

Parallel Transmission Technologies for Imaging the Brain and Body with 7 Tesla MRI



James L. Kent
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Michaelmas 2024

*This thesis is dedicated to
Ellie
for being my best friend, companion and a welcome distraction.
You were the most adorable Bernese mountain dog.*



19th April 2018 - 29th August 2024

Acknowledgements

I would like to take this opportunity to acknowledge some of the many people who have helped me along my journey in academia.

Firstly, I would like to thank my teachers and mentors, who are many, but have all impacted me and shaped me into the person I am today. One person who stands out in particular is Steve Huddleston, my secondary school maths teacher. Your excellent teaching and dedication made maths enjoyable and helped me to realise my potential. Growing up in a working-class family I probably would not have even considered going to university (let alone doing a DPhil) without your initial support and encouragement and for that I am immensely grateful.

I owe a huge amount to my core supervisors Aaron, Laci and Iulius. Your belief in me has helped my progress tremendously and I am so grateful for all your time, effort and support you have devoted to me over these years. Aaron, you are truly an encyclopedia of knowledge, experience and most impressively, ideas and despite this I think you are one of the most humble people I have ever met. Laci, I am amazed by your ability to make time for me and I am always grateful for your feedback, insights and attention to detail. I do not know how you both manage to do it all especially whilst juggling the trials of having young children. Iulius, thank you for your help, support, humour and enthusiasm as well as helping to organise my internship in Erlangen. You have been the best mentors I could have asked for.

The WIN physics group is a very kind, friendly, inclusive and diverse team. I am happy to have known such a collection of humble and intelligent people and I would like to thank everyone for their contributions and discussions which have helped shape and improve my work. This group is a beacon for what a lab should be. I would also like to thank the radiographers at FMRIB and OHBA, Jon, Mike, Juliet, David and Nicky, for their support and training to become a competent and confident scanner operator. Mathew, I really enjoyed our chats during the late-night scanning sessions.

I am so pleased to have met so many colleagues through the various conferences and ISMRMs who I hope will remain friends. The MRI community is full of intelligent and dedicated people doing great work to advance MRI, and it constantly amazes me how much the field continues to evolve. I am proud to have been part of such a great community.

I would also like to thank the Siemens Healthineers ultra-high field team in Erlangen where I completed my internship. In particular Patrick, Jürgen, Gabriela and Nico, you helped make my time in Erlangen both enjoyable and productive. I miss the weißwurstfrühstück, lebkuchen, sushi boat lunches and Christmas market outings.

My colleagues and friends in Oxford have made my time here very special and I hope to keep in touch with you all. In particular, I would like to give special mention to Thijs, Daniel, Xi and Cristiana, who adopted me into their year group when I was the only one in my year to start during the COVID-19 pandemic. Silei, Ziyu, Qijia, Zhiyu, Xinyu, Camille and Izzy as well as past and present visiting students and newcomers. You all helped in some way to keep me motivated whether that was simply lunchtime conversations, punting or pub trips. I am certain you will all do great things and I wish you all the best on your individual journeys.

I am thankful for my close friends Simi, Lucy, Emily, Beth, Peter, Ted, Amy, Dan, Helen and Aabha for always being supportive and interested in talking with me about my work. I greatly enjoy the time we spend together and it has been an enormous privilege to see you all flourishing in your own lives and careers.

My family for all their love and unconditional support throughout the years. I could not have done this without you.

Last, but by no means least, I would like to thank my partner Ben for his kindness and unwavering support. You have been my cornerstone. I am excited to see what you achieve on your own DPhil journey.

Thank you all.

Abstract

The push to higher-and-higher static magnetic field strengths for magnetic resonance imaging systems leads to the radiofrequency fields, that are used for image acquisition, becoming less-and-less uniform. This results in undesirable changes in image contrast which are not due to anatomy. It is, however, possible to mitigate for this radiofrequency field non-uniformity using technology known as parallel transmission. Although, in order to provide a robust subject-specific improvement the non-uniform radiofrequency field must first be mapped in a process known as “ B_1^+ mapping”.

To begin, I will go through some essential magnetic resonance physics and cover some of the key concepts used in this thesis. I will then introduce ultra-high field magnetic resonance imaging and discuss why it is of interest and some of the challenges it currently faces.

The first research chapter explores a new way of mapping this non-uniform radiofrequency field using a novel adaption to a previously used pulse sequence. This method, referred to as "Sandwich", was successful in producing absolute B_1^+ maps across the entire thorax at 7 T and in particular in the heart where typically many methods fail due to blood flow. Using extended phase graphs and Bloch simulations I investigated how this new method compares to those currently in use. I also investigated whether coil-cycling the acquisition of relative maps was beneficial and I found that coil-cycling is especially important when measuring 3D relative magnitude maps.

Mapping many individual transmit channels is time-consuming, hence a method to accelerate the acquisition of B_1^+ maps for the entire transmit array over a large volume was also explored. I did this by combining two methods known as "TxLR" and "B1TIAMO" with the Sandwich method mentioned previously. I showed that accurate B_1^+ maps can be reconstructed from low-resolution and sparse k -space data. I demonstrated this using pulse sequence capable of acquiring 3D single-channel maps for an 8-channel transmit array in only 14 seconds in the brain or 26 seconds in the body.

Finally, subject motion has the ability to degrade the performance of parallel transmission technology. Knowledge of the subject's motion can help inform both the acquisition and reconstruction for motion correction. Motion detection can be achieved through various methods, but radiofrequency sensor-based techniques like parallel transmit scattering or pilot tone are of particular interest, primarily due to their low hardware requirements. Hence, the final research chapter explores these two competing radiofrequency sensor-based methods for motion detection. Both parallel transmit scattering and pilot tone showed good sensitivity to head motion and each method was able to accurately predict rigid body head motion in all six degrees of freedom using a linear regression model.

Contents

List of Figures	xii
List of Tables	xvi
List of Terms and Abbreviations	xvii
1 Introduction	1
1.1 Motivation	1
1.2 Outline	2
2 Background	4
2.1 Introduction	5
2.2 A Brief History of MRI	5
2.3 MR Hardware	7
2.3.1 The Main Magnetic Field (B_0)	7
2.3.2 The Magnetic Field Gradients (G_x, G_y, G_z)	7
2.3.3 The Radio-frequency Field ($B_1+/-$)	8
2.3.3.1 Transmit Chain	9
2.3.3.2 Receive Chain	10
2.4 Basic MRI Physical Principles	10
2.4.1 Signal Generation	10
2.4.2 Excitation (B_1+)	12
2.4.3 Relaxation	13
2.4.4 Signal Reception (B_1-)	14
2.4.5 The Bloch Equations	14
2.4.6 Spatial Encoding	15
2.4.6.1 Slice-selective Encoding	15
2.4.6.2 k-space	15
2.4.6.3 In-plane Localisation	17
2.4.7 FLASH Imaging	18
2.4.8 Calibrating RF Pulse Voltage	19
2.4.9 Adiabatic Pulses	20

2.4.10	Spoiling Mechanisms	21
2.4.11	Signal and Noise in MRI	21
2.4.12	Extended Phase Graphs	24
2.5	Image Reconstruction	26
2.5.1	Partial Fourier	26
2.5.2	Zero-padding	26
2.5.3	Parallel Imaging	27
2.5.3.1	SENSE	27
2.5.3.2	GRAPPA	28
2.5.3.3	g-Factor	29
2.5.4	Compressed Sensing	29
2.5.5	Coil Sensitivity Estimation	30
2.5.5.1	ESPIRiT	31
2.5.6	Virtual Coils	32
2.5.7	Low Rank Methods	32
2.5.7.1	TxLR	33
2.6	Imaging the Moving Heart	34
2.7	Summary	36
3	Ultra-high Field MRI	37
3.1	Introduction	38
3.2	A Brief History of UHF MRI	38
3.3	Advantages and Challenges of UHF	38
3.3.1	B1+ Inhomogeneity	39
3.3.2	Specific Absorption Rate	41
3.4	Parallel Transmission	42
3.4.1	pTx Safety Tests	43
3.4.2	SAR Supervision	44
3.4.3	RF Shimming	45
3.4.4	Dynamic Parallel Transmission	45
3.5	B1+ Mapping	48
3.5.1	Relative B1+ Mapping	49
3.5.2	Absolute B1+ Mapping	50
3.5.2.1	Double Angle Method	51
3.5.2.2	Actual Flip Angle Imaging	53
3.5.2.3	Dual Refocusing Echo Acquisition Mode	54
3.5.2.4	SA2RAGE	55
3.5.2.5	MP3RAGE	56
3.5.2.6	Saturation TurboFLASH	57
3.5.2.7	Fingerprinting	59

3.5.2.8	Phase Sensitive	59
3.5.2.9	Bloch-Siegert Shift	60
3.5.2.10	BEAR	61
3.5.3	Off-resonance Effects	62
3.5.4	B1+ Mapping Acquisition and Encoding Strategies	62
3.5.4.1	B1TIAMO	66
3.6	Summary	67
4	Sandwiched SatTFL B1+ Mapping	68
4.1	Introduction	68
4.2	Theory	70
4.3	Methods	72
4.3.1	Numerical Simulations	72
4.3.1.1	Cardiac Numerical Simulations	73
4.3.1.2	Further Technical Development Simulations	74
4.3.2	Parallel Transmit MRI System	75
4.3.3	Sequence Implementation	76
4.3.3.1	2D Protocols	76
4.3.3.2	3D Protocols	77
4.3.4	Image Reconstruction	77
4.3.5	Phantom Experiments	77
4.3.6	<i>In Vivo</i> Experiments	78
4.4	Results	79
4.4.1	Simulation Results	79
4.4.1.1	Cardiac Simulation Results	84
4.4.1.2	Further Tech. Dev. Simulation Results	87
4.4.2	Phantom Results	88
4.4.3	<i>In Vivo</i> Results	94
4.5	Discussion	100
4.6	Summary	103
5	Comparison of B1+ Mapping Methods	104
5.1	Introduction	104
5.2	Simulation Framework for B1+ Mapping Sequence Comparison	105
5.2.1	Methods	105
5.2.2	Results	107
5.2.3	Discussion	112
5.3	FLASH-based B1+ Mapping Comparison	113
5.3.1	Methods	113
5.3.1.1	Parallel Transmit MRI System	113

5.3.1.2	Sequence Implementation	114
5.3.1.3	Image Reconstruction	114
5.3.2	Results	115
5.3.3	Discussion	119
5.4	Summary	119
6	Coil-cycled Relative Maps	120
6.1	Introduction	120
6.2	Methods	123
6.2.1	Numerical Heart Simulations	123
6.2.1.1	Simulated 2D Multi-slice Protocols	123
6.2.1.2	Simulated 3D Protocols	124
6.2.1.3	Synthetic Heart Data as Simulation Input	124
6.2.1.4	Phantom Data as Simulation Input	124
6.2.1.5	Simulation Experiments	125
6.2.2	Parallel Transmit MRI System	125
6.2.3	Sequence Implementation	125
6.2.4	Image Reconstruction	126
6.2.5	Phantom Experiments	126
6.2.6	<i>In Vivo</i> Experiments	127
6.3	Results	127
6.3.1	Numerical Simulation Results	127
6.3.2	Phantom Results	134
6.3.3	<i>In Vivo</i> Results	136
6.4	Discussion	136
6.5	Summary	138
7	Accelerated pTx B1+ Mapping	139
7.1	Introduction	139
7.2	Methods	141
7.2.1	Acquisition Scheme	141
7.2.2	Numerical Heart Simulations	141
7.2.3	Parallel Transmit MRI System	144
7.2.4	Retrospective Undersampling of Brain Data	144
7.2.5	Sequence Implementation	145
7.2.6	Image Reconstruction	148
7.2.7	Phantom Experiments	148
7.2.8	<i>In Vivo</i> Experiments	149
7.3	Results	149
7.3.1	Numerical Heart Simulation Results	149

7.3.2	Retrospective Undersampling of Brain Data Results	154
7.3.3	Phantom Results	154
7.3.4	<i>In Vivo</i> Results	159
7.4	Discussion	165
7.5	Summary	167
8	RF Sensors for Motion Detection	168
8.1	Introduction	168
8.2	Background of Motion Detection	169
8.2.1	Methods of Motion Detection	169
8.2.2	RF Sensor-based Methods of Motion Detection	170
8.2.2.1	pTx Scattering	171
8.2.2.2	Pilot Tone	174
8.2.3	Motivation and Aims	175
8.3	Methods	176
8.3.1	Sequence Implementation	176
8.3.2	Hardware and Sequence Adaptions for RF Sensors	177
8.3.3	RF Sensor Signal Processing	178
8.3.4	Image Reconstruction	181
8.3.5	Post-processing	181
8.4	Results	182
8.5	Discussion	186
8.6	Summary	188
9	Summary	189
9.1	Contributions	189
9.2	Future Directions	191
9.3	Closing Remarks	192
Appendix A		193
A.1	T1 Correction Scheme	193
A.2	VB17 Slave Dropout	196
A.3	VB17 Slave Synchronisation	197
A.4	Ineffective Gradient Spoiling	199
A.5	HS8 Versus HS4	201
A.6	Scattering Matrix Sensitivity	203
References		206

List of Figures

2.1	Energy level diagram of an atom with spin $\frac{1}{2}$ in an external magnetic field. . . .	11
2.2	Diagram illustrating the excitation of the net magnetisation vector.	13
2.3	Diagram illustrating slice-selective excitation.	15
2.4	Illustration of k -space spatial frequencies.	16
2.5	Simplified sequence diagram illustrating a single TR of a FLASH train.	19
2.6	Normalised pulse amplitude and phase for a hyperbolic secant RF pulse.	20
2.7	Demonstration of the SENSE problem for aliased images.	28
2.8	Diagram illustrating the training and application of a GRAPPA kernel.	29
2.9	Example single receive coil images and associated receive coil sensitivities.	31
2.10	Diagram illustrating the formation of the calibration matrix.	32
2.11	Transform of multidimensional k -space data to block-Hankel structured matrix.	33
2.12	Transmit and receive unfoldings of block-Hankel structured matrices.	34
2.13	Illustration of the ECG waveform and the associated imaging intervals.	35
3.1	Effect of B_1^+ inhomogeneity on a GRE image at various field strengths.	40
3.2	Illustrations of transmit chain hardware.	43
3.3	CP versus pTx MPRAGE.	46
3.4	Simplified DAM B_1^+ mapping sequence diagram.	52
3.5	Simplified AFI B_1^+ mapping sequence diagram.	53
3.6	Simplified DREAM B_1^+ mapping sequence diagram.	54
3.7	Simplified SA2RAGE B_1^+ mapping sequence diagram.	56
3.8	Simplified MP3RAGE B_1^+ mapping sequence diagram.	57
3.9	Simplified SatTFL B_1^+ mapping sequence diagram.	57
3.10	Simplified Phase Sensitive B_1^+ mapping sequence diagram.	60
3.11	Simplified BSS B_1^+ mapping GRE sequence diagram.	60
3.12	Simplified BEAR B_1^+ mapping sequence diagram.	61
3.13	Example methods of B_1^+ map encoding.	65
3.14	Illustration of sequence diagram for B1TIAMO.	66
4.1	Simplified pulse sequence diagrams of SatTFL, short TR SatTFL and Sandwich.	71
4.2	Bloch simulation of sinc and HS8 pulses used for saturation.	72
4.3	Diagram illustrating centric versus mirrored centric phase encoding.	75

4.4	Plot of longitudinal magnetisation for the 2D Sandwich sequence.	80
4.5	FWHM of the PSF for both the reference and prepared images.	81
4.6	Effects of nominal flip angle and T_1 on the measured flip angle.	82
4.7	Effects of nominal flip angle and T_1 on the SD of the measured flip angle. . . .	83
4.8	Bias and SD for different $\alpha : \beta$ ratios and number of phase encoding lines. . . .	84
4.9	Histogram of the range of TRs which occurred during simulations.	85
4.10	Simulation of longitudinal magnetisation with ejection fraction and variable TR. .	86
4.11	Effect of ejection fraction on the estimated FA.	86
4.12	Dynamic range efficiency of the Sandwich method versus TR.	87
4.13	Difference between the Sandwich and T_1 correction scheme in simulation. . . .	88
4.14	Flip angle maps from a silicon oil phantom.	89
4.15	Flip angle maps from a water phantom.	90
4.16	Comparison between acquired CP^+ mode and pTx map for a silicon oil phantom. .	91
4.17	Comparison between acquired CP^+ mode and pTx map for a water phantom. . .	92
4.18	3D B_1^+ maps in a the water phantom.	93
4.19	Simulation results for various T_1 values.	94
4.20	Correlation plots for 2D <i>in vivo</i> brain data.	95
4.21	<i>In vivo</i> brain B_1^+ maps.	96
4.22	Single-channel 2D <i>in vivo</i> B_1^+ maps.	97
4.23	Comparison between acquired CP^+ mode and a pTx map in the brain.	97
4.24	Axial, sagittal and coronal slices through 3D B_1^+ maps of the brain.	98
4.25	3D B_1^+ maps in the heart of a volunteer.	99
5.1	FWHM of the PSF for both reference and prepared images.	107
5.2	Simulated effect of T_1 on the measured B_1^+ for five sequences.	108
5.3	Simulated effect of B_0 on the measured B_1^+ for five sequences.	109
5.4	Simulated effect of coherent flow on the measured B_1^+ for five sequences. . . .	110
5.5	Simulated effect of isotropic diffusion on the measured B_1^+ for five sequences. .	111
5.6	Lookup tables used for B_1^+ map calculation.	115
5.7	"Ground Truth" maps from a combination across multiple transmission voltages. .	116
5.8	Correlation and Bland-Altman plots for each method across multiple voltages. .	117
5.9	Correlation and Bland-Altman plots for each method at a single voltage. . . .	118
6.1	Illustration of overlapping coil sensitivities.	121
6.2	Illustration of distributing the acquisition order of relative maps.	122
6.3	Pseudo code for acquiring relative B_1^+ maps channel-wise versus coil-cycled. . .	122
6.4	Locations of tracked voxels and positions of transmit elements.	125
6.5	Magnetisation of tracked voxels channel-wise and coil-cycled.	127
6.6	Relative maps acquired channel-wise and coil-cycled.	128
6.7	Impact of distributing the acquisition order of transmit channel excitations. . .	130

6.8	Flip angle versus Channel-wise and coil-cycled relative map error.	131
6.9	T1 versus Channel-wise and coil-cycled relative map error.	132
6.10	TR versus Channel-wise and coil-cycled relative map error.	133
6.11	Relative map channel-wise acquisitions versus coil-cycled in phantom.	135
6.12	Standard deviation in relative maps across nominal flip angles.	135
6.13	Relative map channel-wise acquisitions versus coil-cycled in the brain.	136
7.1	Sequence diagram for acquiring fast multi-channel B_1^+ maps in the brain/heart.	142
7.2	Fully sampled native resolution (139×178) synthetic images.	143
7.3	Undersampling masks used for accelerated map acquisition.	146
7.4	Comparison between native resolution (139×178) and k -space cropped maps.	150
7.5	Simulation results of accelerated multi-channel parallel transmit B_1^+ maps.	151
7.6	Effect of not acquiring a small calibration region.	152
7.7	Number of TxLR iterations versus the root-mean-square error.	153
7.8	Complex root-mean-square error for channels combined into CP^+	154
7.9	3D multi-channel maps in a spherical water phantom.	156
7.10	3D multi-channel maps in a large drum water phantom.	157
7.11	The shim efficiencies in a drum phantom.	158
7.12	3D multi-channel and combined CP^+ B_1^+ maps acquired in the brain.	159
7.13	The magnitude of 3D multi-channel and combined CP^+ maps for three subjects.	160
7.14	Comparison between an undersampled and a fully-sampled dataset.	161
7.15	Effect of three coil-compression algorithms on RMSE and reconstruction time.	161
7.16	3D multi-channel and combined CP^+ B_1^+ maps in the body.	162
7.17	3D multi-channel and combined CP^+ B_1^+ maps for three subjects in the body.	163
7.18	The shim efficiencies in the body.	164
8.1	Illustration showing the scattering of an 8 channel transmit system.	172
8.2	Frequency division multiplexing to measure the full scattering matrix.	173
8.3	Effect of step-wise motion on scattering matrix elements.	174
8.4	Example of pilot tone signal present in image data.	175
8.5	CAIPI sampling mask used for data acquisition.	177
8.6	Pilot tone spectral density versus power.	178
8.7	Example of pilot tone signal present in hybrid k -space data.	179
8.8	Singular value decomposition of pilot tone and pTx scattering data.	180
8.9	RF sensor processing pipeline from raw input data to output RF signal vectors.	180
8.10	Motion parameters extract from registration.	181
8.11	RF sensor predicted 6 DOF motion.	184
8.12	RF sensor predicted RMS motion.	185
A.1	Flip angle maps from a water phantom.	194
A.2	Flip angle maps in the brain.	195

A.3	VOP data showing the relative phase to the master channel.	197
A.4	Artefacts in pTx maps.	197
A.5	Phase offset between ungated and gated relative maps.	198
A.6	VOP data showing a delay between the master and slave preparation pulses. .	199
A.7	Worsening artefact visible in prepared images for increasing reference voltage. .	199
A.8	Artefact visible in prepared images for various parameters.	200
A.9	Reduction of artefact.	200
A.10	HS8 and HS4 pulse magnitude and phase.	201
A.11	HS8 and HS4 longitudinal magnetisation as a function of B_0 and B_1^+	202
A.12	HS8 and HS4 longitudinal magnetisation as a function of B_0	203
A.13	Effect of slice selection on scattering matrix elements.	204
A.14	Effect of RF pulse truncation on scattering matrix elements.	204
A.15	Effect of flip angle on scattering matrix elements.	205

List of Tables

2.1	Summary of the impact of bandwidth on the MRI protocol and images.	23
3.1	Summary of the advantages and disadvantages of UHF MRI.	39
3.2	Summary of the required B_1^+ information for a specific application.	49
4.1	Comparison of energy for SatTFL, Sandwich and DREAM.	102
5.1	Table of parameters used for sequence simulation.	106
5.2	Dynamic range of the simulated methods.	109
5.3	Average 10 s power, total energy and peak voltage for 5 simulated sequences. . .	111
5.4	Sequence protocol parameters used to acquire the B_1^+ maps.	115
6.1	Relative map NRMSE of various acquisitions.	134
7.1	Differences in <i>in vivo</i> and phantom sequence protocol parameters.	147
8.1	Root-mean-square error for PT, pTxS and a combination.	183

List of Terms and Abbreviations

1D, 2D, 3D	One-, two-, three dimensional
ACS	Auto-calibrated signals
ADC	Analogue-to-digital converter
AFI	Actual flip angle imaging
AO	Acquisition order
B₀	Static polarising field
B₁⁺	Radiofrequency transmit field
B₁⁻	Radiofrequency receive field
B1TIAMO	B ₁ ⁺ time interleaved acquisition of modes
BART	Berkeley advanced reconstruction toolbox
BET	Brain extraction tool
BPT	Beat pilot tone
BSS	Bloch-Siegert shift
BW	Bandwidth
CAIPI	Controlled aliasing in parallel imaging
CC	Coil-cycled
CP	Circularly polarised (A.k.a Quadrature, TrueForm [Siemens])
CS	Compressed sensing
CW	Channel-wise
DAM	Double angle method
DAQ	Digital acquisition
DICO	Directional couplers
DICOM	Digital imaging and communications in medicine (image format)
DOF	Degree(s) of freedom
DR	Dynamic range
DREAM	Dual refocusing echo acquisition mode

EF	Ejection fraction
EPG	Extended phase graphs
EPI	Echo-planar imaging
ESPIRiT	Eigenvalue-based SPIRiT
FA	Flip angle
FFT	Fast Fourier transform
FID	Free induction decay
FLASH	Fast low angle shot
FLIRT	FMRIB's linear image registration tool
FOV	Field of view
FSL	FMRIB software library
FWHM	Full width at half maximum
GLM	General linear model
GRE	Gradient recalled echo
HB	Heartbeat
HRV	Heart rate variability
HS8	Hyperbolic secant of order eight
ICE	Image calculation environment [Siemens]
IDEA	Integrated development environment for MR applications [Siemens]
IT	Image train (e.g. FLASH)
LUT	Lookup table
MPRAGE	Magnetisation prepared rapid gradient echo
MRI	Magnetic resonance imaging
NCO	Numerical crystal oscillator
NRMSE	Normalised root mean square error
OS	Oversampled
PALI	Power absorption limiter
PD	Proton density
PE	Phase encode
PF	Partial Fourier
PSF	Point-spread function

PT	Pilot tone
pTx	Parallel transmission
pTxS	Parallel transmission scattering
RF	Radiofrequency
RFPA	RF power amplifier
RMSE	Root-mean-square error
ROI	Region of interest
RSS	Root sum of squares
Rx	Receive
SA2RAGE	Saturation prepared with two rapid gradient echoes
SAR	Specific absorption rate
SatTFL	Presaturated turbo FLASH
SD	Standard deviation
SE	Spin echo
SNR	Signal-to-noise ratio
SPIRiT	Iterative self-consistent parallel imaging reconstruction
STE	Stimulated echo
STEAM	Stimulated echo acquisition mode
SVD	Singular value decomposition
T₁	Longitudinal magnetisation relaxation time
T₂	Transverse magnetisation relaxation time
TA	Acquisition time
TBP	Time-bandwidth product
TD	Time delay
TE	Echo time
TR	Repetition time
TWIX	Siemens raw data format (.dat)
Tx	Transmit
TxLR	Transmit low rank
UHF	Ultrahigh field (≥ 7 T)

1

Introduction

Contents

1.1	Motivation	1
1.2	Outline	2

This thesis consists of 9 chapters. This first chapter briefly explains the motivation and outline.

1.1 Motivation

The push to higher-and-higher static magnetic field strengths for magnetic resonance imaging systems leads to the radiofrequency fields, that are used for image acquisition, becoming less-and-less uniform. This results in undesirable changes in image contrast which are not due to anatomy. It is, however, possible to mitigate for this non-uniformity using technology known as parallel transmission.

Whilst universal pulses have been shown to work in the heart, these pulses are always outperformed by tailored RF pulses. However, in order to do this subject-specific calibration in the heart, faster, more robust and reliable B_1^+ maps are required which cover a larger field of view.

1.2 Outline

Chapter 2: Serves as an introduction to some MRI principles. This chapter covers an introduction to the hardware, physics and image reconstruction for acquiring MR images as well as some of the common techniques used in this work.

Chapter 3: Introduces the benefits and challenges with imaging at ultra-high field (≥ 7 T) as well as some pTx technologies.

Chapter 4: Details my work behind a new B_1^+ mapping sequence, the Sandwich. This work was presented as a physical poster at the ultra-high field workshop in 2022 in Lisbon, Portugal [1]. Later the same year, this work was also presented orally at the annual ISMRM meeting in London, UK [2] where it obtained a Magna Cum Laude Merit Award [Top 15% of abstracts within a major review category]. The work is now published in the journal Magnetic Resonance in Medicine [3]. During my internship with the ultra-high field team at Siemens Healthineers in late 2023 I worked towards implementing this method in IDEA and ICE onto the 7 T Terra.X system (XA60) in Erlangen, Germany, which is currently the only clinical 7 T pTx MRI system. The 3D Sandwich B_1^+ mapping sequence presented here is due to be released as a product sequence, including compressed sensing acceleration with online DICOM image reconstruction, to be used for a fully automated adjustment step in mid-2024 replacing the previously used 2D SatTFL. [*Siemens Healthineers cannot guarantee the future availability of this sequence*]

Chapter 5: Covers work which aims to gain a better understanding of the trade-offs between some of the commonly used B_1^+ mapping methods at 7 T. This is split into two sections 1) A simulation framework for the comparison of several B_1^+ mapping methods. This work was presented as a digital poster at the 2024 annual ISMRM meeting in Singapore [4]. 2) A FLASH-based comparison of B_1^+ mapping methods in phantom at 7 T. This was also presented at the 2024 annual ISMRM meeting in Singapore as a physical poster [5].

Chapter 6: Explores the benefits of coil-cycling relative maps. This work was presented orally at the 2024 UHF pTx workshop in Glasgow, Scotland.

Chapter 7: Features work to accelerate the acquisition of B_1^+ maps. Initial results from this work were accepted as a digital poster at the 2022 annual ISMRM meeting in London, UK [6]. An extension of this work which included cardiac B_1^+ mapping was presented as a digital poster at the 2023 annual ISMRM meeting in Toronto, Canada [7]. For this work, I was awarded 1st place prize in the trainee poster competition during the ultra-high field study group meeting. The work is now published in the journal Magnetic Resonance in Medicine [8].

Chapter 8: Covers the work undertaken to explore two different RF sensor-based methods of motion detection. This work was presented as a physical poster at the motion detection and correction workshop in 2022 in Oxford, UK as well as at the 2023 annual ISMRM meeting in Toronto, Canada [9].

Chapter 9: Summarises the work done in this thesis and gives an idea of what future directions could be explored.

Working alongside Sir Peter Mansfield on MRI was hugely exciting, I don't think any of us could foresee the transformative tool it was to become. MRI has this habit of reinventing itself every 5 or 10 years with yet something else it can do that no-one had thought about and that continues to be a certainty.

-Professor Peter Morris

2

Background

Contents

2.1	Introduction	5
2.2	A Brief History of MRI	5
2.3	MR Hardware	7
2.3.1	The Main Magnetic Field (B_0)	7
2.3.2	The Magnetic Field Gradients (G_x, G_y, G_z)	7
2.3.3	The Radio-frequency Field ($B_1+/-$)	8
2.3.3.1	Transmit Chain	9
2.3.3.2	Receive Chain	10
2.4	Basic MRI Physical Principles	10
2.4.1	Signal Generation	10
2.4.2	Excitation (B_1+)	12
2.4.3	Relaxation	13
2.4.4	Signal Reception (B_1-)	14
2.4.5	The Bloch Equations	14
2.4.6	Spatial Encoding	15
2.4.6.1	Slice-selective Encoding	15
2.4.6.2	k-space	15
2.4.6.3	In-plane Localisation	17
2.4.7	FLASH Imaging	18
2.4.8	Calibrating RF Pulse Voltage	19
2.4.9	Adiabatic Pulses	20
2.4.10	Spoiling Mechanisms	21
2.4.11	Signal and Noise in MRI	21
2.4.12	Extended Phase Graphs	24
2.5	Image Reconstruction	26
2.5.1	Partial Fourier	26
2.5.2	Zero-padding	26
2.5.3	Parallel Imaging	27
2.5.3.1	SENSE	27

	2.5.3.2	GRAPPA	28
	2.5.3.3	g-Factor	29
2.5.4		Compressed Sensing	29
2.5.5		Coil Sensitivity Estimation	30
	2.5.5.1	ESPIRiT	31
2.5.6		Virtual Coils	32
2.5.7		Low Rank Methods	32
	2.5.7.1	TxLR	33
2.6		Imaging the Moving Heart	34
2.7		Summary	36

2.1 Introduction

Magnetic resonance imaging (MRI) is used routinely to perform medical imaging of patients since it uses non-ionising radiation and gives excellent soft-tissue contrast. MRI is therefore usually preferred to other imaging techniques, such as computed tomography (CT), when imaging the brain. There are, of course, a few exceptions such as when imaging an intracranial haemorrhage where CT is the gold standard. Hundreds of millions of MRI scans are performed worldwide every year and they provide important information for both medical diagnoses and research.

In order to produce a measurable signal, the nuclei inside the sample being studied need to be in a high enough abundance. Hydrogen (^1H) is the most commonly imaged nuclei in MRI since it constitutes around 63% of the human body, by atomic percent, producing the largest signal compared to other nuclei. MRI is unique in that the contrast between different tissues can be altered to produce the most desirable images and the rate at which contrast changes allows for specific tissue properties to be measured.

2.2 A Brief History of MRI

The development of MRI was the culmination of decades of research. Isidor Rabi is credited with the discovery of the principle behind MRI known as nuclear magnetic resonance (NMR). Isidor Rabi demonstrated that a transition between nuclear spin states could be achieved using an oscillating magnetic field [10]. In 1946, Felix Bloch and Edward Purcell both independently applied Rabi's discovery to condensed matter [11, 12]. Initially, NMR was solely used in Chemistry as an analytical tool for producing chemical spectra, however it was later developed in the 1970s as a method for producing images of the human body.

The discovery of the spin echo (SE) by Erwin Hahn [13] became what is now the fundamental basis of many pulse sequences in MRI. Paul Lauterbur is attributed with the invention of MRI by demonstrating the encoding of spatial information in the magnetic resonance signal by using magnetic field gradients [14]. Lauterbur gave this method of imaging the name "zeugmatography" and was later known as NMR imaging. For fear of negative public association, the term "nuclear" was dropped leading to the technique being known today simply as MRI.

The first MRI scan of pathology was performed by John Mallard in 1974 who imaged a dead mouse [15]. This pioneering work lead John Mallard to obtain a grant from the Medical Research Council which enabled his team to design and build the world's first human whole-body MRI scanner known as the "Mark-1", based on a 0.04 T magnet which was delivered to Aberdeen in 1977.

The first human scan was performed by Sir Peter Mansfield and his team in 1976. This was a cross-section through Mansfield's PhD student Andrew Maudsley's finger and took around 23 minutes to acquire but helped to convince the Medical Research Council to fund the construction of a 0.1 T whole-body MRI scanner [16]. Ultimately, this led to the first commercial MRI device being available for the medical community in 1983, manufactured by Toshiba and Siemens. Later, functional MRI using diffusion tensor imaging (DTI) and blood oxygenation level-dependant (BOLD) techniques were developed allowing the measurement of brain connectivity and activity [17].

The first human MRI scanners in the 1980's were made from permanent or resistive electromagnets which limited field strength to far below 1 T. Modern scanners use superconducting magnets with the majority of clinical MRI scanners being 1.5 T or 3 T. Today, ultra-high field scanners of 7 T are cleared for clinical use.

Sir Peter Mansfield pioneered imaging techniques such as slice selection [18] and echo-planar imaging (EPI) [19]. These methods led to vastly improved image acquisition times than was previously achievable and was imperative to MRI's success in medical imaging. Hence, in 2003 Paul Lauterbur and Sir Peter Mansfield jointly won the Nobel prize in Medicine for their pivotal work. Although, Raymond Damadian, an American physician, disputed that he deserved the prize as he had hypothesised a method to discriminate between cancerous and healthy tissue by their difference in relaxation times [20]. Mansfield wrote in his autobiography that "the person who really missed out on the Noble prize" was Erwin Hahn for his work on spin echoes [21].

