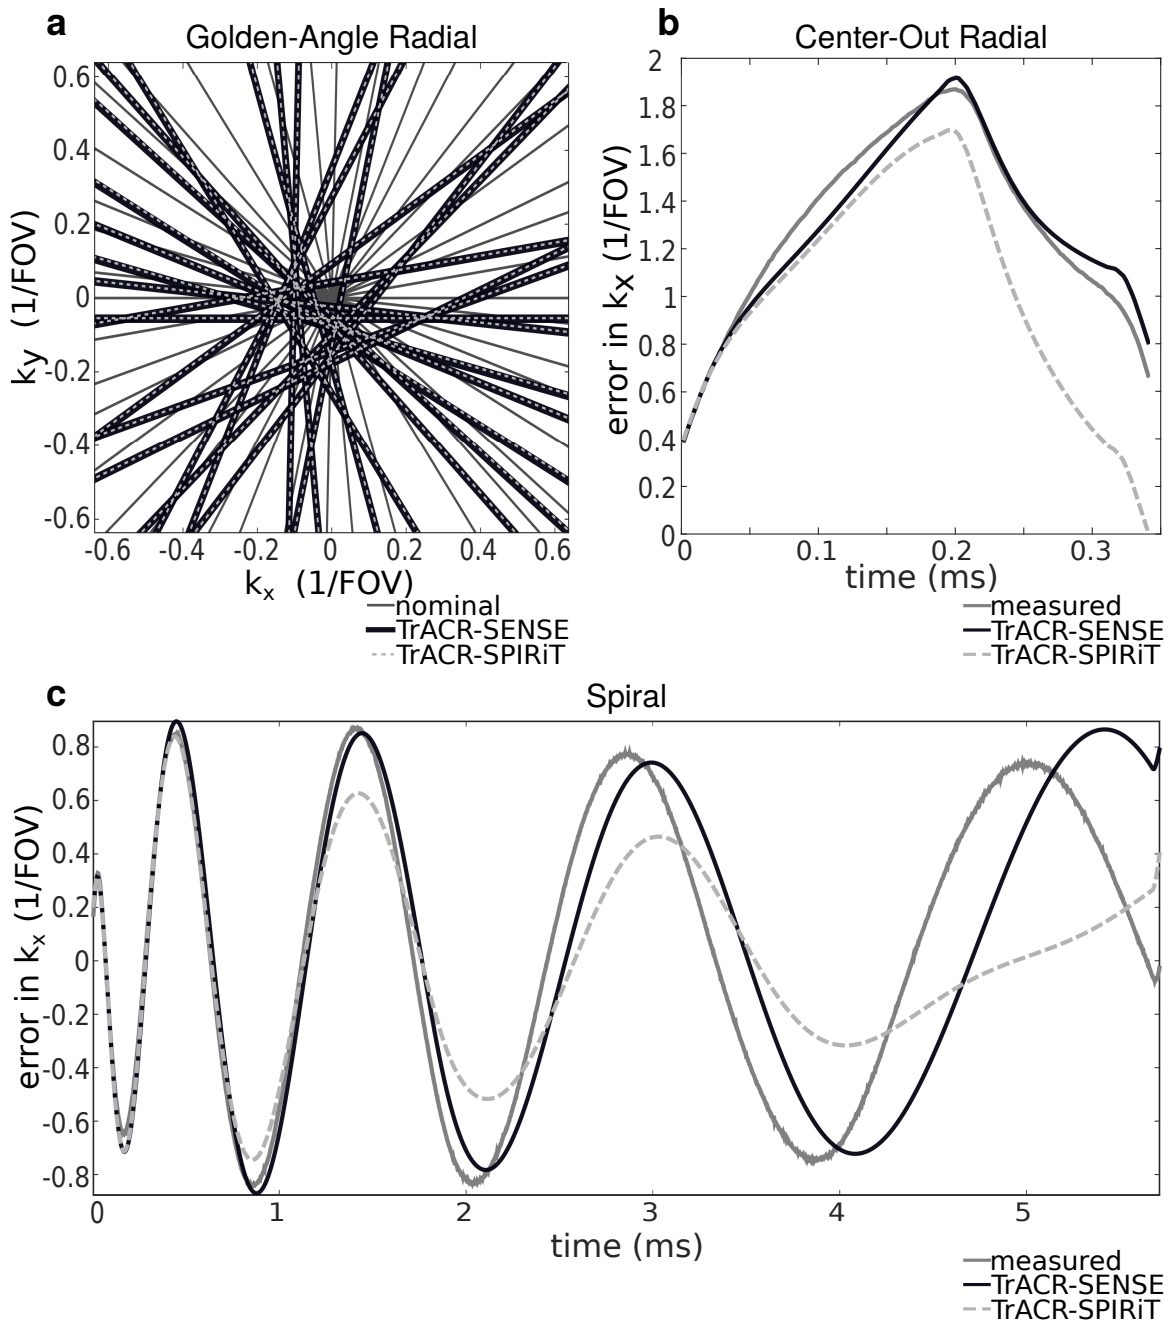
Figure 2.5: Trajectory errors for the image reconstructions in Figs 1-3. (a) A subset of nominal golden-angle radial projections and their corresponding TrACR-SENSE and TrACR-SPIRiT projections in the center of k-space. The TrACR-SENSE and TrACR-SPIRiT projections coincide. (b) Measured, TrACR-SENSE and TrACR-SPIRiT center-out radial k-space trajectory error curves as a function of time, for one projection. (c) The same curves in (b) for the $k_x(t)$ waveform of one shot of the spiral trajectory. Trajectories and errors are plotted in units of multiples of 1/FOV.

SPIRiT. Although the trajectories estimated by TrACR-SENSE and TrACR-SPIRiT differ in the high spatial frequencies in the center-out radial and spiral cases, those differences did not result in significant differences in the final reconstructed images in Figures 2.3 and 2.4.

To investigate the dependence of the trajectory error on k-space acceleration and the number of receive channels, the golden-angle radial TrACR-SENSE reconstructions in this subject were repeated for acceleration factors between 1 and 8, and for numbers of receive channels between 1 and 32. Acceleration was realized by uniformly dropping projections, and the number of channels was varied using SVD coil compression. Figure 2.6a shows fully-sampled and 4x-accelerated (50 projection) image reconstructions using CG-SENSE and the final TrACR trajectories (in this case images were reconstructed using `lsqr` with a stopping tolerance of $10^{-1}$). With $4\times$ acceleration, there is an apparent loss of SNR but no noticeable aliasing artifacts. The same figure also plots the root-mean-square (RMS) trajectory error for each acceleration factor, referenced to the fully-sampled 32-channel TrACR-SENSE result. The errors were calculated after subtracting off the mean k-space trajectory shift, and were low for all acceleration factors, increasing only slightly with acceleration. Figure 2.6b plots the trajectory errors across numbers of receive channels, referenced to the fully-sampled 32-channel TrACR-SENSE result. As the number of coils used for reconstruction increased, the trajectory error decreased for both acceleration factors. For less than 10 coils, the error was higher with both full sampling and $4\times$ acceleration than it was at any acceleration factor with 15 coils, indicating that in the golden-angle radial case the trajectory error depends more on the number of coils than on the acceleration factor.

Figure 2.7 shows the evolution of the k-space error estimates and images over TrACR outer loop iterations, for center-out radial TrACR-SENSE. Images are shown on top, with center-out radial k-space estimates shown at the same TrACR iteration numbers on the bottom and the measured trajectory error provided for reference. The image improves rapidly with early TrACR iterations as the lower k-space locations are corrected, whereas
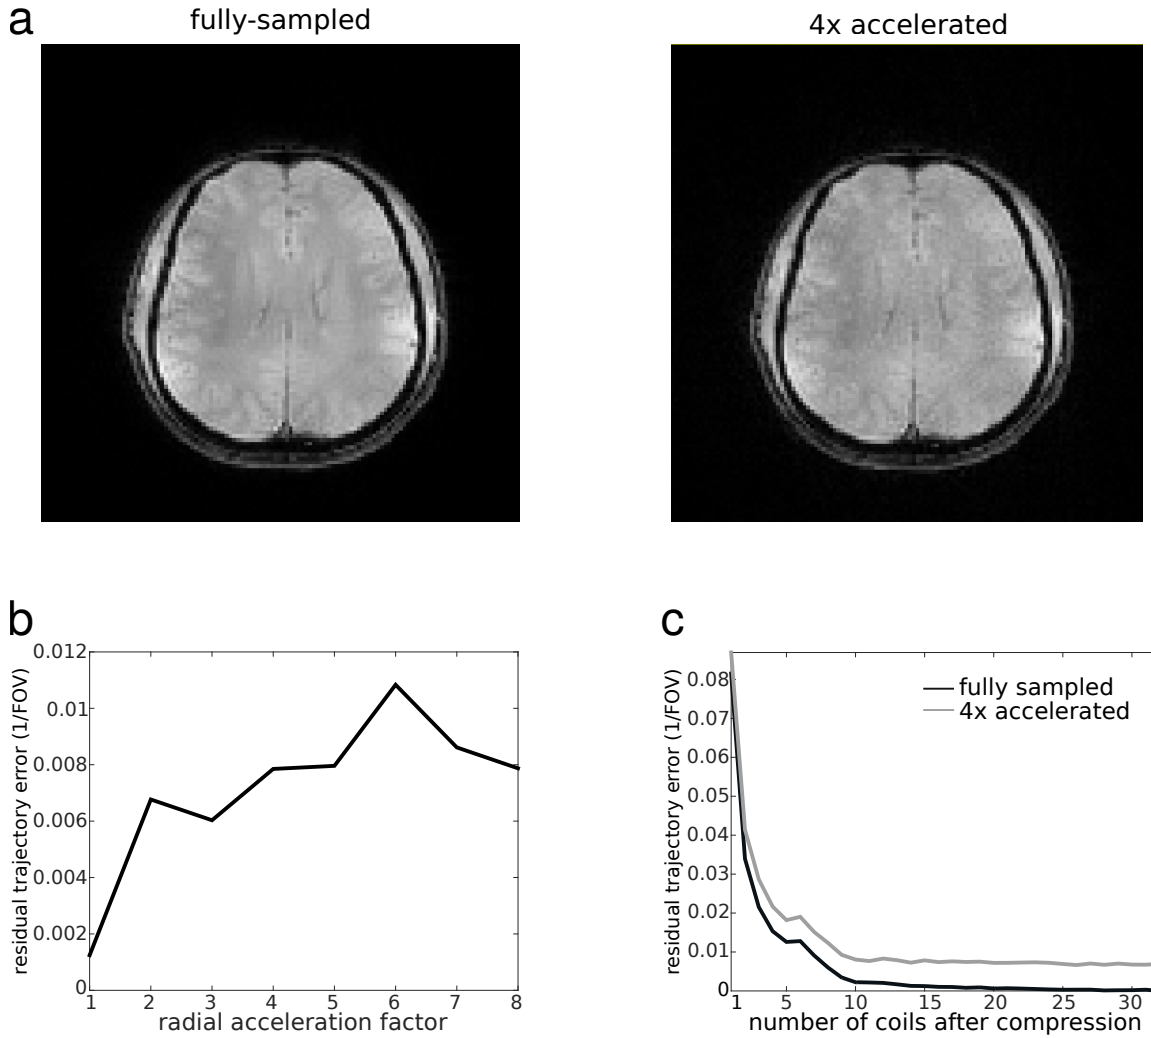
Figure 2.6: Error vs. radial acceleration. (a) TrACR-SENSE corrected CG-SENSE reconstructions for full sampling and $4\times$ acceleration. (b) RMS k-Space trajectory error versus radial acceleration factor for GA radial TrACR-SENSE reconstructions with 15 coils. (c) Error versus number of coils used for TrACR-SENSE, for full sampling and $4\times$ acceleration. All errors are expressed as multiples of 1/FOV and are referenced to the fully-sampled 32-channel TrACR-SENSE trajectory estimate.
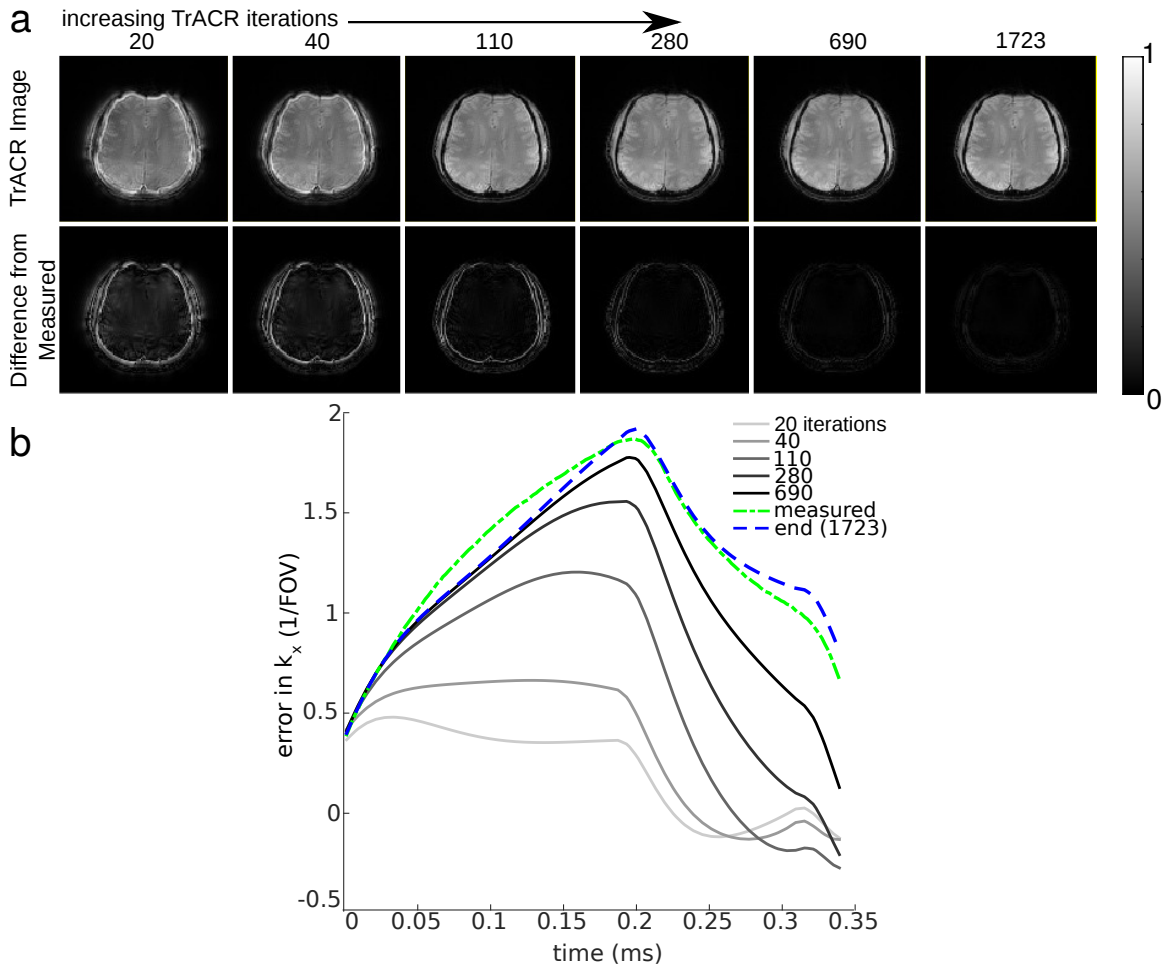
Figure 2.7: Evolution of TrACR-SENSE images and trajectory error estimates versus TrACR outer loop iteration, for a center-out radial reconstruction. (a) Images and magnitude differences between the TrACR image and an image reconstructed using a measured k-space trajectory, versus number of TrACR iterations. (b) Corresponding k-space error estimates, plotted with the final TrACR trajectory error estimate and the measured trajectory error.

high frequency corrections build up more slowly.

Figure 2.8 summarizes the performance of TrACR across the 5 subjects, 3 trajectories, and SENSE and SPIRiT formulations. Figure 2.8a shows how much lower (in percent; higher numbers are better) the final TrACR trajectories' cost functions (Eqs. 2.1 and 2.2) were, compared to no correction. All instances of TrACR significantly reduced the uncorrected image cost, with a median cost reduction across subjects and trajectories of 76%. The same figure also shows the SENSE and SPIRiT cost reductions for the measured trajectories, which in all cases were not markedly higher than the TrACR cost reductions. Figure 2.8b shows the increase in normalized image gradient squared for each case, compared to the uncorrected images. The normalized image gradient squared is an image quality metric that has been reported as having a high correlation with observer image quality rating [56]. It was calculated from the final SENSE and SPIRiT image reconstructions. All reconstructions resulted in increased normalized image gradient squared, which were comparable to values for images reconstructed using the measured k-space trajectories.

Lastly, Fig. 2.9 shows that the TrACR-SPIRiT reconstructions are relatively insensitive to the SPIRiT regularization parameter $\lambda$, at least over two orders of magnitude for each trajectory. Figure 2.10 shows that a measured off-resonance map can be incorporated into the signal model for TrACR reconstruction. This may be desirable in body imaging where the range of off-resonances can be larger than in the brain, which may result in blurring that confounds trajectory error estimation. The spiral trajectory had the longest readout duration, and incorporating the measured map resulted in marked signal recovery in the front of the brain, but did not significantly affect the trajectory error estimate.

