Computational Approaches to Parallel Transmission MRI

by

Christopher Mirfin, MMathPhys.

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

December 2019

Acknowledgements

I would like to thank Professor Richard Bowtell for carefully reviewing this thesis and several preceding ISMRM abstracts, all while leading the department. I am grateful for the funding provided for this project by the EPSRC, MRC and Philips.

Additionally, I am indebted to Dr Shaihan Malik from Kings' College London for many insightful conversations, and for getting me started with parallel transmission in practice.

Thank you to Professor Penny Gowland, Professor Peter Jezzard, and Dr Paul Glover for accepting me onto the ONBI doctoral programme. I am grateful for the friends and colleagues I have met through the programme. In particular, Vaanathi Sundaresan and Hanna Nowicka for teaching me Python. I would also like to thank Professor Mark Jenkinson from Oxford for expert tuition in image processing and providing encouragement.

Thanks to the staff and students within the Sir Peter Mansfield Imaging Centre. Emma Doran for constant support and encouragement. Dr Simon Shah and Dr Matthew Clemence (Philips) for an introduction to pulse programming. Thank you to Tommy Deo and Dr Elena Kleban for the many climbing sessions.

Additionally, I would like to thank Dr Peter Gibbs and Dr Martin Lowry for my introduction to MRI research several years ago at the Hull Royal Infirmary.

Finally, I would like to recognise the support of my family who fostered and nurtured my interest in science.

Contents

List of Figures **5**

List of Tables **13**

List of Algorithms **14**

List of Abbreviations **15**

Abstract **17**

1 Introduction **18**

1.1 History . . . 18

1.1.1 Nuclear Magnetic Resonance . . . 18

1.1.2 Spatial Encoding . . . 19

1.2 Bloch Equation . . . 21

1.2.1 Rotational Frame . . . 21

1.2.2 Spin Domain Bloch Equation . . . 22

1.2.3 Bloch-Riccati Equation . . . 23

1.3 Electromagnetic Fields in Tissue at 7 T . . . 25

1.3.1 Static Field . . . 25

1.3.2 Gradient & Shim Fields . . . 25

1.3.3 B_1^+ Fields . . . 27

1.3.4 Safety considerations . . . 28

1.4 Typical MRI sequences . . . 29

1.4.1 Gradient-Echo Imaging . . . 29

1.4.2 Echo-Planar Imaging . . . 30

1.4.3 Field Mapping . . . 31

1.5 Numerical & Computational Methods . . . 33

1.5.1 Least square problems . . . 33

1.5.2 Interior-Point Algorithm . . . 34

1.5.3 Evolutionary Algorithms . . . 35

1.5.4 Graphical Processing Unit . . . 35

1.6 Summary . . . 36

2 Existing Parallel Transmit Design Methods 38

2.1 RF Shimming . . . 38
2.1.1 Static RF shimming . . . 38
2.1.2 Dynamic RF shimming . . . 40
2.2 Amplitude Modulated Pulse Design . . . 41
2.2.1 Time Optimal *k*-space Trajectories . . . 43
2.2.2 *k*T-Points . . . 45
2.2.3 Frequency Domain methods . . . 46
2.2.4 Large-tip angle . . . 47
2.3 Frequency Modulated Pulse Design . . . 47
2.3.1 Description of FM pulses . . . 48
2.3.2 Adiabatic Pulses . . . 49
2.4 Calibration-Free Parallel Transmission . . . 51
2.5 Summary . . . 51

3 GPU Joint Design of 3D Excitation pTx pulses and k-space trajectories 53

3.1 Introduction . . . 53
3.2 Theory . . . 55
3.2.1 Problem Formulation . . . 55
3.2.2 *k*-space trajectory parameterisations . . . 55
3.3 Methods . . . 56
3.3.1 pTx RF waveform optimisation (*inner-loop*) . . . 57
3.3.2 *k*-space optimisation (*outer-loop*) . . . 57
3.3.3 Initial *k*-space trajectory population . . . 58
3.3.4 GIRF measurement . . . 59
3.3.5 Experiments . . . 59
3.4 Results . . . 60
3.4.1 Computational Time . . . 62
3.4.2 Offline optimisation . . . 62
3.4.3 Online optimisation . . . 63
3.4.4 Measured GIRF . . . 64
3.5 Discussion . . . 64
3.6 Conclusions . . . 68

4 Dynamic Parcellated B_0 shimming for 3D EPI 73

4.1 Introduction . . . 73
4.2 Theory . . . 74
4.2.1 Dynamic B_0 shimming . . . 74

4.2.2 Parcellation schemes . . . 75
4.2.3 Dynamic pTx selective pulses . . . 78
4.3 Methods . . . 80
4.3.1 UK7T Network Dataset . . . 80
4.3.2 Evaluation metrics . . . 80
4.3.3 Experimental Demonstration . . . 81
4.3.4 Modifications to scan control software . . . 82
4.4 Results . . . 82
4.4.1 UK7T Network Dataset . . . 82
4.4.2 Experiments . . . 83
4.5 Discussion . . . 83
4.6 Conclusions . . . 91

5 Sub-Pulse Optimisation for Multidimensional Adiabatic Inversion 92
5.1 Introduction . . . 92
5.2 Theory . . . 94
5.2.1 Sub-pulse design . . . 94
5.3 Methods . . . 98
5.3.1 Numerical Simulation . . . 98
5.3.2 Gradient Design . . . 98
5.3.3 Experimental Validation . . . 99
5.3.4 Modifications to scan control software . . . 101
5.4 Results . . . 101
5.4.1 Numerical Simulation . . . 101
5.4.2 Experimental Validation . . . 102
5.5 Discussion . . . 104
5.6 Conclusions . . . 111

6 Conclusions 112
6.1 Summary . . . 112
6.2 Recent Literature . . . 113
6.3 Future Work . . . 114

A Control Software 115
A.1 System Configuration . . . 115
A.2 Host (Client) Tools . . . 116
A.3 Server Tools . . . 119

List of Figures

1.1 Spatially selective RF pulses together with linear gradients permit reduced Field-of-Excitation (rFOX) of spatial magnetisation (shown here as a solid orange colour). The most commonly used example of this is 1D rFOX or *slice-selection*. For simplicity 2D and 3D rFOX are shown here as cuboidal, but in general the volume of spatial excitation can be arbitrary in shape. 20

1.2 Stereographic projection into the transverse plane (red line) of the bulk magnetisation vector (black line) on the Bloch sphere in the rotating frame. The projection vector intersects the points $(0, 0, -1)$ and the instantaneous magnetisation vector on the Bloch sphere. 24

1.3 Spherical harmonics of order and degree $\{n, m\}$, visualised on the surface of a unit cylinder aligned on axis with the direction of the static-field. . . 26

1.4 Relative transmit sensitivities (magnitude) in the head at 7 T, for central axial, coronal and sagittal slices (rows), for 8 transmit channels. As expected, the highest sensitivity values are adjacent to each coil element. 27

1.5 Gradient-echo pulse timing sequence (a) and the effective k-space traversed throughout (b). After each RF pulse with slice-selective gradient G_z, a prephasing gradient and phase-encoding step gradient G_y are used to set the initial point of each line in k-space. A subsequent read-out gradient G_x is used to traverse a whole line of k-space while the signal is discretely sampled. 30

1.6 EPI pulse timing sequence (a) and the effective k-space traversed throughout (b). After each RF pulse with slice-selective gradient G_z, prephasing gradients are used to set the initial point in k-space. All of k-space is traversed after a single RF pulse, with a trapezoidal readout gradient G_x and phase-encoding gradient G_y. Sampling data along the readout direction is only limited by the discretisation process of signal sampling, whereas sampling along the phase-encoding direction is limited by the T_2^* of the tissue. 31

1.7 Figure adapted from Stockmann & Wald (*32*). EPI distortion at 7 T for 5 axial slices. First column shows the off-resonance map after global shimming up to 2nd order has been applied. The second column shows a reference gradient echo acquisition. The EPI geometric distortion is visible (blue arrows) in the third and fourth columns with the phase-encoding first in the A-P direction, and then again in the reversed P-A direction (2 mm isotropic resolution, parallel imaging factor = 2). Signal dropout is highlighted with green arrows. 32

1.8 Schematic diagram of CPU-GPU layout. Required data for optimisations and simulations (e.g. field maps) can remain in GPU global memory for the duration (or beyond) of multi-step calculations. In this way, multiple transfers across the relatively slow PCIe connection can be avoided. . . . 36

2.1 Static RF Shimming for circular ROI. Polar plots (top row) display complex shim values for 8 channels each denoted by a different colour, and the resulting predicted B_1^+ amplitude (middle row) and phase map (bottom row). Each RF shim setting is power matched to the quadrature solution. The first column represents default quadrature, second column shows result of relative phase nulling, last column shows the maximum efficiency solution. 39

2.2 Relative standard deviation of $|B_1^+|$-field, normalised by values for quadrature volume mode for dynamic slice-by-slice RF shimming across the human brain when using Nova 8-channel transmit coil. Error bars demonstrate variation between 6 volunteers. Uniform power matched RF shimming refers to optimised RF fields, with matched power to quadrature mode. For comparison, uniform RF shimming with no power regularisation is also shown. Adapted from (*101*). 40

2.3 Acceleration with 2x undersampled echo-planar trajectory. Aliasing (phase-encoding is left-to-right here) is observed in each of the excitation patterns for each coil element. Combination of both coils is sufficient to remove aliasing artefact. 43

2.4 Example 2D spiral k-space sampling pattern, resulting pulse duration of 10.3 ms. Example gradient echo magnitude images (axial slice) for a spherical phantom with a radius of 10 cm. RF pulses designed using spatial domain method (*108*) scaled for 90° flip angle (see Equation 2.5). . 43

2.5 kT-Points trajectory design: disordered sparse k-space locations (a); minimum pointwise distance found using genetic algorithm (b); time optimal trajectory using the method described in (*122*, *123*), (c). 46

2.6 Two useful rotating frames: primed (the effective frame), and double primed. The primed frame rotates at the instantaneous RF frequency ω_{RF}, and exhibits an effective magnetic field $\omega_{\mathrm{eff}}(t)$ that is a combination of the transverse component $\omega_1(t)$ and a longitudinal component $\Delta\omega(t)$. The double primed frame rotates at an instantaneous frequency defined by the magnitude of the effective field in the primed frame. For a magnetisation vector initially parallel with the z' axis to remain aligned with the effective field it is required that $|\dot{\psi}(t)| \ll |\omega_{\mathrm{eff}}(t)|$. Adapted from (*151*) 49

2.7 The hyperbolic secant pulse for three values (different colours) of ω_{max} plotted on the phase-plane (a): the pulse starts at the origin and spirals out to ω_{max}, before reversing along the same path to the origin. Adiabatic Full-Passage pulses' insensitivity to $|B_1^+|$-value is explained when viewed in a sweep diagram (b) (*158*). 50

2.8 Neural network training using supervised learning. The training data consists of a library of patient field maps and a reference *gold standard* method to design RF pulses in each case. Backpropagation is used to minimise the difference in the error between the network and the optimal solution. The trained network can then be deployed for new input field maps as a simple feed-forward calculation. Adapted from (*176*) . . . 52

3.1 Trajectory parameterisations with their first and second derivatives. 2nd-order B-splines with uniform spacing (adapted from (*185*)) (a); 2nd-order B-splines with non-uniform spacing (b); Discrete Sines (c) for $N_p = 5$ parameters. The colours here are used to differentiate each individual localised spline curve. 57

3.2 Examples of six initial *classic* 3D k-space trajectories, each with duration of 6.6 ms and subject to a maximum slew rate of 180 $\mathrm{Tm^{-1}s^{-1}}$. 59

3.3 Measures of parameterisation error as a function of a number of parameter values per axis for each of the three parameterisations considered: B-splines, (Log) B-splines, and Discrete Sines. 61

3.4 Execution time for inner-loop and outer-loop (interior-point) optimisations, as a function of problem size using CPU and GPU implementations, when using 64 parameters per axis. 62

3.5 Cost curves for objective function stated in Equation 3.1 for GPU accelerated Interior-Point Algorithm (a) and genetic algorithm (b). 63

3.6 Excitation error (NRMSE), total power, and peak power; before and after optimisation using Interior-Point algorithm for 2nd-order B-spline parameterisation. Results show mean average ± standard deviation across five subjects. Comparative results for the genetic algorithm are also shown. 65

3.7 Bloch simulations of 3D excitation profiles (10° flip angle) shown for central sagittal, coronal and axial slices, before and after optimisation using interior-point method with *stack-of-spirals* initialisation. Line profiles across slice through centre of region-of-excitation shown above each image. 66
3.8 Cross-comparison of excitation error (NRMSE) between optimised k-space trajectories across six initialisations and five subjects. Each of these trajectories was re-optimised on a per subject basis before being applied across each subject. In each case, RF waveforms were re-optimised for each individual subject for a given trajectory. 67
3.9 Cross-comparison of Fréchet distance between optimised k-space trajectories across six initialisations and five subjects. 68
3.10 Example full field of view images for 3D GRE. Sagittal, coronal and axial slices shown for regular excitation (a); reduced Field-of-Excitation (b); reduced Field-of-Excitation windowed to 20% of maximum intensity (c). Field maps from the same phantom are also shown (d). 69
3.11 Example full field of view images for 3D GRE. Sagittal, coronal and axial slices shown for regular excitation (a); inner-volume excitation (b); inner-volume excitation windowed to 20% of maximum intensity (c). . . 70
3.12 For each subject, optimised excitation k-space trajectories per axis (left), and magnitude plot of associated multichannel RF waveforms (right). . . 71
3.13 Measured GIRF function (*Philips 7 T Achieva, Nottingham*), magnitude and phase displayed in the range -30 to 30 kHz. 71
3.14 Example of gradient waveforms and corresponding excitation k-space trajectory, before and after the application of measured GIRF. The low-pass nature of the system response leads to an undershoot the peak gradient amplitude. 72

4.1 Visualisation of the parcellation methods discussed in Section 4.2.3, for both arbitrary 3D sub-volumes, 3D prisms and 3D half-prisms with a total of 16 sub-volumes each. 79
4.2 Schematic diagram of workflow for parcellated dynamic shimming design. Field maps (both ΔB_0 and B_1^+) are acquired and pre-processed (a), a suitable volume parcellation is devised (b), for each sub-volume, shim currents are optimised and a predicted off-resonance field is generated (c), for each sub-volume, rFOX pulses are designed with updated off-resonance field (d). 80

4.3 Shim efficacy for 10 parcellation methods and 32 sub-volumes. Mean and standard deviation across 10 subjects shown. For both rigid and non-rigid methods, only the largest shim efficacy parcellations are displayed. 84
4.4 Shim efficacy as a function of circularity (left) for 3D prism parcellations, and (right) sphericity for arbitrary 3D parcellations. Mean values averaged across 10 subjects are shown for 32 sub-volumes. Regression lines are plotted for both 1st and 2nd order shimming. 85
4.5 Visualisation of 3D arbitrary parcellations (left column), shimmed off-resonance field in Hz (centre column), associated mean-centred histograms (right column). 86
4.6 3D EPI for 6 parcellated sub-volumes (central axial slices shown) with 3mm isotropic resolution (top row = relative magnitude, bottom row = phase with overlaid shim sub-volumes displayed). Also shown is zoomed sub-volume (1.5 mm^2 in-plane resolution) for sub-volume 4. 87
4.7 Example central axial slices from five separate 3D EPI acquisition, each with four dynamics, for five distinct parcellations. 88
4.8 Example central axial slices from subject 5. 3D EPI acquisition with 3mm isotropic resolution each with four dynamics. First row (from top) - default global shim acquisition; second row - dynamic relative magnitude image; third row - dynamic relative magnitude image thresholded at 15% intensity level; forth row - dynamic phase image with overlaid red box to highlight shim box; fifth row - overlapped magnitude image. Yellow arrows in second and third rows highlight excitation error in outer volume region. 89
4.9 Simulated rSENSE g-factor maps (central axial slice) for three parcellations with five different SENSE acceleration factors using a 32-channel receive head coil. Whereas the g-factor becomes unacceptably large for high acceleration factors for a large sub-volume size, more compact parcellations can maintain low g-factors even for high acceleration factors. 90

5.1 Schematic diagram to display the geometry of the pulse design problem. The top row shows the Blipped HS pulse with three different sub-pulses (hard-pulse, 1D spatially-selective, and 2D spatially-selective). Adding gradient blips in the frequency sweep direction generates selectivity in a third direction (bottom row). Due to the discrete nature of the waveform, replicants will be equally spaced along this dimension. 93
5.2 Blipped Hyperbolic Secant (HS) (solid line) and continuous variant (dashed line). Pulse amplitude $|RF|$, phase ϕ_{RF}, and slice-selective gradient G_z. for 15 segments. 94

5.3 Ideal functional form of transition function $h(x)$ and series approximation $h_n(x)$ for $n = \{1, \cdots, 9\}$. 97
5.4 Inversion Recovery Sequence with 2D selective prepulse inversion. High slew-rate spiral gradients split into 15 sub-pulses are used to permit 2D selectivity in the plane orthogonal (e.g. $x - y$) to the frequency-sweep direction (z). If 3D selectivity is desired blipped, z-gradients can be applied. A crusher is used to ensure destruction of remnant transverse magnetisation, before a user-defined delay time preceding the readout module. Inset figure shows one sub-pulse (multichannel amplitudes) and 2D waveforms. 100
5.5 TRMSE over entire volume-of-interest using 2D slice optimisation (a) or full volume optimisation (b). Range of regularisation parameter used: $\lambda \in [10^2, 10^5]$. 102
5.6 Inversion NRMSE over entire volume-of-interest using 2D slice optimisation (a) or full volume optimisation (b). Range of regularisation parameter used: $\lambda \in [10^2, 10^5]$. 103
5.7 Convergence of the (Threshold) Variable Exchange method for MLS and TMLS objective functions. Representative of 3D optimisation with $5{\times}10^4$ spatial voxels. As the objective functions are non-convex, only a local minima can be assumed. 103
5.8 Gradient-echo central axial slice to show STA for 10° flip angle target excitation profile for individual sub-pulse (pulse duration = 1.0 ms (+prephasing), $k_{\max}$ = 80 rad m^{-1}, λ = 1.0). Relative magnitude profile comparison also shown bottom right. 105
5.9 Magnitude images from TFE inversion recovery with different relative peak prepulse inversion powers $\omega_{\max}$ (up to a maximum 15 μT). Central axial slice shown through phantom, using inversion delay of 400 ms. Yellow arrow indicates dark transition region between selected and unselected inversion regions. 106
5.10 Relative difference from signed TFE images with and without prepulse. Prepulse inversion peak powers are denoted $\omega_{\max}$, (up to a maximum 15 μT). Central axial slice shown through phantom using inversion delay of 200 ms. A positive value indicates inverted signal. It is noted that the difference method also leads to spurious edge enhancement in regions where no inversion occurred. 107
5.11 Magnitude (a) and phase (b) variation through axial slices of 3D TFE readout (FOV = 192 × 192 × 26 mm with 1 × 1 × 3 mm resolution, TR = 30 ms, TE = 2.0 ms) with optimised 2D selective adiabatic inversion. Inversion delay time 400 ms. 108

5.12 Inversion recovery curve from single-slice TFE readout with selective adiabatic inversion prepulse in vivo. Plotted for two regions of interest (inner and outer regions), displayed on corresponding magnitude images. Error bars reflect standard deviation in signal across ROIs. 109
5.13 Inversion recovery curve from single-slice TFE readout with selective adiabatic inversion prepulse, for spherical phantom of water. Plotted for two regions of interest (inner and outer regions), displayed on corresponding magnitude images. Error bars reflect standard deviation in signal across ROIs. 110

A.1 Loading neurological data from PAR REC file. Brain extraction (*219*) (shown in the pop-up central window) is automatically performed on first gradient echo image from Dual-TE image set (bottom-left panel). For non-neurological imaging, brain extraction can be replaced with simple background masking. 117
A.2 Loading and preprocessing multichannel relative transmit sensitivity maps (central window panel). By default, single-channel encoding is assumed, but other channel encoding are permissible. Default median filtering is applied with kernel size [3, 3, 3]. The processed maps are displayed in the upper-left panel within the main interface (the view between magnitude and phase can be toggled by right-clicking). 118
A.3 ROI selection of varying size and shape are defined in the transverse plane for a given number of slices. Additional ROIs can be selected for multiple regions if required. First a *Draw* method is selected (left panel), then the user creates the mask (middle panel), and finally non-brain tissue is excluded using the predetermined BET mask (right panel). 119
A.4 RF shimming routines are selected from the *Design -> RF shimming* tab. The resulting channel weights (amplitude or relative phase) are displayed as a bar chart in the upper-middle window panel. The predicted transmit profile map (magnitude or phase) is shown in the lower-middle window panel. In this case a relative phase nulling algorithm was used. 120
A.5 Static ΔB_0 shimming (up to second-order) is selected from the *Design -> Static B0 shimming* tab. A histogram displaying the off-resonance frequencies within the selected ROI is shown in the upper-right window panel. 121
A.6 RF pulse optimisation routines (e.g 2D spiral) are selected from the *Design -> RF Pulses* tab. A separate panel for input pulse properties is used to define further pulse specifications. 122

A.7 After pulse optimisation completes on the remote server, the status switches from ‘busy’ to ’all done’. Resulting RF pTx waveforms are shown in the upper-middle window panel. A Bloch simulation of the transverse (or longitudinal) magnetisation is shown in the lower-middle window panel. Pulses are written to file to be loaded into the vendor spectrometer software. 123
A.8 Application of Control Software for RF shimming in MR Spectroscopy. . 124
A.9 Application of Control Software for RF shimming in Renal Angiography. 125

List of Tables

1.1 Limits for global and local SAR averaged over 6 minute time window. Adapted from International Electrotechnical Commission 2015 (*48*). . . . 28

3.1 Excitation error (NRMSE), total power, and peak power; before and after optimisation using Interior-Point algorithm with each of the three parameterisations for each initial trajectory. Results show mean average ± standard deviation across five subjects. Comparative results for the genetic algorithm are also shown. 64

List of Algorithms

1 Conjugate Gradient Least Squares 34
2 Genetic Algorithm 35
3 Orthogonal Matching Pursuit 45
4 Non-rigid parcellation 76
5 Split & Merge parcellation 77
6 Adaptive Voronoi parcellation 78
7 Threshold Variable Exchange method 96

List of Abbreviations

2D two-dimensional.

3D three-dimensional.

AFI Actual Flip-Angle Imaging.

AM Amplitude Modulated.

CGLS Conjugate Gradient Least Squares.

EPI Echo-Planar Imaging.

EVI Echo-Volumar Imaging.

FM Frequency Modulated.

FOV Field of View.

FOX Field of Excitation.

GA Genetic Algorithm.

GIRF Gradient Impulse Response Function.

GPU Graphical Processing Unit.

GRE Gradient-Recalled Echo.

HS Hyperbolic Secant.

IVI Inner-Volume Imaging.

LLS Linear Least Squares.

LSQR Least Squares Quadratic Reduction.

LTI Linear Time-Invariant.

MLS Magnitude Least Squares.

MRI Magnetic Resonance Imaging.

MSLS Magnitude Squared Least Squares.

NMR Nuclear Magnetic Resonance.

NRMSE Normalised Root Mean Squared Error.

pTx Parallel Transmit.

RF Radio Frequency.

rFOV reduced Field-of-View.

rFOX reduced Field-of-Excitation.

SAR Specific Absorption Rate.

STA Small-Tip Angle.

TFE Turbo Field Echo.

TMLS Threshold Magnitude Least Squares.

TMSLS Threshold Magnitude Squared Least Squares.

TRMSE Threshold Root Mean Squared Error.

Computational Approaches to Parallel Transmission MRI

Christopher Mirfin

Abstract

This thesis is concerned with the design of spatially-selective RF pulses accelerated by parallel transmission. These pulses perturb the magnetisation in an object within a 2D or 3D sub-volume, whilst leaving magnetisation outside of these regions unaffected. As such, they are shown to be useful in certain dynamic field-shimming strategies.

The first chapter reviews the equation of motion for a magnetisation vector and alternative compact formulations; electromagnetic fields in tissue at 7 T; requisite imaging sequences; and some numerical and computational methods.

The second chapter reviews some existing design strategies and capabilities of Parallel Transmit. It also includes a derivation of the adiabatic condition, which is critical for the design of Frequency Modulated pulses.

In Chapter 3 the design of short Parallel Transmit pulses with optimised k-space trajectories to excite an arbitrary 3D volume is presented. As part of this, the parameterisation of the excitation k-space trajectory is considered under alternative linear transformations subject to realistic gradient performance. A Graphical Processing Unit (GPU) is used to reduce computational time of an existing *local* and a novel second *global* optimisation algorithm.

In Chapter 4 a novel acquisition strategy is developed and implemented to improve the overall static field homogeneity by utilising reduced field of excitation with acceleration via parallel transmission in conjunction with dynamic first-order local shimming.

In the preceding chapters the use and design of spatially-selective pTx RF pulses was strictly limited to Amplitude Modulated (AM) waveforms. In Chapter 5 a spatially-selective pTx Frequency Modulated (FM) pulse that is resilient to a level of transmit-field inhomogeneity is implemented. An optimisation scheme is sought that exploits the adiabatic condition to relax the spatial transmit-field magnitude and phase patterns.

Chapter 6 discusses the work presented in this thesis, and highlights possible directions for future projects.

An Appendix is appended to this thesis, highlighting a software application that was developed to ameliorate some of the experimental workflow challenges associated with pTx.

Chapter 1

Introduction

1.1 History

1.1.1 Nuclear Magnetic Resonance

The phenomenon of Nuclear Magnetic Resonance (NMR) arises because certain nuclei with non-zero spin, placed within a *strong* magnetic field absorb electromagnetic energy at a specific frequency. The foundations of NMR originated with the discoveries of Rabi et al (*1*) for molecular and atomic beams from 1938, followed 8 years later with the first measurement of the magnetic moment of the proton by Bloch et al (*2*) and Purcell et al (*3*). Although the property of spin for nucleons is derived from their constituent quarks and gluons, *color-confinement* forbids the isolation of these sub-nuclear particles.

Feynman et al (*4*) demonstrated the equivalence of a two-state quantum system (i.e. a single spin-$\frac{1}{2}$ nuclei) and the classical precession of a magnetic moment in a magnetic field, permitting the use of geometry to understand the behaviour of the model. For an ensemble of nuclei in the presence of an external *static* magnetic field, a thermal equilibrium will be reached whereby the state populations are governed by Boltzmann statistics. In the *high-temperature* approximation, the net magnetisation will be proportional to the difference in energy of the two states and hence is proportional to the static field strength. At room temperature, this energy difference is several orders of magnitude smaller than the available thermal energy and the limited size of the resulting population difference drives the adoption of higher static field systems.

From an understanding of a classical magnetic moment in a magnetic field, the magnetisation vector will precess about the static field B_0 at a rate known as the Larmor frequency,

$$\omega_0 = \gamma B_0 \tag{1.1}$$

where the gyromagnetic ratio γ is specific to the nuclei of interest. Application of a finite-duration, *on-resonance* oscillating magnetic field perpendicular to the main field will perturb the net magnetisation vector off-axis. A coil placed close to the sample will enclose an area of changing flux, and hence a signal will be induced in accordance with

Faraday's law. The rate at which this signal decays gives an insight into the thermodynamic properties of the system.

Magnetic Resonance Spectroscopy makes use of broadband RF pulses to *excite* signal from nuclei across a range of frequencies (a range which is inversely proportional to pulse duration). Fourier analysis of the signal can identify contributions at each frequency in the spectrum from different nuclei in atoms in differing chemical environments (*5*).

1.1.2 Spatial Encoding

The pivotal step to move from NMR Spectroscopy to Magnetic Resonance Imaging (MRI) was made in the 1970s with the understanding that a 1D spectrum of nuclei in multiple chemical environments, could be related to the spectrum from a single-species of nuclei with their spatial distribution *encoded* through use of a linear gradient. In 1973, Lauterbur reconstructed the first MR image from multiple projections (*6*), followed in 1977 when Mansfield and Maudsley demonstrated the first human in vivo image utilising a Fourier reconstruction (*7*).

In an ideal system, the signal emanating from a voxel at location $\mathbf{r}$ with spin-density ρ can be modelled as:

$$S(\mathbf{r}, \mathbf{k}_\mathbf{R}(t)) \propto \int_V \rho(\mathbf{r}) e^{i\mathbf{k}_\mathbf{R}(t)\cdot\mathbf{r}}\, dV \tag{1.2}$$

where the spatial-frequency or *k-space* vector (rad m^{-1}) is defined as:

$$\mathbf{k}_\mathbf{R}(t) = \gamma \int_0^t \mathbf{G}(\tau) d\tau \tag{1.3}$$

Assuming the main static field is orientated along the z-direction, three additional orthogonal linear *gradient* fields $\mathbf{G}(t) = [\frac{\partial B_z}{\partial x}, \frac{\partial B_z}{\partial y}, \frac{\partial B_z}{\partial z}]$, contribute to the net magnetic field $B_0 + \mathbf{G}\cdot\mathbf{r}$, at each spatial position. The three gradient fields are generated independently from dedicated coil-windings situated within the bore of the scanner. It should be noted that whenever a gradient coil-winding is driven, additional *concomitant* fields with a nonlinear spatial dependence are also simultaneously generated, as a consequence of Maxwell's equations. Although the combined magnitude of these additional fields is much smaller than the magnitude of ideal gradient fields, they distort the linearity assumption between applied gradient fields and MR signal phase present in Equation 1.2.

Discretely sampling k-space by sampling the signal acquired while modulating the gradient waveforms $\mathbf{G}(t)$ permits spatial signal localisation by Fourier transformation of Equation 1.2. As such, the sampling step-size Δk is inversely proportional to the Field of View (FOV) in the spatial domain. Signal aliasing will occur when Δk defines a FOV smaller than the the extent of the object from which signal is excited in a given direction. Parallel imaging, whereby signal intensity is simultaneously detected by mul-

tiple independent receiver coils can be used to provide additional spatial information in undersampled data (*8–10*). Undersampling is particularly desirable in MRI as this often directly reduces imaging time. Additional and complimentary techniques to parallel imaging exist in *sparse* signal sampling (*11*).

The maximum extent in k-space is inversely proportional to the voxel size in the spatial domain. Consequently, to exploit intrinsic signal-to-noise ratio gains at ultra-high field, gradient systems capable of higher peak amplitudes are desirable to fully realise increases in resolution (*12*, *13*).

Linear gradient fields can also be used to move some of the spatial encoding from reception to transmission using spatially-selective RF pulses (see Figure 1.1). The most common example of this is 1D slice-selection (*14*). Its use generally allows the reduction of the dimensionality of the read-out step to 2D. Concatenating three 1D-selective RF pulses with orthogonal gradients can be used to select a cuboidal volume (*15*, *16*).

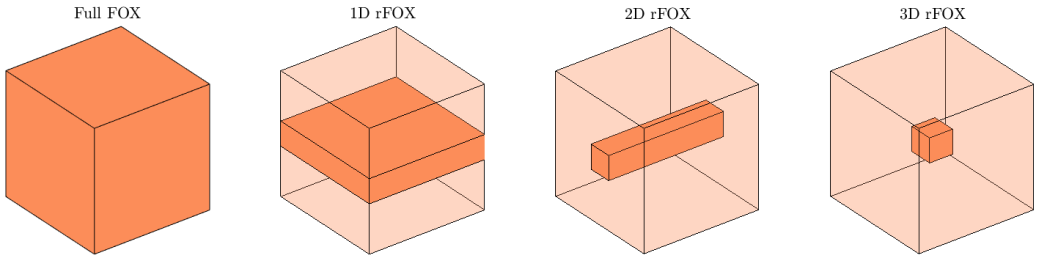

Figure 1.1: Spatially selective RF pulses together with linear gradients permit reduced Field-of-Excitation (rFOX) of spatial magnetisation (shown here as a solid orange colour). The most commonly used example of this is 1D rFOX or *slice-selection*. For simplicity 2D and 3D rFOX are shown here as cuboidal, but in general the volume of spatial excitation can be arbitrary in shape.

Pauly et al (*17*), developed a method of also analysing selective excitation in terms of spatial frequencies or *transmit k-space*. As will be discussed in Chapter 2, the notion of sampling transmit k-space by modulating gradient waveforms is extremely powerful, allowing the design of arbitrary-shape multidimensional selective RF pulses in principle.

The increasing utilisation of ultra-high field (7 T or higher) systems has been driven by gains in the signal-to-noise and contrast-to-noise ratios at the detriment of non-uniform static and transmit fields, and increased power deposition (*18*). Parallel Transmit (pTx) using an array of independently controlled RF transmit coils was developed to allow increased tailoring of the transmit field (*19*, *20*). Motivated by parallel imaging, Katscher et al (*21*) were the first to utilise the additional spatial information provided by multiple transmit coils to undersample transmit k-space, allowing a reduction in pulse duration.

1.2 Bloch Equation

In the classical limit, for non-interacting spins, the bulk magnetisation $\mathbf{m}$ evolves according to the equation of motion:

$$\frac{d\mathbf{m}}{dt} = \gamma \mathbf{m} \times \mathbf{B}\,, \tag{1.4}$$

where γ is the gyromagnetic ratio. Thermodynamic processes can be introduced through two tissue-specific and field-dependent relaxation times T_1 and T_2, yielding the eponymous *Bloch equation* for interacting spins (*22*).

$$\frac{d\mathbf{m}}{dt} = \gamma \mathbf{m} \times \mathbf{B} - \begin{pmatrix} m_x/T_2 \\ m_y/T_2 \\ (m_z - m_0)/T_1) \end{pmatrix} \tag{1.5}$$

T_1-relaxation models the regrowth of longitudinal magnetisation towards thermal equilibrium, due to stochastic processes. T_2-relaxation models the irreversible dephasing of magnetisation in the transverse plane, also due to stochastic processes. In practice, the effective transverse relaxation parameter T_2^* is shorter than T_2 due to local field inhomogenities. As these inhomogenities are time-independent, for nuclei that are not moving, this additional dephasing can be refocused after application of a π-rotation pulse.

In general, the time-scale over which RF pulses typically act is much shorter than T_1, T_2 and can thus be neglected for pulse design.

1.2.1 Rotational Frame

Rabi et al (*23*) demonstrated the utility of a rotating frame of reference for analysing NMR problems. From (*24*), consider a magnetisation vector in an inertial frame (unprimed) and a rotating frame of reference (primed) defined by the vector $\mathbf{\Omega}$:

$$\begin{aligned} \mathbf{m}(t) &= m_x(t)\hat{\mathbf{x}} + m_y(t)\hat{\mathbf{y}} + m_z(t)\hat{\mathbf{z}} \\ \mathbf{m}'(t) &= m_{x'}(t)\hat{\mathbf{x}}'(t) + m'_y(t)\hat{\mathbf{y}}'(t) + m_{z'}(t)\hat{\mathbf{z}}'(t) \end{aligned} \tag{1.6}$$

and differentiating each with respect to time and equating gives:

$$\frac{d\mathbf{m}}{dt} = \underbrace{\frac{dm_{x'}}{dt}\hat{\mathbf{x}}' + \frac{dm_{y'}}{dt}\hat{\mathbf{y}}' + \frac{dm_{z'}}{dt}\hat{\mathbf{z}}'}_{=:(d\mathbf{m}/dt)'} + \underbrace{m_{x'}\frac{d\hat{\mathbf{x}}'}{dt} + m_{y'}\frac{d\hat{\mathbf{y}}'}{dt} + m_{z'}\frac{d\hat{\mathbf{z}}'}{dt}}_{\mathbf{\Omega}\times\mathbf{m}}\,. \tag{1.7}$$

The first part of this equation represents the rate of change of the magnetisation with respect to the rotating frame. Comparison with the Bloch Equation 1.4, leads to the

identification:

$$\left(\frac{d\mathbf{m}}{dt}\right)' = \gamma \mathbf{m} \times \mathbf{B}_{\text{eff}} \,, \tag{1.8}$$

where the effective magnetic field is given by:

$$\mathbf{B}_{\text{eff}} = \mathbf{B} + \frac{\boldsymbol{\Omega}}{\gamma} \,. \tag{1.9}$$

Unless otherwise specified it will be assumed $\boldsymbol{\Omega} = -\gamma B_0 \hat{\mathbf{z}}$, to remove the rotational effect of the static field.

1.2.2 Spin Domain Bloch Equation

The non-interacting Bloch Equation (1.4) constrains the magnetisation vector to a unit-sphere centred on the origin. The instantaneous angular velocity of the magnetisation vector $\omega = (\omega_x, \omega_y, \omega_z)$ is simply dependent on any off-resonance terms and RF fields. Like the real-valued Euler angles, the complex-valued Cayley-Klein parameters (α, β) can uniquely describe any rotation in 3D space. Following the notation of Pauly et al (*25*), the unitary rotation matrix Q_j is:

$$Q_j = \begin{pmatrix} \alpha_j & -\beta_j^* \\ \beta_j & \alpha_j^* \end{pmatrix}. \tag{1.10}$$

For a rotation of ϕ_j about the axis defined by $\hat{n}_j = (n_x, n_y, n_z)_j$, the Cayley-Klein parameters for the jth time step are:

$$\begin{aligned} \alpha_j &= \cos(\phi_j/2) - i n_{z,j} \sin(\phi_j/2) \\ \beta_j &= -i(n_{x,j} + i n_{y,j}) \sin(\phi_j/2) \,. \end{aligned} \tag{1.11}$$

Discretising the angular velocity $\omega(t)$ using the sample time Δt, the cumulative effect of N rotational steps is given by:

$$Q_N = \prod_j^{N-1} Q_j \,. \tag{1.12}$$

The final transverse magnetisation $m_{xy} = m_x + i m_y$, (assuming thermal equilibrium initially) can be calculated in $\mathbb{R}^3$ through the relation $m_{xy} = 2\alpha_N \beta_N m_0$.

The *hard-pulse* approximation decomposes the rotation defined by each Q_j matrix into sequential RF induced rotation effects and gradient-free precession as:

$$Q_j \approx \begin{pmatrix} c_j & -s_j^* \\ s_j & c_j^* \end{pmatrix} \begin{pmatrix} z_j & 0 \\ 0 & z_j^* \end{pmatrix} , \tag{1.13}$$

where,

$$\begin{aligned} c_j &= \cos(|\omega_{1,j}|\Delta t/2) \\ s_j &= ie^{i\varphi_j}\sin(|\omega_{1,j}|\Delta t/2) \\ z_j &= e^{i\omega_{z,j}\Delta t/2}\,. \end{aligned} \tag{1.14}$$

The approximation is increasingly accurate for decreasing step-size Δt. This compact form is used as the basis of an efficient implementation of the Bloch Equation, for simulation of the magnetisation response to given RF and gradient waveforms. As spatial locations ($\mathbf{r}$) are independent, the Bloch Equation (and in particular the Spin Domain formulation) are well suited to GPU acceleration.

1.2.3 Bloch-Riccati Equation

Equations 1.4 and 1.12 both describe the rotation of parameters related to the bulk magnetisation vector under the influence of an externally applied magnetic field and can be straightforwardly calculated. A far more challenging problem is presented when trying to invert the Bloch Equation, to find the external magnetic fields that are required to produce the desired behaviour of the system. To understand the inherent nonlinearity present in the Bloch Equation and justify the Small-Tip Angle (STA) spatial domain method in Chapter 2, a further modified formulation is presented.

Using the definition $\omega_1 = -(B_x + iB_y)$, then Equation 1.4 can be expressed as:

$$\begin{aligned} \dot{m}_{xy} + im_{xy}\omega_z + i\omega_1 m_z &= 0 \\ \dot{m}_z + \frac{i}{2}(\omega_1^* m_{xy} - \omega_1 m_{xy}^*) &= 0\,. \end{aligned} \tag{1.15}$$

Using the definition of the stereographic projection (*26*) of $\mathbf{m}$ as:

$$\Gamma = \frac{m_{xy}}{1+m_z}, \quad (m_z \neq -1)\,. \tag{1.16}$$

Differentiating Equation 1.16 with respect to time and using the expressions in Equation 1.15 leads to formulation of the Bloch Equation in *Riccati* form:

$$\frac{d\Gamma}{dt} = -\frac{i}{2}\omega_1 + i\omega_z\Gamma + \frac{i}{2}\omega_1^*\Gamma^2\,. \tag{1.17}$$

Equation 1.17 is a first-order nonlinear differential equation, but is equivalent to the coupled differential in Equation 1.4. As the magnetisation vector is fixed on the Bloch sphere, the dynamics can be equivalently described by a point in the complex-plane. Figure 1.2 shows the stereographic projection graphically. It should be noted that the dynamics of a magnetisation vector passing through $m_z = -1$, cannot be defined by the

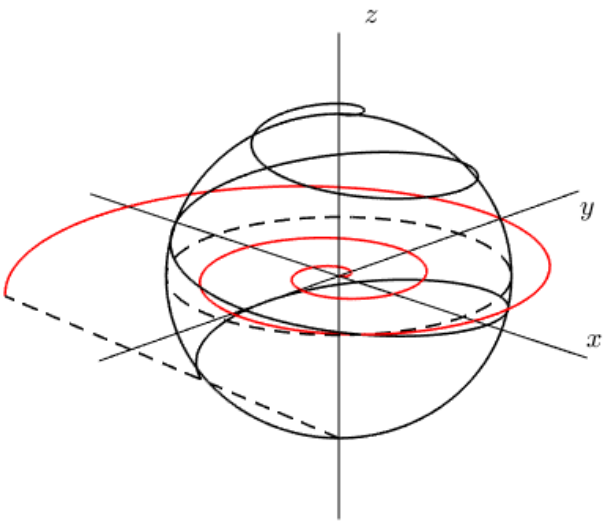

Figure 1.2: Stereographic projection into the transverse plane (red line) of the bulk magnetisation vector (black line) on the Bloch sphere in the rotating frame. The projection vector intersects the points $(0, 0, -1)$ and the instantaneous magnetisation vector on the Bloch sphere.

stereographic projection.

The nonlinearity in Equation 1.17 results from the final quadratic term in $\Gamma(t)$. With application of an integrating factor, a modified formulation is:

$$\frac{d\Gamma'}{dt} = -\frac{i}{2}\omega_1 e^{-i\int \omega_z dt} + \frac{i}{2}\omega_1^* e^{i\int \omega_z dt}\Gamma'^2 \,, \tag{1.18}$$

where,

$$\Gamma' = \Gamma e^{-i\int w_z dt} \,. \tag{1.19}$$

In the Small-Tip Angle (STA) case ($m_z \approx 1$), so $\Gamma' \approx 0$ and Equation 1.18 simply represents a linear Fourier relationship. Application of this method leads to magnetisation errors when linearly scaled to *large*-tip angles. Various techniques have been introduced to extend the linear domain of Equation 1.18 through use of nonlinear transformation functions such as hyperbolic-tangent (*27*) or power-series (*28*). Very recently, a non-trivial mapping of the nonlinear Bloch Equations to a coupled linear spring system was discovered by generalising the stereographic projection in Equation 1.16 to a dynamic projection (*29*). The Bloch-Riccati Equation permits an analytical solution for a special FM pulse as will be discussed in Chapter 2.

1.3 Electromagnetic Fields in Tissue at 7 T

1.3.1 Static Field

There are now over 70 human 7 T scanners worldwide, with this number expected to grow due to recent approval of the Siemens Terra as a medical device (*18*). A field strength of 7 T leads to an enhanced susceptibility effect, which provides a contrast mechanism to study microscopic structures and vasculature (*30*, *31*). Susceptibility changes between macroscopic media (typically tissue and air) also deform the static magnetic field and result in an off-resonance field $\Delta B_0(\mathbf{x})$ (*32*). One of the effects of macroscopic off-resonance is intravoxel dephasing (modelled by an additional parameter T_2'), leading to an increase in the signal decay rate. Strategic placement of passive ferromagnetic material has been traditionally used to ameliorate the field uniformity but is unable to correct for patient-specific conditions. Due to this, an active system utilising current driven *shim* coils is usually required to improve the field over a given region-of-interest.

1.3.2 Gradient & Shim Fields

The magnetic field in free space obeys the Laplace Equation, which has a general solution that can be expressed as a linear combination of spherical harmonics $Y^{\{n,m\}}$:

$$B_z(\mathbf{r}) = \sum_{n=0}^{\infty} \sum_{m=-n}^{n} a^{\{n,m\}} Y^{\{n,m\}}(\mathbf{r}) \,. \tag{1.20}$$

Although this solution is only strictly valid in the absence of tissue, the basis of spherical harmonics is universally applied for field expansions in MRI for convenience. Figure 1.3 displays up to 3rd order spherical basis functions on the surface of a cylinder. A total of $m = 2n + 1$ degrees exist for each order n.

Active shimming utilises a set of current-driven coil windings embedded in the bore of the scanner to produce low-order spherical harmonic field patterns. Practically, this is limited up to 2nd or 3rd order spherical harmonics due to cost and space limitations. To negate a given field pattern $\Delta B_0(\mathbf{x})$ requires the solution of an overdetermined linear system with current constraints (*33*):

$$\hat{I}_{\text{shim}} = \underset{I \in \mathbb{R}}{\arg\min} \, ||B_{\text{shim}}(\mathbf{x}) I_{\text{shim}} - \Delta B_0(\mathbf{x})||^2_{\mathbf{w}}, \qquad I^{\{n,m\}} < I^{\{n,m\}}_{\text{max}} \tag{1.21}$$

where the vector of currents is given by $I_{\text{shim}} = [I^{\{0,0\}}, I^{\{1,0\}}, I^{\{1,1\}}, \ldots, I^{\{n,m\}}]$, and $B_{\text{shim}}(\mathbf{x}) = [B^{\{0,0\}}(\mathbf{x}), B^{\{1,0\}}(\mathbf{x}), B^{\{1,1\}}(\mathbf{x}), \ldots, B^{\{n,m\}}(\mathbf{x})]$ and $B^{n,m}$ is the $\{n, m\}$ shim field per unit current. The maximum current permitted depends on the resistance of each specific coil.

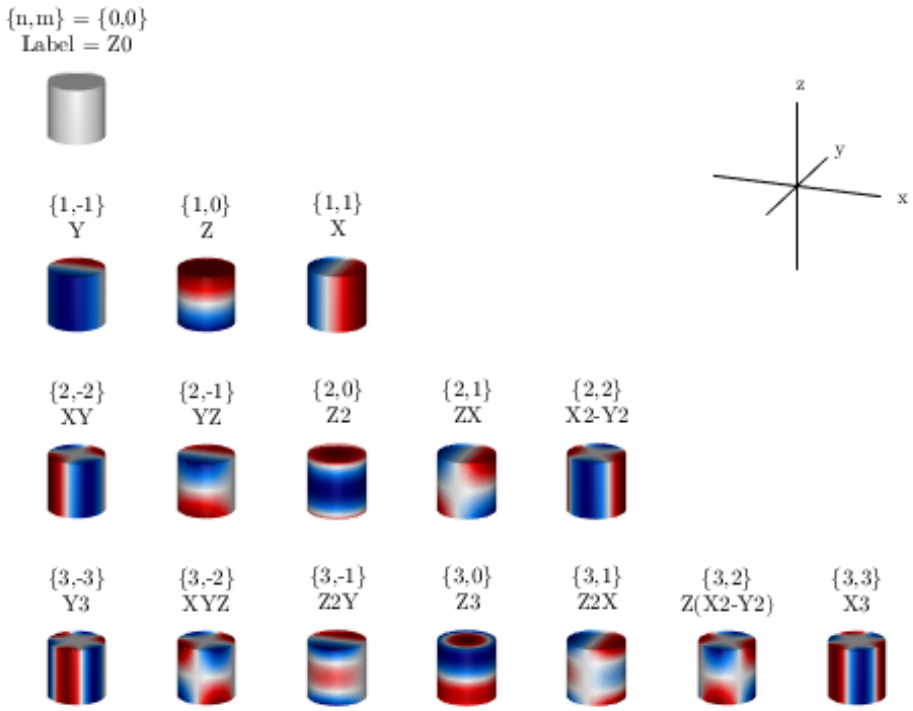

Figure 1.3: Spherical harmonics of order and degree $\{n, m\}$, visualised on the surface of a unit cylinder aligned on axis with the direction of the static-field.

Gradient & Shim System Components

The zeroth-order field adjustment is often performed by applying a frequency-offset rather than with a dedicated coil. The three first-order or *gradient* fields vary linearly with position and are the foundation of spatial encoding methods. As such, gradient coils require rapid switching in the kHz range. Gradient amplifiers are voltage-driven with current control feedback to generate the gradient waveform $G(t)$ required. To act as first-order shims, a separate DC offset is applied to each gradient coil. Higher-order shims require additional coil layers to approximate each spherical harmonic. Shim-coil amplifiers are generally required to provide constant current levels, rather than kilohertz switching (*34*).