2.3 MR Hardware

There are three main components of an MRI. Firstly, the superconducting electromagnet which produces the strong main static magnetic field and polarises the nuclear magnetic moments (e.g. Hydrogen or other X-nuclei). Secondly, there are the gradient coils which produce linearly varying magnetic fields in the x , y , and z directions and allow spatial localisation of the signal. And thirdly, the radiofrequency (RF) coils to induce precession of these magnetic moments and receive the resulting signal. The following will briefly introduce these components.

2.3.1 The Main Magnetic Field (B_0)

A conventional MRI scanner has a static magnetic field which is generated by a superconducting electromagnet. For a 3 Tesla MRI this magnetic field is on the order of 6×10^4 times stronger than the Earth's magnetic field. This magnetic field conventionally points in z -direction down the bore in the longitudinal axis. Whilst the static B_0 field itself typically has good homogeneity, the introduction of subjects into the static field causes local B_0 field changes. The B_0 field can be shimmed to improve these inhomogeneities using a B_0 map measured from a phase-sensitive scan. The optimal shim currents for the shim coils which are integrated into the scanner are then calculated. B_0 effects are particularly noticeable at air-tissue interfaces such as the frontal sinus, where geometric distortion and signal dropout typically occur. Additionally, if the local field changes are dynamically changing with physiological processes such as respiration then this can be even more troublesome.

2.3.2 The Magnetic Field Gradients (G_x , G_y , G_z)

Three orthogonal gradient field components also exist which can be modulated in time and are typically used to spatially encode the MRI signal. The maximum gradient amplitude is usually around 40 mT m^{-1} with slew rates up to $200 \text{ mT m}^{-1} \text{ s}^{-1}$ per gradient axis. Although, recent advances in gradient technology has produced a Connectome scanner capable of gradient amplitudes up to 300 mT m^{-1} and slew rates up to $200 \text{ mT m}^{-1} \text{ s}^{-1}$ per gradient axis [22]. The next generation Connectome 2.0 system is capable of 500 mT m^{-1} gradient amplitudes and slew rates up to $600 \text{ mT m}^{-1} \text{ s}^{-1}$ along each gradient axis [23]. However, peripheral nerve stimulation from the rapidly switching electric fields in the body becomes a substantial safety consideration.

2.3.3 The Radio-frequency Field ($B_1^{+/-}$)

The main component responsible for RF transmission (B_1^+) and reception (B_1^-) is an RF coil. Various kinds of coils exist such as surface coils which have a more inhomogeneous excitation field but with a high surface sensitivity or volume/body coils which are larger coils that can produce a more homogeneous and larger excitation field. Whatever the coil geometry, it must be tuned and matched to the resonant Larmor frequency in order for optimal resonant performance.

The wavelength, λ , associated with the frequency, f , of electromagnetic radiation is given by the electromagnetic properties of the subject, namely the permittivity, ϵ , and permeability, μ , as shown in Eq. 2.1 where v_p is the phase velocity of light which is usually less than the speed of light in a vacuum ($c = 2.99 \times 10^8 \text{ m s}^{-1}$). ϵ_0 is the permittivity of free space ($\epsilon_0 = 8.85 \times 10^{-12} \text{ F m}^{-1}$) and μ_0 is the permeability of free space ($\mu_0 = 1.26 \times 10^{-6} \text{ N A}^{-2}$). In a complete vacuum the relative permittivity, ϵ_r , and relative permeability, μ_r , are equal to one. Hence, in a complete vacuum, the wavelength associated with the ^1H Larmor frequency at 3 T ($f_L = 128 \text{ MHz}$) is 2.34 m. However, the relative permittivity of human tissue (assuming $\epsilon_r \approx 50$) reduces this down to 33 cm.

$$v_p = \frac{1}{\sqrt{\epsilon_0 \epsilon_r \mu_0 \mu_r}} = \lambda f \quad (2.1)$$

Although conventional MRI scanners have only a single transmit (Tx) channel, they generally have several transmit elements. The RF fields generated by each transmit element will superimpose and either constructively or destructively interfere. This will lead to a spatially varying B_1^+ field distribution. B_1^+ inhomogeneity results in altered image contrast and is especially problematic for making quantitative measurements such as when T_1 mapping. B_1^+ inhomogeneity is insignificant at field strengths below 1.5 T but the problem is especially complicated at field strengths $\geq 7 \text{ T}$. At 7 T the wavelength associated with the ^1H Larmor frequency ($\approx 13 \text{ cm}$) is comparable to the size of the subject and can result in regions of complete signal drop-out and localised heating [24–26].

The B_1 field in MRI is only active orthogonal to the B_0 field. This is made up of two orthogonal components, typically B_{1x} and B_{1y} which make up the transmit (B_1^+) and receive (B_1^-) fields. These fields are defined as

$$B_1^+ = \frac{B_{1x} + iB_{1y}}{2} \quad (2.2)$$

$$B_1^- = \frac{(B_{1x} - iB_{1y})^*}{2} \quad (2.3)$$

Depending on the orientation of the B_0 field, these equations may be swapped.

2.3.3.1 Transmit Chain

A transmit "channel" describes an completely independent electronic chain. The important components in the RF transmission chain are:

1. Frequency Synthesizer - The frequency synthesizer produces the continuous sinusoidal carrier wave at the Larmor frequency (298 MHz at 7 T). This device contains a numerical crystal oscillator (NCO) that keeps precise track of the frequency and phase. Additionally, the NCO is used in the receive (Rx) chain for demodulation of the MR signal.
2. Modulator - The output carrier wave is passed down the RF transmit chain to the pulse wave modulator which shapes the RF pulse with a low-frequency envelope (i.e. a sinc or rect function).
3. RF Power Amplifier - The RF pulse is then amplified by specifying a reference voltage. Within the RF power amplifier there is a directional coupler (DICO). The function of the DICO is to split off a sample of the (complex) forward and reflected signal for specific absorption rate (SAR) monitoring (using the concepts of virtual observation points) and fault detection. The standard MR receivers acquire the DICO signal during transmission every 10 μ s. During reception of the MR signal, the MR receivers are switched back to the RF coil elements by a switch-matrix.
4. Quadrature Coupler - For a single transmit channel system, after RF amplification, the output is typically split into two by a quadrature coupler device. The resultant outputs are 90° out of phase with each other and are inputs to the two ports of the quadrature transmit coil (e.g. a birdcage coil). The two outputs of the coupler are commonly known as I and Q, standing for 'in phase' and 'quadrature', respectively.
5. T/R Switch - For transmit coils that are also used as receive coils, a T/R switch prevents the high power transmit chain from destroying the sensitive receive circuitry by decoupling the receive circuitry during transmission.

2.3.3.2 Receive Chain

The important components in the RF receive chain are:

1. T/R Switch - For receive coils that are also used as transmit coils, a T/R switch prevents the high power transmit chain from destroying the sensitive receive circuitry by decoupling the receive circuitry during transmission.
2. Preamplification - The MRI signal detected is exceptionally small and must be immediately amplified to prevent noise contamination.
3. ADC - The amplified analog signal is then digitised using a high frequency analog-to-digital converter. On most modern MRI systems this takes place on the coil itself.
4. Demodulation - To remove the RF carrier wave the signal is demodulated using the pure carrier wave from the RF frequency synthesizer and low-pass filters. The complex output signal is then sent on to the image reconstruction system.

2.4 Basic MRI Physical Principles

2.4.1 Signal Generation

The magnetic spin quantum number, m_I , for a particle of spin $\frac{1}{2}$ (such as ^1H) in a magnetic field can take one of two possible states, as shown in Figure 2.1. The particle can find itself either aligned parallel or anti-parallel to the external magnetic field¹, B_0 , pointing in the z -direction. For the proton, these result in the two energy levels that can be occupied due to the change of their magnetic moments, known as the Zeeman effect.

When the external field is zero, there is no energy difference between the two states. For a non-zero external magnetic field, a transition between these states can be caused by the absorption or emission of a photon with an energy, E_{photon} , equal to the difference in energy between two states, ΔE . The energy corresponding to a photon with angular frequency, ω , is given by the Planck-Einstein relation

$$E_{\text{photon}} = \Delta E = \hbar\omega, \quad (2.4)$$

¹Actually, in the accurate quantum mechanical representation particles exist in a superposition of states. But for our purposes, classical mechanics works just fine.




Figure 2.1: Energy level diagram of an atom with spin $\frac{1}{2}$ in an external magnetic field. There are two possible values of spin quantum number (m_I) with an energy difference, ΔE , between them that varies linearly with magnetic field.

where $\hbar$ is the reduced Planck constant ($h/2\pi$). The frequency at which the angular momentum of the particle precesses around the external field is known as the Larmor frequency, which has the following relation

$$\omega_L = \gamma B_0, \quad (2.5)$$

where γ is the gyromagnetic ratio, which is particle specific and the ratio of the magnetic moment to its angular momentum. The gyromagnetic ratio of the hydrogen proton is $\gamma = 42.57 \text{ MHz T}^{-1}$. Transitions between these states can be induced by applying electromagnetic radiation using a RF coil with oscillating currents. The energy that must be absorbed/emitted to cause a change of state is therefore given by

$$\Delta E = \hbar \gamma B_0. \quad (2.6)$$

The net magnetisation at equilibrium, $\mathbf{M}_0$, for an ensemble of spin $\frac{1}{2}$ particles is proportional to the difference between the number of spins in either spin up, $N_\uparrow$, ($m_I = +\frac{1}{2}$) or spin down, $N_\downarrow$, ($m_I = -\frac{1}{2}$) and the external magnetic field strength. The ratio of the number of spins in either state is described by Boltzmann statistics

$$\frac{N_\uparrow}{N_\downarrow} = \exp\left(\frac{-\Delta E}{k_B T}\right) = \exp\left(\frac{-\gamma \hbar B_0}{k_B T}\right), \quad (2.7)$$

where T is the temperature of the ensemble and k_B is the Boltzmann constant. In the absence of a magnetic field, there are equal populations of spins in either state and therefore the net magnetisation is zero. A strong magnetic field will cause a population imbalance and hence cause a net magnetisation which points in the same direction as B_0 , due to the preferred lower energy parallel state. However, the

lower energy state is only *very* slightly more favoured. For every 1 million protons in the higher state, there are one million and five in the lower energy state ($T = 298\text{ K}$, $B_0 = 1.5\text{ T}$). Despite this, the MR signal is sensitive to this difference in populations but a large magnetic field is required in order to maximise the signal from a sample. Due to this relationship between MR signal and proton density (PD), PD is one of the main mechanisms of contrast in MRI.

2.4.2 Excitation (B_1^+)

A pulse applied to the transverse plane of strength $B_{1,max}$ for a pulse duration, t_p , can cause the net magnetisation to tip, and change the angle of the magnetisation with respect to the external field, reducing the component of magnetisation in the longitudinal plane

$$M_z = \mathbf{M}_0 \cdot \hat{\mathbf{z}} = M_0 \cos(\alpha) \quad (2.8)$$

where, α , is the flip angle (FA) and is determined by the properties of the applied RF pulse and $\hat{\mathbf{z}}$ is the unit vector in the z direction. For a rectangular (non-selective) pulse this relationship is

$$\alpha = \gamma B_{1,max} t_p. \quad (2.9)$$

However, this equation is only valid for small FAs close to the resonant frequency. In general,

$$\alpha = \int_0^{t_p} \gamma B_{1,max}(t) dt. \quad (2.10)$$

A perturbation of the net magnetisation away from the z -axis will produce a component of the magnetisation in the transverse xy -plane, $M_{xy} = M_0 \sin(\alpha)$. In this complex notation, M_x is the real component and M_y is the imaginary component of M_{xy} , where $M_{xy} = M_x \hat{\mathbf{x}} + i M_y \hat{\mathbf{y}}$ where $i = \sqrt{-1}$ and $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are unit vectors in the x and y directions, respectively. If you imagine the net magnetisation in the static frame, the magnetisation vector appears to be spiralling away from B_0 at the Larmor frequency. To simplify things, we can imagine we are in the rotating frame of reference. In this frame of reference, we are rotating around with the magnetisation vector with the angular frequency, ω , such that the B_1^+ vector appears static and simply tips the magnetisation down into the transverse plane, as in Figure 2.2.




Figure 2.2: Diagram illustrating the excitation of the net magnetisation vector, M , by application of a B_1^+ field in the transverse (x - y) plane whilst in presence of a static magnetic field, B_0 , along the longitudinal z -axis in **a)** the laboratory frame of reference and **b)** the rotating frame of reference.

2.4.3 Relaxation

Once the B_1^+ excitation is no longer present, the net magnetisation vector will revert back to equilibrium and once again align with the static external B_0 field. This relaxation is characterised by the decay constants T_1 and T_2 and are important sources of contrast in MR images.

The decay constant T_1 describes the spin-lattice relaxation time. It is characteristic of the rate at which the longitudinal component of net magnetisation recovers at a slow exponential rate back to thermodynamic equilibrium. This decay is caused by Brownian motion since random fluctuations in the motion of particles causes a change in the local magnetic field. Consequently, this influences spin transitions which is the reason this decay constant is partly temperature dependent.

T_2 is the spin-spin relaxation time constant, caused by the ensemble of spins dephasing and reducing the transverse magnetisation. This dephasing is caused by time varying processes, which generates fluctuating differences in local magnetic field that cannot be rephased. T_1 processes also contribute to some T_2 decay hence why $T_1 > T_2$. This time constant can be combined with the time constant T_2' , which is caused by static magnetic field inhomogeneity to produce an effective transverse decay constant T_2^* , given by

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}. \quad (2.11)$$

T_2' depends on the homogeneity of the external magnetic field because a spin in a different location in space will feel a slightly different magnetic field and

hence precesses at a slightly different frequency, which can be rephased. T_2^* is therefore responsible for dephasing the initially coherent transverse magnetisation after excitation. The resulting signal is called the free induction decay (FID) which is a sine wave damped by T_2^* of the form $\sin(\omega_L t) e^{-\frac{t}{T_2^*}}$.

2.4.4 Signal Reception (B_1^-)

When the precessing magnetisation is within the sensitive region of the receiver coil, that is tuned to the Larmor frequency, the changing magnetic flux induces an electric current according to Faraday's law of induction. The voltage change is recorded as a complex MR signal.

2.4.5 The Bloch Equations

The Bloch equations describe the rate of change of the components of the macroscopic magnetisation, $\mathbf{M}$, as a result of external magnetic fields, $\mathbf{B}$, which could be comprised of B_0 and B_1^+ , as well as relaxation processes (T_1 and T_2). If we initially consider a situation where there is no relaxation, the Bloch equations are

$$\frac{d\mathbf{M}(t)}{dt} = \gamma(\mathbf{M}(t) \times \mathbf{B}(t)) \quad (2.12)$$

where $\mathbf{M}$ is comprised of the magnetisation components in the x , y and z directions.

The full Bloch equations split into each magnetisation component including relaxation are

$$\frac{dM_x(t)}{dt} = \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_x - \frac{M_x(t)}{T_2}, \quad (2.13)$$

$$\frac{dM_y(t)}{dt} = \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_y - \frac{M_y(t)}{T_2}, \quad (2.14)$$

$$\frac{dM_z(t)}{dt} = \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_z - \frac{(M_z(t) - M_0)}{T_1}. \quad (2.15)$$

Assuming $\mathbf{B} = 0$ (i.e. in the rotating frame of reference we are on resonance and there is no excitation RF field) these equations have general solutions. The decay of the transverse magnetisation, M_{xy} , at a time, t , is given by

$$M_{xy}(t) = M_{xy}(t=0) e^{-\frac{t}{T_2}}. \quad (2.16)$$

Similarly, the recovery of the longitudinal magnetisation, M_z , is given by

$$M_z(t) = M_0 - (M_0 - M_z(t=0)) e^{-\frac{t}{T_1}}. \quad (2.17)$$

Therefore, in order for M_0 to recover to 99.3% of its original value requires a time of approximately $5T_1$.

2.4.6 Spatial Encoding

In the previous sections, I have described how an MR signal is generated and measured but I have not discussed how the signal provides spatial information. This section will touch on two aspects of image formation; slice-selective encoding and gradient encoding which can provide in-slice localisation (frequency and phase encoding).

2.4.6.1 Slice-selective Encoding

As mentioned earlier, MRI scanners are equipped with three gradient coils which are capable of producing linearly changing magnetic fields along the x , y and z axes. Since magnetisation is only excited if the RF pulse is on-resonance with the Larmor frequency, this can be used to only excite a band of frequencies corresponding to either a thin-slice for 2D imaging or a thick slab for 3D imaging. This is done by applying a gradient field and a RF pulse simultaneously. A transverse, sagittal or coronal slice can be selected by using a gradient field along either the z , x or y direction, respectively. An oblique slice can be selected in any direction by applying a linear combination of these gradient fields.




Figure 2.3: Diagram illustrating slice-selective excitation using a linear gradient in the z direction to excite a thin slice in the brain.

2.4.6.2 k -space

Discretely digitised MR signals are stored in k -space which represents the spatial frequencies in an image. The centre of k -space encodes the low spatial frequencies like general shapes and contrast, whereas the periphery of k -space encodes the high spatial frequencies like edges, as illustrated in Figure 2.4. As k -space represents the raw acquired MRI data a 2D inverse Fourier transform must be performed to reconstruct the final image.




Figure 2.4: An image of a brain is shown (left) along with the corresponding k -space data (right). If only the centre of k -space is retained then only the broad structure of the brain is seen. Whereas if only the periphery of k -space is retained then this only preserves the edges of the brain image. FFT; fast Fourier transform.

k -space is essentially continuous and infinite but it is not necessary to measure k -space completely. Therefore a portion of k -space must be sampled discretely. To do this, the majority of sequences in MRI sample k -space on a Cartesian grid. Cartesian sampling gives images well-defined properties such as resolution and field of view (FOV). The FOV for a Cartesian sampled k -space is given by the inverse

of the resolution to which k -space is measured, Δk .

$$\text{FOV} = \frac{1}{\Delta k} \quad (2.18)$$

Since higher k -space values represent higher spatial frequencies, the resolution in image space, Δx , is given by the maximum k -space value that is sampled, k_{max} . This is given by

$$\Delta x = \frac{1}{k_{max}} \quad (2.19)$$

2.4.6.3 In-plane Localisation

Using a slice-selective excitation restricts the transverse magnetisation to within a slice of interest. However, the signal that is received is the integral of all transverse magnetisation across the entire slice. In order to localise the signal in-plane two additional encodings must be performed.

Gradient Encoding Again using the gradient coils, a linearly changing magnetic field, G , is produced which causes every nuclei along the specific gradient axis to experience a slightly different magnetic field strength, ΔB . This leads to precessional frequencies of the nuclei that depend on the spatial location, r .

$$\omega(r) = \gamma(B_0 + \Delta B) = \gamma(B_0 + G \cdot r) \quad (2.20)$$

As a result, there is a phase shift of the transverse magnetisation. The amount of phase shift depends on both the gradient amplitude and duration, which is called the "gradient moment". The phase shift at any time at any point is given by

$$\phi(r, t) = \gamma \int_0^t G(\tau) \cdot r d\tau \quad (2.21)$$

Once the gradient is switched off, the accumulated phase shift remains until either the transversal magnetisation decays (T_2) or another gradient is applied. If a gradient with identical moment but opposite polarity is then applied, the magnetisation can be rephased at a specific time known as the echo time (TE). This is the basis of a core pulse sequence known as gradient-recalled echo (GRE).

Frequency and Phase Encoding To fill k -space, different combinations of gradients must be turned on to move around k -space and sample different spatial frequencies. The first in-plane axis is sampled using frequency encoding. Before data acquisition (readout) a large negative G_x gradient is applied to move us in k -space to $-k_{max}$. A positive G_x gradient is then applied and k -space samples are acquired as G_x traverses from $-k_{max}$ to $+k_{max}$. Modern scanners are capable of extremely fast digitisation, much more than the typical resolution required. Therefore, the frequency encoding direction (i.e. readout) will, usually by default, oversample by a factor of two. This results in the base resolution in the frequency encoding direction being twice that requested on the scanner. This is useful for preventing aliasing in this dimension and comes with no penalty.

The second in-plane axis is sampled using phase encoding (PE). Before data acquisition a G_y gradient is applied for a specific time to move us in k -space to a certain k_y position. The G_y gradient is then turned off to fix the k_y location during the frequency encoding readout. Phase encoding typically occurs over multiple shots and thus is generally referred to as the “slow” encoding, resulting in this axis being more susceptible to artefacts.

If slice-selective excitations are not performed and instead slab-selective excitations are used, an additional phase encoding direction is necessary in order to localise the signal. This is called 3D imaging since the RF pulse will excite a large slab instead of thin slices. 3D imaging is slower than 2D slice-selective imaging but due to the additional phase-encoding dimension there is an inherent gain in the signal-to-noise ratio (SNR). Additionally, 3D imaging does not suffer from some of the problems with 2D imaging such as slice-profile effects (since in reality a perfect slice cannot be excited). Non-selective pulses can also be used which enable a shorter TE than slab-selective excitations.

2.4.7 FLASH Imaging

Fast low angle shot (FLASH) is based on a spoiled GRE sequence that use small FA excitations [27–29]. A basic FLASH sequence diagram is shown in Figure 2.5. FLASH imaging preserves more of the longitudinal magnetisation than conventional SE sequences and therefore allows for much reduced imaging times (both TE and TR) therefore also reducing sensitivity to motion artefacts. The term “turbo-factor” describes the number of phase encoding lines within a single FLASH imaging train. As a result of the small FAs and short TR the FLASH sequence would have a

low SNR and poor T_1 contrast. Both the SNR and contrast can be substantially improved by performing a magnetisation preparation pulse prior to imaging. This is typically a saturation (e.g. 90°) or inversion (e.g. 180°) pulse. Large turbo-factors lead to a large T_1 filtering effect in the phase encoding direction and depends on how much the longitudinal magnetisation is changing throughout the acquisition [30], although this has been improved in the past using a k -space filter [31].




Figure 2.5: Simplified sequence diagram illustrating a single TR of a FLASH train.

2.4.8 Calibrating RF Pulse Voltage

The relationship between the voltage of the RF pulse and the FA that is executed requires calibration of the system through what is known as B_1^+ mapping. Improvements in B_1^+ mapping techniques is an active area of research even for current clinical field strengths, for example B_1^+ mapping in the presence of metallic hardware [32].

By definition on Siemens scanners, the reference voltage set on the scanner, V_{ref} , is the voltage required for a 1 ms rectangular pulse to achieve a 180° FA. Hence, to calculate the voltage required for an arbitrary (including phase modulated) RF pulse of a specified FA, α , and duration, τ , of a number of samples, N_{samples} , is given by

$$V_{\text{pulse}} = \frac{\alpha V_{\text{ref}} N_{\text{samples}}}{\gamma B_{1,\text{ref}} \tau \text{AmpInt}} \quad (2.22)$$

where the amplitude of the reference pulse, $B_{1,\text{ref}}$, is 500 Hz/11.7 μ T and AmpInt is the amplitude integral of the normalised pulse. For a non-phase-modulated pulse this is given by

$$\text{AmpInt} = \sqrt{\Re^2 + \Im^2} \quad (2.23)$$

where $\Re$ is the sum of the real parts of the RF pulse waveform and $\Im$ is the sum of the imaginary parts.

2.4.9 Adiabatic Pulses

Adiabatic pulses are a special class of RF pulse that are insensitive to the transmit field (over the adiabatic threshold) and hence produce a uniform FA in spite of B_1^+ inhomogeneity. An adiabatic pulse can be designed to be either an excitation (half-passage), a refocusing, or an inversion pulse (full-passage). This is achieved by both amplitude and frequency modulations throughout the duration of the pulse. Their design is usually complicated and formed of various hyperbolic trigonometric functions [33–37]. The effective field is much slower than the bulk magnetisation precession, which means the bulk magnetisation vector is able to follow the effective field despite variations in the bulk magnetisation amplitude. Hence, adiabatic pulses are able to generate uniform FA distributions even when B_1^+ is not uniform. An example hyperbolic secant of order eight (HS8) pulse is shown in Figure 2.6.




Figure 2.6: Normalised pulse amplitude (left) and phase (right) for a 5 ms hyperbolic secant (order 8) RF pulse with time-bandwidth product of 15.

Unlike conventional non-selective (e.g. rectangular) pulses, the FA is no longer proportional to the B_1^+ magnitude and pulse duration. Instead, it depends on how the amplitude and phase of the B_1^+ varies during the pulse. Hence, adiabatic pulses cannot be simply scaled e.g. performing a 180° adiabatic pulse at one-third amplitude will not equate to a 60° pulse. The main disadvantage of this kind of RF pulse is they typically have a large SAR burden. The large SAR is because adiabatic

pulses require a large transmit field amplitude to satisfy the adiabatic condition and may also have a broad bandwidth (BW). Additionally, adiabatic pulses usually have a longer duration than conventional pulses (typically several ms) which could be a problem for minimising TR. The most common adiabatic pulse is the hyperbolic secant pulse, but others exist like Lorentz, Gauss and Hanning among others [38].

2.4.10 Spoiling Mechanisms

Ideally, it would be possible to completely eliminate the transverse magnetisation prior to each repetition which would allow for pure T_1 -weighted images. However, in practice, it is not always completely possible. There are three mechanisms behind spoiling which are commonly used.

Long TR Spoiling If the TR is much greater than T_2^* then the transverse magnetisation will decay to zero before the next repetition. Hence, a sequence using TR values of several hundred milliseconds or longer will be spoiled.

Gradient Spoiling Spoiler gradients are played out at the end of each TR, just before the next RF pulse. The amplitude of the spoiler gradient is either linearly or semi-randomly varied between TRs to ensure that magnetisation is spatially non-uniform.

RF Spoiling RF spoiling uses a constant gradient spoiler and also increments the phase of the RF pulse (as well as the numerical crystal oscillator for the DAQ) for each RF pulse. RF spoiling is superior to gradient spoiling as it is spatially invariant and does not generate eddy currents. The phase of the n -th pulse, $\phi(n)$, is given by linearly increasing the phase by an RF-spoiling increment Φ from an initial phase, ϕ_0 .

$$\phi(n) = \phi_0 + n\Phi \quad (2.24)$$

Typical values for Φ are 50° , 117° or 123° [39, 40] but varies between vendors and the optimal value depends on the T_1 of the tissue being scanned.

2.4.11 Signal and Noise in MRI

An important quality metric in MRI is the SNR which is defined as

$$\text{SNR} = \frac{\text{Signal}}{\sigma_{\text{noise}}} \quad (2.25)$$

where σ_{noise} is the standard deviation (SD) of the noise.

Generally, there is little perceived improvement in image quality for SNR values above 20. The fundamental factor influencing the size of the signal is the number of polarised protons within each voxel. Hence, the SNR is proportional to both B_0 , which is generally fixed, and the voxel volume, ΔV , i.e. image resolution. Hence, any parameters that effect the voxel volume such as the FOV, matrix size, slice thickness etc. all have an impact on the SNR. This means adjusting the FOV and matrix size based on anatomy and resolution requirements is key to retaining SNR and ensuring good image quality. Also, since the signal for tissue close to the receive coil is larger due to higher coil sensitivity, signal is proportional to the coil sensitivity, S_{Rx} . A standard step in MRI image processing is removal of the receive field contribution to provide more homogeneous image intensity. On Siemens systems this is called "prescan normalise". Hence, the MRI signal is proportional to the B_0 field, the coil sensitivity, and the voxel volume.

$$\text{Signal} \propto B_0 \cdot S_{Rx} \cdot \Delta V \quad (2.26)$$

Another parameter that affects the MRI signal is the FA. When the FA is increased, the SNR of the image also increases, up to a certain point. This is because increasing the FA causes more protons to be excited and produce a stronger MRI signal. However, if the FA is too high, the image can become saturated, meaning that all the protons have been excited and there is no further signal increase. In this case, further increasing the FA will not result in a higher SNR and may even decrease the image quality due to artefacts such as stimulated echoes. An appropriate FA is crucial to achieving optimal SNR in MRI. The optimal FA depends on various factors such as the tissue being imaged, the imaging sequence used, and the desired trade-off between SNR and image contrast.

Increasing the TR also increases the SNR due to increased T_1 relaxation, hence more longitudinal magnetisation is available for excitation. For a GRE sequence, the FA that achieves the optimal SNR is given by what is known as the Ernst angle, α_E , given by Equation 2.27 [41]. Although, the Ernst angle does not give the optimal contrast.

$$\alpha_E = \arccos\left(e^{-\frac{TR}{T_1}}\right) \quad (2.27)$$

Noise in the MRI image is primarily due to the thermal "Johnson" noise of the imaging system. This comes from a combination of the coils own resistance, electronic noise in the preamplifier and the resistance of the subject themselves. This noise can be modelled as complex Gaussian and zero-mean. Structured noise resulting from undersampling is discussed briefly in Section 2.5.3. Noise due to physiological motion is discussed in Section 8.2. Hence, increasing the time spent acquiring samples, T_{total} , can improve the SNR by signal averaging, since the MRI signal is cumulative but random noise is not. This means acquiring more repeats or acquiring more phase encoding lines etc. all have a positive impact on SNR.

$$\sigma_{noise} \propto \frac{1}{\sqrt{T_{total}}} \quad (2.28)$$

Since noise is evenly distributed across the frequency spectrum, increasing the BW increases how much noise is mixed into the signal from frequencies outside of the signal of interest. Alternatively, increasing the bandwidth can be viewed as decreasing T_{total} , hence reducing SNR.

$$SNR \propto B_0 \cdot S_{Rx} \cdot \Delta V \cdot \sqrt{T_{total}} \quad (2.29)$$

Changing the bandwidth has other effects on the imaging protocol, as summarised by Table 2.1.

Lower BW	Higher BW
More Chemical Shift Longer minimum TE/TR Increased Susceptibility Artefacts Worse metallic artefacts	Lower SNR Larger minimal FOV

Table 2.1: Summary of the impact of bandwidth on the MRI protocol and images.

The aim of optimising a sequence protocol is to balance the trade-offs between SNR, image resolution and acquisition time. Although, this is not the only consideration for achieving optimal image quality. Contrast is also an important factor to consider and the potential for artefacts in MRI are abundant and therefore require careful mitigation.

2.4.12 Extended Phase Graphs

The Bloch equations can be solved for the evolution of magnetisation for an isolated isochromat. However, in reality, a voxel will contain an ensemble of spins at slightly different frequencies. In order to accurately model the real MRI signal, many Bloch simulations should be performed and averaged to represent the ensemble but these simulations can be extremely time-consuming.

Extended phase graph (EPG) simulations offer advantages over Bloch simulations by using a Fourier representation of the magnetisation in "configuration states" [42–44]. A configuration state is defined as a state that corresponds to a discrete multiple of dephasing from a single unit gradient that gives a phase twist of one cycle (2π) across a voxel. This formalism allows us to easily simulate the generation of echoes and offers additional benefits in simulation speed and has been applied to a variety of applications such as modelling RF spoiling in GRE sequences [45] and optimising sequence design [46]. Open-source MATLAB code is freely available online from Brian Hargreaves: <http://web.stanford.edu/~bah/software/epg>.

The states F_n^+ and F_n^- denote the dephasing and rephasing transverse magnetisation states of order n , which are simply phase twists in direction of the sign of n . Whereas the states Z_n denote the longitudinal magnetisation states of order n , which represent sinusoids of increasing frequency. The relationship between the F^+ , F^- and Z_0 configuration states and the transverse and longitudinal magnetisation is as follows

$$F_n^+ = \int_0^1 M_{xy}(z) e^{-2\pi i n z} dz \quad (2.30)$$

$$F_n^- = \int_0^1 M_{xy}(z) e^{2\pi i n z} dz \quad (2.31)$$

$$Z_n = \int_0^1 M_z^*(z) e^{-2\pi i n z} dz \quad (2.32)$$

In the EPG formalism unit gradients are implemented through a simple "shift" operator which raises or lowers the configuration states and can accurately predict effects such as stimulated echoes (STE). For similar reasons, this formalism also allows the effects of coherent flow and diffusion [47] to be easily studied. The configuration states F^+ , F^- and Z_0 are stored within a matrix Ω

$$\Omega = \begin{bmatrix} F_0 & F_1 & F_2 & \dots \\ F_0^* & F_{-1} & F_{-2} & \dots \\ Z_0 & Z_1 & Z_2 & \dots \end{bmatrix} \quad (2.33)$$

The state matrix can be evolved in time under various conditions such as gradients, RF pulses, relaxation etc. which can be described by the simple matrix multiplication with the corresponding matrix operator.

Applying a gradient will cause an increase or decrease in the F_n^+ or F_n^- state number n

$$\mathbf{G}(\Delta n)F_n = F_{n+\Delta n} \quad (2.34)$$

$$\mathbf{G}(\Delta n)Z_n = Z_n \quad (2.35)$$

Applying an RF pulse will result in a mixing of coefficients between Z_n , F_n^+ and F_n^- . The excitation operator, $\mathbf{E}(\alpha, \phi)$ is a function of the FA, α , which has a phase, ϕ , and is given by

$$\mathbf{E}(\alpha, \phi) = \begin{bmatrix} \cos^2(\frac{\alpha}{2}) & \exp(2i\phi) \sin^2(\frac{\alpha}{2}) & -i \exp(i\phi) \sin(\alpha) \\ \exp(-2i\phi) \sin^2(\frac{\alpha}{2}) & \cos^2(\frac{\alpha}{2}) & i \exp(-i\phi) \sin(\alpha) \\ -\frac{i}{2} \exp(-i\phi) \sin(\alpha) & \frac{i}{2} \exp(i\phi) \sin(\alpha) & \cos^2(\alpha) \end{bmatrix} \quad (2.36)$$

Relaxation has the effect of attenuating various coefficients of the state matrix. T_2 relaxation will cause F_n coefficients to become attenuated whilst T_1 relaxation will attenuate Z_n coefficients and cause Z_0 to recover to equilibrium. This is summarised by the following relaxation operator $\mathbf{R}(\tau, T_1, T_2)$

$$\mathbf{\Omega}(t + \tau) = \mathbf{R}(\tau, T_1, T_2)\mathbf{\Omega}(t) = \begin{bmatrix} E_2 & 0 & 0 \\ 0 & E_1 & 0 \\ 0 & 0 & E_1 \end{bmatrix} \mathbf{\Omega}(t) + \begin{bmatrix} 0 \\ 0 \\ 1 - E_1 \end{bmatrix} \quad (2.37)$$

where $E_1 = \exp\left(\frac{-\tau}{T_1}\right)$ and $E_2 = \exp\left(\frac{-\tau}{T_2}\right)$. In Eq. 2.37 the first term of the final expression describes the attenuation whilst the final term allows for T_1 recovery to equilibrium.