## 2.5   Discussion

In vivo experiments demonstrated TrACR's ability to correct image artifacts caused by k-space trajectory errors in non-Cartesian acquisitions. TrACR corrections made significant visible improvements (reduced streaking and blurring, and enhancement of fine
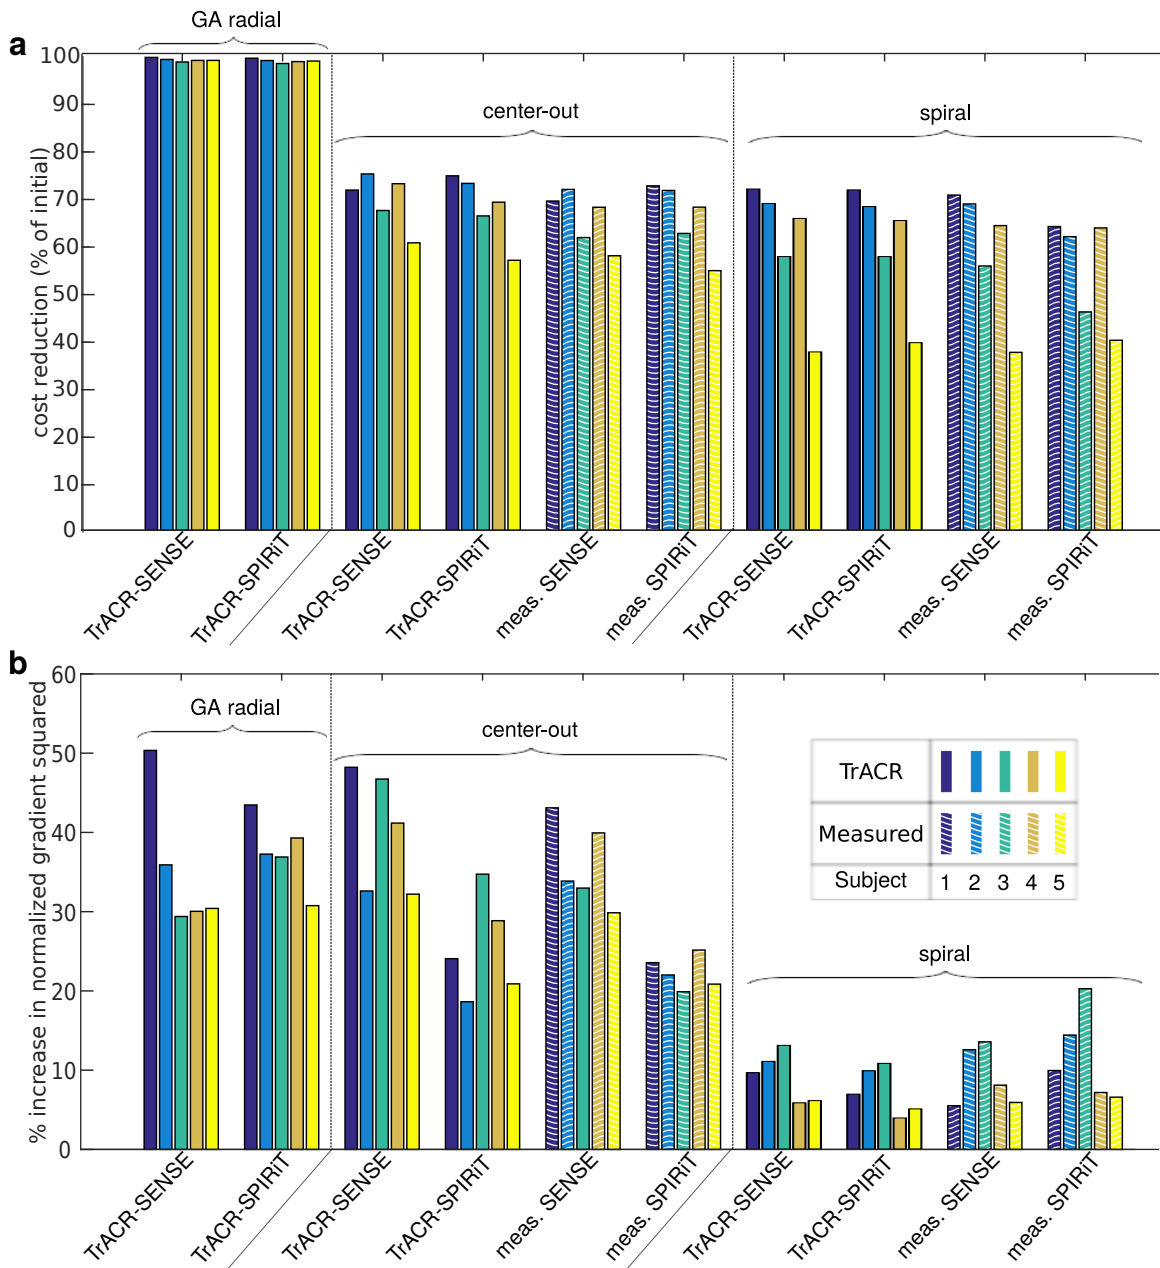
Figure 2.8: Numerical TrACR-SENSE and -SPIRiT results across 5 subjects and the three trajectories. (a) Cost function (Eqs. 1 and 2) reduction as a percentage of the uncorrected (initial) cost. (b) Percentage increase in the normalized image gradient squared, versus no correction. Metrics for reconstructions using measured trajectories are also shown for the center-out radial and spiral cases.
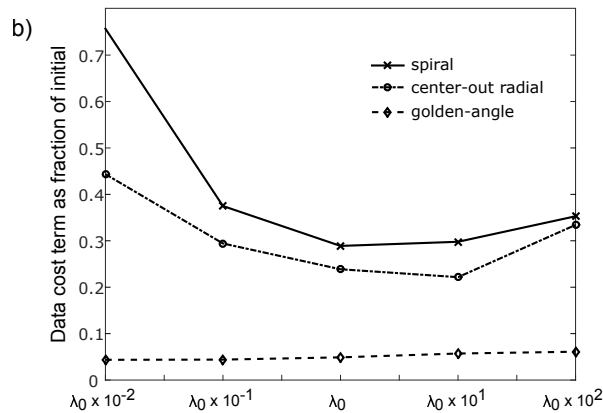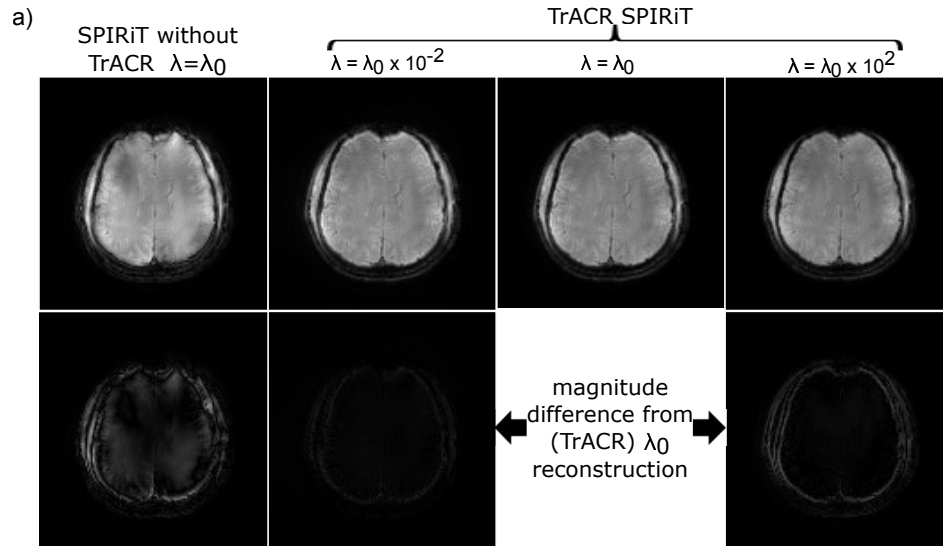
Figure 2.9: TrACR performance with varying SPIRiT regularization. (a) (Left) Spiral image reconstruction before TrACR with the regularization set to $\lambda_0$ (10% of the median of the absolute value of the k-space data - the same value used in the rest of this chapter), and (Right) spiral image reconstructed with TrACR with $\lambda = \lambda_0 \times 10^{-2}$, $\lambda = \lambda_0$, and $\lambda = \lambda_0 \times 10^2$. Images and errors are on the same color scale. (b) The data fidelity term of the cost function in Eq. 2 as a fraction of initial, versus SPIRiT regularization parameter for all 3 trajectories.

43
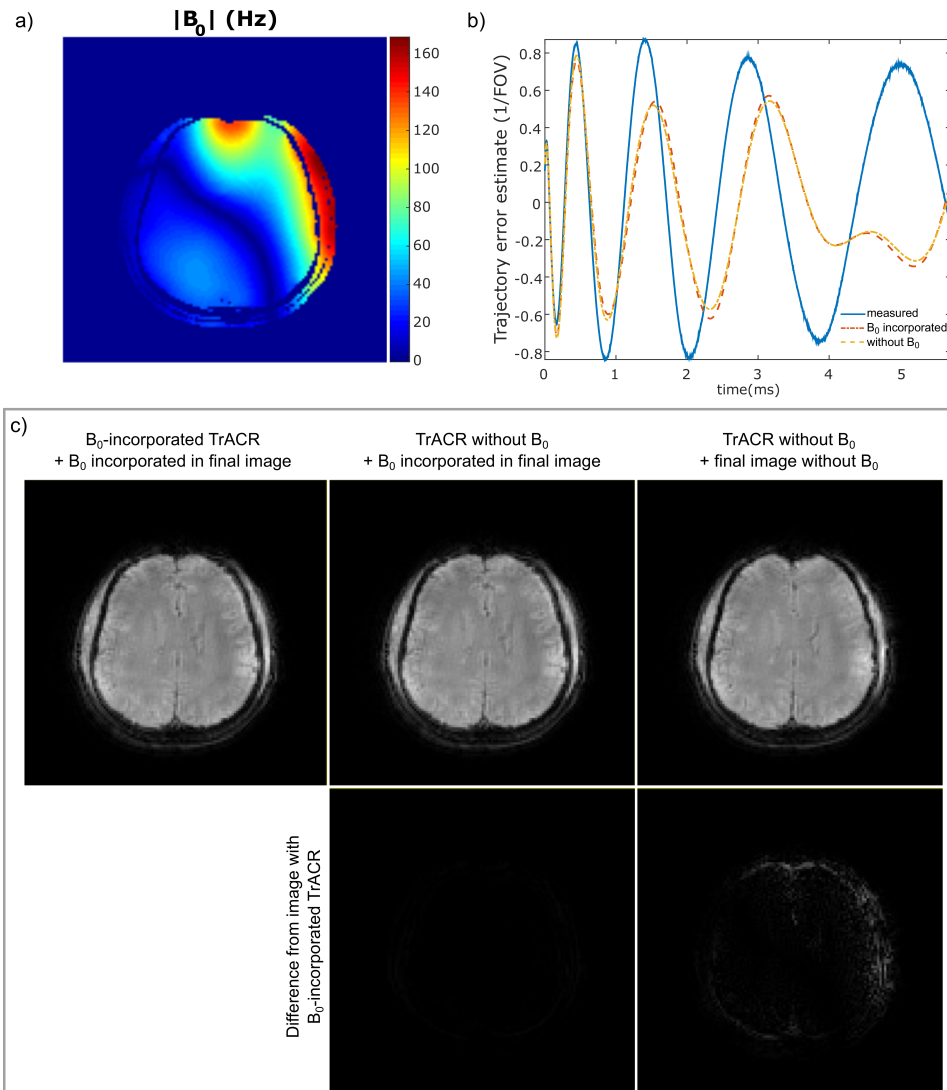
Figure 2.10: TrACR with off-resonance compensation. (a) The magnitude of the measured off-resonance map in Hz for one subject. The map is masked for display. The approximate 150 Hz offset in the front and right side of the brain would result in approximately 1 cycle of phase across the spiral readouts. (b) Measured spiral trajectory error (solid blue), TrACR-estimated error with off-resonance incorporated (red dot-dashed), and TrACR-estimated error without off-resonance incorporated (yellow dashed). (c) Final reconstructed spiral images: (left) image reconstructed with the measured off-resonance map incorporated both in the TrACR and the final reconstruction, (center) image reconstructed with the off-resonance map incorporated in the final image reconstruction system matrix but not in the TrACR estimation, and (right) image reconstructed without off-resonance compensation in TrACR or the final reconstruction. Magnitude differences are also shown from the left reconstruction, on the same color scale. Off-resonance was incorporated in the reconstruction system matrices using time-segmentation [1].

details) to the reconstructed images in the in vivo experiments, with comparable image quality to images reconstructed using measured k-space trajectories. Golden-angle radial reconstructions across acceleration factors demonstrated that TrACR-estimated trajectories were less accurate at higher radial acceleration factors; however, the errors remained relatively low across acceleration factors due to the large signal magnitude and oversampling at the center of k-space even with sub-Nyquist radial sampling. The golden-angle radial reconstructions with varying numbers of coils demonstrated that the method benefits from parallel imaging due to the data redundancy it provides, since error increased as the number of coils decreased. The algorithm performed consistently across five subjects, in terms of the amount by which the SENSE and SPIRiT cost functions were reduced, and in terms of the increase in normalized image gradient squared.

TrACR reconstructions were able to correct most of the measured errors in the center-out radial and spiral trajectories, as shown in Fig. 2.5. Due to the higher signal and higher sampling density in the center of k-space, the algorithm preferentially corrected trajectory errors there, and converged with somewhat higher residual errors at the higher spatial frequencies. Though it was not observed in our reconstructions, it is possible that the TrACR-estimated trajectory will have higher error than the uncorrected trajectory at the high spatial frequencies. This potential problem could be mitigated by multiplying the error basis functions with a window that decreases to zero at the high frequencies. We have tested this approach with the spiral reconstructions (results not shown) and found that it performed similarly to the unwindowed reconstructions, with negligible image differences when the windows truncated the error functions at approximately 75% of the maximum k-space radius. Windowing the error basis may also accelerate algorithm convergence, since (as demonstrated in Fig. 2.7) the low frequencies are fitted early in the TrACR iterations.

The center-out radial and spiral results suggest that accurate corrections at high spatial frequencies may be precluded for cases in which there is a long separation in time between sampling the center of k-space and the end of the trajectory. The algorithm's success also

depends on the provision of a suitable error basis. In this work, error bases for spiral and center-out radial trajectories were derived from eddy current models, and this construction approach can be applied to any existing readout trajectory. However, error basis functions derived from eddy current models may not be effective in capturing other sources of error, such as errors due to gradient amplifier non-linearity and long time constant eddy currents that persist between TRs. Developing suitable error bases in those cases may require the incorporation of hysteresis models (for amplifier non-linearity) and whole-sequence eddy current modeling (for long-time constant eddy currents). Furthermore, while an eddy current error basis can be constructed for any trajectory, properties of the trajectory itself may still preclude effective corrections.

An important consideration in the TrACR-SPIRiT reconstruction is the choice of images used to calibrate the SPIRiT operator. We have found that, in many cases, the algorithm will converge to an acceptable solution if the operator is initially calibrated using low-resolution images reconstructed with the nominal trajectory, and is periodically re-calibrated during the TrACR iterations using the latest trajectory error estimate (results not shown). However, due to model inconsistencies inherent in that approach, it is possible for the iterations to diverge or converge to an unacceptable solution. Therefore a more cautious alternative is to calibrate the operator using Cartesian images of the same geometry, as described in [33]. This is the approach that was used here. Another consideration that may affect both TrACR-SENSE and -SPIRiT performance is the density compensation. In this work, density compensation weights for all TrACR reconstructions were calculated using the nominal trajectory, and were held fixed over the iterations. The weights were then updated using the TrACR error estimate for the final image reconstruction. Slightly better trajectory estimates may be possible by either periodically updating the density compensation, or by not using density compensation at all, which would require increasing the number of iterations used in each image update.

The computation times for the algorithm were shortest for golden-angle radial, and

longest for center-out radial. This result was expected since the golden-angle radial trajectory error basis functions were uniform across each projection, so determining their weights could likely be performed using only the center of k-space, and it was found that the algorithm preferentially corrected trajectory errors there first. The large difference between the computation times for the spiral and center-out radial trajectories is likely due to the fact that the initial center-out radial RMS trajectory errors neared 1/FOV, or the Nyquist sample spacing, whereas the spiral and golden-angle radial trajectory errors were about half as large. Overall, the reported computation times of several minutes (golden-angle radial) to a few hours (center-out radial) were not compatible with online use. However, in the current implementation, the TrACR algorithm was stopped when the backtracking line search returned a zero step size in the first iteration. The motivation for this stopping criterion was to demonstrate the very best possible trajectory correction with the algorithm. In practice we have found that the reconstructed images stop changing significantly well before this stopping criterion is satisfied, and that a more practical criterion that is predictive of this may be to stop the algorithm when the difference between consecutive cost function values falls below 0.1% of the current cost. Using this criterion resulted in approximately 75% shorter computation times/fewer iterations, with worse trajectory errors at high spatial frequencies compared to measurements in the center-out radial and spiral cases, but with negligible final image differences (results not shown). The algorithm's computations could be accelerated using parallel computing [57, 58], and its convergence may be accelerated by jointly (rather than alternately) updating the images and trajectory errors each iteration. The alternating update approach used here was chosen primarily for its flexibility in decoupling the k-space error and image update codes. While CG is widely accepted as an efficient method for MR image reconstruction, algorithms other than CG may work better for the k-space error updates, such as Newton or Gauss-Newton methods. These could accelerate convergence at the cost of increased computational cost per iteration compared to CG.

## 2.6  Conclusions

In this chapter, the TrACR approach to auto-correct non-Cartesian images for k-space trajectory errors was described and validated in vivo for three non-Cartesian trajectories. It is a more general formulation than existing methods, and can be extended to any non-Cartesian trajectory for which a suitable trajectory error basis can be derived. It does not require trajectory measurements or prior calibration data and exploits data redundancy provided by oversampling in non-Cartesian acquisitions and parallel imaging. The method can be used in conjunction with multiple parallel imaging reconstruction techniques. In the following chapter, TrACR is extended for EPI trajectories.

Chapter 3

Echo-Planar Imaging

## 3.1 Introduction

In this chapter, a method is presented to automatically correct trajectory delays and line-to-line phase errors in echo-planar imaging (EPI) via a new extension of the TrACR algorithm introduced in Chapter 2. EPI is a fast MRI technique in which multiple lines of k-space are measured per excitation. It is widely used in functional and diffusion MRI. However, EPI images contain ghosting artifacts due to trajectory delays and phase errors between adjacent k-space lines that result from eddy currents created by rapidly switched readout gradients.

The most common methods to correct EPI ghosting artifacts are based on the collection of calibration data from which delays and phase errors can be estimated and applied in image reconstruction [30, 59–63]. Usually this data comes from a separate acquisition without phase encoding gradient blips, acquired before the imaging scan. Corrections can also be made by re-acquiring EPI k-space data that is offset by one k-space line so that odd k-space lines become even and vice versa [30, 64]. The gradient impulse response function can also be measured and applied to predict errors [65]. However, these methods do not address dynamic errors caused by effects such as gradient coil heating. Dynamic errors can be compensated by measuring calibration data within the imaging sequence itself, for example by reacquiring the center line of k-space within a single acquisition [29]. However, these approaches result in a loss of temporal resolution. Alternatively, dynamic errors can be measured during a scan without modifying the sequence using field-probe measurements [28, 39, 66]. However, the hardware required to make those measurements can take up valuable space in the scanner bore and is not widely available at the time of writing.

As an alternative to separate calibration measurements, many retrospective methods attempt to correct ghosting based on the EPI data or images themselves. The image-based methods [67–70] rely on the assumption that some part of the initial image contains no ghosted signal. Another group of methods makes corrections based on finding phased array combinations that cancel ghosts [71–75]. Several methods use parallel imaging to separately reconstruct images from odd and even lines and then combine them, and these have further been combined with a dynamically alternating phase encode shift or direction [72–74, 76–78]. However, relying on undersampled data for calibration weights may make these approaches unstable, and some methods reduce temporal resolution. Importantly, almost all these calibration-free retrospective methods are either incompatible or have not been validated with multi-shot EPI, and most are also either incompatible with parallel imaging acceleration or have only been validated with small acceleration factors of $2\times$ or less.

In the following section, a flexible EPI-trajectory auto-corrected image reconstruction (EPI-TrACR) is proposed that alleviates ghosting artifacts by exploiting data redundancy between adjacent k-space lines in multicoil EPI data. It is an extension of the method for automatic non-Cartesian trajectory error correction (TrACR-SENSE) [79], described in Chapter 2, to the joint estimation of images and line-to-line delays and phase errors in EPI. In the following, we describe the method, including an efficient segmented FFT algorithm for delayed EPI k-space trajectories. The method is then validated in vivo at 7 Tesla, at multiple acceleration and multishot factors and in a time series. It is demonstrated that EPI-TrACR reduces dynamic ghosting and is compatible with multishot EPI and acceleration. Furthermore, the method benefited from initialization with calibration data but did not require it at moderate acceleration and multishot factors.

## 3.2  Theory

### 3.2.1  Problem Formulation

EPI-TrACR jointly estimates images, delays and phase errors by fitting an extension of the SENSE MR signal model [31] to EPI k-space data:

$$y_c[m,n] = \sum_{i=1}^{N_s} e^{-\iota 2\pi((k_m^x + \Delta k_n^x)x_i + k_n^y y_i)} e^{\iota \Delta \phi_n} s_{ci} f_i, \qquad (3.1)$$

where $y_c[m,n]$ is the signal measured in coil $c$ at the $m$th time point of the $n$th phase-encoded echo, $k_m^x$ is the k-space coordinate in the readout/frequency encoded dimension and $\Delta k_n^x$ is the trajectory delay in that dimension for the $n$th echo (out of $N$ echoes), $k_n^y$ is the $n$th echo's k-space coordinate in the phase-encoded dimension, $\Delta \phi_n$ is the phase error of the $n$th echo resulting from zeroth-order eddy currents, $s_{ci}$ is coil $c$'s measured sensitivity at $(x_i, y_i)$, $f_i$ is the image at $(x_i, y_i)$, and $N_s$ is the number of pixels in the image. The unknown parameters in this model are the image $f$ and the delays and phase errors $\{(\Delta k_n^x, \Delta \phi_n)\}_{n=1}^N$, which are determined by fitting the model to measured data $\tilde{y}_c[m,n]$ by least squares. Assuming additive Gaussian noise in the k-space data, this corresponds to a maximum likelihood (ML) estimation of the parameters. The delays and phase errors are constrained so that a single delay and phase error pair applies to all of a shot's odd echoes and another pair applies to all of its even echoes, with separate parameters for each shot. Without loss of generality, the first shot's odd echoes serve as a reference and are constrained to have zero delay and phase error. Overall, a total of $2(2N_{shot} - 1)$ delay and phase error parameters are fit to the data along with the image. We note that EPI-TrACR implicitly estimates phase error maps for each shot's even and odd echoes that are constrained to contain only zeroth and first order spatial variations; this represents an alternative to estimating spatially-resolved phase maps for each set of echoes, which would typically require spatial regularization.