The system response to a given input can strongly deviate from the desired behaviour due to number of mechanisms including amplifier performance and mechanical vibrations (*35*). Eddy currents are induced in conductors by any changing magnetic field within the bore of the scanner, and lead to opposing magnetic fields. Gradient coils are usually actively shielded with a secondary layer to reduce the gradient stray field and therefore limit induced eddy currents (*36*). Higher-order shim coils are usually unshielded, resulting in the generation of long-lived eddy currents when they are switched.

Imperfections in the gradient system have recently been modelled as a linear time-invariant system (*37*), with the system characterised by a Gradient Impulse Response Function (GIRF). The technique can be used to actively filter input waveforms (pre-emphasis) (*38*) or provide information for post-processing image data (*39*). Application of this approach for selective excitation will be discussed further in Chapter 3.

1.3.3 B_1^+ Fields

The shorter RF wavelength dictated by 7 T is comparable with the dimensions of the imaging volume leading to patient-specific interference patterns in both the transmit $B_1^+(\mathbf{r})$ and receive fields $B_1^-(\mathbf{r})$ (*19*). Single-channel transmission with a coil-array is performed in *quadrature*, whereby input signals are phase-shifted to achieve circular polarisation for maximal NMR coupling (*40*). For neurological imaging at 7 T, this results in the well known phenomena of *central-brightening* in the transmit field. Parallel Transmit (pTx) systems have been developed, consisting of multiple transmit coils driven independently using multiple RF amplifiers (*20*). In the rotating frame of reference, the B_1^+ spatial distribution is a superposition of the fields from each transmit channel:

$$B_1^+(\mathbf{r}, t) = \sum_{c=1}^{N_c} s_c(\mathbf{r}) b_c(t) \,, \tag{1.22}$$

where $s_c(\mathbf{r})$ is the *transmit sensitivity* of the cth channel with associated RF waveform $b_c(t)$. The sensitivity maps smoothly vary throughout the imaging volume but are localised to individual coil elements (see Figure 1.4). Antenna designs and coil-configurations

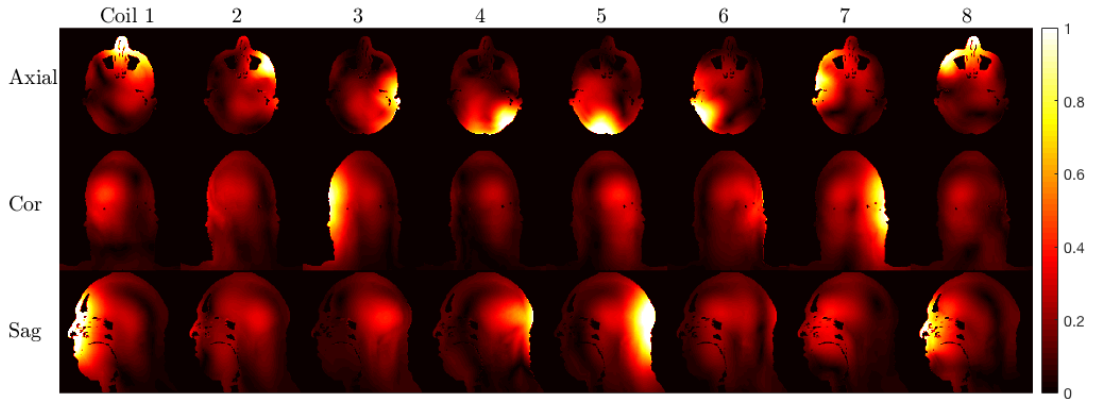

Figure 1.4: Relative transmit sensitivities (magnitude) in the head at 7 T, for central axial, coronal and sagittal slices (rows), for 8 transmit channels. As expected, the highest sensitivity values are adjacent to each coil element.

for head-only and body applications has been a highly-active research area in the last decade (*41–43*) and this is likely to continue, with the recent demonstration of a 32-channel insert body array (*44*).

Operating mode	Global Body SAR	Global Head SAR	Local Head and Body SAR	Local Extremity SAR
Normal	2 W/kg	3.2 W/kg	10 W/kg	20 W/kg
1st Level Controlled	4 W/kg	3.2 W/kg	20 W/kg	40 W/kg
2nd Level Controlled	>4 W/kg	>3.2 W/kg	>20 W/kg	>40 W/kg

Table 1.1: Limits for global and local SAR averaged over 6 minute time window. Adapted from International Electrotechnical Commission 2015 (*48*).

1.3.4 Safety considerations

An extensive review of safety considerations at 7 T is provided in (*45*). There is no evidence that exposure to the static main field at 7 T poses a risk to human health, although it can lead to transitory nystagmus (*46*, *47*).

Gradient system operation is largely independent of field-strength, and is bound by the same slew-rate constraints (*48*) at all field-strengths, to avoid peripheral nerve stimulation. Acoustic noise resulting from gradient systems can approach safety limits (*48*) and is frequently cited in patient-acceptance studies (*49*). Modified coil design (*50*) and vacuum-layers (*51*) are examples of hardware solutions, but software-defined waveform filtering is more commonly applied (*52*).

RF energy is deposited in tissue through the induction of currents in conductive material. Although forward and reflected power can be measured at the coil connections, the deposition of energy is highly non-uniform across the body. The Specific Absorption Rate (SAR) is used to quantify exposure to RF energy in tissue and is defined in units of W/kg. Safe operating limits are defined by the International Electrotechnical Commission, averaged over a 6 minute window (*48*) for both global and local 10-gram tissue averages (see Table 1.1).

For pTx systems, it is convenient to express the instantaneous spatial SAR in quadratic form in terms of model-based matrices $Q_n \in \mathbb{C}^{\{N_c,N_c\}}$ (*53*):

$$\mathrm{SAR}_n(t) = b(t)^H Q_n b(t), \quad n \in \{1,\ldots,N\}, \, , \tag{1.23}$$

The set of N Q-matrices is often reduced to a set of *virtual observation points* using a compression algorithm (*54*, *55*) to an arbitrary degree of peak local SAR overestimation. The benefit of the Q-matrix approach is to precalculate all mass averaging, physical tissue properties and electric field sensitivities; allowing fast evaluation of local SAR for a given multichannel RF waveform $b(t)$ from Equation 1.23. Generation of this set of matrices requires appropriate electromagnetic modelling of the interaction between the patient and the coil (see (*56*, *57*) for further details).

Solving Pennes' bioheat equation (*58*), a compressed set of virtual observation points for temperature rise can be calculated which can permit higher transmit power (*59*, *60*).

Pennes' original bioheat equation models local tissue temperature change in the presence of ideal blood perfusion. An additional term includes energy absorption related to SAR (*61*).

1.4 Typical MRI sequences

An authoritative review of contemporary MRI sequences can be found in reference (*62*). The ambition of this section is to introduce sequences that are utilised in later chapters.

1.4.1 Gradient-Echo Imaging

Gradient-echo imaging (*16*) is typically deployed as a Cartesian-sampled Fourier imaging sequence (*63*). Cartesian sampling of k-space has the unique property that a discrete Fourier transform (and hence Fast Fourier Transform), can be applied directly to the k-space data when evenly sampled.

After the application of a slice-selective RF pulse, the excited magnetisation is assumed to be in-phase across the slice. Application of a *phase-encoding* gradient blip will add a spatially dependent phase-shift across the slice in a given direction. Simultaneously, a second orthogonal gradient blip will additionally dephase the magnetisation across the slice in the *read-out* direction. After a given time, the polarity of this second gradient will be reversed and data acquisition commences concurrently (see Figure 1.5). The signal intensity will reach a maximum when the signal rephases in the read-out direction (forming a *gradient echo*) before dephasing again (with opposite polarity). The free-induction signal envelope decays with T_2^*, meaning the sequence is susceptible to off-resonance effects.

It is preferable to align the larger FOV dimension with the read-out direction due to the higher sampling density achievable. Sampling density along the phase-encoded dimension is dictated by the number of repetitions. The repetition time is the time between each successive RF pulse. Steady-state magnetisation is reached after a number of RF pulses, whereby the longitudinal component returns to the same value in each repetition interval.

Between acquisition of multiple k-space lines, transverse magnetisation can be *spoiled* to give the sequence a simpler T_1-weighting. This can be achieved using either long repetition times, gradient dephasing, or RF phase-cycling (*64*, *65*).

Turbo Field Echo (TFE) imaging is a subset of gradient echo imaging, where a low flip-angle RF pulse is used in concert with a short repetition time. The acquisition of data is performed as the signal approaches steady-state in this case. Magnetisation preparation steps (i.e. before the the first read-out pulse) are used to enhance the contrast of the sequence (*66*).

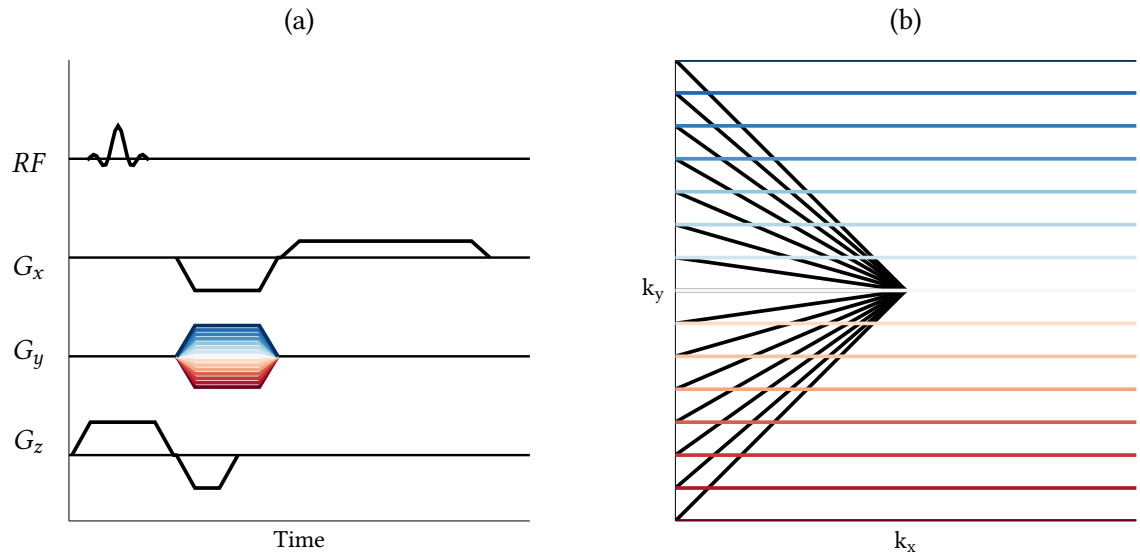

Figure 1.5: Gradient-echo pulse timing sequence (a) and the effective k-space traversed throughout (b). After each RF pulse with slice-selective gradient G_z, a prephasing gradient and phase-encoding step gradient G_y are used to set the initial point of each line in k-space. A subsequent read-out gradient G_x is used to traverse a whole line of k-space while the signal is discretely sampled.

1.4.2 Echo-Planar Imaging

Echo-Planar Imaging (EPI) (*14*) utilises bipolar read-out gradients to acquire N_{ETL} evenly-spaced gradient echoes under the envelope of a single free-induction decay. Each gradient-echo is individually phase-encoded (albeit with different T_2^* weighting), so an entire 2D k-space can be covered in a *single-shot*. The inherent efficiency of the sequence can be utilised to acquire data with high temporal resolution, which is crucial to study physiological processes such as perfusion using arterial spin-labelling. In particular, EPI is pivotal to functional MRI by blood oxygenation level dependent contrast (*67*).

The increased temporal resolution offered by EPI can unfortunately be constrained by a number of imaging artefacts. Nyquist ghosting refers to the appearance of ghost images in the spatial domain offset by FOV/2 in the phase-encoding direction (*68*). As alternate echoes are acquired with different gradient polarity, any mechanism that causes a spatially independent phase error will lead to a relative shift of k-space lines. Using a single polarity read-out gradient (with a rewinding gradient) can avoid this issue.

The extended total sampling time for EPI leads to geometric distortion (see Figure 1.7) in the presence of an off-resonance field $\Delta B_0(\mathbf{r})$ (*69*). Given an echo-spacing τ_{esp}, for a FOV of $N\Delta y$, and parallel imaging acceleration factor R_y, the distortion (*70*) in the phase-encoding direction y is given by:

$$d = \frac{\gamma}{2\pi} \frac{\tau_{\text{esp}} N \Delta y}{R_y} \Delta B_0 \, . \tag{1.24}$$

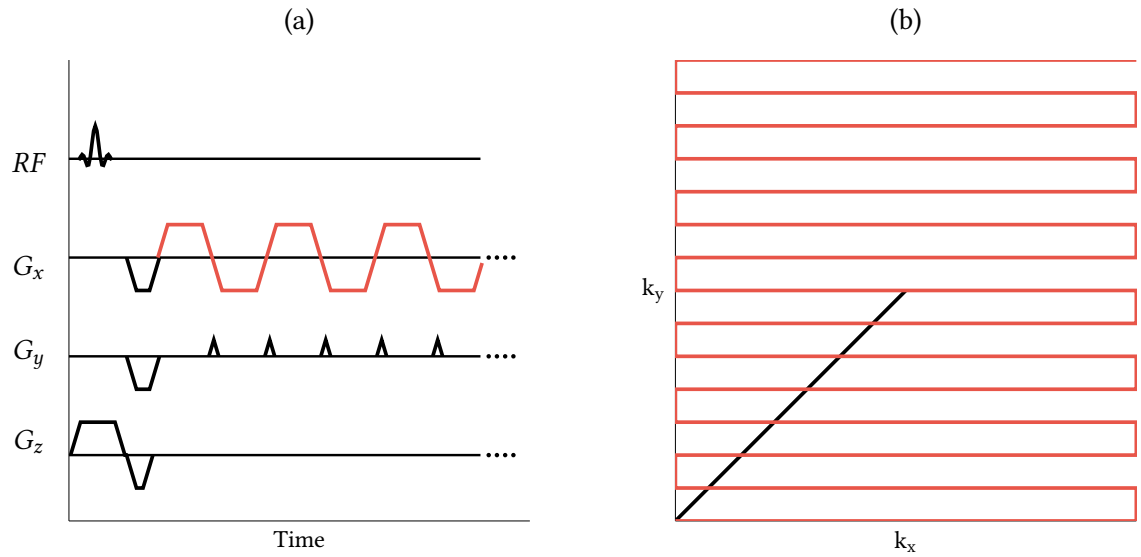

Figure 1.6: EPI pulse timing sequence (a) and the effective k-space traversed throughout (b). After each RF pulse with slice-selective gradient G_z, prephasing gradients are used to set the initial point in k-space. All of k-space is traversed after a single RF pulse, with a trapezoidal readout gradient G_x and phase-encoding gradient G_y. Sampling data along the readout direction is only limited by the discretisation process of signal sampling, whereas sampling along the phase-encoding direction is limited by the T_2^* of the tissue.

Nonlinear warping of the EPI data using an off-resonance map is commonly used to correct for the geometric distortion (*71*, *72*). Nevertheless, it is always preferable to acquire the signal in the absence of any field off-resonance, to avoid signal 'pile up' in voxels.

An additional effect of an off-resonance field for EPI acquisitions is spatially variable signal drop-out due to intravoxel dephasing. Sampling gradient echoes under the envelope of a spin-echo decay (spin-echo EPI) refocuses this dephasing, but this approach has lower functional sensitivity (*73*).

1.4.3 Field Mapping

Accurate measurement of the underlying magnetic field patterns in vivo is essential for the applications of the pulse-design and static-field shim methods discussed in this thesis. This includes a measurement of the spatial B_1^+ distribution (or coil sensitivities) for each of the individual transmit coils, and a separate measurement of the local B_0 field. These measurements form part of a series of preparation steps performed at the beginning of each session alongside traditional power optimisation, centre-frequency optimisation, receive sensitivity estimation, and coil-coupling measurement.

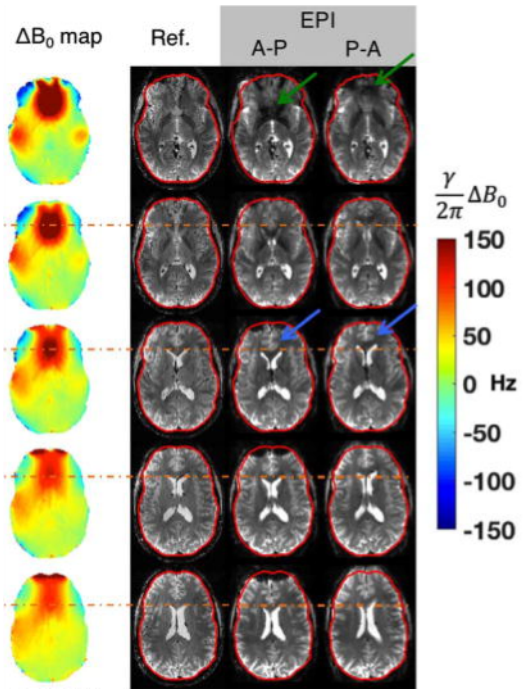

Figure 1.7: Figure adapted from Stockmann & Wald (*32*). EPI distortion at 7 T for 5 axial slices. First column shows the off-resonance map after global shimming up to 2nd order has been applied. The second column shows a reference gradient echo acquisition. The EPI geometric distortion is visible (blue arrows) in the third and fourth columns with the phase-encoding first in the A-P direction, and then again in the reversed P-A direction (2 mm isotropic resolution, parallel imaging factor = 2). Signal dropout is highlighted with green arrows.

Actual Flip-Angle Imaging for B_1^+ mapping

Measurement of the transmit field (and inference of multichannel transmit sensitivities) is a common calibration procedure and vital to methods discussed in this thesis. A number of methods have been proposed for this purpose that use multiple complex images acquired either from repetitive imaging with multiple angles (*74*) or stimulated echoes (*75*, *76*).

The method used in this thesis is the Actual Flip-Angle Imaging (AFI) sequence (*77*), whereby the signal ratio $S(\mathbf{r}) = \mathrm{S}_2(\mathbf{r})/\mathrm{S}_1(\mathbf{r})$, from two spoiled gradient-echo acquisitions with different repetition times ($\mathrm{TR}_1 < \mathrm{TR}_2$) leads to an estimation of the flip-angle as:

$$\theta(\mathbf{r}) = \cos^{-1}\left(\frac{S(\mathbf{r})n - 1}{n - S(\mathbf{r})}\right), \tag{1.25}$$

where $n = \mathrm{TR}_2/\mathrm{TR}_1$. Increasing the ratio of repetition times increases the sensitivity to flip-angle, but increases the duration of the overall sequence. Deviation from the first-order signal approximation used in the derivation of Equation 1.25 occurs for very short T_1 components.

For pTx systems, the large dynamic range of B_1^+ at 7 T can be reduced by finding an optimal linear combination of coil channel elements which is similar across a number of

subjects (*78*). Each linear combination must be distinguishable such that the resulting linear system of equations be well conditioned.

In this manner, the underlying B_1^+ channel profiles can be estimated without large noise amplification.

Relative channel mapping can be achieved by combining a single coil transmit sensitivity, with individual channel gradient echo images sequentially acquired from each of the channels (*79*):

$$s_c(\mathbf{r}) = \frac{S_c^{\mathrm{GRE}}(\mathbf{r})}{\sum_{j=1}^{N_c} S_j^{\mathrm{GRE}}(\mathbf{r})} s_{\mathrm{QUAD}}(\mathbf{r}) \,. \tag{1.26}$$

The spatially variable phase profile of each transmit channel is quantified relative to a fixed channel, to remove receive phase and structural information that is inherent in all acquisitions.

Dual TE Method for B_0 mapping

The phase accrued at each voxel location is proportional to the echo-time and local field strength:

$$\varphi(\mathbf{r}) = \varphi_0(\mathbf{r}) + \gamma \Delta B_0(\mathbf{r}) \cdot \mathrm{TE} \,. \tag{1.27}$$

A minimum of two image acquisitions with different echo-times can therefore be used to extract the local field variation $\Delta B_0(\mathbf{r})$ using a suitable phase unwrapping algorithm (*80*).

1.5 Numerical & Computational Methods

1.5.1 Least square problems

Conjugate Gradient Least Squares (CGLS) (*81*) is an iterative method, useful for the solution of large linear least square problems of the form:

$$\min_x ||Ax - d||^2 \,, \tag{1.28}$$

where the matrix A can be square or rectangular. The generic CGLS algorithm to solve Equation 1.28 is shown in Algorithm 1. The algorithm is dominated by two large matrix-vector products on iteration, which can either be performed implicitly (if a suitable method is available), or explicitly. For the latter, it should be noted that matrix-vector products are well suited to parallel computing (*82*).

A related algorithm is known as Least Squares Quadratic Reduction (*83*), which is equivalent to applying CGLS to the normal equation $A^H Ax = A^H d$ but has better numerical properties if the matrix A is ill-conditioned (*84*).

Algorithm 1: Conjugate Gradient Least Squares (CGLS)
CGLS is an iterative algorithm for estimating the solution of least squares problems of the kind specified in Equation 1.28.

$r \leftarrow d; s \leftarrow A^H r; p \leftarrow s; \nu \leftarrow ||s||^2;$ *// initialise auxiliary parameters*
while $i <$ *max iterations* **do**
  $q \leftarrow Ap$; *// large matrix-vector product*
  $\alpha \leftarrow \nu / ||q||^2$
  $x \leftarrow x + \alpha p$; *// update solution vector*
  $r \leftarrow r - \alpha q$
  $s \leftarrow A^H r$; *// large matrix-vector product*
  $\nu_{\text{old}} \leftarrow \nu;\ \nu \leftarrow ||s||^2$
  $\beta \leftarrow \nu / \nu_{\text{old}}$
  $p \leftarrow s + \beta p$
  $i \leftarrow i + 1$
end

1.5.2 Interior-Point Algorithm

Constrained optimisation refers to optimisation of a function subject to some functional constraints on the design variables (*85*), and can be specified as:

$$\min_{x \in \mathbb{R}^n} f(x), \quad \text{s.t.} \quad g(x) \leq 0\,, \tag{1.29}$$

where $\{f(x), g(x)\}$ are smooth nonlinear functions. The interior-point algorithm is a commonly applied method for solving problems of this kind. The effects of the constraints are included through use of a logarithmic barrier function to approximate the problem statement in Equation 1.29 with:

$$\begin{aligned} &\min_{x \in \mathbb{R}^n, \mu} f(x) - \mu \sum_{i=1}^{m} \log(s_i(x)) \\ &\text{s.t.} \quad g(x) + s(x) = 0\,. \end{aligned} \tag{1.30}$$

The number of auxiliary variables s_i, where $i \in \{1, \ldots, m\}$, comprise the vector $s(x)$ and matches the number of inequality constraints $g(x)$ required. As $\mu \rightarrow 0$, the minimum of Equation 1.30 should approach the minimum of the original problem, Equation 1.29. The Lagrangian of the barrier problem in Equation 1.30 is:

$$\mathcal{L}(x, s, \lambda) = f(x) - \mu \sum_{i=1}^{m} \log(s_i(x)) + \lambda^T (g(x) + s(x))\,, \tag{1.31}$$

where λ is the Lagrangian multiplier for the inequality constraints in Equation 1.29. The algorithm switches between a line search method to a decrease a further merit function and a trust region method that ensures iterations are feasible. An extensive discussion

of the algorithm can be found in (*86*).

1.5.3 Evolutionary Algorithms

Evolutionary algorithms refer to a wide class of population-based optimisation algorithms inspired by biological dynamics and mechanisms. Pioneered in the 1960s, evolutionary algorithms have recently undergone a renaissance due to requirements from Deep Learning, combined with the emergence of large-scale computing resources (*87*).

The genetic algorithm (see Algorithm 2) is motivated by the biological process of natural selection and was introduced by Holland et al (*88*). An initial population (either seeded and/or randomly selected) of *individuals* (a set of parameters) is evaluated using a supplied cost function. Future generations of individuals are created by the processes of multiple genetic operators. The first of these is the *crossover* operator which randomly selects parameter values from a pair of *parent* individuals to create a new individual. The *mutation* operator randomly multiplies one or more parameter values from a parent individual. The *elitism* fraction ensures a predetermined number of parent individuals with the lowest function score are preserved for the next generation of the optimisation.

The genetic algorithm are well suited to global optimisation problems, by naturally exploring the search space in an intelligent manner. The downside of this approach is that it requires many more function evaluations than derivative-based methods.

Algorithm 2: Genetic Algorithm
The Genetic Algorithm is a global optimisation algorithm. A population of solution vectors are evaluated according to a fitness function. Genetic operators of elitism, crossover and mutation create a modified population at each iteration.

$P_0 = \{x_i : i \in \{1, \ldots, N\}\}$; *// initialise population of solutions, size N*
while *stop condition* **do** *// check convergence criteria*
  $f_i \leftarrow f(x_i), \quad \forall i \in \{1, \ldots, N\}$; *// evaluate population fitness*
  $f_i \leftarrow \text{sort}(f_i), \quad \forall i \in \{1, \ldots, N\}$; *// sort fitness values*
  if $i \in \{1, \ldots, (1-\chi)N)$ **then**
    $P_{k+1}(x_i) \leftarrow P_k(x_i)$; *// elitism fraction χ*
  else if $i, j \in \{1, \ldots, (1-\chi)N)$ **then**
    $P_{k+1}(x_i) \leftarrow \text{cross}(x_i, x_j)$; *// crossover function*
  else
    $P_{k+1}(x_i) \leftarrow \text{mutate}(x_i, \mu)$; *// mutation operation*
  end
end

1.5.4 Graphical Processing Unit

Heterogeneous or *hybrid* computing makes use of more than one type of processor, to achieve a higher throughput. This is typically a multiprocessor CPU, and at least one

Graphical Processing Unit (GPU). The basic architecture of a GPU is a series of streaming multiprocessors (see Figure 1.8), typically with 64 cores each and small amounts of cached and shared memory. Groups of cores (typically 32), execute the same instruction, but with different data, which is extremely efficient for element-wise operations.

Repeated transfer of large amounts of data across the CPU-GPU PCIe connection should be avoided where possible, as this can often result in overall slower computational times than executing code sequentially on the CPU. Iterative algorithms that repeatedly perform operations on data residing in GPU global memory effectively hide this latency.

CUDA (Compute Unified Device Architecture) provides an extension of the C/C++ programming languages specifically for GPUs. Higher-level interpreted languages now provide direct support for CUDA enabled GPUs such as PyCUDA for Python and the Parallel Computing Toolbox in MATLAB (MathWorks, Natick, USA).

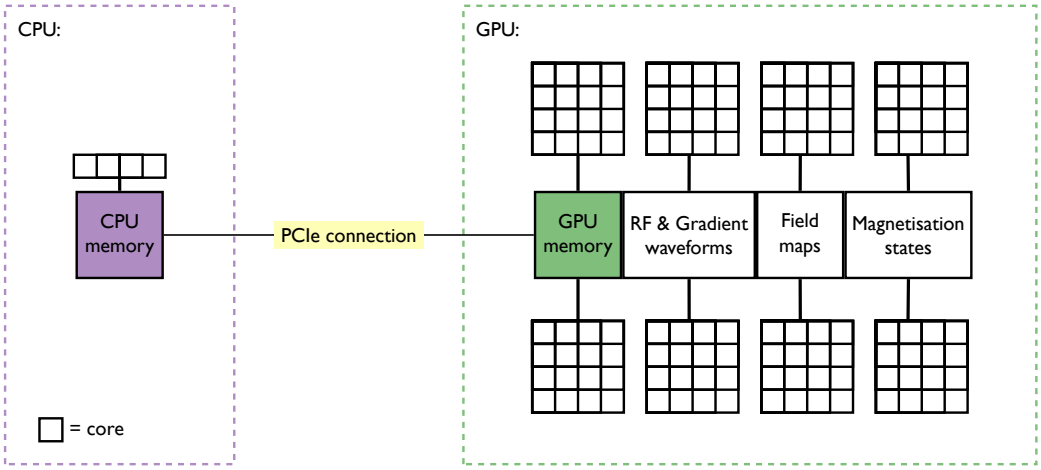

Figure 1.8: Schematic diagram of CPU-GPU layout. Required data for optimisations and simulations (e.g. field maps) can remain in GPU global memory for the duration (or beyond) of multi-step calculations. In this way, multiple transfers across the relatively slow PCIe connection can be avoided.

1.6 Summary

The aim of this chapter was to provide an overview of the key concepts and methods relevant to the work in this thesis. Firstly, the governing equations of the magnetisation vector were described along with other alternative compact formulations. These will be used throughout this thesis as a way to simulate the response of the magnetisation to a user-designed RF input.

As the work within this thesis is performed at a field strength of 7 T, the three requisite magnetic fields used in MRI experiments were discussed, with particular emphasis

on the various field inhomogeneities that are prevalent in both the static and transmit fields at 7 T in vivo. As the mapping of both of these distinct field inhomogeneities is crucial to the countermeasures and approaches proposed in this thesis, a description of field mapping sequences and methods was also discussed for both the static and transmit fields.

Finally, some of the numerical optimisation methods that are utilised in later chapters were described.

Chapter 2

Existing Parallel Transmit Design Methods

The aim of this chapter is to introduce some of the existing design strategies and capabilities of Parallel Transmit (pTx). In particular, it is described how pTx can be utilised to accelerate RF pulses. Analysis and design of Frequency Modulated pulses are also considered with extension for pTx.

2.1 RF Shimming

2.1.1 Static RF shimming

Static RF shimming applies a complex-valued weighting w_c to each of the available transmit channels ($c \in \{1, \ldots, N_c\}$) for the entire acquisition duration. Each channel transmits an amplitude-scaled and phase-shifted version of the same base waveform $b(t)$, such that the instantaneous transmit field is:

$$B_1^+(\mathbf{x}, t) = b(t) \sum_{c=1}^{N_c} s_c(\mathbf{x}) w_c \,, \tag{2.1}$$

where $s_c\mathbf{x}$ are the transmit sensitivity profiles. Many optimisation schemes for the weights in Equation 2.1 have been suggested, but fall into two main contrasting categories. The first promotes constructive interference to improve transmit field uniformity within a region-of-interest (*89*). The second targets increased transmit efficiency (a measure of field strength per unit input power). Approaches to this include nulling the average transmit phase over a small region (*90*) and maximising a measure of transmit efficiency (*91*), numerically or analytically (*92*).

RF shimming has been applied to improve transmit efficiency for spectroscopy in the human brain (*93*, *94*). There are also significant advantages to utilising RF shimming in the body at 7 T, where B_1^+ field interactions are more complicated. Of particular note are applications towards B_1^+ targeted towards small organs, such as renal (*95*, *96*), prostate (*90*), and cardiac applications (*97*). An extensive review of body imaging applications involving RF shimming can be found in (*98*).

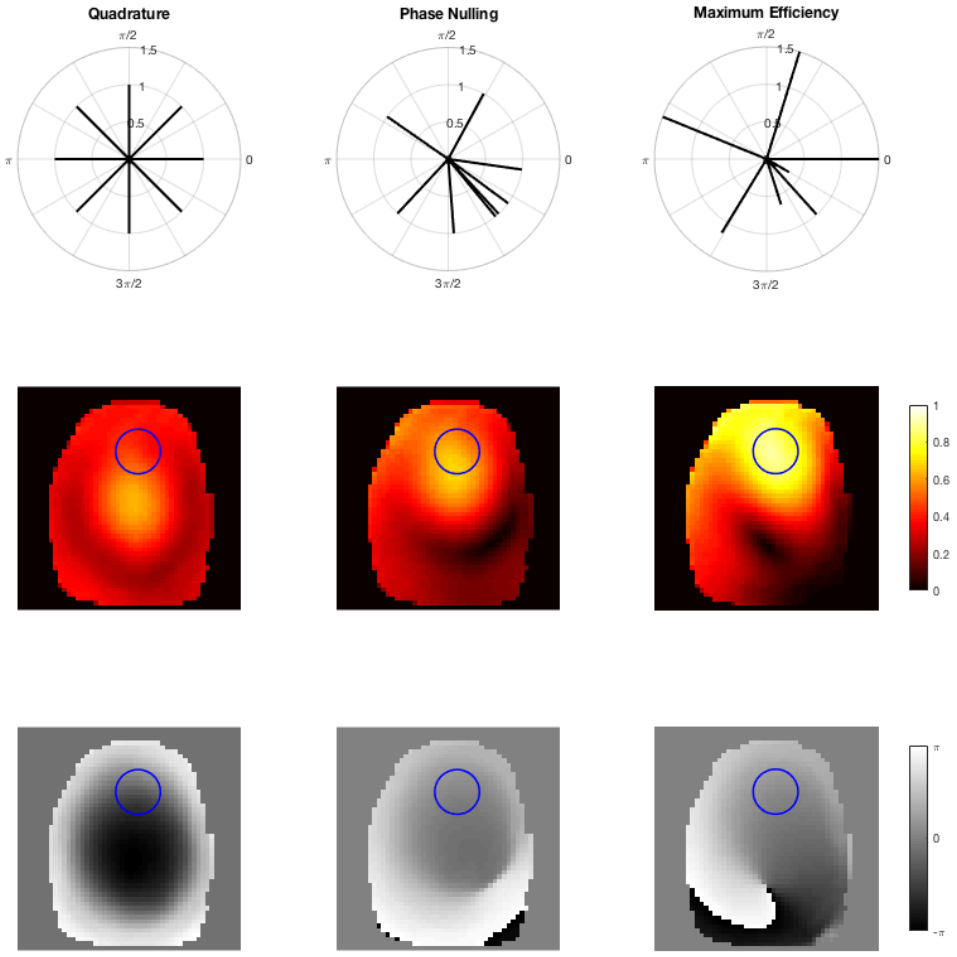

Figure 2.1: Static RF Shimming for circular ROI. Polar plots (top row) display complex shim values for 8 channels each denoted by a different colour, and the resulting predicted B_1^+ amplitude (middle row) and phase map (bottom row). Each RF shim setting is power matched to the quadrature solution. The first column represents default quadrature, second column shows result of relative phase nulling, last column shows the maximum efficiency solution.

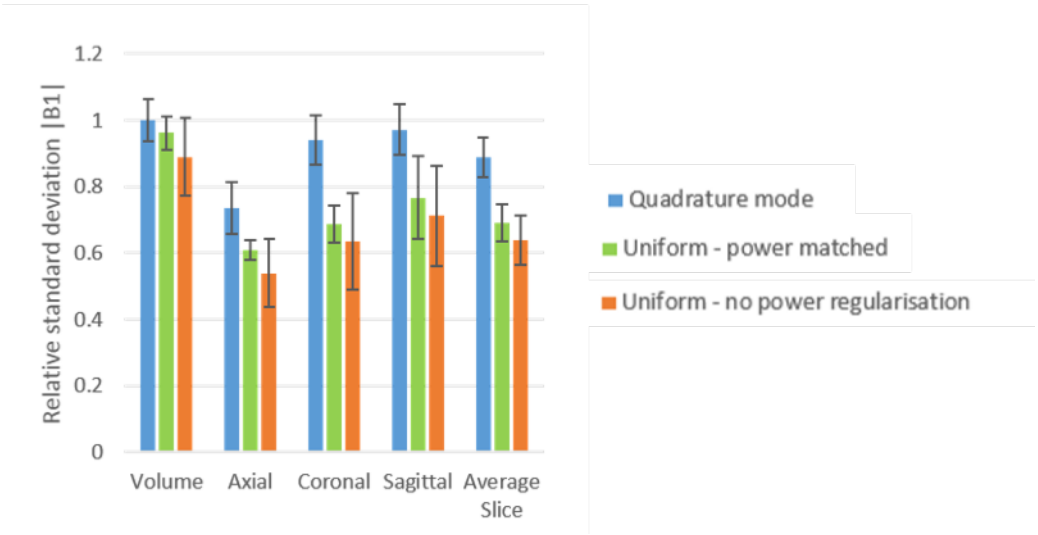

Figure 2.2: Relative standard deviation of $|B_1^+|$-field, normalised by values for quadrature volume mode for dynamic slice-by-slice RF shimming across the human brain when using Nova 8-channel transmit coil. Error bars demonstrate variation between 6 volunteers. Uniform power matched RF shimming refers to optimised RF fields, with matched power to quadrature mode. For comparison, uniform RF shimming with no power regularisation is also shown. Adapted from (*101*).

2.1.2 Dynamic RF shimming

Dynamic RF shimming allows the transmit channel weightings of RF objects to vary at the *sequence level*. The timescale on which dynamic updates to the RF weights occurs is then of the order of the sequence TR. Temporal updates to RF weights also occur for sparse pulse design methods (e.g. *spokes* (*99*)). As these changes occur on a timescale of a few milliseconds for a single RF object, this is not considered as dynamic RF shimming in this thesis.

An obvious application of dynamic RF shimming is multi-slice acquisitions, where the channel weightings can be optimised individually for the transmit field for each slice (*100*). In particular, the author of this thesis investigated slice-by-slice RF shimming with the commonly used Nova 8-channel transmit head coil (*101*). From each orthogonal geometry, slice-by-slice RF shimming on axial slices reduced the standard deviation of the $|B_1^+|$ field by the largest percentage (39.0% ± 3.9%, averaged over 6 volunteers) for the same power as quadrature transmission (see Figure 2.2).

Dynamic RF shimming can also be used to reduce the SAR per unit time of a sequence by trading-off transmit field uniformity and the resulting SAR whilst maintaining image quality (*102*).

For multi-pulse sequences, dynamic RF shimming can be utilised to modulate transmit channel weightings over the duration of the sequence to improve signal properties. A spatially resolved extended phase graph formalism was developed to model the signal evolution at each spatial position over time (*103*, *104*). Optimisation of transmit channel

weights against the desired behaviour of the signal model has been efficiently implemented as an optimal control problem (*105*), and applied for both *FLAIR* (*106*) and fast spin-echo sequences (*107*).

In summary, whereas static RF shimming maintains a fixed amplitude and phase relationship between transmit channels for the entire acquisition, dynamic RF shimming allows this relationship to vary for RF objects at the sequence level. Full pTx control envisages independent RF waveforms on each channel, with microsecond resolution.

2.2 Amplitude Modulated Pulse Design

As discussed in Chapter 1 the Bloch Equation can be linearised for Small-Tip Angle excitations, which is the starting point for many pTx pulse design methods. The *spatial-domain* method (*108*) generalises earlier single-coil methods (*109*), using an iterative optimisation to minimise the excitation error. In the Small-Tip Angle approximation (*17*), the transverse magnetisation $m(\mathbf{x})$ can be decoupled from the longitudinal component $m_z(\mathbf{x})$, and becomes a linear sum of individual transmit channel contributions:

$$m(\mathbf{x}) \approx i\gamma m_0 \sum_{c=1}^{N_c} s_c(\mathbf{x}) \int_0^T b_c(t) e^{i\Delta\omega(\mathbf{x})(t-T)} e^{i\mathbf{x}\cdot\mathbf{k}(t)} dt \,, \tag{2.2}$$

where the definition of transmit k-space is made as:

$$\mathbf{k}(t) = -\gamma \int_t^T \mathbf{G}(\tau) d\tau \,. \tag{2.3}$$

Note that the integration runs forward from the instantaneous time to the end of the duration T, this is in contrast to the definition of readout k-space (see Equation 1.3).

After temporal discretisation with N_t points and step-size Δt, the vector of transverse magnetisation values at multiple spatial locations is expressed as:

$$\mathbf{m} = \sum_{c=1}^{N_c} \mathbf{s}_c \mathbf{F} \mathbf{b}_c \tag{2.4}$$

with system encoding matrix $F_{i,j} = i\gamma m_0 \Delta t e^{i(\Delta\omega_i(j-N_t)\Delta t + \mathbf{x}_i\cdot\mathbf{k}_j)}$. For uniformly sampled k-space points and in the case where off-resonance effects can be neglected, the encoding matrix is a discrete Fourier transform, and hence the forward evaluation of Equation 2.4 can be performed as the summation of N_c fast Fourier transforms, avoiding the computational and memory overhead of constructing $\mathbf{F}$ explicitly. As uniformly sampled k-space trajectories would be extremely limiting, non-uniform fast Fourier transforms can be used where an additional convolutional gridding step is required (*110*, *111*). In the general case where an off-resonance map is used, temporal or frequency interpolation

has to be used to approximate its effect (*112*), although it is often simpler to construct the system encoding matrix directly.

Given a desired spatial transverse magnetisation target $\mathbf{m}_d$, the pulse design can be formulated as a least squares problem:

$$\hat{\mathbf{b}} = \underset{b \in \mathbb{C}^{N_t N_c}}{\arg\min} \{ || \sum_{c=1}^{N_c} \mathbf{s}_c \mathbf{F} \mathbf{b}_c - \mathbf{m}_d ||_{\mathbf{w}}^2 + R(\mathbf{b}) \} \ . \tag{2.5}$$

The regularisation function $R(\mathbf{b})$ can be used to control total power, peak power, or local SAR (*113*). The minimisation problem stated in Equation 2.5 can be routinely solved by the conjugate-gradient method or equivalently by least-squares quadratic reduction (*83*), and can have improved pulse properties compared to direct inversion.

For Tikhonov regularisation $R(\mathbf{b}) = \lambda ||\mathbf{b}||^2$, the Tikhonov parameter λ controls the cost of using large values of the design variable $\mathbf{b}$ (*114*). The optimal value of this regularisation parameter is usually not known before optimisation, and is determined by a secondary method. The popular L-curve method (*115*) solves the optimisation statement in Equation 2.5 for a range of Tikhonov values, and finds the optimal value directly. Sbrizzi et al (*116*) demonstrated an efficient modified multi-shift CGLS algorithm that can solve for a vector of Tikhonov values simultaneously. The multi-shift method exploits the redundancy that the computationally intensive matrix vector products on each iteration step are common across Tikhonov values (*117*).

To evaluate the effectiveness of a designed set of RF pTx pulses, the Normalised Root Mean Squared Error (NRMSE) is defined as:

$$\text{NRMSE} = \frac{||\mathbf{m} - \mathbf{m}_d||}{||\mathbf{m}_d||} \tag{2.6}$$

Acceleration via pTx

To design multidimensional spatially-selective RF pulses in the STA regime, a trajectory should be first designed that adequately samples transmit k-space for a desired spatial resolution Δx and Field of Excitation (FOX). The pulse duration is then determined by the maximum permissible gradient slew-rate. Undersampling k-space when the RF pulse is applied with a single-transmission coil will result in aliasing in the excitation profile. Parallel Transmit can be used to suppress these aliasing contributions, by utilising the additional localised spatial information of each coil element.

A straightforward example to demonstrate this approach is to consider the Echo-Planar trajectory, whereby two-fold acceleration can be achieved by halving the density of the lines. In this manner, the phase-encoded direction is undersampled by a factor of 2, leading to aliasing in the spatial domain with single-channel transmission. For pTx, a minimum of two independently driven coil elements that each have a localised sensi-

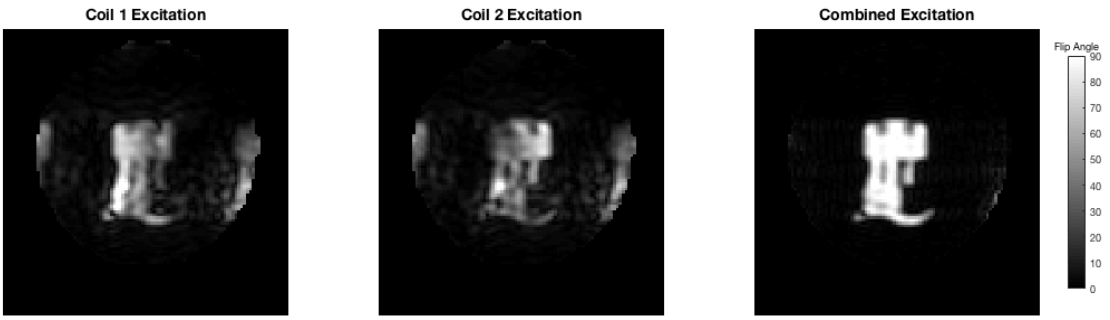

Figure 2.3: Acceleration with 2x undersampled echo-planar trajectory. Aliasing (phase-encoding is left-to-right here) is observed in each of the excitation patterns for each coil element. Combination of both coils is sufficient to remove aliasing artefact.

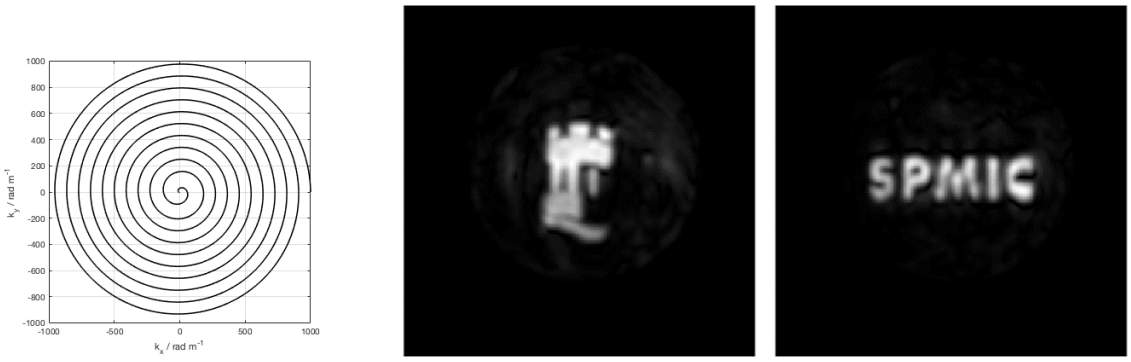

Figure 2.4: Example 2D spiral k-space sampling pattern, resulting pulse duration of 10.3 ms. Example gradient echo magnitude images (axial slice) for a spherical phantom with a radius of 10 cm. RF pulses designed using spatial domain method (*108*) scaled for 90° flip angle (see Equation 2.5).

tivity are required. In the example shown in Figure 2.3, the summed contribution from two diametrically opposed coil elements is sufficient to remove the aliasing excitation.

Higher acceleration factors can be achieved by increasing the number of localised coil elements, although this becomes challenging for high density transmit coil arrays due to significant spatial overlap (*118*). Most commercial systems are currently limited to eight transmit channels due in part to the high cost of RF amplifiers. Array-compressed pTx has recently been demonstrated to overcome this, whereby a small number of transmit channels drive a larger number of coil elements through an in situ optimised array network (*119*, *120*).

2.2.1 Time Optimal k-space Trajectories

Any existing readout k-space trajectory can be equivalently used for excitation. However, as excitation k-space is traversed in a *single-shot*, it is advantageous to omit areas that have negligible weighting for the final magnetisation state. For multidimensional pulses, limits on gradient amplitude slew-rate can led to unsatisfactorily long pulse durations, which are sensitive to off-resonance and even relaxation effects (*121*).

Lustig et al (*122*) addressed the design of time optimal gradient waveforms through the application of a non-iterative algorithm. Gradient-defined k-space trajectories are 1D curves embedded in a 3D space as $\mathbf{k} : [0, T] \rightarrow \mathbb{R}^3$. A convenient parameterisation of curves of this form is the Euclidean arc-length $s(t)$ such that the curve can be expressed as:

$$\mathbf{k}(s) = (x(s), y(s), z(s)) \,, \tag{2.7}$$

where the Euclidean arc-length for the curve is defined in the interval $[0, L]$ and given by:

$$s(t) = \gamma \int_0^t |\mathbf{G}(\tau)| d\tau \,. \tag{2.8}$$

The uniqueness of this parameterisation $k(s)$ is that it has unit speed with respect to the arc-length parameter, and the magnitude of the acceleration is equal to the instantaneous curvature $\kappa(s)$. The time optimal gradient waveform design problem seeks to minimise the duration T, given a desired k-space trajectory, subject to a trade-off between the maximum gradient amplitude $G_{\max}$ or slew-rate $S_{\max}$. The equivalent formulation of this problem in the arc-length parameterisation is to find the optimal velocity as a function of arc-length:

$$\begin{aligned} \min_{s} \quad & T \\ \text{s.t.} \quad & \dot{s}(t) \leq \min\{\gamma G_{\max}, \sqrt{\frac{\gamma S_{\max}}{\kappa(s)}}\} \\ & |\ddot{s}(t)| \leq (\gamma^2 S_{\max}^2 - \kappa^2(s)\dot{s}^4)^{\frac{1}{2}} \\ & s(0) = \dot{s}(0) = 0 \\ & s(T) = L \,. \end{aligned} \tag{2.9}$$

The first two constraints in Equation 2.9 reflect the underlying constraints on the maximum absolute gradient amplitude $G_{\max}$ and slew-rate $S_{\max}$, but also include the geometry of the trajectory itself. The additional constraints are boundary conditions for the arc-length that must be fulfilled for a representative trajectory. The solution to the problem statement in Equation 2.9 is expressed as the minimum of two ordinary differential equations, solved forwards and backwards (see Figure 2 in (*122*)). The gradient waveforms can then be calculated from the relation:

$$\mathbf{G}(t) = \frac{1}{\gamma} \frac{d\mathbf{k}(s(t))}{dt} = \frac{1}{\gamma} \mathbf{k}'(s(t)) \dot{s}(t) \,, \tag{2.10}$$

The method can be extended to include individual hardware constraints per axis (*123*). Similarly, parallel transmission RF constraints can be included as a global upper-bound gradient limit in excitation k-space (*124*). In this manner, the algorithm can be utilised as a variable-rate selective excitation method (*125*, *126*).