Additional matrix operators which simulate the effect of flow and diffusion are defined in Section 5.2.1. Recently, phase distribution graphs which provide quick, differentiable, and spatially encoded Bloch simulations were proposed overcoming some of the limitations of EPG simulations [48].

2.5 Image Reconstruction

MRI image reconstruction by a simple fast Fourier transform (FFT) is quick and straightforward but this assumes a fully-sampled (and Cartesian) k -space which is time-consuming (and unnecessary) to acquire. This section will talk about some popular image reconstruction techniques in MRI. Openly-available tutorial material on image reconstruction methods provided by Mark Chiew can be found at: <https://mchiew.github.io/Teaching.html>.

2.5.1 Partial Fourier

One method to make the acquisition of k -space faster is to use a higher BW, at the cost of reduced SNR, and is limited in practice by hardware constraints. Another option is to simply reduce the number of k -space samples that are acquired. However, once the Nyquist criterion is broken more complicated image reconstruction methods must be employed. Due to the conjugate (Hermitian) symmetry nature of k -space, partial Fourier (PF) methods, in theory, allow up to 50% of k -space to be synthesised. In practice, however, phase PF factors of 75% or above are generally used since there is always some degree of phase error, whether from physiological motion or hardware.

When a PF technique is employed in the readout dimension by only measuring a portion (e.g. 2/3) of an echo this is called asymmetric-, partial- or fractional-echo. This allows the TE to be reduced significantly with only a small SNR penalty and also benefits flow and motion robustness by allowing smaller gradients moments. However, using asymmetric echo does not directly reduce the acquisition time.

2.5.2 Zero-padding

Zero-padding k -space prior to FFT is a quick trick and improves the apparent resolution which is mathematically equivalent to a linear interpolation in image space. Zero-padding results in image blurring, and is not optimal. Typically, zero-padding up to 2 or 3 times the original scan matrix is recommended and further zero-padding results in little-to-no benefit. Additionally, without careful windowing it results in a well-known artefact known as Gibbs ringing. Hence, k -space data should be appropriately filtered (typically Hann or Fermi) before zero-padding to ensure a smooth transition from zeros to actual k -space samples.

2.5.3 Parallel Imaging

Phased array coils which are widely used on most modern MRI scanners consist of multiple receive coils. Each individual coil provides high sensitivity but over a limited region. Hence, a large collection of these is able to provide high SNR over a greater FOV. The use of such phased array coils also allows for a method of acceleration known as parallel imaging. Parallel imaging involves sub-sampling k -space by a factor of R together with knowledge about the individual coil sensitivities to undo aliasing of the image during reconstruction. This effectively subverts the Nyquist sampling limits [49]. In 2D imaging it is only possible to undersample in one direction, which limits the amount of acceleration that can be achieved. 3D imaging can be undersampled in two dimensions and due to the inherent higher SNR allows for higher acceleration factors.

Parallel imaging method can be divided into two types. k -space based such as GRAPPA and image based methods such as SENSE. The following sections will briefly discuss these two fundamental parallel imaging methods.

2.5.3.1 SENSE

Sensitivity encoding [50] (SENSE) is an image domain based method which relies on the concept that the sensitivity of a receive channel further encodes information based on its location with respect to the subject. A requisite of the SENSE method is that it needs the undersampling to be regular and Cartesian, and receive-channel coil sensitivity maps are also required. Sensitivity maps are subject-dependent so either need to be acquired for each scan, or estimated from the data directly (e.g. ESPIRiT [51], see Section 2.5.5.1).

A simple example of SENSE is explained in Figure 2.7 for a Shepp-Logan phantom with $R = 2$. Assuming undersampling from top-to-bottom, $\Delta k_x = 0.5\Delta k_y$, therefore according to Equation 2.18, the FOV along x is half that along y . This causes a problem if the object is larger in the x dimension by $\text{FOV}_x/2$ as is the case here. A pixel in the resulting aliased image, y , is the sum of pixels x_1 and x_2 from the fully sampled unaliased image.




Figure 2.7: Demonstration of the SENSE problem for aliased images.

SENSE enables the ability to tease apart the aliasing by solving the linear set of equations. For a single coil system, there would only be a single equation and this would be under-determined, but thanks to phased array coils, this is usually over-determined, as long as the acceleration factor is not larger than the number of linearly independent coils. In order to obtain the unaliased image, SENSE must solve these equations for every set of aliased pixels in the image. For a two coil system, these equations would be:

$$x_1 \cdot S_{1,1} + x_2 \cdot S_{1,2} = y_1 \quad (2.38)$$

$$x_1 \cdot S_{2,1} + x_2 \cdot S_{2,2} = y_2 \quad (2.39)$$

where $S_{i,j}$ represents the sensitivity of the i^{th} receive coil at a position j .

2.5.3.2 GRAPPA

Generalized autocalibrating partially parallel acquisitions [52] (GRAPPA) is evolved from partially parallel imaging with localized sensitivities [53] (PILS). Whereas SENSE will reconstruct a single (coil-combined) unaliased image, GRAPPA differs in that it is a k -space based method that intends to fill in the missing k -space data directly. An illustrative example of GRAPPA is shown in Figure 2.8. To fill in the missing k -space data requires a region of fully-sampled k -space data known as the autocalibration signal (ACS) lines (grey) with which a small kernel (e.g. 3×3) passes over to learn the required weights from multiple coils (purple). Once the GRAPPA kernel is adequately trained, it can be applied to the missing k -space lines to predict the missing data (orange).




Figure 2.8: Diagram illustrating the training (purple) and application (orange) of a 3×3 GRAPPA kernel from multi-coil data. Filled-in circles represent sampled k -space data whereas unfilled circles represent unacquired data.

The reason GRAPPA is able to recover the missing k -space points is as follows. The signal in each voxel of an image acquired using a receive coil array is weighted by the appropriate coil sensitivity profile. This means that in k -space, there is a smearing of information across the local k -space neighbourhood. Due to this smearing, there is information about a particular k -space point in its local neighbourhood. Hence, as the smearing is the same across the entirety of k -space, if it is possible to learn what the effect of this smearing is, then it should also be possible to obtain information about the missing neighbouring k -space points [54].

2.5.3.3 g-Factor

Typically, undersampling factors of $R = 2$ and $R = 3$ are used with 32-channel receive head coils. The reduction in time spent imaging from undersampling comes at the cost of $\sqrt{R}$ SNR penalty, due to the reduction in the number of samples. Additionally, unlike a fully-sampled acquisition, the noise level changes from pixel-to-pixel with noise correlation between pixels. Hence, the SNR has a spatially varying component known as the g -factor. The g -factor depends of the coil configuration such as the number of receive channels and their sensitivity profiles but also on the undersampling and $g \geq 1$. Generally, g -factor noise increases for higher acceleration factors. The reduced spatially variant $\text{SNR}_{\text{PI}}(r)$ is then

$$\text{SNR}_{\text{PI}}(r) = \frac{\text{SNR}}{g(r)\sqrt{R}} \quad (2.40)$$

2.5.4 Compressed Sensing

Compressed sensing is at the heart of modern image file formats such as JPEG [55], which allows substantial compression of images without any noticeable reduction

in image quality. In parallel imaging, redundancy by virtue of having multiple receive channels is exploited, whereas compressed sensing (CS) exploits the intrinsic redundancy in MR images. Compressed sensing usually results in much higher acceleration factors being achievable than for parallel imaging techniques, due to the lack of structured artefacts that are typically seen for parallel imaging.

A successful CS reconstruction has three requirements: sparsity, incoherence and the use of non-linear reconstruction [56] which ensures data consistency with the acquired samples. Most MR data is sparse in at least one domain. A sparse image is one that has redundant information as the number of voxels is much larger than the degrees of freedom. The artefacts which would appear in the linear reconstruction should appear noise-like (i.e. incoherent) in the sparsifying transform domain. In other words, this means the Fourier components must be randomly sampled. Random sampling is usually performed using a variable density Poisson disc distribution which is weighted to sample more densely closer to the centre of k -space than in the periphery.

2.5.5 Coil Sensitivity Estimation

The multiple receiver coils will each produce their own receive channel image, I_r . Initially the root sum of squares (RSS) method was used routinely to combine these into a single image, I_{RSS} , which is nearly optimal when there is sufficient SNR, but results in a loss of phase information producing magnitude-only images.

$$I_{RSS} = \sqrt{\sum_{r=1}^{N_{Rx}} I_r I_r^*} \quad (2.41)$$

However, the optimal combination of receive channel images needs the sensitivity information for each coil [57]. The receive channel images can then be combined into a single image, I_{sens} , by the multiplication of the receive channel images for each coil with the complex conjugate of the sensitivity profiles for each coil, S_r^* , and summing over all receive coils.

$$I_{sens} = \sum_{r=1}^{N_{Rx}} I_r S_r^* \quad (2.42)$$

Adaptive coil-combination is the current standard on Siemens systems and outperforms RSS when there is lower SNR, and works well for background noise reduction [58–61].




Figure 2.9: Example of single receive coil images for a subset of the receive coils (top row) and the associated receive coil sensitivities (bottom row).

2.5.5.1 ESPIRiT

Eigenvalue-based iterative self-consistent parallel imaging reconstruction (ESPIRiT) [51] estimates the coil sensitivities from k -space data and is used extensively in this work. It can often be difficult to accurately measure the receive coil sensitivities as even small errors can result in visible artefacts in the combined image. However, algorithms based on learned correlations, such as GRAPPA, tend to fail for high acceleration factors, but are much more robust to errors. Hence, ESPIRiT combines the advantages of SENSE with robustness to errors like GRAPPA.

Typically, a fully-sampled region from the centre of k -space is used for sensitivity map estimation, known as the calibration region, which corresponds to low spatial frequencies. Although, to be accurate a sufficiently large region should be used, typically at least 20×20 . The sensitivity profiles are estimated by an eigenvalue decomposition of an operator, $\mathbf{W}$, which is applied patch-wise across the calibration region. $\mathbf{W}$ is composed of two matrix operators: $\mathbf{R}_r$ which takes a patch of N samples from the calibration region centred on a location r as well as the matrix $\mathbf{V}_{||}$ which spans the rows of Block-Hankel representation of the calibration region, known as the calibration matrix, $\mathbf{A}$. The operator $\mathbf{W}$ is then defined as

$$\mathbf{W} = \frac{1}{N} \sum_r \mathbf{R}_r^\dagger \mathbf{V}_{||} \mathbf{V}_{||}^\dagger \mathbf{R}_r \quad (2.43)$$

where $\dagger$ denotes the Hermitian (complex conjugate transpose).

Performing an eigenvalue decomposition on $\mathbf{W}$ yields the coil sensitivities when eigenvectors are chosen with eigenvalue of 1.




Figure 2.10: Diagram illustrating the formation of the calibration matrix from the calibration data. Circles represent k -space samples.

2.5.6 Virtual Coils

With the invention of parallel imaging, larger and larger receive channel arrays have been constructed. Up until relatively recently, the standard number of receive channels was around 8. However, nowadays arrays of 32, 64 or even up to 128 are common. The impact of all this extra data on the reconstruction pipeline can be substantial and, despite similar improvements in compute capabilities, increases in the reconstruction time are significant. Virtual coils can help speed-up reconstruction by compressing the number of receive channels down to a "virtual set" by performing a linear combination in k -space. One common method of doing this is using an singular value decomposition (SVD) [62–64].

2.5.7 Low Rank Methods

If MRI data is indeed linearly predictable (i.e. the Nyquist sampling theorem can be broken because missing samples can be calculated based on a linear combination of measured samples), then it has been shown that by forming the image data into Hankel matrices it will be low-rank [65]. By ensuring a low-rank structure, missing data can be interpolated through structured low-rank matrix completion methods [66, 67]. PRIMO [68] is an extension of ESPIRiT [51] able to generate receive and transmit sensitivity profiles from collected k -space data, by performing a truncation of a SVD to give the roots of the null-space vectors, after concatenating transmit and receive sensitivities along different axes of the Hankel matrix.

2.5.7.1 TxLR

These low rank methods have also been extended to calibrationless multi-transmit mapping using low-rank tensor completion (TxLR) [69], differing to PRIMO in that the transmit rank is simultaneously constrained with the receive rank from two different unfoldings of the k -space data, and was shown to achieve $8\times$ acceleration (equal to the number of transmit channels) from under-sampled calibrationless k -space data.

A 2D k -space kernel of size $m \times n$ (e.g. 3×3) is rastered over the k -space data ($N_{k_x} \times N_{k_y} \times N_{\text{Rx}} \times N_{\text{Tx}}$) to transform the multidimensional data into a block-Hankel structured matrix ($N_1 \times N_2 \times N_{\text{Rx}} \times N_{\text{Tx}}$) where $N_1 = m \times n$ and $N_2 = (N_{k_x} - m + 1) \cdot (N_{k_y} - n + 1)$. The kernel maps to an $N_1 \times 1$ column vector as shown by the dashed boxes in Figure 2.11.




Figure 2.11: Illustration of transforming multidimensional k -space data to block-Hankel structured matrices. The dashed boxes demonstrate how the $m \times n$ (3×3) kernel maps between the k -space data and the block-Hankel structured matrices. Filled in circles represent acquired k -space samples. Figure adapted from Hess et al. [69]

The block-Hankel structured matrices can then be unfolded by either concatenating along the receive dimension resulting in R_c ($N_1 \cdot N_{Rx} \times N_2 \cdot N_{Tx}$). Otherwise, the block-Hankel structured matrices can be unfolded along the transmit dimension resulting in T_c ($N_1 \cdot N_{Tx} \times N_2 \cdot N_{Rx}$), as shown in Figure 2.12.




Figure 2.12: Two different unfoldings of the block-Hankel structured matrices. Filled in circles represent acquired k -space samples. R_c ; receive vertical concatenation matrix, T_c ; transmit vertical concatenation matrix.

The TxLR reconstruction is then formulated as a non-convex problem with strict rank constraints on both R_c and T_c simultaneously. The non-convex optimisation problem is solved by the alternating direction method of multipliers (ADMM) algorithm [70]. The ADMM algorithm is ran for a specific number of iterations or instead a chi-square heuristic can be used for a parameter-free stopping criterion based on input from the receive channel array noise characteristics.

2.6 Imaging the Moving Heart

Motion in cardiac imaging is made up of the vector sum of three parts. Motion caused by the pumping of blood by the heart and by respiration are fairly periodic and relatively predictable. However, additional voluntary or involuntary bulk subject movement (e.g. coughing) occurs randomly and in an unpredictable pattern. Due to this, many scans are restricted to only imaging the myocardium during the diastolic phase, where cardiac motion is minimal [71].

One simple and commonly used method to reduce motion caused by respiration are breath-holds. Healthy subjects can usually hold their breath comfortably for around 30s, but those suffering from breathlessness may last considerably less. Breath-holds are usually performed at the end of expiration as it gives a more reproducible position for subsequent measurements and the blood flow in the heart is less affected [72, 73]. An alternative to breath-hold is to use bellows which monitor the subjects breathing and can be used to gate an acquisition or respectively bin the data. However, this requires additional hardware and can be difficult to setup and uncomfortable for the subject.

Cardiac motion is routinely monitored by electrocardiogram (ECG) and monitoring the R-wave, the biggest ECG signal using three or four electrodes. Prospective synchronisation detects the patients R-wave and data is acquired after a fixed time called the trigger delay, see Figure 2.13. To maximise the available acquisition time a trigger window of zero can be chosen, but a nonzero value is common to restrict the acquisition to diastole, when the heart muscle is at rest.




Figure 2.13: Illustration of the ECG waveform over an R-R interval and the associated imaging intervals. Figure adapted from Berstein et al. [38]

The issue is the ECG trace inside the bore is different to that outside of the bore, so the ECG trace needs to be calibrated inside the bore. This is because conductive blood flowing in the aorta inside the B_0 field induces a voltage, due to the magnetohydrodynamic effect, which is present in the ECG trace and noticeable largely around the T-wave. The solution is vectorcardiography which allows for the R-wave to be much more reliably determined. However, the ECG trace can be

significantly effected at higher field strengths and can be challenging at 7 T. Hence, often a pulse oximeter is used for gating which measures the optical properties of blood. However, the time delay due to measuring the pressure wave pulse delay at the finger (or wherever it is placed) must be estimated and the signal-peak is much broader than for ECG leading to further inaccuracies. More recent alternatives to ECG are discussed in Section 8.

2.7 Summary

In this chapter, I have briefly covered some of the important concepts and tools in MRI and in this thesis. A more comprehensive introduction to MRI Physics can be found in the "green bible of MRI" [74]. The next chapter will explore some of the quirks of ultra-high field MRI.

3

Ultra-high Field MRI

Contents

3.1	Introduction	38
3.2	A Brief History of UHF MRI	38
3.3	Advantages and Challenges of UHF	38
3.3.1	B1+ Inhomogeneity	39
3.3.2	Specific Absorption Rate	41
3.4	Parallel Transmission	42
3.4.1	pTx Safety Tests	43
3.4.2	SAR Supervision	44
3.4.3	RF Shimming	45
3.4.4	Dynamic Parallel Transmission	45
3.5	B1+ Mapping	48
3.5.1	Relative B1+ Mapping	49
3.5.2	Absolute B1+ Mapping	50
3.5.2.1	Double Angle Method	51
3.5.2.2	Actual Flip Angle Imaging	53
3.5.2.3	Dual Refocusing Echo Acquisition Mode	54
3.5.2.4	SA2RAGE	55
3.5.2.5	MP3RAGE	56
3.5.2.6	Saturation TurboFLASH	57
3.5.2.7	Fingerprinting	59
3.5.2.8	Phase Sensitive	59
3.5.2.9	Bloch-Siegert Shift	60
3.5.2.10	BEAR	61
3.5.3	Off-resonance Effects	62
3.5.4	B1+ Mapping Acquisition and Encoding Strategies	62
3.5.4.1	B1TIAMO	66
3.6	Summary	67

3.1 Introduction

Ultra-high field (UHF) concerns any MRI scanner where $B_0 \geq 7$ T. UHF MRI can generate significantly higher SNR than that of conventional clinical 3 T MRI systems. The improved SNR can help reduce the acquisition times of scans or can alternatively be traded-off for higher spatial resolution [75].

3.2 A Brief History of UHF MRI

The first human 7 T MRI research system became available in 1999 [76]. Since then, UHF MRI has become ubiquitous in human neuroimaging research centres worldwide. In 2017, the Magnetom Terra became the first 7 T scanner in the United States which was cleared by the FDA for clinical imaging of the brain and knee, shortly followed by European CE certification. These clearances have paved the way for a shift from exclusive research applications to the utilisation of 7 T MRI systems in both research and clinical settings. Today, there are around 100 UHF MRI scanners worldwide and there are serious efforts to develop 14 T MRI scanners for human neuroimaging [77].

3.3 Advantages and Challenges of UHF

The most notable benefit of UHF MRI is the improved SNR. The stronger the magnetic field, the larger the proportion of nuclei that are aligned. Greater numbers of aligned nuclei correspond to a higher SNR. An exact relationship between SNR and B_0 is difficult to investigate practically due to differences in RF coils and hardware. However, the maximum achievable SNR gain, independent of RF coil geometry, is given by ultimate intrinsic SNR theory [78]. Typically, voxels near the centre of the head will increase SNR with B_0 approximately quadratically, whilst voxels near the surface will increase SNR with B_0 slightly more than linearly, hence SNR is $\propto B_0^{1 \text{ to } 2}$ [79]. The higher SNR at UHF can therefore be traded-off for higher image resolution or faster scan times. However, there are other effects of UHF which come from a change in the relaxation constants. For example, for T_1 tends to increase with field strength for ^1H nuclei, which could be beneficial for some scans but may also result in longer scan times. Decreased T_2 and T_2^* could be beneficial for BOLD sequences. The advantages and disadvantages of UHF are summarised in Table 3.1.

Property	Trend with increased B_0	Advantage	Disadvantages
SNR	$\uparrow$	Higher resolution, reduced scan time, X-nuclei more feasible	None
SAR	$\uparrow$	None	Reduced FA's, reduced slices, longer TRs, additional patient exclusion [†]
Physiological side-effects	$\uparrow$	None	Dizziness, nausea, metallic taste
Relaxation time	$T_1 \uparrow$ $T_2 \downarrow$ $T_2^* \downarrow$	TOF, ASL, VASO SWI, BOLD	Longer scan times DWI, DTI
RF homogeneity	$\downarrow$	Parallel reception, parallel transmission	Spatially dependent flip angle
Susceptibility effects	$\uparrow$	BOLD, SWI, T_2^*	Geometric distortions, intravoxel dephasing, worse metallic implant artefacts
Chemical shift	$\uparrow$	Fat saturation, CEST, MRS	Fat/water and metabolite misregistration

Table 3.1: Summary of the advantages and disadvantages of UHF MRI. [†]Potential for excluding patient populations such as pregnant women, young children, those with poor temperature regulation, implants etc. Reproduced from Ladd et al. [80]

The advantages of UHF have also translated to better clinical outcomes [81]. 7 T MRI has proven to be better at detecting subtle epileptogenic lesions in the brain which would often go undetected using conventional clinical 3 T systems [82], leading to improved patient surgical planning and treatment. The prostate is another potential area of interest for the application of UHF MRI and phosphorous (^{31}P) spectroscopic methods, which benefit from the higher SNR, could identify important biomarkers for prostate cancer aggressiveness [83].

Two main challenges of UHF MRI, namely B_1^+ inhomogeneity and SAR are discussed below, though I will primarily focus on B_1^+ inhomogeneity as it is the main focus of this thesis. Motion is a problem across all field strengths, but particularly for UHF, and I will discuss this in Chapter 8. A more comprehensive evaluation of UHF MRI for human imaging can be found by Ladd et al. [80]

3.3.1 B_1^+ Inhomogeneity

Due to the ^1H Larmor frequency, as explained in Section 2.3.3, at 7 T the associated wavelength is approximately 13 cm. However, the exact wavelength depends on the electrical properties of the tissue, i.e. the relative permittivity and conductivity of

the tissue, which act to attenuate the B_1^+ field. This short wavelength is therefore comparable to the size of the human head and RF interference can result in regions of complete signal drop-out and localised heating [24–26]. Ideally, the RF transmit field would have a uniform spatial response. And indeed, at 1.5 T, the FA requested from the scanner is the FA that is produced mostly everywhere in the brain. At 3 T and above, there is an increasing dynamic range (i.e. ratio between the largest and smallest values) of FAs across the FOV. The effects of this FA distribution on a standard GRE image are demonstrated in Figure 3.1. This inhomogeneity in the B_1^+ field not only poses significant challenges for imaging but particularly for performing quantitative mapping.




Figure 3.1: Simulations of three flip angle distributions at $B_0 = 1.5$ T, 3 T and 7 T demonstrating the effect B_1^+ inhomogeneity on a GRE image. The flip angle distribution is exaggerated due to simulation oversimplifications but represent the expected profiles. Figure adapted from de Buck [84].

As a result, many methods have been proposed to deal with B_1^+ inhomogeneity over the years. One simple and low-cost method involves placing high permittivity dielectric pads around the area where an increase in B_1^+ is required [85, 86]. Metamaterials which behave similarly to dielectric pads are also of significant interest [87].

RF pulses can also be designed to be insensitive to B_1^+ inhomogeneity. Adiabatic pulses are formed from hyperbolic secant functions, which sweep both the magnitude and frequency of the RF pulse, are relatively insensitive to B_1^+ inhomogeneity and have been applied to overcome this problem [34, 35, 88]. However, these pulses are a large burden on SAR so must be used sparingly.

Specialised coil designs such as a spiral volume design [89] and detunable transverse electromagnetic resonators [90] have been proposed to reduce this B_1^+ inhomogeneity, but cannot eliminate it completely.

By far the most promising and widely used method for reducing B_1^+ inhomogeneity is parallel transmission (pTx).¹

3.3.2 Specific Absorption Rate

The SAR is the energy deposited in a specific mass of tissue due to the electric field component of RF pulses and has units of Watts per kilogram. SAR is used as a surrogate metric for the increase in tissue temperature which is the primary safety concern. The relationship between SAR and temperature increase is modelled by the Pennes bioheat equation [93]. Regions of high perfusion, such as the brain, can experience moderately high SAR levels and result in a minimal temperature increase as excess heat is removed by an increase in blood flow. Whereas for regions of lower perfusion, such as in the eyes, the temperature can increase significantly even with relatively low SAR. Poor perfusion can also pose a problem for patients with reduced thermoregulatory capacities such as infants, pregnant women, and the elderly.

The electric field generates electrical currents in conductive tissue resulting in tissue heating. To ensure safety there are strict legal limits on both the local and global SAR that a human can be subjected to. These limits are given for whole-body, partial-body or local (which ever is exceeded first). In normal operating mode whole-body SAR is limited to 2 W kg^{-1} and head local SAR (over a 10 g mass) to 10 W kg^{-1} averaged over six minutes. For short durations over 10s the SAR shall not exceed two times these values. These limits are set by the International Electrotechnical Commission Committee [94] and apply across all field strengths. The purpose of these SAR limits is to prevent a rise in whole body temperature above 0.5°C for normal operating mode and 1°C for first level operating mode.

SAR limits are used by coil manufacturers to define safe power limits which guarantee that the SAR limits will not be exceeded e.g. for the 8Tx/32Rx head coil at our institution the power limits are 5 W average per channel (6 min average) or 10 W per channel over 10 s. Large safety margins are included within these SAR limits to account for deviations between an individual subject and a realistic anatomical human body model.

¹Although other promising techniques such as multi-photon MRI exist [91, 92].

These limits can be easily reached, particularly at UHF as power deposition is proportional to B_0^2 , if the factors affecting the SAR are not carefully considered. Unlike lower field strengths, where SAR is fairly homogeneous, UHF is susceptible to local hot spots. Hence, one of the largest factors that is limiting the wider application of pTx is the issue surrounding the inability to accurately know the local SAR.

3.4 Parallel Transmission

Standard single transmit channel MR systems operate in circularly polarised (CP) (a.k.a quadrature, TrueForm [Siemens nomenclature]) mode. This describes the mode of operation for the RF coil, whereby the complex weights (magnitude and phase) of each transmit channel, w_i , are fixed and evenly spaced in phase (i.e. 45° between channels for an eight-channel system). CP^+ mode generally gives adequate homogeneity when imaging spherically shaped objects such as the brain (kind of) and is usually the default transmit mode for non-pTx imaging.

First investigated in 2000 [95, 96] and demonstrated in [97, 98] pTx evolved from parallel imaging. PTx makes use of multiple transmitter elements in an array each with their own transmission chain. Inhomogeneities in B_1^+ can then be corrected by altering the complex weights, as well as the individual pulse shape of the RF pulse transmitted on each individual transmit channel independently [99]. PTx can also help improve the RF transmit power efficiency (i.e. the flip angle achieved for a given input power) by designing pTx RF pulses to have both low global and local SAR. In addition to homogenising B_1^+ and lowering SAR, pTx has other benefits such as enabling an arbitrary excitation pattern². Arbitrary excitation patterns allow a reduced FOV by only exciting a restricted volume called inner-volume-imaging [100]. This helps to avoid artefacts either from aliasing or by not exciting moving anatomy etc.

However, the added complexity comes at a cost of approximately \$50k per channel, which typically limits systems to using fewer than 16 transmit channels. Figure 3.2 illustrates the additional hardware required for a full pTx system.

²Arbitrary excitation patterns are possible with static pTx, but dynamic pTx makes the RF pulse feasible in terms of length.




Figure 3.2: Illustrations of hardware required for **A)** conventional single-channel transmit, **B)** RF shimming and **C)** full dynamic pTx. Figure adapted from Williams et al. [101]

3.4.1 pTx Safety Tests

There are a number of tests to perform prior to human scanning. These tests must be performed manually on the Magnetom 7T (VB17 Step 2.3, Siemens Healthineers, Erlangen, Germany) in order to check the system is working correctly and as expected. These tests are:

1. Phase adjustment - Determines the relative phases of the modulators and receivers.
2. VOP shutdown test - Checks to make sure RF transmission is stopped when local SAR is exceeded (see Section 3.4.2).
3. RF monitor function test - Checks that both the global and local RF monitors and within the required specification.
4. PALI shutdown test - The per-channel peak power as well as the sliding window mean power over all transmit channels is averaged over 10s and 6 minutes. This test ensures that the power absorption limiter (PALI) disables RF transmission when these value breach the user-defined global SAR limits.

5. System check - Checks the DICO data is as expected.
6. DICO test - Checks RF amplitude and phase is as expected and is performed automatically before each measurement.

3.4.2 SAR Supervision

The whole-body averaged SAR is easily estimated from the forward versus reflected power [102]. However, due to the superposition of many fields from multiple transmit channels in pTx, this field is spatially and temporally dependent and unknown locations of hot spots can occur that may exceed the local SAR limits. This leads to worst-case local SAR estimations usually putting overly conservative limitations on the amount of RF power that can be applied. SAR electromagnetic simulations applied to modelled Tx-Rx arrays can be performed in software such as Sim4Life (ZMT, Zurich, Switzerland) or CST (CST AG, Darmstadt, Germany) using different tissue types on realistic anatomical models. These electromagnetic field solvers use a finite-difference time-domain method which can be extremely computationally consuming, often taking hours if not days for full electromagnetic simulations, and must be repeated for specific coils, anatomies and positions. This results in generally being restricted to using a few modes of excitations such as CP⁺ mode.

Fortunately, virtual observation points (VOPs) have become the default way to monitor local SAR with pTx [103]. This involves simulation of each transmit element and calculation of their interaction matrix (Q -matrix) for each spatial location. The Q -matrices contain all field and subject model information, but not RF pulse information and hence can be pre-calculated [104]. The Q -matrices are then clustered based on their maximum eigenvalues and compressed. The level of compression of these clusters determines how much maximum overestimation there is of the worst-case local SAR. Then, on the scanner, the local SAR can be estimated knowing the sequence parameters such as RF pulse waveform, TR and the Q -matrices.

Ultimately, pTx allows for reduced SAR due to the additional degrees of freedom available to us. A "SAR hopping" algorithm can minimise the local SAR by distributing the hot spots to different locations for each slice or group of slices [105]. Therefore hot spots do not accumulate and local SAR is reduced.

3.4.3 RF Shimming

It is possible to move away from using CP⁺ mode and alter the complex weights associated with each channel because each transmit channel has its own transmission chain. This is termed "RF Shimming" or "static pTx" and is performed to help reduce the B₁⁺ inhomogeneity. Generally, the static complex weights are computed for RF shimming by solving cost functions from acquired B₁⁺ maps, or a set of complex weights might be used which a researcher may have previous experience that it gives a good shim in a particular region of interest (ROI). RF shimming can be a good solution for homogenising B₁⁺ for individual slices at 7 T but static RF shimming is, however, limiting in that B₁⁺ inhomogeneities can be moved and reduced but not entirely eliminated when large volumes are imaged.

Time interleaved acquisition of modes (TIAMO) [106] is a relatively simple method to help overcome B₁⁺ inhomogeneity. TIAMO works by acquiring two or more different modes interleaved using static RF shimming. After each PE readout for the first shim mode, the same PE readout is performed for the second shim mode. Each readout line is then interpreted as a single line but with twice the number of receive channels. This would usually result in a scan-time penalty of 2, but since these two acquisitions can be seen as being received by virtual coils, acceleration from GRAPPA minimises this penalty and only a small loss in SNR is required, ultimately obtaining better homogeneity with noise being uncorrelated between modes since they are acquired at different times. The concept behind TIAMO was subsequently applied to produce B₁⁺ maps, known as B1TIAMO [107] (see Section 3.5.4.1).

3.4.4 Dynamic Parallel Transmission

Dynamic pTx unlocks another degree of freedom by also individually altering the pulse shape on each of the i -th transmitters, $p_i(t)$, in addition to the complex weights, w , that are altered during RF shimming. This allows for specialised subject-tailored pTx pulse design. The B₁⁺ at a particular time, t , and position, r , in space is found by the superposition of N_{Tx} transmission elements and given by

$$B_1^+(r, t) = \sum_{i=1}^{N_{\text{Tx}}} p_i(t) w_i S_i(r), \quad (3.1)$$

where $S_i(r)$ are the individual transmit channel sensitivity profiles. If the underlying transmit sensitivities are known, pTx RF pulses can be designed to overcome B₁⁺ inhomogeneities and reduce the SAR burden. An example of the improvement in

image quality for a pTx magnetisation prepared rapid gradient echo (MPRAGE) over a standard CP⁺ MPRAGE is shown in Figure 3.3.




Figure 3.3: Sagittal MPRAGE showing the difference in image contrast between a **A)** CP versus **B)** pTx acquisition. There is a subtle, but noticeable, B_1^+ contrast in the cerebellum when imaging in CP mode, as indicated by the yellow arrows. This is significantly reduced with a universal pulse pTx MPRAGE.

One of the initial pTx pulse design methods was the spatial domain small tip angle approximation method [108]. However, this relies on the small FA approximation to simplify the design of RF pulses with FAs less than 30° . Nevertheless, the small FA approximation works reasonably well for FAs up to 90° but breaks down for FAs above 90° e.g. inversion or refocusing pulses [109]. The Shinnar Le-Roux algorithm [110, 111] is good for large tip angle non-adiabatic frequency selective RF pulses but this only works for non-pTx RF pulses. Other methods were proposed to achieve larger FAs using an iterative correction using the Bloch equations such as the additive angle method [112]. The optimal control approach is more sophisticated [113] at the cost of increased computational burden. More recent methods by Zheng et al. [114] work similarly to the additive angle method but with a reduced computational burden when compared to the optimal control method, by using perturbation analysis of the Bloch equations. For pTx applications, RF pulse design should also take into account the local SAR as well as per-channel voltage limits, which can be included in the pulse design algorithm [115].