### 3.2.2    Algorithm

The EPI-TrACR algorithm minimizes the data-model error by alternately updating the estimated image $\boldsymbol{f}$, the k-space delays $\{\Delta k_n^x\}_{n=1}^N$, and the phase errors $\{\Delta \phi_n\}_{n=1}^N$. The image is updated with a conjugate-gradient (CG) SENSE reconstruction [32]. The delay and phase error updates are both performed using a nonlinear Polak-Ribière (CG) algorithm [49], which requires computation of the derivatives of the sum of squared data-model errors with respect to those parameters. Denoting the sum-of-squared data-model errors as the function $\Psi$, the derivative with respect to each delay $\Delta k_n^x$ is:

$$\frac{\partial \Psi}{\partial \Delta k_n^x} = \sum_{c=1}^{N_c} \sum_{m=1}^{M} \sum_{i=1}^{N_s} \Re \left\{ -\iota 2\pi x_i e^{-\iota \Delta \phi_n} e^{\iota 2\pi ((k_m^x + \Delta k_n^x) x_i + k_n^y y_i)} s_{ci}^* f_i^* r_{cmn} \right\}, \tag{3.2}$$

and the derivative with respect to each phase error $\Delta \phi_n$ is:

$$\frac{\partial \Psi}{\partial \Delta \phi_n} = \sum_{c=1}^{N_c} \sum_{m=1}^{M} \sum_{i=1}^{N_s} \Re \left\{ \iota e^{-\iota \Delta \phi_n} e^{\iota 2\pi ((k_m^x + \Delta k_n^x) x_i + k_n^y y_i)} s_{ci}^* f_i^* r_{cmn} \right\}, \tag{3.3}$$

where $\Re$ denotes the real part, $^*$ is complex conjugation, and $r_{cmn}$ is the residual error between the measured data and the model given the current parameter estimates, $\hat{f}$, $\Delta \hat{k}_n^x$, and $\Delta \hat{\phi}_n$:

$$r_{cmn} = \tilde{y}_c[m,n] - \sum_{i=1}^{N_s} e^{-\iota 2\pi ((k_m^x + \Delta \hat{k}_n^x) x_i + k_n^y y_i)} e^{\iota \Delta \hat{\phi}_n} s_{ci} \hat{f}_i. \tag{3.4}$$

To constrain the delays and phase errors to be the same for the set of odd or even echoes of each shot, the derivatives above are summed across the echoes in that set, and a single delay and shift pair is determined for the set each CG iteration. The updates are alternated until the sum-of-squared data-model error stops changing significantly.

### 3.2.3    Segmented FFTs

Since a delayed EPI trajectory is non-Cartesian, the model in Equation 3.1 corresponds to a non-uniform discrete Fourier transform (DFT) of the image. Non-uniform fast Fourier

transform (FFT) algorithms (e.g., Ref. [80]) are typically used to efficiently evaluate non-uniform DFTs, but they use gridding, which would result in long compute times in EPI-TrACR, since Equation 3.1 is repeatedly evaluated by the algorithm. Figure 3.1 illustrates a segmented FFT algorithm that applies the delays as phase ramps in the image domain, instead of gridding the delayed data in the frequency domain. In addition to eliminating
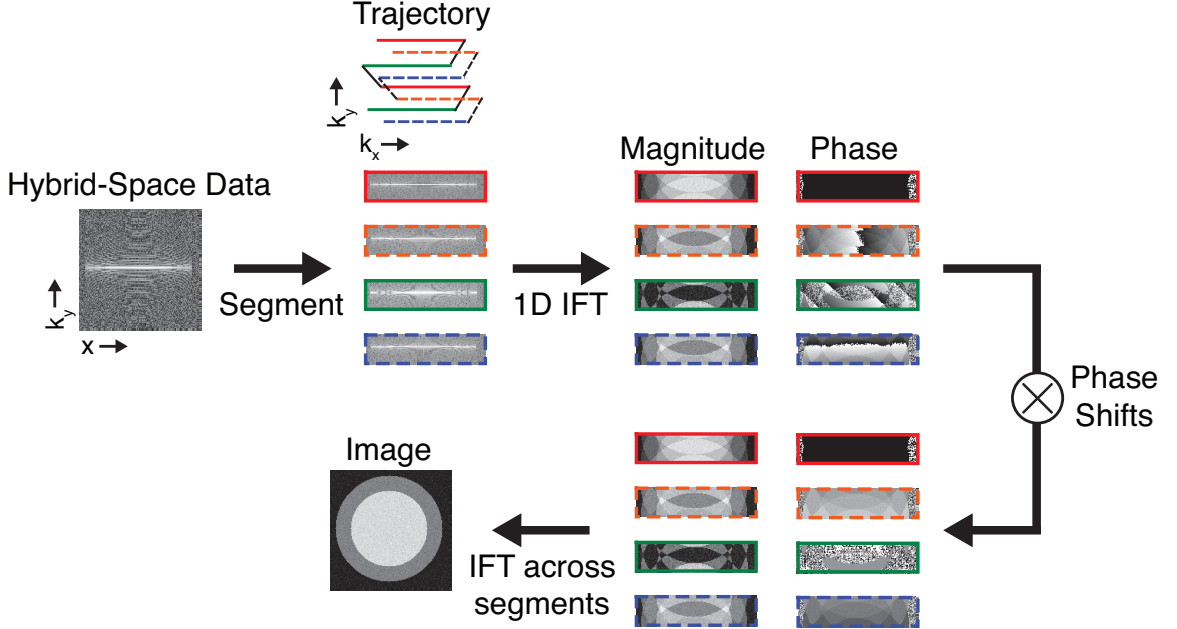


Figure 3.1: Illustration of the inverse segmented FFT, starting with 2-shot $x$-$k_y$ EPI data corrupted by line-to-line delays and phase errors. First the data are segmented into $2N_{shot}$ submatrices and individually inverse Fourier transformed. Then each image-domain sub-matrix is phase shifted to account for its offset in $k_y$, its phase error, and its delay. Finally, an inverse Fourier transform is calculated across the segments, and the result is reshaped into the image.

gridding, this also enables the data to be FFT'd in the frequency-encoded dimension before starting EPI-TrACR, so that the algorithm only needs to compute 1D FFTs in the phase-encoded dimension. The figure shows an inverse segmented FFT (k-space to image space) for a 2-shot dataset with delays and phase errors, which comprises the following steps:

1. The data in each set of odd or even echoes of each shot are collected into $2N_{shot}$ submatrices of size $M \times (N/(2 \times N_{shot}))$, and the 1D inverse FFT of each submatrix is computed in the phase-encoded dimension.

2. The estimated phase errors are applied to each submatrix.

3. A phase ramp is applied in the phase-encoded spatial dimension of each submatrix to account for that set's relative position in the phase-encoded k-space dimension. This is necessary since the inverse FFTs assume all the submatrices are centered in k-space.

4. The phase ramp corresponding to each set's estimated delay is applied to its submatrix in the frequency-encoded spatial dimension.

5. For each submatrix entry, the inverse DFT across submatrices is computed to obtain $2N_{shot}$ subimages of size $M \times (N/(2 \times N_{shot}))$, which are concatenated in the column dimension to form the final $M \times N$ image.

For efficiency, the phase errors of steps 2 through 4 are combined into a single precomputed matrix that is applied to each submatrix by elementwise multiplication. To perform the forward segmented FFT (image space to k-space), the steps are reversed, with the phase ramps and shifts negated. Steps 1 and 5 dominate the computational cost, and respectively require $O(MNN_{shot})$ and $O(MN\log(N/(2N_{shot})))$ operations.

## 3.3 Methods

### 3.3.1 Algorithm Implementation

The EPI-TrACR algorithm was implemented in MATLAB 2016a (The Mathworks, Natick, MA, USA) on a workstation with dual 6-core 2.8 GHz X5660 Intel Xeon CPUs (Intel Corporation, Santa Clara, CA) and 128 GB RAM. For each iteration of the algorithm's outer loop, CG-SENSE image updates were initialized with zeros to prevent noise amplification, and were performed using MATLAB's `lsqr` function and a fixed tolerance of $10^{-1}$, capped at 25 iterations. CG delay and phase updates were each fixed to a maximum of 5 iterations per outer loop iteration, and terminated early if all steps were less

than $10^{-6}$ cm$^{-1}$ (for delays) or $10^{-6}$ radians (for phase errors). The maximum permitted delay step in a single iteration was limited to $1/FOV$, and the maximum permitted phase step in a single iteration was limited to $\pi/10$ radians. Enforcing these maximum step sizes mitigated the effects of phase wraps in the calculated delay derivatives, by preventing the algorithm from taking large steps that may cause it to become stuck in local minima. Outer loop iterations stopped when the change in squared error was less than the previous iteration's error times $10^{-6}$. Code and example data for EPI-TrACR can be downloaded at https://bitbucket.org/wgrissom/tracr.

### 3.3.2 Experiments

A healthy volunteer was scanned on a 7T Philips Achieva scanner (Philips Healthcare, Best, Netherlands) with the approval of the Institutional Review Board at Vanderbilt University. A birdcage coil was used for excitation and a 32-channel head coil was used for reception (Nova Medical Inc., Wilmington, MA, USA). EPI scans were acquired with 24 $\times$ 24 cm FOV, 1.5 $\times$ 1.5 $\times$ 3 mm$^3$ voxels, TR 3000 ms, TE 56 ms, flip angle 60°. They were repeated for 1-4 shots, acceleration factors of 1-4x, and the 2-shot/1x scan was collected with 20 time points. The TE of 56 ms was chosen to facilitate side-by-side comparisons between images with different multishot and acceleration factors by maintaining the same contrast and matrix size between images, and was the shortest possible TE for the single-shot/1x acquisition, which had a readout duration of 102 ms. A calibration scan with phase encodes turned off was acquired in each configuration, and delays and phase errors were estimated from it using cross-correlation followed by an optimization transfer-based refinement [81]. SENSE maps were also collected using the vendor's mapping scan.

Images were reconstructed to 160 $\times$ 160 matrices using CG-SENSE with no corrections, and with phase error and delay estimates from the calibration scans; the latter reconstructions are hereafter referred to as 'conventional calibrated reconstructions.' Images were also reconstructed using EPI-TrACR initialized with the delays and phase errors from

the calibration scans, and with zeros. EPI-TrACR was further compared to the calibration-free PAGE EPI correction method [71]. Since EPI-TrACR's compute time depends on image and data size, the amount of k-space data necessary to estimate delays and phase errors was characterized by repeating the algorithm on 2-shot/1x in vivo data that was truncated in both k-space dimensions, across a range of truncation factors. The reconstructed image matrix sizes within EPI-TrACR were correspondingly reduced, so that the image matrix size matched the data matrix size. The final estimated delays and phase errors were then applied in a full-resolution CG-SENSE reconstruction. In order to characterize the sensitivity of EPI-TrACR to initialization, reconstruction of the 1-shot/1x data was repeated with a range of combinations of initial magnitude and phase error estimates. Except where indicated, displayed images shown are windowed down to 20% of their maximum amplitude for clear display of ghosting, and ghosted signals were measured in all images as the root-mean-square (RMS) signal outside an elliptical region-of-interest that excluded the brain and skull.

A separate experiment was performed in a phantom at 3T (Philips Achieva), using a volume coil for excitation and a 32-channel coil for reception (Nova Medical Inc., Wilmington, MA, USA). Data were collected for a single off-axis slice (5°/20°/30°) using a single-shot EPI scan with 60 dynamics; scan parameters were: $23 \times 23$ cm FOV, $1.8 \times 1.8 \times 4$ mm voxels, TR 2000 ms, TE 43 ms, flip angle 90°. The trajectory was measured for a single dynamic using a modified Duyn method [38, 82]. A SENSE map and a calibration scan with phase encodes turned off were also collected as for the 7T in vivo data. Delays were estimated from the measured trajectory as the average shift between each pair of odd and even lines over the middle quarter of the readout dimension. EPI-TrACR was used to reconstruct the phantom data in the same manner as for the 7T in vivo data. Residual ghosted signal was calculated for all images as the root-mean-square (RMS) signal outside an elliptical region-of-interest masking out the phantom.

## 3.4    Results

Figure 3.2 shows reconstructed images across multishot factors. Ghosting was lowest with EPI-TrACR in all cases. EPI-TrACR achieved lower ghosting than PAGE in all cases, and the differences between zero initialization and calibrated initialization EPI-TrACR images are negligible: averaged across multishot factors, the RMS difference between EPI-TrACR-estimated delays and phase errors with and without calibrated initialization was 0.014%. EPI-TrACR RMS ghosted signals were on average 37% lower than for conventional calibrated reconstructions, and 36% lower than for PAGE reconstructions. In addition, EPI-TrACR suppressed a strong aliased edge inside the brain which appeared in the 4-shot conventional calibrated reconstruction (indicated by the yellow arrow). Due to the long readout duration of the 1-shot acquisition, all the 1-shot reconstructions contain a similar off-resonance-induced geometric distortion at the back of the brain (indicated by the green arrow in the conventional calibrated reconstruction). The uncorrected 1-shot acquisition also contains much dimmer ghosts than the multishot acquisitions, so the corrected 1-shot images are more similar than the multishot corrected images.

Figure 3.3 shows reconstructed 2-shot EPI images with 1-4× acceleration. Compared to conventional calibrated reconstruction, EPI-TrACR with calibrated initialization again reduced ghosting up to 4× acceleration, and RMS ghosted signals were 18% lower on average. Compared to PAGE, EPI-TrACR with calibrated initialization had 44% lower ghosting on average. Furthermore, EPI-TrACR estimates matched with and without calibrated initialization up to 3× acceleration: averaged across factors of 1-3×, the RMS difference between estimated delays and phase errors with and without calibrated initialization was 0.024%.

Figure 3.4a plots RMS ghosted signal across repetitions for the 2-shot/1× scan, for conventional calibrated reconstruction, PAGE, and EPI-TrACR.