Davids et al (*127*) have suggested an alternative approach to the time optimal gradient design problem. By considering an arbitrary k-space trajectory as a set of control points, the exact trajectory shape can then also be included as an additional degree of freedom, allowing a further reduction in overall duration. A separate step is therefore required to select suitable control points.

2.2.2 kT-Points

The kT-points trajectory (*128*) was developed as a 3D non-selective concatenated pulse that deposits RF energy at a small number of k-space locations. The additional spatial frequency information is utilised to improve transmit field uniformity over an extended volume, and can viewed as a generalisation of RF shimming. Parallel transmission is used in this context to reduce the number of k-space locations required for a given level of transmit field homogeneity, as compared to single-channel transmission.

Increasing the number of k-space locations visited on the trajectory increases the selectivity of the eventual pulse, in the so-called *Extended* kT-points method (*129*). It is also more efficient to transmit RF continuously, rather than at discrete locations. To determine the minimum set of k-space locations required for a given excitation error, an *Orthogonal Matching Pursuit* algorithm (*130*) has previously been suggested for multidimensional trajectory design (*131*). This is achieved by greedily selecting k-space locations from a dictionary, to reduce the excitation error at each iteration step. The algorithm is summarised in Algorithm 3.

Algorithm 3: Orthogonal Matching Pursuit
The Orthogonal Matching Pursuit algorithm greedily selects locations in excitation k-space to minimise the excitation error at each iteration step. Updated pTx pulses are generated during each iteration. Additional locations are chosen until the residual error meets a predefined tolerance.

$\hat{k} \leftarrow (0,0,0)$; *// initial point is centre of k-space*
$k_{\text{list}}.\text{insert}(\hat{k})$; *// update k-space list*
$F \leftarrow e^{i\hat{k}\mathbf{x}}$; *// initial phase encoding matrix*
while $||\hat{\mathbf{y}}|| > tolerance$ **do**
  $\hat{k} \leftarrow \underset{k\in\mathbb{R}^3}{\arg\min} ||\hat{\mathbf{y}} - e^{ik\mathbf{x}} \cdot \sum_{c=1}^{N_c} \mathbf{s}_c b_c)||, \quad \forall b_c;$ *// find next k-space location*
  $k_{\text{list}}.\text{insert}(\hat{k})$; *// update k-space list*
  $\mathbf{F} \leftarrow [\mathbf{F}, e^{i\hat{k}\mathbf{x}}]$; *// update encoding matrix*
  $\mathbf{b} \leftarrow \underset{b\in\mathbb{C}^{N_t N_c}}{\min} || \sum_{c=1}^{N_c} \mathbf{s}_c \mathbf{F}\mathbf{b}_c - \mathbf{m}_d||$; *// calculate new RF pulses*
  $\hat{\mathbf{y}} \leftarrow \sum_{c=1}^{N_c} \mathbf{s}_c \mathbf{F}\mathbf{b}_c - \mathbf{m}_d$; *// calculate new residual vector*
end
return k_{list} ; *// return list of k-space locations*

The final output of the Orthogonal Matching Pursuit algorithm is a list of k-space

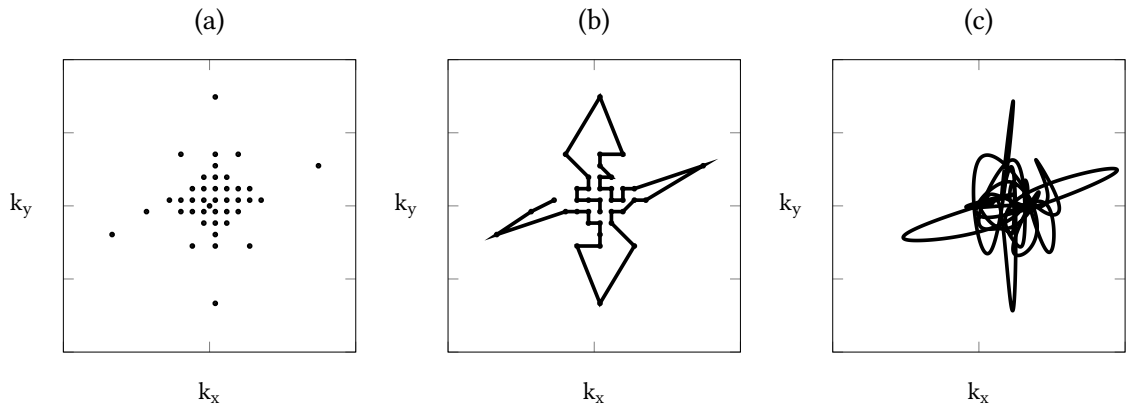

Figure 2.5: *k*T-Points trajectory design: disordered sparse *k*-space locations (a); minimum point-wise distance found using genetic algorithm (b); time optimal trajectory using the method described in (*122*, *123*), (c).

locations that is spatially disordered (*132*). Time ordering of these locations is an example of the well-known *travelling salesperson* problem, an NP-hard problem in combinatorial optimisation (*133*). In addition to finding the shortest path through *k*-space, it is useful to apply a boundary condition such that the *k*-space trajectory is refocused. The genetic algorithm (see Chapter 1), is well suited to these categories of optimisations. An example of *k*-space point ordering is given in Figure 2.5.

The time optimal gradient waveform design algorithm (*122*, *123*) can be used to design efficient waveforms subject to gradient system limits. One of the drawbacks of the existing Extended *k*T-Points method is that these limits are not considered during the selection of each subsequent *k*-space location.

2.2.3 Frequency Domain methods

Three other pTx pulse design methods which have been developed, are formulated in the *frequency domain* (*21*, *134*, *135*). The first of these, known as *Transmit SENSE* (*21*), pioneered the concept of parallel undersampling for transmission. The approach is roughly equivalent to the spatial-domain method in the absence of off-resonance, except multiplications in the spatial-domain take place as convolutions in the frequency domain.

A second approach (*134*), was envisaged as a transmit equivalent of the GRAPPA parallel imaging algorithm (*10*). The method determines the weighted superposition of pTx pulses transmitted for a single line of a transmit *k*-space, to excite a given magnetisation target. Unfortunately, the approach is limited to echo-planar trajectories only.

Finally, a recent suggestion (*135*), known as the *k-space domain* method, is a highly efficient technique to calculate a sparse approximation of the pseudo-inverse of the pTx system encoding matrix (see Equation 2.5). The basis of this is the identification of the fact that transmit sensitivities are sparse in the Fourier domain. Truncation of these maps can lead to computational efficiencies but at the cost of increased excitation errors.

None of these frequency domain methods have matched both the flexibility and performance of the spatial domain method.

2.2.4 Large-tip angle

The aforementioned methods are based on the Small-Tip Angle approximation. General pTx large-tip angle design is challenging due to the inherent nonlinearity in the Bloch Equation (see Equation 1.18).

The linear-class of large-tip angle k-space trajectories, have the property that the designed pulses can scale to any flip angle without nonlinear effects (*136*). The conditions on the k-space trajectory are that it should be inherently refocused, and that it can be decomposed into a sequence of inherently refocused sub-trajectories (e.g. a spiral). The mechanism equally applies for pTx, but the conditions for linearity are likely to be broken when used with high acceleration factors (*137*).

A general multidimensional method was introduced by Xu et al (*138*) using an optimal control framework, whereby the state variable is posed as the terminal transverse magnetisation and the pTx pulses act as control variables. The computational efficiency of the method is heavily dependent on the numerical algorithm and techniques used to approximate second-order derivatives (*139*, *140*). Grissom et al demonstrated a faster (and reduced memory) optimal control method (*141*) formulated in the spinor domain, rather than real space. All optimal control methods require repetitive evaluation of the Bloch Equations and at-least first-order derivatives. An alternative approach to optimal control is to use concatenated additive small-angle rotational pulses which reduces the number of required Bloch evaluations (*142*).

2.3 Frequency Modulated Pulse Design

Frequency modulated (FM) RF pulses involve temporal evolution of the carrier frequency, usually in addition to amplitude modulation. Often though, frequency modulation is performed as a phase modulation in hardware. FM pulses for NMR were developed during the 1980s when Ackerman et al (*143–145*) demonstrated the increased locally adjacent sensitivity of surface coils over traditional volumes coils. The advantage of FM pulses in this regard was a unique property to provide uniform flip angle despite a highly inhomogenous coil sensitivity (*146*).

Spatiotemporal encoding methods also utilise FM pulses to apply a quadratic spatial phase pattern, that in combination with gradient fields can move the resonance condition through space (*147*, *148*). Snyder et al demonstrated steering localised resonance through space along a spiral trajectory in the spatial domain in multiple shots, and this was later acheived in a single-shot by Jang et al (*149*) with parallel transmission. This can allow

local correction to large variations in both B_1^+ and ΔB_0 field conditions, at the expense of relatively high SAR requirements (*150*).

2.3.1 Description of FM pulses

As the carrier frequency is constantly changing throughout the duration of an FM pulse, the most convenient frame of reference $(\hat{x}', \hat{y}', \hat{z}')$, rotates at the angular velocity ω_{RF}, which is the instantaneous RF frequency. When the pulse is off-resonant ($\omega_{\text{RF}} \neq \omega_0$), there is a non-zero magnetic field component aligned with the static field with an angular velocity equal to $\Delta\omega(t) = \omega_0 - \omega_{\text{RF}}(t)$, in addition to a transverse component ω_1 (see Figure 2.6). The effective magnetic field (expressed as an angular velocity) in this frame of reference is given by:

$$\omega_{\text{eff}}(t) = \hat{x}'\omega_1(t) + \hat{z}'\Delta\omega(t)\,, \tag{2.11}$$

where the RF field has been assumed to be aligned along $\hat{x}'$ without loss of generality. The angle subtended between the effective magnetic field and the $\hat{z}'$ axis is:

$$\psi(t) = \tan^{-1}\left(\frac{\omega_1(t)}{\Delta\omega(t)}\right). \tag{2.12}$$

Applying the chain rule, the rate of change of this angle is:

$$\dot{\psi}(t) = \frac{\dot{\omega}_1(t)\Delta\omega(t) - \omega_1(t)\Delta\dot{\omega}(t)}{|\omega_{\text{eff}}(t)|^2}\,. \tag{2.13}$$

Following the analysis by Garwood et al (*151*), and defining a second frame of reference $(\hat{x}'', \hat{y}'', \hat{z}'')$, that rotates with the instantaneous angular velocity of the effective magnetic field $\omega_{\text{eff}}(t)$. In this second frame, the effective field is given by:

$$\tilde{\omega}_{\text{eff}}(t) = \hat{y}''\dot{\psi}(t) + \hat{z}''|\omega_{\text{eff}}(t)|\,. \tag{2.14}$$

The angular velocity $\dot{\psi}(t)$ appears to be the effect of a fictitious force in Equation 2.14, due to the rotation of ω_{eff} about the y' axis. The precession of a magnetisation vector initially parallel with z″ in the double-prime frame, will lead to a maximum angle of $2\eta(t)$ between the magnetisation vector and $\hat{z}''\omega_{\text{eff}}(t)$.

$$\eta(t) = \tan^{-1}\left(\frac{|\dot{\psi}(t)|}{|\omega_{\text{eff}}(t)|}\right) \approx \frac{|\dot{\psi}(t)|}{|\omega_{\text{eff}}(t)|}\,. \tag{2.15}$$

To ensure a magnetisation vector that is initially parallel with $\hat{z}'\omega_{\text{eff}}(t)$ will remain aligned with $\hat{z}'$, the angle $\eta(t)$ is the quantity to minimised as:

$$|\dot{\psi}(t)| \ll |\omega_{\text{eff}}(t)|\,. \tag{2.16}$$

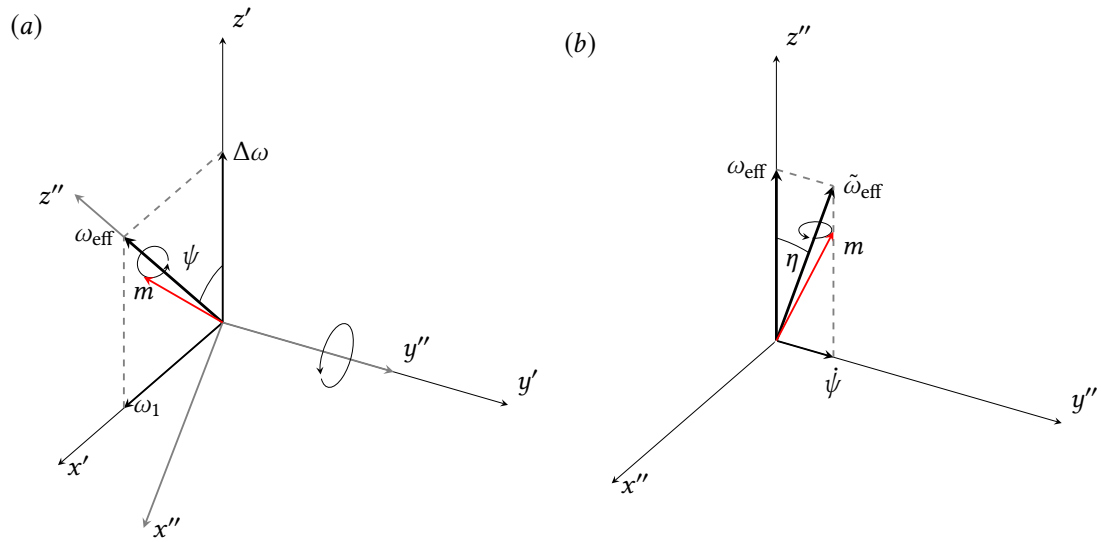

Figure 2.6: Two useful rotating frames: primed (the effective frame), and double primed. The primed frame rotates at the instantaneous RF frequency ω_{RF}, and exhibits an effective magnetic field $\omega_{\mathrm{eff}}(t)$ that is a combination of the transverse component $\omega_1(t)$ and a longitudinal component $\Delta\omega(t)$. The double primed frame rotates at an instantaneous frequency defined by the magnitude of the effective field in the primed frame. For a magnetisation vector initially parallel with the z' axis to remain aligned with the effective field it is required that $|\dot{\psi}(t)| \ll |\omega_{\mathrm{eff}}(t)|$. Adapted from (*151*)

The above statement is referred to as the *adiabatic condition* and ensures magnetisation vectors initially coaxial with the effective field, remain coaxial (or *spin-locked*) for the duration of the pulse. Likewise, components of a magnetisation vector initially in the plane perpendicular to the effective field, remain perpendicular for the duration of the pulse if the adiabatic condition is fulfilled.

2.3.2 Adiabatic Pulses

FM pulses are categorised into Adiabatic Half-Passage (excitation) (*152*, *153*) and Full-Passage (spin inversion) (*154*). For Full-Passage pulses, of particular note is the Hyperbolic Secant (HS) pulse, as introduced by Silver et al (*155*, *156*):

$$\omega_1(t) = \left(\omega_{\max}\mathrm{sech}\left(2\beta\left(t - \frac{T}{2}\right)\right)\right)^{1+i\mu}. \tag{2.17}$$

The Hyperbolic Secant pulse produces an analytical solution of the Bloch-Riccati Equation 1.17, which can be expressed in terms of hypergeometric functions and certain pulse parameters (*156*). The functional form of Equation 2.17 can be seen in Figure 2.7. In the limit $t \rightarrow T$, the value of the fractional longitudinal magnetisation for an on-resonant isochromat is given by:

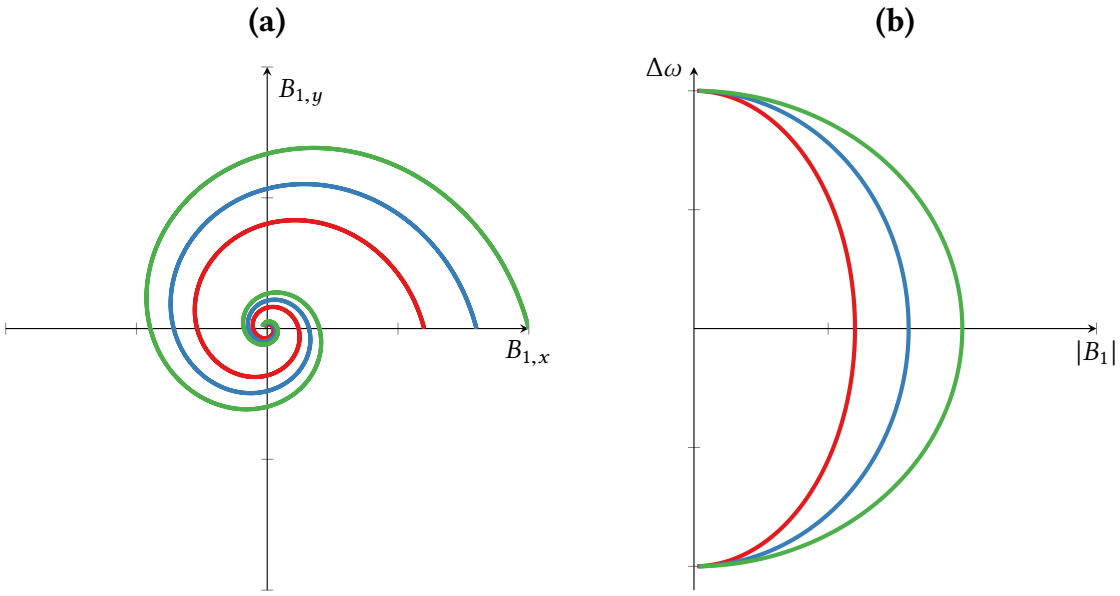

Figure 2.7: The hyperbolic secant pulse for three values (different colours) of $\omega_{\max}$ plotted on the phase-plane (a): the pulse starts at the origin and spirals out to $\omega_{\max}$, before reversing along the same path to the origin. Adiabatic Full-Passage pulses' insensitivity to $|B_1^+|$-value is explained when viewed in a sweep diagram (b) (*158*).

$$\left.\frac{m_z}{m_0}\right|_{t\to T} = \frac{1+\cosh\left(\frac{AT\pi}{2\beta}\sqrt{1-(\omega_{\max}/A)^2}\right)}{\cosh\left(\frac{AT\pi}{4\beta}\right)^2} - 1 \tag{2.18}$$

where A is proportional to the pulse bandwidth (*157*).

Since the introduction of the HS driving pulse, several optimisation schemes have been developed to further improve and tailor the performance of the pulse. Principle contributions have been the development of *numerically optimised modulations* (NOMs), that optimise the adiabatic factor of the generated pulse for a given range of ω_{max} values (*159–161*) by temporal resampling, or *offset-independent* adiabacity (*162*). Popular reshaping functions with few design parameters include the C-FOCI (*163*), and the generalised TR-FOCI (*164*) forms.

Adiabatic pulses have also been developed by extending some of the successful AM pulse design methods. An optimal control framework was also developed, minimising a novel adiabatic cost function in addition to the squared error cost of the final magnetisation state (*165*). The Shinnar–Le Roux algorithm has also recently been adapted for use with adiabatic pulses (*166*, *167*).

A hybrid approach can be envisaged utilising the unique properties of adiabatic pulses and the advantages of multiple transmit channels. RF shimming has been applied to reduce the peak channel power required to satisfy the adiabatic condition at multiple voxel locations using a minimax algorithm (*168*) or simulated annealing optimisation scheme (*169*). Deniz et al (*91*) constructed and numerically optimised a novel Rayleigh

quotient to maximise local B_1^+ per unit input power for use with adiabatic pulses. Subsequently, Carluccio et al (*92*) suggested an analytical method to maximise local transmit efficiency. Both of these methods utilise information from the transmit channel scattering matrix to evaluate power deposition.

2.4 Calibration-Free Parallel Transmission

The direct measurement of field maps and design time using some of the design methods described in the preceding sections inevitably add latency. Furthermore, field mapping is typically performed as an initial preparation step once at the beginning of a scanning session. It is therefore subject to change due to patient motion from both coherent (such as respiratory cycle) and incoherent (such as changes in head pose) sources (*170*, *171*).

Calibration free pTx envisages discarding some or all of these preparation and design stages. Gras et al introduced the notion of *universal pulses* (*172*), whereby the pulses were optimised over a population dataset for a variety of applications (*106*, *173*). Tomi-Tricot et al (*174*) extended this to train a machine learning classifier to select a single pTx k_T-points pulse for uniform abdominal excitation from a set of three, using simple information available from the survey-scan. Ianni et al (*175*) used ridge-regression to predict RF shim weights from head geometry alone; and improved prediction when including low-resolution B_1^+ maps.

Neural network approaches to learning RF weights and k-space sampling locations for large-tip angle spoke pulses were investigated by the author of this thesis (*176*) (see Figure 2.8). Unfortunately, the narrow range of the training data resulted in large excitation errors. Vinding et al (*177*) recently demonstrated a supervised deep network to classify 2D RF pulses with a fixed k-space trajectory, for single-transmission only. In this case, the training data was a library of magnetisation excitation targets derived from ImageNet (*178*), together with a RF pulse for each target designed by optimal control methods. No discussion was presented as to how this training set could be extended to include parallel transmission or arbitrary k-space trajectories.

The increased degrees of freedom available with multidimensional RF pulse design still makes this approaches attractive for *real-time* patient-specific optimisation. However, contributions and inferences deriving from machine learning algorithms will remain an active and important area of research.

2.5 Summary

The aim of this chapter was to introduce some of the important existing design algorithms and capabilities of parallel transmission and RF pulse design. In particular, the spatial domain method, orthogonal matching pursuit algorithm, and the time optimal

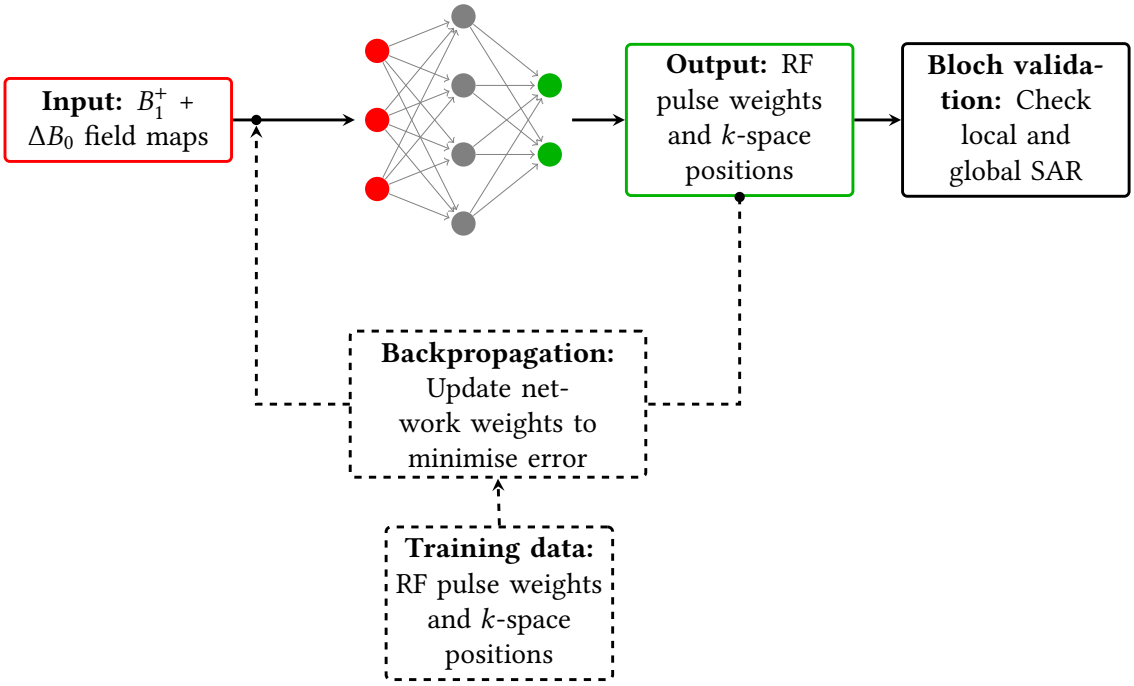

Figure 2.8: Neural network training using supervised learning. The training data consists of a library of patient field maps and a reference *gold standard* method to design RF pulses in each case. Backpropagation is used to minimise the difference in the error between the network and the optimal solution. The trained network can then be deployed for new input field maps as a simple feed-forward calculation. Adapted from (*176*)

gradient design algorithm will be utilised in Chapters 3 and 4. In addition, a derivation of the adiabatic condition for frequency modulated pulses was also given, which will be useful in Chapter 5.

Chapter 3

GPU Joint Design of 3D Excitation pTx pulses and k-space trajectories

In this chapter the design of short pTx pulses with optimised k-space trajectories to excite an arbitrary 3D volume is presented. The parameterisation of the excitation k-space trajectory under alternative linear transformations subject to realistic gradient performance is considered as part of this work. A Graphical Processing Unit (GPU) is used to reduce computational time of an existing *local* optimisation algorithm and a novel *global* optimisation algorithm.

- Initial results from this chapter formed the basis of a conference abstract (*179*)
- Code is published at https://github.com/cmirfin/Joint_Pulse_Design

3.1 Introduction

In certain MRI applications, increasing the degree of spatial encoding from signal readout to signal excitation can be beneficial. In particular, Inner-Volume Imaging (IVI) has been used to reduce acquisition times for a given resolution where a reduced field of view is acceptable for the application (*180*, *181*). In many cases, two-dimensional (2D) selective excitation is sufficient when a Cartesian readout scheme is employed, and the frequency-encoding direction can be bandwidth filtered. In other cases such as non-Cartesian excitation (*182*), or novel parcellated acquisition schemes (*183*), full three-dimensional (3D) spatial excitation is desirable. Parallel transmission is essential in this case to reduce pulse lengths to acceptable durations by providing a degree of spatial encoding in addition to the application of gradients (*21*).

The joint-optimisation of both the RF waveforms and the excitation k-space trajectory has been shown to reduce the excitation error in comparison to optimising RF pulses alone for predefined gradient waveforms (*184*). The motivation is to exploit subject-specific information (both coil sensitivity and off-resonance maps), as well as the desired magnetisation target. Existing methods parameterise the gradient waveforms or equivalently the k-space trajectory using either a set of basis functions, such as splines (*185*)

or sinusoids (*186*), or control points that are connected using the k-space trajectory that can most rapidly be traversed (*187*). Ideally the maximal degree of parameter compression is sought. However an optimal parameterisation for achieving this is unclear, as it must trade off computational time per parameter with the ability to generate refined waveforms. Additionally, the realised gradient waveforms must abide by amplitude and slew-rate constraints.

In the Small-Tip Angle (STA) approximation (*136*), the design of pTx RF pulses for a given excitation k-space trajectory is greatly simplified and can be described as a linear system. Problems can therefore usually be solved with iterative methods such as Conjugate Gradient Least Squares (CGLS) (*108*). This problem is well suited to acceleration by a GPU, whereby numerical operations are run concurrently on potentially thousands of processors. For example, an order-of-magnitude speedup using this approach was reported for a 64×64 grid with 8-transmit channels (*188*).

Conversely, designing the excitation k-space trajectory presents a constrained non-convex optimisation problem. Previous methods (*185*, *187*) typically use local solvers (i.e. they are dependent on the derivatives of the cost function at a given point) that converge to local minima and are hence very sensitive to the initial conditions. Recently, an iterated local search was suggested to reduce the effect of initial conditions, and implemented on a multicore CPU (*189*). Typically global optimisation methods such as simulated annealing or local solvers with multiple initial conditions are too slow to run on CPUs.

As discussed in Chapter 1, the genetic algorithm GA is a different class of metaheuristic optimisation algorithm motivated by the biological process of natural selection. Due to the stochastic nature of genetic algorithms they are beneficial in avoiding local minima, although they can require significantly more cost function evaluations to run than a traditional local solver. Genetic algorithms have previously been applied in the field of MR pulse design for the design of inversion pulses (*164*), yielding improved results over local based methods. Genetic algorithms have not been used in the context of pTx pulse design due to the inherently longer computational time that they require in comparison to derivative-based methods.

3.2 Theory

3.2.1 Problem Formulation

The joint-optimisation of pTx waveforms ($b \in \mathbb{C}^{N_t N_c}$) and k-space trajectory (represented by the set of parameters $\rho^x, \rho^y, \rho^z \in \mathbb{R}^{N_\rho}$) is formalised as:

$$\begin{aligned} \min_{\substack{b\in\mathbb{C}^{N_t N_c} \\ \rho^x,\rho^y,\rho^z\in\mathbb{R}^{N_\rho}}} \quad & \Big\{ || \sum_{c}^{N_c} S_c F(\rho^x, \rho^y, \rho^z) b_c - m_{\text{target}} ||_{\mathbf{w}}^2 + R(b) \Big\} \\ \text{s.t.} \quad & |\Delta^{(1)} k_j^x| \le \gamma G_{\max} \Delta t, \\ & |\Delta^{(1)} k_j^y| \le \gamma G_{\max} \Delta t, \\ & |\Delta^{(1)} k_j^z| \le \gamma G_{\max} \Delta t, \\ & |\Delta^{(2)} k_j^x| \le \gamma S_{\max} \Delta t^2, \\ & |\Delta^{(2)} k_j^y| \le \gamma S_{\max} \Delta t^2, \\ & |\Delta^{(2)} k_j^z| \le \gamma S_{\max} \Delta t^2, \quad \forall j \in \{1, \dots, N_t\}\,, \end{aligned} \tag{3.1}$$

where $F \in \mathbb{C}^{N_s, N_t}$ represents the system encoding matrix with a sampling interval of Δt for a set of k-space basis coefficients parameters (ρ^x, ρ^y, ρ^z), and off-resonance effects. S_c represents the transmit sensitivity and $b_c \in \mathbb{C}^{N_t}$ the RF waveform for the cth channel respectively ($c \in \{1, \dots, N_c\}$) . The vectorised transverse magnetisation target m_{target} consists of N_s spatial voxels. The regularisation term $R(b)$ can be used to balance global Specific Absorption Rate (SAR), local SAR (*113*), or for our purposes, simply the total integrated transmit power $R(b) = \lambda \sum_c^{N_c} ||b_c||^2$. $\Delta^{(1)}$ and $\Delta^{(2)}$ symbolise the first and second-order difference operators. The maximum gradient amplitude and slew rates are denoted by $G_{\max}$ and $S_{\max}$ respectively.

3.2.2 k-space trajectory parameterisations

With a typical sampling interval on the order of microseconds, the number of time points N_t is of the order of 10^3. Discretising the k-space trajectory at this resolution is unnecessary due to the implicit smoothness requirement in equation (3.1) and is undesirable since it requires extended computational times. Instead we follow an existing approach and parameterise the k-space trajectory as a linear transformation of N_ρ parameters per axis, where $N_\rho \ll N_t$ (*185*, *190*). In reality, the gradients realised by the scanner hardware can be significantly distorted due to eddy-currents and amplifier delays. As discussed in Chapter 1 these deformations can be well approximated by treating the system as a Linear Time-Invariant (LTI) model (*37*) and thus we convolve each basis function $g_k(t)$ with

the Gradient Impulse Response Function (GIRF), $H(t)$ for each gradient axis (x, y, z).

$$\begin{aligned} G_x(t) &= \sum_{n=1}^{N_\rho} \rho_n^x \Big(H^x(t) * g_n(t) \Big) = \sum_{n=1}^{N_\rho} \rho_n^x \; \tilde{g}_n^x(t) \\ G_y(t) &= \sum_{n=1}^{N_\rho} \rho_n^y \Big(H^y(t) * g_n(t) \Big) = \sum_{n=1}^{N_\rho} \rho_n^y \; \tilde{g}_n^y(t) \\ G_z(t) &= \sum_{n=1}^{N_\rho} \rho_n^z \Big(H^z(t) * g_n(t) \Big) = \sum_{n=1}^{N_\rho} \rho_n^z \; \tilde{g}_n^z(t) \; . \end{aligned} \tag{3.2}$$

Therefore, the predicted k-space trajectory in the x-direction is given by:

$$k^x(t) = -\gamma \sum_{n=1}^{N_\rho} \rho_n^x \int_t^T \tilde{g}_n^x(\tau) \, d\tau \; , \tag{3.3}$$

where the summation has been taken outside of the integral (similar expressions exist for y and z gradient axes). After discretisation, Equation (3.3) can be viewed as a linear transform with modified basis functions. The GIRF needs to be measured once using either a field camera with NMR-active probes (*191*) or, alternatively, an image based approach (*192*, *193*).

The best choice of linear transformation is unclear and will likely depend on the target excitation pattern. In this work we consider three parameterisations of the k-space trajectory: (a) a piecewise second-order B-spline with uniform spacing as suggested in (*185*); (b) a piecewise second-order B-spline with non-uniform spacing; (c) a discrete sine transformation motivated by the necessity to compress the number of free parameters. In the special case of second-order B-splines, the function derivatives are piecewise linear, reducing the number of gradient constraint checks required. We introduce non-uniform piecewise spacing (see Figure 3.1), using a logarithmic function motivated by the knowledge that the highest fidelity trajectory will be required at the centre of k-space. By construction, each of the parameterisations lead to k-space trajectories that finish at the centre of k-space, making them suitable for excitation.

3.3 Methods

The optimisation of Equation (3.1) is implemented as an alternate minimisation consisting of an *outer-loop* for k-space parameters and a Conjugate Gradient Least Squares (CGLS) *inner-loop* for the design of pTx RF pulses. A GPU is used to accelerate elements of code where appropriate.

To minimise data transfer across the slow CPU-GPU bridge, all B_1^+, B_0 maps, magnetisation targets, and transformation matrices reside on the GPU in global device memory

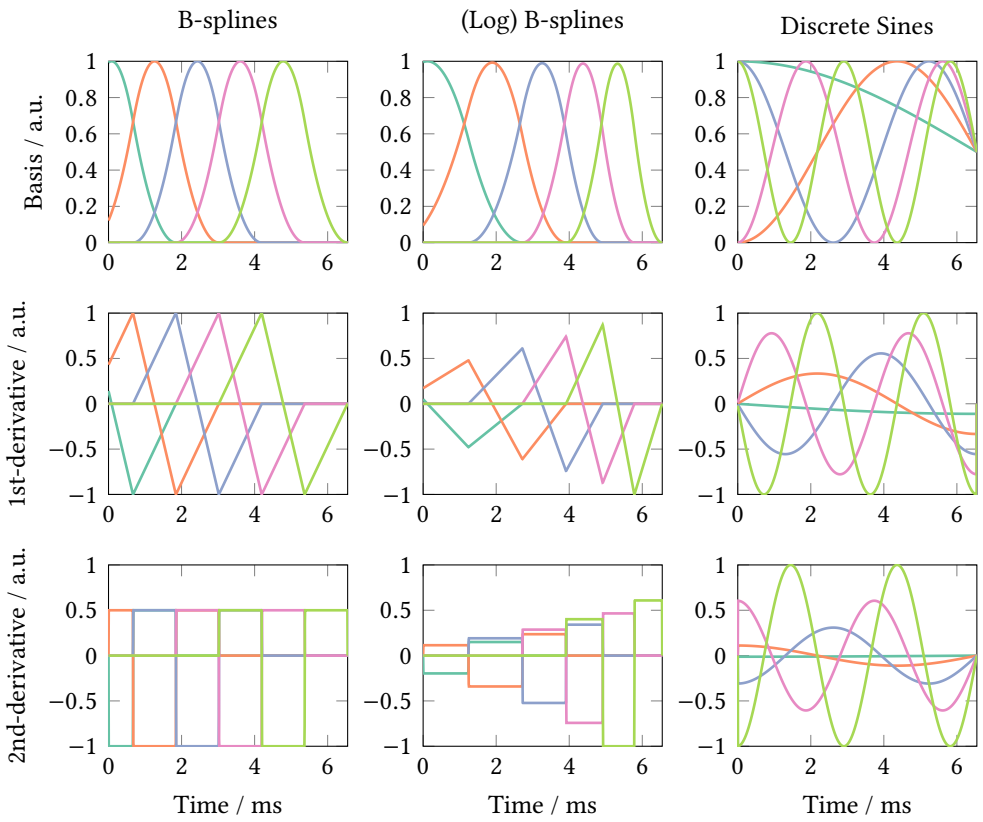

Figure 3.1: Trajectory parameterisations with their first and second derivatives. 2nd-order B-splines with uniform spacing (adapted from (*185*)) (a); 2nd-order B-splines with non-uniform spacing (b); Discrete Sines (c) for $N_p = 5$ parameters. The colours here are used to differentiate each individual localised spline curve.

for the entire duration of the joint-optimisation.

3.3.1 pTx RF waveform optimisation (*inner-loop*)

As discussed in Chapter 2, the *inner-loop* optimisation of short 3D pTx RF waveforms is an over-determined sparse linear system. A CGLS algorithm was implemented using high level MATLAB constructs directly on the GPU. The CGLS algorithm is dominated by two repeated matrix multiplications $m = Ab$ and $b = A^H p$, which are well suited for parallel operations (*188*). Only the residual cost is returned to the host CPU. For our purposes, a fixed-value of the power regularisation parameter $\lambda = 8$ proved sufficient. A maximum of 20 conjugate gradient steps per loop was allowed.

3.3.2 k-space optimisation (*outer-loop*)

In contrast to the *inner-loop*, the *outer-loop* optimisation of the k-space trajectory is a non-convex problem with multiple hard constraints. For this, the use of a *local* derivative-based iterative solver and a *global metaheuristic* solver were investigated.

The *interior-point* algorithm (see Chapter 1) (*86*, *194*), used the method available

within the MATLAB Optimisation toolbox (R2018b, The MathWorks, Inc., Natick, Massachusetts, USA). Analytical gradient and (approximate) analytical Hessian (as described in (*185*)) functions were calculated explicitly, making use of sparse linear algebra where permissible. After each outer-loop update, the new k-space parameters were transferred to the GPU and uncompressed with the appropriate transformation. The normalised system matrix A was then calculated directly, including an off-resonance term.

The genetic algorithm, as provided within the MATLAB Global Optimisation toolbox was used as an alternative search algorithm. Genetic operations were performed strictly on the host-CPU, with only scalar values returned from the cost evaluation of each new k-space trajectory from the GPU. Although parallel genetic algorithms are becoming more prevalent (*195*), it is not feasible to execute a joint-optimisation algorithm of this size exclusively on a GPU due to memory restrictions. The crossover and elitism fractions (as discussed in Chapter 1, see Algorithm 2) were set at 80% and 5% respectively.

For both algorithms, linear slew-rate constraints were enforced with a maximum of 180 $\mathrm{Tm}^{-1}\mathrm{s}^{-1}$ for each axis. Initial studies found it was unnecessary to apply gradient amplitude constraints due to the low resolution of the excitation targets.

An NVIDIA Titan X Pascal GPU was used for all work consisting of 3584 thread processors (arranged as 28×128 streaming multiprocessors) with 12 Gb of global device memory. The host-CPU was a Lenovo P510 desktop with 32 Gb of RAM.

3.3.3 Initial k-space trajectory population

Figure 3.2 shows the initial population of k-space trajectories, which consisted of variants of five existing *classic* trajectories and a subject-specific trajectory. These were: stack-of-spirals (SoS) (*196*); sampling on concentric spheres (Shells) (*197*); echo-volumar (EVI) (*198*); a cross trajectory (*187*); extended spiral-nonselective (SPINS) (*199*). Each trajectory was also rotated around each axis in 45° steps to yielding of 5×4^3 trajectories. In addition, an extended kT-points (*131*, *185*), (an extension to the original kT-points (*200*)), where phase locations are greedily chosen (and optimally ordered) such that the lowest residual excitation error is achieved was implemented (see Chapter 2). In total 321 trajectories were selected as separate initialisations for the interior-point method, and the total population for the genetic algorithm.

Gradient waveforms for each of the trajectories were generated using the time-optimal method (*122*) and accordingly adjusted to a duration of 6.6 ms with a dwell-time of 6.4 μs. A least-squares fit was used to find the necessary parameter values, which represent each initial trajectory for a given transformation. In addition, the number of free parameters per axis was varied in the range 16 to 128 to observe the change in residual fitting error.

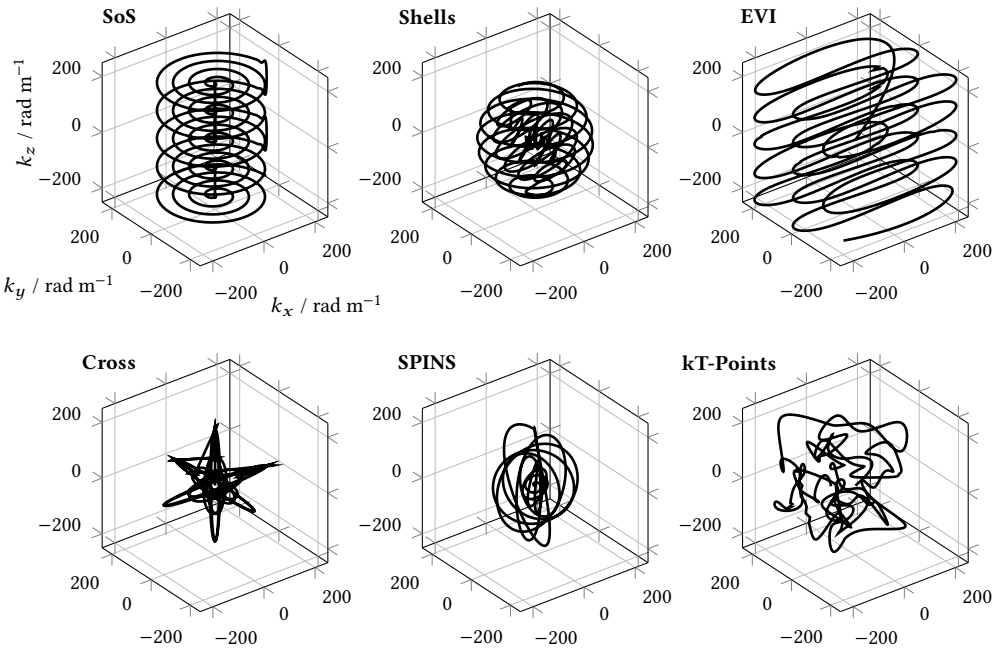

Figure 3.2: Examples of six initial *classic* 3D k-space trajectories, each with duration of 6.6 ms and subject to a maximum slew rate of 180 Tm^{-1}s^{-1}.

3.3.4 GIRF measurement

A dedicated field monitoring system (Skope, Zurich, Switzerland) consisted of 16 NMR active Fluorine probes each with an acquisition bandwidth of 1 MHz. The probes were evenly distributed on the surface of a cylinder closed at one end. The probe positions were accurately mapped by measuring frequency offsets under constant gradient fields of 2.5 mT m $^{-1}$ applied consecutively along each axis (*35*).

A series of four linear chirp waveforms, each of length 10 ms and bandwidths of 10, 20, 30 and 40 kHz respectively were executed on each gradient axis. These waveforms were shared by the author in (*201*). The effective T_2^* lifetime of the NMR probes determined a measurement window of 20 ms, equating to a frequency resolution of 50 Hz. For each gradient channel, a total of 100 averages were performed for each waveform which comprised of 50 positive and 50 negative linear frequency sweeps to ensure sufficient signal at each frequency within the range. The measurements were combined using a weighted sum.

3.3.5 Experiments

All experiments were performed using a Philips 7 T Achieva (Best, Netherlands) scanner with an 8-channel transmit, 32-channel receive head-only coil (Nova Medical, Massachusetts, USA). Multichannel transmit sensitivities were derived from a relative combination of a quadrature-mode sensitivity map and 15° flip angle gradient-echo images

from each channel. The quadrature-mode sensitivity map was measured using the AFI method (*77*) with a flip angle of 60° and repetition times of 30 ms and 150 ms. Off-resonance maps were acquired using the dual-TE method with echo-times of 20 ms and 21 ms. All field maps covered the whole brain and were acquired at a resolution of 4×4×8 mm^3 with 15 axial slices, before being interpolated to an isotropic voxel size of 5.0 mm.

A cubical magnetisation target with side 5 cm, smoothed with a 3D Gaussian filter of width 0.5 cm, was used for all experiments. In each case, an optimised 3D k-space trajectory and 8-channel pTx RF waveforms of duration 6.6 ms with a dwell-time of 6.4 μs were desired. From pilot work, a total of 64 parameters per axis were used to represent the k-space trajectories. Using the optimisation methods discussed in the previous section, a simulation of excitation error was evaluated over five subjects, three parameterisations and 321 initialisations.

A cross-comparison between each of the optimised k-space trajectories was performed using the discrete Fréchet distance,

$$F(k_a(t), k_b(t)) = \inf_{\alpha,\beta} \sup_{t\in 0,T} ||k_a(\alpha(t)) - k_b(\beta(t))|| \tag{3.4}$$

over the duration T, with the non-decreasing parameters $\alpha, \beta \in [0, T]$ for two trajectories $k_a, k_b : \mathbb{R} \rightarrow \mathbb{R}^3$.

Furthermore, a second cross-comparison between the optimised k-space trajectories was evaluated using the Normalised Root Mean Squared Error (NRMSE) as defined in Chapter 2. The RF waveforms were re-optimised in each case, to observe how well the optimised k-space trajectories would generalise across subjects.

Separately, a comparison of the execution time for the *inner-loop* and *outer-loop* optimisation on both the GPU and CPU was made by varying the design grid size (up to the GPU memory limit of 10^4 voxels) and averaging over 100 runs.

Second-order B_0 field shimming was performed prior to joint-optimisation over the whole brain. Imaging experiments were performed using a 3D gradient-echo with frequency-encoding in the foot-head direction with (15 slices, 10° flip angle, TE = 1.58 ms, TR = 100 ms, voxel size 3 × 3 × 8 mm), with and without optimised 3D excitation. The Shells trajectory (*197*) was used as the initialisation for the joint-optimisation for *online* optimisation. All data handling, ΔB_0 shimming, and pulse design was handled via custom in-house software (see Appendix 1 for details).

3.4 Results

Figure 3.3 shows the decrease in residual fitting error as a function of the number of parameters per axis. Similar performance between the three parameterisations was observed, although logarithmically-spaced B-splines displayed higher fitting error when

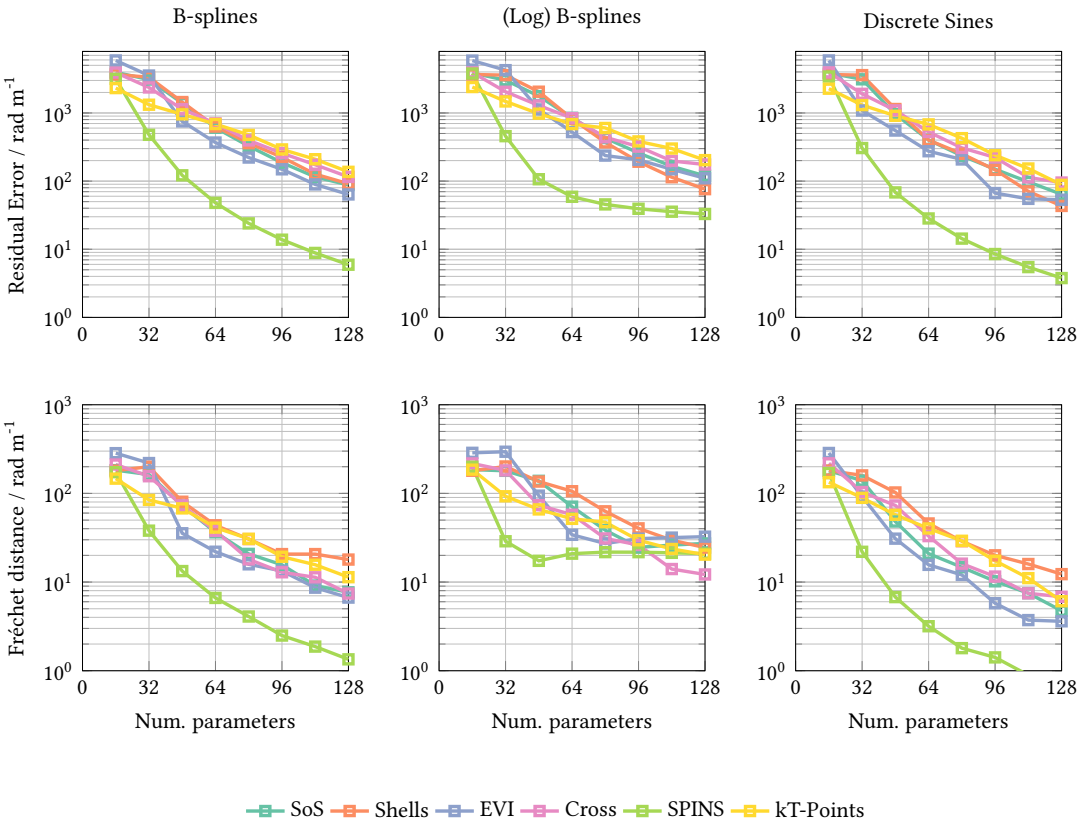

Figure 3.3: Measures of parameterisation error as a function of a number of parameter values per axis for each of the three parameterisations considered: B-splines, (Log) B-splines, and Discrete Sines.