As mentioned, one of the desirable qualities of pTx is to image reduced FOVs. However, to do this long RF pulses are required as the RF pulse k -space trajectory needs to reach high spatial frequencies. The Transmit SENSE [97, 98] approach helps to overcome this problem and is based in the k -space domain making use of transmit field sensitivities, comparable to parallel imaging. Using Transmit SENSE the duration of pTx RF pulses can be shortened because their excitation path in k -space can be reduced without sacrificing robustness. This shortening of the pTx pulses not only makes them fast to perform but also reduces relaxation and off-resonance effects.

k_T -points is a non-slice-selective excitation [116] where RF energy is only deposited at a small number of k -space locations, usually less than 10, which helps overcome B_1^+ inhomogeneities and means a pulse duration can be sub-millisecond. RF power is only deposited during stationary periods in k -space, where no gradient is applied. Slice-selective ‘fast k_z ’ or ‘2D spokes’ [117–119] homogenise a slice even though the spokes are a 3D excitation. However, constraints imposed by the gradients and SAR can result in pulse durations of several milliseconds [120]. For spokes and k_T -points the pTx pulses consist of individual RF sub-pulses. These sub-pulses are typically fixed and use sinc pulses for spokes and rectangular pulses for k_T -points and therefore the only optimisation required is the per-channel weights, w , of RF sub-pulses and gradient blips between sub-pulses. As mentioned earlier, individual channels can have their own RF waveforms and this is utilised by pTx pulses such as spiral non-selective (SPINS) [121]. SPINS is a non-selective pulse design using a spiral pulse design with a duration on the order of a millisecond. k_T -points, 2D spokes and SPINS are good alternatives to adiabatic pulses, however, the underlying B_1^+ must be known in order to design these pulses to work well in a particular subject.

Regardless of the RF pulse optimisation algorithm chosen, the process of acquiring B_1^+ maps and calculating the individual channels RF waveform can take a considerable amount of time, potentially tens of minutes. The universal pulses (UPs) [122] technique is based on the idea that most subjects anatomy is relatively similar, so by acquiring a database of B_1^+ maps in participants, RF pulses can be designed to work across the whole dataset to mitigate RF inhomogeneity and can be applied to a general population without requiring individual B_1^+ maps. UPs have been shown to be highly effective. Recently, the application of the UPs to TIAMO has been demonstrated in the heart at 7 T and the pelvis at 10.5 T called “universal M modes” [123]. Initially, it was uncertain how well UPs would work for imaging the body and regions which are more variable in their composition, as up until recently they had only been applied in the brain. UPs have now been shown to work well at imaging the heart at 7 T, designed from a library of only 20 participants [124, 125].

However, fast online-customised pTx pulses (FOCUS) [126] has been shown to utilise UPs in a way that can generate individually optimised pulses requiring less than a minute in additional preparation time by acquiring fast B_1^+ maps and have superior performance over UPs alone. Additionally, FOCUS pulses generalise better than UPs when used with a different coil than what the B_1^+ database was collected with [127].

3.5 B_1^+ Mapping

The invention of pTx has caused a resurgence in the development of new B_1^+ mapping techniques. The goal of B_1^+ mapping is to obtain a measurement of the B_1^+ field which is smoothly varying and should be free of any anatomical structural contrast.³ The most basic use of B_1^+ mapping is to calibrate the transmit gain of the scanner, otherwise known as the reference voltage. This is defined, on Siemens systems, as the voltage required by a 1 ms rectangular pulse to achieve a 180° inversion. Aside from this, B_1^+ mapping is useful for RF shimming, quantitative MR, and pTx RF pulse calculation and may also help to better inform SAR predictions [129].

For the multi-array coils used in pTx systems at UHF a B_1^+ map must be obtained for each transmit element. This is generally difficult due to the limited dynamic range of a particular single transmit coil element and can be extremely time-consuming since a map for N_{Tx} channels must be acquired. Since scan time is precious and B_1^+ mapping is largely only used for scanner calibration, *in vivo* imaging requires that B_1^+ mapping be both fast and accurate. Additionally, B_1^+ mapping for single-channel systems does not need to be concerned about the phase of the B_1^+ . However, at UHF with pTx coils the relative phase between each of the transmit channels must also be measured.

In certain circumstances, knowledge of the B_1^+ field relative to another transmit channel may be all that is necessary and this is called "relative B_1^+ mapping" [130]. The conditions for whether relative or absolute maps are required for a specific application are outlined in Table 3.2.

³A caveat of this is a well-known perturbation of B_1^+ in the ventricles of the brain. Simulations suggest that this is a real effect, not a T_1 -related artefact, and due to the large contrast in electrical conductivity between cerebrospinal fluid and the surrounding brain tissue which leads to the ventricles being visible in B_1^+ maps [128].

Aim	Relative B_1^+ Phase	Absolute B_1^+ Magnitude
Phase only shim	✓	
Magnitude and phase shim	✓	✓
FA Map		✓
pTx FA Map	✓	✓

Table 3.2: Summary of the required B_1^+ information for a specific application.

In this section, I will explore a few conventionally used methods of B_1^+ mapping.

3.5.1 Relative B_1^+ Mapping

When knowledge of the absolute B_1^+ field is not required, relative B_1^+ maps may be acquired which are much quicker to acquire and typically uses a simple GRE acquisition. For simple phase-only shimming, since the phase is independent of T_1 , TE, TR and FA GRE images can be acquired at the Ernst angle for optimal SNR [130]. However, if relative B_1^+ magnitude is required there are a few constraints. For example, small FAs ($\leq 10^\circ$) must be used for the signal intensity to be approximately linearly proportional to B_1^+ . Additionally, to reduce bias caused by T_1 the TR should be carefully chosen.

Multiple relatives maps at increasing reference voltages can be acquired to increase the dynamic range, ensuring a small FA, and combining the images using a maximum likelihood estimation [131]. Relative maps can be updated to an absolute map by additionally acquiring an absolute map [117, 130], using any of the methods mentioned in the next section.

A set of relative maps for each transmit channel element, n_c , can be obtained from fast low tip angle GRE images [132], $I_{n_c}(r)$.

$$B_{1,n_c,rel}^+(r) = \frac{I_{n_c}(r)}{\sum_{n_c=1}^{N_C} |I_{n_c}(r)|} \quad (3.2)$$

Flip-angle maps are much sparser than the images acquired to estimate B_1^+ since the acquired image contains not just the transmit sensitivities but also T_1/T_2 weighting, PD weighting and Rx sensitivities. Hence, directly estimating B_1^+ from k -space means a subset of data can be collected which would otherwise not fully resolve the anatomy that is scanned [133]. Hess et al. demonstrated that in order to map the dominant transmit sensitivity in the heart only a 24×24 scan matrix is required [69].

Conventionally, fast low-tip angle GRE images are acquired for relative mapping channel-wise sequentially, i.e. one Tx channel is imaged at a time, before moving onto the next neighbouring Tx channel until all images are acquired. Alternatively, relative maps may be acquired coil-cycled [131], where one k -space line is acquired on each Tx channel. This is suggested to speed up the acquisition of the relative channel-wise maps by reducing the effect of neighbouring coil saturation. The importance of coil-cycling relative maps and the mechanism behind the improvements seen are investigated in Chapter 6.

Deep learning has been investigated to obtain relative maps from localiser scans, providing maps at high resolution and with no additional scan time [134].

Small FA GRE images have also been used to estimate single-channel absolute B_1^+ maps for a transceiver array [132]. Relative single-channel maps are scaled to absolute units by acquiring a single absolute map using a method of your choice. This work utilised two underlying assumptions. Firstly, the Tx and Rx sensitivities are distinguishable from their twisted patterns in opposite rotations. Secondly, RF coil arrays are distributed axially around the ROI with multiple symmetries in the axial plane. And thirdly, that the sum of magnitude of all Tx B_1^+ sensitivities resembles the sum of magnitude of all Rx B_1^- sensitivities. The resulting maps are, however, biased by the square root of proton density which leads to some slight PD-weighted anatomical contrast. However, since the PD in tissue is usually fairly constant, whereas B_1^+ varies considerably, this can be neglected.

3.5.2 Absolute B_1^+ Mapping

Absolute B_1^+ mapping is trickier. The MR signal depends on several variables and at least two images of a different transmit field contrast, to encode B_1^+ , must be obtained with typically a sinusoidal dependence between the signal intensity and B_1^+ . These images usually vary in that they use two different FAs or different repetition times (TR) to encode different B_1^+ contrast. The terms "FA map" and " B_1^+ map" are often used interchangeably. B_1^+ maps may be reported in units of $[\mu\text{T}]$, $[\text{Hz V}^{-1}]$ or simply as a FA $[\circ]$.

B_1^+ mapping in the body is particularly challenging due to cardio-respiratory motion and blood flow which results in phase shifts between successive acquisitions. The requirement of two images with different B_1^+ contrast for B_1^+ mapping therefore often leads to erroneous B_1^+ maps. Ideally, both images would be acquired within the same cardiac phase and respiratory position to acquire a reliable B_1^+ map. Additionally, the large size of the body requires a large dynamic range and it is not possible for a single transmit configuration to measure B_1^+ in the whole body.

The information of the actual FA is encoded within either the magnitude or the phase of these images, and so there are broadly speaking, two different approaches to absolute B_1^+ mapping. There is, however, some overlap. For example, some methods use the magnitude of two images to determine the FA but then also consider the phase of the images to extend the dynamic range. The dynamic range can be a problem because the methods used to measure α are only accurate within a certain range of FAs. Generally, these methods exhibit bias at high α and tend to highly underestimate the FA. At low FAs ($\approx \leq 20^\circ$) measurement noise means a wide range of FAs could be incorrectly measured.

Phase-based methods are more susceptible to B_0 inhomogeneities due to off-resonance, such as susceptibility jumps in anatomy like the brain stem, but unlike magnitude-based methods are not prone to errors caused by T_1 . There are many different methods and to date, no single method is dominant as they all have their pros and cons. A majority of methods perform poorly in larger regions of the body and where there is motion.

Generally, there are fewer phase-based methods than magnitude-based, but I will briefly cover the most prominent ones. Several methods are compared in more detail in the literature [135, 136].

3.5.2.1 Double Angle Method

The double angle method (DAM) is perhaps the most widely used and simplest approach to mapping the B_1^+ field. It is a 3D approach which overcomes slice profile errors and uses two nominally applied FAs of α_{nom} and $2\alpha_{nom}$ with a sufficiently long TR time (typically $5T_1$), to allow complete relaxation and hence there is no T_1 dependence.




Figure 3.4: Simplified double angle method (DAM) B_1^+ mapping sequence diagram.

The popularity of this method largely stems from the ability to use any basic SE or GRE sequence [137], which can be found on any scanner. This gives a good measure of the B_1^+ inhomogeneity but takes a considerably long time to acquire, equal to $2 \times 5T_1 \times N_{PE}$. The FA is then derived from a simple relationship between the two images (S_0 , S_1) from

$$\alpha = \cos^{-1} \left(\left| \frac{S_1}{2S_0} \right| \right), \quad (3.3)$$

for GRE methods or otherwise

$$\alpha = \sec^{-1} \left(\sqrt[3]{\frac{8S_0}{S_1}} \right). \quad (3.4)$$

for SE sequences.

Alternatives to reduce the long TR time include driven recovery, the authors dub this the "Compensated DAM" [138] method. Another method of TR reduction, at the expense of SNR, is pre-saturated DAM [139], where all signal is effectively nulled before the acquisition, then after a time delay, TD_{sat} , the imaging RF pulse is applied, so the signal is always constant across the sequence and independent of T_1 . There are multiple options for this saturation pulse such as composite pulses or adiabatic pulses like BIR-4 (B_1^+ insensitive rotation). These are not, however, perfect and may require multiple applications to successfully saturate the signal and are highly SAR inefficient. Alternatively, to a non-selective saturation, a catalysed RF pulse train may be applied to achieve a similar effect [140]. The dynamic range of DAM can be extended from $0 - 90^\circ$ to $0 - 180^\circ$ by also examining the phase of the images.

3.5.2.2 Actual Flip Angle Imaging

Actual flip angle imaging (AFI) uses two images formed from a dual-TR sequence, TR_1 and TR_2 where $TR_2 > TR_1$ which are interleaved and acquire the same PE line per dual-TR. It is typically used as a 3D sequence, but 2D versions have been proposed where only single slice B_1^+ maps are required [141].




Figure 3.5: Simplified actual flip angle imaging (AFI) B_1^+ mapping sequence diagram.

The ratio, r , of these two images is found and then α is calculated from

$$\alpha = \cos^{-1} \left(\frac{rn - 1}{n - r} \right), \quad (3.5)$$

where $n = \frac{TR_2}{TR_1}$ and require both TR_1 and $TR_2 \ll T_1$ [142]. Typical values of TR_1 and TR_2 might be 30 ms and 150 ms ($n = 5$), respectively. However, any value may be chosen as long as there is sufficient difference between them to provide enough sensitivity. This method is relatively fast and insensitive to T_1 but requires highly efficient spoiling of the transverse magnetisation between TRs [143]. It also requires the use of many dummy pre-scans to get the magnetisation into a steady state, unless a special preparation module is performed to allow a faster approach to a steady state [144]. As a 3D sequence, this has become one of the most well-used methods and is often used in literature as the "ground truth", as long as the sequence is well-spoiled and maintains a steady state [145]. However, due to the use of a high number of large FA imaging RF pulses ($\alpha_{nom} \approx 60^\circ$), AFI also has a relatively high SAR. The dynamic range of AFI can also be extended to larger FAs by considering the phase of the images, extending the theoretical range from approximately $0 - 102^\circ$ ($r = 0$ and assuming $n = 5$ in Eq. 3.5) to $0 - 180^\circ$.

AFI requires a steady-state is achieved through efficient spoiling of the transverse magnetisation which may or may not be an accurate assumption in the body. The presence of fat in subjects has been reported to cause large errors with B_1^+ maps using Cartesian AFI, due to insufficient spoiling [146]. This is not as much of an issue when imaging regions of lower fat content (i.e. the brain) but is problematic when imaging the body. The presence of fat causes ghosting artefacts which are responsible for both local and global inaccurate B_1^+ estimates and requires the use of longer TRs for improved spoiling or fat suppression [146].

Alternatively, Samsonov et al. demonstrated that in order to keep the short TRs, which are desirable for a clinical implementation of AFI, the acquisition requires multiple echoes where fat and water are out of phase [146]. The B_1^+ maps can then be calculated based only on the water signal images. Radial phase encoded acquisitions which are typically used in the body to overcome motion artefacts [147] could also improve fat signal spoiling due to the non-Cartesian trajectories [148].

Recent work by Egger et al. has shown an accelerated radial density-adapted [149] AFI sequence with randomised x-y-spoiling [150] which performed well in the heart and prostate with an acquisition time of only 30 s ($R = 12$) [151].

3.5.2.3 Dual Refocusing Echo Acquisition Mode

Dual refocusing echo acquisition mode (DREAM) is unlike other B_1^+ mapping sequences in that both images are taken quasi-simultaneously from the use of a STE [152, 153].




Figure 3.6: Simplified dual refocusing echo acquisition mode (DREAM) B_1^+ mapping sequence diagram. Square brackets indicate an imaging pulse train of length N_{PE} .

One image (S_{FID}) is the result of a FID decay whilst the other (S_{STE}) is the result of stored longitudinal magnetisation from higher order configuration states of the magnetisation (see Section 2.4.12). The images are taken using a stimulated echo acquisition mode (STEAM) imaging sequence [154, 155].

$$\alpha = \tan^{-1} \left(\sqrt{\frac{2S_{\text{STE}}}{S_{\text{FID}}}} \right), \quad (3.6)$$

DREAM has the additional benefit of providing a B_0 phase map and a transceive phase map (sum of the phases of the transmit and receive fields) with no extra scans required. DREAM is very insensitive to motion and is only weakly sensitive to T_1 and T_2 [156]. It has a theoretical measurement range of 0-90° but has been shown in practice to only be bias-free and accurate between 20-50° [157]. However, the dynamic range can be increased by taking multiple DREAM acquisitions at varying nominal α FAs (or voltages) [158]. Despite this, the inherently limited range of measurable FAs may not be suitable for the large dynamic range found in the body.

Additionally, DREAM is inherently sensitive to flow due to its use of STE's and so it is not applicable for mapping in regions such as the blood chambers of the heart [159–162] and suffers from a factor of 2 SNR penalty.

3DREAM [163] demonstrated the acquisition of a whole-brain map in just a couple of seconds using a single-shot method. However, this leads to a very long echo train which results in image blurring, ringing artefacts at boundaries as well as other issues resulting from a difference in the evolution of the FID and STE signals. The use of spiral readouts has helped to mitigate this blurring [164, 165].

3.5.2.4 SA2RAGE

Developed from the MP2RAGE scheme, saturation prepared with two rapid gradient echoes (SA2RAGE) uses two gradient echo trains which follow an $\alpha_{\text{nom}} = 90^\circ$ saturation pulse with different imaging FAs ($\beta_1 \approx 4^\circ$, $\beta_2 \approx 11^\circ$) and inversion times ($\text{TI}_1 \approx 56 \text{ ms}$, $\text{TI}_2 \approx 1800 \text{ ms}$) are performed in a steady-state manner [166]. SA2RAGE differs from MP2RAGE in that a saturation pulse replaces the inversion and the sequence parameters are chosen to maximise sensitivity to B_1^+ and minimise sensitivity to T_1 .




Figure 3.7: Simplified saturation prepared with 2 rapid gradient echoes (SA2RAGE) B_1^+ mapping sequence diagram. Square brackets indicate an imaging pulse train of length N_{PE} .

This allows fast 2D sampling, interleaved multi-slice, or 3D. The FA map is determined from a simulated dictionary using a lookup table (LUT), although a nominal T_1 must be assumed. This was later further adapted into NO2RAGE by omitting the preparation pulse and removing the delays [167]. This allows NO2RAGE to acquire higher-resolution maps in a shorter time.

Recently, a SA2RAGE whole-brain single-shim absolute B_1^+ map in under 30 seconds was demonstrated using a compressed sensing reconstruction, down from a 2-minute acquisition. Undersampling factors of 15 showed less than a 5% deviation from a GRAPPA 2 reference acquisition [168].

3.5.2.5 MP3RAGE

Quantitative T_1 mapping requires a measurement of B_1^+ to unbias T_1 measurements. Hence, simultaneous T_1 and B_1^+ maps are desirable for this purpose. Recently, an MP3RAGE [169] scheme was shown to obtain 1 mm isotropic T_1 and B_1^+ maps in under 7 minutes in the brain at 7T.




Figure 3.8: Simplified magnetisation prepared with 3 rapid gradient echoes (MP3RAGE) B_1^+ mapping sequence diagram. Square brackets indicate an imaging pulse train of length N_{PE} .

3.5.2.6 Saturation TurboFLASH

Saturation TurboFLASH was originally proposed by Klose et al. [170] which acquired around 20 images of increasing transmitter amplitude. This was further developed by Chung et al. [171] who showed that only two FLASH images were required. It is often referred to in the literature by several names (pre-pulse FLASH, pre-RF FLASH, pre-conditioning TurboFLASH or SatTFL) and is typically a 2D method. SatTFL also uses two TurboFLASH images [27], one unprepared (S_0) and one prepared using a large FA ($\alpha_{nom} \approx 90^\circ$) pulse (S_1). SatTFL is a very fast method compared to DAM and requires only a time of $5T_1$ between the two scans, hence approximately 20s for a single shim mode 2D slice.




Figure 3.9: Simplified saturation TurboFLASH (SatTFL) B_1^+ mapping sequence diagram. Square brackets indicate an imaging pulse train of length N_{PE} .

The FA maps are found from a simple arccosine relationship

$$\alpha = \cos^{-1} \left(\frac{S_1}{S_0} \right). \quad (3.7)$$

The benefits of this method are its inherent insensitivity to motion from the fast acquisitions. This method has become particularly popular at UHF, largely due to the relatively low SAR. Unfortunately, despite centric k -space reordering, SatTFL does have some sensitivity to short T_1 values due to the readout train. Using a variable FA train has been shown to further reduce the effect of T_1 blurring [172]. This method typically has a measurable range of 0-90°, however, by incorporating phase information the theoretical measurement range can be extended to 0-180°.

Equation 3.7 assumes no T_1 relaxation between the preparation pulse and the readout train. However, neglecting T_1 relaxation does not hold for the entire readout train. Blurring caused by long image trains makes it challenging to perform automatic masking for pTx pulse calculation. To reduce the effect of T_1 throughout the readout train a centric reordering of k -space is usually performed. The centre of k -space is acquired first after the preparation RF pulse and for the outer parts of k -space T_1 relaxation leads to an increasingly unreliable estimate of the FA.

Recently, work has been done to overcome the large underestimation for high B_1^+ values by performing a linear regression against more accurate, but slower, AFI maps. However, this produced a correction which was spatially dependent [173]. More recently, a linear correction using deep learning has been demonstrated offering a more reliable and less spatially dependent correction [174].

The total scan time to acquire a non-selective volumetric multichannel map with sequential slices would be $2 \times 5T_1 \times N_{Tx} \times N_{\text{slice}}$. Hence, 24 slices for an 8-channel transmit system would take on the order of 1 hour. However, recent research has shown 3D single-shot versions of SatTFL using a spiral-centric readout with variable FAs [172, 175]. More commonly, however, instead of using a non-selective saturation a selective preparation pulse is used and the dead time between images is better used by interleaving slices. Off-resonance effects are typically ignored by ensuring the pre-saturation pulse covers a broader slice ($2\times$ thickness) than the imaging excitations. The broadening of the pre-saturation pulse results in an increased bandwidth of the RF pulse, which typically limits the peak transmit power available, leading to either re-scaled pulses or them being clipped. This also means that interleaving slices would result in slice profile effects as adjacent slices will have previously experienced another saturation pulse thus altering the magnetisation history of that slice, but this can be largely avoided by acquiring an odd-then-even slice order and allowing full relaxation between neighbouring slices.

Another trick to reduce the acquisition time is using another saturation pulse to destroy the longitudinal magnetisation to operate in saturation recovery. This allows a reduction in the delay from $5T_1$ to approximately T_1 at a cost to the SNR, but it is ultimately impractical to achieve perfect saturation.

3.5.2.7 Fingerprinting

MR fingerprinting is a relatively new technique which can obtain many tissue properties from a single sequence, such as T_1 , T_2 etc. by dictionary matching to simulations [176]. MR fingerprinting has been shown to acquire high-resolution B_1^+ maps [177]. This may be useful to derive a tissue's electrical properties, including conductivity and permittivity, also known as electrical properties tomography [178]. MRF has been successfully demonstrated for 2D channel-wise absolute B_1^+ mapping in the abdomen at 7 T in under one minute [179–181].

3.5.2.8 Phase Sensitive

The Phase Sensitive method is a 3D technique that only considers the phase of the image, not the signal amplitude. The method works by first applying a non-selective pulse of FA 2α about the x axis which is immediately followed by a pulse of FA α about the y axis [182]. As this method only relies on the phase, it should be insensitive to T_1 since the phase does not change with T_1 . However, a pre-saturation pulse is usually applied to enable shorter TRs. This is because for some FAs the longitudinal magnetisation will be negative which will invert the phase for the following excitation. To remove phase effects from other sources, such as B_0 , a second acquisition is performed where the 2α pulse is instead applied along the $-x$ axis. Hence, this method does not suffer from B_0 dependence but does have a quadratic phase sensitivity and requires phase unwrapping for FAs in the range $0-45^\circ$ and $135-180^\circ$ [182]. A similar method also uses orthogonal RF pulses but instead with equal amplitude [183]. Imperfections in the RF spoiling could cause systematic errors for these orthogonal RF pulse phase sensitive methods [184].




Figure 3.10: Simplified Phase Sensitive B_1^+ mapping sequence diagram.

3.5.2.9 Bloch-Siegert Shift

Bloch-Siegert Shift (BSS) B_1^+ mapping is a novel phase-based sequence that uses an off-resonance Fermi RF pulse in the kilohertz range to shift the frequency of spin precession proportional to B_1^+ squared [185, 186]. The nutation of the spin can be neglected. The BSS sequence can be acquired using either a GRE or SE readout; Figure 3.11 shows the BSS for a GRE sequence.




Figure 3.11: Simplified Bloch-Siegert shift (BSS) B_1^+ mapping GRE sequence diagram.

The phase difference, $\Delta\Phi_{BS}$, between two images acquired with opposite offset frequencies, $\Delta\omega$, is given by

$$\Delta\Phi_{BS} = 2B_{1,\text{peak}}^2 \int_0^T \frac{(\gamma B_{1,\text{norm}}(t))^2}{2\Delta\omega} dt, \quad (3.8)$$

where $B_{1,\text{norm}}(t)$ is the normalised shape of the Bloch-Siegert pulse and $B_{1,\text{peak}}$ is the peak magnitude of $B_1(t)$.

Khalighi et al. [187, 188] demonstrated BSS at 7 T by acquiring a 2D 32×32 B_1^+ map in the brain in only 16 s for BSS whereas this took DAM 330 s. BSS has also been shown to work well with interferometric encoding approaches at 7 T [189] and has been employed at 11.7 T in a preclinical scanner [190].

BSS is relatively robust to different T_1 values, chemical shift, B_0 , and has a good dynamic range. However, the main limitation of BSS is the high SAR at UHF. To overcome this high SAR, BSS can be paired with a STEAM readout (BSS-STEAM) to acquire both images with only one BSS pulse [191]. Like SatTFL, this requires the use of a long time between images of $5T_1$. Bearing in mind the general trend at higher B_0 is an increase in T_1 , the sequences limited by a TR of $5T_1$ become increasingly unfeasible for clinical UHF use.

3.5.2.10 BEAR

B_1^+ estimation using adiabatic refocusing (BEAR) [192] is another phase-based method. A SE method which uses two adiabatic full passage pulses. The concept of the sequence is that the phase of an echo will depend linearly on B_1^+ when the relative magnitude of two adiabatic full passage refocusing pulses is altered. In this method, a lookup table is generated to relate the phase to the B_1^+ .




Figure 3.12: Simplified B_1^+ estimation using adiabatic refocusing (BEAR) B_1^+ mapping sequence diagram.

To my knowledge, this method has not been demonstrated at 7 T. This is likely due to the high SAR and since BEAR must reach the adiabatic condition, this makes pTx B_1^+ mapping difficult since the large dynamic range at 7 T will make this highly unlikely to be achieved everywhere in the FOV unless specific encoding schemes are used (see Section 3.5.4).

3.5.3 Off-resonance Effects

The FA not only depends on B_1^+ but also on the local B_0 and so α is only proportional to B_1^+ for spins that are on-resonance. This can cause erroneous measurements if B_0 is not accounted for [193, 194]. Three ways to overcome off-resonance effects are:

- Slice-selective pulses - If the α -encoding RF pulse is separate from the slice-selection pulse, then a broader slice for the α -encoding pulse can be applied and so it can usually be safely assumed that the FA within the imaging slice is fairly constant. This does, however, have slice profile implications for neighbouring slices. Additionally, for large FAs $>160^\circ$, this is not sufficient and a correction for the slice profile needs to be performed from Bloch simulations [195].
- Broadband RF Pulses - Off-resonance effects can be neglected if an RF pulse is applied with a much larger bandwidth than the sample that is excited.
- Model $\alpha(\omega, B_1^+)$ - Using the Bloch equations, the frequency response for a given RF pulse can be numerically simulated. Therefore, if a B_0 map is also acquired, the measured FA can be corrected.

3.5.4 B_1^+ Mapping Acquisition and Encoding Strategies

As I briefly alluded to, one method to overcome the issue of a large dynamic range is to acquire B_1^+ maps at multiple voltages [196]. Thereby, ensuring every signal-contributing voxel is mapped at least once within the sensitive range of the chosen method, at the expense of running additional scans.

If the α -encoding is performed separately from the image encoding (not always the case, i.e. it is for SatTFL but not for AFI or DAM), then all channels can be excited for imaging, with each channel in turn transmitting the α -encoding pulse [197]. This means a single reference image can be used for all transmit channels, therefore reducing the number of image acquisitions by a factor of $(N_{Tx} + 1)/(2N_{Tx})$. Another advantage is the increased sensitivity to low B_1^+ . The problem with this approach is it requires a homogeneous baseline field from which to make these maps, which is not always the case. Also, another dataset for the relative phases must be acquired but this can be done quickly.

Another method of overcoming the dynamic range associated with only a single transmit channel is to use a linear combination of the channels [198, 199]. Instead of mapping a single array element at a time, combinations of elements are mapped, and then these maps can be unencoded back to obtain channel-wise maps since it is known what encoding was used. Several approaches have been published on how best to produce these linear combinations, from Fourier/Butler to all-but-one [200]. Examples of these encodings are shown in Figure 3.13.

Termed "interferometric mapping", a baseline field distribution is measured and the interference of each element with this baseline constitutes a "mode", m [199]. This enables the calculation of relative maps for each transmit channel element, n_c . To ensure the linear system is not under-determined the total number of modes, M , must be greater than or equal to the total number of transmit channel elements, N_C . When $M > N_C$ the system is over-determined and the additional encoded measurements can be used to help suppress erroneous measurements. As I mentioned, the simplest option is to only excite a single channel at a time, in this case the encoding matrix, $\mathbf{E}$, is of size $[M \times N_C]$ equal to the identity matrix and as a result produces very large dynamic ranges but there is no need to invert this matrix. Another way that produces better homogeneity is starting with the standard CP^+ mode but with one channel having its phase inverted per mode, the elements of $\mathbf{E}$ would be defined as

$$(\mathbf{E})_{m,n_c} = \begin{cases} -\exp\left(\frac{2\pi i(n_c-1)}{N_C}\right), & m = n_c; \\ \exp\left(\frac{2\pi i(n_c-1)}{N_C}\right), & m \neq n_c. \end{cases} \quad (3.9)$$

Another option is to turn off one channel per mode, with the encoding matrix

$$(\mathbf{E})_{m,n_c} = \begin{cases} 0, & m = n_c; \\ \exp\left(\frac{2\pi i(n_c-1)}{N_C}\right), & m \neq n_c. \end{cases} \quad (3.10)$$

In Fourier/Butler encoding [201] each transmit channel undergoes a regular phase step, intending to move any regions of low B_1^+ around the subject. This has been shown to work well as it is highly likely that regions of low FA are not in fixed locations but are moved per configuration state, ensuring at least one reliable measurement of B_1^+ everywhere. This has the benefit that no calibration steps need to be performed before the acquisition, as is the case with some linear encodings where some prior knowledge of the transmit field is necessary. The Fourier encoding matrix is defined

$$(\mathbf{E})_{m,n_c} = \exp\left(\frac{2\pi i(m-1)n_c}{M}\right). \quad (3.11)$$

The individual channel maps are returned by the product of the pseudoinverse of the encoding matrix and the encoded measurements, on a pixel-by-pixel basis, as described by

$$B_{1,n_c}^+(\vec{r}) = \sum_{m=1}^M \left((\mathbf{E}^\dagger \mathbf{W} \mathbf{E})^{-1} \mathbf{E}^\dagger \mathbf{W} \right)_{n_c, m} B_{1,m}^+(\vec{r}) \quad (3.12)$$

where a weighting matrix, $\mathbf{W}$, can be useful when the measurements are over-determined when $M > C$, for suppressing erroneous measurements such as those measured from a low FA. This weighting matrix may be defined, for example from a soft thresholding function,

$$\mathbf{W}_{m,m'} = \begin{cases} (1 + e^{-q(J_m - p)})^{-1}, & \text{if } m' = m; \\ 0, & \text{if } m' \neq m, \end{cases} \quad (3.13)$$

where $J_m = I_{1,m}/\max(I_1)$ with $I_{1,m}$ representing the signal of the limiting image train for a particular mode and p, q are empirically determined constants. If no weighting is required or the system is not over-determined then $\mathbf{W}$ is simply the identity matrix, $\mathbf{I}$. For the case of Fourier encoding when $\mathbf{W} = \mathbf{I}$, this is equivalent to performing the inverse Fourier transform across the modes. This method has been demonstrated to acquire whole-liver B_1^+ maps in combination with a "PE-DREAM" sequence, using 13 encodings each across one breath hold [202].




Figure 3.13: The magnitude of the B_1^+ for some example methods of B_1^+ map encoding using simulated body data with 8 transmit channels and 8 transmit modes.

3.5.4.1 B1TIAMO

Following from TIAMO [106] (see Section 3.4.3), Brunheim et al. [107, 203] proposed a new method for acquiring B_1^+ maps using a coil-cycled GRE sequence for relative mapping following by SatTFL absolute maps for two shim modes (CP⁺ and CP²⁺ [90° phase offset between channels]). They demonstrated the feasibility of B1TIAMO by acquiring single-channel B_1^+ maps for a 32 Tx channel system over a large FOV 2D slice through the body in only 16 s. The B1TIAMO sequence is shown in Figure 3.14.




Figure 3.14: Illustration of sequence diagram for B1TIAMO[107]. Fast coil-cycled maps (green) are acquired first. This is followed by SatTFL absolute mapping (red) using a single reference image and two prepared images for two shim modes all separated by a time delay (TD) to allow for longitudinal relaxation between images. Note that the shim mode applies to the rectangular pre-saturation pulse only and not the image train. The square brackets indicate repeated sequence blocks.