The signal levels are normalized to that of the first repetition's EPI-TrACR reconstruction. On average, residual ghosted signals in the conventional calibrated and PAGE recon-
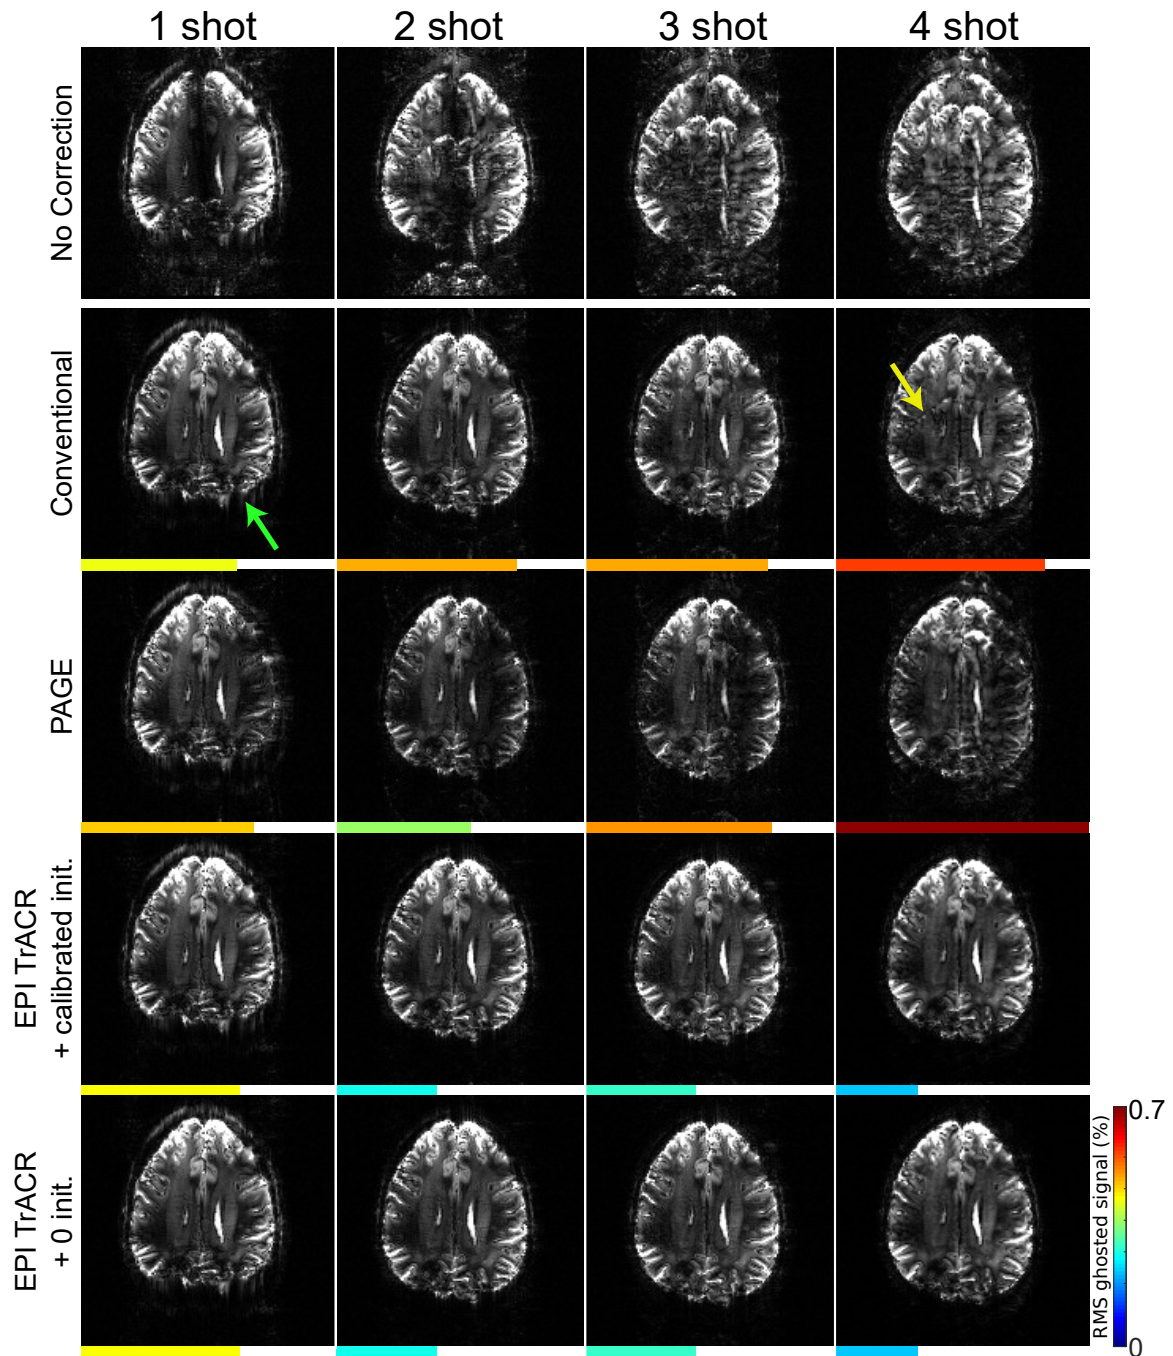
Figure 3.2: Multishot echo-planar images (no acceleration) reconstructed with no correction, conventional calibrated reconstruction, PAGE, EPI-TrACR with calibrated initialization, and EPI-TrACR with zero initialization. The length and color of the horizontal bars beneath each image represent the residual RMS ghosted signal as a percentage of maximum image intensity, as defined by the color scale on the right. The green arrow in the conventional calibrated 1-shot reconstruction indicates off-resonance-induced geometric distortion at the back of the head which appears in all of the 1-shot reconstructions. The yellow arrow in the conventional calibrated 4-shot reconstruction indicates an edge that aliased into the brain, which is not visible in the EPI-TrACR reconstructions.
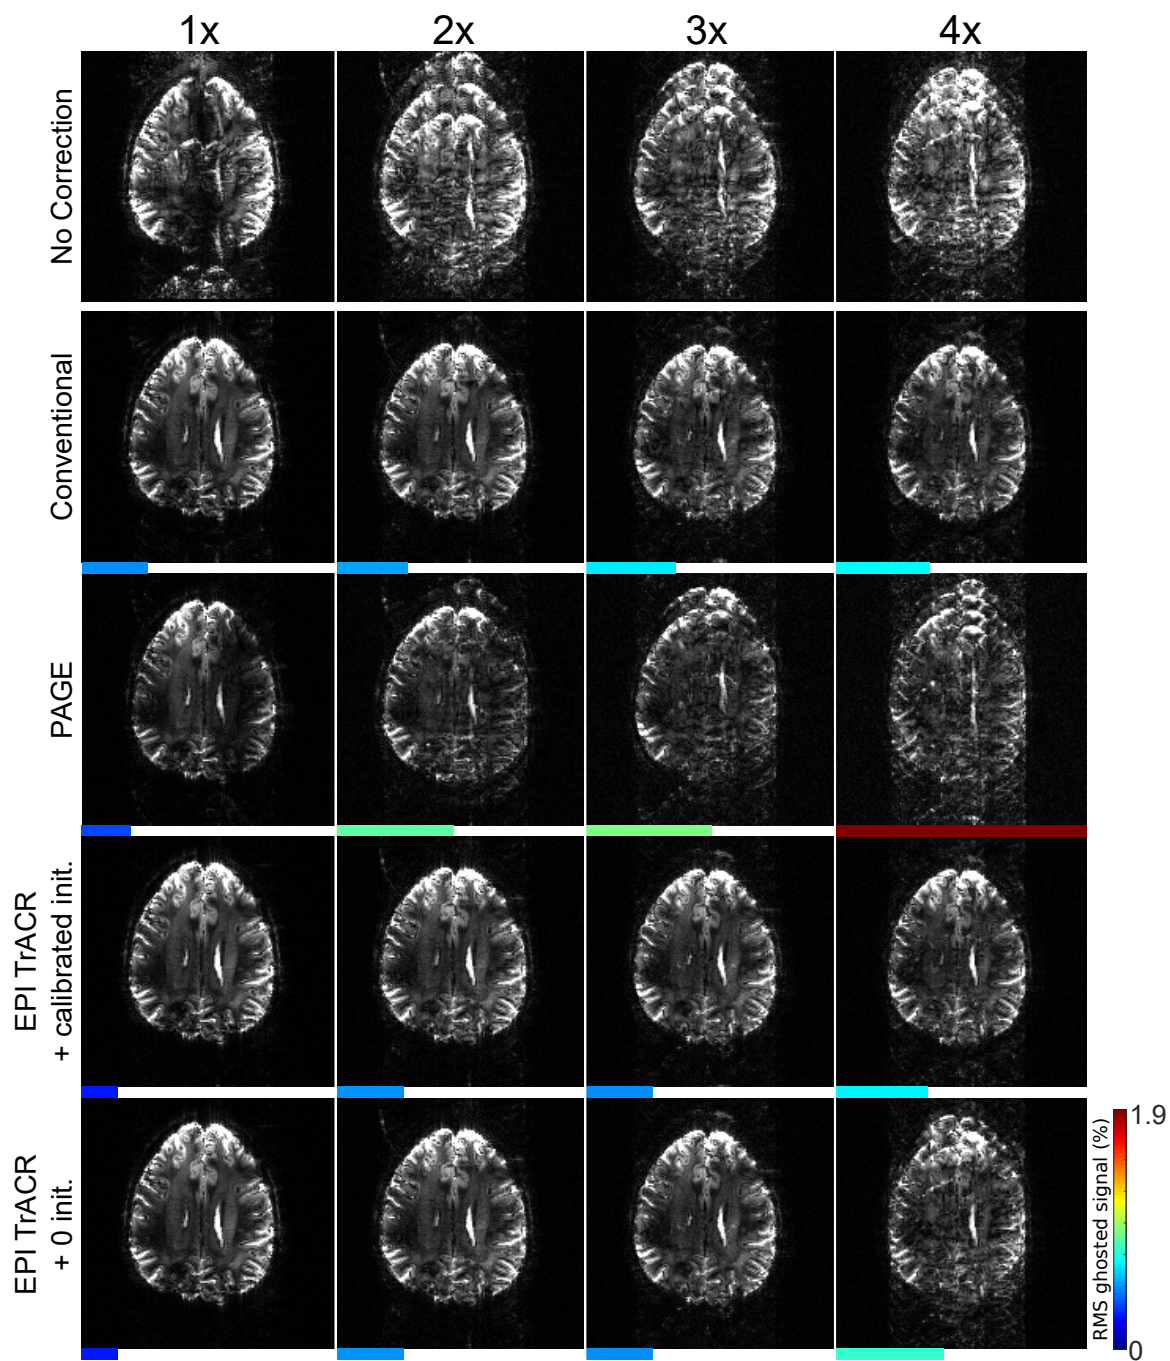
Figure 3.3: 1x-4x 2-shot echo-planar images reconstructed with no correction, conventional calibrated reconstruction, PAGE, EPI-TrACR with calibrated initialization, and EPI-TrACR with zero initialization. The length and color of the horizontal bars beneath each image represent the residual RMS ghosted signal as a percentage of maximum image intensity, as defined by the color scale on the right.
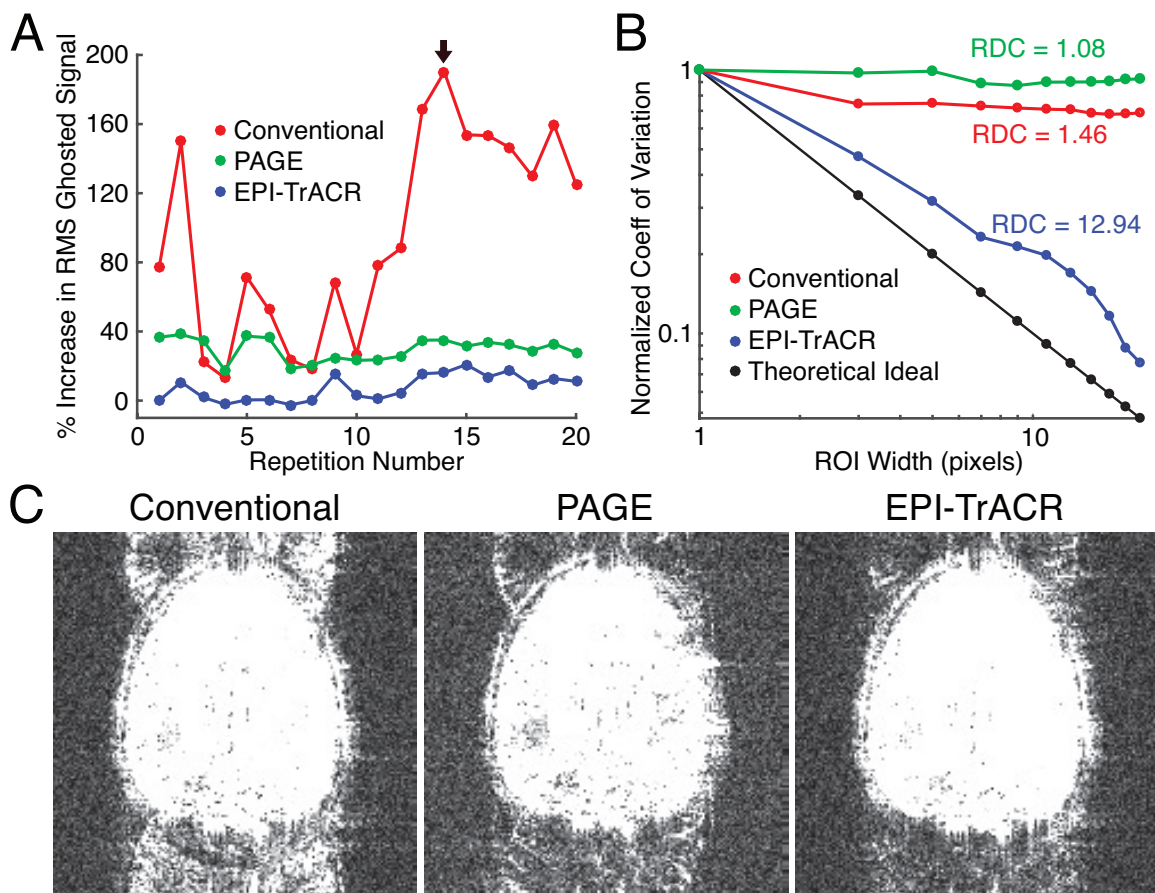
Figure 3.4: 2-shot echo-planar images over 20 repetitions reconstructed using conventional calibrated reconstruction, PAGE, and EPI-TrACR with zero initialization. (a) Percentage increase in RMS ghosted signal versus repetition, normalized to that of the EPI-TrACR reconstruction of the first repetition. (b) Weisskoff plot showing the normalized coefficient of variation over repetitions for an ROI of increasing size, for conventional calibrated reconstruction, PAGE, and EPI-TrACR compared to the theoretical ideal. (c) Windowed-down conventional calibrated reconstruction, PAGE, and EPI-TrACR reconstructions, at the 14th repetition (indicated by the arrow in (a)).

structions were respectively 80% and 20% higher than in the EPI-TrACR reconstructions. Figure 3.4b shows a Weisskoff plot [83] for all three reconstructions compared to the theoretical ideal; the coefficient of variation over repetitions is plotted for an ROI of increasing size. The EPI-TrACR time series had a radius of de-correlation (RDC) of 12.94 which was much higher than for the conventional calibrated and PAGE reconstructions, which had RDC's of 1.46 and 1.08, respectively. This indicates that, while PAGE reduced ghosting compared to conventional calibrated reconstruction, it was less temporally stable; i.e. residual ghosting was lower, but the image less consistent over time.

Figure 3.4c shows conventional calibrated reconstruction, PAGE, and EPI-TrACR (with zero initialization) images at the 14th repetition. The conventional, PAGE, and EPI-TrACR images at the 14th repetition respectively have 190%, 35%, and 16% higher RMS ghosted signal compared to the first repetition's EPI-TrACR reconstruction.

The truncated 2-shot EPI-TrACR results are shown in Figure 3.5. Figure 3.5a shows that delay and phase error estimation errors relative to full-data EPI-TrACR estimates are low up to very high truncation factors, and Figure 3.5b shows that compute time can be reduced up to 90% by truncating the data by 90%. Figures 3.5c and d show that images reconstructed with full data and 90%-truncated data delay and phase estimates are indistinguishable: RMS ghosted signal was 8% higher in the truncated EPI-TrACR image versus the full-data reconstruction, but still 40% lower than the conventional calibrated reconstruction (which appears in Figure 3.2). For greater than 90% truncation though, the compute time starts to increase again due to increasing iterations. For full data, EPI-TrACR reconstruction times using the described segmented FFT's ranged from one minute (for 1 shot, $1\times$ acceleration, and calibrated initialization) to 88 minutes (for 2 shots, $4\times$ acceleration, and zero initialization). In comparison, reconstructions using NUFFTs [80] in place of the segmented FFTs ranged from 8 minutes (for 1 shot, $1\times$ acceleration, and calibrated initialization) to 269 minutes (for 2 shots, $4\times$ acceleration, and zero initialization). Fig. 3.6 shows errors in final delay and phase shift estimates and ghosting, as a function of initial
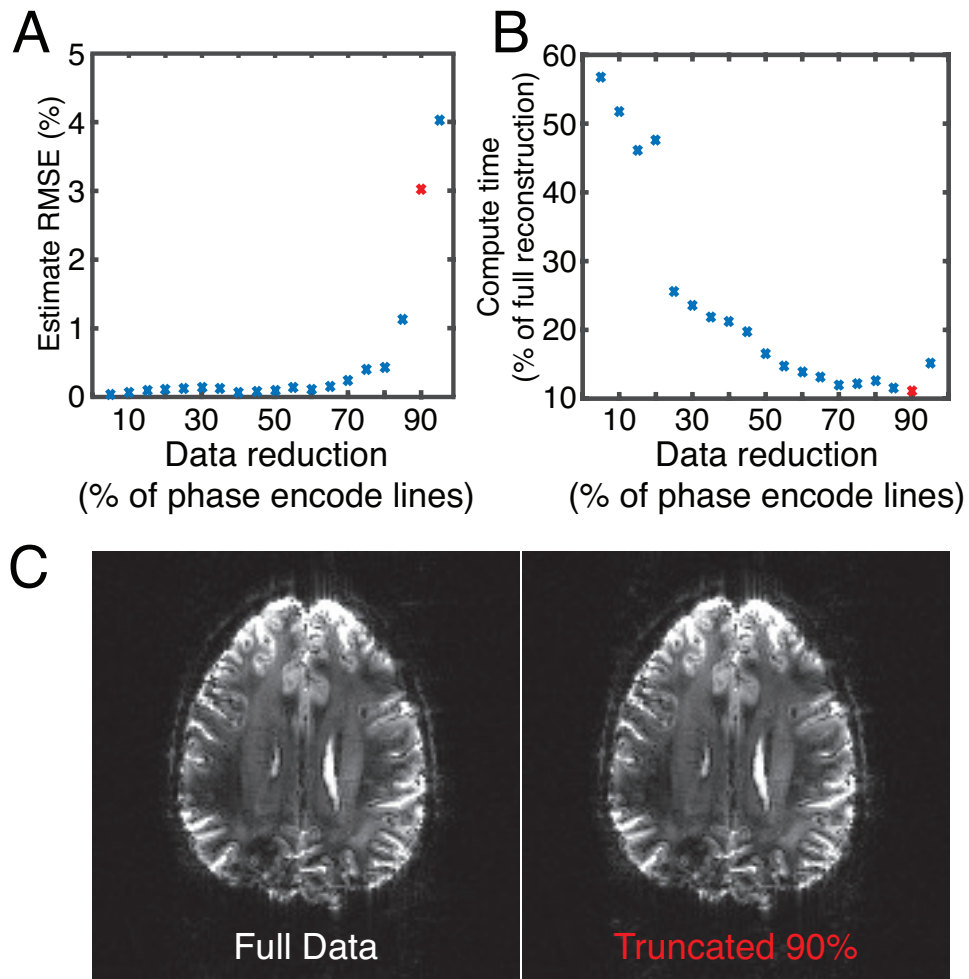
Figure 3.5: 2-shot, 1x-accelerated EPI-TrACR reconstructions from truncated data. The plots show (a) combined root mean square error (RMSE) in the estimates of DC and linear phase shifts compared to full-data EPI-TrACR estimates and (b) compute time as a percentage of a full-data EPI-TrACR compute time; both are shown as a function of percentage of degree of data reduction in each dimension. (c-d) Images reconstructed using the full data EPI-TrACR estimates using 90%-truncated EPI-TrACR estimates (16 PE lines) (d) (the red data point in a-b).
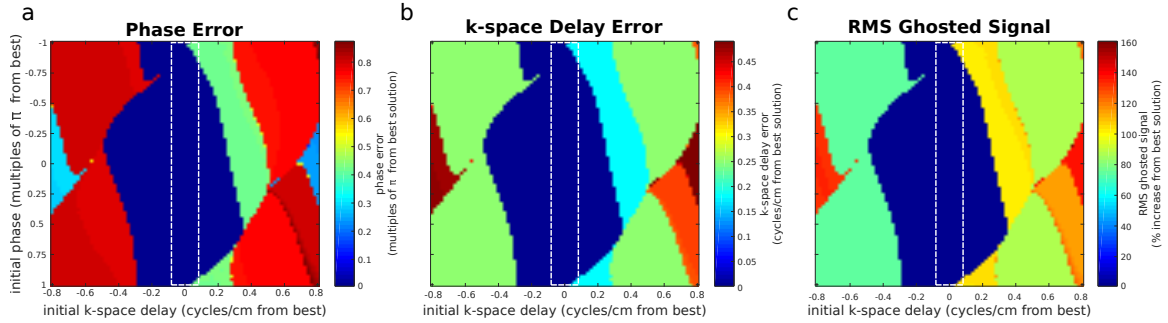
Figure 3.6: Sensitivity of EPI-TrACR to initialization. Here, the 1-shot/1x data was reconstructed using EPI-TrACR across combinations of erroneous initial phase errors and delays. Shown are the resulting final (a) phase error (multiples of $\pi$), (b) k-space delay (cycles/cm), and (c) RMS ghosted signal, for each initialization. The delays and phase shifts are expressed relative to the actual EPI-TrACR solution, which comprised a phase offset of -2.96 radians and a k-space delay of 0.075 cycles/cm. The white dashed boxes indicate the range of observed k-space delay and phase offsets in this work (across all multishot and acceleration factors).

values for the 1-shot/1× data.

Fig. 3.7a and b show boxplots of the line-to-line trajectory delays and DC phase errors, respectively, measured in the phantom at 3T. EPI-TrACR estimates (green) and conventional calibration estimates (dashed black) for the delays and phase errors are superimposed. The bulk even/odd line shift estimated was approximately 13% different between the two trajectories. Fig. 3.7c shows the corresponding images for the first repetition, reconstructed without correction, with the trajectory corrected by conventional calibration, with the trajectory and phase estimated by EPI-TrACR (with calibrated initialization), and with the measured trajectory. RMS image ghosting is 19% lower in the TrACR image than in the measured image. The conventional calibrated reconstruction did not correct for the large amount of ghosting in the uncorrected image. Both EPI-TrACR and measured-trajectory reconstructions had lower ghosting than the conventional calibration reconstruction.
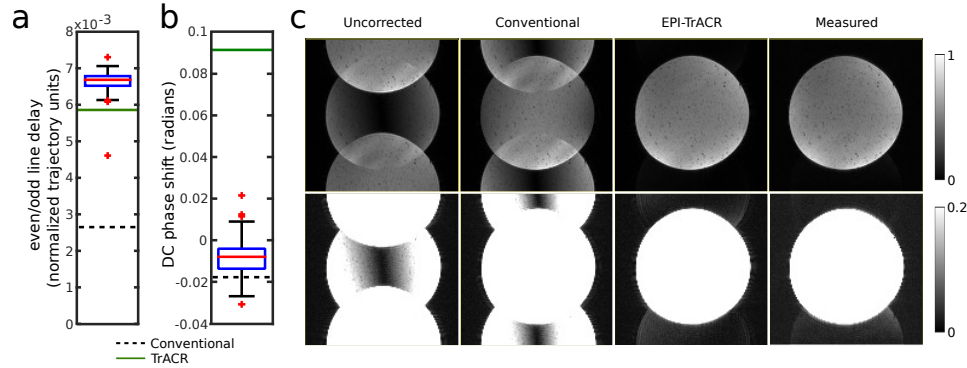
Figure 3.7: Shown in this figure are boxplots of the measured line-to-line trajectory delays in the readout dimension (a) and DC phase errors (b), with lines superimposed to mark the conventional calibrated (dashed black) and EPI-TrACR (solid green) estimates. (c) CG-reconstructed images of the first dynamic of phantom data using the uncorrected trajectory, the trajectory corrected by conventional calibration, the trajectory estimated by EPI-TrACR (with calibrated initialization), and the measured trajectory. Images are shown at full magnitude (top) and windowed to 20% (bottom).

## 3.5    Discussion

EPI-TrACR is an iterative maximum likelihood algorithm that jointly estimates EPI echo delays and phase errors, along with images that are compensated for them. Compared to conventional calibrated corrections, EPI-TrACR consistently reduced image ghosting across multishot factors, acceleration factors, and a time series, by 27% on average. It also reduced image ghosting compared to PAGE in all cases, by 40% on average. In most cases it was able to do so without being initialized with calibrated delays and phase errors. Because EPI-TRACR leverages data redundancy between nearby lines of k-space, its performance is expected to degrade as the distance between k-space lines increases with increasing acceleration factor, which was observed here in the zero-initialized $4\times$-accelerated case. However, when initialized with calibrated delays and phase errors, the method always reduced ghosting compared to conventional calibrated reconstruction. As can be seen in Fig. 3.6, when the initial delay values are too far from the global minimum, due to phase wraps in the calculated derivatives the algorithm becomes stuck in a suboptimal local minimum with high ghosting. However, the figure also indicates the ranges of phase shifts and

delays encountered in our data, which are within the region that converged to the global minimum. It may also be possible to expand the region of convergence using phase unwrapping or regularization [84]. It is demonstrated that EPI-TrACR results in a similar reconstructed image to the measured-trajectory reconstruction, as shown in Fig. 3.7, which provides additional confidence in the EPI-TrACR estimates. Residual ghosting apparent in both measured-trajectory and EPI-TrACR reconstructions may be attributed in part to the off-axis slice, which yielded particularly large trajectory and line-to-line phase shifts. The measured trajectory accounted for additional errors, such as slightly reduced k-space extent in the readout dimension, which are not captured in the EPI-TrACR reconstruction; however, these errors did not significantly degrade the EPI-TrACR reconstruction.

We chose the Polak-Ribière conjugate-gradient (CG) algorithm to update the delays and phase errors because it efficiently finds a local minimizer of the data-model error with respect to these parameters. Related methods could be used, such as gradient or steepest descent, which may have simpler formulations but would generally converge more slowly, and would still require the derivatives to be computed. Global optimization approaches such as genetic algorithms may be more robust to local minima than derivative-based methods, but would be impractical for more than one or two shots due to the number of dimensions that must be searched over while jointly estimating the image. Furthermore, we showed that the CG-based updates robustly converged to solutions that reduced image ghosting compared to conventional methods, especially when initialized with calibrated values.