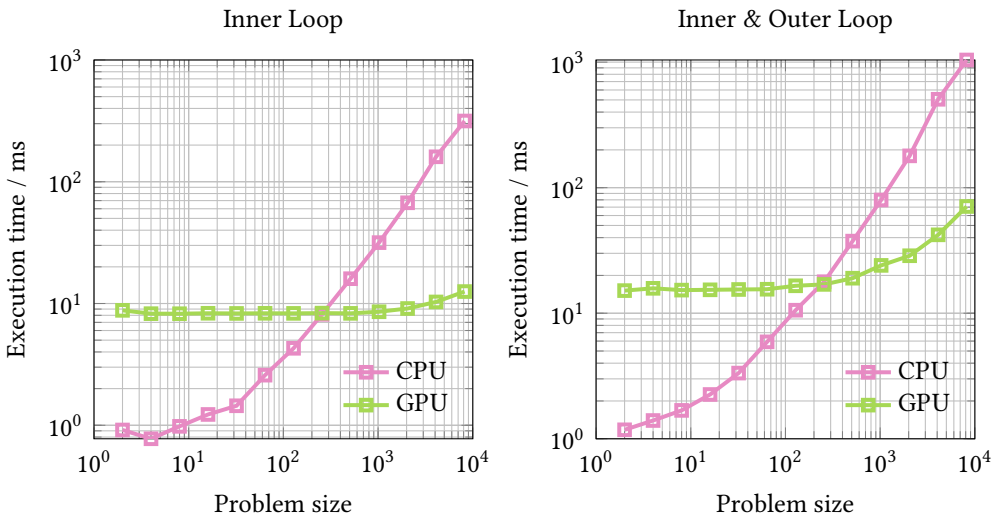

Figure 3.4: Execution time for inner-loop and outer-loop (interior-point) optimisations, as a function of problem size using CPU and GPU implementations, when using 64 parameters per axis.

using a large (≥ 64) number of parameters. A value of 64 parameters per axis was chosen for the rest of the work in this chapter as a trade-off between trajectory fidelity and computational complexity.

3.4.1 Computational Time

Figure 3.4 displays the execution times for one iteration of the inner and outer loop optimisations as a function of problem size. Although, the GPU implementation is slower than the CPU implementation for problems with fewer than 100 voxels, its relative performance improves for large problems where it is up to $\sim$ 50 times faster for 10^4 voxels. For problems with fewer than 100 voxels, the GPU implementation is actually slower than the CPU implementation due to the latency of transferring data across the slow GPU-CPU connection.

The genetic algorithm was automatically terminated after one-hour of computation in each case. Visualisation of the cost curves confirms that this was justified, with the mean population cost approaching the minimum population cost (see Figure 3.5). For the interior-point algorithm, 100 iterations were requested, resulting in a computational time of less than two-minutes per initialisation.

3.4.2 Offline optimisation

Optimisation reduced the excitation error in all cases. Table 3.1 summarises the excitation error and predicted RF power; a plot of these results for the Discrete Sine transformation is shown in Figure 3.6. Under the interior-point algorithm, there was a mean 22% $\pm$ 6% reduction in NRMSE across all initialisations. Total and peak power were also

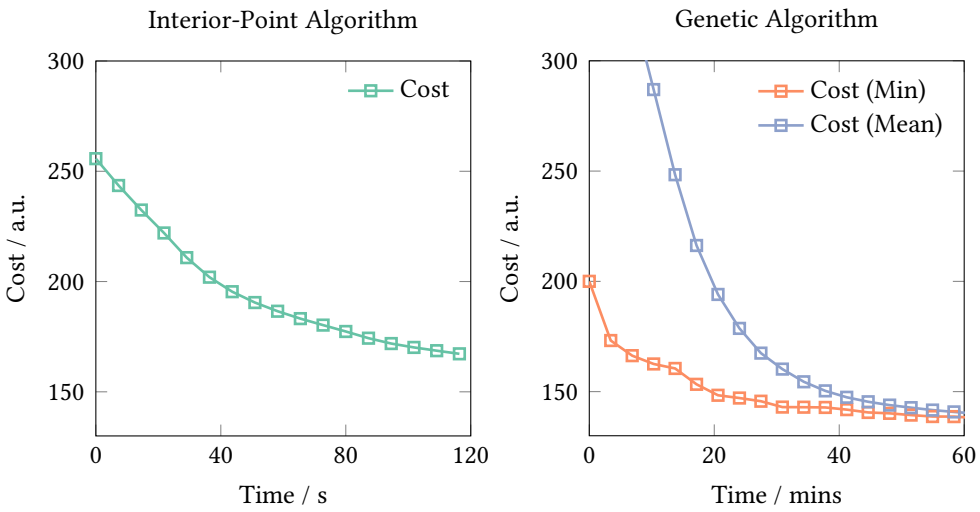

Figure 3.5: Cost curves for objective function stated in Equation 3.1 for GPU accelerated Interior-Point Algorithm (a) and genetic algorithm (b).

reduced by 24% ± 8% and 24% ± 7%, respectively.

There was negligible difference in the results between the three parameterisations using the interior-point algorithm. In contrast, the Discrete Sine transform produced results with lower excitation error across all subjects as compared to the B-spline parameterisations using the genetic algorithm. This coincided with elevated total and peak RF power.

A cross-comparison of excitation error for each of the optimised k-space trajectories across five subjects is shown in Figure 3.8, where the pTx RF waveforms were re-optimised for each subject and trajectory.

The Fréchet distance between optimised k-space trajectories is also shown as a cross-comparison matrix in Figure 3.9. Optimised variants of the five *classic* trajectories (SoS, Shells, EVI, Cross, and SPINS), have a Fréchet distance of less than 150 rad m^{-1}. For an individual subject, the optimised trajectories are clearly distinct, with Fréchet distances greater than 300 rad m^{-1}. The genetic algorithm produced distinct trajectories for each subject with Fréchet distances in the range 208 : 290 rad m^{-1} between subjects. Despite this, a cross-comparison suggests that these trajectories would be effective across subjects, albeit with an elevated excitation error (see Figure 3.8). Indeed, the intra-subject NRMSE = 0.23 ± 0.01, while inter-subject NRMSE = 0.25 ± 0.02.

3.4.3 Online optimisation

Example whole head in vivo images for a basic 3D GRE sequence are shown in Figure 3.11. Inner-volume excitation is demonstrated, but excited signal is clearly visible in the outer-volume when the images are windowed below 20%.

Figure 3.10 summarises the corresponding result for a spherical phantom. The corresponding off-resonance map displays local variations in field in the range -250 to 250

		NRMSE / a.u.		Total Power $/(\mu T)^2$		Peak Power $/\mu T$	
Trajectory	Parameterisation	Before	After	Before	After	Before	After
SoS	Spline	0.32 ± 0.02	0.27 ± 0.02	30.4 ± 14.4	21.6 ± 10.0	14.5 ± 6.1	9.2 ± 3.4
	(Log)-Spline	0.32 ± 0.02	0.26 ± 0.01	26.8 ± 12.3	18.7 ± 8.5	11.4 ± 4.4	8.1 ± 3.1
	DST	0.31 ± 0.02	0.26 ± 0.01	26.9 ± 12.8	22.0 ± 10.4	10.9 ± 4.4	9.2 ± 3.9
Shells	Spline	0.35 ± 0.03	0.29 ± 0.01	12.1 ± 5.9	11.4 ± 5.3	7.3 ± 3.1	6.0 ± 2.3
	(Log)-Spline	0.35 ± 0.03	0.29 ± 0.01	12.5 ± 6.2	10.9 ± 5.2	6.7 ± 2.4	4.7 ± 1.5
	DST	0.34 ± 0.03	0.29 ± 0.02	12.0 ± 5.8	11.3 ± 5.5	5.5 ± 1.9	5.5 ± 2.3
EVI	Spline	0.33 ± 0.02	0.31 ± 0.03	49.9 ± 20.8	43.5 ± 17.6	18.1 ± 4.2	16.5 ± 4.1
	(Log)-Spline	0.34 ± 0.03	0.30 ± 0.03	48.0 ± 20.0	38.6 ± 15.3	17.8 ± 5.2	15.0 ± 4.5
	DST	0.33 ± 0.02	0.30 ± 0.03	49.3 ± 20.4	40.5 ± 16.6	18.9 ± 4.9	16.1 ± 5.2
Cross	Spline	0.40 ± 0.03	0.31 ± 0.02	17.5 ± 8.3	11.4 ± 5.2	9.1 ± 4.0	7.1 ± 2.5
	(Log)-Spline	0.40 ± 0.03	0.32 ± 0.03	16.2 ± 7.6	12.1 ± 5.8	10.1 ± 4.3	7.2 ± 2.8
	DST	0.40 ± 0.03	0.32 ± 0.03	15.7 ± 7.4	11.3 ± 5.3	9.1 ± 3.9	6.6 ± 2.4
SPINS	Spline	0.42 ± 0.04	0.30 ± 0.01	13.1 ± 7.0	10.7 ± 4.8	5.8 ± 2.1	4.4 ± 1.2
	(Log)-Spline	0.42 ± 0.04	0.30 ± 0.01	13.1 ± 7.0	10.5 ± 4.8	5.8 ± 2.1	4.2 ± 1.2
	DST	0.42 ± 0.04	0.31 ± 0.01	13.1 ± 6.9	10.4 ± 4.8	5.7 ± 2.0	4.3 ± 1.3
kT-Points	Spline	0.43 ± 0.05	0.33 ± 0.02	30.6 ± 14.9	18.0 ± 7.7	10.2 ± 3.2	6.6 ± 1.3
	(Log)-Spline	0.43 ± 0.05	0.33 ± 0.01	29.0 ± 14.2	19.9 ± 8.5	10.9 ± 3.5	8.3 ± 2.2
	DST	0.42 ± 0.04	0.35 ± 0.04	28.7 ± 12.3	23.5 ± 9.2	9.8 ± 3.0	8.5 ± 2.5
GA	Spline	0.40 ± 0.05	0.27 ± 0.01	26.1 ± 17.4	11.2 ± 4.5	11.4 ± 6.2	5.5 ± 2.0
	(Log)-Spline	0.40 ± 0.05	0.29 ± 0.03	26.4 ± 17.9	13.5 ± 7.3	11.5 ± 6.3	6.1 ± 2.2
	DST	0.40 ± 0.05	0.23 ± 0.01	26.0 ± 17.7	22.2 ± 16.0	11.0 ± 6.1	9.6 ± 6.2

Table 3.1: Excitation error (NRMSE), total power, and peak power; before and after optimisation using Interior-Point algorithm with each of the three parameterisations for each initial trajectory. Results show mean average ± standard deviation across five subjects. Comparative results for the genetic algorithm are also shown.

Hz after second-order shimming. These field variations visually correlate with areas of signficant outer-volume excitation, despite the inclusion of the off-resonance map in the design process.

3.4.4 Measured GIRF

Figure 3.13 displays the measured GIRF for the 7 T Achieva gradient system. This result is in agreement with that reported for similar hardware at other sites (*201*). Zero phase at 0 kHz suggests the system is well calibrated with no net delay. The measured GIRF can be physically interpreted as a low-pass filter on input signals. The effect of this can be seen in Figure 3.14, where gradient waveforms undershoot their desired peak, consequently distorting the expected k-space trajectory.

3.5 Discussion

The pulse duration used in this work was predefined as 6.6 ms, to be comparable with that of 2D excitation. The resulting trajectories are highly-undersampled. Joint-optimisation of the trajectory and RF waveforms allows input of the target excitation, sensitivity maps, and off-resonance maps to ameliorate the effect of this undersampling. Ultimately, due to slew-rate constraints, trajectories of this duration will inevitably lead to some excitation error.

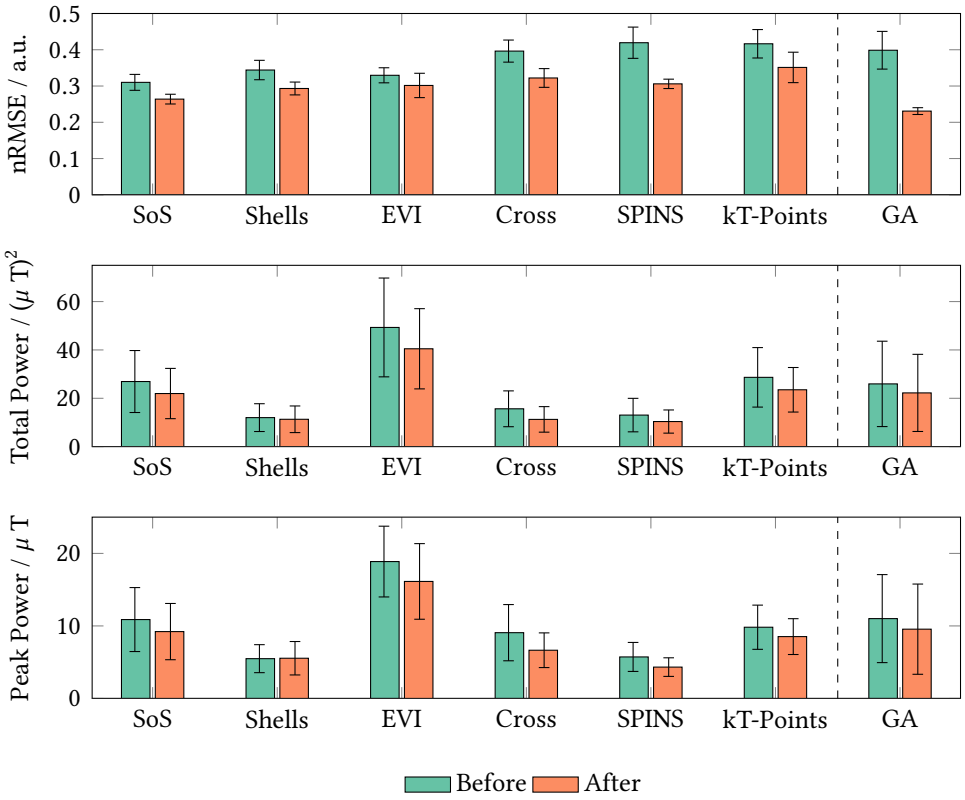

Figure 3.6: Excitation error (NRMSE), total power, and peak power; before and after optimisation using Interior-Point algorithm for 2nd-order B-spline parameterisation. Results show mean average ± standard deviation across five subjects. Comparative results for the genetic algorithm are also shown.

The objective function in Equation 3.1 specifies the general case where the target excitation profile is complex-valued. In applications such as GRE sequences where some spatial phase variation is acceptable, a Magnitude Least Squares (MLS) optimisation would be beneficial in place of our inner-loop. For Fast Spin-Echo sequences, Malik et al (*202*) demonstrated an additional CPMG based signal suppression mechanism, by biasing the optimisation of the imaginary component of the excitation profile.

For Inner-Volume Imaging, any outer-volume signal will alias in the phase-encoding directions, degrading the signal-to-noise ratio. Spatial weighting of the outer-volume could be used to reduce this, at the expense of inner-volume excitation error. Additionally, switching to a 2D excitation and band-limiting the frequency-encoding imaging direction would produce lower excitation error. These techniques will be considered further in Chapter 4.

The use of GPUs was vital to reduce the computational-time for the optimisation methods used in this work. Indeed, in the case of the interior-point method, this allows for *online* use within an acceptable time window. One drawback at present is the limited device memory available on current GPU hardware. This led to the use of a spatial design grid with large voxel dimensions of 5 mm^3, which consequently led to a

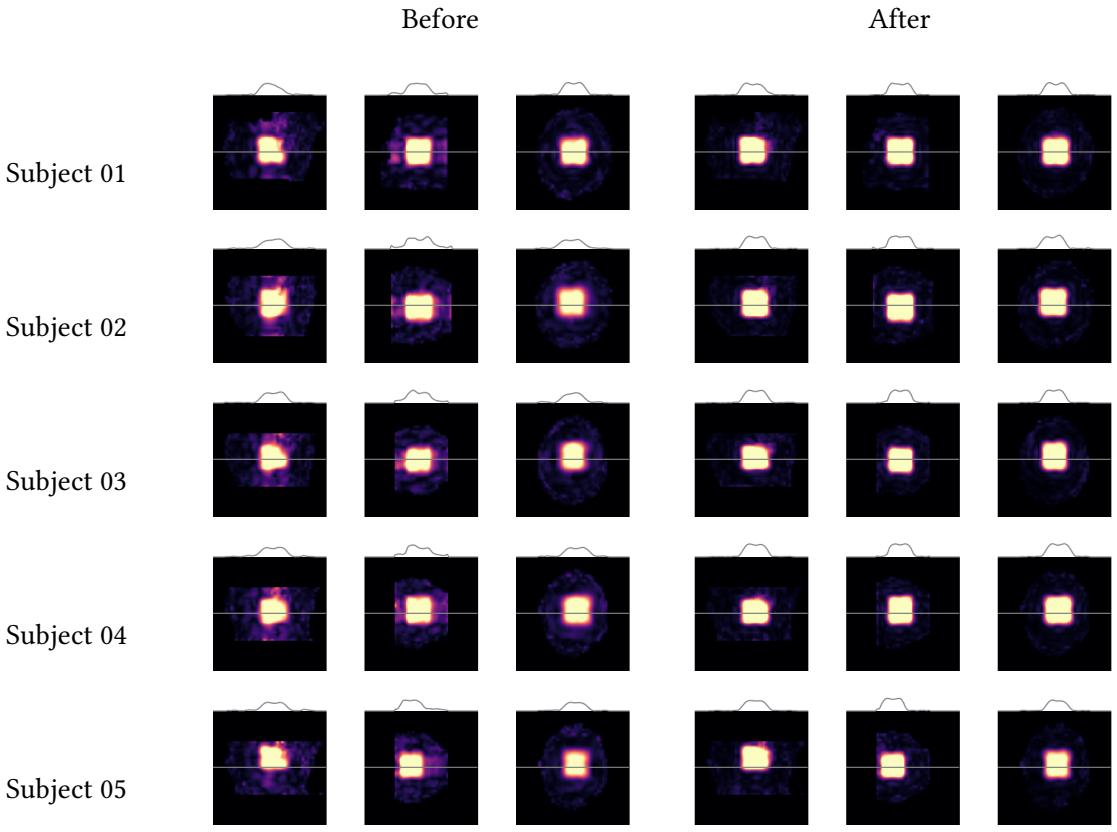

Figure 3.7: Bloch simulations of 3D excitation profiles (10° flip angle) shown for central sagittal, coronal and axial slices, before and after optimisation using interior-point method with *stack-of-spirals* initialisation. Line profiles across slice through centre of region-of-excitation shown above each image.

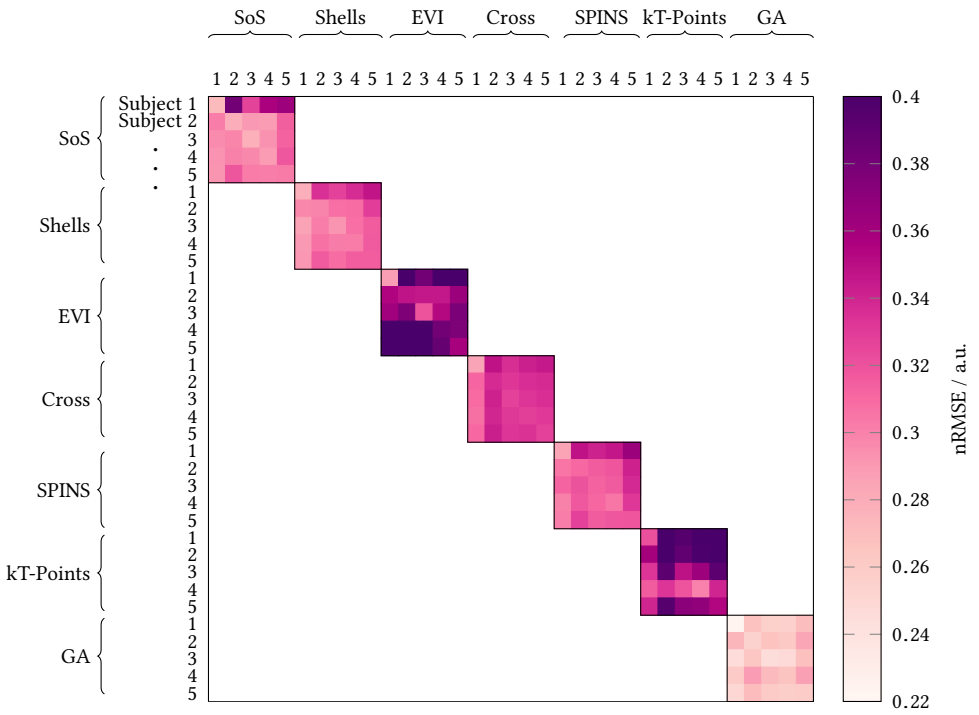

Figure 3.8: Cross-comparison of excitation error (NRMSE) between optimised k-space trajectories across six initialisations and five subjects. Each of these trajectories was re-optimised on a per subject basis before being applied across each subject. In each case, RF waveforms were re-optimised for each individual subject for a given trajectory.

loss of information contained within the measured field maps. In addition, the design grid was only considered on-resonance, when ideally a *spectro-spatial* solution would be sought (*187*). Improvements in GPU hardware are ongoing, with developments in joint GPU/CPU memory. An immediate solution would be to use multiple GPU cards with the design grid split across each card (*188*).

Although there was very little to distinguish between the performance of the three parameterisations in this work, it is important to note that only one target excitation pattern was considered. It would be interesting to observe how well the parameterisations here perform across a variety of excitation targets. As previously noted, second-order B-splines have the useful property of linear piecewise derivatives, so make a computationally efficient choice.

The genetic algorithm produced solutions that were distinct across subjects and had the lowest predicted excitation error. Although the current computational time ($\geq$ 1 hour) is impractical for *online* use, the solutions may be useful as improved initialisations to the interior-point method.

In this work, it was found that correcting for distortions in the gradient system was essential for accurate 3D excitation (and also 2D excitation in pilot work). Traditionally, gradient pre-distortion is applied once a desired k-space trajectory is designed. Equivalently, the realistic performance of the gradient system was included at each stage of the

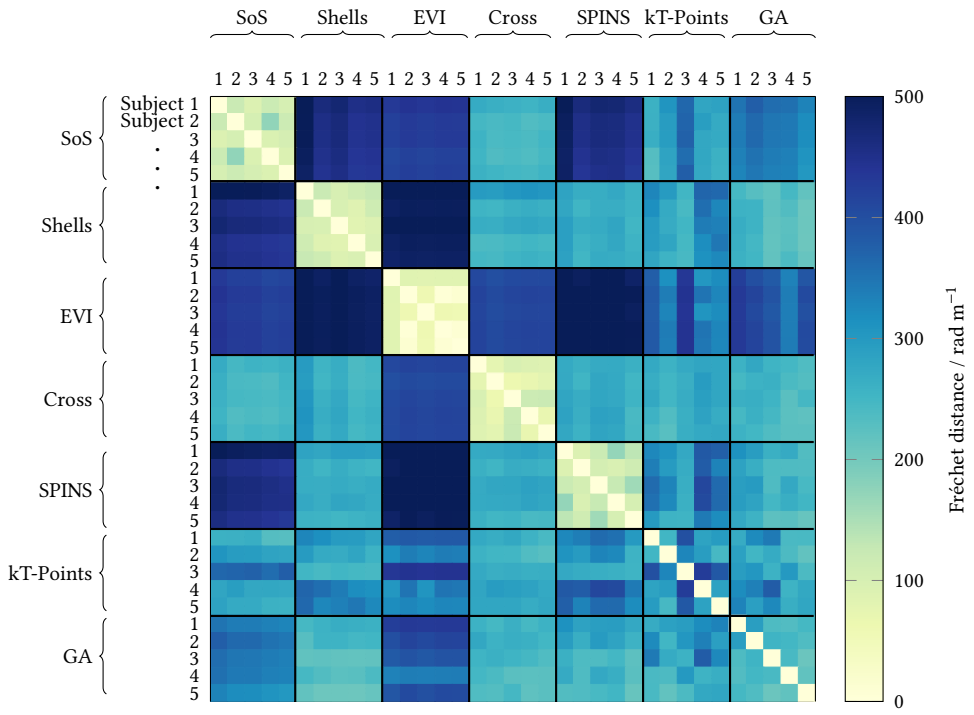

Figure 3.9: Cross-comparison of Fréchet distance between optimised k-space trajectories across six initialisations and five subjects.

optimisation by distorting the linear basis set (see Equation 3.3). Slew-rate constraints were evaluated assuming an ideal system response for computational efficiency. This is a valid assumption as the system response is causal, leading to an over-estimation of the realised slew-rates.

3.6 Conclusions

In this chapter a GPU accelerated joint design of short 3D excitation pTx pulses and k-space trajectories was implemented and tested. A linear transformation was used to parameterise an arbitrary 3D k-space trajectory, which included distortion due to gradient-chain imperfections. Optimisation methods produced novel k-space trajectories that resulted in reduced excitation error.

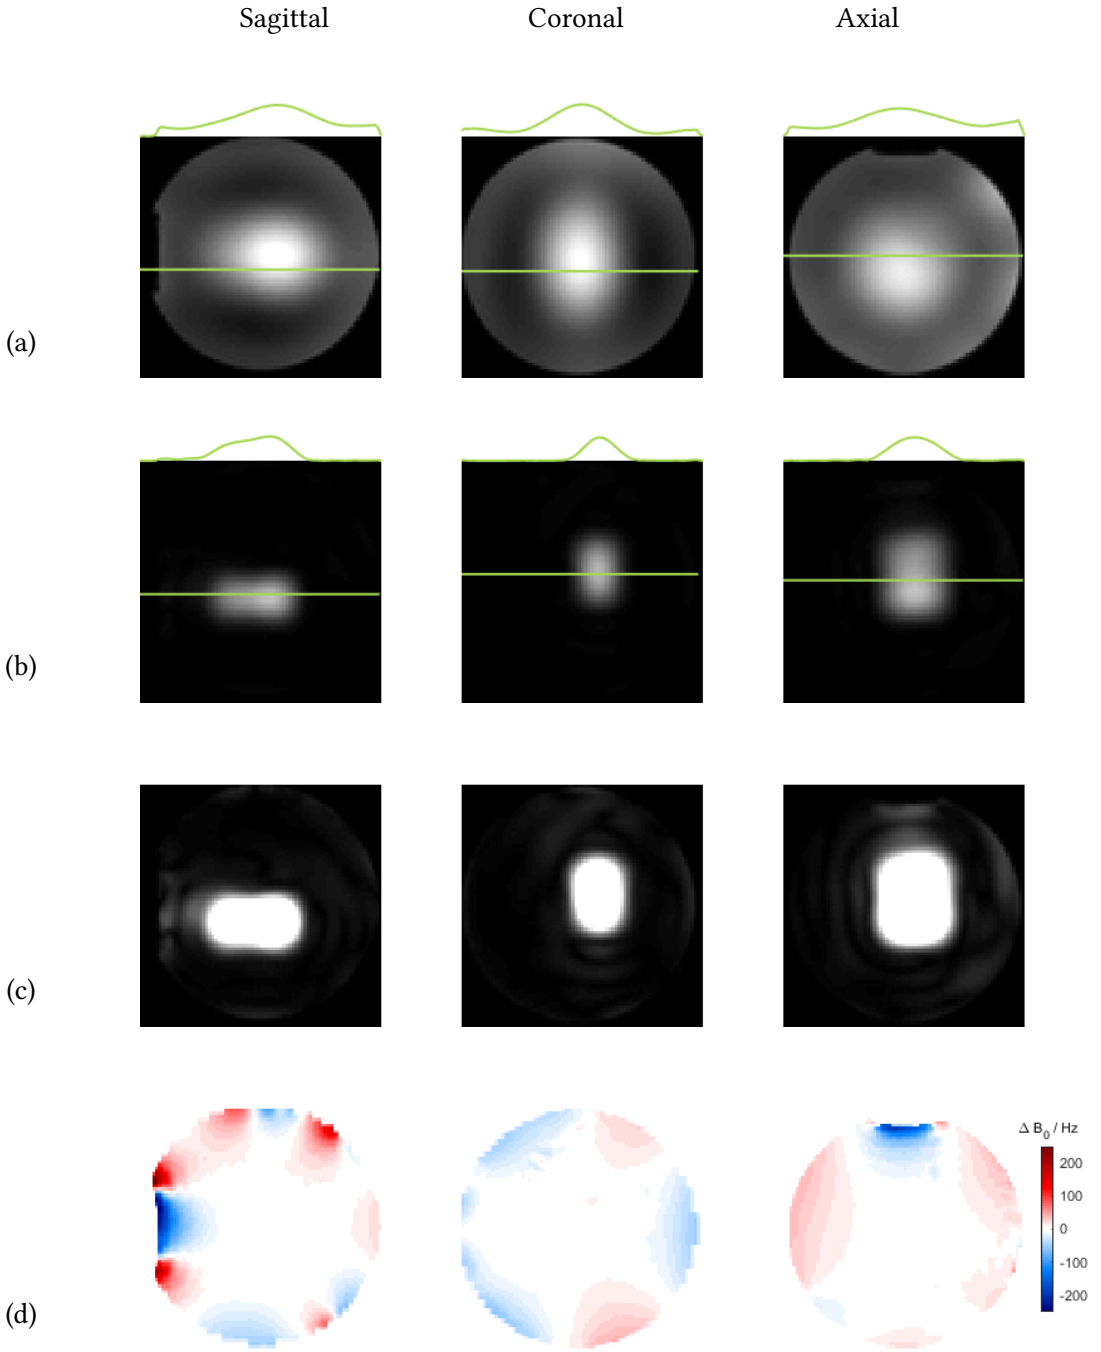

Figure 3.10: Example full field of view images for 3D GRE. Sagittal, coronal and axial slices shown for regular excitation (a); reduced Field-of-Excitation (b); reduced Field-of-Excitation windowed to 20% of maximum intensity (c). Field maps from the same phantom are also shown (d).

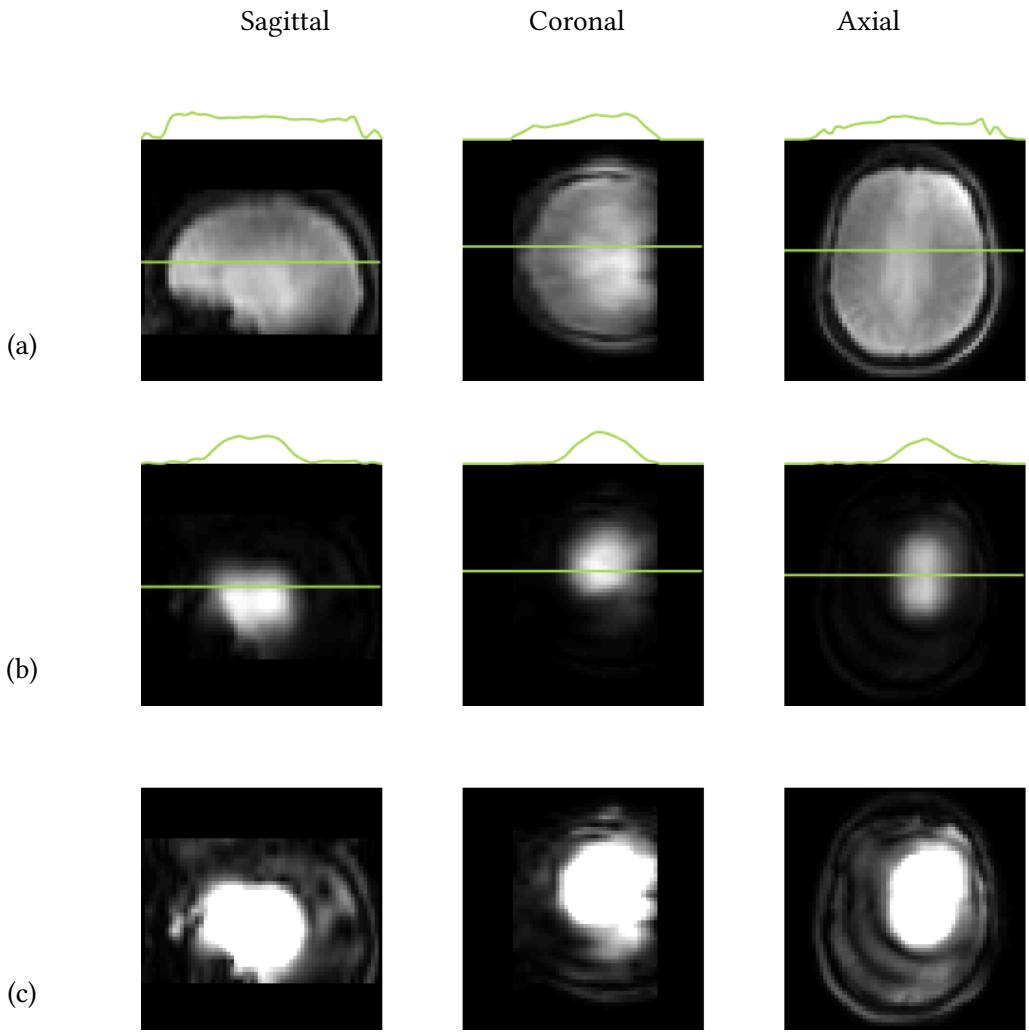

Figure 3.11: Example full field of view images for 3D GRE. Sagittal, coronal and axial slices shown for regular excitation (a); inner-volume excitation (b); inner-volume excitation windowed to 20% of maximum intensity (c).

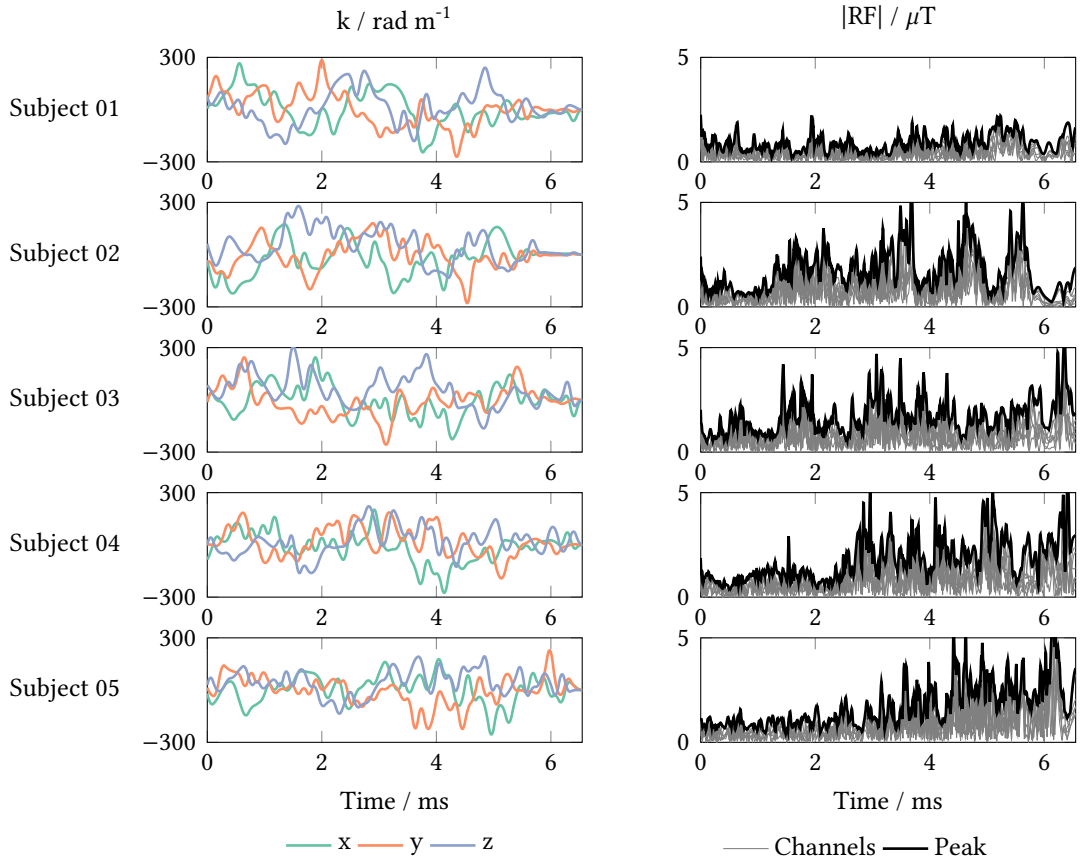

Figure 3.12: For each subject, optimised excitation k-space trajectories per axis (left), and magnitude plot of associated multichannel RF waveforms (right).

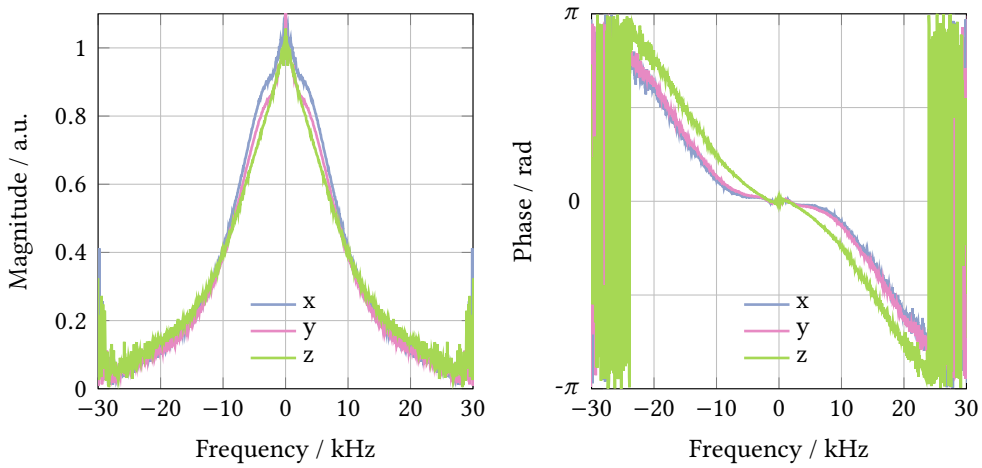

Figure 3.13: Measured GIRF function (*Philips 7 T Achieva, Nottingham*), magnitude and phase displayed in the range -30 to 30 kHz.

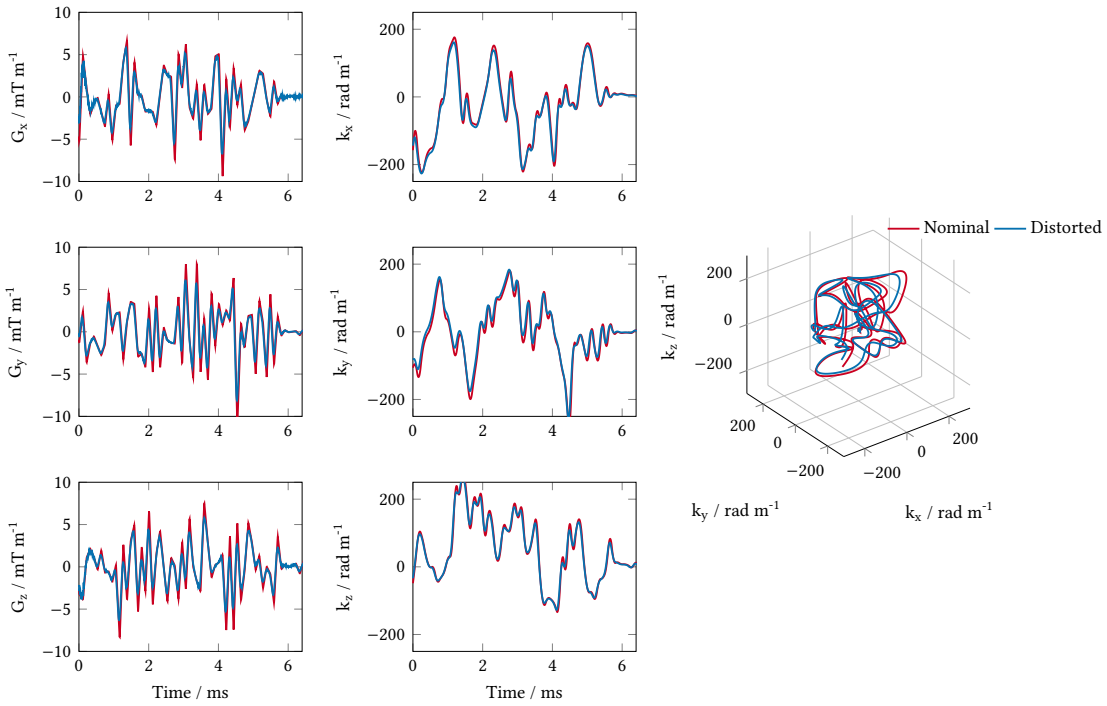

Figure 3.14: Example of gradient waveforms and corresponding excitation k-space trajectory, before and after the application of measured GIRF. The low-pass nature of the system response leads to an undershoot the peak gradient amplitude.

Chapter 4

Dynamic Parcellated B_0 shimming for 3D EPI

In this chapter a novel acquisition strategy is developed and implemented to improve the effective overall static field homogeneity by utilising reduced field of excitation with acceleration via parallel transmission, in conjunction with dynamic first-order local shimming. Dynamic shimming on a slice-by-slice basis has been shown to improve shimming efficiency (*203*, *204*). Previous work based on simulations has shown that further gains in shimming efficiency can obtained by dividing the global volume into compact parcellated sub-volumes rather than slices (*183*). Algorithms are introduced to identify arbitrary parcellations with sub-volumes of varying size and shape to further improve effective shimming efficacy.

- Initial results from this chapter formed the basis of a conference abstract (*205*)
- The parcellation code is published at https://github.com/cmirfin/ParcellatedB0

4.1 Introduction

Fast three-dimensional (3D) imaging has been central to the development of MRI since its inception. The developers of Echo-Volumar Imaging (EVI) envisioned the acquisition of a whole volume following a single excitation in a fraction of second (*198*, *206*). Ultimately, limitations on gradient slew rates due to physiological concerns and the requirement for higher resolutions have limited the application of EVI. The two-dimensional (2D) variant, Echo-Planar Imaging (EPI) has proved much more applicable. Imaging a whole volume is made possible by either multi-slice methods (*207*), phase-encoding in a second dimension (*208*), or by simultaneous multi-slice excitation (*209–212*). As all of these sequences acquire multiple, sequentially phase-encoded gradient echoes after each excitation, they are vulnerable to geometric distortion and signal dropout due to a spatially-varying B_0 fields (*32*).

As described in Chapter 1, producing a uniform B_0 field at 7 T is extremely challenging over an extended volume. As a typical first stage, scan preparation steps are used to determine the centre frequency offset due to the enhanced susceptibility effects. A

global shim optimises the shim coil currents to reduce spatial variation of the field over the whole volume-of-interest. These currents then remain fixed for the entire sequence.

As an alternative, *dynamic* shimming is performed by dynamically-updating the shim coil currents on a slice-by-slice basis (*203*). Implementations of dynamic shimming up to third-order have been successfully demonstrated at 7 T (*204*). One of the primary applications of dynamic shimming at 7 T has been for Magnetic Resonance Spectroscopy in the human brain (*34*, *213*, *214*), where field variations lead to loss of sensitivity and spectral resolution. Juchem et al (*213*) noted that even with up to third-order shim coils, the low order spherical harmonic fields cannot fully compensate the higher-order field inhomogeneities observed in vivo.

In previous work, the concept of dynamic shimming was extended from slices to *parcellated* cuboidal sub-volumes (*183*). For the human brain, reduced variation in the static field was found when shimming over a regular lattice of cuboidal sub-volumes than when shimming dynamically over the equivalent number of slices. Although a trend was identified between increasing sub-volume *compactness* and shim efficacy, the underlying field morphology was not used in the design of each volume parcellation.

Despite this, no dynamic implementation of parcellated shimming has thus far been demonstrated due to experimental practicalities. Ideally, the application of parcellated dynamic shimming would incur a minimal time penalty as compared to a standard 3D acquisition. Naively, the acquisition of N sub-volumes with a full Field of View would result in an N-fold increase in imaging time.

To avoid saturating the signal from subsequent sub-volumes, it is requisite to use reduced Field-of-Excitation (rFOX) to excite signal only within a given sub-volume. Note, for these reasons outer-volume suppression is incompatible with this approach. To sequentially excite sub-volumes for a given parcellation scheme requires the tailored redesign of rFOX pulses. In this work, dynamic parallel transmit rFOX pulses are used in addition to first-order local ΔB_0 shimming. (see schematic Figure 4.2).

4.2 Theory

4.2.1 Dynamic B_0 shimming

The determination of N-dynamic shim coil currents is posed as a series of N-constrained, linear, least-squares optimisations over N sub-volumes ($V_n \in \mathbb{R}^3$) of the whole volume $V \in \mathbb{R}^3$. The shim currents for the nth dynamic shim are denoted by: $I_n = [I_n^{\{0\}}, I_n^{\{1,-1\}}, I_n^{\{1,0\}}, I_n^{\{1,1\}}, \ldots]^T$ up to a predefined order $\{m\}$.

$$\min_{\substack{I_n \\ \{V_1,V_2,\ldots,V_N\}}} \sum_{n=1}^{N} ||B_{\text{shim}}(\mathbf{x})\, I_n - \Delta B_0(\mathbf{x})||^2_{V_n} \quad \text{s.t.} \quad |I_n^{(o)}| \leq I_{\text{max}}^{(o)}$$
$$V_n \subseteq V, \quad \bigcup_{n=1}^{N} V_n = V, \quad V_n \bigcap V_m = \varnothing \,. \tag{4.1}$$

Here $\Delta B_0(\mathbf{x})$ represents the measured off-resonance field for each point in the whole volume. Shim coil fields are concatenated into the matrix B_{shim}. Each dynamic shim is solved separately using each specified sub-volume V_n as a region mask in the optimisation. The current constraints $I_{\text{max}}^{(m)}$ are related to the maximum current output of each shim coil amplifier. As the zeroth-order shim is usually applied as a global frequency offset, there is accordingly no current constraint. For a given volume parcellation $\{V_1, V_2, \ldots, V_n\}$, each constrained linear system is readily solved using an interior-point method (see Chapter 1).

4.2.2 Parcellation schemes

As denoted in Equation 4.1, a suitable volume parcellation $\{V_1, V_2, \ldots, V_n\}$ should be dense (i.e. no gaps between sub-volumes) and non-overlapping. For practical imaging terms, each sub-volume should be a *connected* region and have a volume that is greater than some minimum value. With these constraints in mind, three general parcellation approaches can be categorised as:

1. Regular sub-volumes, regular packing
2. Regular sub-volumes, irregular packing
3. Irregular sub-volumes, irregular packing

The first of these consists of regular cuboids of dimension $(\Delta x, \Delta y, \Delta z)$, centred on a regular grid, as previously discussed (*183*). The second approach relaxes the condition that the cuboids are arranged on a regular grid. However, to maintain a dense packing, only certain translations are allowed. A search algorithm is described in Algorithm 4.

For the third general parcellation approach of irregular sub-volumes with irregular packing, an algorithm is required to suitably subdivide the volume. So-called *adaptive mesh refinement* is common in the field of numerical analysis, whereby simulation domains are sub-divided around regions of rapid change or instability (*215*). A similar concept referred to as the *split and merge* segmentation algorithm exists in image processing (*216*). For application to parcellated B_0 shimming, the motivation is to minimise the number of sub-volumes required to achieve a field uniformity below a certain threshold. To sub-divide a 3D volume, an octagonal tree decomposition can be used (see Algorithm 5).

Algorithm 4: Non-rigid parcellation

For a given rigid cuboidal parcellation, the algorithm performs an exhaustive search over permitted offsets in the relative position of each parcellated sub-volume. The positional offsets that yield the lowest cumulative residual field error are returned.