The FA maps for each shim mode, $FA_{\text{shim1}}(r)$ and $FA_{\text{shim2}}(r)$, are calculated as standard for SatTFL and divided by the relative maps combined into each shim mode, $Rel_{\text{shim1}}(r)$ and $Rel_{\text{shim2}}(r)$, giving shim mode specific scaling factors $s_1(r)$ and $s_2(r)$,

$$s_1(r) = \frac{FA_{\text{shim1}}(r)}{Rel_{\text{shim1}}(r)}, \quad s_2(r) = \frac{FA_{\text{shim2}}(r)}{Rel_{\text{shim2}}(r)}. \quad (3.14)$$

Then a total scaling factor, $S(r)$, is calculated from a combination of the scaling factors from each shim mode weighted by an exponent m which is determined experimentally.

$$S(r) = \frac{s_1(r) \cdot Ref_{\text{shim1}}(r)^m + s_2(r) \cdot Ref_{\text{shim2}}(r)^m}{Ref_{\text{shim1}}(r)^m + Ref_{\text{shim2}}(r)^m} \quad (3.15)$$

The channel-wise B_1^+ is then calculated by scaling the individual channel-wise relative maps by the total scaling factor,

$$B_{1,n_c,abs}^+(r) = S(r) \cdot B_{1,n_c,rel}^+(r). \quad (3.16)$$

3.6 Summary

In this chapter, I have explored some of the fundamental issues surrounding imaging at UHF. One of these is B_1^+ inhomogeneity. Ultimately, knowledge of the B_1^+ field is required and although a plethora of B_1^+ mapping sequences are in existence, to my knowledge none are really viable for absolute B_1^+ mapping in the body at 7 T in a clinically reasonable time. Hence, this was the motivation for the following chapters.

4

Sandwiched SatTFL B_1^+ Mapping

Contents

4.1	Introduction	68
4.2	Theory	70
4.3	Methods	72
4.3.1	Numerical Simulations	72
4.3.1.1	Cardiac Numerical Simulations	73
4.3.1.2	Further Technical Development Simulations	74
4.3.2	Parallel Transmit MRI System	75
4.3.3	Sequence Implementation	76
4.3.3.1	2D Protocols	76
4.3.3.2	3D Protocols	77
4.3.4	Image Reconstruction	77
4.3.5	Phantom Experiments	77
4.3.6	<i>In Vivo</i> Experiments	78
4.4	Results	79
4.4.1	Simulation Results	79
4.4.1.1	Cardiac Simulation Results	84
4.4.1.2	Further Tech. Dev. Simulation Results	87
4.4.2	Phantom Results	88
4.4.3	<i>In Vivo</i> Results	94
4.5	Discussion	100
4.6	Summary	103

4.1 Introduction

UHF MRI promises superior SNR over conventional clinical MRI systems [204]. However, inhomogeneity in the B_1^+ transmit field, due to the wavelength associated

with the increased Larmor frequency, produces altered image contrast and significant signal dropout. Various approaches of overcoming this B_1^+ inhomogeneity have been proposed in the literature. These include, but are not limited to, the use of adiabatic pulses [34], pTx [98, 99] methods such as kT -points [116], 2D Spokes [117], and TIAMO [106, 107].

Parallel transmit enables spatial and temporal RF-shimming overcoming local B_1^+ inhomogeneity, but prior knowledge of the channel-wise B_1^+ is required for the optimal RF pulse design. B_1^+ mapping at UHF is more challenging than at lower fields due to the increased dynamic range across a FOV and the limited range of FAs over which methods are accurate [135]. Additionally, the acquisition duration increases with the number of transmit channels, necessitating faster mapping schemes. *in vivo* imaging requires fast and robust B_1^+ calibration, however, current methods suffer from limitations of the SAR and inaccuracy due to T_1 relaxation times.

There is currently no gold standard for *in vivo* B_1^+ mapping, particularly for the heart, where this is compounded by cardio-respiratory motion. AFI [142, 145] is accurate as well as being insensitive to T_1 but requires efficient spoiling to eliminate unwanted signal from residual transverse magnetisation. AFI employs a dual-TR steady-state gradient-refocused echo acquisition and consequently requires many dummy TRs before imaging to ensure operation in a steady state. Bloch-Siegert shift [186] is phase-based and has been shown to obtain motion-robust cardiac B_1^+ maps at 3 T [205] but suffers from high SAR at UHF and is affected by susceptibility. DREAM [152] offers an improvement in speed and robustness to motion by acquiring two images quasi-simultaneously using a FID and stimulated echo combination [154]. DREAM is fast, has low SAR and in practice has a dynamic range of 7 (10 - 70°). Although, whilst it performs well even in the myocardium DREAM suffers from dropout in the blood pool of the heart due to flow sensitivity [153, 159, 160]. SA2RAGE [166] has a large dynamic range reported of 10 but requires long delay times between two TurboFLASH readouts following a saturation pulse making it unsuitable for cardiac imaging.

B_1^+ mapping using SatTFL [170, 171] is extremely SAR efficient [136], flow insensitive and intrinsically robust to motion. However, it requires long TRs to allow for full T_1 relaxation between a reference and magnetisation-prepared acquisition. Moreover, SatTFL suffers from T_1 sensitivity due to relaxation during the TurboFLASH readout train. Chung et al. [171] proposed reducing the time

delay (TD) of SatTFL down from $5T_1$ to approximately T_1 , at the cost of reduced SNR, by implementing a ‘saturation recovery module’ using adiabatic pulses to produce a consistent magnetisation state between acquisitions. However, this requires uniform saturation across the whole subject, which is difficult to achieve in practice and highly SAR intensive at 7 T. Additionally, since B_1^+ errors can compound and cause B_0 maps to be inaccurate (and vice versa), a method of B_1^+ mapping which is B_0 insensitive would be ideal. This could therefore make B_0 shimming more robust if performed after B_1^+ shimming.

The aim of this work was to design and evaluate a novel adaptation to SatTFL to mitigate the T_1 sensitivity and shorten the acquisition time (TA), while keeping all the beneficial characteristics typical of B_1^+ mapping using SatTFL. To achieve this, the reference and prepared images are ‘sandwiched’ together, thus removing the $5T_1$ time delay completely and acquiring these images back-to-back [2]. This adaptation was evaluated in a segmented approach to k -space acquisition, multi-transmit FA mapping and as a 3D sequence. This Sandwich method is compared to a SatTFL and short TR SatTFL method in simulation, phantoms and *in vivo* at 7 T for the brain and heart.

4.2 Theory

The SatTFL B_1^+ mapping sequence [170] measures a series of rapid FLASH images, each with a different magnitude preparation pulse. Chung et al. [171] demonstrated that in practice only two FLASH images are needed; one reference (S_0) and one magnetisation-prepared (S_1) with a nominal preparation FA of α (typically of order 60-120°), as depicted in Figure 4.1a. The FA of the preparation pulse is determined as in Eq. 3.7.




Figure 4.1: Simplified pulse sequence diagrams of SatTFL and the adapted schemes using a preparation pulse FA, α , and image train (IT) FAs, β , for a single repetition time (TR) along with three example T_1 recovery curves (short T_1 blue, mid T_1 yellow, long T_1 green) not to scale. **a)** SatTFL with a long relaxation time between a reference and prepared IT (with N phase encodes) allowing for full T_1 relaxation. **b)** Short TR: IT acquired segmented in 4 (with N' phase encodes; reference images (S_0) first followed by prepared images (S_1)) with time delay (TD) of approximately 1 s and TR of 2 s. **c)** Sandwiching: IT segmented into four shots and TD minimised, images acquired pairwise. Square brackets indicate a TurboFLASH repetition period (TR_{FLASH}) for phase encoding. Each phase-encoding step and preparation RF pulse is followed by gradient spoiling in addition to RF spoiling. Note **b)** and **c)** start from incomplete recovery of the longitudinal magnetisation due to prior TRs.

Chung et al. [171] recommended having a TD between S_0 and S_1 to enable full relaxation or using a "reset pulse" to image faster as reducing this delay introduces a T_1 dependence as shown in Figure 4.1b. In this work, an alternative approach to enable a much faster SatTFL acquisition was evaluated. The time delay was set to $TD = 0$ s such that the total duration of S_0 and S_1 is much shorter than T_1 . This results in the T_1 contrast in S_0 being matched to S_1 as depicted in Figure 4.1c. The matched T_1 contrast removes the sensitivity to T_1 and enables the use of Eq. 3.7 to calculate the FA. Here, this timing is referred to as the Sandwich scheme. In these experiments, it was found that the image excitation (β) should

be subjected to the same B_1^+ field as the preparation pulse α , and in the Sandwich scheme experiments below they are linearly proportional.

Due to the large B_0 inhomogeneities in a 3D volume at 7 T, a 5 ms non-selective broad-band full-passage HS8 pulse was used with a time-bandwidth product (TBP) of 15, operating below the adiabatic limit, as the preparation pulse [206] in the 3D schemes (Figure 4.2). The HS8 pulse is insensitive to B_0 over a broad frequency range determined by the TBP. Additionally, the HS8 pulse offers a low peak RF voltage making it suited to regions of low B_1^+ and high B_0 inhomogeneity. Correction for the non-linearity of the HS8 pulse response to B_1^+ must be performed to maintain accuracy at high B_1^+ .




Figure 4.2: Bloch simulation of **a)** sinc pulse (5 ms) and **b)** HS8 (5 ms) pulses used for saturation of the magnetisation-prepared image. HS8 pulse is highly insensitive to changes in B_0 as well as 20% of the peak B_1^+ compared to the equivalent time-bandwidth product sinc pulse.

4.3 Methods

4.3.1 Numerical Simulations

All sequences were simulated using the extended phase graph formalism [42] in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States) to explore the effects of α , β , repetition time between segments (i.e. TR), T_1 and the length of the TurboFLASH imaging train (N') on the measured B_1^+ for each method. The effect of noise was investigated using Monte Carlo simulation. Nominal FAs α_{nom} and β_{nom} were varied from 1-180° in 1° increments and 1-20° in 0.1° increments, respectively. Similarly, the effect of T_1 relaxation on the calculated value of α was investigated by simulating a range of T_1 times 0.5 s to 3 s in increments of 50 ms.

Simulations were based on Brian Hargreaves EPG code <http://web.stanford.edu/~bah/software/epg>. Each complete simulated image train consists of 36 phase-encoding steps, image train TE = 1.78 ms, TR_{FLASH} = 4.2 ms, low-angle, β_{nom} , simulating k -space readouts with spoilers throughout the image train to eliminate transverse magnetisation before each imaging RF pulse. The transverse magnetisation immediately following the preparation pulse was effectively spoiled using a quadratically increasing number of ‘unit’ gradients. RF spoiling was also implemented using a linearly increasing phase increment of $\Delta\Phi=50^\circ$. The pulse schemes were simulated as described in the theory section with two additional dummy TRs to help reach a steady-state magnetisation for the segmented schemes. The T₂ for all simulations was kept fixed at 25 ms.

The image trains are reordered to mimic a centrically acquired k -space, zero-filled to produce an array of length 64 and then inverse fast Fourier transformed to obtain the point-spread function (PSF). The relative effects of noise across the different schemes were evaluated using a Monte Carlo simulation with 1000 repetitions, with a fixed level of zero-mean complex Gaussian noise added independently to each image train. Following this, α is calculated using the centre of the PSF and Eq. 3.7. Figure generating code can be found at: <https://github.com/jameslewiskent/sat2TFL>.

4.3.1.1 Cardiac Numerical Simulations

There are additional considerations to be made to account for imaging the heart.

Heart Rate Variability A normal sinus rhythm for a healthy heart has a heart rate variability (the change in the time intervals between consecutive heartbeats; HRV) of approximately 50 ms. This means that when a sequence is cardiac triggered, slowly varying changes in heart rate lead to a TR which is not precisely consistent. This is especially true for subjects with arrhythmia where the heart can beat irregularly (HRV > 120 ms). Additionally, false and missed triggers have the potential to extend individual TRs. It is therefore important that a sequence that is to be cardiac gated is largely unaffected by a variable TR. The effect of a variable TR will manifest as signal variations in the second phase encode direction but should have little impact in the first phase encoding direction. To test the potential impact this may have on the Sandwich sequence, variable TRs were simulated with a mean of 1 s and SD of 0.1 s. A TR which fell below a minimum TR of 700 ms was treated as a missed trigger. This effectively means skipping a heartbeat which results in an effective TR of 2 heartbeats. An assumed TR of 1 s was used for FA calculation.

Ejection Fraction During the cardiac cycle, a portion of the blood (typically $>50\%$ in a healthy heart) within the chambers of the heart is ejected and replenished by fresh blood, known as the ejection fraction (EF). This fresh blood will not have seen any previous RF pulses and hence will have no magnetisation history. Simulations were also performed to test the effect of this (including the effects of HRV) on the sequence. To do this, a two compartment pool was simulated. In compartment one, magnetisation was unaffected and in a second, the longitudinal magnetisation was reset back to equilibrium each TR. To simulate a range of EFs a mixture of these two compartments were taken (0-100% of compartment two in 1% increments).

4.3.1.2 Further Technical Development Simulations

Extensive simulations were also performed to best optimise the design of the Sandwich sequence. This section describes some of the simulations performed to accomplish this.

Dummy RF Pulses Dummy RF pulses (approximately 1000) are sometimes used in conventional SatTFL in order to establish a steady state prior to acquiring the reference image. This may be beneficial for the Sandwich sequence, especially to reduce the impact of signal instability causing image artefacts due to large ($\alpha > 180^\circ$) FAs close to the transmitter - particularly troublesome for body imaging. Due to the interleaved reference and prepared images and the short TR it is not possible to perform 1000 dummy RF pulses prior to each reference image for the Sandwich sequence without drastically increasing the TR. Additionally, for body imaging there are limitations as to how many dummy RF pulses can be performed during the trigger delay (unless they are performed continuously). Therefore, to test whether this may be beneficial for the Sandwich sequence, this was tested using 25 dummy RF pulses (≈ 100 ms) prior to the reference image train.

Encoding Direction Conventionally, the SatTFL sequence is acquired with centric phase encoding to acquire the centre of k -space immediately following the preparation pulse. However, whilst this is sensible for SatTFL, the encoding scheme for the Sandwich sequence should acquire the centre of k -space just prior to the saturation pulse for the reference image. The proposed mirrored centric encoding scheme is illustrated in Figure 4.3. The impact of using both centric and mirrored centric encoding was evaluated.




Figure 4.3: Diagram illustrating **a)** centric versus **b)** mirrored centric phase encoding for the Sandwich sequence reference and prepared images, S_0 and S_1 , respectively.

TR Time The Sandwich TR value of 1 s was chosen to coincide with a sensible resting heart rate of 60 bpm. However, in the brain this is not necessary and a more optimal and potentially shorter value of TR for a given dynamic range (DR) is possible. To identify such a value a range of TR values were simulated from 0.5 s to 2 s in 50 ms intervals. For a fair evaluation, a LUT generated with the same TR was used for FA calculation. In order to evaluate this, strict criteria for the DR was required, as although in theory a large range of FAs can be measured in practice only a much smaller region provides a reliable measurement. The DR was defined as the region of FAs measured, on resonance, where the mean FA and 1.96 SD (95% of observations) across all T_1 values were within 10% of the nominal simulated FA value.

T_1 Correction Scheme An additional modification to the proposed Sandwich scheme was tested in order to further reduce sensitivity to T_1 relaxation. This involved a third readout train, S_2 , which is performed directly after S_1 . Equation 3.7 must be adapted

$$\alpha = \cos^{-1} \left(\frac{S_1 + 0.5(S_1 - S_2)}{S_0} \right). \quad (4.1)$$

This modification will be referred to as the " T_1 Correction" scheme.

4.3.2 Parallel Transmit MRI System

A 7 Tesla MRI (Magnetom VB17 Step 2.3, Siemens Healthineers, Erlangen, Germany) equipped with parallel transmit was used for phantom and *in vivo* imaging. Phantom and brain images were acquired using the same 8-channel transmit, 32-channel receive head coil (Nova Medical, Wilmington MA, USA) in CP⁺ mode at magnet isocentre. To image the heart an 8-channel transmit/receive dipole array (MR Coils BV, Zaltbommel, Netherlands) was used with cardiac triggering using a pulse oximeter.

4.3.3 Sequence Implementation

Both 2D and 3D Sandwich pulse sequences were written in C++ and compiled for use on the scanner inside a VB17 IDEA environment on a Windows XP VirtualBox machine.

4.3.3.1 2D Protocols

On the scanner, 2D acquisitions used an $\alpha : \beta$ ratio of 10:1, based on simulations, with nominal α/β FAs 90/9°. 2D B_1^+ maps were measured at four different reference voltages of 30 V, 60 V, 90 V and 120 V achieving α FAs of approximately 40°, 80°, 120° and 160° in the centre of the brain to assess the linearity and the dynamic range. k -space was segmented into 4 shots (readout trains) for all but the conventional SatTFL method each with $N' = 9$ phase encoding lines.

The TD between consecutive images and the TR to acquire the pair of S_0 and S_1 images per segment (TD/TR) was 9.85 s/20 s for SatTFL and 0.96 s/2 s for short TR SatTFL, with $TR_{\text{fill}} = TD$. The Sandwich method had a TD/TR of 0 s/1 s, with a TR_{fill} of 0.92 s. TurboFLASH TE = 1.78 ms and TurboFLASH repetition time (TR_{FLASH}) 4.2 ms, FOV = $250 \times 250 \text{ mm}^2$, slice thickness = 8 mm, scan matrix = 50×36 (including asymmetric echo, 75% PF and 75% phase resolution), bandwidth = 490 Hz px⁻¹, centric phase-encoding, and RF spoiling. B_1^+ information is encoded using a 0.5 ms non-selective rectangular pulse, α , and imaging with a sinc pulse, β . The short TR SatTFL sequence acquires the 4 segments of S_0 followed by the 4 segments of S_1 , unlike the Sandwich scheme where these segmented images are acquired pairwise.

Additionally, two dummy TR periods were implemented for the Sandwich scheme prior to data acquisition ensuring consistent magnetisation. A reference B_1^+ map was acquired using a segmented (4 segments with a readout duration of 40 ms) SatTFL map with TD/TR = 9.96 s/20 s. The shorter imaging train used is expected to minimise the T_1 sensitivity of the SatTFL method. The reference map was acquired at two voltage levels (60 V and 90 V) to obtain a best estimate of the ground truth and a mean reference map was calculated for comparison to the other methods. The TA was 80 s for the reference scheme, 20 s for SatTFL, 8 s for short TR SatTFL and 6 s (including two dummy TR periods) for the Sandwich scheme.

4.3.3.2 3D Protocols

3D B_1^+ maps were also acquired with k -space segmented into 18 (brain) or 24 (heart) shots of $N' = 36$ phase encoding lines per shot, this is one partition per TR. The scan matrix was $50 \times 36 \times 18/24$ (brain/heart). Non-selective excitation with a TurboFLASH TE = 0.87 ms and TR_{FLASH} = 2.7 ms was used and a HS8 pulse as the preparation pulse.

Based on simulations, it was determined that for the longer TurboFLASH image train of the 3D sequences, a higher $\alpha : \beta$ ratio of 20:1, using nominal α/β FAs 130/6.5°, was preferable and improves the linear dynamic range. Non-selective excitation used a larger spoiler as spoiling due to slice selection is not present and helps to suppress artefacts from undesirable coherence pathways (See Appendix A).

4.3.4 Image Reconstruction

Image reconstruction was performed offline in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States) using the raw TWIX data format. Binary files (.dat) were read-in using an open-source parsing function (*mapVBVD.m*) [207]. The k -space data were zero-padded to 64×64 (2D) or $64 \times 64 \times 18/24$ (3D brain/heart) and a Hanning filter was applied to reduce Gibbs ringing.

Separate receive channel images were then formed from an inverse fast Fourier transform and combined using coil sensitivity profiles calculated from ESPIRiT [51] to produce complex images S_0 and S_1 . The resulting maps were then calculated using Eq. 3.7. The B_1^+ maps in phantom and *in vivo* were masked using a simple value-thresholding approach informed from the reference image and slightly eroded; except in the 3D cases where masks were manually drawn.

4.3.5 Phantom Experiments

Two spherical phantoms containing silicon oil/water were used to evaluate the sequences. 2D single-shim as well as 2D individual pTx and 3D B_1^+ maps were obtained for each method.

Additionally, the "T₁ Correction" scheme using an additional third FLASH train, as described in Section 4.3.1.2, was tested in a water phantom. The results of this are shown in Figure A.1.

The water phantom was a 16.5 cm diameter spherical MRS ‘Braino’ phantom (General Electric Medical Systems) containing 10 mmol creatine, 3 mmol choline, 5 mmol Lac, 1 mL L⁻¹ Gd-DPTA (Magnevist), 12.5 mmol Glu, 7.5 mmol myo-Ins, 12.5 mmol NAA, 0.1% sodium azide, 56 mmol sodium hydroxide and 50 mmol potassium phosphate monobasic.

The silicon oil phantom was a 16.5 cm diameter sphere containing silicon-based organic polymer (Polydimethylsiloxane, T₁ = 750 ms, T₂ = 750 ms) which is non-polar and hence has a more uniform excitation profile.

4.3.6 *In Vivo* Experiments

Six healthy volunteers were scanned in accordance with Oxford University’s ethical practices. 2D B₁⁺ maps were measured in the head using each of the methods for four different reference voltages. Voxel-wise correlation and Bland-Altman plots compare each method to a reference map.

Additionally, the "T₁ Correction" scheme using an additional third FLASH train, as described in Section 4.3.1.2, was also tested *in vivo*. The results of this are shown in Figure A.2.

A proposed benefit of the Sandwich method is that the paired acquisition of a reference and prepared image provides insensitivity to magnetisation history, as each S₁ image has a matching S₀ image. This facilitates rapid multi-transmit mapping since you can switch to a different B₁⁺ shim without waiting for full relaxation of the longitudinal magnetisation. To test this, the Sandwich scheme was evaluated in one volunteer for rapid 2D single-shot (non-segmented, readout duration of 151 ms) B₁⁺ mapping of individual transmit channels in the brain and compared to SatTFL and short TR SatTFL. Due to the longer readout train, a higher $\alpha : \beta$ ratio was used of 20:1, using nominal α/β FAs 130/6.5°.

The SatTFL schemes acquired one S₀ image and 8 S₁ images transmitting with β in CP⁺ mode and α one channel at a time. The TA was 90 s for SatTFL and 9 s for short TR SatTFL scheme. For the Sandwich scheme, to ensure a constant $\alpha : \beta$ ratio (within a TR period), the same B₁⁺ shim was used for both α and β . RF pulses were transmitted one channel at a time for a total TA of 8 s, no dummy pulses were performed since the B₁⁺ shim changes between TR periods. In addition to these absolute maps, an interleaved coil-cycled gradient echo sequence was performed to obtain relative maps and used to rephase and weight the total B₁⁺.

To evaluate the single channel maps, they were combined using the phase of relative transmit maps, to form a synthetic CP^+ mode map and compared the result to a map measured in CP^+ mode. The acquisition parameters for the interleaved coil-cycled gradient echo sequence to obtain these relative maps were: matched FOV and 64×64 scan matrix, FA 7° , TE = 1.9 ms, TR_{GRE} = 4.3 ms for a total TA of 2 s.

In another healthy volunteer, the Sandwich scheme was evaluated as a 3D B_1^+ mapping acquisition in the brain and compared to a 3D sequential SatTFL and short TR SatTFL. A scan matrix of $50 \times 36 \times 18$ was used with a FOV of $250 \times 250 \times 180 \text{ mm}^3$, 10 mm partition thickness, with a reference voltage of 60 V. For the SatTFL methods the TD/TR was 9.9 s/20 s for SatTFL and 0.9 s/2 s for short TR SatTFL, with TR_{fill} = TD. The Sandwich method had a TD/TR of 0 s/1 s, with a TR_{fill} of 0.8 s. The total TA for SatTFL was 360 s, the short TR SatTFL scheme took 36 s and the Sandwich scheme took 20 s (including two dummy TRs).

Lastly, to evaluate the Sandwich scheme for whole-heart, single breath-hold, B_1^+ mapping a 3D Sandwich B_1^+ map was measured in the heart of a healthy volunteer. For reference, a 2D SatTFL B_1^+ map was acquired along the horizontal long axis of the heart. The acquisition parameters for the 3D Sandwich were a scan matrix of $50 \times 36 \times 24$ with a FOV of $384 \times 384 \times 240 \text{ mm}^3$, 10 mm partition thickness and a 200 V reference voltage. The total TA was 20 heartbeats for a single 2D SatTFL B_1^+ map (minimum TD = 10 s) and 26 heartbeats for the 3D Sandwich scheme, each acquired within a breath-hold. A slice through the 3D Sandwich data was oriented and resampled, using a linear interpolation in FSL [208], to match the 2D SatTFL for comparison.

4.4 Results

4.4.1 Simulation Results

The longitudinal magnetisation for the 2D segmented Sandwich sequence is shown in Figure 4.4. Figure 4.4 shows that after 2 TR periods (2 s) the longitudinal magnetisation is largely in a steady state. This is beneficial to reduce signal variations in the phase encode direction which may cause undesirable artefacts particularly when the sequence is segmented. Hence, 2 dummy TR periods were used going forward.




Figure 4.4: The longitudinal magnetisation for the 2D Sandwich sequence plotted for three nominal preparation pulse flip angles of 30°, 90° and 150°.

The full-width at half-maximum of the point spread function for the 3D Sandwich scheme is shown in Figure 4.5. This indicates intervoxel blurring in both images S_0 and S_1 and in both phase encode directions is low across various T_1 times and FAs, with an average full width at half maximum (FWHM) of the PSF of 1.14 voxels. The FWHM of the PSF peaks at 3 voxels in S_1 due to the signal null.

Figure 4.6 shows the mean bias for different simulated nominal α and β and the effect of T_1 on the measured value of α . Dark red signifies large positive overestimation of α and dark blue large underestimations. The SatTFL scheme shows little variation in bias for different ratios of $\alpha : \beta$ which is not the case for the other schemes. Figure 4.6b reports large under- and overestimations for short TR SatTFL across a range of preparation pulse FAs for all nominal $\alpha : \beta$ ratios, which is also T_1 -dependent. This error in α is due to there being insufficient time for the longitudinal magnetisation to fully relax to equilibrium. As a result, the longitudinal magnetisation is not consistent prior to imaging. Figure 4.6c shows the Sandwich scheme performs remarkably well using a TR of 1 s due to the segmentation and ‘sandwiching’ of the two image trains. The bottom row of Figure 4.6 shows that at a fixed $\alpha : \beta$ ratio of 10:1 there is some T_1 -dependent bias introduced which increases with the preparation FA for both SatTFL schemes. However, this bias due to T_1 is greatly reduced for the Sandwich scheme.




Figure 4.5: FWHM of the PSF for both the reference image (S_0) and prepared image (S_1) along both phase encode (PE) directions for the Sandwich sequence.

As expected, SatTFL exhibits relatively low bias in α and the least sensitivity to noise (Figure 4.7) due to the long TD of $5T_1$ between the reference and prepared images. Contrastingly, noise is more significant in the Sandwich scheme (Figure 4.7) due to the reduced longitudinal magnetisation in the steady state and depleted magnetisation following S_0 (resulting in a lower SNR in S_1).




Figure 4.6: The effects of a nominally applied preparation pulse, α_{nom} , and image train FA, β_{nom} , on the mean bias in the FA map (top row) and the effect of T_1 on the measured FA (bottom row) for the three different 2D B_1^+ mapping sequences. Simulations performed using a T_1 of 2 s (top row) and fixed $\alpha : \beta = 10:1$ (bottom row) with a fixed noise level added across all schemes using a Monte Carlo method with 1000 repeats, standard deviation shown in Figure 4.7. Four lines of constant preparation pulse to image train FA ($\alpha : \beta$) ratio are plotted and shown in the common legend in **a**). (Note: bias $> 20^\circ$ is colour saturated to give a greater dynamic colour range to bias in range of interest i.e. in range -20° to 20°).




Figure 4.7: The effects of a nominally applied preparation pulse, α_{nom} , and image train FA, β_{nom} , on the standard deviation of the bias in the FA map (top row) and the effect of T_1 on the measured FA (bottom row) for the three different B_1^+ mapping sequences. Simulations performed using a T_1 of 2 s (top row) and fixed $\alpha : \beta = 10:1$ (bottom row) with a fixed noise level added across all schemes using a Monte Carlo method with 1000 repeats. Four lines of constant preparation pulse to image train FA ($\alpha : \beta$) ratio are plotted and shown in the common legend in **a**).

It was therefore identified that using an $\alpha : \beta$ ratio of 10:1 may yield the best compromise between a large range of low-bias α and low sensitivity to noise for 2D measurements; thus, all 2D experiments (except for 2D multi-transmit mapping) use this ratio.

The effect of the readout train length and $\alpha : \beta$ ratio on the mean bias and standard deviation is shown in Figure 4.8. Higher $\alpha : \beta$ ratios are likely to be beneficial for 3D sequences where the length of the TurboFLASH imaging trains are longer.




Figure 4.8: Bias **a)** and standard deviation **b)** at different $\alpha : \beta$ ratios and number of phase encoding lines within a TurboFLASH readout (N') shown for a nominal α of 120° . The parameters which correspond to the 2D ($\alpha : \beta = 10:1$, $N' = 9$) and 3D ($\alpha : \beta = 20:1$, $N' = 36$) sequences used in this manuscript are marked by crosses.

4.4.1.1 Cardiac Simulation Results

Heart Rate Variability The impact of HRV on the Sandwich sequence was found to be negligible. There was no systemic error in the measured FA across all repeats, although the standard deviation of the measured FA did increase slightly. However, this was much less than 1° across all T_1 values of 0.5s to 3s. Figure 4.9 shows a histogram of the HRV which occurred in the simulations. The mean TR variability across all repeats was 144ms which would represent severe arrhythmia and worse than expected in reality.

Figure 4.10 demonstrates the impact of both EF (50%) and HRV (1.0 ± 0.1 s) on the longitudinal magnetisation plotted for three nominal FAs, 30° , 90° and 150° . The HRV leads to variable sequence duration.




Figure 4.9: A histogram showing the range of TRs which occurred when a variable TR was simulated (blue) and the resulting heart rate variability (HRV; orange).

Ejection Fraction Figure 4.11 shows the results from simulation of the effect of EF on the measured FA. Figure 4.11a) shows the relative error in FA across the dynamic range averaged across T_1 . Figure 4.11b) shows the relative error in FA across the range of T_1 values simulated, for the mean in range 30° - 140° (as indicated by the dashed lines in 4.11a)). On average, an EF tends to lead to an overestimation in the measured FA.

In healthy humans EF ranges between 50% to 70%. Assuming an EF of 50% and a T_1 of blood of approximately 2s [209], these results suggest that the B_1^+ within the blood chambers of the heart are likely to be overestimated by 4.2% on average, i.e. a 90° FA in the chambers of the heart is measured to be 94° . This error induced by the EF on the measured FA could be reduced by including an expected average EF in the simulation of the LUT.




Figure 4.10: The longitudinal magnetisation for the 3D Sandwich sequence when both a 50% ejection fraction and variable TR of 1.0 ± 0.1 s is used. Longitudinal magnetisation is plotted for three nominal preparation pulse FAs of 30° , 90° and 150° .




Figure 4.11: The effect of ejection fraction on the estimated flip angle. **a)** mean bias due to flip angle across T_1 values in the range 0.5 s to 3 s. **b)** mean bias due to T_1 for flip angles the range 30° - 140° (as indicated by the dashed lines in **a**)). FA; flip angle.

4.4.1.2 Further Technical Development Simulation Results

Dummy RF Pulses Performing dummy RF pulses prior to the reference image in the Sandwich sequence did not provide any additional benefit and, in fact, resulted in a minor reduction in DR with the added cost of additional SAR. This is expected as the number of dummy RF pulses is not enough to reach a steady state in the FLASH train. Hence, no additional dummy RF pulses were implemented on the scanner.

Encoding Direction The impact of using a mirrored centric phase encoding scheme over a centric encoding scheme resulted in a minor improvement in the DR of the Sandwich scheme from 4.8 to 5.0, where the DR is as previously defined in Section 4.3.1.2. Hence, a mirrored centric phase encoding scheme was implemented on the scanner.

TR Time The DR efficiency (DR/TR) versus TR is shown in Figure 4.12. A TR value of around 1 s appears to give adequate DR efficiency as indicated by the elbow of the scatter plot. Longer TR values lead to an increased DR whilst the reduction of the TR below 1 s provides a lower DR due to decreased SNR. The TR was kept constant at 1 s for the remaining experiments.




Figure 4.12: Scatter plot showing the dynamic range efficiency of the Sandwich method versus TR.

T₁ Correction Figure 4.13 shows the effect of the additional FLASH train for T₁ bias correction. The T₁ correction scheme was successful in reducing the variation in measured FA due to T₁, extending the upper bounds of the DR. However, this ultimately comes at the cost of an overall slightly reduced DR (5.0 to 4.9) due to decreased signal for low FAs. This modification also had an additional SAR penalty of 16% compared to the standard proposed Sandwich sequence. The T₁ correction scheme was further tested on the scanner in phantom and *in vivo* and the results can be found in Appendix A.1.




Figure 4.13: Simulation results showing the difference between the a) Sandwich and b) T₁ correction schemes.

4.4.2 Phantom Results

Figures 4.14 and 4.15 show results from a silicon oil and water phantom, respectively. The first two rows of each figure show the magnitude of the images used to calculate the FA maps in the third row. The difference between each of the maps at 60 V and a mean reference map, calculated from three reference voltages for the reference scheme and scaled by the voltage are shown in the fourth row, along with the root-mean-square error (RMSE). FA profiles for horizontal and vertical profiles through the centre of the phantom are also shown. Note that the reduced time to acquire Sandwich was 6 s compared to 80 s for the reference, 20 s for SatTFL and 8 s for a short TR SatTFL.




Figure 4.14: Flip angle maps from a silicon oil phantom at reference voltage 60 V for various sequences. **i)** segmented SatTFL method with TR = 20 s **ii)** SatTFL **iii)** SatTFL acquired segmented and TR reduced to 2 s **iv)** segmented Sandwich scheme with TR = 1 s, $\beta = 0.1\alpha$.

Figure 4.14 shows the 2D silicon oil phantom results. An average FA of approximately 80° was achieved throughout the slice. The FA maps show excellent agreement between the Sandwich scheme and both the reference and SatTFL, with a RMSE of just 0.35° . The short TR SatTFL shows a large overestimation in the measured FA with an RMSE of 2.6° .




Figure 4.15: Flip angle maps from a water phantom at reference voltage 60 V for various sequences. **i)** segmented SatTFL method with TR = 20 s **ii)** SatTFL **iii)** SatTFL acquired segmented and TR reduced to 2 s **iv)** segmented Sandwich scheme with TR = 1 s, $\beta = 0.1\alpha$.