Off-resonance is not currently modeled in EPI-TrACR, and may degrade delay and phase error estimates. While a full study of EPI-TrACR's off-resonance sensitivity is beyond the scope of this work, we note that: a) A measured field map could be incorporated in the signal model of Equation 3.1 [85]; b) Figure 3.5 showed that EPI-TrACR can accurately estimate delays and phase errors from a small number of k-space lines, over which there would be very little phase accrual due to off-resonance; and c) One of the main advantages of EPI-TrACR over previous image data-based methods is that it can be directly applied to

multishot and accelerated acquisitions, which have inherently reduced off-resonance sensitivity.

Compared to reconstruction with a fixed calibrated trajectory, the main tradeoff for EPI-TrACR's improved delay and phase shift estimates is increased computation, but this can be mitigated in several ways. First, we showed that compute time can be reduced by truncating the data matrix down to the low frequencies, without compromising the delay and phase shift estimates. Compute times are also shorter when the algorithm is initialized with calibrated estimates, since fewer iterations are required to reach a solution. The algorithm could be applied in parallel across repetitions or slices, or within the algorithm the FFTs could be parallelized across receive coils.

There are a number of ways the method could be extended. First, in the present work it was assumed that all the echoes within a set of even or odd echoes of a shot had the same delay and phase shift. However, it is also possible to estimate different delays and phase shifts for different echoes within a set by expressing them as a weighted sum of basis functions. We have previously tested this extension by expanding the delays and phase shifts of each set of odd and even k-space lines across multiple triangular basis functions, but found little improvement with our data. Nevertheless, as others may find it useful, this functionality is included in the provided code. Second, the method could be extended to jointly estimate a single set of delays and phase shifts over a whole stack of slices simultaneously, which would increase the effective signal-to-noise ratio for estimation. This could help in particular for highly accelerated acquisitions, in which the method is currently more sensitive to poor initialization. Finally, we note that in its current form, EPI-TrACR is not suitable for correcting shot-to-shot phase errors caused by bulk and physiological motion in multishot diffusion-weighted EPI, since these errors are generally higher than first order. It may however be useful as a preprocessing step to correct even/odd delays and phase shifts within each shot individually, which would then be followed by higher order inter-shot phase correction using a method such as Ref. [86]. It may also be possible to estimate

higher order phase error maps using EPI-TrACR by increasing its polynomial order, or to estimate a spatially resolved phase error map for each shot and set of even or odd lines; the latter approach would likely require spatial regularization of the estimated phase error maps [84].

## 3.6 Conclusions

The EPI-TrACR method alleviates ghosting artifacts by exploiting data redundancy between adjacent k-space lines in multicoil EPI data. It benefits from initialization with calibration data but does not require it at moderate acceleration and multishot factors. EPI-TrACR reduced dynamic ghosting without sacrificing temporal resolution, is compatible with multishot and accelerated acquisitions, and does not rely on a ghost-free image region. It was validated in vivo at 7T, at multiple acceleration and multishot factors and in a time series.

Chapter 4

RF Shim Learning

## 4.1   Introduction

In this chapter, a machine learning algorithm is presented to obviate the need for traditional tailored RF-shimming at high field by predicting RF shims from minimal calibration data. Despite the advantages afforded by a high signal-to-noise ratio, magnetic resonance imaging (MRI) at high field strengths has faced several barriers to widespread adoption. Among these is the increasing spatial inhomogeneity and subject dependence of transmit radiofrequency (RF) ($B_1^+$) fields with increasing field strength, resulting from an RF wavelength on the order of the size of the imaged object [2, 87]. The resulting signal non-uniformity creates spatially-varying contrast, can obscure underlying pathology, and complicates quantitative imaging [3, 5].

Several approaches exist to mitigate the effects of $B_1^+$ inhomogeneity. Frequency-swept adiabatic pulses are relatively insensitive to $B_1$ inhomogeneity [10]. However, they have long durations and high amplitudes, which leads to high specific absorption rate (SAR) and renders them impractical for many scans, and they require a tradeoff in frequency and spatial selectivity. Another approach is to position dielectric pads around the subject [11, 12]. These typically boost RF fields superficially, but they require subject-specific positioning and occupy valuable space in RF coils.

Patient-tailored RF shimming with multiple transmit channels and coils is a highly flexible approach to mitigating $B_1^+$ inhomogeneity, in which each coil in an array is driven with a unique amplitude and phase to achieve a more uniform combined field across a slice or volume [4, 13]. The amplitudes and phases can be tuned for each subject and scan geometry. However, optimizing these variables currently requires the $B_1^+$ fields of each coil

to be measured in each subject. Furthermore, since only the amplitude of the combined $B_1^+$ field needs to be uniform, optimization of the weights requires solution of a non-convex magnitude-least squares (MLS) problem [88]. Several advanced optimization approaches have been proposed to robustly solve this problem [89], but these require considerable computation, and a globally optimal solution cannot be guaranteed in a short computation time, especially for a large number of coils and slices. Furthermore, the fastest methods for full $B_1^+$ mapping currently require approximately 16 s per coil to scan a whole-brain volume at 3-4 mm isotropic resolution [14]. Mapping the 24 coils used in this work would require 6.5 minutes of scan-time, after which RF shim calculation would require at least 30 seconds per slice. For a whole-brain volume then, slice-by-slice mapping and RF shim calculation requires at least 10-15 minutes while the patient lies in the scanner, or approximately 25% of a total scan duration.

As an alternative to patient-specific $B_1^+$ mapping and RF pulse design, Gras et al. [90] have proposed the concept of "universal" multidimensional parallel transmission pulses, which are jointly optimized over a large set of patient brains such that a single pre-optimized pulse could be applied to any adult brain. However, this approach trades homogeneity for broad applicability, since by nature solutions cannot be both universal *and* tailored, and to our knowledge it has only been applied to 3D non-selective multidimensional $k_T$-points excitations [8]. Mirfin et al. explored using a neural network for tailored spokes pulse prediction [91], but this approach was not successful in interpolating over the space of MLS problem solutions, which can have very different phase profiles. This resulted in an inability to produce homogeneous flip angle profiles.

In this chapter, a machine learning method is presented and tested in simulation for the instantaneous prediction of patient-tailored, SAR-efficient RF shims. The method, called RF Shim Prediction by Iteratively Projected Ridge Regression (PIPRR), avoids the weakness of Mirfin et al's approach by merging the design of the training shims with the training of the learner. This makes it possible to interpolate over the training shims using kernelized

ridge regression. Simulation results show that PIPRR can predict SAR-efficient RF shim weights with minimal $B_1^+$ map data and negligible online compute times. In the following, we describe the PIPRR method, characterize it in terms of the amount of training data required, robustness to noise, and required features, and validate it against directly designed shims as well as shims predicted by nearest neighbor using the same features. Preliminary reports of this work were presented in Refs. [92–94].

## 4.2  Theory

### 4.2.1  Magnitude Least-Squares RF Shimming

Patient-tailored RF shimming is typically posed as a magnitude least squares optimization problem [19, 20, 88, 95–97]:

$$\hat{\mathbf{b}} = \underset{\mathbf{b}}{\operatorname{argmin}} \frac{1}{2} \left\| \mathbf{m} - |\mathbf{Ab}| \right\|_{\mathbf{W}}^2 + \frac{\lambda}{2} \mathrm{R}(\mathbf{b}), \tag{4.1}$$

$$= \underset{\mathbf{b},\phi}{\operatorname{argmin}} \frac{1}{2} \left\| \operatorname{diag}\left(e^{\imath\phi_i}\right) \mathbf{m} - \mathbf{Ab} \right\|_{\mathbf{W}}^2 + \frac{\lambda}{2} \mathrm{R}(\mathbf{b}), \tag{4.2}$$

where the length-$N_c$ vector $\mathbf{b}$ contains the complex-valued transmit RF weights for each coil, $\mathbf{m}$ is a vector containing the desired excitation pattern at all spatial locations ($\mathbf{m}$ is typically a vector of ones), $\phi$ is the desired target phase pattern which replaces the absolute value around $\mathbf{Ab}$ and is jointly optimized with the RF weights, $i$ indexes spatial locations, $\mathbf{A}$ is a matrix containing the $B_1^+$ maps for each coil at each spatial location, which are concatenated in the column dimension, the diagonal matrix $\mathbf{W}$ contains spatially-dependent weights that select samples within a tissue mask, $\lambda$ is a regularization parameter, and R is a quadratic regularization function that can be used, e.g., to regularize SAR or RF power. This work focuses on multislice RF shimming, in which a unique shim vector $\mathbf{b}$ is determined for each slice in an imaged volume. A typical iterative approach to solve Equation 4.2 alternates between updating the RF shim weights $\mathbf{b}$ while holding the target phase $\phi$

fixed, and updating the target phase $\phi$ for fixed shim weights **b** by setting it equal to the phase of **Ab**. The shim weights can be updated using a standard solver for regularized least-squares problems, such as regularized pseudoinverse or the conjugate gradients (CG) algorithm.

### 4.2.2 Kernelized Ridge Regression Prediction of RF Shims

Instead of designing RF shims by solving the problem in Equation 4.2, we will apply kernelized ridge regression (KRR) to predict the complex RF weights **b** from a length-$N_{feat}$ complex vector of slice-specific features **f**, which includes a bias entry. The RF shim weight $\hat{b}_j$ for coil $j$ predicted by KRR can be written as:

$$\hat{b}_j = \mathbf{p}_j^T \mathbf{f}, \tag{4.3}$$

where the length-$N_{feat}$ complex-valued feature weight vector $\mathbf{p}_j$ relates each feature to coil $j$'s shim weight. The entries of **f** are normalized by the mean and standard deviation of each feature across the training slices. The method is kernelized because we include non-linear transformations (specifically, first order cross-products) of features in **f**, as will be described later. KRR learns the feature weights contained in the $\mathbf{p}_j$ vectors by fitting a regularized linear model to a set of training shim weights for each coil, as:

$$\hat{\mathbf{p}}_j = \operatorname*{argmin}_{\mathbf{p}_j} \frac{1}{2} \left\| \mathbf{b}_{j,\text{train}} - \mathbf{F}\mathbf{p}_j \right\|^2 + \frac{\beta}{2} \left\| \mathbf{p}_j \right\|^2, \tag{4.4}$$

where the length-$N_{train}$ vector $\mathbf{b}_{i,train}$ contains coil $i$'s shim weights for all the training slices, **F** is an $N_{train} \times N_{feat}$ matrix of features for each training slice, and $\beta$ is the ridge regression regularization parameter. Each column of **F** is normalized to have zero mean and standard deviation one. Given this model, one might directly use a set of training shims that are solutions to Equation 4.2 to solve Equation 4.4 for the feature weights. However, as will be shown later, this results in very poor shims due to large variations in shimmed

phase between MLS problem solutions. In the following, we describe the RF Shim PIPRR algorithm, which solves this problem by merging the design of a set of training shims for KRR with learning the coil-specific feature weights.

### 4.2.3 RF Shim Prediction by Iteratively Projected Ridge Regression (PIPRR) Algorithm

In order to derive a set of training shims that can be fitted with low errors to each slice's features using KRR, we propose to merge the design of the training shim weights for KRR (the $\mathbf{b}_{j,\text{train}}$ vectors) with learning the coil-specific feature weights (the $\mathbf{p}_j$ vectors). This is realized by alternately taking a few iterations towards the solution to Equation 4.2, and projecting the current training shim weights $\hat{\mathbf{b}}_{j,\text{train}}$ onto the space of shim weights that are predictable by KRR, as:

$$\hat{\mathbf{b}}_{j,\text{train}} \leftarrow \mathbf{F}(\mathbf{F}^H\mathbf{F} + \beta\mathbf{I})^{-1}\mathbf{F}^H\hat{\mathbf{b}}_{j,\text{train}}, \tag{4.5}$$

for each coil, where $H$ denotes Hermitian transpose. Note that the regularized projector matrix $\mathbf{F}(\mathbf{F}^H\mathbf{F} + \beta\mathbf{I})^{-1}\mathbf{F}^H$ does not change and is calculated once, before the iterations start. The steps of the training algorithm are summarized as follows, and are also shown in the yellow box of Figure 4.1:

---

**PIPRR: Training shim design and weight learning**

---

1: Set target phase $\phi$ to circularly polarized mode (CP) phase.

2: **repeat**

3:     Update training shims $\hat{\mathbf{b}}$ to minimize Equation 4.2 for each training slice individually, using a few CG iterations (Equation 4.2).

4:     Project training shim weights onto the set that is predictable by KRR (Equation 4.5).

5:  Update target phase $\phi$ to that produced by current shims $\hat{\mathbf{b}}$.

6: **until** change in cost (Equation 4.2 summed across training slices) $<$ tolerance

7: Calculate final feature weights for each coil by solving Equation 4.4, using regularized pseudoinverse.
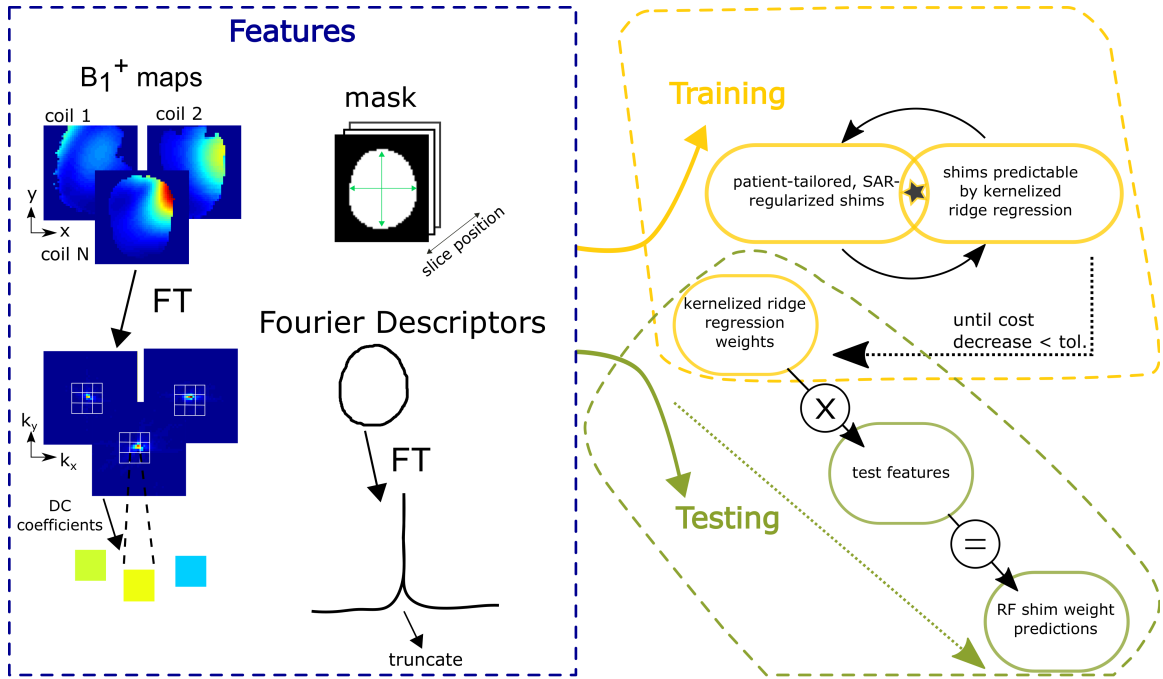


Figure 4.1: The RF shim Prediction by Iteratively Projected Ridge Regression (PIPRR) algorithm. (Blue box) Features for each slice include DC coefficients of the coils' $B_1^+$ maps and tissue mask metrics, including mask centroid, standard deviation of $x$ and $y$ coordinates within the brain mask, slice position, Fourier shape descriptors of the mask contour, and all first-order cross-terms of these features. (Yellow box) The training stage consists of feeding the features, $B_1^+$ maps and SAR matrices into an alternating minimization targeting SAR-efficient, homogeneous RF shim solutions that are predictable via kernelized ridge regression. (Green box) Testing involves predicting RF shims for new subjects by applying the kernelized ridge regression weights learned in the training stage to the new subject's features.

Overall, this procedure corresponds to an alternating minimization over a non-convex set (the set of solutions to Equation 4.2) and a convex set (the vector space spanned by the feature projector matrix) [98].

## 4.3    Methods

### 4.3.1    Electromagnetic Simulations and Features

A 100-subject 3D in silico phantom dataset of $B_1^+$ maps was generated by simulating a 24-element loop coil array with shielding at 7T in XFDTD (Remcom Inc., State College, PA, USA). The coils were arranged in 3 rows of 8, as shown in Figure 4.2. The total
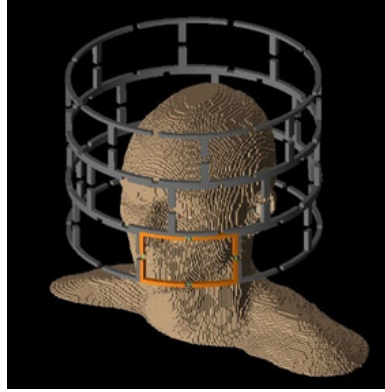


Figure 4.2: The 24 channel loop coil, simulated in XFDTD to obtain $B_1^+$ maps and SAR matrices. The loops were arranged in 3 rows of 8 elements each. The total height of the array was 20.5 cm, and its diameter was 30 cm.

height of the array was 20.5 cm, and the diameter was 30 cm. The 100 subjects were generated by magnifying Ella and Duke head phantoms from the Virtual Family [99] in three dimensions, according to a population-based normal distribution [100]. Magnification factors ranged from 0.93x to 1.10x, 0.87x to 1.08x, and 0.92x to 1.21x in the left-right, anteroposterior, and craniocaudal dimensions, respectively. This corresponds to head sizes ranging from 13.5-16.5 cm left-right and 17.7 to 21.9 cm anterior-posterior. Figures 4.3a and b show the range of independent scaling factors of original Duke and Ella models in each dimension. The resolution of the simulated maps was 5 mm isotropic, with a slice-gap of 0 mm. The dataset contained 31 axial slices for each subject, centered in the middle of the coil; slices containing fewer than 20 voxels were discarded. To regularize local and global SAR, virtual observation point (VOP) SAR matrices were calculated for each simulated head using a parameter $u$ equal to 20% of the maximum spectral norm [101], and
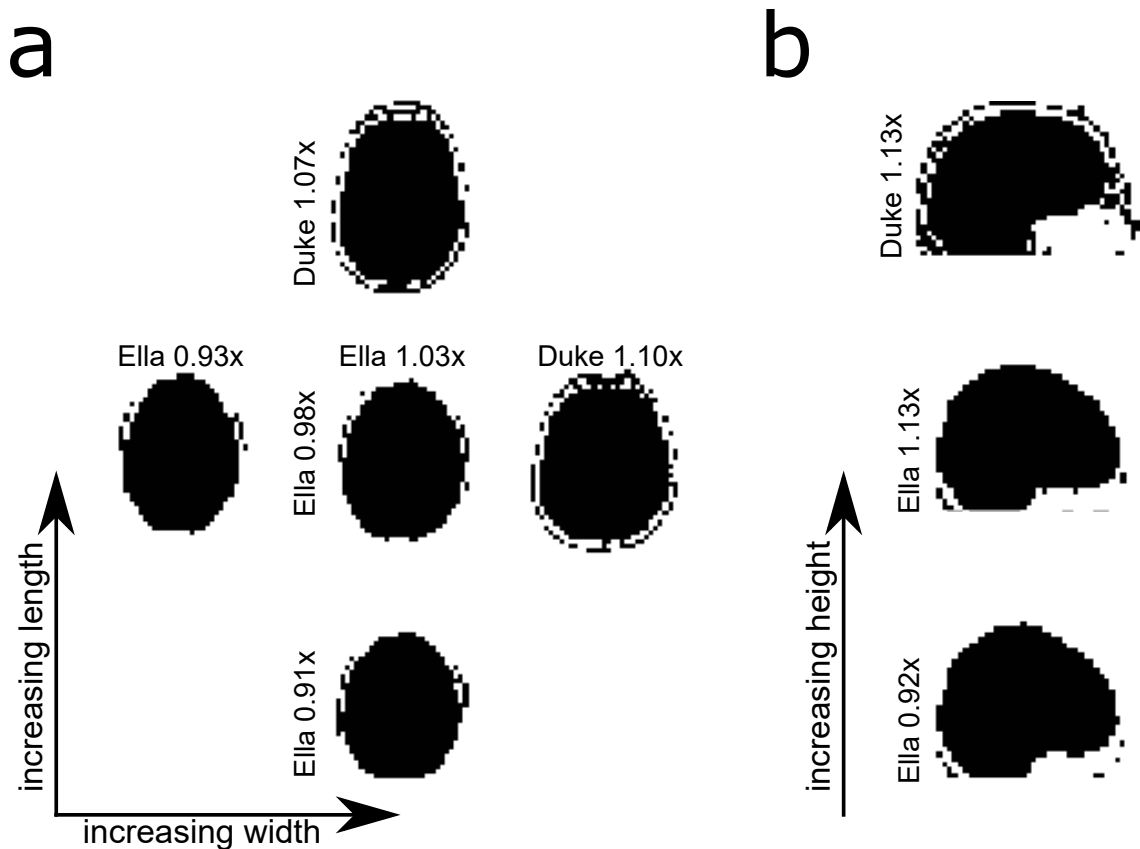
Figure 4.3: a) Tissue masks for the center transverse slice in five subjects, to demonstrate the variation in head sizes across all subjects. Maximum and minimum head widths and lengths are shown, as well as the median size (middle). b) Central sagittal tissue masks for the subjects with maximum, median, and minimum head height. The numbers next to the names indicate the amplification factor applied to the original Duke or Ella model in the corresponding dimension.

were summed with a global SAR matrix which was scaled by 3.1 to reflect the tighter limit on global head SAR. With this construction, the regularization term in Equation 4.2 had the form: $R(\mathbf{b}) \triangleq \mathbf{b}^{\mathbf{H}}\mathbf{C}\mathbf{b}$, where $\mathbf{C}$ is the total SAR regularization matrix. For comparison with RF Shim PIPRR, directly-designed shims were computed for every slice of every subject by solving Equation 4.2 using 100 random target phase initializations, which were obtained by generating 100 random sets of shim weights. The result that minimized the cost function of Equation 4.2 was taken as the 'Direct Design' shim for each slice.