Given parcellation sub-volume dimensions $(\Delta x, \Delta y, \Delta z)$ over domain $\Omega \subset \mathbb{R}^3$; choose slice direction;

$(X_k, Y_k, Z_k),\ \ k \in \{1, \ldots, N_v\}$; *// Initial parcellation sub-volume centres.*

for $n_z \leftarrow 1$ **to** N_z **do** *// Loop over slices N_z*

  for $n_x \leftarrow 1$ **to** N_x **do** *// Loop over N_x positions*

    for $n \leftarrow 1$ **to** N **do** *// Loop over y-shifts*

      $Y'_k(n_x, :, n_z) \leftarrow Y_k(n_x, :, n_z) + n\Delta y/N;$

      $V'_k \leftarrow \{(x, y, z) \in \Omega \mid |x - X_k(n_x, :, n_z)| \leq \Delta x \wedge |y - Y'_k(n_x, :, n_z)| \leq \Delta y \wedge |z - Z_k(n_x, :, n_z)| \leq \Delta z\},\ \ \forall k \in \{1, \ldots, N_y\};$

      $m_{y,n} \leftarrow \sum_{k=1}^{N_y} \min ||B_{\text{shim}}(\mathbf{x}) I^{\{k\}}_{\text{shim}} - \Delta B_0(\mathbf{x})||^2_{V_k};$ *// sub-volume with largest residual*

    end

    $m_{x,y} \leftarrow m_{x,y} + \min m_{y,n}$; *// add minimum residual for this x-offset*

  end

  for $n_y \leftarrow 1$ **to** N_y **do** *// Loop over N_y positions*

    for $n \leftarrow 1$ **to** N **do** *// Loop over x-shifts*

      $X'_k(:, n_y, n_z) \leftarrow X_k(:, n_y, n_z) + n\Delta x/N;$

      $V'_k \leftarrow \{(x, y, z) \in \Omega \mid |x - X'_k(:, n_y, n_z)| \leq \Delta x \wedge |y - Y_k(:, n_y, n_z)| \leq \Delta y \wedge |z - Z_k(:, n_y, n_z)| \leq \Delta z\},\ \ \forall k \in \{1, \ldots, N_x\};$

      $m_{x,n} \leftarrow \sum_{k=1}^{N_x} \min ||B_{\text{shim}}(\mathbf{x}) I^{\{k\}}_{\text{shim}} - \Delta B_0(\mathbf{x})||^2_{V_k};$ *// sub-volume with largest residual*

    end

    $m_{y,x} \leftarrow m_{y,x} + \min m_{x,n}$; *// add minimum residual for this y-offset*

  end

  if $m_{x,y} < m_{y,x}$ **then**

    Pick x-offsets for n_z slice;

  else

    Pick y-offsets for n_z slice;

  end

end

Algorithm 5: Split & Merge parcellation

The split & merge algorithm recursively sub-divides regions until every sub-volume has a residual shim error that is below a given tolerance, or the sub-volume is of a minimum dimension permitted. A second stage performs region growing to iteratively merge adjacent sub-volumes that have the lowest residual shim error (up to a given number of sub-volumes).

$x, y, z \leftarrow 1$; *// Initial point, dimension length Δ with minimum dimension Δ_0*
while *length*$(x) > 0$ ***and*** $\Delta \geq \Delta_0$ **do**
  for $k \leftarrow 1$ **to** *length*(x) **do** *// Loop over all regions to be processed*
    $V_k \leftarrow \{(x, y, z) \in \Omega \mid |x - x_k| \leq \Delta \wedge |y - y_k| \leq \Delta \wedge |z - z_k| \leq \Delta\}$;
    $m \leftarrow \min_{|I|<I_{\max}} ||B_{\text{shim}}(\mathbf{x}) I_{\text{shim}}^{\{k\}} - \Delta B_0(\mathbf{x})||_{V_k}^2$;
    if $m >$ *tolerance* **then** *// Check region uniformity*
      $x'_k \leftarrow [x_k, x_k, x_k + \Delta/2, x_k + \Delta/2, x_k, x_k, x_k + \Delta/2, x_k + \Delta/2]$; *// Octree*
      $y'_k \leftarrow [y_k, y_k + \Delta/2, y_k, y_k + \Delta/2, y_k, y_k + \Delta/2, y_k, y_k + \Delta/2]$;
      $z'_k \leftarrow [z_k, z_k, z_k, z_k, z_k + \Delta/2, z_k + \Delta/2, z_k + \Delta/2, z_k + \Delta/2]$;
    else *// Remove kth point*
      $x'_k, y'_k, z'_k \leftarrow \varnothing$
    end
  end
  $\Delta \leftarrow \Delta/2$; *// reduce dimension length for next iteration*
  $x \leftarrow x',\ y \leftarrow y',\ z \leftarrow z'$; *// assign new points*
end
while *length*$(x) > N_v$ **do**
  merge adjacent sub-volumes that have lowest residual errors
end

An alternative approach for an arbitrary 3D parcellation is to use a Voronoi tessellation, whereby all grid points in the domain are associated to the nearest point in a finite set of site points. Again, an algorithm is required to define this finite set of points and a proposed method is described in Algorithm 6.

Algorithm 6: Adaptive Voronoi parcellation
The algorithm iteratively divides a volume into an increasing number of sub-volumes. At each iteration, the sub-volume with the largest residual shim error is divided into two further Voronoi polyhedrons by selecting the point most distant from the current sub-volume origin.

$p_0 \in X$; *// initial point in domain*
for $n \leftarrow 1$ **to** $N_v - 1$ **do** *// add new sub-volume up to maximum N_v*
  $V_k \leftarrow \{x \in X \mid d(x, p_k) \leq d(x, p_j), \forall j \neq k\}, \ \forall k \in \{1, \ldots, n\}$; *// Euclidean $d(\cdot)$*
  $k^* \leftarrow \underset{\substack{k \in \{1,\ldots,n| \\ |V_k| > V_{\min}\}}}{\arg\min} ||B_{\text{shim}}(\mathbf{x}) I_{\text{shim}}^{\{k\}} - \Delta B_0(\mathbf{x})||_{V_k}$; *// sub-volume with largest residual*
  $p_{n+1} \leftarrow \underset{x \in V_{k^*}}{\arg\max}\, d(x, p_{k^*})$; *// select new Voronoi point as distant point in V_{k^*}*
end

The parcellation methods so far described are referred to as 3D *arbitrary* parcellations. In addition, it is useful to consider parcellations restricted to at least one full dimension (3D *prism*) or half-dimension (3D *half-prism*). These parcellation methods are visualised in Figure 4.1.

4.2.3 Dynamic pTx selective pulses

Selective excitation is required to excite signal in each given parcellated sub-volume. To dynamically achieve this over multiple volumes requires the design of multiple selective excitation pulses. As the ΔB_0 field will vary over each sub-volume, this should be taken account of in the design process. For the jth dynamic with first-order shim coil currents $\mathbf{I}_j^{\{1\}} = (I_0^{\{0,0\}}, I_0^{\{1,1\}}, I_0^{\{1,-1\}})$ and zeroth order offset $\Delta\omega_l^{\{0,0\}}$, the modified effective spatial field is given by:

$$\Delta B_{0,j}(\mathbf{x}) = \Delta B_0(\mathbf{x}) - \Delta\omega_j^{\{0\}}/\gamma - \mathbf{x} \cdot \mathbf{I}_j^{\{1\}} + O(\mathbf{x}^2) + \ldots, \tag{4.2}$$

where $\Delta B_0(\mathbf{x})$ is the unshimmed field and higher-order shim terms have been neglected for simplicity. Expanding the spatial-domain Small-Tip Angle approximation (see Chapter 2, Equation 2.5), the excitation pattern resulting from RF waveforms $b_c(t)$ applied on N_c transmit channels becomes:

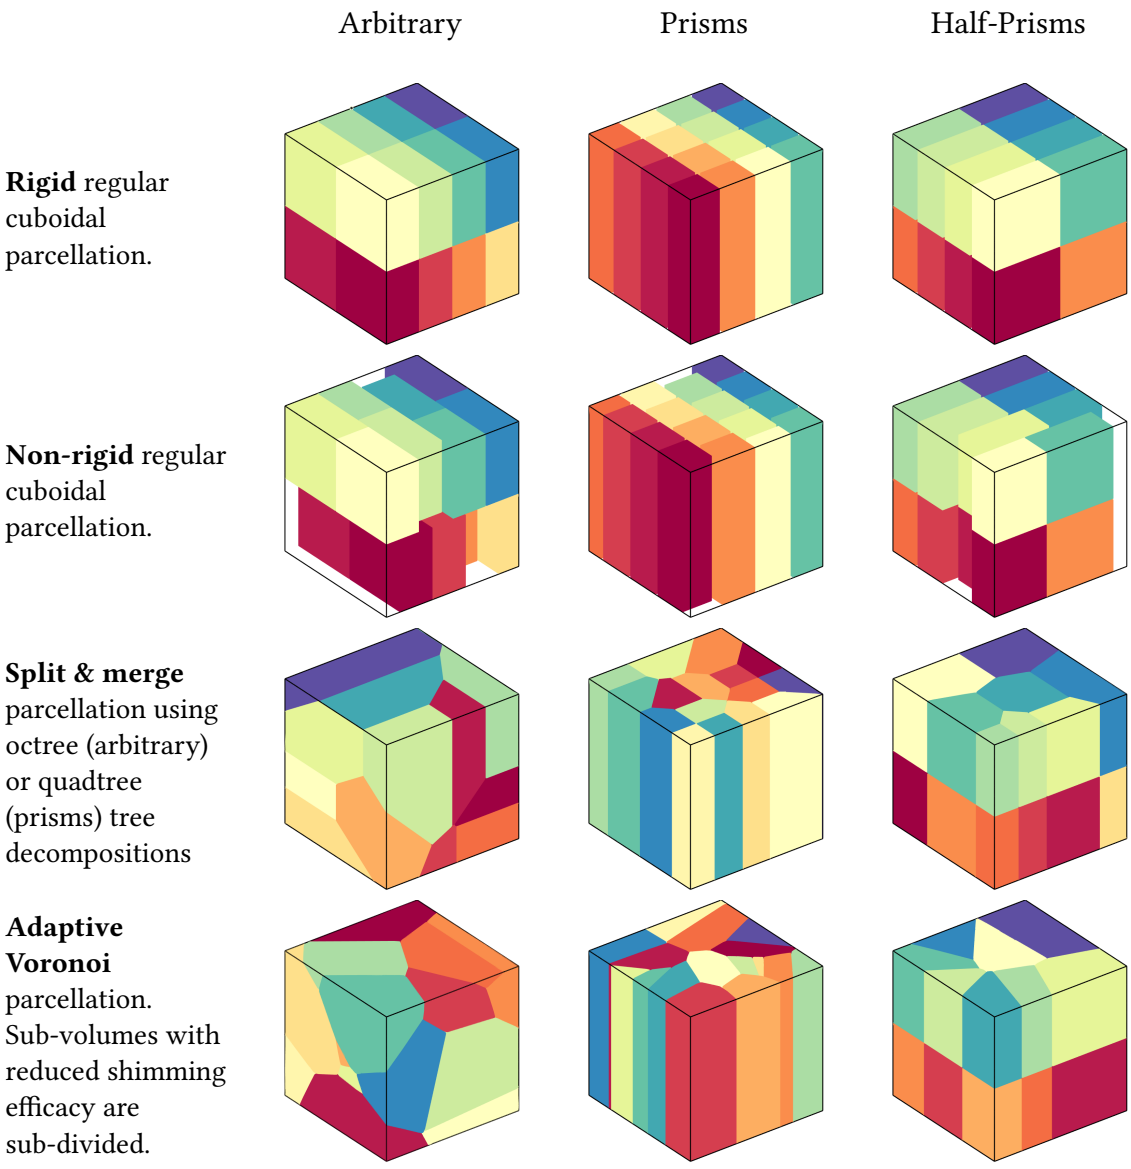

Figure 4.1: Visualisation of the parcellation methods discussed in Section 4.2.3, for both arbitrary 3D sub-volumes, 3D prisms and 3D half-prisms with a total of 16 sub-volumes each.

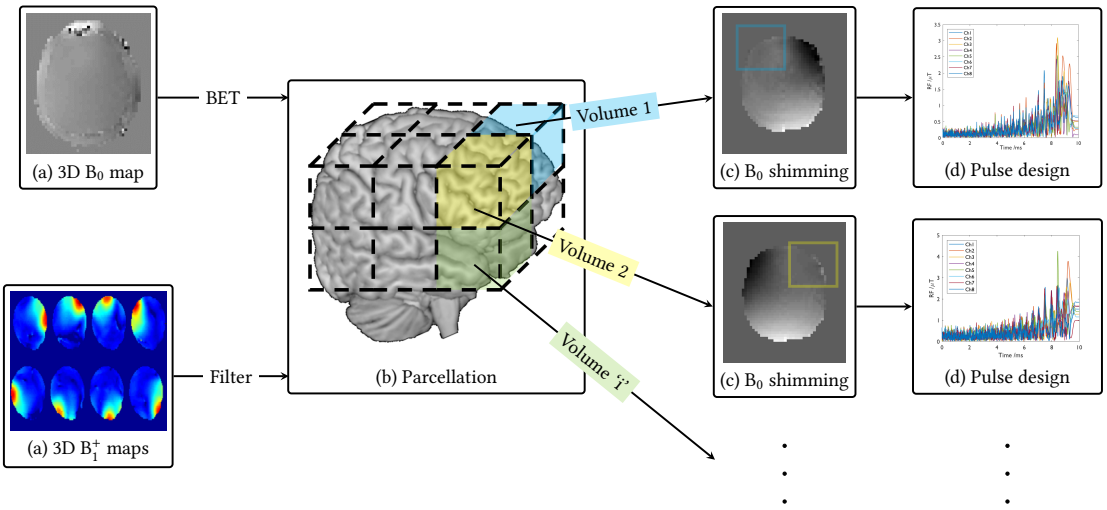

Figure 4.2: Schematic diagram of workflow for parcellated dynamic shimming design. Field maps (both ΔB_0 and B_1^+) are acquired and pre-processed (a), a suitable volume parcellation is devised (b), for each sub-volume, shim currents are optimised and a predicted off-resonance field is generated (c), for each sub-volume, rFOX pulses are designed with updated off-resonance field (d).

$$m(\mathbf{x}) = i\gamma m_0 \sum_{c=1}^{N_c} s_c(\mathbf{x}) \int_0^T b_c(t) \underbrace{e^{i\mathbf{x}\cdot\mathbf{k}(t)}}_{\substack{\text{encoding}\\ \text{matrix}}} \underbrace{e^{i\gamma\Delta B_0(\mathbf{x})(t-T)}}_{\substack{\text{offset from}\\ \text{unshimmed}\\ \text{field}}} \underbrace{e^{-i\Delta\omega_j^{\{0\}}(t-T)}}_{\substack{j\text{th zeroth}\\ \text{offset}}} \underbrace{e^{-\gamma\mathbf{x}\cdot\mathbf{I}_j^{\{1\}}(t-T)}}_{\substack{j\text{th first-order}\\ \text{phase correction}}} dt, \quad (4.3)$$

Assuming a fixed excitation k-space trajectory $\mathbf{k}(t)$, the first two exponents in Equation 4.3 are common across dynamics and hence can be precomputed. The two final exponents are dependent on the jth shim values.

4.3 Methods

4.3.1 UK7T Network Dataset

To explore the theoretical performance of the parcellated ΔB_0 shimming schemes suggested in Section 4.2.2, an evaluation was performed using field maps from the UK7T Network study (*217*). The dataset consisted of scans from 10 volunteer subjects, but was restricted to data acquired from the Philips 7 T Achieva scanner in Nottingham. This consisted of whole-brain field maps with an isotropic resolution of 4 mm. Each field-map was interpolated to a grid size of $32 \times 32 \times 32$.

4.3.2 Evaluation metrics

The notion of *compactness* was found to be a key parameter in previous work (*183*). For the purposes here, compactness of an arbitrary 3D shape is defined using the measure

of *sphericity* (*218*):

$$\Theta_3 = \frac{\pi^{\frac{1}{3}}(6V)^{\frac{2}{3}}}{A_S}, \tag{4.4}$$

where V and A_S are the volume and surface area of the shape respectively. By construction, Θ_3 is bounded on the unit interval, with $\Theta_3 = 1.0$ for a sphere of any radius. Shapes which are more elongated in any one dimension produce a reduced measure. For prisms a different measure is used: the isoperimetric quotient, or *circularity* is defined as the ratio of cross-sectional area $A_\times$ of the prism to the area of a circle with the equivalent perimeter L:

$$\Theta_2 = \frac{4\pi A_\times}{L^2}, \tag{4.5}$$

Shimming efficacy ε quantifies the improvement in field uniformity after shimming and is defined as:

$$\varepsilon = 1 - \frac{\sum_{n=1}^{N} ||B_{\text{shim}}(\mathbf{x})\, I_{\text{shim}}^{\{n\}} - \Delta B_0(\mathbf{x})||_{V_n}}{||\Delta B_0(\mathbf{x})||}, \tag{4.6}$$

where $||\cdot||$ represents the l_2-norm. A value of $\varepsilon = 1$ corresponds to perfect shimming, while $\varepsilon = 0$ indicates no improvement in field homogeneity on shimming.

4.3.3 Experimental Demonstration

All experiments were performed using an 8-channel transmit/32-channel receive head coil (Nova Medical Inc, Wilmington). The experimental protocol is described in schematic form in Figure 4.2. 3D off-resonance maps with $(4 \times 4 \times 8)\text{mm}^3$ resolution were acquired using the dual TE method (flip angle = 8°, TR = 3.8 ms, TE = 1.52 ms, ΔTE = 1 ms). 3D B_1^+ sensitivity maps were calculated from an Actual Flip-Angle (*77*) quadrature transmit mode (flip angle = 60°, TR_1/TR_2 = 30 ms / 150 ms) and eight individual channel gradient echoes (flip angle = 15°, TR= 7.1 ms).

All developed shimming, RF and gradient waveform design was performed online using in-house written software (see Appendix A). Brain extraction using BET (*219*) was performed using the first magnitude image obtained from the Dual TE method. B_1^+ sensitivity maps were median filtered and appropriately thresholded. Shim currents were optimised up to first-order using the interior-point algorithm from the MATLAB optimisation toolbox (Natick, Massachusetts). A predefined spiral k-space trajectory (duration 6 ms, acceleration 3.5) was used for all pulse design with gradient pre-emphasis.

3D EPI readout was used with 3 mm isotropic resolution and a whole brain FOV (flip angle = 10°, TR = 30 ms, TE = 27 ms) with a SENSE factor of 2 in the in-plane phase encoding direction.

As a first demonstration, each sub-volume from a parcellation was imaged separately. The parcellation consisted of 6 cuboidal sub-volumes arranged firstly over the superior brain region (two volunteers) and secondly over the inferior brain region (two volun-

teers). A dynamic implementation (TR_{dyn} = 30ms) used a parcellation of 4 cuboidal sub-volumes across the mid-brain region (two-volunteers). The dynamic implementation was also tested with a spherical water phantom using multiple parcellations to validate the correct rFOX selection.

4.3.4 Modifications to scan control software

The existing scan control software for the Philips 7 T Achieva (Best, Netherlands) can load a single set of pTx excitation pulses and associated gradient waveforms from file. This utility was extended to allow multiple sets of pTx pulses to be externally defined. Care was taken to disable internal global SAR calculations for all but the set of pTx pulses with the largest integrated power, ensuring an upper-bound on the estimated sequence global SAR per dynamic.

In the sequence control software, each of the pulses for all dynamics were generated and overlaid at the same sequence time point with a unique identifier. At run-time, for each distinct dynamic, a dynamic loop counter was updated and used to disable all but one of the pTx RF objects. Due to memory restrictions on the data acquisition system, a maximum of four sets of dynamic pulses were allowed. For the same reason, an identical excitation k-space trajectory was used for all dynamics.

Similarly, an existing mechanism to load shim current values from file was extended to allow for N-dynamics. Before the execution of each dynamic sequence, the frequency offset and first-order shim currents were updated accordingly. These updates are performed on the order of milliseconds, and so produced minimal sequence delays.

4.4 Results

4.4.1 UK7T Network Dataset

Parcellating the brain region into an increasing number of sub-volumes led to an increase in shim efficacy across all (n = 10) subjects, for both 1st and 2nd order shimming. This applied to all of the methods described in the previous section (see Figure 4.3). The greatest rate of increase in shim efficacy with increasing number of shim volumes was found for low numbers of sub-volumes (less than 20). Increasing the number of sub-volumes beyond this number is of limited gain. Indeed, for 2nd order shimming with 3D prisms, the flatness of the curves in Figure 4.3(d) beyond $N_v = 20$ suggests the field changes along each prism cannot be characterised using up to 2nd order shim fields alone.

For the case of global shimming ($N_v = 1$), efficacies of $\varepsilon = 0.13 \pm 0.07$ and $\varepsilon = 0.19 \pm 0.06$ were achieved using 1st and 2nd order shimming respectively. Here the mean average and standard deviations are taken over subjects. For the special case of axial slice shimming ($N_v = 32$), the efficacies increased to $\varepsilon = 0.26 \pm 0.04$ and $\varepsilon = 0.32 \pm 0.06$

using 1st and 2nd order shimming respectively. Using the same number of sub-volumes, but rearranged into different geometries, much higher efficacies were obtained. For 32 sub-volumes, the largest shim efficacy for a 3D arbitrary parcellation was found using the adaptive Voronoi algorithm with values of $\varepsilon = 0.54 \pm 0.05$ and $\varepsilon = 0.73 \pm 0.05$ for 1st and 2nd order shimming respectively.

For both 3D geometries (arbitrary and prisms), the difference between the parcellation methods (Rigid, non-rigid, split & merge, adaptive Voronoi) was reduced using 2nd order as compared to 1st order shimming.

Figure 4.4 displays the relationship between measures of sub-volume compactness against shim efficacy for a fixed number of sub-volumes. In the general 3D case, shim efficacy was found to linearly increase with mean sub-volume sphericity. Similarly, in the 3D prisms geometry, shim efficacy was found to linearly increase with mean circularity (of the 2D cross-sectional area).

Figure 4.5 shows example parcellations and resulting B_0 fields for each of the parcellation methods, with associated mean-centred histograms.

4.4.2 Experiments

When parcellated 2nd order shimming was applied statically with 6 cubical regions, the shim efficacy of the ΔB_0 field over the combined superior sub-volumes was increased by 22% (subject 1) and 23% (subject 2) as compared to global shimming. Similarly, the shim efficacy of the B_0-field inhomogeneity over the combined inferior sub-volumes was increased by 26% (subject 3) and 32% (subject 4). Figure 4.6 displays the 3D EPI magnitude and phase for subject 1. The signal is localised within each desired sub-volume, and the low spatial variation of the field over selected-volumes is evident. The ROI used for B_0 field shimming slightly extended beyond the excitation ROI for each sub-volume.

Figure 4.7 displays 3D EPI magnitude images acquired using the dynamic approach (dynamic 1st order shimming and dynamic Parallel Transmit excitation pulses) in a spherical phantom. Although limited to four dynamics, multiple different parcellations demonstrate the flexibility of the approach. A sharp transition band was prioritised in the design of each excitation pulse. A spatial weighting could be used to further improve the magnitude of the B_1^+ within each reduced Field-of-Excitation.

Figure 4.8 shows the same experiment with four dynamics applied in vivo. The total time required to calculate shim values and pTx pulses for four dynamics was less than 2 minutes in all cases, with no additional workflow compared to standard pTx experiments.

4.5 Discussion

Global shimming up to 2nd order over the whole brain is not able to fully remove regions of field inhomogeneity. Slice-by-slice dynamic shimming is also limited in efficacy. Par-

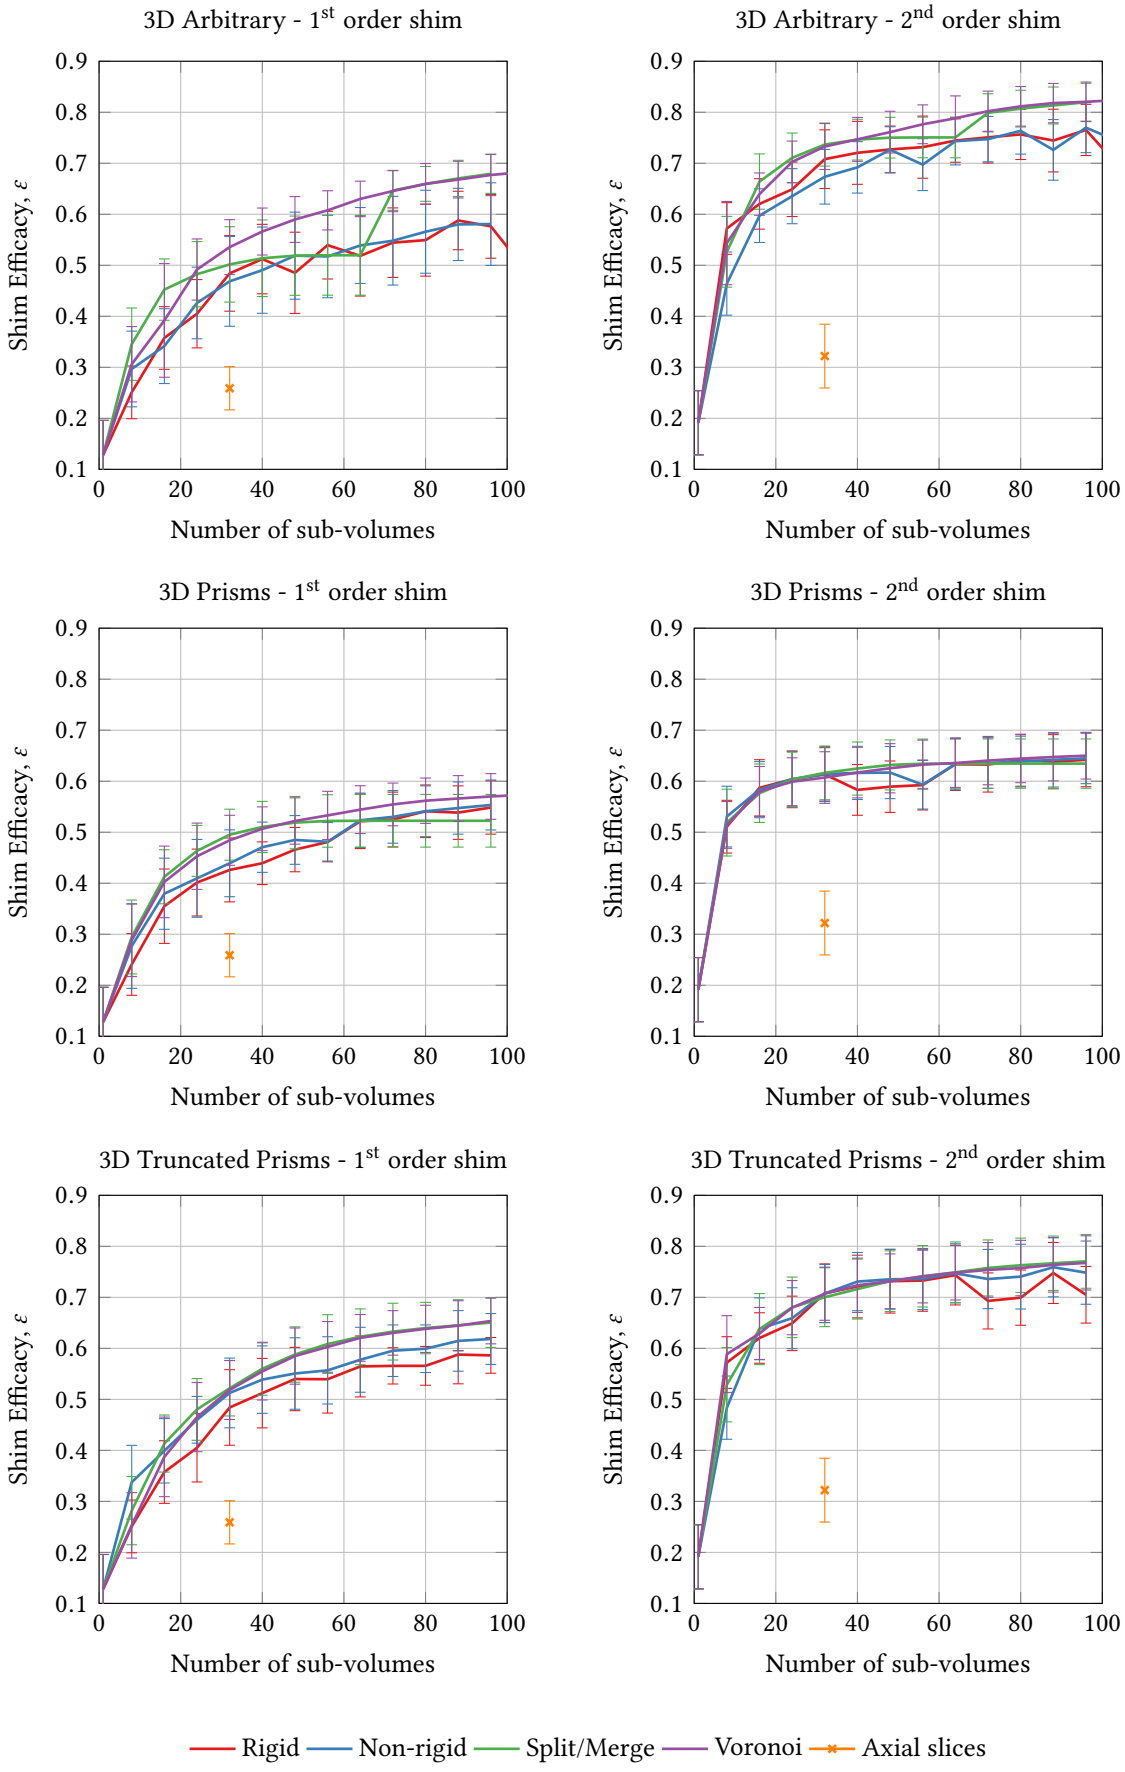

Figure 4.3: Shim efficacy for 10 parcellation methods and 32 sub-volumes. Mean and standard deviation across 10 subjects shown. For both rigid and non-rigid methods, only the largest shim efficacy parcellations are displayed.

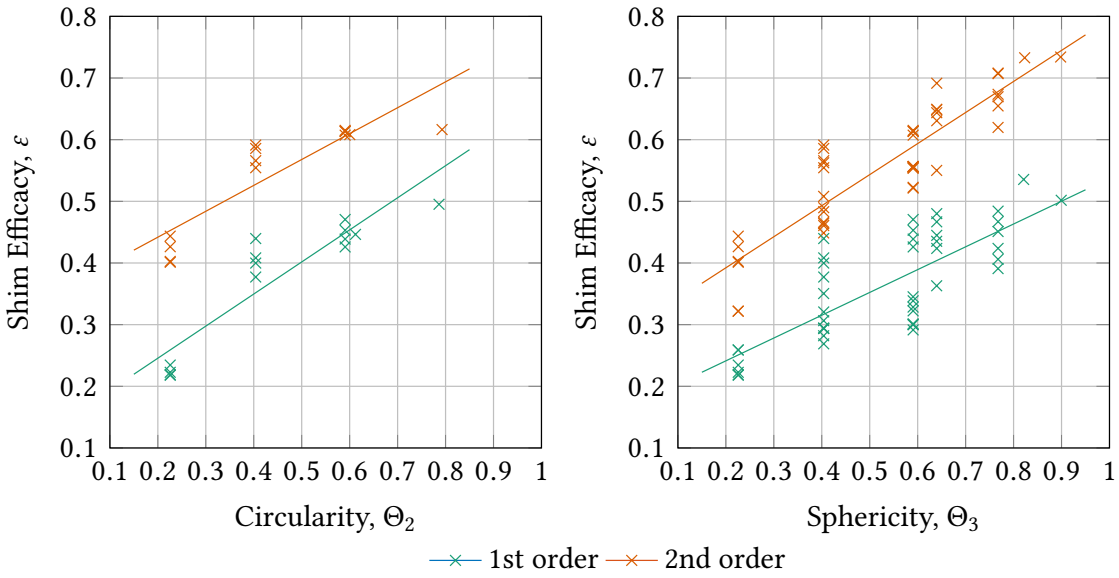

Figure 4.4: Shim efficacy as a function of circularity (left) for 3D prism parcellations, and (right) sphericity for arbitrary 3D parcellations. Mean values averaged across 10 subjects are shown for 32 sub-volumes. Regression lines are plotted for both 1st and 2nd order shimming.

cellated shimming affords a method to ameliorate field inhomogeneity. In particular, parcellated shimming with compact sub-volumes is increasingly effective. The explanation for this is that locally (within a compact region), the field variations can be better represented by 2nd order (or lower) spherical harmonic fields. Less compact sub-volumes, such as slices, are likely to encounter field variations that would require higher orders of spherical harmonic fields. A straightforward approach to parcellated shimming is therefore to choose highly compact sub-volumes, such as cubes.

However, due to the underlying morphology, the field inhomogeneity can be further ameliorated by relaxing the condition on using sub-volumes of equal size and shape. The two algorithms proposed for this purpose were the adaptive Voronoi and split and merge algorithms. Both attempt to segment the field based on the residual field error resulting from shimming over each of the sub-volumes. The split and merge algorithm is initially based on a cubical octree grid, which is recursively sub-divided until below a threshold (or minimum size) before a merging process. Using a coarse resolution (matrix size $32 \times 32 \times 32$), the resultant parcellations still exhibited many cuboidal sub-volumes. The adaptive Voronoi algorithm produces arbitrary convex polyhedrons that achieve higher shim efficacy than regular cuboids. This difference is particularly pronounced when using 1st order shimming (see Figure 4.3).

Restricting the parcellations to densely packed prisms limits the shim efficacy achievable as compared to arbitrary 3D parcellations, for the previously discussed notion of compactness. Despite this, the obvious advantage of using 3D prisms is that the region can be excited through the use of 2D selective pulses rather than full 3D excitation. Consequently, reduced pulse durations and higher excitation fidelity can be achieved. For

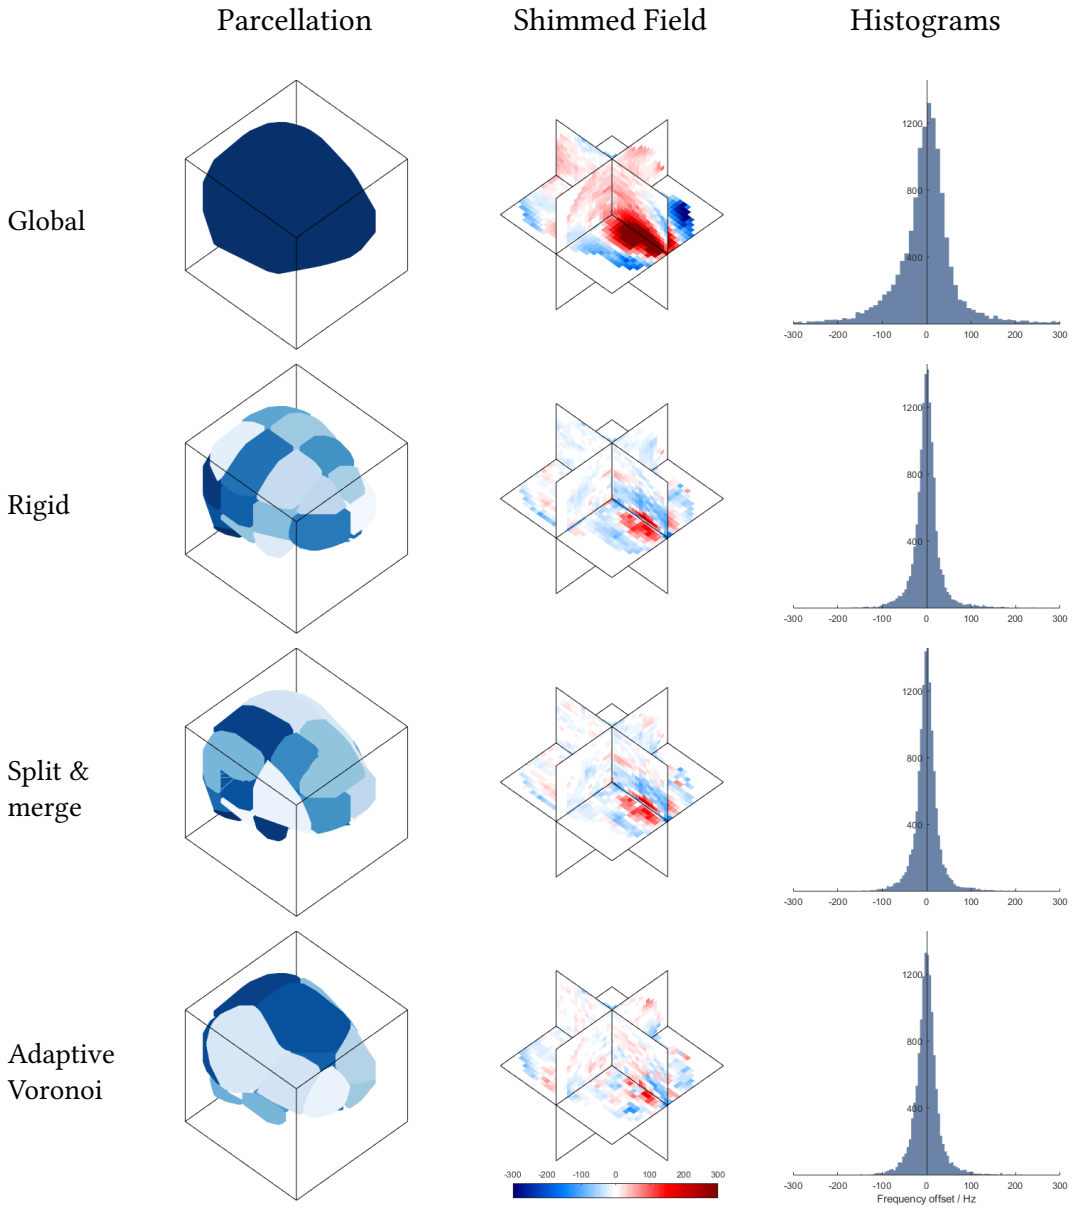

Figure 4.5: Visualisation of 3D arbitrary parcellations (left column), shimmed off-resonance field in Hz (centre column), associated mean-centred histograms (right column).

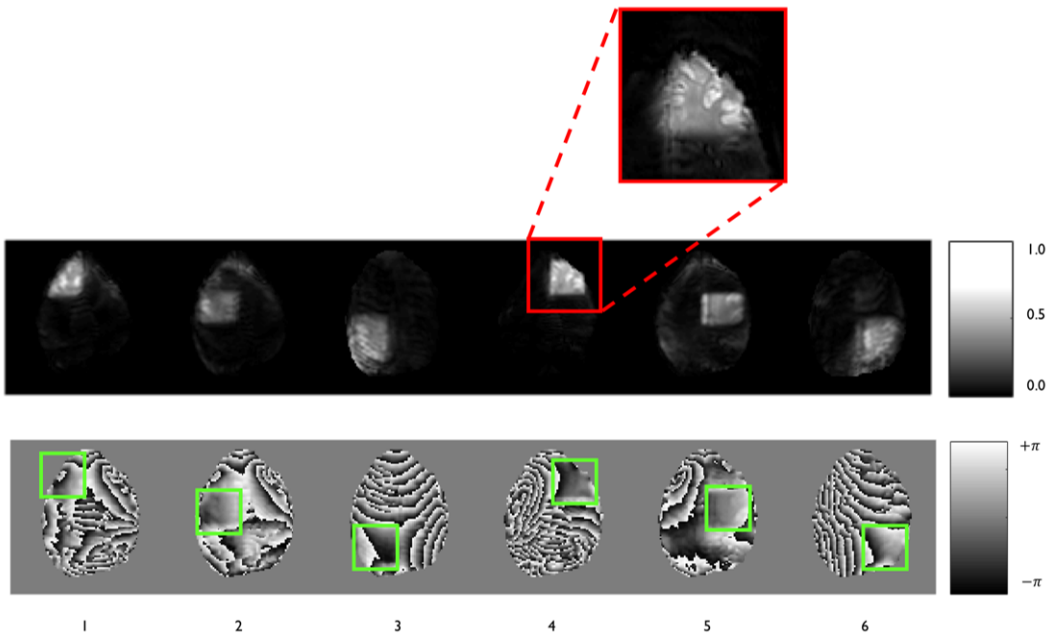

Figure 4.6: 3D EPI for 6 parcellated sub-volumes (central axial slices shown) with 3mm isotropic resolution (top row = relative magnitude, bottom row = phase with overlaid shim sub-volumes displayed). Also shown is zoomed sub-volume (1.5 mm^2 in-plane resolution) for sub-volume 4.

a 3D EPI readout, it is efficient to align the read-direction along the major axis of the prism.

An alternative hybrid approach could be to use the half-prism parcellation (with higher shim efficacy), but still to use a 2D excitation. As the magnetisation within the full-height prism would be saturated when exciting an individual half-prism, a period of T_1 recovery would be necessary before it would be possible to image the superior or inferior half-prism. A suitable acquisition cycle would need to be devised such that adjacent half-prisms could be imaged during this relaxation period.

The experimental implementation of dynamic parcellated shimming was limited by the available memory on the spectrometer's single board computer that is used for the control of each RF amplifier. More modern systems should not have this problem. Dynamic 2nd order shimming was not possible in this work, as the rise time of the associated amplifiers was on the order of seconds.

In order to avoid a large phase shift at the edge of each sub-volume experimentally (Figure 4.6), the ROI used for B_0 shimming slightly extended beyond the ROI of excitation in each case. Further improvements in the excitation robustness in the presence of off-resonance could be achieved by optimising the RF pulses at multiple offset frequencies (*187*). This approach was not utilised in this work due to limited computational memory.

Future development work should focus on reducing the additional time penalty for using dynamic parcellated shimming with reduced Field-of-Excitation (rFOX). Primarily in this work a full FOV was used for each parcellated sub-volume, increasing the

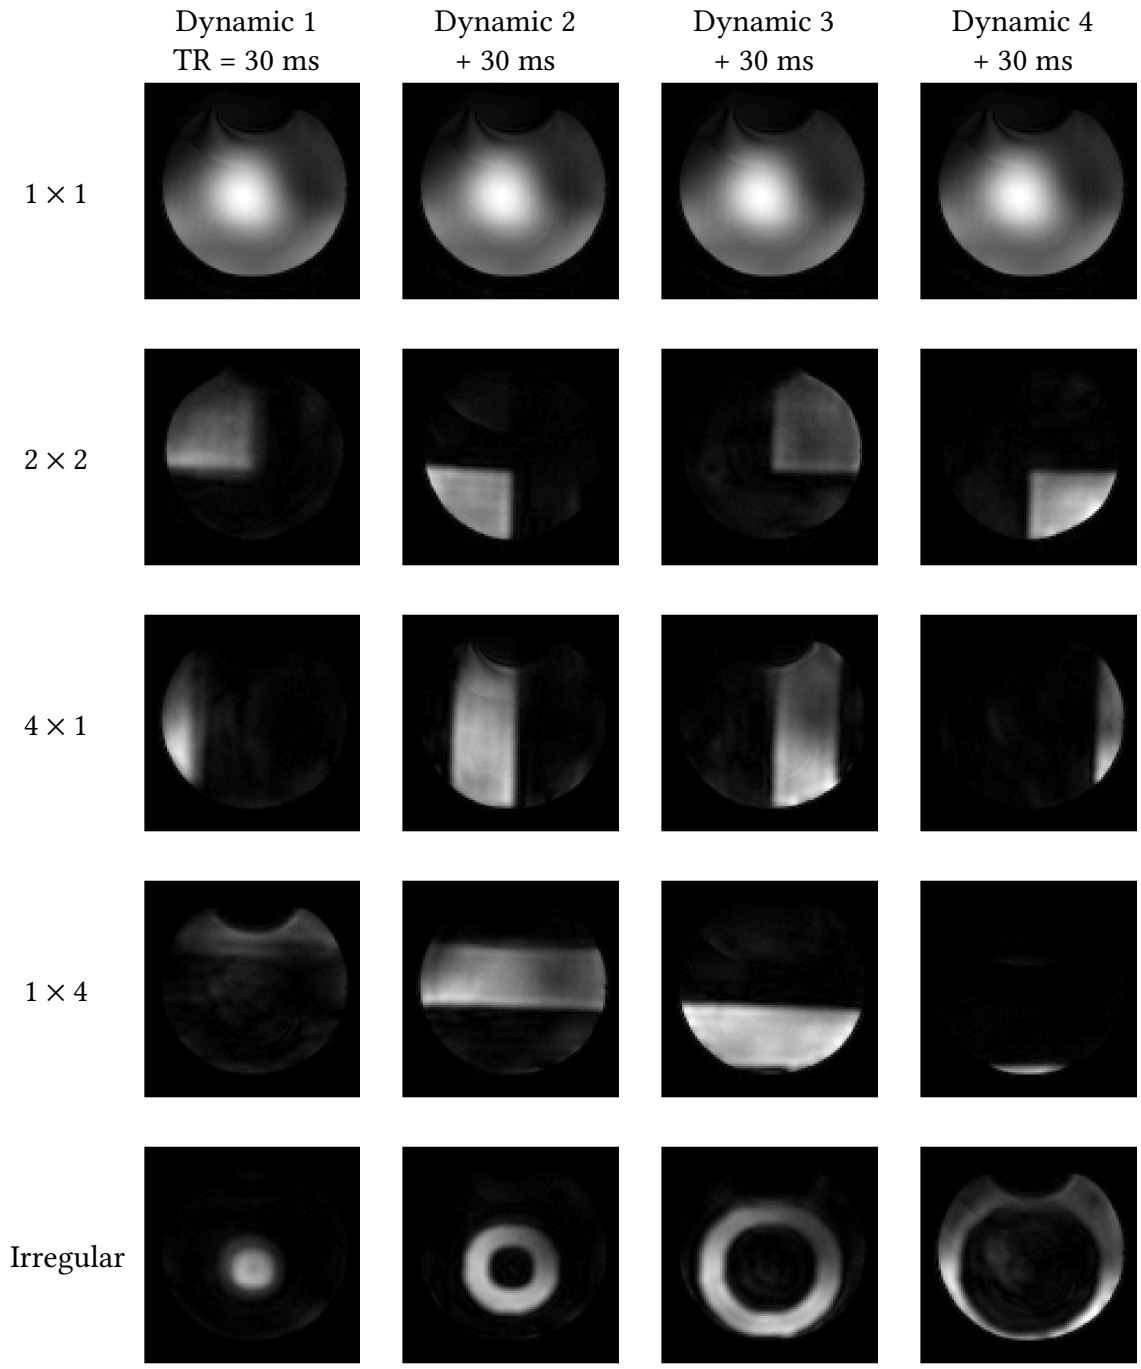

Figure 4.7: Example central axial slices from five separate 3D EPI acquisition, each with four dynamics, for five distinct parcellations.