Figure 4.15 shows the 2D water phantom results. Note that the dielectric properties of the water phantom give rise to much more inhomogeneous FA profiles than for the (non-polar) silicon oil phantom. The FA maps show excellent agreement between the Sandwich scheme and both the reference and SatTFL, with a RMSE of just 1.3°. The short TR SatTFL shows a large overestimation in the measured

FA with an RMSE of 9.7° and a reduced dynamic range.

The difference between an acquired CP^+ mode map to a pTx combined map is shown in Figures 4.16 and 4.17 for silicon oil and water phantoms, respectively. The Sandwich scheme appears to underestimate the FA towards the centre of the waterphantom, when compared to the acquired CP^+ map, whereas SatTFL overestimates. There is good agreement between the acquired and combined CP^+ maps for both phantoms. In the silicon oil phantom the RMSE is 5.6° for SatTFL and 5.5° for the Sandwich scheme. In the water phantom the RMSE is 2.7° for SatTFL and 3.9° for the Sandwich scheme.




Figure 4.16: Comparison between acquired CP^+ mode flip angle map and an eight-channel pTx map for a silicon oil phantom.

The 3D results in a water phantom are shown in Figure 4.18. The average RMSE of the Sandwich scheme ($\text{TA} = 20\text{ s}$) was 2.3° over all partitions when compared to SatTFL ($\text{TA} = 360\text{ s}$).




Figure 4.17: Comparison between acquired CP⁺ mode flip angle map and an eight-channel pTx map for a water phantom.




Figure 4.18: 3D B_1^+ maps for 18 partitions (10 mm) of the water phantom acquired using the 3D sequences in CP^+ mode. The top three rows show maps acquired using SatTFL, short TR SatTFL and Sandwich schemes for comparison. The difference and corresponding RMSE between SatTFL and the other sequences is shown below. 3D data acquired in 20 s for the Sandwich scheme and 36 s for short TR SatTFL whereas this took 360 s with SatTFL using a TR of 20 s.

4.4.3 In Vivo Results

The top row in Figure 4.19 shows simulation results for various T_1 values and the bottom row shows Bland-Altman plots of measurements for each scheme from 2D B_1^+ maps acquired at 4 reference voltages to cover a larger dynamic range. For both SatTFL and short TR SatTFL, short T_1 s result in a spread of measured FAs deviating by up to 12° compared to a long T_1 (evaluated at 120°), while this spread is reduced to $<2^\circ$ with the Sandwich scheme. The short TR SatTFL scheme shows especially large T_1 -dependent bias. There is good agreement between the Sandwich and reference methods. However, the limits of agreement vary with α and the Sandwich scheme average measured α deviates $>5^\circ$ above a nominal α of 115° .




Figure 4.19: Simulation results for various T_1 values (0.5 s to 3 s) (top row) and Bland-Altman plots for each scheme for *in vivo* 2D brain data (bottom row).

Simulation shows the insensitivity of the Sandwich scheme to variation in T_1 . The Bland-Altman plots show data aggregated from 5 volunteers for 4 B_1^+ maps at reference voltages of 30 V, 60 V, 90 V and 120 V for greater dynamic range. The 120 V datasets are removed from the Bland-Altman analysis due to oversaturation beyond the linear region but remain in the figure to show the response. Dotted lines in Bland-Altman plots represent ± 1.96 standard deviations from the mean. RPC; reproducibility coefficient, CV; coefficient of variation.

A correlation plot for one subject can be seen in Figure 4.20 combined with simulation.




Figure 4.20: Correlation plots for 2D *in vivo* brain data for **a)** SatTFL, **b)** short TR segmented SatTFL and **c)** the Sandwich scheme plotted against the measured flip angle from the segmented SatTFL method acting as reference. The data in each plot are aggregated from 4 flip angle maps at varying reference voltages (30 V, 60 V, 90 V and 120 V) for greater dynamic range, each point representing the flip angle measured in a single voxel from the ROI (shown in top left). The shaded region corresponds to the deviation due to T_1 values in the range 0.5 s to 3 s for a fixed noise level.

Figure 4.21 shows 2D FA maps of an axial slice of the brain from a single volunteer for each sequence at a reference voltage of 60 V. The first and second rows show the magnitude of the S_0 and S_1 images. As expected, the SNR associated with the Sandwich method (Figure 4.21c) has approximately halved, in both S_0 and S_1 compared to SatTFL. All maps show qualitatively similar results and large B_1^+ inhomogeneity, with the expected brightening around the region of the lateral ventricles due to imaging in CP^+ mode. Structural contrast is only visible in the B_1^+ maps of short TR SatTFL. The difference between the reference map and the maps acquired with the other schemes is shown in the fourth row along with the RMSE.

2D pTx maps are shown in Figure 4.22 for each method. As expected, short TR SatTFL displays severe overestimation with an average RMSE to conventional SatTFL of 15° . The Sandwich displays similar B_1^+ transmit maps to SatTFL with an average RMSE of less than 3° and acquired in 8 seconds.

The difference between an acquired CP^+ mode map to a pTx combined map is shown in Figure 4.23. The Sandwich scheme appears to underestimate the FA towards the centre of the brain, when compared to the acquired CP^+ map, whereas SatTFL overestimates. There is good agreement between the acquired and combined CP^+ maps. The RMSE is 5.1° for SatTFL and 4.0° for the Sandwich scheme.

The 3D FA maps acquired for 18 partitions in the brain of a volunteer are shown in Figure 4.24.




Figure 4.21: *In vivo* brain B_1^+ maps of one volunteer at a reference voltage of 60 V for various 2D sequences. **a)** SatTFL **b)** segmented short TR (2 s) SatTFL **c)** segmented Sandwich scheme with TR = 1 s, $\alpha : \beta = 10:1$. The first two rows show the magnitude of the images (S_0 and S_1) used to calculate the corresponding flip angle maps in the third row. The difference between each of the maps at 60 V and a mean reference map (calculated from two reference voltages for the reference scheme and scaled by the voltage) are shown in the fourth row, along with the RMSE. Flip angle profiles for horizontal and vertical profiles through the centre of the brain are also shown. Note the reduced time to acquire **c)** was 6 s compared to 20 s for **a)** and 8 s for **b)**.




Figure 4.22: Single-channel 2D *in vivo* B_1^+ maps for an axial slice through the brain of a single volunteer for an 8Tx-32Rx head coil. Top row shows individual maps for the SatTFL scheme with a TD = 10 s for a total acquisition time of 90 s (one reference [β in CP⁺ mode] and 8 prepared images [β in CP⁺ mode and α on 1 channel at a time] [197]). The second row shows individual channel maps for short TR SatTFL with a TD of 1 s, for a total acquisition time of 9 s. The third row shows the Sandwich scheme (non-segmented, $N' = 36$) with a TR = 1 s for a total acquisition time of 8 s [α and β one channel at a time]. The RMSE between each channel map and the SatTFL scheme are shown in the bottom right-hand corner. Acquisition time does not include additional 2 s for relative maps.




Figure 4.23: Comparison between acquired CP⁺ mode flip angle map and an eight channel pTx map in the brain.




Figure 4.24: Axial, sagittal and coronal slices through 3D B_1^+ maps for the brain of a volunteer acquired using each method along with a T_1 -weighted MPRAGE. 3D data acquired in 20 s (including two dummy TRs) for the Sandwich scheme, 36 s for short TR SatTFL and 360 s for SatTFL. B_1^+ maps resampled with a linear interpolation using FSL [208] to match the MPRAGE resolution. The RMSE between map and the SatTFL scheme are shown in the bottom right-hand corner.

Figure 4.25 shows 3D B_1^+ maps acquired in the heart for the Sandwich scheme with a 2D SatTFL slice along the horizontal long axis of the heart. The high-resolution CINE demonstrates the inhomogeneous image contrast across the heart at 7 T.




Figure 4.25: B_1^+ maps in the heart of a volunteer acquired using the 3D Sandwich sequence in CP^+ mode. 3D data acquired in 26 heartbeats (including 2 dummy TRs) for the Sandwich scheme compared to 20 heartbeats for a single conventional 2D SatTFL B_1^+ map. A frame from a 2D GRE CINE is shown with the region of interest in green and the red line showing the short axis (SA) slice through the heart.

The 3D Sandwich dataset has been orientated and resampled to match the horizontal long axis (HLA) view of the 2D SatTFL. All images acquired during breath-hold and gated.

4.5 Discussion

Simulations in Figure 4.6 show that removing the delay between S_0 and S_1 provides a considerable decrease in TA without substantial loss in the performance for short image trains. Noise is the dominant cause of error at low α , and data acquired in the high FA region using SatTFL shows the most sensitivity to T_1 . This is demonstrated by the increase in spread of measured FAs. This bias due to T_1 in the SatTFL sequence cannot be corrected for unless the local T_1 is known. Figure 4.6b conveys why shortening of the SatTFL TR period is unwise since it introduces significant T_1 -dependent bias. The Sandwich scheme, as expected from simulations, shows a reduced T_1 sensitivity. The Sandwich scheme also performs similarly well to SatTFL in the range 40-120°. On the other hand, compared to SatTFL, the dynamic range is slightly reduced in the Sandwich scheme, and two dummy TRs must be employed to ensure consistent magnetisation between partitions. The maximum measurable α shifts from 180° for SatTFL to approximately 160° for the Sandwich scheme. However, SatTFL is rather inaccurate when measuring at this extreme due to T_1 .

The simple arccosine model holds when both the imaging excitation FA is small and the readout train short, i.e. if one is increased the other should be reduced, this is demonstrated in Figure 4.8. Since the 3D acquisition used a longer readout train, the $\alpha : \beta$ ratio was increased to 20. This maintained a low T_1 bias and the dynamic range at the expense of a lower SNR. The SNR loss due to this higher ratio is however compensated by the nature of the 3D acquisition. In both 2D and 3D cases, the Sandwich scheme departs from the arccosine relationship (Eq. 3.7) at large α , this can be mitigated by replacement with a LUT.

The discussed novel method has been proven to obtain *in vivo* maps which perform similarly to the SatTFL method with a significant decrease in TA and less sensitivity to T_1 , as shown in Figure 4.19. The segmented short TR SatTFL method in Figure 4.21 does not show the same SNR loss as the Sandwich scheme since the segmented S_0 is fully acquired before the segmented S_1 , however the B_1^+ measurements are invalid. Cerebrospinal fluid in the ventricles appears dark in S_0 of the Sandwich scheme due to the long T_1 , unlike SatTFL where it appears bright. As seen in Figure 4.21, for a further reduced scan time than short TR SatTFL, the Sandwich scheme has only 0.5° higher RMSE than SatTFL, largely due to chemical shift effects and remains accurate within the brain.

A benefit of the paired acquisition of a reference and prepared image provides insensitivity to magnetisation history which also enables rapid multi-transmit mapping. 2D single-transmit maps were acquired for each scheme, as shown in Figure 4.22. The Sandwich method provided a substantial reduction in TA whilst maintaining a small RMSE to SatTFL. The pTx combined map in Figure 4.23 shows that the Sandwich scheme tended to underestimate at the centre of the brain whereas SatTFL tended to overestimate, which was also seen in Figure 4.17 for the water phantom. However, the SatTFL pTx acquired map in Figure 4.23 shows much larger overestimation than seen in phantom and shows structural artefacts likely due to sensitivity to T_1 , not seen for the Sandwich scheme. Improvements in the multi-transmit maps could be achieved using a linear combination of transmit sensitivities such as Fourier encoding [200, 201].

The 3D Sandwich scheme performs similarly well with low RMSE to SatTFL (Figure 4.24), allowing a 3D B_1^+ map to be acquired in approximately the same time as a single TR of SatTFL (26 s vs 20 s). In Figure 4.25, both SatTFL and the Sandwich scheme show similar B_1^+ map profiles in the heart. Recent work has applied 3D FLASH readouts to SatTFL [172, 210] to enable single-shot acquisitions of an entire brain volume. Although this works well in the brain, this is not suitable for the heart as the readout is much longer than diastole and T_1 effects become a concern.

Aside from the two additional dummy TR periods, the Sandwich method does not necessarily change the energy deposited compared to the conventional SatTFL method (Table 4.1). Hence, due to the much shorter duration the Sandwich method shows a large increase in SAR, which is directly proportional to the reduction in TA. In particular, the use of two dummy TR periods and reduction in scan time accounts for approximately a factor of 22-fold increase in SAR over conventional SatTFL.

	SatTFL (2D)	SatTFL (3D)	SatTFL (3D)	Sandwich (3D)	Sandwich (3D)	DREAM
α type/ FA/ t_{dur} / V_{peak}	Rect./ 90°/ 0.5 ms/ 60 V	Sinc/ 90°/ 5 ms/ 180 V	HS8/ 90°/ 5 ms/ 37 V	Sinc/ 90°/ 5 ms/ 180 V	HS8/ 90°/ 5 ms/ 37 V	Sinc/ 40°/ 0.7 ms/ 80 V
β type/ FA/ t_{dur} / V_{peak}	Rect./ 4°/ 0.1 ms/ 14 V	Rect./ 4°/ 0.1 ms/ 14 V	Rect./ 4°/ 0.1 ms/ 14 V	Rect./ 4°/ 0.1 ms/ 14 V	Rect./ 4°/ 0.1 ms/ 14 V	Sinc/ 15°/ 0.7 ms/ 30 V
TR	20 s	20 s	20 s	1 s	1 s	1 s
Acquisition Time	480 s	480 s	480 s	26 s (incl. 2 dummy TRs)	26 s (incl. 2 dummy TRs)	24 s
Average 10 s Power (Per Channel)	0.001 W (S_0) or 0.005 W (S_1)	0.001 W (S_0) or 0.012 W (S_1)	0.001 W (S_0) or 0.012 W (S_1)	0.135 W	0.137 W	0.141 W
Total Energy (Per Channel)	1.54 J	3.25 J	3.28 J	3.52 J	3.56 J	3.39 J

Table 4.1: Comparison of 10s energy and total energy for SatTFL, the Sandwich scheme and DREAM using a reference voltage of 60 V and a matrix size of $50 \times 36 \times 24$. Power limits for the Nova head coil (8Tx/32Rx) were 12 W short-term (10s average) and 24 W long-term (6 minute average). Sinc and HS8 preparation pulses have matched bandwidths of ± 3 kHz. Sinc preparation pulses are likely to be hardware restricted by peak forward voltage.

Table 4.1 also shows that compared to a matched acquisition DREAM sequence, there is a similar average power deposition for the 3D Sandwich scheme. Using a non-selective broad-band full-passage HS8 pulse as the preparation pulse does contribute significant additional energy over the standard rectangular pulse although it is still $<10\%$ of power limits. The HS8 has equivalent RF energy but with only 20% of peak B_1^+ for an equivalent bandwidth and FA sinc pulse, reducing the effect of hardware constraints on deliverable peak B_1^+ . Additionally, the HS8 is robust to large B_0 inhomogeneities over a bandwidth ± 1.5 kHz (Figure 4.2), whilst maintaining a monotonic relationship to B_1^+ . In my experience, the HS8 bandwidth covers the range of B_0 expected in both the head and the heart. The HS8 pulse, however, is non-linear for higher FAs and a LUT correction should be used to account for this. The non-linearity of the HS8 at high FAs, with the use of a LUT, potentially facilitates an increase in the measurable B_1^+ up to 200° .

An adaptation to the commonly used SatTFL B_1^+ mapping sequence has been introduced and the performance of the method in simulations and in experiments, both in phantom and *in vivo*, has been demonstrated. The intention here has been to show the ability to produce absolute volumetric FA maps for the purpose of body

imaging, with validation in the brain and heart. The adaptations retain the relatively low SAR, are less T_1 dependent, require no special saturation pulses and are more robust to motion, i.e. able to obtain S_0 and S_1 within the same cardiac cycle.

Sandwiching the images overcomes the issues associated with a short TR SatTFL sequence by minimising T_1 relaxation between S_0 and S_1 . The Sandwich method achieves similar measurement times as DREAM but without the flow artefacts (i.e. apparent depleted B_1^+ in blood chambers) and may be of potential interest for cardiac imaging, as well as having applications elsewhere. Going forward, it could be possible to include additional echoes within the TurboFLASH train which would allow for simultaneous B_1^+ and B_0 field mapping. Further work could also combine this Sandwich B_1^+ mapping sequence with acceleration methods such as TxLR [69]. TxLR in combination with a B1TIAMO [107] style acquisition could allow acceleration of multi-channel mapping by up to a factor of eight [6], and may allow single-shot volumetric multi-channel mapping of the pTx array in a single breath-hold for cardiac gated applications. This is precisely the topic of Chapter 7.

4.6 Summary

The Sandwich sequence enables fast acquisitions with reduced T_1 bias and facilitates 3D B_1^+ measurements of the whole thorax or brain in 26 s or less.

5

Comparison of B_1^+ Mapping Methods

Contents

5.1	Introduction	104
5.2	Simulation Framework for B1+ Mapping Sequence Comparison	105
5.2.1	Methods	105
5.2.2	Results	107
5.2.3	Discussion	112
5.3	FLASH-based B1+ Mapping Comparison	113
5.3.1	Methods	113
5.3.1.1	Parallel Transmit MRI System	113
5.3.1.2	Sequence Implementation	114
5.3.1.3	Image Reconstruction	114
5.3.2	Results	115
5.3.3	Discussion	119
5.4	Summary	119

5.1 Introduction

Many B_1^+ mapping methods have been developed thus far, but no single method has become the "gold-standard", with many UHF sites opting for their own implementations. Previous works have used Monte Carlo simulations to evaluate sequences [135, 211–213] but these simulations are not open-source, do not consider off-resonance and spin-history effects and do not include newer techniques. Hence, in this chapter I attempt to gain better insight into the trade-offs of some popular methods. This chapter is split into two sections. The first section describes an open-source EPG, Bloch and Monte Carlo simulation framework which allows for an

easy comparison between methods under various conditions such as B_0 , T_1 and flow. The second section describes an experiment comparing three popular FLASH-based methods of B_1^+ mapping, namely SatTFL, SA2RAGE and Sandwich.

5.2 A Simulation Framework for B_1^+ Mapping Sequence Comparison

5.2.1 Methods

The simulation framework was based on the EPG simulations performed for Chapter 4 with the inclusion of flow and diffusion as well as additional sequences. Using EPG, Bloch and Monte Carlo methods simulations for 2D multi-slice and 3D B_1^+ mapping sequences were simulated in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States). The five sequences chosen for comparison were:

- (a) 2D Multislice SatTFL
- (b) 3D Sandwich
- (c) 3D SA2RAGE
- (d) 3D AFI
- (e) single-shot 3DREAM

The sequences were simulated as closely as possible to match the literature using protocols as shown in Table 5.1 including RF spoiling and slice profile effects using a slice select gradient amplitude of 5 mT m^{-1} and a 10 mm slice thickness. A large DR of B_1^+ was simulated from 0 to 3 times the nominal FA. B_0 off-resonances from 0 kHz to 1 kHz were simulated along with T_1 times of 0.5 s to 3 s. The T_2 for all simulations was kept fixed at 25 ms. The sensitivity of each method to flow and diffusion was evaluated by simulating a range of coherent flows with a velocity in the range of 0 m s^{-1} to 0.2 m s^{-1} and an isotropic free water diffusion coefficient of $3 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$.

The effect of blood flow was simulated by assuming coherent motion. The effect of coherent flow on the configuration state matrix is given by the flow operator, $\mathbf{F}$, for a particular velocity, v , with angle, γ , between the flow and gradient direction of duration, τ , is given by

$$\text{diag}(\mathbf{F}) = \exp \left(-i \left[n + \frac{\Delta n}{2}, n + \frac{\Delta n}{2}, n \right] v \tau \cos(\gamma) \right), \quad (5.1)$$

where n is the wave vector magnitude and Δn is the change in n due to an applied gradient. Note that for the Z -state $\Delta n = 0$ because there is no additional dephasing of this state during a gradient.

Sequence	TE (ms)	TR (ms)	TR _{tot} (ms)	Nominal Imaging Excitation	Nominal Saturation	Other	Acq. Time (s)
MS SatTFL	1.78 ^a	4.2 ^a	10000 ^b	5° 2 ms TBP 4 wind. Sinc	90° 5 ms TBP 9 Sinc V _{peak} =54 V	Odd then even slice order	20
Sandwich	1.78 ^a	4.2 ^a	1000	5° 0.1 ms Rect	90° 5 ms TBP 15 HS4 V _{peak} =32 V	2 Dummy TR _{tot} LUT ^c	34
SA2RAGE	1.13 ^a	2.8 ^a	2400	4°/11° 0.1 ms Rect	90° 0.5 ms Rect V _{peak} =60 V	TD ₁ =50ms TD ₂ =1800ms 4 Dummy TR _{tot} LUT ^c	87
AFI	5	20/100	120	60° 0.5 ms Rect V _{peak} =60 V	N/A	20 Dummy TR _{tot}	126
3DREAM	1.2/1.9 ^d	3.2	3.2	6° 0.2 ms Rect	60° 0.4 ms Rect V _{peak} =50 V	N/A	4

Table 5.1: Table of parameters used for each sequence simulation. All sequences simulated for $N_{PE1} \times N_{PE2} = 32 \times 32$. ^aTurboFLASH TE/TR. ^bMinimum TR between reference/prepared. ^cLookup table (LUT) generated using a nominal T_1 of 1.5 s. ^dSTE/FID. TBP; Time-bandwidth product. TD; Delay time.

The diffusion of blood within the chambers of the heart may lead to erroneous measurements of B_1^+ . In order to explore the effects of flow diffusion these were included in the EPG simulations using an isotropic diffusion coefficient, D . Similarly to as above for flow, isotropic diffusion was implemented into the simulation framework using the isotropic diffusion operator, $\mathbf{D}$, which acts on the configuration state matrix and is defined, following Weigel et al. [47], as

$$\text{diag}(\mathbf{D}) = \exp \left(- \left[\left(n + \frac{\Delta n}{2} \right)^2 + \frac{\Delta n^2}{12}, \left(n + \frac{\Delta n}{2} \right)^2 + \frac{\Delta n^2}{12}, n^2 \right] \tau D \right). \quad (5.2)$$

The relative effects of noise across the different schemes were evaluated using a Monte Carlo method with 1000 repetitions, with a fixed level (peak SNR 60 dB) of zero-mean complex Gaussian noise added independently to each image. The FA was calculated as reported in the literature and all methods in this section consider the phase to extend the DR. The Sandwich and SA2RAGE sequences FAs were calculated using a LUT with a global T_1 assumed to be 1.5 s.

Since the methods use different nominal FAs, results were normalised by their DR for plotting and evaluation. The DR of 1 was defined as the lowest B_1^+ measured, on resonance, where the difference in the mean FA and 1.96 standard deviations (95% of observations) across all T_1 values was less than 10% of the simulated value. Simulation and figure generating code can be found at: <https://github.com/jameslewiskent/B1-mapping-simulations>.

5.2.2 Results

Figure 5.1 shows the FWHM of the PSF for each of the five sequences simulated. All methods show intervoxel blurring is low in the first image (S_0) whilst intervoxel blurring is higher in the second image (S_1) particularly for SA2RAGE and Sandwich above the nominal preparation FA. The PSF peaks where the signal crosses a null point.




Figure 5.1: The full width at half maximum of the point spread function for both reference (S_0) and prepared (S_1) images of each B_1^+ mapping sequence.

Figure 5.2 shows the response of the measured B_1^+ to different T_1 values. All methods using long readout trains (a, b, c) suffer from large T_1 dependence, which can be reduced by shortening the readout train (decreasing the TurboFactor), minimising the image train TR or using variable FAs. Sandwich and SA2RAGE are most affected by noise as demonstrated by the larger shaded regions ($\pm 1.96\text{SD}$). AFI shows the least T_1 dependence and only a slight underestimation (2.5° underestimation at 50°) of the measured FA which varies with the applied B_1^+ which could be corrected with a LUT. 3DREAM also displays a systematic error with an increasing underestimation of B_1^+ .

Table 5.2 displays the calculated DR of the five methods. SatTFL, Sandwich and SA2RAGE show similar DRs of between 5 to 6. AFI displays by far the largest DR of around 11 and able to map the lowest FAs. 3DREAM has the lowest DR of 4.3. However, since the DR can be improved by additionally acquiring maps with




Figure 5.2: The effect of T_1 values in the range of 0.5s to 3s on the measured B_1^+ for the five sequences simulated. Results are plotted in units of dynamic range (DR) and the linear component is removed. The applied DR is defined as the nominally applied flip angle divided by the lower bound flip angle of the DR (i.e. $DR = 1$). Shaded regions show ± 1.96 standard deviations.

multiple voltages or Tx modes, another metric to evaluate the methods is the DR efficiency ($DR/$ Acquisition Time). A method with low DR but high DR efficiency can be repeated to artificially increase the DR. Hence, a higher DR efficiency is desirable. Out of the methods simulated, 3DREAM has by far the highest DR efficiency and SA2RAGE the lowest. Although, it is likely that an additional delay would be required for 3DREAM of a few seconds between shots. A 4s delay between 3DREAM shots would decrease the DR efficiency to $0.54s^{-1}$, which is still the highest DR efficiency of any of the methods investigated here.

Sequence	Dynamic Range Bounds	Dynamic Range [a.u.]	Dynamic Range Efficiency [s ⁻¹]
MS SatTFL	27.0° - 174.6°	6.47	0.32
Sandwich	28.8° - 143.1°	4.97	0.15
SA2RAGE	29.7° - 162.0°	5.45	0.06
AFI	16.2° - 180.0°	11.10	0.09
3DREAM	19.8° - 85.2°	4.30	1.08

Table 5.2: Dynamic range of the simulated methods. The dynamic range was defined as the B_1^+ measured, on resonance, where the difference in the mean flip angle and 1.96 standard deviations across all T_1 values was less than 10% of the simulated value. DR efficiency = DR/Acquisition Time.

Figure 5.3 shows the response of the measured FA to B_0 off-resonances of 0 Hz, 100 Hz, 500 Hz, and 1 kHz. All methods show good B_0 insensitivity up until 100 Hz. Although, Sandwich sequence is the only sequence with robustness up to ± 1 kHz, thus requiring no additional B_0 correction. Whereas all other methods above 100 Hz display substantial systematic errors due to B_0 which require correction from a B_0 map.




Figure 5.3: The effect of B_0 values in the range 0 kHz to 1 kHz on the measured B_1^+ for the five sequences simulated for $T_1 = 1.5$ s. Results are plotted in units of dynamic range (DR) and the linear component is removed. The applied DR is defined as the nominally applied flip angle divided by the lower bound flip angle of the DR (i.e. DR = 1).

Figure 5.4 shows the effect of flow on each of the methods. 3DREAM and AFI display the most significant sensitivity to flow. 3DREAM has complete dropout for even the lowest flow rate simulated as expected due to the use of a stimulated echo. Whereas AFI has a complicated relationship to flow but still maintains reasonably accurate measurement of B_1^+ with a small flow-related bias. SA2RAGE is also somewhat impacted by flow and SatTFL and Sandwich display no such impact.




Figure 5.4: The effect of coherent flow in the range 0 ms^{-1} to 0.2 ms^{-1} on the measured B_1^+ , on resonance, for the five sequences simulated for $T_1 = 1.5 \text{ s}$. Results are plotted in units of dynamic range (DR) and the linear component is removed. The applied DR is defined as the nominally applied flip angle divided by the lower bound flip angle of the DR (i.e. $\text{DR} = 1$).

The impact of diffusion of the five sequences are shown in Figure 5.5. Although diffusion is not expected to play a significant role in the accuracy of B_1^+ maps, particularly for the resolutions acquired here, diffusion could complicate B_1^+ in the blood pools of the heart. Therefore, the less sensitive a method is to diffusion the better. These results show all methods suffer little bias due to isotropic free water diffusion.




Figure 5.5: The effect of isotropic free water diffusion with apparent diffusion coefficient of $3 \times 10^{-9} \text{m}^2 \text{s}^{-1}$ on the measured B_1^+ , on resonance, for the five sequences simulated for $T_1 = 1.5 \text{s}$. Results are plotted in units of dynamic range (DR) and the linear component is removed. The applied DR is defined as the nominally applied flip angle divided by the lower bound flip angle of the DR (i.e. $\text{DR} = 1$).

Table 5.3 shows the average 10 s and total energy of the five sequences. 3DREAM has the lowest acquisition time and total energy of any of the methods tested here. Conversely, AFI has the longest acquisition time and highest total energy. Aside from the energy, Sandwich has the lowest peak voltage which is beneficial when mapping regions of low B_1^+ .

Sequence	Average 10 s Power (Watts per channel)	Total Energy (Joules per channel)	Peak Voltage (Volts)
MS SatTFL	0.61	1.20	54
Sandwich	1.10	3.73	32
SA2RAGE	0.56	4.80	60
AFI	2.67	33.41	60
3DREAM	1.43	0.45	50

Table 5.3: Average 10 s power, total energy and peak voltage for the five simulated sequences.

5.2.3 Discussion

The choice of a single B_1^+ mapping technique, as originally implemented, seems to be impossible. However, the simulation framework provides better understanding of the trade-offs between the different B_1^+ mapping sequences and the impact of various factors on their accuracy to help to inform better pulse sequence design and parameter optimisation.

During this work I found that using TurboFactors significantly larger than 32 will result in severe T_1 bias and are to be avoided for SatTFL, Sandwich and SA2RAGE. These simulations used literature parameters, however, if Sandwich were to use the same TE/TR_{FLASH} as SA2RAGE this would further reduce the T_1 sensitivity and hence increases the DR of the Sandwich to 5.4 (27.9° - 151.2°) which is on par with SA2RAGE but 40% of the acquisition time.

In my experience, the majority of B_0 off-resonances in the cerebrum are expected to be below 100 Hz where all compared sequences work well in this range. However, off-resonances can be as much as 1.2 kHz when including the brain stem and neck [214]. Across the heart, off-resonances can be on the order of a few hundred Hz and B_1^+ mapping must be robust to these off-resonances [206]. The large B_0 inhomogeneities found in the body may make AFI unsuitable for B_1^+ mapping in the body. Therefore, the Sandwich sequence appears to be the best candidate for gated volumetric cardiac B_1^+ mapping due to its insensitivity to B_0 , flow as well as its ability to have a non-fixed variable TR_{tot} time. The impact of heart rate variability on sequences was not explored here. SA2RAGE could also benefit from using an HS8 pulse for better B_0 insensitivity. A potential additional benefit of B_0 robust B_1^+ mapping sequences is that the acquisition of B_1^+ maps could be performed before B_0 mapping which may also help improve B_0 shimming [206]. However, this has not been investigated here.

If desired, additional sequences could be easily added to this open-source simulation framework and other effects, such as magnetisation transfer, could be investigated. The effects of magnetisation transfer could be implemented into the simulation framework by following Malik et al. and simulating multiple compartments [215], similar to how I investigated ejection fraction effects on the Sandwich sequence in Section 4.3.1.1. I expect magnetisation transfer to have a negligible effect on the Sandwich and 3DREAM schemes due to the negligible time between the reference and prepared images as well as the phase encoding scheme that

was employed. Whereas, magnetisation transfer is likely to be more of a problem for SatTFL, AFI and SA2RAGE where there is more time for free water and protons bound to macromolecules to exchange magnetisation. I am unable to find any quantification of how much of an impact this has on any B_1^+ mapping scheme in the literature but given the resolution of the B_1^+ maps I expect this to be low.

5.3 FLASH-based B_1^+ Mapping Comparison

FLASH-based absolute B_1^+ mapping is suitable for UHF and is flow-insensitive and robust to motion. 2D B_1^+ sequences, e.g. SatTFL, are often used in the cerebrum due to their fast acquisition times. However, the compromise is anisotropic voxels, slice profile effects, and smaller coverage, and it may not be feasible in other body regions such as the heart. Additionally, the use of slice-selective saturation pulses can be SAR-limiting due to their high peak voltage. A 3D B_1^+ mapping sequence is desirable, but a 3D sequential SatTFL acquisition is too slow. In this section, I assessed three FLASH-based B_1^+ mapping methods in a head phantom at 7 T.

5.3.1 Methods

B_1^+ maps were acquired using 2D SatTFL, 3D Sandwich and 3D SA2RAGE using parameters closely matching those from the literature. 3D Sandwich and 3D SA2RAGE sequences were implemented by making minor pulse sequence modifications to a FLASH product sequence in the C++ XA60 IDEA environment on a Windows 10 Professional machine.

5.3.1.1 Parallel Transmit MRI System

A 7 Tesla MRI (Magnetom Terra.X, Siemens Healthineers, Erlangen, Germany) equipped with pTx was used for phantom imaging. The imaged phantom was an ‘Ella’ head phantom with realistic dielectric properties (61% distilled water, 36% polyvinylpyrrolidone, 0.7% NaCl, 1.5% Agarose). It is worth noting that this phantom contains obvious cracking defects which can be seen in the localiser as well as signal voids where glass rods had previously been used to insert temperature probes. All phantom images were acquired using the same 8-channel Tx, 32-channel Rx head coil (Nova Medical, Wilmington MA, USA) in CP⁺ mode at magnet isocentre.

5.3.1.2 Sequence Implementation

A 2D SatTFL protocol was chosen to closely match that of a default product protocol. Appropriate Sandwich and SA2RAGE parameters were chosen from the literature [2, 168]. All sequences acquired a fully sampled $24 \times 32 \times 32$ scan matrix (including asymmetric echo) for a large $256 \times 256 \times 256 \text{ mm}^3$ FOV covering the entire head phantom. All sequences were linearly encoded with a bandwidth of 500 Hz px^{-1} , a TurboFactor of 32 and were RF spoiled. Sequence-specific parameters and acquisition times are summarised in Table 5.4. On the scanner, shorter FLASH TE/TR times of 0.9 ms/2.7 ms than used previously for Sandwich (1.8 ms/4.2 ms) were obtainable which are expected to improve the DR by reducing T_1 sensitivity. B_1^+ maps were acquired at 9 reference voltages (channel-combined) between 50 V and 450 V in 50 V increments and combined to produce a common "ground truth" across the entire FOV following the work of de Buck et al. [214].

5.3.1.3 Image Reconstruction

Image reconstruction was performed offline in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States) and coil sensitivities were estimated using ESPIRiT [51]. Images were masked based on the reference image intensity. As in the literature, B_1^+ maps were calculated using an arccosine function for SatTFL (Eq. 3.7), and using a LUT for Sandwich and SA2RAGE which were generated from EPG and Bloch simulations with an assumed T_1 of 1.5 seconds, shown in Figure 5.6.