A vector of 561 features was calculated for each simulated brain slice. This included the DC Fourier coefficient of each coil's $B_1^+$ map, the tissue mask centroid, standard devi-

ations of the within-mask x and y coordinates, slice z-position, the central 3 Fourier shape descriptors [102] of the slice mask, all 1st-order cross-terms of these features, and a bias term. The features are illustrated in the blue box of Figure 4.1.

### 4.3.2 Algorithm Implementation

PIPRR was implemented in MATLAB R2015a (The Mathworks, Natick, MA, USA) using Vanderbilt University's Advanced Computing Center for Research and Education (ACCRE) cluster to parallelize computation across k-folds. The RF shim weight updates used 3 CG iterations, the SAR regularization parameter $\lambda$ was 0.01, and the KRR regularization parameter $\beta$ was 1. The target excitation pattern vectors **m** were set to one inside the tissue masks for each slice. The PIPRR stopping criterion was set as described below.

### 4.3.3 Experiments

#### 4.3.3.1 Comparison to Other RF shim Designs

To evaluate the PIPRR algorithm, a 10-fold random cross-validation was performed across all the simulated heads, with 90 of the phantom heads used for training and 10 used for testing; slices from the same head were not split between the sets. Shims were predicted for slices in the test sets using the learned feature weights from the training set. The first k-fold of test data was used to determine a stopping tolerance for PIPRR, to prevent overfitting in the other k-folds. The tolerance was set equal to the maximum of $10^{-6}$ and the consecutive difference in training shim costs (the sum of the cost in Equation 4.2 across all training slices) immediately before the test shim costs started increasing. The first k-fold was then omitted from all results, and the other k-folds used the determined stopping tolerance. PIPRR shims were compared to shims obtained by nearest-neighbors (NN), directly designed shims, and CP mode shims, in terms of shimmed $B_1^+$ standard deviation and the calculated SAR regularization term R(**b**).

### 4.3.3.2 Required Training Data and Features, and Noise Sensitivity

Four additional experiments were performed to characterize the PIPRR algorithm, in terms of required training data, sensitivity to noise, importance of different feature classes, and the amount of $B_1^+$ mapping data required. First, to characterize the amount of training data required by the algorithm, PIPRR training was repeated using 10 to 90 randomly-selected training heads from each fold, and the shimmed test $B_1^+$ standard deviation was calculated for each number of training heads. The effect of additive noise in the test features was studied by adding noise over a range of signal-to-noise ratios (SNRs) to the test $B_1^+$ maps in each fold, calculating the corresponding SNRs of the $B_1^+$ map DC Fourier coefficients, and then adding noise to the remaining features to achieve the same SNRs. A total of 10 $B_1^+$ map SNR levels between 10 and 100 were simulated, and the shimmed $B_1^+$ standard deviation was calculated for each level. The importance of each feature class was determined by calculating the norm of the set of weights assigned to each feature in PIPRR's final KRR weight learning step, while varying the regularization parameter $\beta$ over several orders of magnitude. The weight norms were combined across folds, and the feature classes were ranked according to their norms. The effect of each feature class on the quality of test shim predictions was then calculated by dropping all feature classes from the model and then individually re-incorporating them in their order of importance. Finally, the effect of further reducing the required $B_1^+$ mapping data for PIPRR was characterized by leaving out coils' $B_1^+$ DC Fourier coefficients from the PIPRR training and predictions for each fold. This ranged from the full set of 24 coils, down to zero coils, with approximately equidistant spacing of retained coils.

**4.4.0.1 Comparison to Other RF Shim Designs**

Figure 4.4 shows the best-, median-, and worst-case $B_1+$ patterns across all slices and folds in terms of $B_1^+$ standard deviation for CP mode, NN, Direct Design, and PIPRR training and test shims. Figure 4.5a shows box plots of $B_1^+$ inhomogeneity for each method,
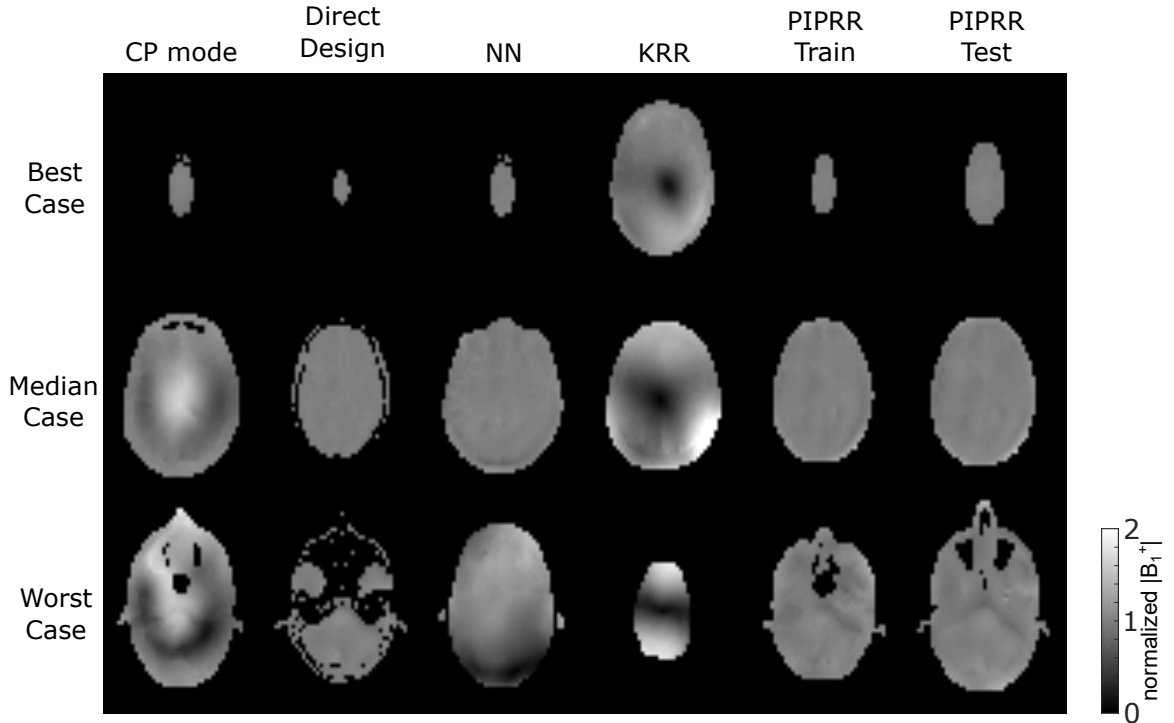


Figure 4.4: Shimmed $B_1^+$ patterns for the best-, median-, and worst-case (in terms of shimmed $B_1^+$ inhomogeneity) slices across all test slices, for circularly polarized (CP) mode, direct design, nearest neighbors (NN), kernelized ridge regression (KRR) applied to the Direct Design shims, and PIPRR (PIPRR Test). PIPRR training shim patterns are also shown.

across all test slices. Mean $|B_1^+|$ standard deviation was 21.3% for CP mode shims, 2.1% for Direct Design shims, 4.1% for NN shims, 38.9% for KRR shims, and 3.1% for PIPRR-predicted training and test shims. Thus, direct design produced the most uniform shims, followed by PIPRR. There is also little difference between the PIPRR training shims and the Direct Design shims, indicating that the projection step did not significantly degrade the quality of the training shims. NN also produced relatively uniform shims on average,
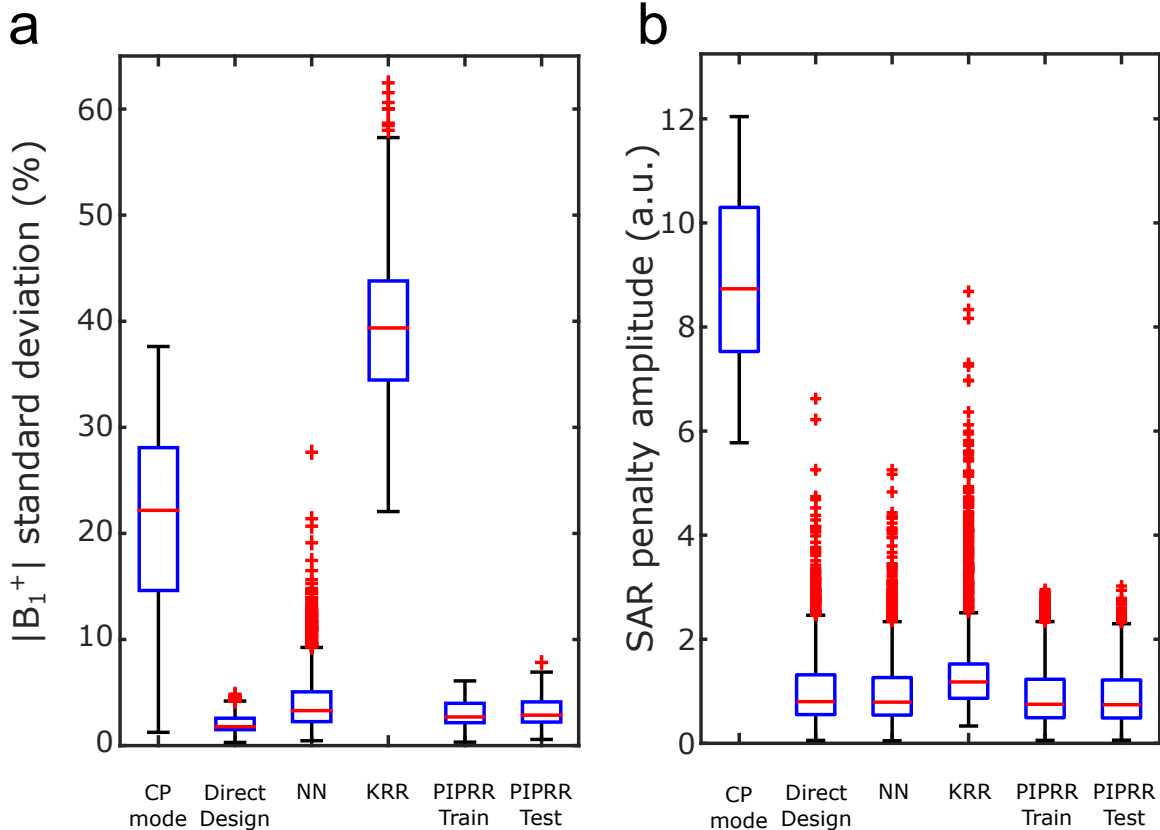
Figure 4.5: a) $B_1^+$ pattern inhomogeneity across all test slices for circularly polarized (CP) mode, Direct Design, Nearest Neighbors (NN), kernelized ridge regression (KRR) applied to the Direct Design shims, and PIPRR (PIPRR Test). PIPRR training shims are also shown. b) The corresponding SAR penalty terms across all test slices. The values are normalized to the mean SAR penalty of the Direct Design shims. Blue box edges delineate the 25th and 75th percentiles, and medians are indicated by the red bars. Red crosses indicate outliers (values that exceeded the 75th percentile by greater than $1.5 \times$ the difference between the 75th and 25th percentiles). The black whiskers indicate the extent of data not considered outliers.

but also produced outliers with very poor shims. Conventional KRR applied to the Direct Design shims failed because of the large variation in phase between the training shims. Figure 4.5b shows the SAR penalty terms for each test slice, normalized to the mean of the Direct Design SAR penalties. The mean SAR penalty term, normalized to that of the Direct Design shims, was 8.78 a.u. for CP mode, 0.98 a.u. for NN, 1.40 a.u. for KRR, 0.89 a.u. for PIPRR training shims and 0.88 a.u. for PIPRR test shims. The PIPRR SAR penalty values are similar to or lower than the Direct Design values, and are lowest among the predicted

shims. The average time to calculate PIPRR shims for a single new test slice was 4.92 ms (4.86 ms for feature calculation and 0.06 ms for prediction via multiplication with the feature weights).

#### 4.4.0.2 Required Training Data and Features, and Noise Sensitivity

Figure 4.6 shows box plots of one fold's test set $B_1^+$ inhomogeneity when the final KRR weights are learned using a varying number of randomly-chosen heads from the training set. A steep drop-off in predicted profile inhomogeneity occurs between 50 and 60 heads.
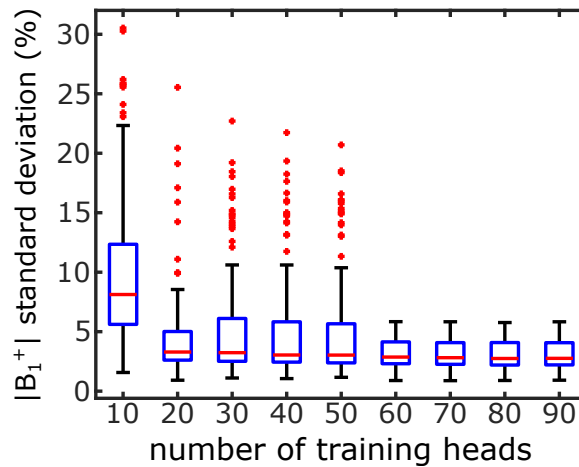


Figure 4.6: Shimmed $B_1^+$ inhomogeneity of one fold's test set slices, when varying the number of heads included in the final KRR weight learning. The homogeneity of the predicted shim profiles is comparable to those predicted with the full 90-head training set when at least 60 heads are included in the weight learning.

Inhomogeneity of predicted shims was comparable when 60 heads was used for training versus the full 90-head training set. Figure 4.7 shows that PIPRR predictions tolerate a moderate amount of noise in the features without significant degradation in homogeneity of the shim predictions. Box plots of test $B_1^+$ pattern inhomogeneity are shown with increasing feature SNR, which is reported as a function of the equivalent $B_1^+$ map SNR. Linearly increasing feature SNR results in approximately exponentially decreasing inhomogeneity. A feature SNR of 40 or above produces approximately the same level of homogeneity in the predictions as noise-free features, at 3.4% and 3.1%, respectively. Figure 4.8 shows the
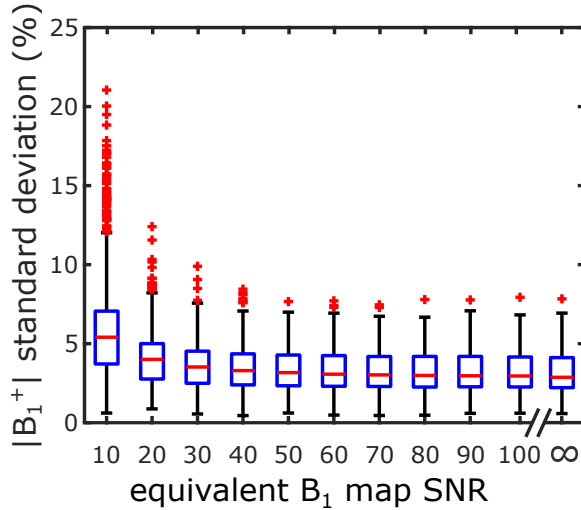
Figure 4.7: Shimmed $B_1^+$ inhomogeneity of one fold's test set slices, with noise of varied amplitude added to the features used for PIPRR prediction. Noise level is reported in terms of equivalent $B_1$ map SNR. The no-noise case is indicated by SNR $= \infty$.

test set $B_1^+$ inhomogeneity versus the feature classes included in the KRR model, beginning with the feature class of highest importance, which is the product of the mask centroids and the $B_1^+$ map DC Fourier coefficients, and adding in the other feature classes in order of importance, which was determined as described in the Methods. The inhomogeneity levels off after five feature classes are included, and the most important classes involve products of $B_1^+$ map DC coefficients and tissue mask features. Figure 4.9 shows how the homogeneity of PIPRR shims depends on the number of coils whose $B_1^+$ map DC coefficients are included as features. Average $B_1^+$ inhomogeneity increased approximately linearly with a reduced number of coils, but remained less than 6% even with no coils.

## 4.5  Discussion

### 4.5.1  Summary and Implications of Results

Computational experiments demonstrated that RF shim PIPRR can predict SAR-efficient tailored RF shims with homogeneous $B_1^+$ patterns in 100 simulated heads with population-representative dimensions. The predicted shims were consistently more homogeneous than
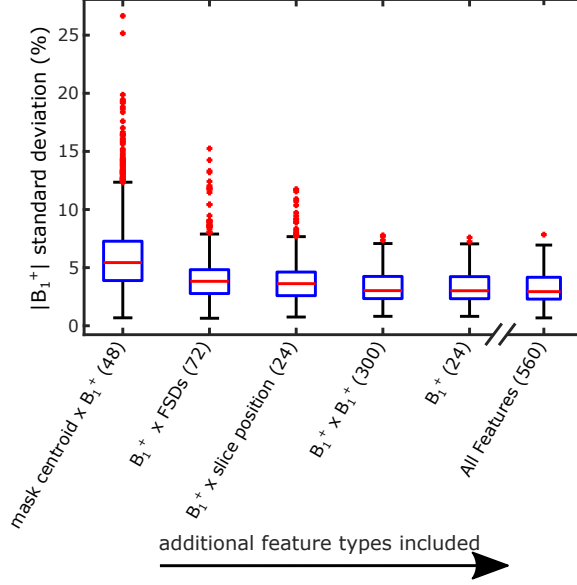
Figure 4.8: Analysis of feature importance. The box plot shows $|B_1^+|$ standard deviation of one fold's test set slices, as feature groups are accrued into the final KRR weight learning and testing, in order of importance. The number of features in each class is reported in parentheses next to each class. Importance was measured as the norm of the KRR weights on each feature class over a range of KRR regularization parameters. The number of features included in each class is shown in parenthesis next to the feature group. Cross-terms of mask centroids, $B_1^+$ DC Fourier coefficients, Fourier shape descriptors (FSDs), and slice position were the most important features.