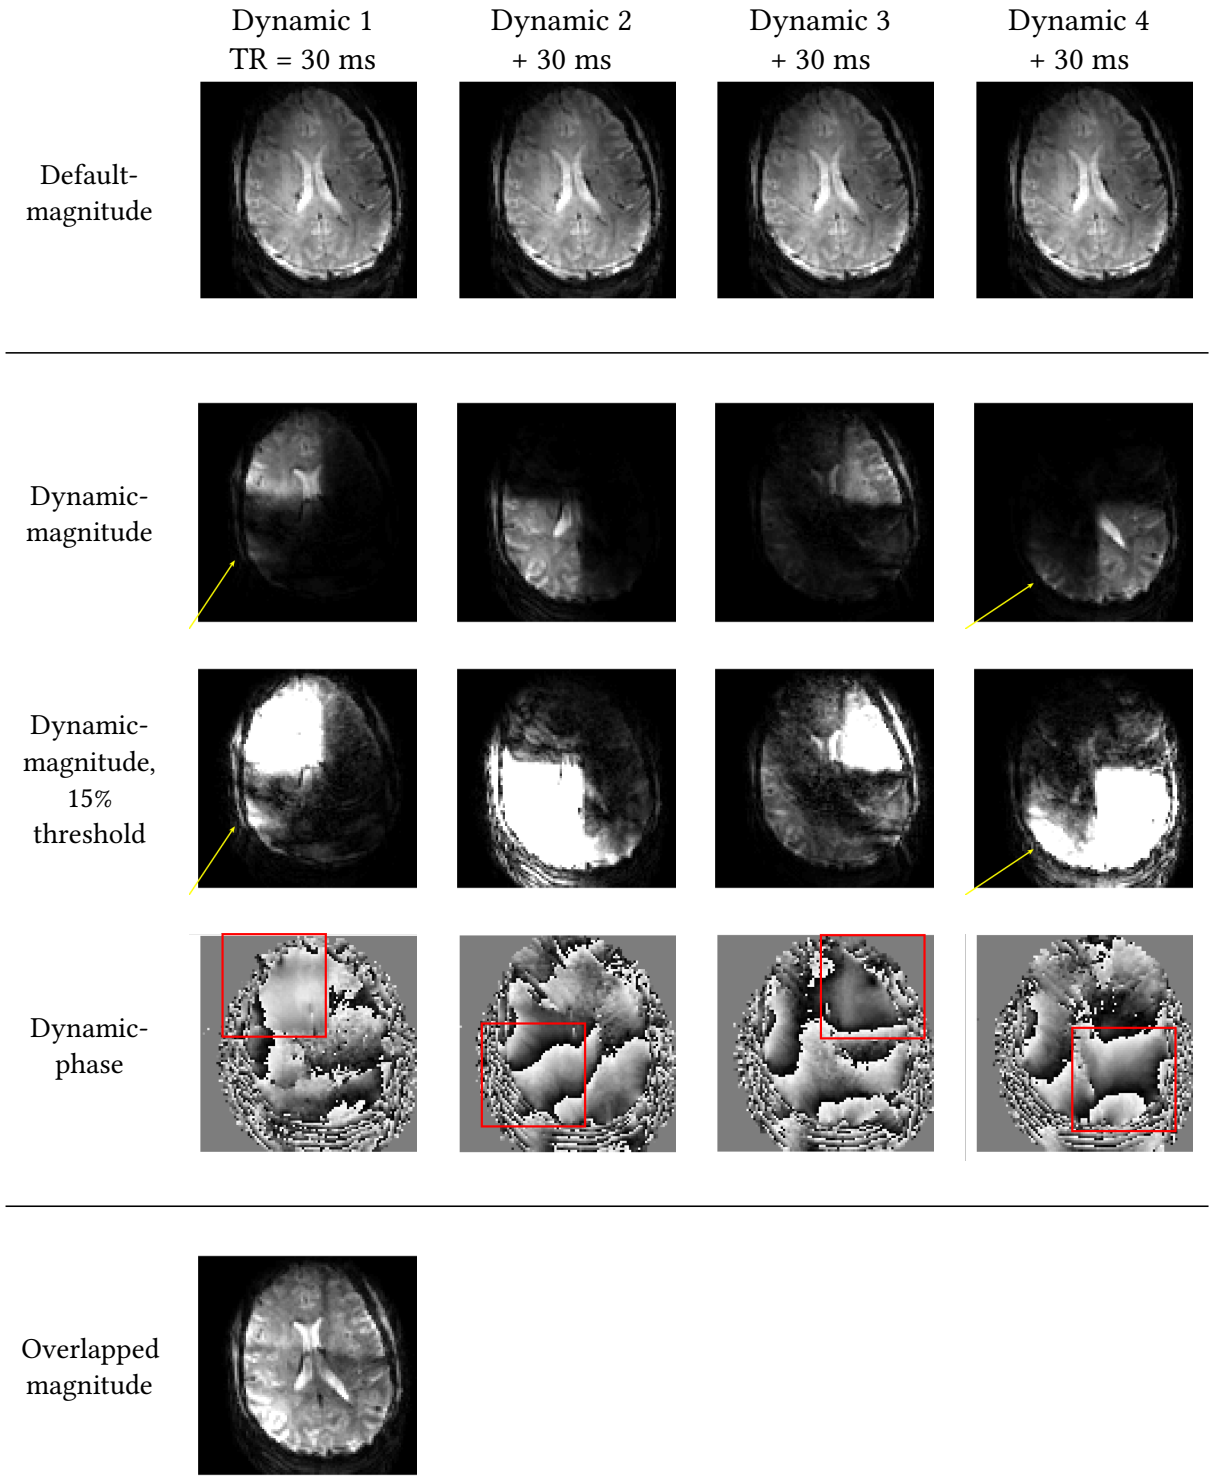

Figure 4.8: Example central axial slices from subject 5. 3D EPI acquisition with 3mm isotropic resolution each with four dynamics. First row (from top) - default global shim acquisition; second row - dynamic relative magnitude image; third row - dynamic relative magnitude image thresholded at 15% intensity level; forth row - dynamic phase image with overlaid red box to highlight shim box; fifth row - overlapped magnitude image. Yellow arrows in second and third rows highlight excitation error in outer volume region.

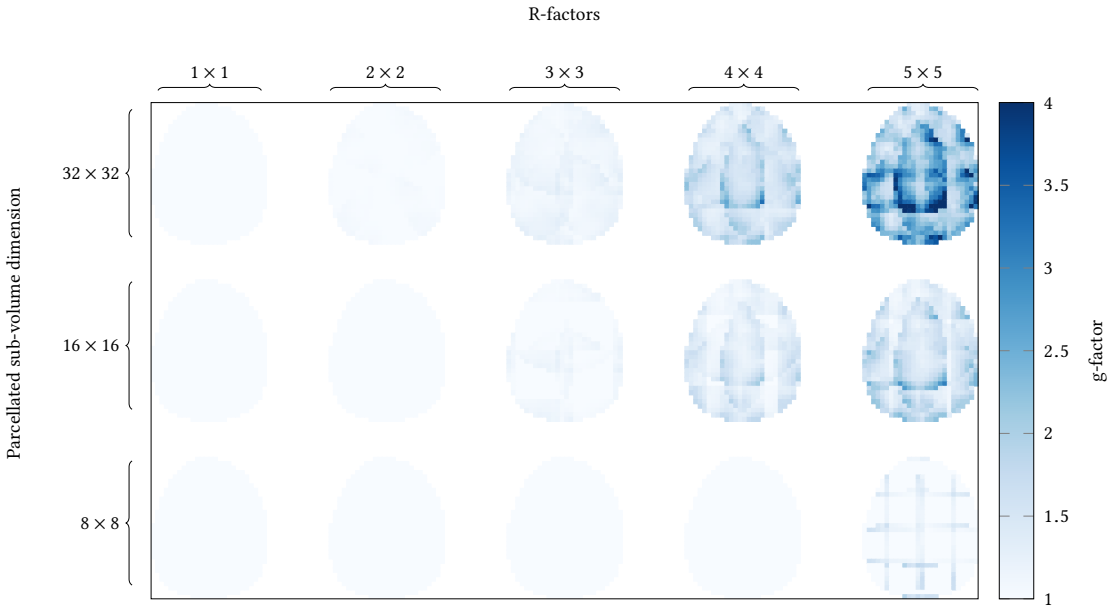

Figure 4.9: Simulated rSENSE g-factor maps (central axial slice) for three parcellations with five different SENSE acceleration factors using a 32-channel receive head coil. Whereas the g-factor becomes unacceptably large for high acceleration factors for a large sub-volume size, more compact parcellations can maintain low g-factors even for high acceleration factors.

acquisition time unnecessarily. An obvious method to reduce this would be to use reduced Field-of-View (rFOV) imaging in combination with rFOX, so called Inner-Volume Imaging (see Figure 4.6). This approach would only be compatible with regular cuboidal parcellations. Furthermore, using IVI with a compact FOV would limit the applicability of parallel imaging and require updates to the geometry of each dynamic. A more efficient approach known as *reduced* SENSE (rSENSE) uses knowledge of the limited rFOX in the SENSE reconstruction (*220*). Figure 4.9 shows example simulated g-factor maps for various parcellations and acceleration factors for a 32-channel receive head coil. The g-factor map for a single parcellated sub-volume was calculated by considering every voxel that would alias into the sub-volume rFOV, (see Equation 23 in (*9*)). Using a full Field of Excitation, using a reduced Field-of-View would result in significant aliasing. In an ideal case, using a reduced Field-of-Excitation can effectively remove *silent* voxels from the calculation of the g-factor map, as they will not contribute aliasing signal. As each parcellated sub-volume would be acquired separately, individual g-factor maps for each sub-volume were first calculated before being concatenated with every other sub-volume.

In the original paper proposing rSENSE, receive coil sensitivities were measured using the rFOX pulse used in the subsequent imaging acquisition. For dynamic parcellated imaging, it would be simpler to mask the receive coil sensitivities using a given rFOX for each sub-volume.

4.6 Conclusions

In this chapter a novel acquisition strategy was developed and implemented to improve the overall static field homogeneity by utilising reduced field of excitation with acceleration via parallel transmission in conjunction with dynamic first-order local shimming. Algorithms were introduced to further develop arbitrary parcellations with sub-volumes of varying size and shape to further improve shim efficacy. As the robust separation and identification of biochemicals in NMR spectra fundamentally depends on the homogeneity of the static field conditions (*34*), an application of the work in this chapter could first focus on MR spectroscopy at 7 T.

Chapter 5

Sub-Pulse Optimisation for Multidimensional Adiabatic Inversion

In the preceding chapters the use and design of spatially-selective pTx RF pulses was strictly limited to Amplitude Modulated (AM) waveforms. In this chapter, a spatially-selective pTx, Frequency Modulated (FM) pulse that is resilient to a significant level of B_1^+ inhomogeneity is designed and implemented. An optimisation scheme is sought that exploits the adiabatic condition to relax the spatial B_1^+ magnitude and phase patterns.

- Optimisation routines are published at https://github.com/cmirfin/Adiabatic3D

5.1 Introduction

Parallel Transmit affords the possibility to spatially tailor the phase profile. This is commonly exploited in RF shimming to generate transmit efficient (*90*, *91*) or spatially-uniform magnitude transmit fields (*221*). Furthermore, spatial phase variation is the basis of the Magnitude Least Squares (MLS) pulse design method, especially pertinent for gradient-echo imaging. Rather than specifying a target spatial phase (e.g. linear or quadratic), the spatial phase design is iteratively approached via the Variable Exchange method (*222*). This iterative approach also allows the pulse-designer to specify a degree of spatial phase smoothing to counteract intravoxel dephasing. The phase retrieval algorithm described herein as the Variable Exchange method is actually equivalent to the Gerchberg-Saxton algorithm (*223*), as noted by (*224*) with applications in crystallography (*225*) and optical imaging (*226*).

Frequency modulated (FM) RF pulses allow temporal evolution in the carrier frequency, usually in addition to amplitude modulation. The development of these pulses has been driven by use of inhomogeneous transmit field surface coils and the move to ever higher static fields (*146*).

As discussed in Chapter 2, adiabatic pulses are a special class of FM pulse whereby the frequency and amplitude modulations are designed to ensure the magnetisation vector remains spin-locked to the trajectory of the effective magnetic field. This occurs when

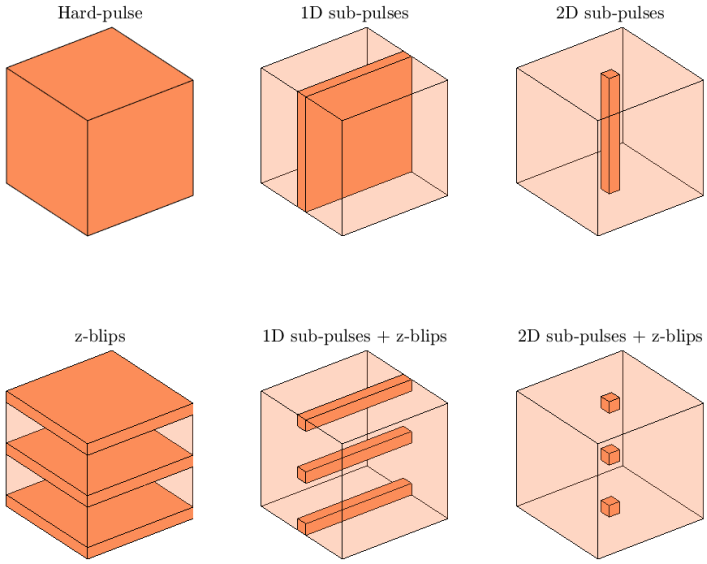

Figure 5.1: Schematic diagram to display the geometry of the pulse design problem. The top row shows the Blipped HS pulse with three different sub-pulses (hard-pulse, 1D spatially-selective, and 2D spatially-selective). Adding gradient blips in the frequency sweep direction generates selectivity in a third direction (bottom row). Due to the discrete nature of the waveform, replicants will be equally spaced along this dimension.

the precession frequency of the magnetisation vector about the effective field is much greater than the rate of change of the effective field itself (*227*), the so called *adiabatic condition* (see Equation 2.16 in Chapter 2).

To fully invert a magnetisation vector from $+M_z$ to $-M_z$ requires a frequency sweep from positive to negative; with an appropriate amplitude modulation to maintain the adiabatic condition throughout the pulse. A popular driving function to achieve this is the continuous Hyperbolic Secant (HS) pulse (*155*, *156*), parameterised by three constants $(\omega_{max}, \beta, \mu)$ and defined as:

$$\omega_1(t) = (\omega_{\max}\text{sech}(\beta\tau))^{1+i\mu}, \tag{5.1}$$

where $\tau = 2t/T - 1$ is the zero-centred normalised time parameter. Application of a slice-selective gradient with magnitude G_z yields an inverted region of thickness $\Delta z = \mu\beta/\pi\gamma G_z$. A *blipped* version of the HS pulse is shown in Figure 5.2, with 15 piecewise segments. The discretisation produces side-bands in the reciprocal (physical) space.

Jang et al (*228*) demonstrated the first volume selective adiabatic pulse, by complex multiplying each piecewise segment of a blipped HS pulse with a suitable 2D spatially selective small-tip angle pulse. In this manner, pulse amplitudes are scaled by the driving parent pulse (HS pulse), whilst the phases of the parent and sub-pulse are added. In addition, this approach is consistent with acceleration through parallel transmit, whereby

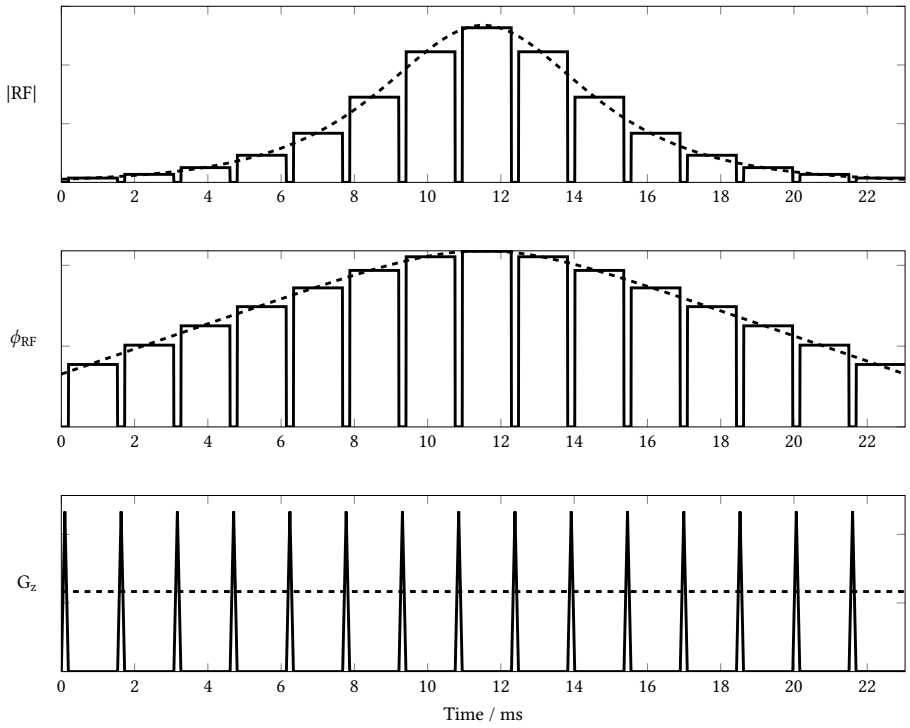

Figure 5.2: Blipped Hyperbolic Secant (HS) (solid line) and continuous variant (dashed line). Pulse amplitude $|RF|$, phase ϕ_{RF}, and slice-selective gradient G_z. for 15 segments.

each of the sub-pulses can be undersampled and hence reduced in length. The problem of multidimensional adiabatic pulse design is therefore simplified into one frequency-modulated dimension (parent pulse) and at most two amplitude-modulated dimensions (sub-pulses) (see Figure 5.1). If the blipped gradients are removed, the overall pulse is 2D selective (infinite prism), but the magnetisation vector follows a frequency-sweep trajectory. Conolly et al (*229*) earlier demonstrated a 2D variant with one frequency-modulated dimension (parent pulse) and a 1D sinc sub-pulse.

For this work, a basic HS parent pulse was assumed; and the design of 2D sub-pulses was considered. Variants on the existing Magnitude Least Squares optimisation scheme are proposed to relax the spatial B_1^+ magnitude and phase patterns on the sub-pulse design, and show that this is consistent with the design of a multidimensional adiabatic inversion pulse.

5.2 Theory

5.2.1 Sub-pulse design

The sub-pulse design used by Jang et al (*228*) utilised the Linear Least Squares (LLS) spatial domain method (*108*) to produce a target transverse magnetisation profile over

the central slice of the volume-of-interest:

$$\min_{b\in\mathbb{C}^{N_t\times N_c}} \sum_{m=1}^{N_m} (A_m b - d_m e^{i\varphi_m})^2 + R(b)\ . \tag{5.2}$$

where A is the concatenated system matrix and b represents the channel concatenated complex-valued waveforms to be designed, subject to regularisation function $R(b)$. The magnetisation target has been decomposed into complex polar form, where d_m is the magnitude transverse target at voxel m. Correspondingly, φ_m is the phase at voxel m.

The Magnitude Least Squares (MLS) method (*221*) relaxes the condition on the spatial target phase, allowing reduced magnitude target error to be achieved with reduced power,

$$\min_{b\in\mathbb{C}^{N_t\times N_c}} \sum_{m=1}^{N_m} (|A_m b| - d_m)^2 + R(b)\ . \tag{5.3}$$

For the purposes of adiabatic inversion sub-pulse design, the actual spatial phase pattern is unimportant as long as it is maintained across each sub-pulse. It could therefore be beneficial to exploit this flexibility to potentially improve inversion efficiency, either through improved selectivity or reduced transmit power.

Uniquely, in addition to phase relaxation, there is no requirement that the magnitude target d need be spatially uniform, as adiabatic pulses are robust to a certain level of inhomogeneity. The parameters and properties of the parent pulse define the minimum $|B_1^+|$, above which the adiabatic condition is satisfied. To avoid the use of hard constraints on the spatial profile of the magnetisation response, a modified least squares approach is desired.

A transition function $h : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is introduced to deweight inputs higher than some critical value x_0. The ideal functional form of $h(x)$ is stated as:

$$h(x) = \begin{cases} x, & \text{if } x < x_0 \\ x_0, & \text{if } x \geq x_0\ . \end{cases} \tag{5.4}$$

The function is used to define the Threshold Magnitude Least Squares (TMLS) minimisation problem as:

$$\min_{b\in\mathbb{C}^{N_t\times N_c}} \sum_{m=1}^{N_m} (h(|A_m b|) - d_m)^2 + R(b)\ . \tag{5.5}$$

Equation 5.5 is non-convex and the function $h(x)$ is non-differentiable at the transition point x_0, providing difficulties for gradient-based optimisation schemes. As such, the approach taken here is to iteratively relax the target magnitude in addition to the phase patterns during an alternating optimisation process. This straightforward extension of the original Variable Exchange method (*222*) is thus referred to as the Threshold Variable

Exchange method, and is summarised in Algorithm 7.

Algorithm 7: Threshold Variable Exchange method
Given an initial magnitude and phase spatial distribution the algorithm iteratively relaxes the spatial phase target. In addition, target magnitude voxels are updated with a new value if the current predicted magnitude exceeds a critical value.

$d_m, \varphi_m \quad m \in 1, \ldots, N_m$; *// initialise magnitude and phase target voxels*
while $(\tilde{\varepsilon} - \varepsilon)/\varepsilon > tolerance$ **do** *// check convergence criteria*
  $\varepsilon \leftarrow \tilde{\varepsilon}$;
  $\tilde{b} \leftarrow \underset{b \in \mathbb{C}^{N_t \times N_c}}{\arg\min} \sum (A_m b - d_m e^{i\varphi_m})^2 + R(b)$; *// solve LLS*
  $\tilde{\varepsilon} \leftarrow \sum (A_m \tilde{b} - d_m e^{i\varphi_m})^2 + R(\tilde{b})$; *// residual*
  $\varphi_m = \mathrm{Arg}(A_m \tilde{b})$; *// update phase target*
  if $|A_m \tilde{b}| > \theta_{min}$ **and** $m \in rFOX$ **then**
    $d_m = |A_m \tilde{b}|$; *// update magnitude target*
  end
end

Briefly, on each iteration, target magnitude voxels (d_m) within the reduced Field-of-Excitation (rFOX) that have an existing value greater than $\theta_{\min}$ are replaced with the expected predicted flip angle ($|A_m b|$). The critical flip angle $\theta_{\min} = \gamma |B^+_{1,\min}|$ is used for convenience. The ambition here is to guide the optimisation process towards a local minima where excitation inside the rFOX is above a given threshold, whereas the excitation is effectively suppressed outside.

An everywhere differentiable alternative to the objective function in 5.5 can also be formulated. As a first step, the Magnitude Squared Least Squares (MSLS) minimisation (*222*) can be expressed as:

$$\min_{b \in \mathbb{C}^{N_t \times N_c}} \sum_{m=1}^{N_m} (|A_m b|^2 - d_m^2)^2 + R(b) \; . \tag{5.6}$$

As the transition function $h(x)$ presented in Equation 5.4 is non-differentiable, a series approximation up to order N is used instead:

$$h_N(x) = \sum_{n=1}^{N} a_n \tanh(x/n) \; . \tag{5.7}$$

The functional forms of $h(x)$ and $h_N(x)$ are displayed in Figure 5.3. The Threshold Magnitude Squared Least Squares (TMSLS) minimisation is then defined as:

$$\min_{b \in \mathbb{C}^{N_t \times N_c}} \sum_{m=1}^{N_m} (h_N(|A_m b|^2) - d_m^2)^2 + R(b) \; . \tag{5.8}$$

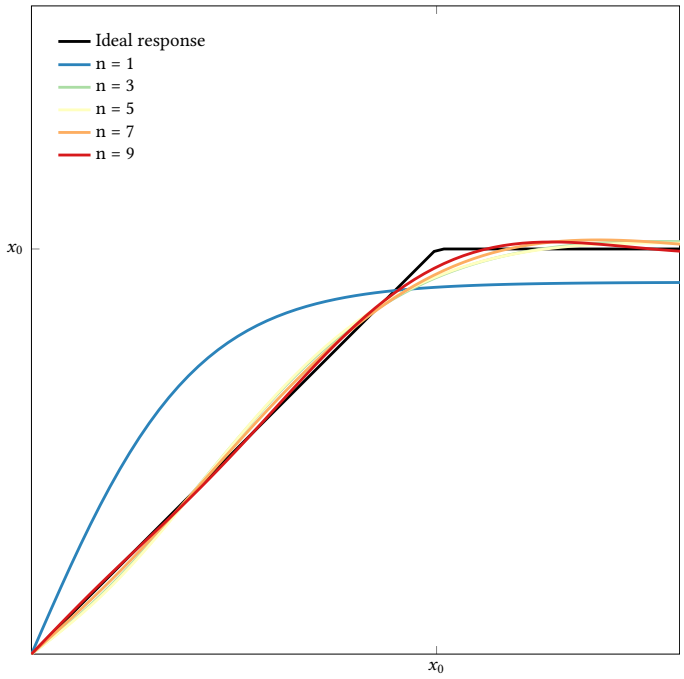

Figure 5.3: Ideal functional form of transition function $h(x)$ and series approximation $h_n(x)$ for $n = \{1, \cdots, 9\}$.

Tikhonov regularisation ($R(b) = \lambda b^H b$) can be used to constrain the amplitude of b. As Equation 5.8 is defined using complex variables, it is useful to transform to real variables for derivative calculations (*222*, *224*).

$$\min_{u \in \mathbb{R}^{2N_t \times N_c}} f(u) = \min_{u} \sum_{m=1}^{N_m} (h_N(u^T M_m M_m^T u) - d_m^2)^2 + \lambda u^T u \,, \tag{5.9}$$

where,

$$u := \begin{bmatrix} \mathrm{Re}(b) \\ \mathrm{Im}(b) \end{bmatrix}$$

$$M_m := \begin{bmatrix} \mathrm{Re}(A_m^H) & -\mathrm{Im}(A_m^H) \\ \mathrm{Im}(A_m^H) & \mathrm{Re}(A_m^H) \,. \end{bmatrix}$$

The gradient $\nabla f(u)$ can be calculated as:

$$\nabla f(u) = 4 \sum_{m=1}^{N_m} (h_N(u^T M_m M_m^T u) - d_m^2) h'_N(u^T M_m M_m^T u) M_m M_m^T u + 2\lambda u \,, \tag{5.10}$$

where,

$$h'_N(x) = \sum_{n=1}^{N} \frac{a_n}{n} \left(1 - \tanh^2\left(\frac{x}{n}\right)\right) . \tag{5.11}$$

The objective function in Equation 5.9 can minimised through use of a Quasi-Newton method, which uses the BFGS formula (*230–233*) to approximate the Hessian matrix.

5.3 Methods

5.3.1 Numerical Simulation

A simulation study was undertaken to invert an elliptical target region with major and minor axes of 3.2 and 2.4 cm respectively. To summarise, the four objective functions of interest were: Linear Least Squares (LLS); Magnitude Least Squares (MLS); Threshold Magnitude Least Squares (TMLS); and Threshold Magnitude Squared Least Squares (TMSLS). The pulses were optimised separately across both a 2D central axial slice and the full 3D volume with 3 mm isotropic resolution. Bloch simulations were performed to estimate the subsequent inversion profiles.

The parent pulse used in each case was derived from a normalised hyperbolic secant waveform discretised into 15 sub-pulses, each of length 1.2 ms. The shape parameters were fixed as $\beta = 400$ rad s^{-1} and $\mu = 3.9$.

To evaluate the performance of the optimised inversion pulse, only the longitudinal component of the magnetisation M_z is of interest (any residual transverse magnetisation will be crushed). In this case, the Normalised Root Mean Squared Error (NRMSE) is defined as:

$$\text{NRMSE} = \frac{||M_z - M_{z,d}||}{2||M_{z,d}||}, \tag{5.12}$$

where $M_{z,d}$ is the target inversion profile, and $||\cdot||$ represents the Euclidean norm.

To understand the behaviour of the various optimisation schemes with relation to the transition function $h(x)$, a Threshold Root Mean Squared Error (TRMSE) is defined as:

$$\text{TRMSE} = ||h(|Ab|) - d||, \tag{5.13}$$

5.3.2 Gradient Design

An Archimedean spiral-in excitation k-space trajectory consisting of 3.5 revolutions from $k_{\max} = 80$ rad m^{-1} was used for all experiments. A smooth prephasing gradient G_p was constructed to re-centre k-space prior to each sub-pulse. To maximise the smoothness of the waveforms, the prephasing gradients for each axis were found by

minimising the function first-derivatives subject to specific equality constraints:

$$\begin{aligned} G_p &= \arg\min||\dot{G}_p(t)|| \\ \text{s.t.} \quad &\int_0^{t_p} G_p dt = -\int_{t_p}^{T} G dt \\ &G_p(t=0) = 0 \\ &G_p(t=t_p) = G(t=t_p) \\ &\dot{G}_p(t=0) = 0 \\ &\dot{G}_p(t=t_p) = \dot{G}(t=t_p)\,. \end{aligned} \tag{5.14}$$

The first constraint ensures that the phase accrued under the prephasing gradient reverses the phase accrued due to the main gradient waveform $G(t)$. The additional four constraints are boundary constraints at the start and end of the prephasing waveform. The duration of the prephasing gradient t_p was manually adjusted to obey maximum slew-rate constraints.

A measured GIRF was used to account for the characteristics of the gradient system (*37*) (as described in Chapter 3). The maximum gradient amplitude and slew rate did not exceed 5 mTm^{-1} and 180 $\mathrm{Tm}^{-1}\mathrm{s}^{-1}$. Each sub-pulse was replicated and concatenated into 15 segments (see Figure 5.4).

5.3.3 Experimental Validation

To demonstrate the spatial selectivity of the optimised adiabatic inversion pulses, an inversion recovery experiment (see Figure 5.4) was performed across three healthy volunteers using a Philips 7 T Achieva (Best, Netherlands) connected with an 8-channel transmit, 32-channel receive head-coil (Nova Medical, Wilmington, USA). Each optimised pulse was used as a prepulse inversion, followed by spoiler gradient. The inversion delay time was incremented between 400 and 2000 ms in steps of 100 ms, and between 2000 and 5000 ms in steps of 500 ms. A Turbo Field Echo (TFE) readout was used with 2 mm isotropic resolution, 10° flip-angle, TR = 30 ms, TE = 2.03 ms. Lines of k-space were from acquired from low to high.

Whole-brain multichannel transmit sensitivities were derived from a relative combination of a quadrature-mode sensitivity map and 15° flip angle gradient-echo images from each channel as described in Section 1.5.3. The quadrature-mode sensitivity map was measured using the AFI method (*77*) with a flip angle of 60° and repetition times of 30 ms and 150 ms. Off-resonance maps were acquired before and after 2nd-order B_0 shimming using the dual-TE method with echo-times 1.5 ms and 2.5 ms.

The experiments were also repeated on a spherical water filled phantom, whereby the peak amplitude $\omega_{\max}$ of the adiabatic prepulse was varied between 0.0 and 15.0 μT.

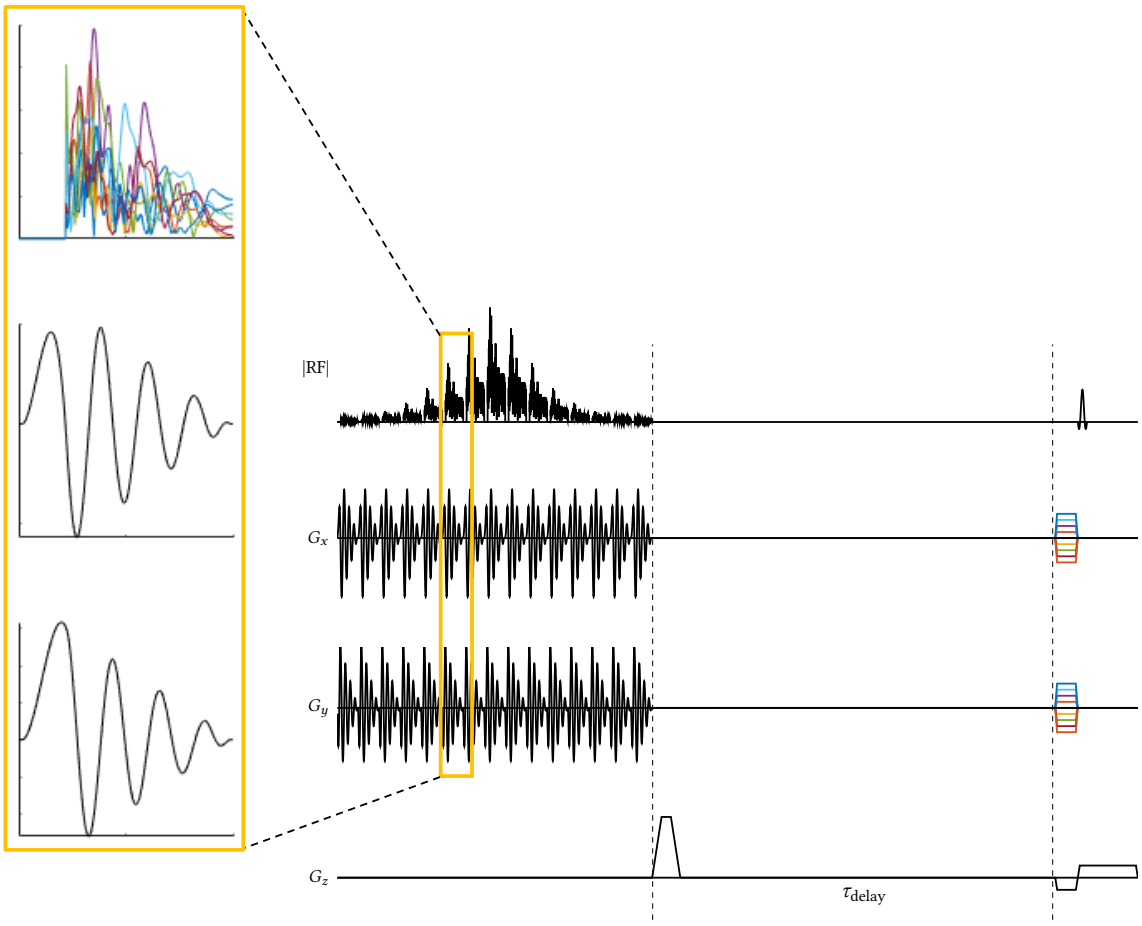

Figure 5.4: Inversion Recovery Sequence with 2D selective prepulse inversion. High slew-rate spiral gradients split into 15 sub-pulses are used to permit 2D selectivity in the plane orthogonal (e.g. $x - y$) to the frequency-sweep direction (z). If 3D selectivity is desired blipped, z-gradients can be applied. A crusher is used to ensure destruction of remnant transverse magnetisation, before a user-defined delay time preceding the readout module. Inset figure shows one sub-pulse (multichannel amplitudes) and 2D waveforms.

Real-valued signed images (normalised in the range [-1,1]) were calculated by reversing the sign of the magnitude image where a step change was present in the phase image. Difference images were then calculated by subtracting this from a Turbo Field Echo (TFE) magnitude image acquired without prepulse preparation.

Finally, individual sub-pulses for each of the optimisation approaches were used as Small-Tip Angle excitation pulse in a gradient-echo readout. Single-slice axial magnitude and phase images were acquired with 1 mm in-plane resolution using a nominal 10° flip angle, TR = 7.8 ms, TE = 4.93 ms. As such, the length of the (sub)-pulse duration in each case was 1.2 ms.

5.3.4 Modifications to scan control software

The existing scan control software for the Philips 7 T Achieva can load a set of pTx excitation pulses and associated gradient waveforms from a text file. New RF and gradient objects were defined to allow this functionality to be extended to prepulses if requested. In this manner all user-defined pTx pulses required for a given sequence (e.g. prepulse, excitation, and echo) are maintained in a single text file that can be supplied by an external program (see Appendix A for further details). Frequency-modulated pulses were implemented as phase modulations.

5.4 Results

5.4.1 Numerical Simulation

A sub-pulse duration of 1.2 ms led to a total pulse length of 18 ms including prephasing blips. The Threshold Root Mean Squared Error as a function of relative total integrated power is displayed in Figure 5.5, comparing the performance of LLS, MLS, TMLS and TMSLS optimisations. The trade-off between total integrated power and TRMSE was evaluated for eight values of the regularisation parameter λ on a logarithmic scale. Optimisation across a single-slice showed very little variation between each of the cost functions across the entire volume. Volume optimisation produced more diffuse results, with the TMLS achieving the lowest TRMSE across all regularisation values.

These results were similarly reflected in the NRMSE values shown in Figure 5.6. A NRMSE value of 0.23 using the TMLS objective function was the lowest across all optimisations, representing a 18% improvement compared to LLS. Again, this difference was only apparent when the optimisation was performed over the whole volume rather than a single slice. As can be seen in Figure 5.6, each curve does not converge to a horizontal asymptote with increasing transmit power, and instead have a minimum point. This behaviour is symptomatic of the trade-off between inversion error within and outside of the passband. As the parent driving pulse is adiabatic, inversion error within the

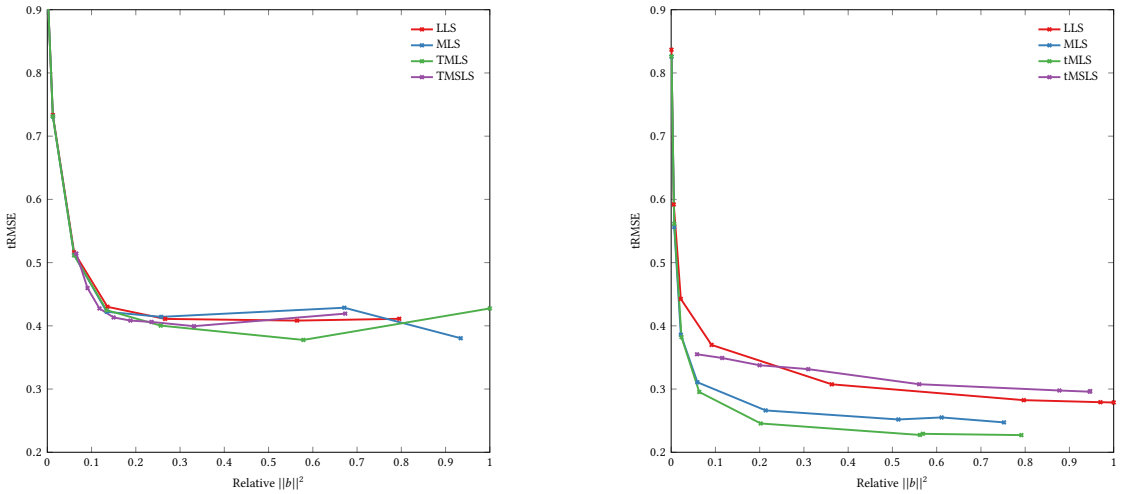

(a) Slice optimised **(b)** Volume optimised

Figure 5.5: TRMSE over entire volume-of-interest using 2D slice optimisation (a) or full volume optimisation (b). Range of regularisation parameter used: $\lambda \in [10^2, 10^5]$.

passband cannot be improved beyond a certain point with increasing power. However, due to the inherent low resolution of designed subpulses, increasing the total transmit power (and hence B_1^+ field) will lead to excitation in the stopband and an increase in the overall Normalised Root Mean Squared Error.

The convergence of the (Threshold)- Variable Exchange method for optimising the MLS and TMLS objective functions is shown in Figure 5.7. All 3D optimisations were completed in less than three minutes (which represented 5×10^4 voxels), although this was often much less for convergence of the TMLS function (see Figure 5.7). For comparison, this is generally 50-times more computational time than standard LLS due to additional phase retrieval iterations.

5.4.2 Experimental Validation

Gradient-echo magnitude and phase images for a 1.2 ms sub-pulse optimised with each of the optimisation approaches are displayed in Figure 5.8. Whereas LLS demonstrates uniform phase within the target region, each of the subsequent methods exhibit spatial phase variation within the target region. The mean signal intensity across the target region was 38% higher using TMLS rather than LLS, and 21% greater than MLS alone.

Figure 5.9 displays inversion recovery TFE magnitude images from a spherical water phantom acquired using a range of peak prepulse inversion powers with a 200 ms inversion delay. The transition region between the selected and unselected region is highlighted as a null signal band, where the signal sign is flipped. The first image in Figure 5.9 was acquired without a prepulse inversion ($\omega_{\max} = 0.0$), creating a reference TFE image. Difference signed (magnitude) images from this reference are shown in Figure 5.10, which exhibit positive values in regions of inversion. The image difference method

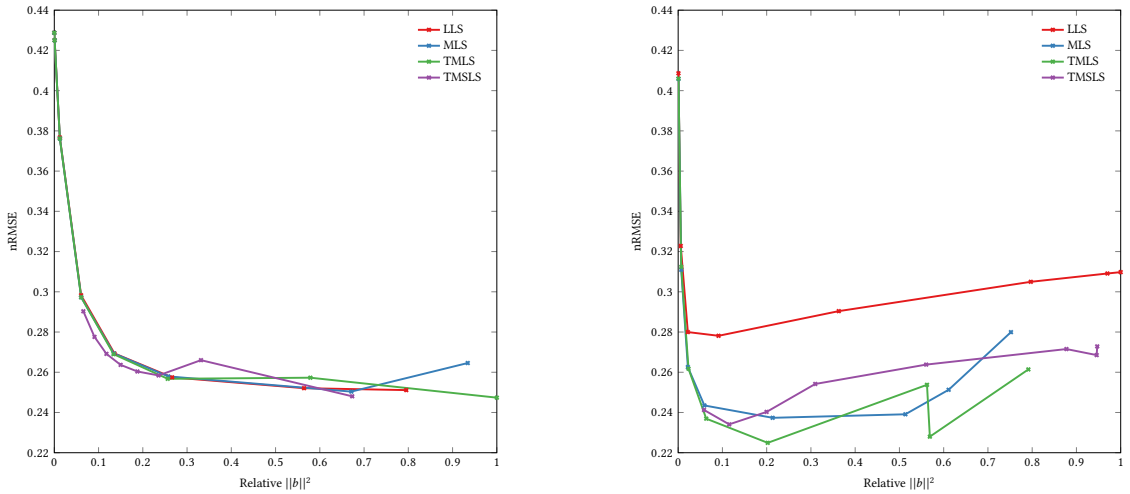

(a) Slice optimised **(b)** Volume optimised

Figure 5.6: Inversion NRMSE over entire volume-of-interest using 2D slice optimisation (a) or full volume optimisation (b). Range of regularisation parameter used: $\lambda \in [10^2, 10^5]$.

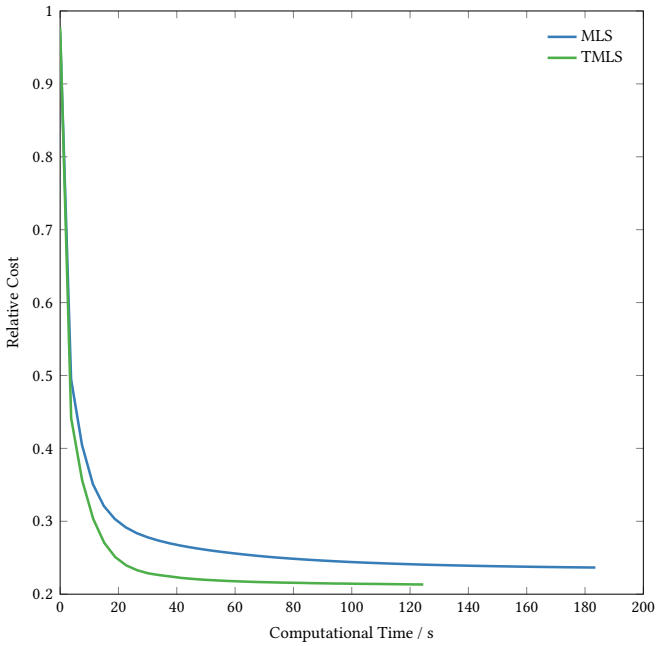

Figure 5.7: Convergence of the (Threshold) Variable Exchange method for MLS and TMLS objective functions. Representative of 3D optimisation with 5×10^4 spatial voxels. As the objective functions are non-convex, only a local minima can be assumed.

also leads to spurious edge enhancement.

In vivo, 3D multi-slice magnitude and phase TFE images are shown in Figure 5.11 using a 2D selective adiabatic prepulse inversion with 400 ms delay time. In-plane selection is demonstrated and maintained across each of the slices. A change in contrast is noticeable between the selected and unselected regions.

Figures 5.12 5.13 display the inversion recovery curves generated by varying the delay time between the prepulse inversion and TFE readout in vivo and separately in a spherical water phantom. In each case, two regions-of-interest were manually defined inside and outside of the target inversion region, encompassing multiple tissue types for the in vivo experiment. Although the signal outside of the target inversion area remained strictly positive, some signal suppression is evident at short delay times. The signal curve from within the target inversion region exhibited a zero-crossing at approximately 650 ms for the in vivo experiment, and approximately 1100 seconds for the spherical water phantom.

5.5 Discussion

A new method for designing multichannel, multidimensional sub-pulses for an adiabatic inversion prepulse was conceived and implemented. By iteratively relaxing the spatial magnitude target in addition to spatial phase relaxation, reduced inversion NRMSE can be achieved over an extended 3D volume. To this end, applying the Threshold Variable Exchange method to the TMLS objective function proved most effective.

To achieve this, a (non-linear) transition function was used to deweight inputs higher than the critical value determined by the adiabatic condition, the parameters of which are set by the parent-pulse (in this case a simple HS pulse). Threshold regression (and relatedly logistic regression) are commonly used in multiple areas of science and economics (*234*)). It should be noted that this approach is distinct from weighted least-squares approaches, whereby spatial voxels are weighted a priori, with no knowledge of the response of the system to a given input.

Directly optimising the thresholded magnitude squared (TMSLS) function did not result in either reduced threshold error (TRMSE) or inversion errors (NRMSE). Bloch simulations and gradient-echo imaging (Figure 5.8 for the TMSLS optimisation revealed ringing in the unselected regions that is present in the TMLS solution. One problem is that the series approximation $h_N(x)$ is applied to the squared value $|A_m b|^2$, subtly changing the scaling of the algorithm in comparison with TMLS. Notwithstanding, as both the TMLS and TMSLS optimisations are susceptible to local minima it could be worthwhile investigating the outcome of multiple initialisations.

High SAR inversion pulses limit the shortest inversion delay possible and can increase the sequence TR, reducing temporal SNR. Tikhonov regularisation was used to

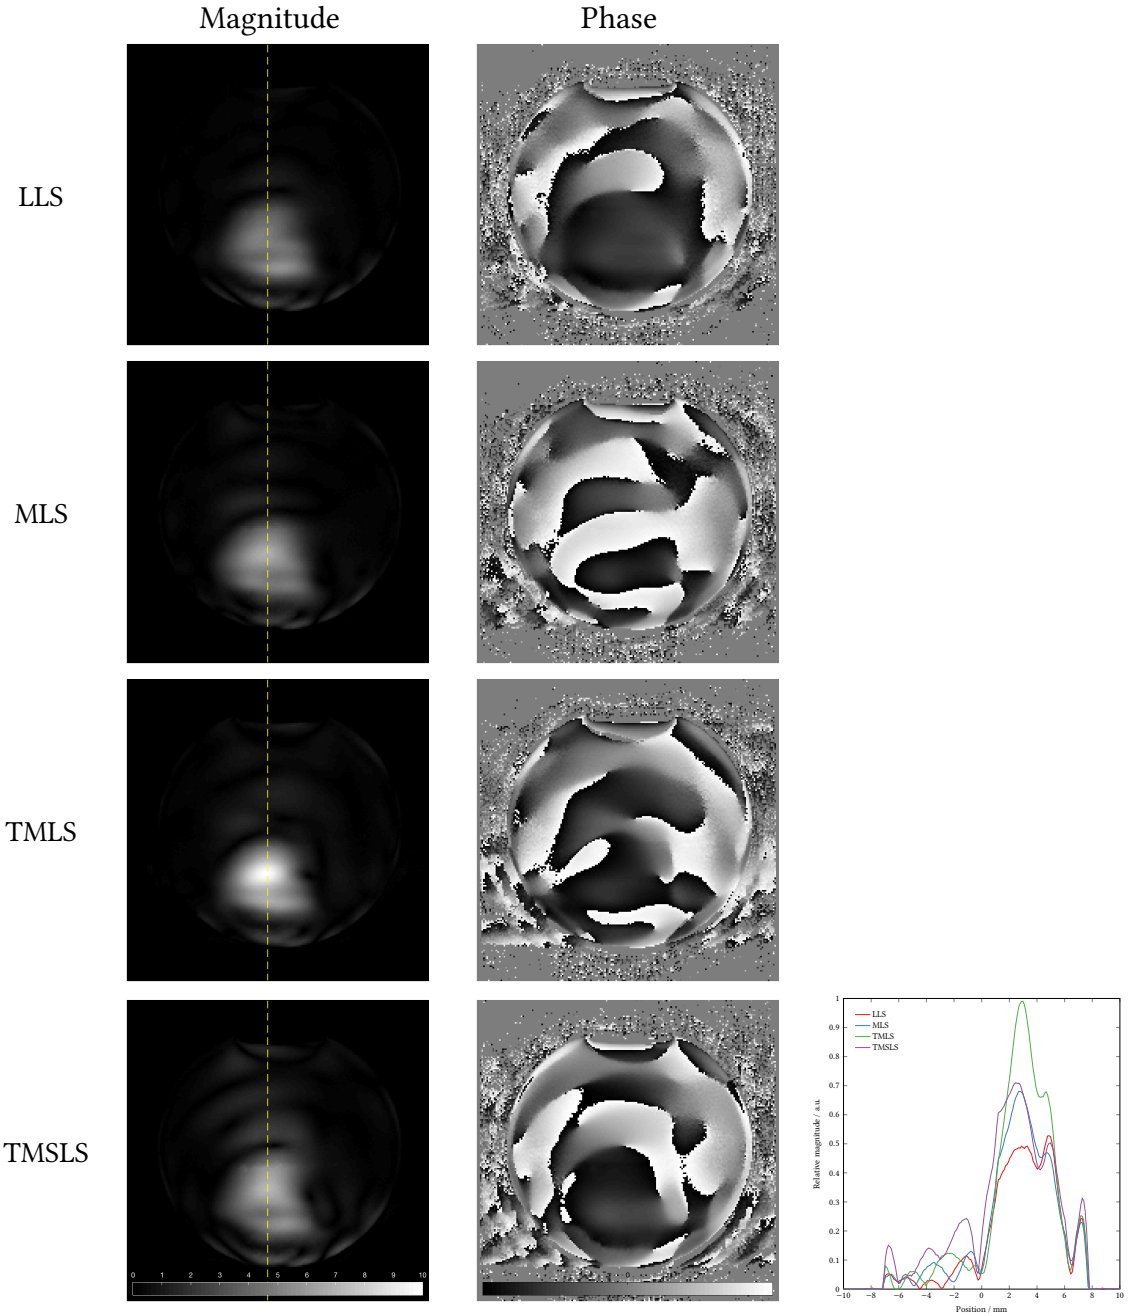

Figure 5.8: Gradient-echo central axial slice to show STA for 10° flip angle target excitation profile for individual sub-pulse (pulse duration = 1.0 ms (+prephasing), k_{max} = 80 rad m^{-1}, λ = 1.0). Relative magnitude profile comparison also shown bottom right.