Maps were additionally masked based on two simple criteria; firstly, that there is a minimum amount of signal in the reference images and secondly, that the measured FA should be increasing with voltage, otherwise the voxel was excluded from the 'ground truth'. This means that data acquired where the RF pulse was clipped, due to exceeding the maximum peak B_1^+ achievable, did not contribute to the 'ground truth' estimation. Reconstruction code and data are openly available at: <https://github.com/jameslewisKent/FLASH-B1-Mapping>.




Figure 5.6: Lookup tables generated from simulations which were used to calculate the B_1^+ maps from the prepared/reference image ratio.

Sequence	FLASH TE/TR (ms)	TR _{tot} (s)	Nominal Imaging Excitation	Nominal Saturation	Other	Scan Time (s)
2D SatTFL	1.38/3.4	10 ^a	5° 0.4 ms Sinc	90° TBP 6 1.4 ms Sinc	1000 Dummy RF Pulses Interleaved Slice Order 6.5 mm Slice Thickness 1.5 mm Slice Distance	21
3D Sandwich	0.87/2.72	1	4° 0.1 ms Rect	90° 0.5 ms Rect ^b	2 Dummy TR _{tot} 's LUT ^c	34
3D SA2RAGE	0.87/2.74	2.4	4°/11° 0.1 ms Rect	90° 0.5 ms Rect	TD ₁ =53 ms TD ₂ =1800 ms 4 Dummy TR _{tot} 's LUT ^c	86

Table 5.4: Sequence protocol parameters used to acquire the B_1^+ maps. ^aMinimum TR between reference and prepared images. ^bConventionally, Sandwich uses a HS8 however this was not used here for a fairer comparison to SA2RAGE which could also benefit from this saturation pulse. ^cLookup table generated from extended phase graph and Bloch simulations using a nominal T_1 of 1.5 s.

5.3.2 Results

The "ground truth" maps are shown in Figure 5.7 along with the localiser which clearly demonstrates the B_1^+ shading artefact. Figure 5.7a) to c) had a mean FA per volt across the entire head of $0.25 \pm 0.02^\circ/\text{V}$ for SatTFL, $0.26 \pm 0.03^\circ/\text{V}$ for SA2RAGE and $0.26 \pm 0.01^\circ/\text{V}$ for Sandwich. Streaking artefacts are most noticeable in the transverse view of the SatTFL "ground truth" maps. SatTFL had the greatest signal magnitude due to the long (10 s) recovery time between the two images. SA2RAGE had an average reference image magnitude of $2.3\times$ Sandwich, whereas Sandwich had an average prepared image magnitude of $1.1\times$ SA2RAGE.




Figure 5.7: "Ground Truth" B_1^+ maps calculated from a combination across multiple transmission voltages shown along three orthogonal planes in a human head phantom for each of the three FLASH sequences. Note the phantom contains some cracking defects and signal voids which are visible in the localisers.

Figure 5.8 shows correlation and Bland-Altman plots for each sequence. SA2RAGE and Sandwich show good agreement with similar DRs of 20°-140°, whereas SatTFL shows a reduced DR of 20°-100°. Reference voltages for SatTFL above 275 V lead to clipping of the saturation RF pulse, as the high pulse voltages cannot be supplied by the RFPA, hence voxels from measurements above this were not included in the 'ground truth' calculation. Both SatTFL and Sandwich tend to underestimate high FAs, whereas SA2RAGE overestimates.

Figure 5.9 shows a single measurement at 200 V. The departure from a linear relationship occurs for SA2RAGE around 120° and leads to an increasing over-estimation of B_1^+ . SatTFL had slightly fewer voxels remaining after automated masking. The standard deviation from the "ground truth" was 1.9°, 1.7° and 1.3° for SatTFL, SA2RAGE and Sandwich, respectively.




Figure 5.8: Correlation and Bland-Altman plots for each method showing voxel-wise flip angle measurements from across the entire field of view for transmission voltages of 50 V to 450 V plotted in different colours. Clipping of the RF saturation pulse above 275 V in **a)** leads to erroneous B_1^+ measurements which were not included in ground truth estimation.




Figure 5.9: Correlation and Bland-Altman plots for each method showing voxel-wise flip angle measurements from across the entire field of view for a single transmission voltage of 200 V. The dotted lines in Bland-Altman plots represent ± 1.96 standard deviations from the mean (dashed lined). n; number of voxels in analysis, RPC; reproducibility coefficient, CV; coefficient of variation.

5.3.3 Discussion

SatTFL displays significant disadvantages, in particular poor accuracy in regions of low B_1^+ necessitating higher voltages which may not be achievable due to RF clipping. This is likely a problem for body imaging. Using a LUT for SatTFL instead of the arccosine function did not improve the SatTFL results. Sandwich and SA2RAGE both enable higher reference voltages to reach areas of low B_1^+ by using non-selective RF pulses. Typically, the Sandwich sequence makes use of a non-adiabatic HS8 pulse as the saturation, giving better B_0 insensitivity than a rectangular pulse. However, since SA2RAGE could also benefit from this saturation pulse and it was not available to SA2RAGE at the time of data acquisition, I have not used this to allow for a fairer comparison.

Recently, a SA2RAGE acquisition accelerated with compressed sensing achieved up to 15-fold undersampling with $<5\%$ relative error in B_1^+ [168]. Hence going forward, it would be interesting to investigate how Sandwich and SA2RAGE compare when accelerated, given the slight differences in SNR of the two methods. It is likely that given the higher SNR of SA2RAGE that higher acceleration factors would be achievable. However, since the acquisition time for Sandwich in these experiments was 40% of that for SA2RAGE, to achieve a similar acquisition time, SA2RAGE would need an acceleration 2.5 times greater than that of Sandwich. Additionally, when acquiring low resolution B_1^+ maps the push to higher acceleration factors brings diminishing returns in terms of acquisition times. An accelerated version of the Sandwich sequence is explored in Chapter 7.

In conclusion, both SA2RAGE and Sandwich enable non-selective RF pulses with lower transmit power demand maintaining accuracy for low B_1^+ regions.

5.4 Summary

The simulation framework has proven a useful tool to compare B_1^+ mapping methods under various scenarios and allow for easy sequence optimisation. The Sandwich sequence appears to outperform other FLASH-based techniques, both in simulation and in phantom, in terms of acquisition time, robustness to B_0 off-resonances and the peak B_1^+ achievable.

6

Coil-cycled Relative Maps

Contents

6.1	Introduction	120
6.2	Methods	123
6.2.1	Numerical Heart Simulations	123
6.2.1.1	Simulated 2D Multi-slice Protocols	123
6.2.1.2	Simulated 3D Protocols	124
6.2.1.3	Synthetic Heart Data as Simulation Input	124
6.2.1.4	Phantom Data as Simulation Input	124
6.2.1.5	Simulation Experiments	125
6.2.2	Parallel Transmit MRI System	125
6.2.3	Sequence Implementation	125
6.2.4	Image Reconstruction	126
6.2.5	Phantom Experiments	126
6.2.6	<i>In Vivo</i> Experiments	127
6.3	Results	127
6.3.1	Numerical Simulation Results	127
6.3.2	Phantom Results	134
6.3.3	<i>In Vivo</i> Results	136
6.4	Discussion	136
6.5	Summary	138

6.1 Introduction

Conventionally, relative maps are acquired *channel-wise* (CW) *sequentially*. This means that one channel is fully imaged (i.e. channel-wise) before moving to the next neighbouring coil (i.e. sequentially). "Sequentially" describes the acquisition

order (AO) in which Tx channels are mapped. Acquiring relative maps CW sequentially works well enough for phase-only shimming as the phase of a GRE image is independent of T_1 , TE, TR and FA so GRE images can be acquired at the Ernst angle. However, in order to maintain accuracy for magnitude-based shimming, typically two conditions are usually recommended for relative magnitude mapping [130]. These are:

1. Small FAs ($FA \leq 10^\circ$) so that signal is linearly proportional to FA.
2. TR chosen to reduce T_1 relaxation bias.

However, spin history effects from CW sequential acquisitions are also potentially problematic and may bias relative B_1^+ magnitudes. This is because spins within overlapping regions of Tx sensitivity, as illustrated by Figure 6.1, will have had their magnetisation altered.




Figure 6.1: Illustration to demonstrate the overlapping sensitivities of transmit coils when imaging the body. The blue area represents the transmit field from Tx 4 whereas the pink area represents the transmit field from Tx 5. The overlapping region is marked where spins will have previously been excited when transmission is moved from Tx 4 to Tx 5.

Typically, for CW mapping a delay is introduced between the mapping different coils of approximately 4s to allow for some relaxation before moving to the next channel. In order to increase the delay between the mapping of neighbouring coils, instead of mapping coils sequentially in an AO 1→8, the order in which channels are imaged could be distributed by moving to the coil that maximises the delay between neighbouring coils i.e. an AO of [1,4,7,2,5,8,3,6]. In this example, this would result in an additional effective delay between mapping neighbouring channels

of $2 \text{ TR}_{\text{shot}}$ s, therefore tripling the effective delay time to 12s. An illustration of the acquisition of sequential versus distributed relative maps is shown in Figure 6.2.




Figure 6.2: Illustration of the acquisition order for acquiring relative B_1^+ maps sequential (red) versus distributed (blue). Sequential maps have an acquisition order of [1,2,3,4,5,6,7,8] whereas distributed maps have an acquisition order of [1,4,7,2,5,8,3,6].

A complimentary method to distributing the AO is a coil-cycled (CC) acquisition. This is done by only acquiring a single k -space line on a particular Tx channel before moving to the next channel [107]. Not only does coil-cycling avoid spin history effects, but it allows for the reduction in the TR as demonstrated by Brunheim et al. [107] where a 2D CC acquisition of 6s was shown to give equivalent results to a 2D channel-wise sequential acquisition that was acquired shot-wise (one Tx channel per shot) of 32s for a 32 Tx system. Psuedo code for acquiring relative maps channel-wise sequential versus CC distributed is shown in Figure 6.3.




Figure 6.3: Pseudo code for acquiring relative B_1^+ maps channel-wise sequential versus coil-cycled distributed. A simple reordering of the acquisition loops is required. The acquisition order is indicated in the square brackets. PE_N ; number of phase encoding lines.

The aim of this chapter was to explore how much of an effect coil-cycling and a distributed AO had on relative magnitude maps and how critical it is to implement. To do this EPG simulations were performed using a synthetic body model with realistic Tx field sensitivities as well as a phantom experiment.

6.2 Methods

6.2.1 Numerical Heart Simulations

The EPG simulation framework from Chapter 5 was extended to include multi-Tx mapping simulations in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States) using realistic Tx sensitivities including slice-profile effects.

Both 2D and 3D sequences were simulated for CW and CC single-channel relative maps, with both a sequential and distributed AO. Sequential maps had an AO of [1,2,3,4,5,6,7,8]. Whereas distributed maps had an AO of [1,4,7,2,5,8,3,6] to maximise the duration between neighbouring Tx coil excitations. T_2 was fixed at 25 ms. A fixed level (peak SNR 60 dB) of zero-mean complex Gaussian noise was added to the simulated signal before analysis. The effect of B_0 inhomogeneity was not included in these simulations and on-resonance is assumed. Simulation and figure generating code can be found at: <https://github.com/jameslewiskent/B1-mapping-simulations>.

6.2.1.1 Simulated 2D Multi-slice Protocols

The acquisition of multi-slice CW and CC single-channel relative maps, with both sequential and distributed AOs were simulated using a standard 2D GRE sequence. The simulated sequences used a nominal FA of 5° , 2 ms sinc pulse (TBP = 4) for excitation with TE = 1.31 ms, TR = 3.14 ms, RF spoiling and 100 dummy RF pulses per Tx channel. 32 slices with 32 PE lines were simulated with an even-then-odd slice acquisition encoding to reduce slice-profile effects from neighbouring slices. The 2D CW scheme acquired all PE lines for all slices before a delay and then moving to the next Tx channel in the AO. Therefore, using a $TR_{\text{shot}} = 10$ s gave a delay between shots of approximately 6.5 s (shot duration is $[100 + 32 \times 32] \times 3.14$ ms = 3.5 s). This lead to a total acquisition time of 80 s for the CW scheme and 28.2 s for the CC scheme. The acquisition times include the 100 dummy RF pulses per Tx channel for a total dummy duration of 2.5 s.

6.2.1.2 Simulated 3D Protocols

The acquisition of 3D CW and CC single-channel relative maps, with both sequential and distributed AOs, were simulated using a standard 3D GRE sequence. The simulated sequences used a nominal FA of 5° , 0.1 ms RECT pulse for excitation with TE = 1.31 ms, TR = 3.14 ms, RF spoiling and 100 dummy RF pulses per Tx channel. 32 PE lines were simulated in both PE directions with a linear ascending encoding. The 3D CW scheme acquired all PE lines before a delay and then moving to the next Tx channel in the AO. Therefore, using a $TR_{\text{shot}} = 10$ s gave a delay between shots of approximately 6.5 s (shot duration including dummy RF pulses is $[100 + 32 \times 32] \times 3.14 \text{ ms} = 3.5 \text{ s}$). This lead to a total acquisition time of 80 s for the CW scheme and 28.2 s for the CC scheme. The acquisition times include the 100 dummy RF pulses per Tx channel for a total dummy duration of 2.5 s.

6.2.1.3 Synthetic Heart Data as Simulation Input

The 2D multi-slice and 3D protocols were repeated using the Tx fields from synthetic body data. The *in silico* body $B_1^{+/-}$ fields were simulated in Sim4Life 3.4 (ZMT, Zurich, Switzerland) using an 8-channel Tx/Rx dipole array [216] centred over the heart of Duke (Virtual Population, ITIS Foundation, Zurich, Switzerland) at 7 T [69]. The FOV was $278 \times 356 \text{ mm}^2$ (AP/LR) and data was re-sampled onto a uniform grid of $2 \times 2 \text{ mm}^2$. Synthetic body T_1 values were arbitrarily defined based on tissue type except for heart tissue and blood which were given T_1 values of 1.925 s and 2.29 s, respectively, (see Figure 6.4). For the purpose of plotting, the magnetisation in 10 voxels around the Duke synthetic body, including two in Duke's heart (one in the myocardium, one in the blood pool) were tracked. The locations of the tracked voxels are also shown in Figure 6.4.

6.2.1.4 Phantom Data as Simulation Input

The 2D multi-slice and 3D protocols in Sections 6.2.1.1 and 6.2.1.2, respectively, were repeated using the Tx fields from the CC phantom data acquired on the scanner using a 8Tx/32Rx head coil, as outlined in Section 6.2.5. Simulations used a fixed T_1 of 1.5 s.




Figure 6.4: Locations of tracked voxels ($\times$) in the synthetic Duke body model and positions of the transmit elements (Tx 1-8). The arbitrary T_1 values used for simulation are shown.

6.2.1.5 Simulation Experiments

The behaviour of CW sequential versus CC distributed acquisitions were assessed as 2D and 3D sequences using the synthetic body data and phantom data from the head coil under various protocol parameters. A range of nominal FAs was simulated from 1° to 15° in 1° increments, as well as a range of T_1 values from 0.4 s to 4 s in 200 ms increments (replacing the synthetic T_1 values) and TRs of 5 ms to 250 ms in 20 ms increments. The normalised root mean square error (NRMSE) was evaluated for the central slice of the 2D multi-slice sequence and the central partition for the 3D sequence. Plots for the Duke dataset show the NRMSE using both the whole-body and heart only masks.

6.2.2 Parallel Transmit MRI System

A 7 Tesla MRI (Magnetom Plus VE12U_AP01, Siemens Healthineers, Erlangen, Germany) equipped with pTx was used for phantom imaging. Phantom and brain images were acquired using the same 8-channel Tx, 32-channel Rx head coil (Nova Medical, Wilmington MA, USA) at magnet isocentre.

6.2.3 Sequence Implementation

The default product sequence "tfl_rfmap" (SatTFL) was used to acquire 2D CW sequential maps both shot-wise and single-shot. Whereas the Sandwich sequence was coded for VE12U_AP01 and used to obtain 2D CC distributed maps. Both

sequences acquired single channel maps only (i.e. encoding matrix equal to the identity matrix).

6.2.4 Image Reconstruction

Image reconstruction was performed offline in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States) using the raw TWIX data format. Binary files (*.dat*) were read-in using an open-source parsing function (*mapVBVD.m*) [207]. The k -space data was Hann filtered and zero-padded to 64×64 . The Tx coil sensitivity profiles were then calculated using Tx-ESPIRiT from the TxLR toolbox [69] which helps to produce smooth maps without holes (although this was not a problem for this experiment). For relative maps acquired with the "tfl_rfmap" sequence, where single-channel maps are by default acquired with CP⁺ mode 45° phase offsets, the CP⁺ mode phase offsets were removed. The relative maps were then masked using a simple value-thresholding approach informed from the CW shot-wise images and the same mask was applied to all data.

6.2.5 Phantom Experiments

On the scanner, single-slice 2D CW shot-wise, CW single shot and CC acquisitions were acquired in a water phantom transversely in CP⁺ mode. All protocols were repeated with four nominal excitation FAs of either 2°, 5°, 10°, 15° and 20°. It is worth noting that although these nominal FAs were used, FAs closest to coil elements could be an order of magnitude higher. Three iterative B₀ shims were performed prior to imaging and full relaxation was allowed between every acquisition by waiting a minimum of 15 s. All sequence variants acquired a fully-sampled scan matrix of 32×32 with a FOV of $192 \times 192 \text{ mm}^2$ and slice thickness of 6 mm hence an isotropic resolution of 6 mm^3 . The sequences used identical sinc pulses for excitation with TE = 1.87 ms, TR = 3.9 ms, RF spoiling and 100 dummy RF pulses (per channel only for the CC acquisition). All acquisitions used a 297 V reference voltage (channel combined), a bandwidth of 500 Hz px⁻¹ and a linear ascending PE direction. The CC scheme was the only sequence acquired with a distributed AO.

The shot-wise acquisition had a TR_{shot} of 4 s which lead to a duration of 32 s. The single shot acquisition had a 1.4 s duration and the CC acquisition had a duration of 4.1 s (including 3.1 s of dummy RF pulses). The maps were compared to a reference low FA shot-wise acquisition which had a long TR_{shot} of 10 s and a duration of 80 s. Although the saturation effects for this sequence are expected to be low between shots, it is not a perfect reference as saturation will occur within the shot during the TurboFLASH readout train.

6.2.6 *In Vivo* Experiments

The sequences used in the phantom experiment were repeated in the brain of one subject with identical parameters except from a 284 V reference voltage (channel combined).

6.3 Results

6.3.1 Numerical Simulation Results

Figure 6.5 shows the longitudinal magnetisation in 10 voxels in the synthetic body model for 3D sequences. Figure 6.5A) shows the CW sequential acquisition of relative maps leads to inconsistent magnetisation. Whereas Figure 6.5B) shows the magnetisation reaches a steady state after approximately 2 s for the CC scheme (hence the dummy excitations).




Figure 6.5: Longitudinal magnetisation of tracked voxels in the Duke synthetic body for both 3D **A)** channel-wise sequential and **B)** coil-cycled distributed sequences. Colours represent the different tracked voxels shown in Figure 6.4. The red line represents a voxel tracked in the myocardium and the orange line represents a voxel tracked in the blood pool of the heart.

The impact of a CW versus CC acquisition on the measured 3D relative B_1^+ maps is shown in Figure 6.6. The percentage difference between the ground truth simulated B_1^+ fields and the simulated relative B_1^+ maps is shown in sub figures 6.6C) and 6.6D) for CW and CC distributed acquisitions, respectively. A CW acquisition lead to the largest percentage error where the Tx sensitivity is highest due to a larger saturation effect. However, there is still considerable error within

the heart. There is also a noticeable T_1 dependence to the error. Whereas CC distributed maps show little error with the majority of error only due to the fixed level of zero-mean Gaussian noise.




Figure 6.6: Relative maps acquired in the Duke synthetic body model **A)** channel-wise sequential and **B)** coil-cycled distributed using the simulated 3D GRE sequence. The relative error for each acquisition of shown in **C)** and **D)**.

The effect of sequential and distributed acquisitions is shown in Figure 6.7 for CW shot-wise and CC schemes, for four different sequence and coil configurations. A distributed AO had a negligible effect on both CW and CC schemes across all sequence and coil configurations that were simulated. The main improvement in NRMSE comes from coil-cycling which consistently shows the lowest NRMSE when compared to CW acquisitions and excellent linearity. This figure also demonstrates the significant overlapping of coil sensitivities for the head coil with the maximum relative B_1^+ contribution of 45% for the phantom maps versus 83% for the synthetic body maps using the body dipole array.

The effect of nominal FA on the NRMSE for the four sequence/coil configurations simulated are shown in Figure 6.8, where each data point represents the NRMSE from one of the eight relative maps. Across all configurations investigated CW (purple) acquisitions perform worse for an increasing FA, due to increased saturation effects. Of all the configurations simulated CW performed the least worst for a 2D multi-slice sequence in phantom. This is likely due to the extent to which the Tx sensitivities overlap using the head coil reducing the variation in longitudinal magnetisation between CW acquisitions. Both CW and CC (green) sequences show an increase in NRMSE at low nominal FAs due to noise, with the 2D multi-slice sequences showing a larger increase than the 3D sequences due to the inherent improvement in SNR when imaging in 3D. The CC maps show significantly lower NRMSE and only begin to show an increase in NRMSE for the 3D Duke configuration in 6.8C) above 10° where the signal deviates from the linear relationship to FA. Analysis of the SNR found no significant difference between sequential and distributed AOs.

The effect of T_1 on the NRMSE for the four sequence/coil configurations simulated are shown in Figure 6.9, for a nominal FA of 5° , where each data point represents the NRMSE from one of the eight relative maps. For CC acquisitions the NRMSE does not vary with T_1 due to imaging in a steady-state. This is not the case for the channel-wise acquisitions, where there is a large T_1 dependent bias in the NRMSE.

The effect of TR on the NRMSE for the four sequence/coil configurations simulated are shown in Figure 6.10, where each data point represents the NRMSE from one of the eight relative maps. The channel-wise NRMSE decreases rapidly for longer TRs in the 3D sequences whereas longer TRs have little impact on CC acquisitions.




Figure 6.7: The impact of distributing the acquisition order of transmit channel excitations. Correlation plots of simulated B_1^+ versus ground-truth for four different acquisition schemes and four different sequence/coil configurations using a nominal flip angle of 5° . **A)** 2D multi-slice Duke, **B)** 2D multi-slice phantom, **C)** 3D Duke and **D)** 3D Phantom. **i)** channel-wise sequential, **ii)** channel-wise distributed, **iii)** coil-cycled sequential and **iv)** coil-cycled distributed. Sequential maps (red) had an acquisition order of $[1,2,3,4,5,6,7,8]$ whereas distributed maps (blue) had an acquisition order of $[1,4,7,2,5,8,3,6]$.




Figure 6.8: NRMSE in relative maps versus flip angle. Data shows channel-wise and coil-cycled NRMSE for four different sequence and coil setups. **A)** 2D multi-slice (Duke), **B)** 2D multi-slice (Phantom), **C)** 3D (Duke), **D)** 3D (Phantom).

Channel-wise use a sequential acquisition order and coil-cycled use a distributed acquisition order. For the Duke dataset, whole-body masked NRMSE is indicated by a dot ($\bullet$) and heart masked NRMSE is indicated by a circle ($\circ$). Each data point represents the NRMSE from one transmit channel. Phantom used fixed $T_1 = 1$ s.




Figure 6.9: NRMSE in relative maps versus T_1 . Data shows channel-wise and coil-cycled NRMSE for four different sequence and coil setups. **A)** 2D multi-slice (Duke), **B)** 2D multi-slice (Phantom), **C)** 3D (Duke), **D)** 3D (Phantom).

Channel-wise use a sequential acquisition order and coil-cycled use a distributed acquisition order. For the Duke dataset, whole-body masked NRMSE is indicated by a dot ($\bullet$) and heart masked NRMSE is indicated by a circle ($\circ$). Each data point represents the NRMSE from one transmit channel. Nominal flip angle of 5° .




Figure 6.10: NRMSE in relative maps versus TR. Data shows channel-wise and coil-cycled NRMSE for four different sequence and coil setups. **A)** 2D multi-slice (Duke), **B)** 2D multi-slice (Phantom), **C)** 3D (Duke), **D)** 3D (Phantom).

Channel-wise use a sequential acquisition order and coil-cycled use a distributed acquisition order. For the Duke dataset, whole-body masked NRMSE is indicated by a dot ($\bullet$) and heart masked NRMSE is indicated by a circle ($\circ$). Each data point represents the NRMSE from one transmit channel. Phantom used fixed $T_1 = 1$ s.

6.3.2 Phantom Results

The NRMSE of CW acquisitions to the CC acquisition is shown in Table 6.1 for a water phantom. Increasing the nominal FA leads to an increase in the NRMSE for all acquisitions, as expected from simulations, since larger FAs perturb the longitudinal magnetisation to a greater extent. The single shot acquisitions show a larger error due to the lack of an approximately 4 s delay, which is typical of the conventional approach. The CW shot-wise scheme shows the lowest error, but this is likely due to the reference used which although has a long delay between shots, will still suffer from saturation effects during the TurboFLASH readout.

Nominal FA [°]	Relative Map NRMSE to Reference [%]		
	CW Shot-wise (TA = 32 s)	CW Single Shot (TA = 1.4 s)	Coil-cycled (TA = 4.1 s)
2	0.6	1.6	2.4
5	0.9	1.9	2.5
10	1.4	4.4	2.9
15	2.3	9.4	3.4
20	3.9	15.9	4.1

Table 6.1: The NRMSE of relative maps acquired in a water phantom shot-wise (TR = 4 s) and single shot (both sequentially) as well as coil-cycled distributed acquisition across nominal flip angles of 2°, 5°, 10°, 15° and 20°. Reference was a shot-wise acquisition with a TR of 10 s. FA; flip angle, CW; channel-wise, TA; acquisition time.

Figure 6.11 shows the relative maps for the three sequence variants at a nominal FA of 20° in the water phantom. Whilst all maps appear visually similar, there is some obvious hyper-intensity for the single shot acquisition close to the Tx element. The difference between the shot-wise and single shot acquisitions to the CC acquisition is shown in the bottom two rows. For single shot mapping, there are underestimations (blue) where the Tx sensitivity was high for the previous coil mapping as well as overestimations of the current Tx element. Hence, the first coil (Tx 1) only shows overestimation (red). The same is true for the shot-wise mapping but to a much lesser extent due to the relaxation period.

Figure 6.12 shows the SD across the range of nominal FAs for each sequence variant. CC acquisitions were consistent across nominal FAs as demonstrated by the low SD. On the other hand, the CW acquisitions share a similar difference profile across all nominal FAs, indicating a systematic error due to magnetisation history effects, as expected.




Figure 6.11: Relative maps and difference to the coil-cycled acquisition for a nominal FA of 20° in a water phantom. CW; channel-wise, SW; shot-wise (32 s), SS; single-shot (1.4 s), CC; coil-cycled (4.1 s).




Figure 6.12: The standard deviation in relative maps across nominal FAs of 5° , 10° , 15° and 20° . CW; channel-wise, CC; coil-cycled.

6.3.3 *In Vivo* Results

Figure 6.13 shows the relative maps for the three sequence variants at a nominal FA of 20° in the brain. At this nominal FA, the error in the single shot maps is also obvious. It is noted that for the first few Tx channels maps are mostly overestimated (red) for the channel-wise single shot acquisition whereas the final few Tx channels are mostly underestimated (blue) when compared to the CC acquisition.




Figure 6.13: Relative maps acquired in the brain and difference to the coil-cycled acquisition for a nominal FA of 20° . CW; channel-wise, SW; shot-wise (32 s), SS; single-shot (1.4 s), CC; coil-cycled (4.1 s).

6.4 Discussion

In this work I have demonstrated through EPG simulations the impact of not coil-cycling the acquisition of relative maps using small FA GREs. The issue is caused by the longitudinal magnetisation changing considerably between acquisitions. Hence, unless more dummy pulses are performed to reach a steady state, a channel-wise acquisition will not provide accurate or precise relative magnitudes. This means that a distributed AO over a sequential acquisition order makes little difference as a steady-state is still not achieved. The improvement comes from reordering the acquisition and acquiring one k -space PE line on a particular Tx channel before moving to the next channel and repeating until all PE lines have been

acquired (i.e. coil-cycling). Coil-cycling allows a steady-state to be achieved and ensures that Tx sensitivity and/or T_1 does not bias the relative maps. Therefore, implementing a CC acquisition allows for accurate and precise 3D relative magnitude maps. This is beneficial as 3D sequences inherently have higher SNR than a 2D multi-slice sequence and allows for higher acceleration as undersampling can be performed in both PE directions.

I found coil-cycling to be especially important when acquiring 3D relative maps and/or imaging over large FOVs (i.e. body imaging) where the B_1^+ in a particular voxel largely originates from a single coil rather than being more of a combination from multiple coils (as is the case in the head coil). Performing 100 dummy pulses per Tx channel, totalling 2.5 s ($8 \times 100 \times 3.14$ ms), was found to be sufficient to provide an adequate steady state for imaging. Coil-cycling appears less essential (although much faster) for 2D CW acquisitions and this is in agreement with Brunheim et al. who found no observable difference between a 2D CW and a CC acquisition [107]. This is likely due to the relatively low simulated resolution ($N_{PE} = 32$) hence less saturation since the magnetisation experiences fewer RF pulses. I expect coil-cycling to be more important for higher resolution 2D relative magnitude maps or when using larger FAs, as demonstrated in simulations.

The errors due to a CW acquisition are worst close to the Tx elements but still propagate into the ROI. The percentage error is approximately the same within the heart as in the rest of the body as demonstrated by Figure 6.6. The errors from the sequential acquisition could be reduced by decreasing the FA, as seen in Figure 6.8, but this would lead to a reduction in the SNR, and hence DR. Alternatively, the TR could be increased, as seen in Figure 6.10, but this results in an undesirably long acquisition time.

I have also demonstrated the results from simulations in a phantom and brain experiment. Single slice 2D relative maps were acquired which are expected to show the least sensitivity to these effects, however, the errors were still present and, as expected, most obvious for the single shot acquisition at a nominal FA of 20° . Although the errors in the channel-wise single shot method show slightly different behaviour in simulations versus in phantom/*in vivo* this is likely due to the different experimental setup and sequence parameters. For example; slight differences in sequence timings, the number and timings of the dummy scans, the nominal FA as well as the T_1 of the phantom differing to what was simulated (1.5 s). Additionally, the use of the long TR channel-wise shot-wise acquisition as a reference is not ideal.

In this chapter I have only considered single-channel mapping. Alternatively, other encoding schemes could be used such as Fourier encoding, leave-one-out etc. (see Section 3.5.4). However, all of these encoding schemes will ultimately suffer the same saturation problem. Additionally, using these encoding schemes in the presence of noise (as is always the case when performing real-world experiments) leads to an under-determined system unless more modes are acquired than the number of Tx channels. Therefore, I believe that these encoding methods are best kept for absolute mapping, but I expect the benefits of coil-cycling demonstrated here would also apply to these other acquisition schemes.

6.5 Summary

Coil-cycling is *essential* for acquiring accurate and precise relative magnitude maps for multi-Tx arrays using a small FA GRE sequence. Coil-cycling is optimal to ensure a consistent steady-state during imaging and helps to reduce the acquisition time by shortening the TR. I propose the two conditions for acquiring relative maps should instead be:

1. Small FAs ($FA \leq 10^\circ$) so that signal is linearly proportional to FA.
2. Acquire relative maps coil-cycled to remove magnetisation history effects.

7

Accelerated pTx B_1^+ Mapping

Contents

7.1	Introduction	139
7.2	Methods	141
7.2.1	Acquisition Scheme	141
7.2.2	Numerical Heart Simulations	141
7.2.3	Parallel Transmit MRI System	144
7.2.4	Retrospective Undersampling of Brain Data	144
7.2.5	Sequence Implementation	145
7.2.6	Image Reconstruction	148
7.2.7	Phantom Experiments	148
7.2.8	<i>In Vivo</i> Experiments	149
7.3	Results	149
7.3.1	Numerical Heart Simulation Results	149
7.3.2	Retrospective Undersampling of Brain Data Results	154
7.3.3	Phantom Results	154
7.3.4	<i>In Vivo</i> Results	159
7.4	Discussion	165
7.5	Summary	167

7.1 Introduction

UHF MRI promises superior SNR over conventional clinical MR systems [80]. However, at 7 Tesla, B_1^+ inhomogeneity is large across the brain and severe across the heart limiting the utility of ultra-high field quantitative cardiac imaging. This problem can be solved [126, 217, 218] by pTx [98, 99] but requires multi-channel B_1^+ maps for optimal RF pulse design and RF shimming. Multi-channel B_1^+ maps in the

heart are difficult to acquire in a clinically reasonable timeframe and are confounded by issues of cardiorespiratory motion, blood flow and large B_0 off-resonances. Since B_1^+ maps are generally only required for calibration before scanning, their acquisition should be as fast as possible and ideally well under a minute. While accurate maps are essential, high-resolution maps are not usually necessary since the wavelength associated with the 1H Larmor frequency is approximately 13 cm at 7 T in the human brain. Free-breathing approaches have been demonstrated [219] but these require acquisition and reconstruction times of several minutes. The efficacy of measuring B_1^+ maps during diastole and applying the calculated dynamic RF pulses throughout the entire cardiac cycle has been shown [217].

Relative B_1^+ maps are generally fast to acquire taking less time than absolute B_1^+ mapping [132]. Recent approaches have shown the ability to estimate relative maps from localisers [134], but to achieve a specific FA an absolute measure of B_1^+ is required. Absolute B_1^+ mapping using 3DREAM [163–165] works well in the head and can acquire high-resolution 8-channel B_1^+ maps in approximately one minute. However, DREAM [152, 153] based B_1^+ mapping is undesirable for use in the heart due to sensitivity to flow [159, 160]. Presaturation TurboFLASH [171] (i.e. SatTFL) is ideal for B_1^+ mapping in the heart due to its inherent insensitivity to flow but has a strong T_1 -dependency [210]. Additionally, SatTFL suffers from long acquisition times due to the long time delay of $5 T_1$ required between the reference and prepared images. In Chapter 4 I showed that using short imaging trains and ‘sandwiching’ the reference and prepared images together overcomes both of these problems. The use of a non-adiabatic HS8 pulse for presaturation also achieves better B_0 insensitivity over a conventional sinc or rectangular preparation pulse thereby requiring no additional B_0 correction [206]. These simple pulse sequence modifications allowed for faster acquisition of the absolute B_1^+ maps and, as a result, in Chapter 4 I demonstrated single breath-hold whole-heart 3D B_1^+ mapping for a single shim mode.