CP mode and NN-predicted shims, and were only slightly less homogeneous than Direct Design shims (3.1% versus 2.1% average $|B_1^+|$ standard deviation), with 12% lower SAR penalty. Compared to Direct Design shims, which require full $B_1^+$ maps of all coils in a slice, RF shim prediction with PIPRR required only a tissue mask and the DC coefficients of the coils' $B_1^+$ maps in a slice. A tissue mask can be derived from a fast low-flip angle gradient-recalled echo scan, which is commonly performed as a prescan step for $B_0$ shimming or receive sensitivity mapping. The $B_1^+$ map DC coefficients could be measured with a single excitation for each coil and slice, which would be approximately two orders of magnitude faster than full $B_1^+$ mapping since no in-slice phase encoding would be required. It was further shown that inhomogeneity degraded slowly as the number of $B_1^+$ map coefficients was reduced, indicating that coils could be skipped altogether. Additionally, the $B_1^+$ inhomogeneity of PIPRR-predicted shims also rose slowly as the SNR of the features
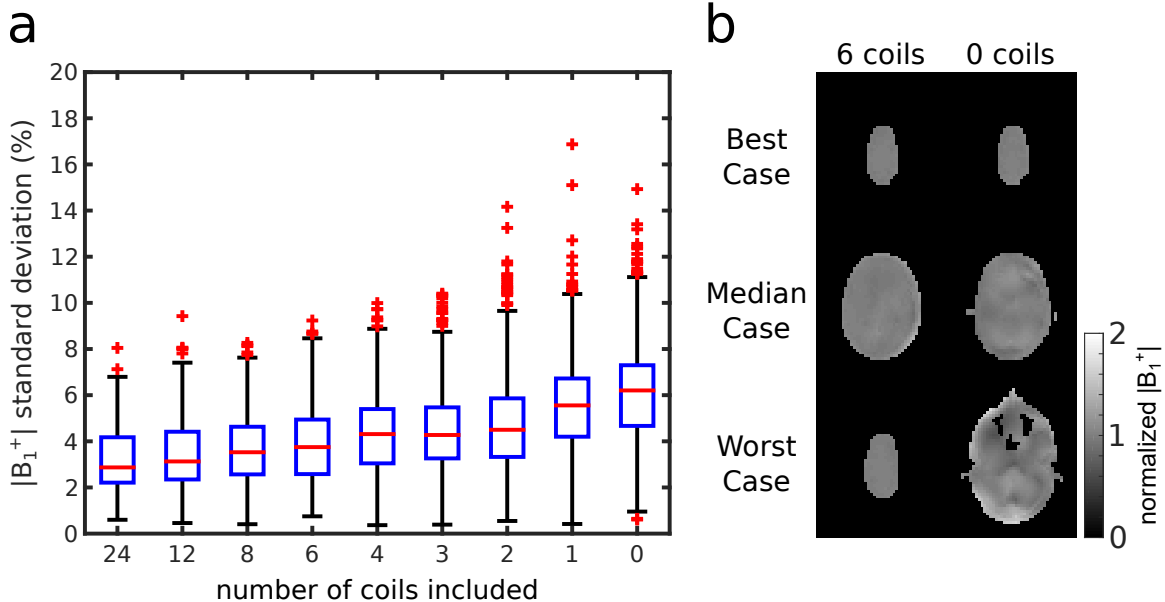
Figure 4.9: PIPRR test performance with reduced sets of $B_1^+$ map DC coefficients. a) Shimmed $B_1^+$ inhomogeneity of one fold's test set slices, versus the number of coils whose DC Fourier coefficients were included in the final KRR weight learning. b) Shimmed $B_1^+$ patterns for the best-, median-, and worst-case predicted test slices when 6 coils' coefficients were included, and when zero coils' coefficients were included (i.e., only the tissue mask size, shape, and position features were used for prediction).

was decreased, indicating that it is not highly sensitive to noise. The computational burden of making a set of RF shim predictions with PIPRR is also much lower than for Direct Design of RF shims (approximately 5 ms to both compute the features and make a prediction). Overall, the PIPRR method could alleviate the prescan and computational burdens of slice-specific RF shimming for high-field MRI.

### 4.5.2 Importance of Iterative Training

PIPRR is fairly unique among machine learning techniques for its iterative re-design of its own training data. The need for this was demonstrated by the very poor performance of shims predicted by conventional KRR when it was trained on Direct Design shims, as was shown in Figures 4.4 and 4.5. Even the best KRR slice had a null in the field pattern, and the shims were overall less homogeneous than CP mode shims. This resulted from

the fact that the MLS shim design problem is non-convex with many local minima and infinite solutions (since any phase-shifted solution is also a solution), which makes it very difficult for any ML method to learn the relationships between the features and the shims, without a very high order model. Early in this work, we investigated whether MLS shim solutions in a training set could be made more similar by removing the average phase shifts between them prior to learning the KRR weights, but this did not significantly improve the KRR training error or predictions (results not shown). An alternative approach may be to regularize differences in target phase patterns or RF shims between slices and subjects, but this would make the MLS problem more complicated and may overly constrain the designs without significantly improving KRR training error. PIPRR overcomes this problem by regularizing the training shims so that they can be more easily fit by KRR. It is possible to use other ML methods in the same iterative framework introduced by PIPRR; KRR was chosen for this work because it is a simple, fast learning method that predicted shims of similar quality to Direct Design shims. With non-linear feature transforms (which were shown to be the most important feature classes), it was able to accurately relate features to shims and model the variations in shims between slices and subjects. More sophisticated learning methods, such as random forests or neural networks, could also be used within this framework, but would in most cases require more training data and would require significantly more computation than KRR, which may be prohibitive for iterative training.

### 4.5.3 Extensions and Future Work

This study used simulations to validate and characterize the PIPRR algorithm in 100 simulated human heads, using a 24-element transmit coil array at 7T. The next step in development is to implement it in vivo. Training based on numerical simulations of a coil array should be compared to training based on $B_1^+$ map scans of a population of subjects with the same transmit coil. Numerical simulation-based training would be more convenient since the learner could more easily incorporate changes in the coil or its electromagnetic

environment. The method should also be extended to predict shims for slices with arbitrary orientations (coronal, sagittal, and angled), which will require the addition of slice angles as features, as well as for subjects whose heads are not perfectly centered and oriented in the coil.

Another important consideration for human implementation is SAR. As shown in Figure 4.5b, PIPRR predictions had somewhat lower SAR penalty values than Direct Design shims. In this work, SAR was implemented as a regularization term in the training data, and not as an explicit constraint, since the shims were not designed for a specific pulse sequence. While it is possible to implement strict constraints in the training shim design for a specific flip angle and TR, it is not yet clear how to constrain the predictions to also meet these constraints, without dilating RF pulses or increasing TR. One approach might be to predict shims as a weighted average of SAR-constrained shims, where the weights collectively have norm one. This could be achieved using, e.g., a softmax function at the output of a neural network.

Finally, PIPRR could be extended to predict not just RF shims, but full RF parallel transmission waveforms. For example, spokes [103, 104] and $k_T$-points [8] pulses are effectively trains of RF shims, so it should be possible to predict all the shims in the train. It may also be possible to predict parameterized gradient waveforms for such pulses, such as the gradient moments between the subpulses. Designing these pulses requires considerably more computation than RF shimming, especially for large-tip-angles, so the ability to predict them rapidly would have a greater impact on the parallel transmission workflow. More finely sampled waveforms such as spirals could also be predicted, though it is not yet clear how the target pattern could be flexibly resized or positioned in the imaged volume.

## 4.6  Conclusions

This simulation study showed that the RF shim PIPRR method predicts SAR-efficient tailored slice-specific RF shims. Because it does not require full $B_1^+$ mapping, it would

save considerable prescan time compared to directly designed RF shims, and computation of predicted shims requires only feature calculation and a single matrix multiply.

Chapter 5

Contributions and Future Work

Overall, this research has introduced enabling techniques for faster, more homogeneous imaging at high field strengths. Optimization and machine learning techniques that use minimal calibration data were employed to automatically correct for trajectory errors and excite uniform flip angle profiles with parallel imaging and transmission. The specific contributions of this dissertation are:

- TrACR, an automatic trajectory correction method which jointly reconstructs images and non-Cartesian k-space trajectory errors without calibration data. It is the first method to correct non-Cartesian trajectory errors without additional measurements or hardware that is applicable to multiple trajectory types. (Chapter 2)

- An extension of the TrACR method to EPI trajectories, enabling the correction of trajectory and phase errors in accelerated, multishot, and dynamic acquisitions. It is novel in its demonstrated capability to successfully correct dynamic trajectory and phase errors in multishot and/or accelerated EPI acquisitions, without any additional measurements or hardware. (Chapter 3)

- A new machine learning algorithm, PIPRR, that rapidly predicts patient-tailored RF shims for fast, uniform parallel excitations at high field strengths. PIPRR is the first machine learning method to successfully predict patient-tailored RF-shims for parallel transmission, or any other RF pulse design. (Chapter 4)

This chapter will summarize contributions from these bodies of research and conclude by suggesting avenues for building off of the techniques presented herein.

## 5.1 Non-Cartesian Trajectory Correction

In Chapter 2, the TrACR algorithm was introduced and shown to reduce image artifacts from k-space trajectory errors due to gradient eddy currents and delays. TrACR images were comparable to those reconstructed with trajectories measured using a modified Duyn method, for both center-out radial and spiral reconstructions. It was also shown that TrACR is capable of incorporating off-resonance into the reconstruction, as well as reconstructing accelerated data. Importantly, TrACR can make use of SPIRiT, SENSE, or any other parallel imaging reconstruction, and has been validated for golden-angle radial, center-out radial, and spiral trajectories, but is a general model that can be extended to many trajectories. As part of this work, software for implementing TrACR has been provided online, and TrACR has already been implemented to improve reconstructions in several thermometry applications. [105, 106]. Additionally, this work has inspired a new joint image and trajectory reconstruction for wave-CAIPI trajectories [107].

All non-Cartesian acquisitions suffer from trajectory errors, and TrACR could be an enabling factor for some applications. For example, trajectories for compressed sensing require fast readouts with fast switching of gradients. As Feng et al. note [108], an extension of TrACR for compressed sensing could improve those reconstructions. Though non-Cartesian TrACR has been validated only in golden-angle radial, center-out radial and spiral trajectories, the eddy current model used to develop error basis functions is generally applicable; it only requires knowledge of the nominal gradient waveforms. The reconstructions in this work took advantage of the rotational symmetry inherent in the center-out radial and spiral trajectories to simplify the trajectory error models, as well as the empirical knowledge that the x- and y- gradients are similarly calibrated. There are certainly trajectories that do not meet this criteria, and for these, one may need to find another way to reduce the number of parameters to estimate or otherwise artificially increase redundancy in the data. For example, one could optimize jointly over several slices, which should exhibit similar errors and introduce additional data redundancy that allows estimation of more

error parameters. TrACR for golden-angle radial trajectories can faithfully estimate many parameters (1 for each line and dimension) due to its robust oversampling of the center of k-space without changing gradients there, but this is not true of many non-Cartesian trajectories. (PROPELLER trajectories [109] are an exception to this and therefore TrACR should be easily extended for this application).

The main weakness of TrACR is the compute time required for the iterative reconstruction. Compute times for reconstructions in this work ranged from several minutes for the golden-angle radial trajectory, to several hours for center-out radial and spiral trajectories, and are prohibitively long for practical integration with imaging workflows. Future work should therefore focus on methods to accelerate the reconstruction. It may be possible to truncate the data to leave only the center of k-space and reconstruct low-resolution images within TrACR, as was demonstrated with EPI-TrACR in Chapter 3. This approach should be straightforward for golden-angle radial trajectories in particular. For center-out radial and spiral trajectories, it may be necessary to gradually re-introduce the higher frequency data to the reconstruction as it progresses. As was shown in Chapter 2, TrACR exhibits a tendency to estimate trajectory errors for the lower frequencies first, so this seems a feasible approach, and should decrease the time required for reconstruction. Depending on the application, it may also be possible to terminate the algorithm at an earlier time point, since later iterations do not significantly change the image.

If the compute time can be significantly brought down, more possibilities open up for this method; if TrACR reconstructions can be performed online, it should also be possible to introduce a tailored online pre-emphasis, for example in a dynamic series. This might work by performing an initial TrACR reconstruction at the beginning of a dynamic series and feeding the estimated error into the design of gradient waveforms pre-emphasized to compensate for those errors. These waveforms would be played out on the next dynamic, for which an image is also reconstructed with TrACR (and with new error basis waveforms based on an eddy-current model and derived from the new gradient waveforms). If this can

be done quickly enough that gradient coil heating is not a relevant factor, then one may find a stable, eddy-current-optimal regime. In concept, this would be like an acquisition version of gradient iterative predistortion (GrIP) [110].

## 5.2 EPI Trajectory and Phase Correction

Chapter 3 presented an extension of TrACR and demonstrated its capability to correct trajectory delays and line-to-line phase errors in EPI acquisitions. EPI-TrACR outperformed both PAGE and the conventional correction method at acceleration factors up to 4x and up to an EPI shot factor of four, as well as in a dynamic series. Its flexibility in application to multishot and accelerated datasets makes it unique among uncalibrated EPI correction methods, and may allow for more accurate and aggressive parallel acquisitions, particularly in applications like functional imaging or DTI that require many repetitions.

As for non-Cartesian TrACR, future work for EPI-TrACR should focus on a fast implementation, to make it practical for clinical use. Although it is much faster than non-Cartesian TrACR, since it does not use NUFFTs and much of the development focused on speed, EPI-TrACR could be more practical in an online implementation. There are several possible ways to gain additional speed-ups. Firstly, in multislice acquisitions, the same trajectory is used for multiple slices which exhibit similar errors; in initial experiments, EPI-TrACR was able to jointly estimate errors for multiple slices. After initial joint correction, it may take less time to separately fine-tune corrections for each individual slice, since empirically, EPI-TrACR takes less time to converge with better initialization. In a similar vein, it may be practical for some applications to perform an initial calibration to initialize EPI-TrACR, but then use EPI-TrACR to correct for dynamic errors that occur in long time series. Better initialization (and faster performance) can also be achieved by initializing each dynamic with the estimates from the last; in practice, in a small sample set, this resulted in about a 50% decrease in compute time. However further experiments are needed to determine whether initializing reconstructions in this manner is reasonably stable. Ad-

ditionally, it would be interesting to see whether it is possible to gain some improvement in speed by time-windowing a dynamic series to perform joint corrections over several repetitions, and separately fine-tuning them (similar to multi-slice corrections) in a multi-temporal resolution approach. Ultimately, as for non-Cartesian TrACR, pursuing a tailored online pre-emphasis to mitigate EPI trajectory delays would be a worthwhile future direction for this project. For EPI, line-to-line and shot-to-shot phase errors could prove more difficult to compensate prospectively by pre-emphasis, since they are not very stable from one dynamic to the next; likely these would still need to be corrected in post-processing. However, it will take less time to perform phase-only corrections using EPI-TrACR than to perform full delay and phase correction.

While the largest trajectory errors in EPI are delays that occur in the readout dimension, EPI trajectories also exhibit smaller errors in the phase encode dimension, as well as some non-uniform readout-dimension shifts at high frequencies [28, 82]. Future development of EPI-TrACR should extend this method to incorporate these additional types of errors in cases for which it benefits the reconstruction. These errors primarily manifest as eddy-current errors induced during the fast switching of the gradients at these high frequencies, and therefore should be well-captured by an eddy current model similar to that used for TrACR in the center-out radial and spiral error basis generation. The non-Cartesian TrACR reconstructions found difficulty in estimating trajectory errors which occur at high spatial frequencies, and therefore one might think that this formulation would suffer similarly. However, an eddy-current error basis model for EPI would constrain the error estimates such that the estimated delay at the center of k-space is dependent on accurate estimates at high-frequencies; since TrACR is robust for low frequency estimates, this should prevent the algorithm from finding poor solutions. EPI-TrACR could be further improved by incorporating $B_0$ maps into the reconstruction, as was demonstrated with non-Cartesian TrACR, particularly for longer, unaccelerated acquisitions that exhibit significant $B_0$ distortion.

Finally, though the current EPI-TrACR model is restricted to estimating line-to-line

phase shifts, it should be possible to extend this method to estimate higher-order phase maps to correct phase errors from patient motion and respiration, which is a requirement for diffusion imaging. This would require some sort of regularization to require estimated phase maps to be smooth. This may be done with a finite-differencing regularization as in the work of Zhao et al. [86], or by a total-variation (TV)-based regularizer, since TV is already widely used in motion-correction for diffusion imaging with EPI [111]. However, new variability will be introduced by the requirement to tune a regularization parameter, particularly within an iterative method, therefore it may be desirable to estimate these motion phase errors as weights on several smooth functions, thereby implementing a constrained model rather than regularization.

## 5.3    Fast Prediction of RF Shims

Chapter 4 introduced PIPPR and validated its performance in simulated head phantoms to predict tailored RF shims that produce uniform excitation profiles in a SAR-efficient manner. PIPRR brings the compute time required to implement tailored shimming down several orders of magnitude, to just several milliseconds. Furthermore, it all but eliminates the $B_1^+$ mapping required for RF shimming, since it is only necessary to have the central Fourier coefficient from the maps of as few as $1/4$ of the transmit coils. This is another increase in speed of several orders of magnitude over current RF shimming procedures, representing a major advance in the pTx workflow, and could ultimately allow parallel transmission and RF shimming to transition to the clinic.

The most obvious next step for this research is to implement PIPRR in vivo. In vivo implementation will always introduce new sources of variability, and it is important to ensure that training sets incorporate variation in all forms – noise (in $B_1^+$ maps and other measurements), patient position (rotation, translation), anatomical variation, scan volume prescription, etc. In vivo implementation may therefore require additional training data over what was needed in the simulations presented in this dissertation, in order to make

predictions for a more varied dataset. SAR is a more important concern in this domain; particularly since PIPRR implements this as a regularization term, solutions are not guaranteed to meet SAR constraints. However, SAR is already calculated before implementing directly designed RF shims in an in vivo setting; this would also be done before implementing PIPRR predictions, and pulses could simply be dilated if they do not meet constraints. Currently, PIPRR does not attempt to restrict RF power or amplitude, but this could be implemented in a manner analogous to the current SAR regularization.

Another area for future development is in the extension of PIPRR for full RF design. Dynamic shimming with a fixed trajectory should be a straightforward extension of this method, since kT-points [8] or spokes [103, 104] pulses are simply a train of RF shims. One could also imagine predicting tailored parameterized gradient waveforms to go along with these pulses, either jointly or via an alternating optimization. Additionally, while it is a more complicated problem, full RF waveform prediction for more sophisticated pulses might also be possible through an extension of PIPRR. One challenge in implementing these methods, or in transitioning PIPRR to an in vivo method, is that the training process is lengthy. For the data presented in this dissertation, training requires several days compute time in a cluster implementation. Some future work in optimizing the training process to reduce compute time will allow for faster development of the next-generation of PIPRR extensions.

BIBLIOGRAPHY

[1] J A Fessler, S Lee, V T Olafsson, H R Shi, and D C Noll. Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity. *IEEE Trans Sig Proc*, 53(9):3393–3402, 2005.

[2] Q X Yang, J Wang, X Zhang, C M Collins, M B Smith, H Liu, X Zhu, J T Vaughan, K Ugurbil, and W Chen. Analysis of wave behavior in lossy dielectric samples at high field. *Magn Reson Med*, 47(5):982–989, 2002.

[3] W A Grissom. Improving high-field MRI using parallel excitation. *Imaging Med*, 2(6):675–693, 12 2010.

[4] W Mao, M B Smith, and C M Collins. Exploring the limits of RF shimming for high-field MRI of the human head. *Magn Reson Med*, 56(4):918–922, 2006.

[5] J T Vaughan, M Garwood, C M Collins, W Liu, L DelaBarre, G Adriany, P Andersen, H Merkle, R Goebel, M B Smith, and K Ugurbil. 7T vs. 4T: RF power, homogeneity, and signal-to-noise comparison in head images. *Magn Reson Med*, 46(1):24–30, 2001.

[6] J Hua, Q Qin, P C M van Zijl, J J Pekar, and C K Jones. Whole-brain three-dimensional T2-weighted BOLD functional magnetic resonance imaging at 7 Tesla. *Magn Reson Med*, 72, December 2013.

[7] C J Wargo and J C Gore. Localized high-resolution DTI of the human midbrain using single-shot EPI, parallel imaging, and outer-volume suppression at 7T. *Magn Reson Imag*, 31(6):810–819, 2013.

[8] M A Cloos, N Boulant, M Luong, G Ferrand, E Giacomini, D Le Bihan, and

A Amadon. k T-points: Short three-dimensional tailored RF pulses for flip-angle homogenization over an extended volume. *Magn Reson Med*, 67(1):72–80, 2012.

[9] O Kraff, A Fischer, A M Nagel, C Mönninghoff, and M E Ladd. MRI at 7 Tesla and above: Demonstrated and potential capabilities. *J Magn Reson Imaging*, 33:13–33, 2014.

[10] R S Staewen, A J Johnson, B D Ross, T Parrish, H Merkle, and Michael Garwood. 3-D FLASH imaging using a single surface coil and a new adiabatic pulse, BIR-4. *Invest Radiol*, 25(5):559–67, 1990.

[11] K Haines, N B Smith, and A G Webb. New high dielectric constant materials for tailoring the distribution at high magnetic fields. *J Magn Reson*, 203(2):323–327, 2010.