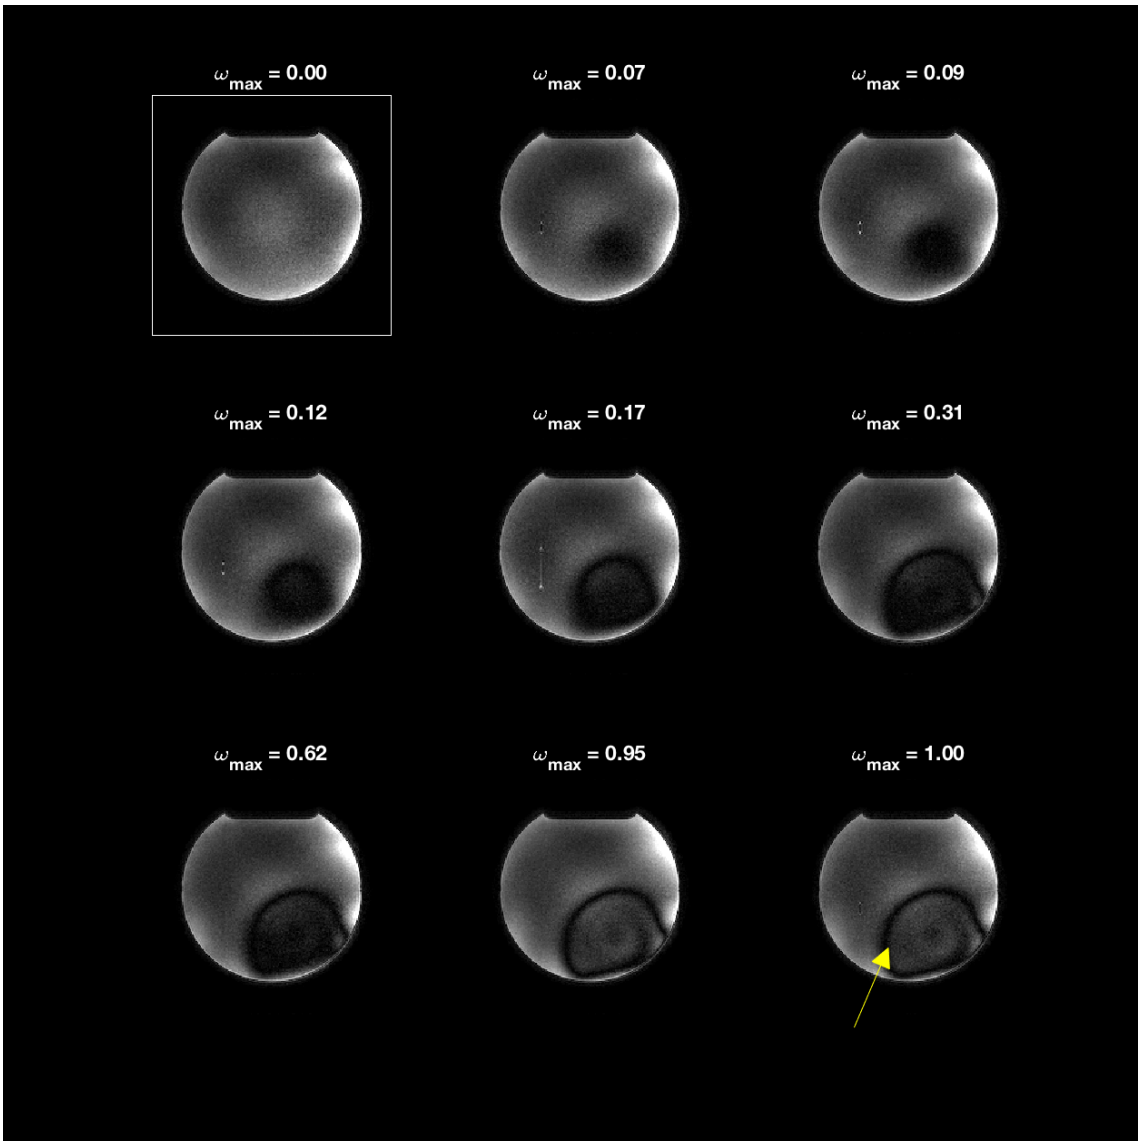

Figure 5.9: Magnitude images from TFE inversion recovery with different relative peak prepulse inversion powers ω_{max} (up to a maximum 15 μT). Central axial slice shown through phantom, using inversion delay of 400 ms. Yellow arrow indicates dark transition region between selected and unselected inversion regions.

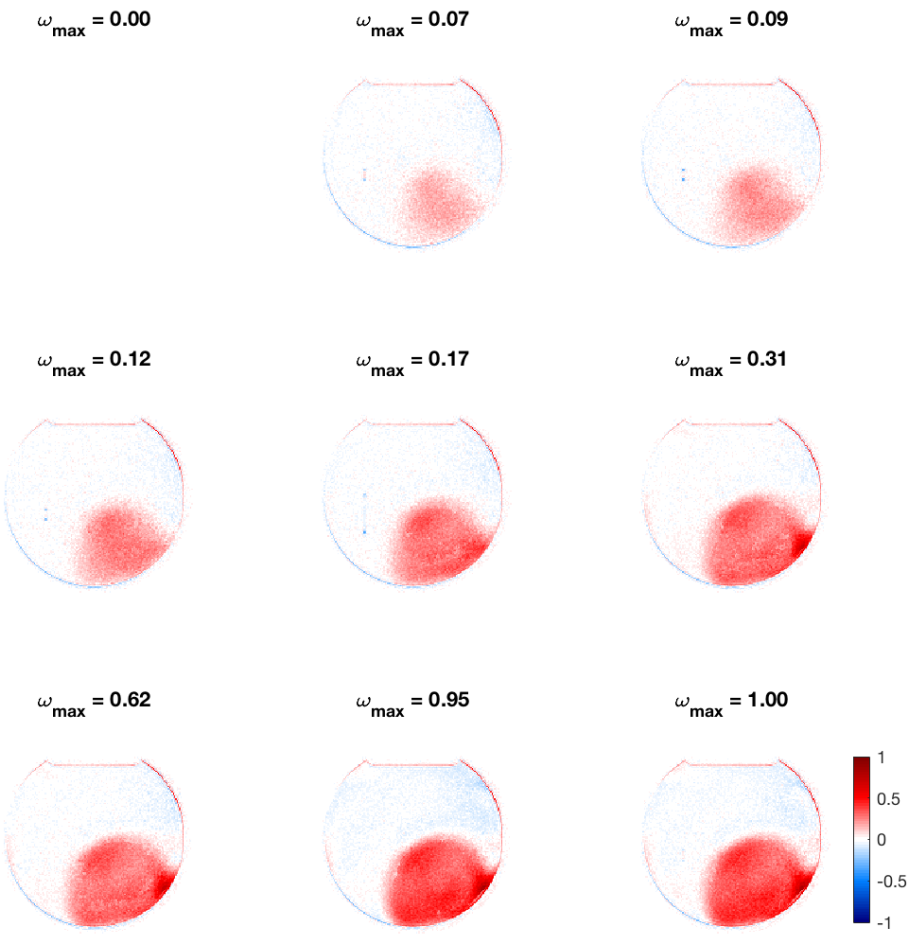

Figure 5.10: Relative difference from signed TFE images with and without prepulse. Prepulse inversion peak powers are denoted $\omega_{\max}$, (up to a maximum 15 μT). Central axial slice shown through phantom using inversion delay of 200 ms. A positive value indicates inverted signal. It is noted that the difference method also leads to spurious edge enhancement in regions where no inversion occurred.

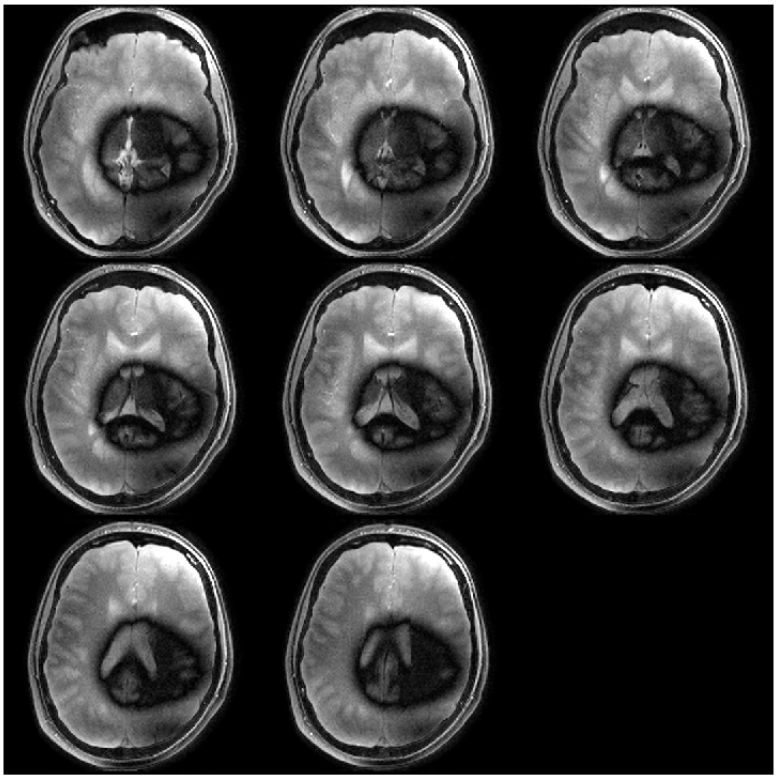

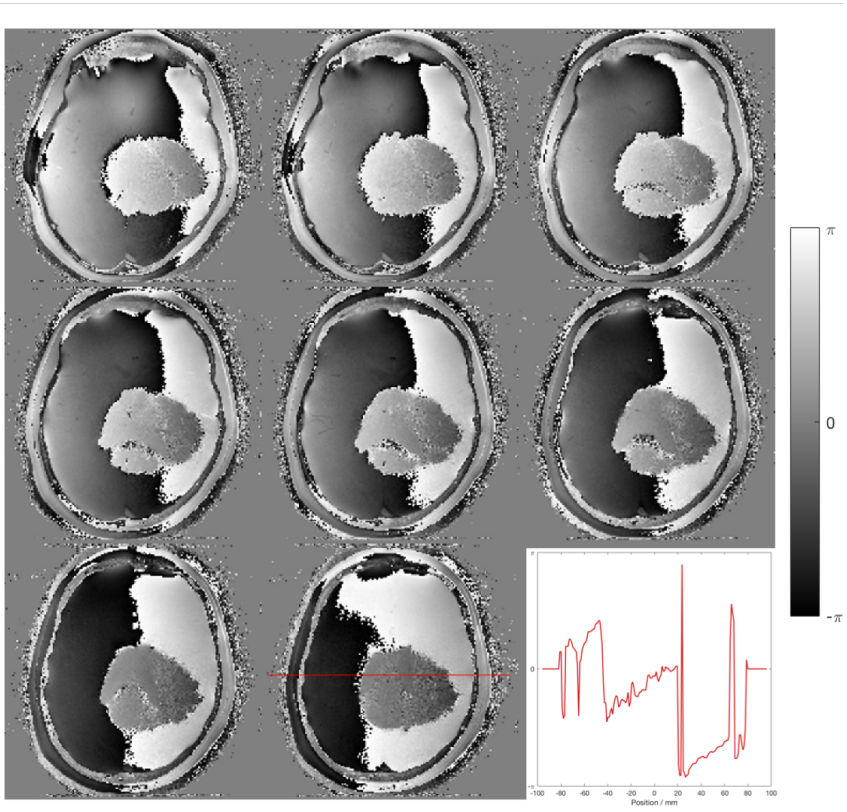

Figure 5.11: Magnitude (a) and phase (b) variation through axial slices of 3D TFE readout (FOV = 192 × 192 × 26 mm with 1 × 1 × 3 mm resolution, TR = 30 ms, TE = 2.0 ms) with optimised 2D selective adiabatic inversion. Inversion delay time 400 ms.

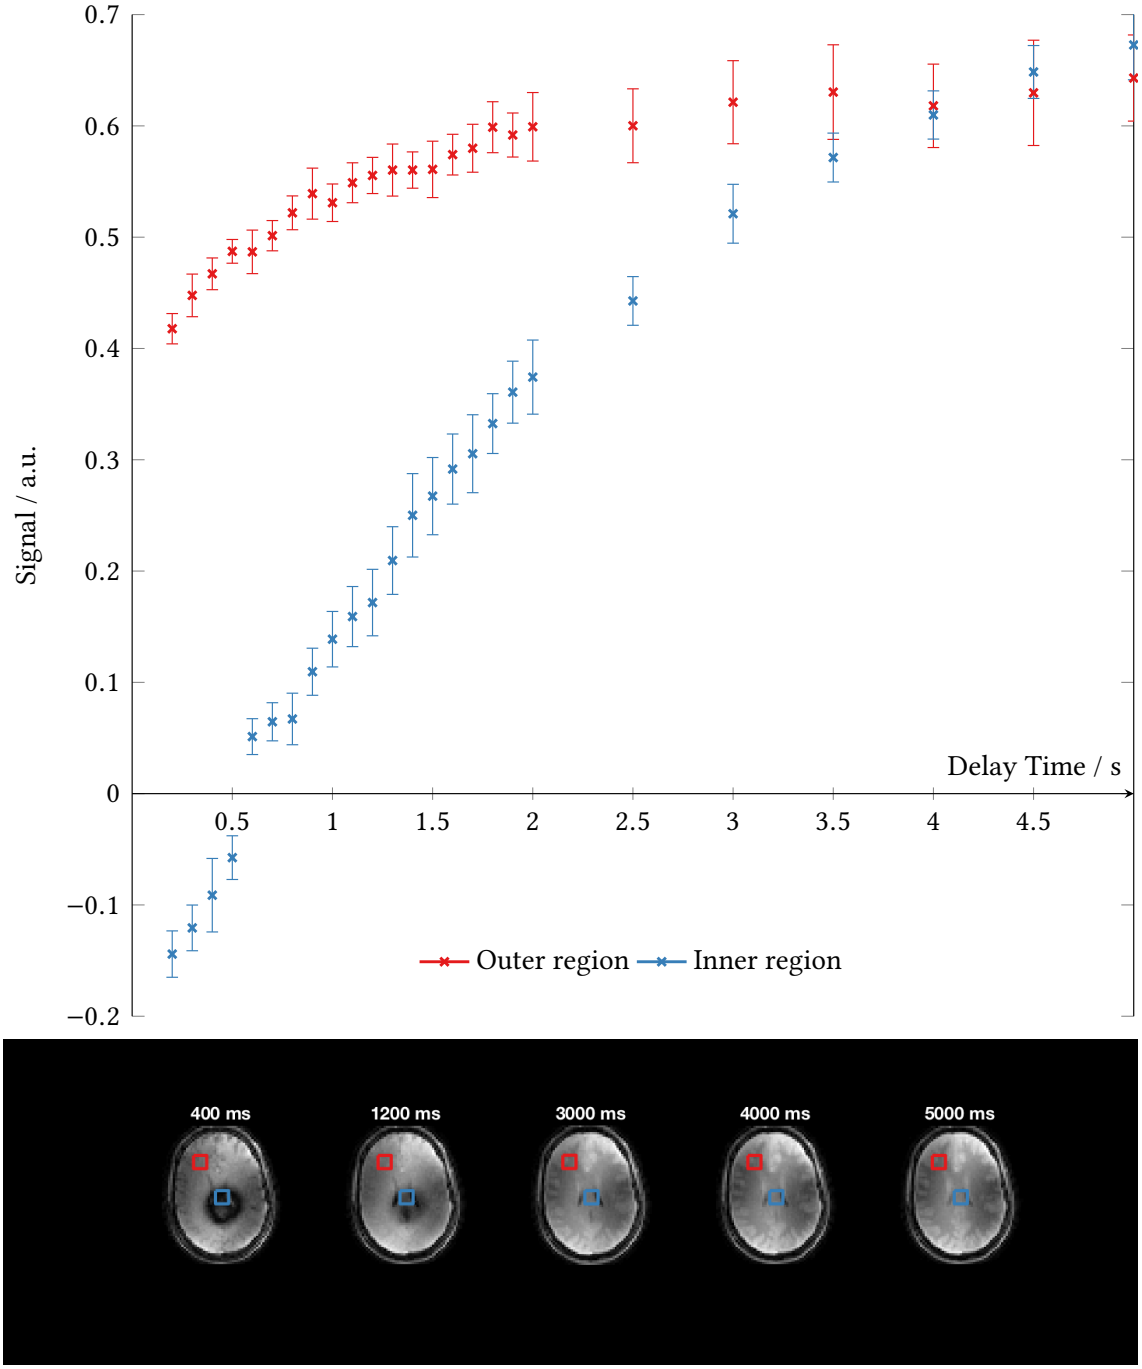

Figure 5.12: Inversion recovery curve from single-slice TFE readout with selective adiabatic inversion prepulse in vivo. Plotted for two regions of interest (inner and outer regions), displayed on corresponding magnitude images. Error bars reflect standard deviation in signal across ROIs.

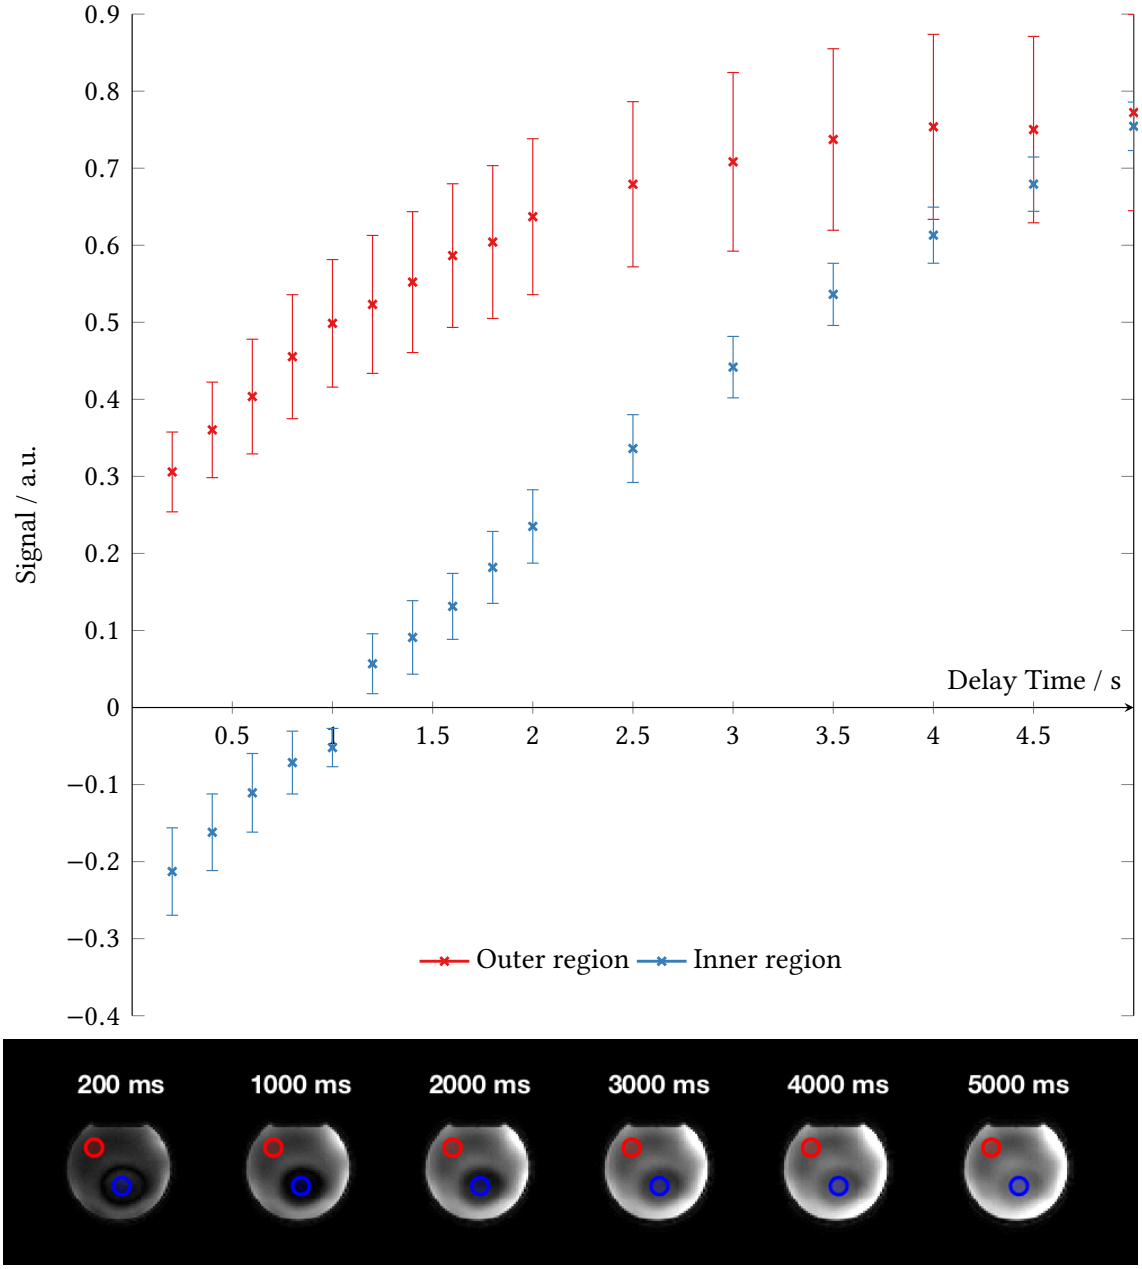

Figure 5.13: Inversion recovery curve from single-slice TFE readout with selective adiabatic inversion prepulse, for spherical phantom of water. Plotted for two regions of interest (inner and outer regions), displayed on corresponding magnitude images. Error bars reflect standard deviation in signal across ROIs.

control total transmit power of the designed sub-pulses as a proxy for SAR. However, it would be straightforward to constrain absolute channel powers and local SAR limits if necessitated (*113*). Due to the nonlinear nature of these constraints, it would be expected that this would add significantly to the optimisation duration (*224*). In particular, the additional computational complexity of this approach is dependent on the degree of overestimation factor used in the generation of the Virtual Observation Point set itself (*54*, *55*).

The short duration of the sub-pulses dictated the maximum extent permissible in transmit k-space due to slew-rate constraints, limiting the effective spatial resolution in the spatial domain. Increasing the duration of the sub-pulses to accommodate an increase in transmit spatial resolution would increase the entire pulse duration to unrealistic lengths. There is no foreseeable alternative when using whole-body gradient coils due to physiological limitations. Smaller, head-only or localised gradient coils could allow higher slew-rates (*235*, *236*), but are not widely available at this time. Regardless, a measured GIRF was required to correct for the low-pass nature of the gradient amplifiers.

It would be interesting to see how the sub-pulse optimisation suggested in this work would compare to optimising the entire pulse. The advantage of such an approach would be to exploit the additional temporal degrees of freedom to potentially increase inversion efficiency and reduce power deposition. This would require the use of full Bloch operator and an appropriate optimisation framework. Rosenfeld et al (*165*) demonstrated an optimal control approach to single-channel adiabatic inversion pulse design, which utilised a novel cost function to encourage adiabatic pulse properties. A pTx extension could be realised, although the obvious drawback of this approach would be large increase in computational complexity. A basic HS pulse was used as the basis parent-pulse in this work due to its easy implementation and robustness. It would be straightforward to replace this with C-FOCI (*163*), TR-FOCI (*164*) or HSn class (*162*) adiabatic pulses if desired.

The modifications made to the scan software allowed a fully customisable pTx prepulse. In this work, a regular quadrature mode excitation pulse was used in the TFE readout. It would be advisable to apply RF shimming or a tailored pulse (e.g.(*237*), (*238*)) to homogenise the excitation B_1^+ field.

5.6 Conclusions

In this chapter a Threshold Magnitude Least Squares (TMLS) optimisation was introduced to design spatially-selective 2D pTx sub-pulses compatible with existing multidimensional adiabatic inversion pulses. Relaxing both the magnitude and phase transmit patterns can lead to reduced inversion error over a 3D volume.

Chapter 6

Conclusions

6.1 Summary

This thesis has been concerned with the optimisation of multidimensional pulse design with parallel transmission, and the computational approaches to achieve this.

At the beginning of the project, there was limited capability provided by the 7 T vendor host system to handle the unique workflow presented by pTx. As such, a user-interface and algorithms were first developed to provide the functionality of RF shimming routines. Static RF shimming is extremely robust in experimental conditions, and has the advantage that it doesn't change the underlying sequence timings. However, due to the low-frequency content of transmit coil sensitivity patterns, the full benefits of parallel transmission are more fully exploited when variation over time is permitted.

For 3D excitation, despite acceleration due to parallel transmission, the topology and velocity of the volumetric k-space trajectory will determine the pulse duration due to limited absolute gradient slew-rates. For this reason, joint optimisation of parallel transmission RF pulses and k-space trajectory is advantageous in reducing pulse duration. Expressing the gradient waveforms (or equivalently the k-space trajectory) as a linear combination of basis functions is essential to reduce the number of free parameters in the optimisation. The alternate minimisation strategy used in Chapter 3 to solve this joint optimisation problem was shown to be well suited to heterogeneous computing systems (CPU + GPU).

Gradient system imperfections led to distortions in the excitation profile in the initial stage of this work. This was effectively removed by measurement and application of a Gradient Impulse Response Function (GIRF). Due to linearity considerations, the GIRF could actually be applied directly to the gradient waveform basis functions.

Optimised pTx RF pulses operate on timescale of milliseconds. A second timescale for pTx exists at the inter-sequence level, that is RF objects across multiple TRs or dynamics. In Chapter 4, dynamic pTx RF pulses were implemented to allow the update of inter-sequence pulses. This facilitated the deployment of a novel parcellated dynamic B_0 shimming method. In particular, multidimensional RF pulses have the advantage that

they can be designed to excite an arbitrary shape, whilst leaving outer volume magnetisation unperturbed. B_0 shimming parcellation algorithms were developed to exploit this flexibility. As multidimensional RF pulses typically have long durations (even with undersampling), degradation due to off-resonance effects is more pronounced. Including a predicted shimmed off-resonance map in the spatial domain pulse design method for each separate dynamic accounted for this.

In Chapter 5, a Threshold Magnitude Least Squares (TMLS) optimisation was introduced to design 2D pTx sub-pulses compatible with a parent adiabatic function. This optimisation was based on the recognition that the target phase of the sub-pulse design problem is unimportant. In addition, the target magnitude within the reduced Field-of-Excitation for a given sub-pulse need only exceed a given critical value. Modifications were made to the scan control software to allow the definition of a parallel transmit inversion prepulse to be specified by a text-file. The ability to additionally specify pTx excitation pulses (or RF shimming) was retained, but not investigated in this work.

6.2 Recent Literature

Some significant contributions related to ultrahigh field MRI research have emerged during the time frame of the work presented in this thesis.

Wang et al recently introduced a *time-resolved* Echo-Planar Imaging technique at 3 T (*239*). The method uses a continuous multi-shot readout and uses a highly undersampled trajectory to sample the joint *phase encoding-time space*. The reconstruction algorithm utilises both receive coil sensitivity and ΔB_0 information to approximate signal loss. It will be interesting to see how the method performs at 7 T with increased ΔB_0 variations.

A complementary technique to this that has grown in significance has been simultaneous multi-slice acquisitions (*209–212*). Tailored slice-selective RF pulses are combined with a phase modulation function to excite signal across multiple slices. Reconstruction methods use additional information from multiple receive coils to separate out the separate slices. With reference to EPI, whereas traditional parallel imaging techniques reduce the readout time (and hence ΔB_0 effects), simultaneous multi-slice instead reduces the total number of slice acquisitions. (*240*).

More generally is the concept of time optimal signal encoding in inhomogenous fields (both B_1^+ and ΔB_0). Recent advances in Magnetic Resonance Fingerprinting (*241*) and nonlinear tomography (*242*) have sought to apply this for multi-parametric imaging. Traditionally, parallel transmission has been used for field homogenisation (as well as transmit k-space undersampling). An alternative role for pTx has been suggested as an additional spatio-temporal encoding mechanism to help distinguish tissue specific signal responses (*243*).

6.3 Future Work

From a computational approach, there has been some recent advances in communication connections between multiple GPU devices that offer the possibility to quickly access global memory shared across devices (*244*). Large system encoding matrices that are required for high resolution 3D spatially-selective RF pulse design could be distributed across multiple GPU devices in this manner.

A further study into compact k-space parameterisations for 3D excitation would be of interest, to observe how well each perform across a library of excitation targets. There may also be other representations that were not considered in this work that are of interest, such as integer control points or wavelet decomposition.

For the dynamic, parcellated B_0 shimming for 3D EPI, validation of the Reduced SENSE reconstruction algorithm (*220*) is necessary to retain temporal efficiency if arbitrarily shaped rFOX were to be used. This would involve storing a binary mask for each separate sub-volume excitation (e.g. thresholding the predicted reduced Field-of-Excitation), and utilising this information in the reconstruction algorithm.

The efficacy of B_0 shimming increases with increasing sub-volume compactness. For small sub-volumes it becomes viable to utilise Echo-Volumar Imaging (EVI) due to the reduced readout duration, as compared to a full 3D FOV. This would require an EVI sequence to be implemented in the host spectrometers scan control software. Consideration should then be given to the trade-off in sub-volume dimensions, achievable B_0 uniformity, parallel imaging factors, and total imaging time.

The experimental work in this thesis has focused on parallel transmission in the human brain, largely due to the maturity of the hardware available for brain imaging at 7 T. However, there are significant opportunities for utilising pTx in the abdomen. RF shimming translates straightforwardly into the abdomen (see Appendix A). Some of the development work related to parcellated dynamic B_0 could be of use for targeted areas of the abdomen where susceptibility artefacts are caused by air in the gut. With regard to multidimensional adiabatic inversion, a perfusion experiment could be generated using a spatially localised inversion pulse. Patient motion would need to be considered in all of these abdominal applications.

Appendix A

Control Software

Commercial vendors have implemented (to varying degrees) Parallel Transmit in to their products operaing at and above 3 T field strengths. However, as the technology and applications for pTx are still emerging, there is still a requirement for customised software to simplify the workflow and additionally provide enhanced functionality.

Open-source toolboxes have played a crucial role in the development and deployment of advanced reconstruction algorithms, e.g. the Berkeley Advanced Reconstruction Toolbox (*245*), Michigan Image Reconstruction Toolbox (*246*), and the Gadgetron project (*247*). As many of the algorithms and tools described in these toolboxes are also applicable for transmission, it will be interesting to see if explicit support is added in the future for pTx. Two useful open-source repositories for pTx that are currently available, are:

- KCL group, maintained by Shaihan Malik (https://github.com/mriphysics/)
- Vanderbuilt group, maintained by Will Grissom (https://github.com/wgrissom)

The aim of this appendix is to describe a customised user-interface, computational setup, and specific functionality that was created in aid of the research presented in this thesis. MATLAB functions for file read/write from the Philips system were shared by Shaihan Malik.

The software has been used by other researchers within the Sir Peter Mansfield Imaging Centre, primarily for static RF and ΔB_0 shimming at this time. Two such examples are given at the end of this appendix for Magnetic Resonance Spectroscopy in the brain (see Figure A.8), and renal Magnetic Resonance Angiography using an abdominal body coil (see Figure A.9).

A.1 System Configuration

Alongside the existing vendor host (client) computer, a dedicated Linux-based server (Ubuntu 16.04, 32 Gb RAM) is configured to process compute intensive pTx operations,

equipped with an Nvidia Titan X Pascal using CUDA Toolkit 9.0 (*248*). Communication between the host and server is performed as a secure file copy using the SSH protocol.

A.2 Host (Client) Tools

Client-side tools read Philips PAR/REC files and store them as binary MATLAB-format files. The vendor system generates request text files with information such as pulse object identifier (e.g. sequence excitation pulse). For gradient waveforms, additional system information is passed in the request file such as maximum gradient amplitude and slew-rate for each gradient coil. This request format was extended optionally to specify prepulse inversion pulses and/or variant pTx pulses per dynamic.

The client-side user-interface was designed in MATLAB (with some additional Java Swing components) and subsequently compiled as an executable for Windows, requiring a standalone set of MATLAB Runtime shared libraries to be installed on the host computer.

Loading of input field maps is performed under *File -> Load* B_0 and *File -> Load Relative* B_1. Each load action generates a separate window panel where additional preprocessing steps can be performed. For neurological imaging, brain-extraction is performed and used to define a region mask (*219*). B_1^+ maps can be median filtered with a user-defined kernel size. A present limitation is the assumption that ΔB_0 and B_1^+ maps use the same geometry and same resolution, but this is usually acceptable in practice.

An intuitive graphical user-interface displays an anatomical magnitude image (the first gradient echo image from Dual TE images) and B_1^+ maps that can be toggled to show magnitude or phase. Scrolling the mouse wheel updates the slice view of all windows.

Single or multiple 3D regions-of-interest can be defined through either drawing or by using predefined drag-and-drop shapes. A separate window panel is enabled under *Design -> ROI* to do this (see Figure A.3). The desired ROI is automatically bounded by the background mask.

The *Design* tab allows easy selection of the following:

- RF Shimming.
  - Phase Nulling
  - Maximum Efficiency
  - MLS
- RF Pulse Design
  - 2D Spiral
  - 2D Spiral Constrained

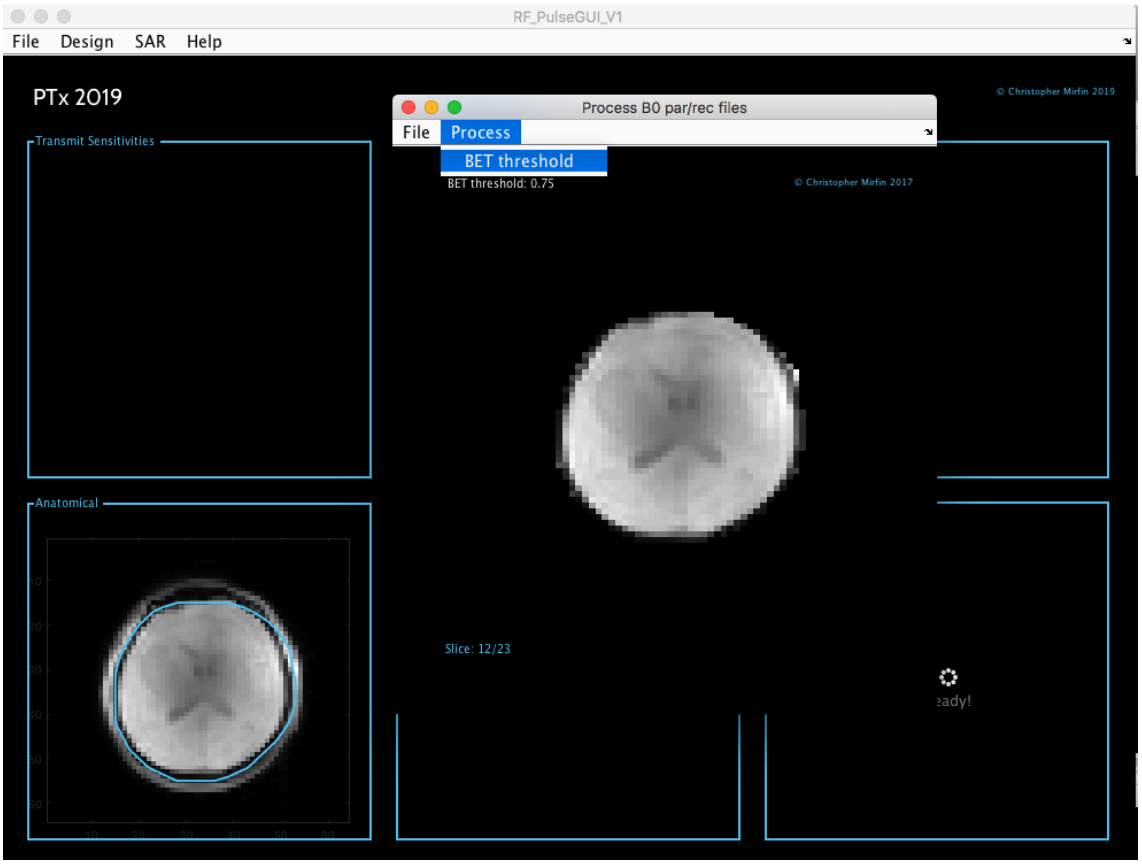

Figure A.1: Loading neurological data from PAR REC file. Brain extraction (*219*) (shown in the pop-up central window) is automatically performed on first gradient echo image from Dual-TE image set (bottom-left panel). For non-neurological imaging, brain extraction can be replaced with simple background masking.

  - 3D Shells
  - 3D Inversion (TFEPP)

- B_0 Shimming
  - Static B_0 Shimming
  - Dynamic B_0 Shimming

- Parcellated

Design -> RF Shimming presents several optimisation cost functions as described in Chapter 2. As these are computational non-intensive they are performed locally on the host computer. A data text file is generated as a series of real-valued RF shim values (amplitudes and phases) that is picked up and automatically loaded by the vendor host software.

ΔB_0 shimming is controlled through *Design -> B_0 shimming*. Two choices of either *static* or *dynamic* shimming are presented. The former calculates shim coil currents up to second-order for the specified ROI. Dynamic shimming automatically parcellates the

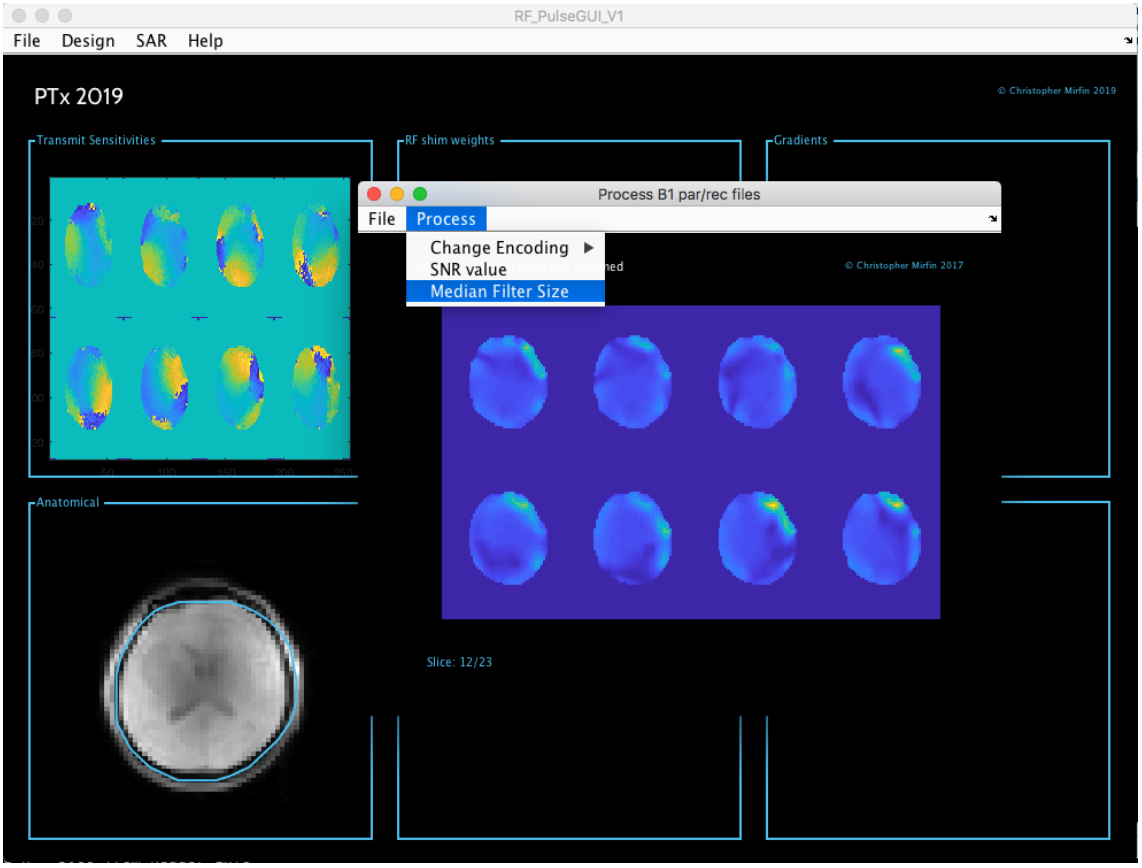

Figure A.2: Loading and preprocessing multichannel relative transmit sensitivity maps (central window panel). By default, single-channel encoding is assumed, but other channel encoding are permissible. Default median filtering is applied with kernel size [3, 3, 3]. The processed maps are displayed in the upper-left panel within the main interface (the view between magnitude and phase can be toggled by right-clicking).

FOV into four quadrants (or other shapes by editing a configuration file). A histogram of the ΔB_0 distribution for both unshimmed and predicted shimmed field is subsequently shown.

Selecting any of the *Design -> RF Pulse Design* options creates a secondary window panel to input specific design information, such as pulse object identifier, target flip angle and power constraints (see Figure A.6). *Spiral* pulses are implemented with either total power regularisation, or maximum channel power constraints. Selecting *Design -> RF Pulse Design -> 3D Inversion* creates a prepulse object, with a custom gradient reshaping parameter. All pulse computations are performed on the server-side.

RF pTx waveforms are displayed in the central window pane with a further window-panel displaying the expected (STA approximation) flip angle map (see Figure A.7). For the prepulse option, this is replaced by a Bloch simulated longitudinal magnetisation map scaled to between $[-1, 1]$ to show the transition between inverted and non-inverted regions.

Finally the option *Design -> Parcellated*, automatically parcellates the FOV, creates a

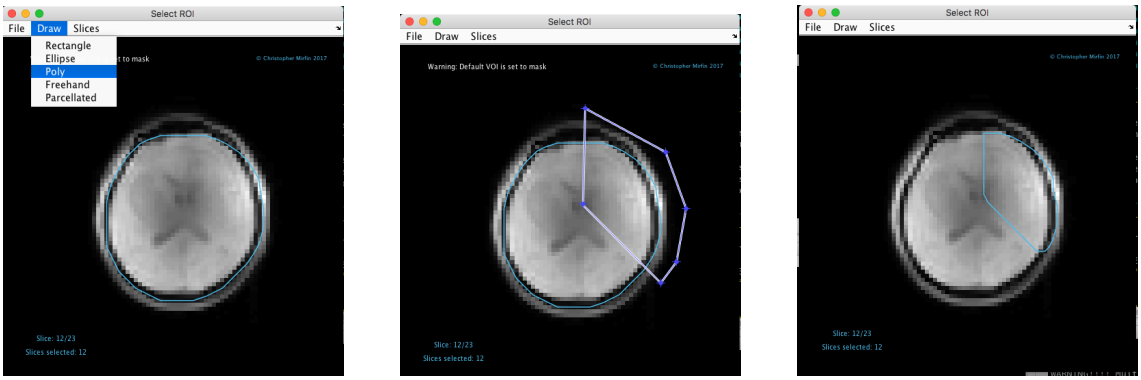

Figure A.3: ROI selection of varying size and shape are defined in the transverse plane for a given number of slices. Additional ROIs can be selected for multiple regions if required. First a *Draw* method is selected (left panel), then the user creates the mask (middle panel), and finally non-brain tissue is excluded using the predetermined BET mask (right panel).

set of dynamic shim currents (up to first-order), and associates each of these shim sets with a distinct pTx selective excitation pulse. By default the FOV is parcellated into four sub-volumes in the transverse plane.

A.3 Server Tools

Server applications are controlled by a series of Bash shell scripts. FSL command line tools are utilised for phase unwrapping if required in addition to brain extraction (*71*). Numerical optimisation routines are implemented in MATLAB and again are executed from a shell script. Optimised RF and gradient waveforms are returned to the host computer as text files that are automatically read-in by vendor software. A 'busy' indication is shown on the graphical user-interface while computations are performed on the server-side, so that the process is not interrupted.

For GPU computations, all field-maps are stored as a *gpuArray* data-type in MATLAB, which essentially acts as a pointer to GPU global memory. Suitable computations (e.g. element-wise) can then benefit from a large computational speed-up (usually at least one order-of-magnitude) as compared to sequential CPU calculations. It should be noted that for most low-priced or graphics dedicated GPUs have limited support for double precision operations.

Several overloaded functions can be directly used for GPU operations (e.g. matrix calculations) that use underlying accelerated libraries (*249*). For specific applications such as Bloch simulations, underlying CUDA code can be compiled into *cuda-mex* files that can be called directly from existing MATLAB programs.

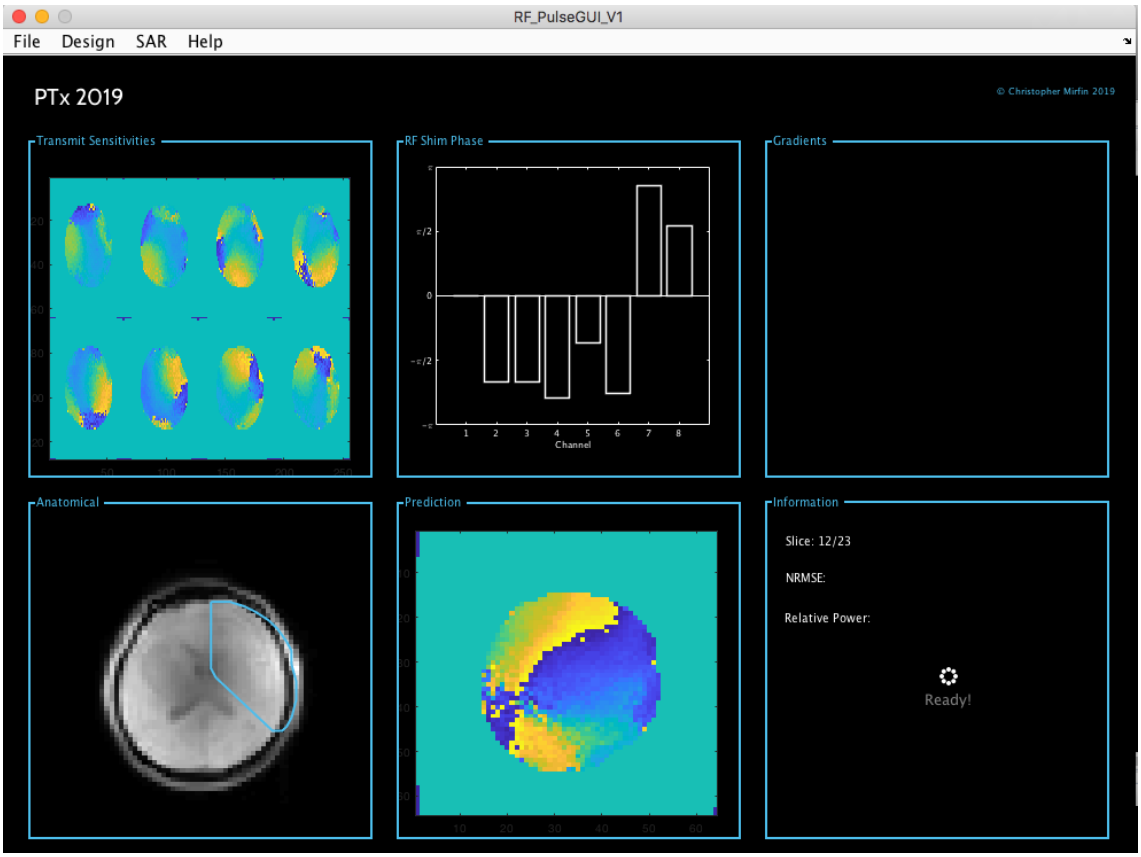

Figure A.4: RF shimming routines are selected from the *Design -> RF shimming* tab. The resulting channel weights (amplitude or relative phase) are displayed as a bar chart in the upper-middle window panel. The predicted transmit profile map (magnitude or phase) is shown in the lower-middle window panel. In this case a relative phase nulling algorithm was used.

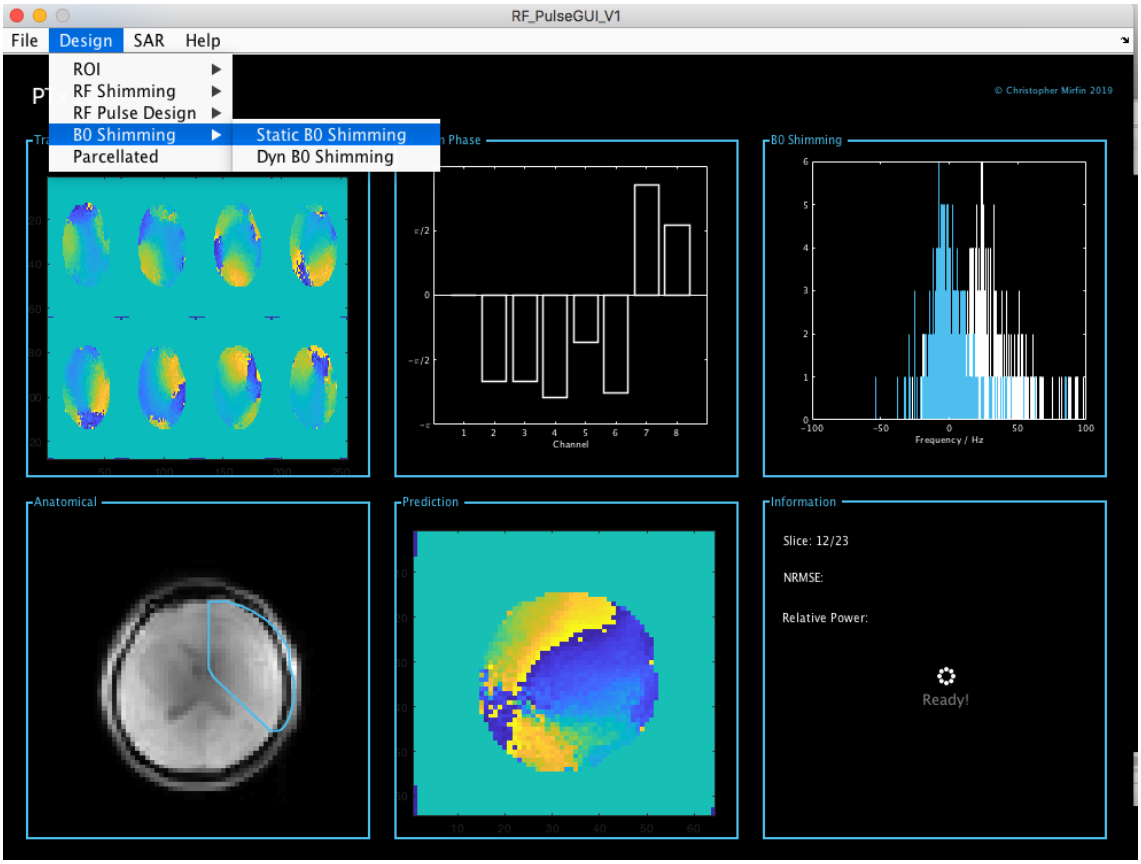

Figure A.5: Static ΔB_0 shimming (up to second-order) is selected from the *Design -> Static B0 shimming* tab. A histogram displaying the off-resonance frequencies within the selected ROI is shown in the upper-right window panel.

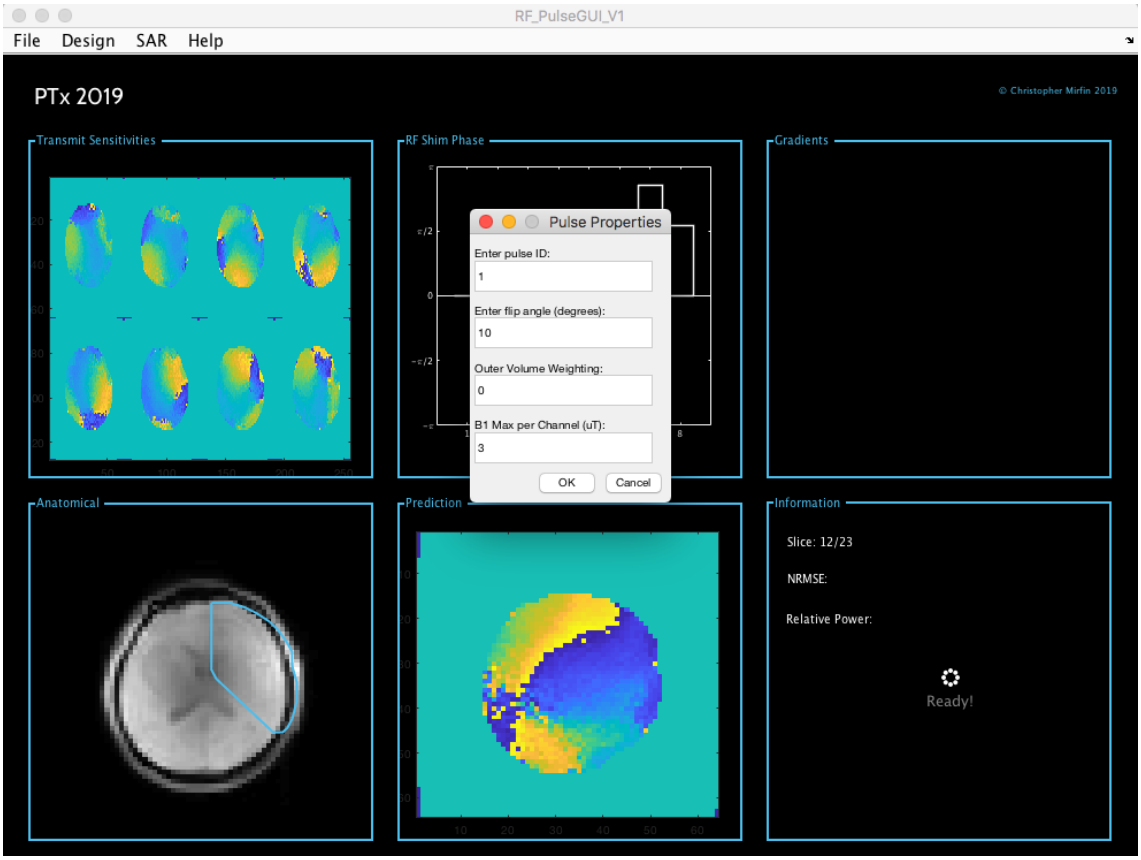

Figure A.6: RF pulse optimisation routines (e.g 2D spiral) are selected from the *Design -> RF Pulses* tab. A separate panel for input pulse properties is used to define further pulse specifications.

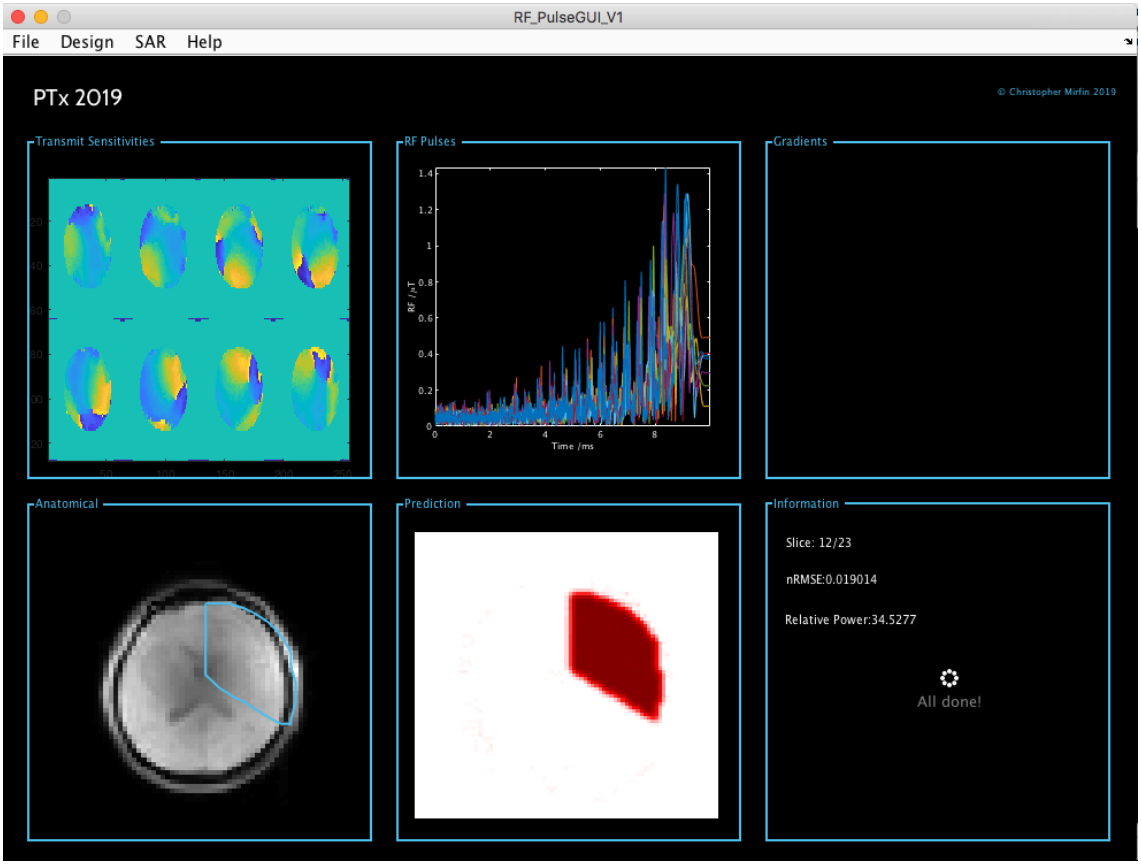

Figure A.7: After pulse optimisation completes on the remote server, the status switches from 'busy' to 'all done'. Resulting RF pTx waveforms are shown in the upper-middle window panel. A Bloch simulation of the transverse (or longitudinal) magnetisation is shown in the lower-middle window panel. Pulses are written to file to be loaded into the vendor spectrometer software.

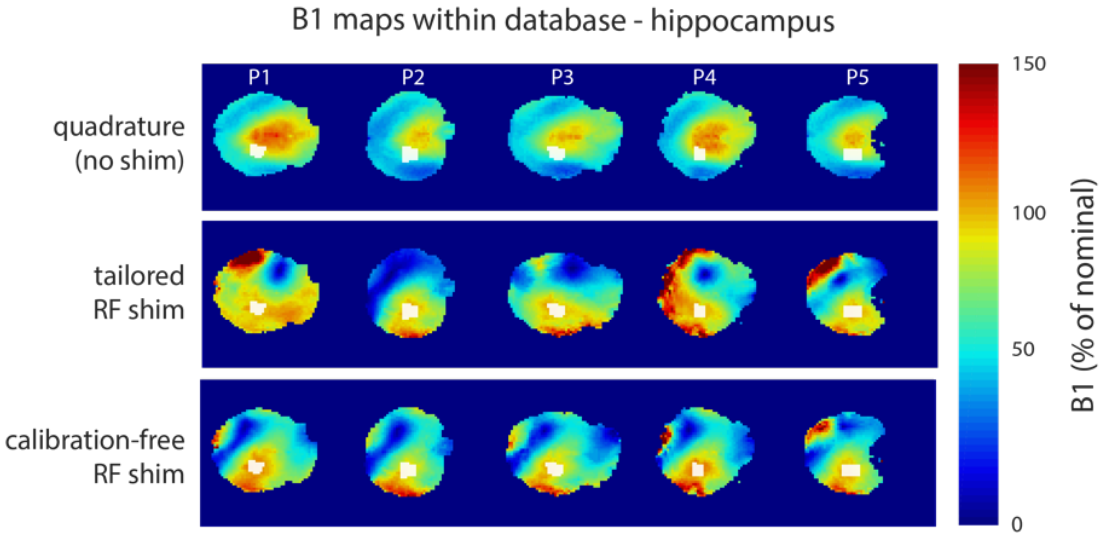

(a) Predicted $|B_1^+|$ maps (on a single axial slice) for five volunteers (P1 - P5) after RF shimming in the hippocampus, for quadrature, tailored RF shim, and calibration-free database RF shim for region.

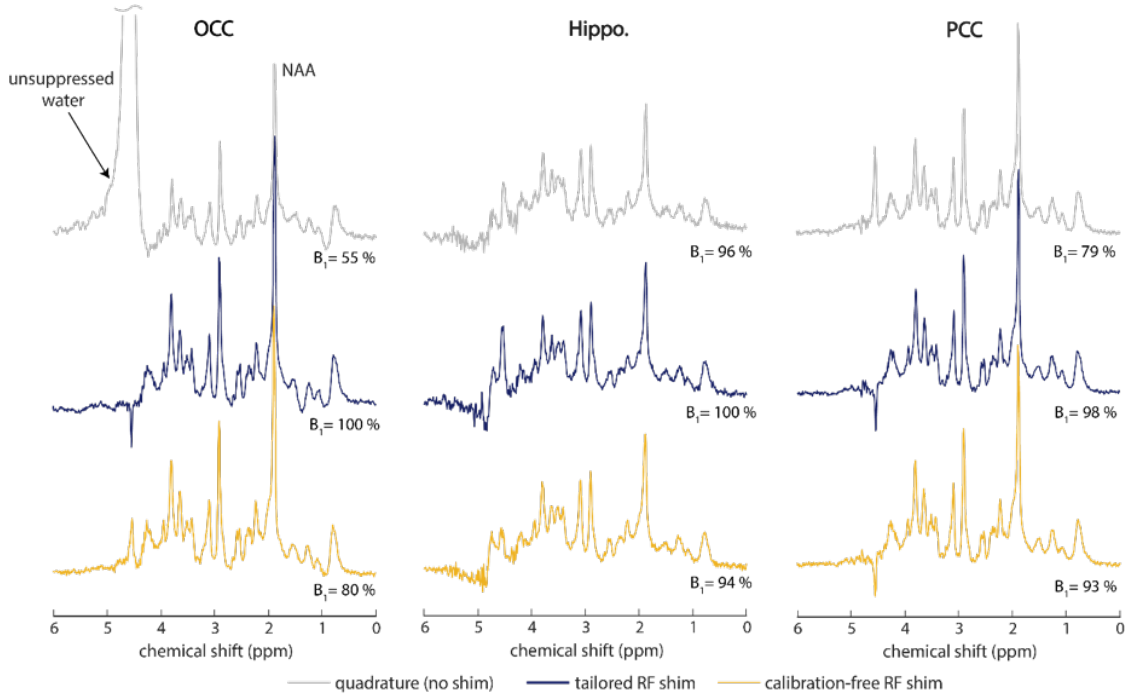

(b) Example single-voxel spectra from occipital cortex (OCC) (20 mm isotropic), hippocampus ($30 \times 15 \times 12$ mm^3), and posterior cingulate cortex (PCC) (20 mm isotropic).

Figure A.8: Application of Control Software for RF shimming in MR Spectroscopy. Adapted with permission from author Adam Berrington, in submission to the Proceedings of ISMRM 2020.

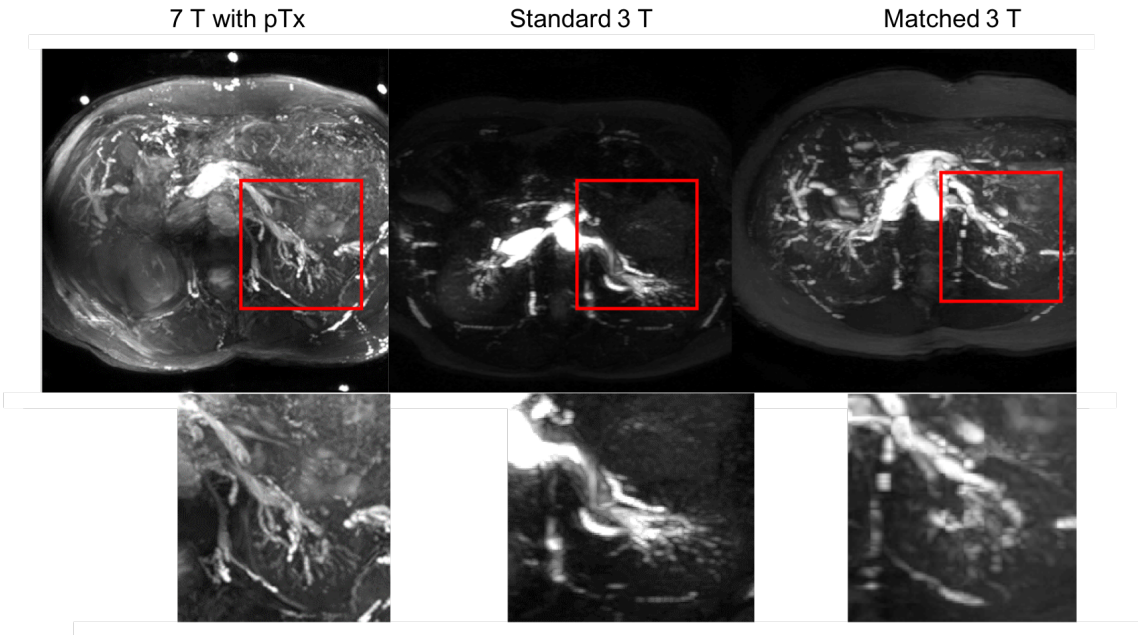

Figure A.9: Application of Control Software for RF shimming in Renal Angiography. 2D gradient echo maximum intensity-projection Time-of-flight renal angiography at 7 T with phase nulling RF shimming over right kidney to detect the arterial vasculature of the kidney and delineate the intrarenal vasculature network. Matched to parameters at 3 T. Incomplete background suppression at 7 T due to power limts. Adapted with permission from author Emma Doran (*250*).

Bibliography

1. I. I. Rabi, J. R. Zacharias, S. Millman, P. Kusch, *Physical Review* **53**, 318 (1938).
2. F. Bloch, W. W. Hansen, M. Packard, *Physical Review* **69**, 127 (1946).
3. E. M. Purcell, H. C. Torrey, R. V. Pound, *Physical Review* **69**, 37–38 (1946).
4. R. P. Feynman, F. L. Vernon, R. W. Hellwarth, *Journal of Applied Physics* **28**, 49–52 (1957).
5. R. R. Ernst, W. A. Anderson, *Review of Scientific Instruments* **37**, 93 (1966).
6. P. C. Lauterbur, *Nature* **242**, 190–191 (1973).
7. P. Mansfield, A. A. Maudsley, *British Journal of Radiology* **50**, 188–194 (1977).
8. D. K. Sodickson, W. J. Manning, *Magnetic Resonance in Medicine* **38**, 591–603 (1997).
9. K. P. Pruessmann, M. Weiger, M. B. Scheidegger, P. Boesiger, *Magnetic Resonance in Medicine* **42**, 952–962 (1999).
10. M. A. Griswold *et al.*, *Magnetic Resonance in Medicine* **47**, 1202–1210 (2002).
11. M. Lustig, D. Donoho, J. M. Pauly, *Magnetic Resonance in Medicine* **58**, 1182–1195 (2007).
12. K. Setsompop *et al.*, *NeuroImage* **80**, 220–233 (2013).
13. K. Uğurbil *et al.*, *NeuroImage* **80**, 80–104 (2013).
14. P. Mansfield, A. A. Maudsley, *Journal of Magnetic Resonance* **27**, 101–119 (1977).
15. R. J. Ordidge, R. M. Bowley, G. McHale, *Magnetic Resonance in Medicine* **8**, 323–331 (1988).
16. J. Frahm, K. D. Merboldt, W. Hänicke, *Journal of Magnetic Resonance* **72**, 502–508 (1987).
17. J. Pauly, D. Nishimura, A. Macovski, *Journal of Magnetic Resonance* **81**, 43–56 (1989).
18. M. E. Ladd *et al.*, *Progress in Nuclear Magnetic Resonance Spectroscopy* **109**, 1–50 (2018).
19. T. S. Ibrahim, L. Robert, A. M. Abduljalil, B. A. Baertlein, P. M. L. Robitaille, *Magnetic Resonance Imaging* **19**, 219–226 (2001).

20. D. I. Hoult, *Journal of Magnetic Resonance Imaging* **12**, 46–67 (2000).

21. U. Katscher, P. Börnert, C. Leussler, J. S. Van den Brink, *Magnetic Resonance in Medicine* **49**, 144–150 (2003).

22. F. Bloch, *Physical Review* **70**, 460 (1946).

23. I. I. Rabi, N. F. Ramsey, J. Schwinger, *Reviews of Modern Physics* (1954).

24. R. W. Brown, Y. C. N. Cheng, E. M. Haacke, M. R. Thompson, R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design: Second Edition* (John Wiley and Sons Ltd, 2014).

25. J. Pauly, D. Nishimura, A. Macovski, P. L. Roux, *IEEE Transactions on Medical Imaging* **10**, 53–65 (1991).

26. D. E. Rourke, M. J. Prior, P. G. Morris, J. A. Lohman, *Journal of Magnetic Resonance* **107**, 203–214 (1994).

27. D. L. Jaggard, Y. Kim, *Journal of the Optical Society of America A* **2**, 1922–1930 (1985).

28. N. Boulant, D. I. Hoult, *Magnetic Resonance in Medicine* **67**, 339–343 (2012).

29. J. S. Li, J. Ruths, S. J. Glaser, *Nature Communications* **8**, 446 (2017).

30. J. P. Marques, R. Bowtell, *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* **25B**, 65–78 (2005).

31. J. H. Duyn, *NeuroImage* **168**, 152–161 (2018).

32. J. P. Stockmann, L. L. Wald, *NeuroImage* **168**, 71–87 (2018).

33. H. Wen, F. A. Jaffer, *Magnetic Resonance in Medicine* **34**, 898–904 (1995).

34. C. Juchem, R. A. de Graaf, *Analytical Biochemistry* **529**, 17–29 (2017).

35. S. J. Vannesjo *et al.*, *Magnetic Resonance in Medicine* **78**, 1607–1622 (2017).

36. P. Mansfield, B. Chapman, *Journal of Magnetic Resonance* **66**, 573–576 (1986).

37. S. J. Vannesjo *et al.*, *Magnetic Resonance in Medicine* **69**, 583–593 (2013).

38. A. E. Campbell-Washburn, H. Xue, R. J. Lederman, A. Z. Faranesh, M. S. Hansen, *Magnetic Resonance in Medicine* **75**, 2278–2285 (2016).

39. S. J. Vannesjo *et al.*, *Magnetic Resonance in Medicine* **76**, 45–58 (2016).

40. G. H. Glover *et al.*, *Journal of Magnetic Resonance* **64**, 255–270 (1985).

41. G. Adriany *et al.*, *Magnetic Resonance in Medicine* **53**, 434–445 (2005).

42. G. Adriany *et al.*, *Magnetic Resonance in Medicine* **59**, 590–597 (2008).

43. A. J. Raaijmakers *et al.*, *Magnetic Resonance in Medicine* **75**, 1366–1374 (2016).

44. S. Orzada *et al.*, *PLoS ONE* **14**, e0222452 (2019).

45. M. J. van Osch, A. G. Webb, *Current Radiology Reports* **2**, 61 (2014).

46. D. C. Roberts *et al.*, *Current Biology* **21**, 1635–1640 (2011).

47. O. S. Mian, Y. Li, A. Antunes, P. M. Glover, B. L. Day, *PLoS ONE* **8**, e78748 (2013).

48. International Electrochemical Commission, *Medical electrical equipment: part 2-33. Particular requirements for the safety of magnetic resonance equipment for medical diagnosis*, 2015.

49. J. Rauschenberg *et al.*, *Investigative Radiology* **49**, 249–259 (2014).

50. P. Mansfield, B. Haywood, *Magnetic Resonance Materials in Physics, Biology and Medicine* **10**, 147–151 (2000).

51. W. A. Edelstein *et al.*, *Magnetic Resonance Imaging* **20**, 155–163 (2002).

52. M. Segbers, C. V. Rizzo Sierra, H. Duifhuis, J. M. Hoogduin, *Magnetic Resonance in Medicine* **64**, 546–553 (2010).

53. F. Bardati, A. Borrani, A. Gerardino, G. A. Lovisolo, *IEEE Transactions on Biomedical Engineering* **42**, 1201–1207 (1995).

54. G. Eichfelder, M. Gebhardt, *Magnetic Resonance in Medicine* **66**, 1468–1476 (2011).

55. J. Lee, M. Gebhardt, L. L. Wald, E. Adalsteinsson, *Magnetic Resonance in Medicine* **67**, 1566–1578 (2012).

56. M. De Greef, O. Ipek, A. J. Raaijmakers, J. Crezee, C. A. Van Den Berg, *Magnetic Resonance in Medicine* **69**, 1476–1485 (2013).

57. M. Le Garrec *et al.*, *Magnetic Resonance in Medicine* **78**, 1217–1223 (2017).

58. H. H. Pennes, *Journal of Applied Physiology* **1**, 93–122 (1948).

59. N. Boulant, A. Massire, A. Amadon, A. Vignaud, *Magnetic Resonance in Medicine* **72**, 679–688 (2014).

60. N. Boulant *et al.*, *Magnetic Resonance in Medicine* **75**, 249–256 (2016).

61. T. M. Fiedler, M. E. Ladd, A. K. Bitz, *NeuroImage* **168**, 33–58 (2018).

62. M. A. Bernstein, K. F. King, X. J. Zhou, *Handbook of MRI Pulse Sequences* (Elsevier, 2004).

63. W. A. Edelstein, J. M. Hutchison, G. Johnson, T. Redpath, *Physics in Medicine and Biology* **25**, 751 (1980).

64. Y. Zur, S. Stokar, *Journal of Magnetic Resonance* **71**, 212–228 (1987).

65. J. H. Duyn, *Magnetic Resonance in Medicine* **37**, 559–568 (1997).

66. M. Brant-Zawadzki, G. D. Gillan, W. R. Nitz, *Radiology* **182**, 769–775 (1992).

67. K. K. Kwong *et al.*, *Proceedings of the National Academy of Sciences of the United States of America* **89**, 5675–5679 (1992).

68. M. H. Buonocore, L. Gao, *Magnetic Resonance in Medicine* **38**, 89–100 (1997).

69. P. Jezzard, R. S. Balaban, *Magnetic Resonance in Medicine* **34**, 65–73 (1995).

70. L. L. Wald, *NeuroImage* **62**, 1221–1229 (2012).

71. S. M. Smith *et al.*, *NeuroImage* **23**, S208–S219 (2004).

72. P. Jezzard, *NeuroImage* **62**, 648–651 (2012).

73. R. M. Panchuelo, D. Schluppeck, J. Harmer, R. Bowtell, S. Francis, *Brain Topography* **28**, 62–65 (2014).

74. C. H. Cunningham, J. M. Pauly, K. S. Nayak, *Magnetic Resonance in Medicine* **55**, 1326–1333 (2006).

75. K. Nehrke, P. Börnert, *Magnetic Resonance in Medicine* **68**, 1517–1526 (2012).

76. K. Nehrke, A. M. Sprinkart, P. Börnert, *Magnetic Resonance Materials in Physics, Biology and Medicine* **28**, 185–194 (2015).

77. V. L. Yarnykh, *Magnetic Resonance in Medicine* **57**, 192–200 (2007).

78. S. J. Malik, D. J. Larkman, J. V. Hajnal, *Magnetic Resonance in Medicine* **62**, 902–909 (2009).

79. P.-F. Van de Moortele *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **15**, 1676 (2007).

80. R. Cusack, N. Papadakis, *NeuroImage* **16**, 754–764 (2002).

81. M. Hestenes, E. Stiefel, *Journal of Research of the National Bureau of Standards* **49**, 409–436 (1952).

82. K. Fatahalian, J. Sugerman, P. Hanrahan, *Proceedings of the SIGGRAPH/Eurographics Workshop on Graphics Hardware*, 133–137 (2004).

83. C. C. Paige, M. A. Saunders, *ACM Transactions on Mathematical Software* **8**, 43–71 (1982).

84. C. C. Paige, M. A. Saunders, *ACM Transactions on Mathematical Software* **8**, 195–209 (1982).

85. D. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods* (Academic Press Inc., 1982).

86. R. A. Waltz, J. L. Morales, J. Nocedal, D. Orban, *Mathematical Programming* **107**, 391–408 (2006).

87. S. R. Young, D. C. Rose, T. P. Karnowski, S. H. Lim, R. M. Patton, *Proceedings of Machine Learning in High-Performance Computing Environments* **4** (2015).

88. J. H. Holland, *Adaption in natural and artificial systems* (The University of Michigan, 1975).

89. P. F. Van De Moortele *et al.*, *Magnetic Resonance in Medicine* **54**, 1503–1518 (2005).

90. G. J. Metzger *et al.*, *Magnetic Resonance in Medicine* **59**, 396–409 (2008).

91. C. M. Deniz *et al.*, *Magnetic Resonance in Medicine* **69**, 1379–1388 (2013).

92. G. Carluccio, C. M. Collins, D. Erricolo, *Magnetic Resonance in Medicine* **71**, 432–439 (2014).

93. H. P. Hetherington, N. I. Avdievich, A. M. Kuznetsov, J. W. Pan, *Magnetic Resonance in Medicine* **63**, 9–19 (2010).

94. U. E. Emir *et al.*, *NMR in Biomedicine* **25**, 152–160 (2012).

95. G. J. Metzger *et al.*, *Magnetic Resonance in Medicine* **69**, 114–126 (2013).

96. A. de Boer *et al.*, *Magnetic Resonance Materials in Physics, Biology and Medicine* **29**, 417–433 (2016).

97. C. J. Snyder *et al.*, *Magnetic Resonance in Medicine* **61**, 517–524 (2009).

98. M. A. Erturk, X. Li, P.-F. Van de Moortele, K. Ugurbil, G. J. Metzger, *Topics in Magnetic Resonance Imaging* **28**, 101–124 (2019).

99. S. Saekho, F. E. Boada, D. C. Noll, V. A. Stenger, *Magnetic Resonance in Medicine* **53**, 479–484 (2005).

100. A. T. Curtis, K. M. Gilbert, L. M. Klassen, J. S. Gati, R. S. Menon, *Magnetic Resonance in Medicine* **68**, 1109–1116 (2012).

101. C. Mirfin, S. Bawden, P. Glover, P. A. Gowland, R. Bowtell, *Proceedings of the International Society for Magnetic Resonance in Medicine* **25**, 5204 (2017).

102. H. Homann *et al.*, *Magnetic Resonance in Medicine* **65**, 350–357 (2011).

103. S. J. Malik, F. Padormo, A. N. Price, J. V. Hajnal, *Magnetic Resonance in Medicine* **68**, 1481–1494 (2012).

104. S. J. Malik, R. P. A. Teixeira, J. V. Hajnal, *Magnetic Resonance in Medicine* **80**, 767–779 (2018).

105. A. Sbrizzi *et al.*, *Magnetic Resonance in Medicine* **77**, 361–373 (2017).

106. A. Beqiri, H. Hoogduin, A. Sbrizzi, J. V. Hajnal, S. J. Malik, *Magnetic Resonance in Medicine* **80**, 1533–1545 (2018).

107. S. J. Malik, A. Beqiri, F. Padormo, J. V. Hajnal, *Magnetic Resonance in Medicine* **73**, 951–963 (2015).

108. W. Grissom *et al.*, *Magnetic Resonance in Medicine* **56**, 620–629 (2006).

109. C. Y. Yip, J. A. Fessler, D. C. Noll, *Magnetic Resonance in Medicine* **54**, 908–917 (2005).

110. J. A. Fessler, B. P. Sutton, *IEEE Transactions on Signal Processing* **51**, 560–574 (2003).

111. L. Greengard, J. Y. Lee, *SIAM Review* **46**, 443–454 (2004).

112. B. P. Sutton, D. C. Noll, J. A. Fessler, *IEEE Transactions on Medical Imaging* **22**, 178–188 (2003).

113. B. Guerin, M. Gebhardt, S. Cauley, E. Adalsteinsson, L. L. Wald, *Magnetic Resonance in Medicine* **71**, 1446–1457 (2014).

114. S. Boyd, L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).

115. P. C. Hansen, in *Computational Inverse Problems in Electrocardiology* (WIT Press, 2000), vol. 4, pp. 119–142.

116. A. Sbrizzi *et al.*, *Magnetic Resonance in Medicine* **67**, 824–834 (2012).

117. J. Van den Eshof, G. L. Sleijpen, *Applied Numerical Mathematics* **49**, 17–37 (2004).

118. W. Mao, M. B. Smith, C. M. Collins, *Magnetic Resonance in Medicine* **56**, 918–922 (2006).

119. Z. Cao, X. Yan, W. A. Grissom, *Magnetic Resonance in Medicine* **76**, 1158–1169 (2015).

120. X. Yan, Z. Cao, W. A. Grissom, *Magnetic Resonance in Medicine* **75**, 2545–2552 (2016).

121. P. Kellman, D. A. Herzka, M. S. Hansen, *Magnetic Resonance in Medicine* **71**, 1428–1434 (2014).

122. M. Lustig, S. J. Kim, J. M. Pauly, *IEEE Transactions on Medical Imaging* **27**, 866–873 (2008).

123. S. Vaziri, M. Lustig, *Proceedings of the International Society for Magnetic Resonance in Medicine* **20**, 2284 (2012).

124. D. Lee, M. Lustig, W. A. Grissom, J. M. Pauly, *Magnetic Resonance in Medicine* **61**, 1471–1479 (2009).

125. S. Conolly, D. Nishimura, A. Macovski, G. Glover, *Journal of Magnetic Resonance* **78**, 440–458 (1988).

126. B. A. Hargreaves, C. H. Cunningham, D. G. Nishimura, S. M. Conolly, *Magnetic Resonance in Medicine* **52**, 590–597 (2004).

127. M. Davids, M. Ruttorf, F. G. Zöllner, L. R. Schad, *IEEE Transactions on Medical Imaging* **34**, 564–577 (2015).

128. M. A. Cloos *et al.*, *Magnetic Resonance in Medicine* **67**, 72–80 (2012).

129. H. Sun, J. A. Fessler, D. C. Noll, J.-F. Nielsen, *Proceedings of the International Society for Magnetic Resonance in Medicine* **22**, 1498 (2014).

130. S. G. Mallat, Z. Zhang, *IEEE Transactions on Signal Processing* **41**, 3397–3415 (1993).

131. D. Chen *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **17**, 171 (2009).

132. C. M. Deniz *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **19**, 4434 (2011).

133. P. M. Pardalos, D. Z. Du, R. L. Graham, *Handbook of Combinatorial Optimization* (Springer, 2013).

134. M. A. Griswold, S. Kanneggiesser, M. Muller, P. M. Jakob, *Proceedings of the International Society for Magnetic Resonance in Medicine* **13**, 2435 (2005).

135. W. Grissom, *Proceedings of the International Society for Magnetic Resonance in Medicine* **26**, 3396 (2018).

136. J. Pauly, D. Nishimura, A. Macovski, *Journal of Magnetic Resonance* **82**, 571–587 (1989).

137. D. Xu, K. F. King, Y. Zhu, G. C. McKinnon, Z. P. Liang, *Magnetic Resonance in Medicine* **58**, 326–334 (2007).

138. D. Xu, K. F. King, Y. Zhu, G. C. McKinnon, Z.-P. Liang, *Magnetic Resonance in Medicine* **59**, 547–60 (2008).

139. M. S. Vinding, I. I. Maximov, Z. Tošner, N. C. Nielsen, *Journal of Chemical Physics* **137**, 054203 (2012).

140. M. S. Vinding *et al.*, *Magnetic Resonance Materials in Physics, Biology and Medicine* **30**, 29–39 (2017).

141. W. A. Grissom, D. Xu, A. B. Kerr, J. A. Fessler, D. C. Noll, *IEEE Transactions on Medical Imaging* **28**, 1548–1559 (2009).

142. W. A. Grissom, C. Y. Yip, S. M. Wright, J. A. Fessler, D. C. Noll, *Magnetic Resonance in Medicine* **59**, 779–787 (2008).

143. J. J. Ackerman, T. H. Grove, G. G. Wong, D. G. Gadian, G. K. Radda, *Nature* **283**, 167–170 (1980).

144. J. L. Evelhoch, M. G. Crowley, J. J. Ackerman, *Journal of Magnetic Resonance* **56**, 110–124 (1984).

145. M. G. Crowley, J. L. Evelhoch, J. J. Ackerman, *Journal of Magnetic Resonance* **64**, 20–31 (1985).

146. M. Garwood, K. Uğurbil, *Journal of Magnetic Resonance* **291**, 84–93 (2018).

147. Y. Shrot, L. Frydman, *Journal of Magnetic Resonance* **172**, 179–190 (2005).

148. J. N. Dumez, L. Frydman, *Journal of Magnetic Resonance* **226**, 22–34 (2013).

149. A. Jang *et al.*, *Magnetic Resonance in Medicine* **76**, 1364–1374 (2016).

150. G. Liberman, L. Frydman, *Magnetic Resonance in Medicine* **77**, 1959–1965 (2017).

151. M. Garwood, L. DelaBarre, *Journal of Magnetic Resonance* **153**, 155–177 (2001).

152. K. Ugurbil, M. Garwood, M. Robin Bendall, *Journal of Magnetic Resonance* **72**, 177–185 (1987).

153. R. S. Staewen *et al.*, *Investigative Radiology* **25**, 559–567 (1990).

154. C. J. Hardy, W. A. Edelstein, D. Vatis, *Journal of Magnetic Resonance* **66**, 470–482 (1986).

155. M. S. Silver, R. I. Joseph, D. I. Hoult, *Journal of Magnetic Resonance* **59**, 347–351 (1984).

156. M. S. Silver, R. I. Joseph, D. I. Hoult, *Physical Review A* **31**, 2753–2755 (1985).

157. J. Zhang, M. Garwood, J. Y. Park, *Magnetic Resonance in Medicine* **77**, 1630–1638 (2017).

158. S. Conolly, D. Nishimura, A. Macovski, *Journal of Magnetic Resonance* **83**, 549–564 (1989).

159. K. Uğurbil, M. Garwood, A. R. Rath, *Journal of Magnetic Resonance* **80**, 448–469 (1988).

160. D. Rosenfeld, Y. Zur, *Magnetic Resonance in Medicine* **36**, 124–136 (1996).

161. D. Rosenfeld, S. L. Panfil, Y. Zur, *Journal of Magnetic Resonance* **126**, 221–228 (1997).

162. A. Tannús, M. Garwood, *NMR in Biomedicine* **10**, 423–434 (1997).

163. R. J. Ordidge, M. Wylezinska, J. W. Hugg, E. Butterworth, F. Franconi, *Magnetic Resonance in Medicine* **36**, 562–566 (1996).

164. A. C. Hurley *et al.*, *Magnetic Resonance in Medicine* **63**, 51–58 (2010).

165. D. Rosenfeld, Y. Zur, *Magnetic Resonance in Medicine* **36**, 401–409 (1996).

166. P. Balchandani, J. Pauly, D. Spielman, *Magnetic Resonance in Medicine* **64**, 843–851 (2010).

167. P. Balchandani, M. M. Khalighi, G. Glover, J. Pauly, D. Spielman, *Magnetic Resonance in Medicine* **67**, 1077–1085 (2012).

168. K. Setsompop, L. L. Wald, E. Adalsteinsson, *Proceedings of the International Society for Magnetic Resonance in Medicine* **15**, 1687 (2007).

169. P. Balchandani *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **19**, 2907 (2011).

170. E. Kopanoglu, A. J. D. Plumley, E. M. Arcan, C. M. Deniz, R. G. Wise, *Proceedings of the International Society for Magnetic Resonance in Medicine* **27**, 4686 (2019).

171. M. Piek, N. G. Lee, A. Marsman, V. O. Boer, E. T. Petersen, *Proceedings of the International Society for Magnetic Resonance in Medicine* **26**, 3400 (2018).

172. V. Gras, A. Vignaud, A. Amadon, D. Le Bihan, N. Boulant, *Magnetic Resonance in Medicine* **77**, 635–643 (2017).

173. V. Gras *et al.*, *Magnetic Resonance in Medicine* **80**, 53–65 (2018).

174. R. Tomi-Tricot *et al.*, *Magnetic Resonance in Medicine* **82**, 2016–2031 (2019).

175. J. D. Ianni, Z. Cao, W. A. Grissom, *Magnetic Resonance in Medicine* **80**, 1871–1881 (2018).

176. C. Mirfin, P. Glover, R. Bowtell, *Proceedings of the International Society for Magnetic Resonance in Medicine* **26**, 3388 (2018).

177. M. S. Vinding, B. Skyum, R. Sangill, T. E. Lund, *Magnetic Resonance in Medicine* **82**, 586–599 (2019).

178. J. Deng *et al.*, *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (2009).

179. C. Mirfin, P. Glover, R. Bowtell, *Proceedings of the International Society for Magnetic Resonance in Medicine* **27**, 4627 (2019).

180. D. A. Feinberg *et al.*, *Radiology* **156**, 743–747 (1985).

181. P. Mansfield, R. J. Ordidge, R. Coxon, *Journal of Physics E: Scientific Instruments* **21**, 275–280 (1988).

182. C. Y. Yip, J. A. Fessler, D. C. Noll, *Magnetic Resonance in Medicine* **56**, 1050–1059 (2006).

183. M. Poole, R. Bowtell, *Magnetic Resonance Materials in Physics, Biology and Medicine* **21**, 31–40 (2008).

184. J. T. Schneider *et al.*, *Magnetic Resonance in Medicine* **69**, 1367–1378 (2013).

185. S. Hao, J. A. Fessler, D. C. Noll, J. F. Nielsen, *IEEE Transactions on Medical Imaging* **35**, 468–479 (2016).

186. C. J. Hardy, P. A. Bottomley, M. O'Donnell, P. Roemer, *Journal of Magnetic Resonance* **77**, 233–250 (1988).

187. M. Davids, L. R. Schad, L. L. Wald, B. Guérin, *Magnetic Resonance in Medicine* **76**, 1170–1182 (2016).

188. W. Deng, C. Yang, V. A. Stenger, *Magnetic Resonance in Medicine* **65**, 363–369 (2011).

189. J.-F. Nielsen, H. Sun, D. C. Noll, *Proceedings of the International Society for Magnetic Resonance in Medicine* **24**, 1013 (2016).

190. C. Y. Yip, W. A. Grissom, J. A. Fessler, D. C. Noll, *Magnetic Resonance in Medicine* **58**, 598–604 (2007).

191. N. De Zanche, C. Barmet, J. A. Nordmeyer-Massner, K. P. Pruessmann, *Magnetic Resonance in Medicine* **60**, 176–186 (2008).

192. J. H. Duyn, Y. Yang, J. A. Frank, J. W. Van Der Veen, *Journal of Magnetic Resonance* **132**, 150–153 (1998).

193. N. G. Papadakis, A. A. Wilkinson, T. A. Carpenter, L. D. Hall, *Magnetic Resonance Imaging* **15**, 567–578 (1997).

194. R. H. Byrd, J. C. Gilbert, J. Nocedal, *Mathematical Programming* **89**, 149 (2000).

195. P. Pospichal, J. Jaros, J. Schwarz, *Applications of Evolutionary Computation. EvoApplications* **6024**, 442–451 (2010).

196. J. M. Pauly, B. S. Hu, S. J. Wang, D. G. Nishimura, A. Macovski, *Magnetic Resonance in Medicine* **29**, 2–6 (1993).

197. S. T. S. Wong, M. S. Roos, *Magnetic Resonance in Medicine* **32**, 778–784 (1994).

198. P. Mansfield, *Journal of Physics C: Solid State Physics* **10**, L55–L58 (1977).

199. S. J. Malik, S. Keihaninejad, A. Hammers, J. V. Hajnal, *Magnetic Resonance in Medicine* **67**, 1303–1315 (2012).

200. M. A. Cloos *et al.*, *Magnetic Resonance in Medicine* **67**, 72–80 (2012).

201. M. Çavuşoğlu, R. Mooiweer, K. P. Pruessmann, S. J. Malik, *NMR in Biomedicine* **30**, e3697 (2017).

202. S. J. Malik, J. V. Hajnal, *Magnetic Resonance in Medicine* **76**, 848–861 (2016).

203. R. A. De Graaf, P. B. Brown, S. McIntyre, D. L. Rothman, T. W. Nixon, *Magnetic Resonance in Medicine* **49**, 409–416 (2003).

204. S. Sengupta *et al.*, *Magnetic Resonance Imaging* **29**, 483–496 (2011).

205. C. Mirfin, S. Shah, P. Glover, R. Bowtell, *Proceedings of the International Society for Magnetic Resonance in Medicine* **26**, 3387 (2018).

206. P. Mansfield, A. M. Howseman, R. J. Ordidge, *Journal of Physics E: Scientific Instruments* **22**, 324 (1989).

207. A. M. Howseman *et al.*, *NeuroImage* **9**, 363–376 (1999).

208. B. A. Poser, P. J. Koopmans, T. Witzel, L. L. Wald, M. Barth, *NeuroImage* **51**, 261–266 (2010).

209. D. J. Larkman *et al.*, *Journal of Magnetic Resonance Imaging* **13**, 313–317 (2001).

210. D. A. Feinberg *et al.*, *PLoS ONE* **5**, e15710 (2010).

211. K. Setsompop *et al.*, *Magnetic Resonance in Medicine* **67**, 1210–1224 (2012).

212. D. A. Feinberg, K. Setsompop, *Journal of Magnetic Resonance* **229**, 90–100 (2013).

213. C. Juchem *et al.*, *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* **212**, 280–288 (2010).

214. V. O. Boer, D. W. Klomp, C. Juchem, P. R. Luijten, R. A. De Graaf, *Magnetic Resonance in Medicine* **68**, 662–670 (2012).

215. M. J. Berger, P. Colella, *Journal of Computational Physics* **82**, 64–84 (1989).

216. S. L. Horowitz, T. Pavlidis, *Journal of the ACM (JACM)* **23**, 368–388 (1976).

217. W. T. Clarke *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **26**, 317 (2018).

218. H. Wadell, *The Journal of Geology* **43**, 250–280 (1935).

219. S. M. Smith, *Human Brain Mapping* **17**, 143–155 (2002).

220. R. Mooiweer *et al.*, *Magnetic Resonance in Medicine* **78**, 88–96 (2017).

221. K. Setsompop *et al.*, *Magnetic Resonance in Medicine* **56**, 1163–1171 (2006).

222. P. Kassakian, *IEEE International Conference on Acoustics, Speech and Signal Processing* **3**, 53–56 (2005).

223. R. W. Gerchberg, W. O. Saxton, *Optik* **34**, 275–284 (1971).

224. A. Hoyos-Idrobo, P. Weiss, A. Massire, A. Amadon, N. Boulant, *IEEE Transactions on Medical Imaging* **33**, 739–748 (2014).

225. R. W. Harrison, *Journal of the Optical Society of America A* **10**, 1046–1055 (1993).

226. G.-z. Yang, B.-z. Dong, B.-y. Gu, J.-y. Zhuang, O. K. Ersoy, *Applied Optics* **33**, 209–218 (1994).

227. K. Ugurbil, M. Garwood, M. Robin Bendall, *Journal of Magnetic Resonance* **72**, 177–185 (1987).

228. A. Jang, X. Wu, E. J. Auerbach, M. Garwood, *Magnetic Resonance in Medicine* **79**, 701–710 (2018).

229. S. Conolly, J. Pauly, D. Nishimura, A. Macovsiu, *Magnetic Resonance in Medicine* **24**, 302–313 (1992).

230. C. G. Broyden, *IMA Journal of Applied Mathematics* **6**, 76–90 (1970).

231. R. Fletcher, *The Computer Journal* **13**, 317–322 (1970).

232. D. Goldfarb, *Mathematics of Computation* **24**, 23–26 (1970).

233. D. F. Shanno, *Mathematics of Computation* **24**, 647–656 (1970).

234. M. H. Seo, O. Linton, *Journal of Econometrics* **141**, 704–735 (2007).

235. T. P. Wade *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **22**, 4851 (2014).

236. S. A. Winkler *et al.*, *NeuroImage* **168**, 59–70 (2018).

237. M. A. Cloos *et al.*, *NeuroImage* **62**, 2140–2150 (2012).

238. S. J. Malik, S. Keihaninejad, A. Hammers, J. V. Hajnal, *Magnetic Resonance in Medicine* **67**, 1303–1315 (2012).

239. F. Wang *et al.*, *Magnetic Resonance in Medicine* **81**, 3599–3615 (2019).

240. M. Barth, F. Breuer, P. J. Koopmans, D. G. Norris, B. A. Poser, *Magnetic Resonance in Medicine* **75**, 63–81 (2016).

241. D. Ma *et al.*, *Nature* **495**, 187–192 (2013).

242. A. Sbrizzi *et al.*, *Magnetic Resonance Imaging* **46**, 56–63 (2018).

243. M. A. Cloos *et al.*, *Nature Communications* **7**, 12445 (2016).

244. A. Li *et al.*, *IEEE Transactions on Parallel and Distributed Systems* **31**, 94–110 (2020).

245. M. Uecker *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **23**, 2486 (2015).

246. C. Y. Lin, J. A. Fessler, *IEEE Transactions on Computational Imaging* **5**, 17–26 (2018).

247. M. S. Hansen, T. S. Sørensen, *Magnetic Resonance in Medicine* **69**, 1768–1776 (2013).

248. J. Nickolls, I. Buck, M. Garland, K. Skadron, *Queue-GPU Computing* **6**, 40–53 (2008).

249. W. Bosma, J. Cannon, C. Playoust, *Journal of Symbolic Computation* **24**, 235–265 (1997).

250. E. Doran *et al.*, *Proceedings of the International Society for Magnetic Resonance in Medicine* **26**, 4602 (2018).