[12] J E M Snaar, W M Teeuwisse, M J Versluis, M A van Buchem, H E Kan, N B Smith, and A G Webb. Improvements in high-field localized MRS of the medial temporal lobe in humans using new deformable high-dielectric materials. *NMR Biomed*, 24:873–879, 2011.

[13] T S Ibrahim, R Lee, B A Baertlein, A M Abduljalil, H Zhu, and P M L Robitaille. Effect of RF coil excitation on field inhomogeneity at ultra high fields: a field optimized TEM resonator. *Magn Reson Imag*, 19(10):1339–47, 2001.

[14] K Nehrke and P Börnert. DREAM–a novel approach for robust, ultrafast, multislice $B_1$ mapping. *Magn Reson Med*, 68(5):1517–26, November 2012.

[15] L I Sacolick, F Wiesinger, I Hancu, and Mika W Vogel. $B_1$ mapping by Bloch-Siegert shift. *Magn Reson Med*, 63(5):1315–22, May 2010.

[16] A Sharma, S Tadanki, M Jankiewicz, and W A Grissom. Highly-accelerated Bloch-

Siegert $|B_1^+|$ mapping using joint autocalibrated parallel image reconstruction. *Magn Reson Med*, 71(4):1470–7, April 2014.

[17] W A Grissom, D Xu, A B Kerr, J A Fessler, and D C Noll. Fast large-tip-angle multidimensional and parallel RF pulse design in MRI. *IEEE Trans Med Imaging*, 28(10):1548–1559, 2009.

[18] S Saekho, C Y Yip, D C Noll, F E Boada, and V A Stenger. Fast-kz three-dimensional tailored radiofrequency pulse for reduced $B_1$ inhomogeneity. *Magn Reson Med*, 55(4):719–724, 2006.

[19] A B Kerr, Y Zhu, and J M Pauly. Phase constraint relaxation in parallel excitation pulse design. In *Proc Intl Soc Magn Reson Med, 15*, page 1694, 2007.

[20] U Katscher, P Vernickel, I Graesslin, and P Börnert. RF shimming using a multi-element transmit system in phantom and in vivo studies. In *Proc Intl Soc Mag Reson Med, 15*, page 1693, 2007.

[21] U.S. Department of Health Administration, Human Services: Food, and Drug. Guidance for industry: Guidance for the submission of premarket notifications for magnetic resonance diagnostic devices, 1998.

[22] D C Noll, J D Cohen, C H Meyer, and W Schneider. Spiral k-space MR imaging of cortical activation. *J Magn Reson Imag*, 5(1):49–56, 1995.

[23] K S Nayak, C H Cunningham, J M Santos, and J M Pauly. Real-time cardiac MRI at 3 Tesla. *Magn Reson Med*, 51(4):655–660, 2004.

[24] A V Barger, W F Block, Y Toropov, T M Grist, and C A Mistretta. Time-resolved contrast-enhanced imaging with isotropic resolution and broad coverage using an undersampled 3D projection trajectory. *Magn Reson Med*, 48:297–305, 2002.

[25] F E Boada, J S Gillen, G X Shen, S Y Chang, and K R Thulborn. Fast three dimensional sodium imaging. *Magn Reson Med*, 37(5):706–715, 1997.

[26] D J Tyler, M D Robson, R M Henkelman, I R Young, and G M Bydder. Magnetic resonance imaging with ultrashort TE (UTE) PULSE sequences: Technical considerations. *J Magn Reson Imag*, 25(2):279–89, March 2007.

[27] J J Van Vaals and A H Bergman. Optimization of eddy-current compensation. *J Magn Reson (1969)*, 90(1):52 – 70, 1990.

[28] L Kasper, S Bollmann, S J Vannesjo, S Gross, M Haeberlin, B E Dietrich, and K P Pruessmann. Monitoring, analysis, and correction of magnetic field fluctuations in echo planar imaging time series. *Magn Reson Med*, 409:396–409, 2014.

[29] A Jesmanowicz, E C Wong, and J S Hyde. Phase correction for EPI using internal reference lines. In *Proc Soc Magn Reson Med, 3*, page 1239, 1993.

[30] X Hu and T H Le. Artifact reduction in EPI with phase-encoded reference scan. *Magn Reson Med*, 36(1):166–171, 1996.

[31] K P Pruessmann, M Weiger, M B Scheidegger, and P Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med*, 42(5):952–962, 1999.

[32] K P Pruessmann, M Weiger, P Börnert, and P Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med*, 46:638–651, 2001.

[33] M Lustig and J M Pauly. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med*, 64(2):457–71, August 2010.

[34] M A Griswold, P M Jakob, R M Heidemann, M Nittka, V Jellus, J Wang, B Kiefer, and A Haase. Generalized autocalibrating partially parallel acquisitions GRAPPA. *Magn Reson Med*, 47(6):1202–1210, 2002.

[35] G F Mason, T Harshbarger, H P Hetherington, Y Zhang, G M Pohost, and D B Twieg. A method to measure arbitrary k-space trajectories for rapid MR imaging. *Magn Reson Med*, 38(3):492–6, September 1997.

[36] J H Duyn, Y Yang, J A Frank, and J W van der Veen. Simple correction method for k-space trajectory deviations in MRI. *J Magn Reson*, 132(1):150–3, May 1998.

[37] Y Zhang, H P Hetherington, E M Stokely, G F Mason, and D B Twieg. A novel k-space trajectory measurement technique. *Magn Reson Med*, 39(6):999–1004, June 1998.

[38] P Gurney, J Pauly, and D G Nishimura. A simple method for measuring $B_0$ eddy currents. *Proc Int Soc Magn Reson Med*, page 866, 2005.

[39] C Barmet, N De Zanche, and K P Pruessmann. Spatiotemporal magnetic field monitoring for MR. *Magn Reson Med*, 60(1):187–97, July 2008.

[40] H Tan and C H Meyer. Estimation of k-space trajectories in spiral MRI. *Magn Reson Med*, 61(6):1396–404, June 2009.

[41] M Takizawa, H Hanada, K Oka, T Takahashi, E Yamamoto, and M Fujii. A robust ultrashort TE (UTE) imaging method with corrected k-space trajectory by using parametric multiple function model of gradient waveform. *IEEE Trans Med Imaging*, 32(2):306–16, February 2013.

[42] I C Atkinson, A Lu, and K R Thulborn. Characterization and correction of system delays and eddy currents for MR imaging with ultrashort echo-time and time-varying gradients. *Magn Reson Med*, 62(2):532–7, August 2009.

[43] N O Addy, H H Wu, and D G Nishimura. Simple method for MR gradient system characterization and k-space trajectory estimation. *Magn Reson Med*, 68(1):120–9, July 2012.

[44] A Deshmane, M Blaimer, F Breuer, P M Jakob, J L Duerk, N Seiberlich, and M A Griswold. Self-calibrated trajectory estimation and signal correction method for robust radial imaging using GRAPPA operator gridding. *Magn Reson Med*, 75(2):883–896, 2016.

[45] D S Smith and E B Welch. Self-calibrated gradient delay correction for golden angle radial MRI. In *Proc Intl Soc Magn Reson Med, 22*, page 229, 2014.

[46] T Wech, J Tran-Gia, T A Bley, and H Köstler. Using self-consistency for an iterative trajectory adjustment (SCITA). *Magn Reson Med*, 73(3):1151–1157, 2015.

[47] L Ying and J Sheng. Joint image reconstruction and sensitivity estimation in SENSE (JSENSE). *Magn Reson Med*, 57(6):1196–1202, 2007.

[48] B P Sutton, D C Noll, and J A Fessler. Dynamic field map estimation using a spiral-in/spiral-out acquisition. *Magn Reson Med*, 51(6):1194–1204, 2004.

[49] W H Press. *Numerical recipes in C*. Cambridge University Press, 2nd ed., v edition, 1993.

[50] L Greengard and J Y Lee. Accelerating the nonuniform fast Fourier transform. *Soc Ind Appl Math Rev*, 46(3):443–454, 2004.

[51] N R Zwart, K O Johnson, and J G Pipe. Efficient sample density estimation by combining gridding and an optimized kernel. *Magn Reson Med*, 67(3):701–10, March 2012.

[52] S Boyd and L Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[53] B Hargreaves. Variable-density spiral design functions, http://mrsrl.stanford.edu/~brian/vdspiral/. Accessed April 6, 2015.

[54] F Huang, S Vijayakumar, Y Li, S Hertel, and G R Duensing. A software channel compression technique for faster reconstruction with many channels. *Magn Reson Imag*, 26(1):133–141, 2008.

[55] M A Bernstein, K F King, and X J Zhou. *Handbook of MRI pulse sequences*. Elsevier Academic Press, 2004.

[56] K P McGee, A Manduca, J P Felmlee, S J Riederer, and R L Ehman. Image metric-based correction (autocorrection) of motion effects: Analysis of image metrics. *J Magn Reson Imag*, 11(2):174–81, February 2000.

[57] M Murphy, M Alley, J Demmel, K Keutzer, S Vasanawala, and M Lustig. Fast $\ell_1$-SPIRiT compressed sensing parallel imaging MRI: Scalable parallel implementation and clinically feasible runtime. *IEEE Trans Med Imaging*, 31(6):1250–1262, June 2012.

[58] X L Wu, J Gai, F Lam, M Fu, J P Haldar, Y Zhuo, Z P Liang, W-M Hwu, and B P Sutton. IMPATIENT MRI: Illinois massively parallel acceleration toolkit for image reconstruction with enhanced throughput in MRI. In *Proc Intl Soc Magn Reson Med, 19*, 2011.

[59] E C Wong. Shim insensitive phase correction for EPI using a two echo reference scan. In *Proc Int Soc Magn Res Med, 2*, page 4514, 1992.

[60] X Wan, G T Gullberg, D L Parker, and G L Zeng. Reduction of geometric and intensity distortions in echo-planar imaging using a multireference scan. *Magn Reson Med*, 37(6):932–942, 1997.

[61] S B Reeder, A Z. Faranesh, E Atalar, and E R McVeigh. A novel object-independent 'balanced' reference scan for echo-planar imaging. *J Magn Reson Imaging*, 9(6):847–852, 1999.

[62] N K Chen and A M Wyrwicz. Removal of EPI Nyquist ghost artifacts with two-dimensional phase correction. *Magn Reson Med*, 51(6):1247–1253, 2004.

[63] D Xu, K F King, Y Zur, and R S Hinks. Robust 2D phase correction for echo planar imaging under a tight field-of-view. *Magn Reson Med*, 64(6):1800–1813, 2010.

[64] Q S Xiang and F Q Ye. Correction for geometric distortion and N/2 ghosting in EPI by phase labeling for additional coordinate encoding (PLACE). *Magn Reson Med*, 57(4):731–741, 2007.

[65] A Campbell-Washburn, H Xue, R Lederman, A Faranesh, and M Hansen. Real-time distortion correction of spiral MRI using the gradient system impulse response function. *Proc Intl Soc Magn Reson Med 19*, 22:1–8, 2014.

[66] B J Wilm, B E Dietrich, J Reber, S J Vannesjo, and K P Pruessmann. Gradient response harvesting for continuous system characterization during MR sequences. In *Proc Intl Soc Magn Reson Med, 24*, page 544, 2016.

[67] M H Buonocore and L Gao. Ghost artifact reduction for echo planar imaging using image phase correction. *Magn Reson Med*, 38(1):89–100, 1997.

[68] M H Buonocore and D C Zhu. Image-based ghost correction for interleaved EPI. *Magn Reson Med*, 45(1):96–108, 2001.

[69] D L Foxall, P R Harvey, and J Huang. Rapid iterative reconstruction for echo planar imaging. *Magn Reson Med*, 42(3):541–7, 1999.

[70] K J Lee, D C Barber, M N Paley, I D Wilkinson, N G Papadakis, and P D Griffiths. Image-based EPI ghost correction using an algorithm based on projection onto convex sets (POCS). *Magn Reson Med*, 47(4):812–817, 2002.

[71] P Kellman and E R McVeigh. Phased array ghost elimination. *NMR Biomed*, 19(3):352–361, 2006.

[72] H Li, K Fox-Neff, B Vaughan, D French, J P Szaflarski, and Y Li. Parallel EPI artifact correction (PEAC) for N/2 ghost suppression in neuroimaging applications. *Magn Reson Imag*, 31(6):1022–1028, 2013.

[73] Yoon Chul Kim, J F Nielsen, and K S Nayak. Automatic correction of Echo-Planar Imaging (EPI) ghosting artifacts in real-time interactive cardiac MRI using sensitivity encoding. *J Magn Reson Imaging*, 27(1):239–245, 2008.

[74] V B Xie, M Lyu, and E X Wu. EPI Nyquist ghost and geometric distortion correction by two-frame phase labeling. *Magn Reson Med*, 77(5):1749–1761, 2017.

[75] W S Hoge, H Tan, and R A Kraft. Robust EPI Nyquist ghost elimination via spatial and temporal encoding. *Magn Reson Med*, 64(6):1781–1791, 2010.

[76] F Hennel, M Buehrer, C von Deuster, A Seuven, and K P Pruessmann. SENSE reconstruction for multiband EPI including slice-dependent N/2 ghost correction. *Magn Reson Med*, 76(3):873–879, 2015.

[77] J Lee, K H Jin, and J C Ye. Reference-free single-pass EPI Nyquist ghost correction using annihilating filter-based low rank Hankel matrix (ALOHA). *Magn Reson Med*, 76(6):1775–1789, 2016.

[78] V B Xie, M Lyu, Y Liu, Y Feng, and E X Wu. Robust EPI Nyquist ghost removal by incorporating phase error correction with sensitivity encoding PEC-SENSE Magn Reson Med 2017. doi: 10.1002/mrm.26710. *Magn Reson Med*, 2017.

[79] J D Ianni and W A Grissom. Trajectory auto-corrected image reconstruction. *Magn Reson Med*, 76(3):757–768, 2016.

[80] J A Fessler and B P Sutton. Nonuniform fast Fourier transforms using min-max interpolation. *IEEE Trans Sig Proc*, 51(2):560–574, 2003.

[81] A K Funai, J A Fessler, D T B Yeo, V T Olafsson, and D C Noll. Regularized field map estimation in MRI. *IEEE Trans Med Imaging*, 27(10):1484–1494, October 2008.

[82] E B Welch and K D Harkins. Robust k-space trajectory mapping with data readout concatentation and automated phase unwrapping reference point identification. In *Proc Intl Soc Magn Reson Med, 26*, page 1387, 2017.

[83] R M Weisskoff. Simple measurement of scanner stability for functional NMR imaging of activation in the brain. *Magn Reson Med*, 36(4):643–645, 1996.

[84] F Zhao, D C Noll, J F Nielsen, and J A Fessler. Separate magnitude and phase regularization via compressed sensing. *IEEE Trans Med Imag*, 31(9):1713–1723, September 2012.

[85] B P Sutton, D C Noll, and J A Fessler. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE Trans Med Imaging*, 22:178–188, 2003.

[86] M Many, M Jacob, D Kelley, and V Magnotta. Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS). *Magn Reson Med*, 78(2):494–507, 2017.

[87] P F Van de Moortele, C Akgun, G Adriany, S Moeller, J Ritter, C M Collins, M B Smith, J T Vaughan, and K Ugurbil. $B_1$ destructive interferences and spatial phase patterns at 7 T with a head transceiver array coil. *Magn Reson Med*, 54(6):1503–1518, 2005.

[88] K Setsompop, L L Wald, V Alagappan, B A Gagoski, and E Adalsteinsson. Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels. *Magn Reson Med*, 59(4):908–915, 2008.

[89] A Hoyos-Idrobo, P Weiss, A Massire, A Amadon, and N Boulant. Variant strategies to solve the magnitude least squares optimization problem in parallel transmission pulse design and under strict SAR and power constraints. *IEEE Trans Med Imaging*, 33(3):739–48, 2014.

[90] V Gras, A Vignaud, A Amadon, D Le Bihan, and N Boulant. Universal pulses: A new concept for calibration-free parallel transmission. *Magn Reson Med*, 2016.

[91] C Mirfin, P Glover, and R Bowtell. Optimisation of parallel transmission radiofrequency pulses using neural networks. *Proc 2nd UK7T Ultra High Field Network Symposium*, April 2017.

[92] J D Ianni, Z Cao, and W A Grissom. RF shim design by supervised learning. In *Proc Intl Soc Magn Reson Med Ultra High Field Workshop*, page 12, 2016.

[93] J D Ianni, Z Cao, and W A Grissom. Predicting multi-coil RF shims via machine learning. In *Proc i2i Workshop, New York University*, page 6, 2016.

[94] J D Ianni, Z Cao, and W A Grissom. SAR-efficient RF shim prediction via machine learning. In *Proc Intl Soc Magn Reson Med, 25*, page 3992, 2017.

[95] P W Kassakian. *Convex Approximation and Optimization with Applications in Magnitude Filter Design and Radiation Pattern Synthesis*. Ph.D. dissertation, University of California, Berkeley, CA, 2006.

[96] X Wu, S Schmitter, E J Auerbach, S Moeller, K Uğurbil, and P F Van de Moortele. Simultaneous multislice multiband parallel radiofrequency excitation with independent slice-specific transmit $B_1$ homogenization. *Magn Reson Med*, 70(3):630–638, 2013.

[97] X Wu, S Schmitter, E J Auerbach, K Uğurbil, and P F Van de Moortele. A gener-

alized slab-wise framework for parallel transmit multiband RF pulse design. *Magn Reson Med*, 75(4):1444—1456, 2016.

[98] A S Lewis and J Malick. Alternating projections on manifolds. *Math Oper Res*, pages 216–234, 2008.

[99] A Christ, W Kainz, E G Hahn, K Honegger, M Zefferer, E Neufeld, W Rascher, R Janka, W Bautz, J Chen, B Kiefer, P Schmitt, H P Hollenbach, J Shen, M Oberle, D Szczerba, A Kam, J W Guag, and N Kuster. The Virtual Family–development of surface-based anatomical models of two adults and two children for dosimetric simulations. *Phys Med Biol*, 55(2):N23–N38, 2010.

[100] J W Young. Head and face anthropometry of adult U.S. civilians. Technical report, US Department of Transportation Federal Aviation Administration, Civil Aeromedical Institute, 1993.

[101] G Eichfelder and M Gebhardt. Local specific absorption rate control for parallel transmission by virtual observation points. *Magn Reson Med*, 66:1468–1476, 2011.

[102] R C Gonzalez and R E Woods. *Digital Image Processing*. Pearson, 2007.

[103] Z Zhang, C Y Yip, W A Grissom, D C Noll, F E Boada, and V A Stenger. Reduction of transmitter B1 inhomogeneity with transmit SENSE slice-select pulses. *Magn Reson Med*, 57(5):842–7, may 2007.

[104] W A Grissom, M-M Khalighi, L I Sacolick, B K Rutt, and M W Vogel. Small-tip-angle spokes pulse design using interleaved greedy and local optimization methods. *Magn Reson Med*, 68(5):1553–62, nov 2012.

[105] T Chen, F Wang, Z Dong, H Qi, S Wang, H Chen, and K Ying. Evaluation of the effect of trajectory correction with radial sampling on temperature imaging. In *Proc Int Soc Magn Res Med, 24*, page 2109, 2016.

[106] S V Jonathan and W A Grissom. Volumetric MRI thermometry using a three-dimensional stack-of-stars echo-planar imaging pulse sequence. *Magn Reson Med*, 2017. doi 10.1002/mrm.26862.

[107] S F Cauley, K Setsompop, B Bilgic, H Bhat, B Gagoski, and L L Wald. Auto-calibrated wave-CAIPI reconstruction; joint optimization of k-space trajectory and parallel imaging reconstruction. *Magn Reson Med*, 78(3):1093–1099, 2017.

[108] L Feng, T Benkert, K T Block, D K Sodickson, R Otazo, and H Chandarana. Compressed sensing for body MRI. *J Magn Reson Imaging*, 45(4):966–987, 2017.

[109] J G Pipe. Motion correction with PROPELLER MRI: Application to head motion and free-breathing cardiac imaging. *Magn Reson Med*, 42(5):963–969, 1999.

[110] W A Grissom, A B Kerr, P P Stang, M Lustig, G C Scott, and J M Pauly. GrIP : Gradient iterative predistortion for multidimensional and parallel excitation. *Proc Intl Soc Magn Reson Med*, 18:4925, 2010.

[111] N K Chen, A Guidon, H C Chang, and A W Song. A robust multi-shot scan strategy for high-resolution diffusion weighted MRI enabled by multiplexed sensitivity-encoding (MUSE). *NeuroImage*, 72:41–47, 2013.