Recently, an acquisition strategy known as B1TIAMO [106, 107, 203] was proposed which acquires fast relative multi-channel maps and two absolute B_1^+ maps with complementary shim modes. The two absolute shim mode maps are used to help overcome the limited dynamic range by ensuring sufficient B_1^+ in at least one of the shim modes and to update the multi-channel relative maps to absolute units. B1TIAMO has been shown to obtain accurate 2D multi-channel maps for body imaging in 16 seconds for a 32 transmit channel system [107]. Furthermore, Hess et al. recently proposed an algorithm for reconstructing undersampled calibrationless

transmit field maps (transmit low rank; TxLR) [69] that could allow for acceleration factors of up to 8 in the body, matching the number of transmit channels. Hess et al. also showed that the minimum matrix size for capturing the dominant transmit modes in the heart was 24×24 . Hence, TxLR enables 8-channel transmit 3D relative maps with $24 \times 24 \times 8$ lines of k -space data to be acquired in only 2 s or 5 heartbeats ($TR = 3.5$ ms, cardiac window = 400 ms).

Therefore, in this chapter, I evaluated TxLR in combination with a B1TIAMO-style acquisition and fast absolute maps using the Sandwich method [3] to acquire accelerated 3D multi-channel absolute B_1^+ maps in 14 s in the brain or 23 heartbeats in the body [7]. I evaluated this approach using: *in silico* simulations; retrospective undersampling of 3D *in vivo* whole-head data from a library of 9 subjects; and finally prospectively undersampled data acquired in phantoms, the brain of three subjects and the heart of another three subjects. This study investigates a method to accelerate maps and is not proposing a new B_1^+ mapping method. Hence, to avoid repetition of previous work, the evaluation of the proposed method is largely focused on comparing to fully-sampled data.

7.2 Methods

7.2.1 Acquisition Scheme

The proposed acquisition scheme is shown in Figure 7.1. In a B1TIAMO [107, 203] fashion, 3D coil-cycled multi-channel relative maps are acquired using a spoiled low FA GRE sequence followed by two absolute maps with complementary orthogonal shim modes (first and second order circularly polarised, CP^+ and CP^{2+} respectively) using the modified Sandwich presaturation TurboFLASH sequence [3].

7.2.2 Numerical Heart Simulations

Simulations were performed using retrospectively undersampled *in silico* data to establish the achievable acceleration. The *in silico* body $B_1^{+/-}$ fields were simulated in Sim4Life 3.4 (ZMT, Zurich, Switzerland) using an 8-channel transmit/receive dipole array [216] centred over the heart of Duke (Virtual Population, ITIS Foundation, Zurich, Switzerland) at 7 T [69]. A synthetic proton density image was generated by setting the proton density equal to the tissue density, with tissue densities greater than 1200 kg m^{-3} (bone) or less than 400 kg m^{-3} (lung) set to 80 to make them resemble an MR image. Voxels containing no anatomy were set to zero. The




Figure 7.1: Schematic sequence diagram for acquiring fast multi-channel B_1^+ maps in the brain/heart. Firstly, multi-channel relative maps (green) are acquired for a number of coil-cycles (N_{cc}) including 16 dummy coil-cycles. This is followed by absolute maps (red) using Sandwich for complementary shim modes - first and second-order circularly polarised (CP^+ , CP^{2+}) modes with a TurboFactor of N . The absolute maps excitation (β) and 5 ms non-adiabatic hyperbolic secant (HS8) preparation pulse (α) have matched transmit mode. The sequence can be adapted for cardiac B_1^+ mapping by segmenting the relative maps into 9 HBs and acquiring the absolute maps with $TR_{long} = 1$ HB. Brackets indicate repeated sequence blocks. TA, acquisition time; HB, heartbeat; TD, time delay; R, acceleration factor.

FOV was $278 \times 356 \text{ mm}^2$ (AP/LR). Data were re-sampled onto a uniform grid of $2 \times 2 \text{ mm}^2$. The relative ($N_{Tx} = 8$) and reference ($N_{Tx} = 2$) images were formed from a multiplication of these proton density weighted images and either individual transmit B_1^+ or CP^+ and CP^{2+} . The prepared ($N_{Tx} = 2$) images are formed from the multiplication of the reference images with the cosine of the B_1^+ for CP^+ and CP^{2+} . Receive channel images were then formed from the multiplication of the relative, reference and prepared images by the synthetic coil-sensitivities. The synthetic (receive coil-combined) images are shown in Figure 7.2.

The synthetic body images were fast Fourier transformed and k -space corrupted with channel-independent complex Gaussian noise (peak SNR 60 dB). The k -space data was cropped to a matrix size of 24×24 ; which was previously found to be the minimum scan matrix for capturing the dominant transmit modes in the heart [69]. This was re-validated by comparing the relative error between k -space cropped (24×24 , Hann filtered and zero-padded back to 139×178) and native resolution (139×178) ground truth relative and absolute maps.

Homogenous Poisson-disc undersampling masks were generated using the iterative self-consistent parallel imaging reconstruction (SPIRiT) Toolbox v0.3 for effective undersampling factors between 1 and 10. To evaluate the impact of the absence of a calibration region, masks were also generated both with and without a central 4×4 calibration region. A different undersampling mask was used for each transmit




Figure 7.2: The fully sampled native resolution (139×178) synthetic images (coil-combined) used to test various acceleration factors for an 8-transmit 8-receive dipole array chest coil at 7T.

mode, with the reference and prepared images of the absolute maps using the same undersampling mask.

The relative ($N_{Tx} = 8$), reference ($N_{Tx} = 2$) and prepared ($N_{Tx} = 2$) undersampled k -spaces were concatenated to form a $24 \times 24 \times 8 \times 12$ matrix ($N_{k_x} \times N_{k_y} \times N_{Rx} \times N_{Tx}$) and jointly reconstructed using the TxLR algorithm (5×5 kernel, 50 iterations and rank threshold of 50). The reconstructed k -space data were Hann filtered, zero-padded up to 139×178 and inverse fast Fourier transformed back to image space. The receive channel images were combined using receive sensitivity maps estimated from the relative map reconstructed k -space using ESPIRiT (5×5 kernel, 0.02 eigen threshold) [51]. ESPIRiT was also used to calculate the transmit sensitivities. The FA for the absolute maps was found from the arccosine of the reference and prepared image ratio.

The individual multi-channel maps were then calculated based on the method proposed by Brunheim et al. with a total scaling exponent $m = 4$, as defined in the original B1TIAMO paper [107], and weighted based on the reference image intensity. Maps were masked from a simple binary value-thresholding based on the fully-sampled relative images. Simulations were repeated 50 times with each repetition using unique noise and undersampling masks. To assess their accuracy within the heart, the difference in the reconstructed multi-channel maps to the

fully-sampled native resolution (noiseless) ground truth is shown as well as the RMSE for maps combined into CP^+ mode.

7.2.3 Parallel Transmit MRI System

A 7 Tesla MRI (Magnetom VB17 Step 2.3, Siemens Healthcare Ltd, Germany) equipped with parallel transmit was used for phantom and *in vivo* imaging. Brain images were acquired using an 8Tx/32Rx head coil (Nova Medical, USA) using fixed per-channel power limits as provided by the coil manufacturer of 1.5 W per channel. To image the heart an 8Tx/8Rx dipole array (MR Coils BV, The Netherlands) was used with internally approved power limits based on simulation and measurement for a worst-case per channel limit of 3.98 W. Cardiac triggering was performed using a pulse oximetre (Siemens Healthcare Ltd, Germany) with a trigger delay of 100 ms.

7.2.4 Retrospective Undersampling of Brain Data

Using a similar processing pipeline as outlined above, fully-sampled 3D relative multi-channel maps and absolute B_1^+ maps for CP^+ and CP^{2+} in the brain were retrospectively undersampled from a database of 9 subjects (23-56 years old; 7 male/2 female) acquired by de Buck et al. [196, 214]. This database was acquired using a similar sequence as outlined in the following sections but with the following sequence parameters. Both relative and absolute sequences used a $300 \times 225 \times 225 \text{ mm}^3$ FOV with a reference voltage of 50 V (per channel).

The relative maps were acquired using a GRE sequence with $TE_{\text{GRE}} = 1.02 \text{ ms}$, $TR_{\text{GRE}} = 2.90 \text{ ms}$, bandwidth = 500 Hz px^{-1} and RF spoiling. A scan matrix of $40 \times 36 \times 36$ and, non-selective rectangular RF excitation with a nominal FA of 7° . The scan time required for this relative sequence was 30 seconds.

Absolute maps were acquired for two shim modes, which were CP^+ and CP^{2+} . TurboFLASH $TE_{\text{FLASH}} = 1.76 \text{ ms}$, TurboFLASH $TR_{\text{FLASH}} = 3.92 \text{ ms}$, $TR_{\text{long}} = 1 \text{ second}$, bandwidth = 489 Hz px^{-1} , RF spoiling. A scan matrix of $48 \times 27 \times 36$ was used with non-selective rectangular RF excitation (β). A $500 \mu\text{s}$ rectangular RF pulse was used for presaturation (α) (whereas a HS8 pulse is used in the current work) with an $\alpha : \beta$ ratio of 10:1, using nominal α/β FAs $90^\circ/9^\circ$. The scan time required for each shim mode map was 36 seconds for a combined scan time of 1 minute 42 seconds.

7.2.5 Sequence Implementation

Separate sequences were used to acquire the relative and absolute portions of the proposed method. 3D accelerated relative maps were acquired using a GRE sequence, coil-cycled to minimise the influence of magnetisation history (see Chapter 6), with $TE_{GRE} = 1.31$ ms, $TR_{GRE} = 3.14$ ms, bandwidth = 490 Hz px⁻¹, RF spoiling and an acceleration factor of $R = 4$. A scan matrix of $24 \times 24 \times 24$ was used, non-selective rectangular RF excitation with a FA (λ) of 9°. The acquisition time was 4 seconds ($N_{cc} = 160$) of which the first 400 ms were 16 dummy coil-cycles. Undersampling masks were as described previously (see Section 7.2.2) and ordered as in Jaeschke et al.[220] to minimise the eddy currents from a rapid change in the phase encoding. For cardiac mapping during diastole, the relative map k -space was segmented into 9 heartbeats of 16 coil-cycles (400 ms). During the trigger delay, 4 dummy coil-cycles (100 ms) were performed ($N_{cc} = 180$).

3D accelerated absolute maps were acquired for CP⁺ and CP²⁺ using the Sandwich sequence [3] with TurboFLASH $TE_{FLASH} = 1.26$ ms, TurboFLASH $TR_{FLASH} = 3$ ms, bandwidth = 490 Hz px⁻¹ and RF spoiling. A scan matrix of $24 \times 24 \times 24$ was used with non-selective rectangular RF excitation (β). A 5 ms non-adiabatic hyperbolic secant (HS8) [36, 206] RF pulse was used for presaturation (α) using nominal α/β FAs 90°/5°. Different acceleration factors were used in the brain/heart for the acquisition of the absolute maps which were $R = 6/4$, respectively. Each absolute map was segmented into 4/6 ($R = 6/4$), one segment per $TR_{long} = 1$ second, to maintain a readout train of 24 phase encoding lines (TurboFactor = 24, duration of 145 ms) per segment. The acquisition time was 5 seconds ($R = 6$) or 7 heartbeats ($R = 4$) per map, which includes 1 dummy TR_{long} . Hence, a total of 10 seconds or 14 heartbeats for both mode maps. Undersampling masks were as described previously (see Section 7.2.2) and shown in Figure 7.3). A square-spiral acquisition order was used, centre-in for the reference images and centre-out for the prepared images to reduce the impact of T_1 on the image trains.

Although separate sequences were used, the total time to acquire both the relative and absolute maps is 14 seconds ungated or 23 heartbeats gated. Table 7.1 summarises the differences in protocols used in phantom and in vivo.




Figure 7.3: The undersampling masks used for **a)** relative maps with an acceleration factor $R = 4$ **b)** absolute maps in the brain with an acceleration factor $R = 6$ and **c)** absolute maps in the heart with an acceleration factor $R = 4$. The colours indicate the segment during which the readout line was acquired.

Target	Coil	FOV [mm ³]	$R^\dagger$	TR_{long}	Ref. Voltage [V]	Acquisition Time
Braino Phantom	8Tx/32Rx	$250 \times 250 \times 250$	6	1 s	60	14 s
Drum Phantom	8Tx/8Rx	$300 \times 300 \times 300$	4	700 ms	200	16 s or 23 HB
Brain	8Tx/32Rx	$250 \times 250 \times 250$	6	1 s	60	14 s
Heart	8Tx/8Rx	$380 \times 380 \times 285$	4	1 HB	200	23 HB

Table 7.1: Differences in sequence protocol parameters used *in vivo* and in phantom to acquire the accelerated multi-channel B_1^+ maps. † Undersampling factor for absolute maps.

7.2.6 Image Reconstruction

Reconstruction was performed offline in MATLAB (R2021a, MathWorks, Natick, Massachusetts, United States) using the parallel computing toolbox (14 workers) on an Intel Xeon 28×2.40 GHz 128 Gb RAM computer running Rocky Linux 8.5. Relative, reference and prepared k -space data were concatenated to form a $24 \times 24 \times 24 \times (8/32) \times 12$ ($N_{k_x} \times N_{k_y} \times N_{k_z} \times N_{Rx} \times N_{Tx}$) matrix and jointly reconstructed using TxLR slice-by-slice on hybrid x - k_y - k_z data. The k -space data were Tukey filtered (cosine fraction = 0.7), zero-padded to $32 \times 32 \times 32$ (unless mentioned otherwise) and inverse fast Fourier transformed to full image space. Receive channel images were combined using sensitivity maps estimated from relative map k -space using ESPIRiT. ESPIRiT was also used to calculate the transmit sensitivities.

The FA for the absolute maps was calculated from the ratio of the reference and prepared images using a lookup table generated from Bloch and EPG simulations [4], which includes a correction for the non-linearity of the HS8 pulse employed for saturation. The individual multi-channel maps were then calculated based on the method proposed in the B1TIAMO paper [107] with a total scaling exponent $m = 3$, as defined in the original B1TIAMO paper, and weighted based on the reference image intensity. Maps were masked from a simple binary value-thresholding based on the fully-sampled relative images. Voxels in the absolute maps outside of the lookup table were masked out and any holes were linearly interpolated. Voxels which had a maximum shim efficiency below 30%, as calculated from the relative maps, were additionally masked out due to low signal producing erroneous measurements of absolute B_1^+ .

7.2.7 Phantom Experiments

3D multi-transmit B_1^+ maps were acquired in two coil/phantom combinations. The first was the 8Tx/32Rx head coil using a 16.5 cm diameter spherical MRS ‘Braino’ water phantom (General Electric Medical Systems) (Reference voltage 60 V) and the second was the 8Tx/8Rx chest coil using a large custom 27 cm diameter 73 mmol NaCl drum phantom (Reference voltage 200 V). The Braino phantom was acquired with a $TR_{\text{long}} = 1$ s, using a $250 \times 250 \times 250 \text{ mm}^3$ FOV giving $10.4 \times 10.4 \times 10.4 \text{ mm}^3$ resolution. The drum phantom was acquired gated with a simulated RR interval of 700 ms. The FOV used for imaging the drum phantom was increased to $300 \times 300 \times 300 \text{ mm}^3$ giving $12.5 \times 12.5 \times 12.5 \text{ mm}^3$ resolution.

The offline reconstruction time was 5 minutes for the head coil and 1 minute for the chest coil using 20 TxLR iterations. For reference, a 3D fully-sampled B1TIAMO map was acquired with a matched scan matrix and FOV using fully-sampled relative maps (15 s or 36 simulated heartbeats) and gated Sandwich absolute CP^+ and CP^{2+} maps (50 s or 50 simulated heartbeats for both modes).

7.2.8 *In Vivo* Experiments

All subjects were scanned under an institutionally approved technical development standard operating procedure (FMRIB_004_V4). 3D multi-transmit maps were acquired in the brain of three volunteers at a reference voltage of 60 V as described above for the head coil phantom at acceleration factors of $R = 4/6$. The FOV for relative and absolute Sandwich maps was $250 \times 250 \times 250 \text{ mm}^3$ giving $10.4 \times 10.4 \times 10.4 \text{ mm}^3$ acquired resolution. The maps were reconstructed to $64 \times 64 \times 64$ for plotting. A 3D fully-sampled B1TIAMO dataset was acquired for reference in 65 seconds.

I also tested the application of three coil-compression algorithms (see Section 2.5.6) using the acquired brain data. These were geometric, single-value and ESPIRiT-based coil-compression from the SPIRiT Toolbox v0.3 before TxLR reconstruction for a range of virtual coils from 2 to 32.

3D cardiac multi-transmit maps were acquired in three healthy volunteers at a reference voltage of 200 V. The FOV for relative and absolute Sandwich maps was $380 \times 380 \times 285 \text{ mm}^3$ giving $15.8 \times 15.8 \times 11.9 \text{ mm}^3$ resolution. The maps were reconstructed to $64 \times 64 \times 48$ for plotting. The coil-cycled relative maps and absolute maps were acquired in two separate breath-holds due to the sequence preparation time, both at end-expiration to reduce respiratory motion artefacts. However, in total 23 heartbeats (9 segments relative maps + 2 Sandwich dummies + 6 segments Sandwich CP^+ + 6 segments Sandwich CP^{2+}) were required for data acquisition. A fully-sampled set of 3D B1TIAMO maps would have required 86 heartbeats.

7.3 Results

7.3.1 Numerical Heart Simulation Results

Figure 7.4 shows the impact on relative and absolute maps of cropping k -space down to a central 24×24 region compared to the native resolution (139×178) ground

truth. The NRMSE within the heart was 3.1% for relative maps and the RMSE 3.8° (mean FA 67°) for absolute maps with the largest errors close to regions of B_1^+ constructive/destructive interference. This indicates it is possible to capture the dominant transmit modes in the heart using small scan matrices.




Figure 7.4: Comparison between the magnitude of the native resolution (139×178) and k -space cropped (24×24) ground truth for relative and absolute maps (CP⁺ and CP²⁺) of synthetic body data. The mean error between the cropped and native resolution relative maps as a percentage of the total B_1^+ available is 3.1% within the region of interest (green line). The mean error between the cropped and native resolution absolute maps is 3.8° (mean FA in heart of 67°) for the absolute maps within the region of interest (green line).

Simulation results from synthetic 8Tx/8Rx body images for undersampling factors are shown in Figure 7.5. The mean FA and SD over 50 repeats each with different undersampling masks are shown in Figure 7.5a) and b). The mean FA within the heart in CP⁺ mode was 66° (min/max = $6^\circ/137^\circ$, CV = 36%). In Figure 7.5c) large undersampling factors of 10 show increased error in the measured B_1^+ within the heart when compared to the fully-sampled native resolution maps. Figure 7.5d) shows a complex RMSE to the fully-sampled native resolution ground truth of 4.2° combined in CP⁺ mode within the heart at an acceleration factor of 4. The RMSE increases linearly with undersampling factor.




Figure 7.5: Simulation results from synthetic body images showing the mean **a)** and SD **b)** for multi-channel parallel transmit B_1^+ maps with undersampling factors of 1 to 10. The mean FA in the heart in CP^+ mode was 66° (min/max = $6^\circ/137^\circ$, CV = 36%). 50 TxLR iterations and 50 repeats with different undersampling masks were performed. Subfigure **c)** shows the mean absolute error across these repeats to the fully-sampled native resolution (139×178) k -space in the ROI (red box in the top left of subfigure **a)**). **d)** RMSE calculated on the complex values for multi-channel maps combined in CP^+ mode for the 50 repeats. ROI, region of interest; RMSE; root-mean-square error.

Figure 7.6 shows that fully-sampling the central 4×4 region of k -space had a negligible impact on the error in the reconstructed maps at undersampling factors of 4. However, including a fully-sampled central 4×4 region did improve the error in the reconstructed maps at higher undersampling factors, lowering the RMSE from 7.8° to 6.1° at an acceleration factor of 8.

Figure 7.7 indicates that 20 TxLR iterations is optimal for acceleration factors up to 4 in the body, where the RMSE of the combined CP^+ map to the native resolution ground truth was 4.3° in the heart (5.0° whole-body). The fully-sampled had an RMSE of 2.5° in the heart (3.1° whole-body).




Figure 7.6: The effect of not acquiring a calibration region is shown (left) in comparison to acquiring a small 4×4 calibration region (right). Mean **a)** and SD **b)** for multi-channel parallel transmit B_1^+ maps across various undersampling factors (1-16). 50 repeats with different undersampling masks were performed. Subfigure **c)** shows the mean difference across repeats to the fully sampled k -space in a region of interest, the results of which are summarised by the RMSE **d)**. The small red box in the top left of subfigure **a)** indicates the heart ROI used for plots **c)** and **d)**.




Figure 7.7: The number of TxLR iterations versus the root-mean-square error in the combined CP^+ B_1^+ maps to the fully sampled native resolution (139×178) ground truth based on data simulated using synthetic body images. Dashed lines indicate $\pm 1.96SD$.

7.3.2 Retrospective Undersampling of Brain Data Results

The results of the retrospective undersampling of 9 subjects using the head coil are shown in Figure 7.8. The mean FA in the brain in CP^+ was 45° (min/max = $11^\circ/88^\circ$, CV = 29%) with the RMSE increasing with increased acceleration factors. Acceleration factors of 4, 6 and 8 gave an average RMSE for multi-channel maps combined into CP^+ of 1.7° , 1.9° and 2.1° , respectively.




Figure 7.8: Boxplots showing the complex root-mean-square error for channels combined into CP^+ (mean FA in CP^+ was 45° , min/max = $11^\circ/88^\circ$, CV = 29%) for various retrospective undersampling factors. *In vivo* data using a B_1^+ library of 3D brain data from 9 subjects.

7.3.3 Phantom Results

Figure 7.9 shows B_1^+ maps (relative maps $R = 4$ and absolute maps $R = 6$) obtained in the ‘Braino’ phantom and the difference compared to the ground truth fully-sampled maps. The mean FA across the entire phantom was 90° in CP^+ (min/max = $34^\circ/175^\circ$, CV = 21%) with an RMSE of 3.7° to the fully-sampled.

Figure 7.10 similarly shows the synthetically gated B_1^+ maps (relative and absolute maps $R = 4$) obtained in the large drum phantom and the difference to the ground truth fully-sampled maps. The mean FA across the entire drum phantom in CP^+ was 60° (min/max = $0^\circ/214^\circ$, CV = 45%) with an RMSE of 3.5° . A large number of phase ramps are visible in Figure 7.10 due to the high permittivity of the phantom.

There are some artefacts that appear in the B_1^+ maps of Figure 7.10 which are explained by the choice of complementary shims. Figure 7.11 shows the shim efficiencies and confirms that these errors are pronounced where the maximum shim efficiency is below 30%, as indicated by the white arrows.




Figure 7.9: 3D multi-channel maps in a spherical water phantom acquired in 14s using an 8Tx/32Rx head coil. The proposed method is shown for both magnitude **a)** and phase relative to transmit channel 1 **b)**. A slice from the fully sampled reference multi-channel maps is also shown for both magnitude **c)** and phase **d)**. The difference between a central slice (outlined in red in **a)**) is shown for both magnitude **e)** and phase **f)**.




Figure 7.10: 3D multi-channel maps in a large drum water phantom acquired gated in 23 simulated heartbeats (RR interval = 700 ms) using an 8Tx/8Rx dipole-array chest coil. The proposed method is shown for both magnitude **a)** and phase relative to transmit channel 1 **b)**. A slice from the fully sampled reference multi-channel maps is also shown for both magnitude **c)** and phase **d)**. The difference between a central slice (outlined in red in **a)**) is shown for both magnitude **e)** and phase **f)**.




Figure 7.11: The shim efficiencies are shown for a slice from the drum phantom for a) CP^+ mode, b) CP^{2+} mode and c) the maximum efficiency used for additional masking. As indicated by the white arrows, CP^+ and CP^{2+} do not appear to provide sufficient coverage in this phantom.

7.3.4 *In Vivo* Results

The magnitude and phase of the B_1^+ maps acquired in the brain are displayed in Figures 7.12(a) and 7.12(b), respectively. The mean FA in CP^+ in the brain was 86° (min/max = $19^\circ/189^\circ$, CV = 30%). Additional subjects can be seen in Figure 7.13.




Figure 7.12: 3D multi-channel and combined CP^+ B_1^+ maps acquired in the brain of one subject are shown for axial, coronal and sagittal planes for **a)** magnitude and **b)** phase. Data was acquired at a 60 V reference voltage. Maps are reconstructed to $64 \times 64 \times 64$ from an acquired scan matrix of $24 \times 24 \times 24$.

Figures 7.14(a) and 7.14(b) show a correlation and Bland-Altman plot comparing accelerated maps to the fully-sampled reference. The maps show excellent linearity and good agreement with an RMSE of 0.7° across all multichannel maps to the reference and 1.9° when combined in CP^+ .

Coil-compression reduced the reconstruction time from over 6 minutes to 80 seconds using 10 virtual coils and 20 TxLR iterations or 50 seconds for 10 TxLR iterations with only a small (0.6°) impact on RMSE to a fully-sampled non-compressed reference, as seen in Figure 7.15.




Figure 7.13: The magnitude of 3D multi-channel and combined CP^+ B_1^+ maps for three subjects acquired in the brain shown for axial, coronal and sagittal planes. Data was acquired at a 60 V reference voltage. Maps are reconstructed to $64 \times 64 \times 64$ from an acquired scan matrix of $24 \times 24 \times 24$.

Data acquired in the body of one subject is shown in Figure 7.16. The mean FA over the heart in CP^+ is 42° (min/max = $6^\circ/137^\circ$, CV = 29%). The poor cine image intensity in the right atrium is explained by the low B_1^+ magnitude of the combined CP^+ map in this region. Additional subjects can be seen in Figure 7.17.

The shim efficiency for CP^+ , CP^{2+} and the maximum of both shim modes is shown in Figure 7.18. The combination of CP^+ and CP^{2+} appears to be inadequate over the entire heart, particularly in the right atrium.




Figure 7.14: Correlation and Bland-Altman plots can be seen in **a)** and **b)**, respectively, comparing the undersampled 3D dataset to a fully-sampled reference across the entire masked volume in Figure 7.12. Individual channels (black markers) and combined CP+ (red markers) are plotted together. Statistics are shown for the individual (non-combined) maps. The dotted lines represent $\pm 1.96SD$ around the mean (solid line).




Figure 7.15: The effect of three coil-compression algorithms on the RMSE and reconstruction time for **a)** 10 TxLR iterations and **b)** 20 TxLR iterations for *in vivo* brain data acquired with an 8Tx/32Rx head coil.




Figure 7.16: 3D multi-channel and combined CP^+ B_1^+ maps for **a)** magnitude and **b)** phase acquired in the body in 23 heartbeats shown for axial, coronal and sagittal planes. Data was acquired at a 200 V reference voltage. Maps are reconstructed to $64 \times 64 \times 48$ from an acquired scan matrix of $24 \times 24 \times 24$. Body and heart boundaries are overlaid from the cine images.




Figure 7.17: The magnitude of 3D multi-channel and combined $CP^+ B_1^+$ maps for three subjects acquired in the body in 23 heartbeats shown for axial, coronal and sagittal planes. Data was acquired at a 200 V reference voltage. Maps are reconstructed to $64 \times 64 \times 48$ from an acquired scan matrix of $24 \times 24 \times 24$. Body and heart boundaries are overlaid from the cine images. Note: The placement of the posterior elements for subject 3 were unintentionally reversed.




Figure 7.18: The shim efficiencies are shown for axial, coronal and sagittal planes in the body for a) CP^+ mode, b) CP^{2+} mode and c) the maximum efficiency.

7.4 Discussion

In this chapter I have demonstrated the feasibility of accelerating the acquisition of multi-channel absolute B_1^+ maps. I did this by combining two methods from the literature (TxLR [69], B1TIAMO [107]) with my new Sandwich B_1^+ mapping sequence (see Chapter 4) and utilising small 24×24 scan matrices to derive the dominant transmit fields for pTx calibration. The use of such matrices leverages the same underlying physical principles as for receive-coil sensitivity calibration which typically relies on comparably sized scan matrices. Simulations using high-resolution synthetic body data demonstrate that the largest source of error in the cropped maps is due to these small matrices not capturing the highest frequency fluctuations in the absolute B_1^+ from constructive/destructive interference. However, this error is small with an RMSE of 3.8° in CP^+ mode with a mean FA of 66° over the heart. My simulation results using *in silico* heart data and retrospective undersampling of brain data show that it is possible to obtain reliable undersampled B_1^+ maps in the head and body using acceleration factors of 4. I also found from simulations that using a small (4×4) calibration region was important for robust B_1^+ field map estimation at larger undersampling factors.

The results in Figure 7.12 demonstrate the ability to obtain good quality multi-channel B_1^+ maps in the brain in only 14 seconds, free of any structural artefacts except where there are signal voids such as in the sinuses. These maps are in agreement with the fully-sampled reference yielding an RMSE of 1.9° combined in CP^+ and are consistent with retrospective undersampling results. Higher accelerations were achievable with the Nova (8Tx/32Rx) head coil over the dipole array (8Tx/8Rx) due to additional receive channels present. Although, higher accelerations for the relative mapping portion of the sequence were not worthwhile due to only a minor additional scan time of 1 second for $R = 4$ versus $R = 6$.

Figure 7.16 shows the multi-channel B_1^+ maps through axial, coronal and sagittal slices of the body obtained in only 23 heartbeats. A fully-sampled acquisition is unfeasible as it would require around 86 heartbeats and several breath-holds. The current implementation required the acquisition in two separate breath-holds, due to sequence preparation time. Currently, work is underway on an IDEA and ICE implementation for Oxford's updated VE12U system which will acquire both coil-cycled relative and Sandwich absolute maps within a single sequence enabling single breath-hold acquisitions.

To my knowledge, this method enables multi-channel B_1^+ mapping in the heart within a single breath-hold for the first time¹. Further pulse sequence optimisation may need to be performed as 23 heartbeats could still be too long for patients suffering from breathlessness. To overcome this, maps could be acquired in two separate breath-holds of around 17 heartbeats, where additional time could be spent mapping additional modes. The B_1^+ maps shown here were acquired in the heart of subjects with small chests and they have not been directly validated due to the lack of a reliable ground-truth in the body. The best gold standard available is *in silico* body data and the simulations I performed, along with retrospective results in the brain and prospective results in phantoms and the brain, provide us with confidence in the validity of these maps.

I expect the maps acquired using this method to be more robust to motion due to the short acquisition time during diastole. However, there are some artefacts present in the magnitude of the multi-channel maps which are present where the maximum shim efficiency of CP^+ and CP^{2+} is low. Therefore, although CP^+ and CP^{2+} appear to work well in the brain, this work suggests that CP^+ and CP^{2+} shim modes, with the dynamic range of the Sandwich method, are not optimal for 3D B_1^+ mapping in the body and an additional shim mode or additional voltages [196] may be required in the body which warrants further investigation. Alternatively, a universal shim set [123] could be used or as discussed in the original B1TIAMO paper, a dedicated heart shim could be calculated from the relative maps- but this would require additional processing time. Higher acceleration factors could be explored for cardiac coils with more than eight receive channels.

Recent work by Aghaeifar et al. has shown that RFPA drift negatively effects B_1^+ mapping sequences [221, 222]. As suggested by this work, RFPA drift could be corrected for by scaling the RF pulse amplitude within the B_1^+ mapping sequence based on the initial pulse amplitude, as measured by the DICOs, to maintain B_1^+ mapping accuracy.

This method is conceptually simple to apply to larger transmit arrays. In the brain a 16-channel transmit system would require a scan time of approximately 18 seconds and a 32-channel transmit system around 26 seconds, without considering further acceleration. For imaging the heart with larger transmit arrays more heartbeats will be needed in the relative mapping portion of the sequence and may require multiple breath-holds.

¹Since writing this thesis, a free-breathing method of acquiring accelerated B_1^+ maps in the heart was published using a 30s 3D radial density adapted AFI sequence [151].

The main limitation of this acquisition scheme is currently the long reconstruction time with the average offline reconstruction times for the chest/head coils being 1/6 minutes using current hardware. However, this has not been optimised nor is it prohibitive and there are several potential avenues to reduce this. One such possibility is the use of coil compression [64]. Initial tests with a geometric-decomposition coil compression algorithm [62] were able to reduce brain reconstruction times from 6 minutes to around one minute (Figure 7.15). Alternatively, a more efficient method of computing the SVD which accounts for a significant portion of the computation time in the TxLR algorithm or a machine learning approach could be investigated.

7.5 Summary

In this chapter I have demonstrated a method to accelerate the acquisition of 3D multi-channel B_1^+ maps that is fast and has good agreement with a fully-sampled reference in the brain. This method was applied to obtain 3D multi-channel maps in the body and is important to help facilitate cardiac imaging at 7 T.

8

RF Sensors for Motion Detection

Contents

8.1	Introduction	168
8.2	Background of Motion Detection	169
8.2.1	Methods of Motion Detection	169
8.2.2	RF Sensor-based Methods of Motion Detection	170
8.2.2.1	pTx Scattering	171
8.2.2.2	Pilot Tone	174
8.2.3	Motivation and Aims	175
8.3	Methods	176
8.3.1	Sequence Implementation	176
8.3.2	Hardware and Sequence Adaptions for RF Sensors	177
8.3.3	RF Sensor Signal Processing	178
8.3.4	Image Reconstruction	181
8.3.5	Post-processing	181
8.4	Results	182
8.5	Discussion	186
8.6	Summary	188

8.1 Introduction

MRI inherently takes time to acquire images. Hence, MRI is inherently sensitive to motion which occurs during the acquisition. This also explains why motion artefacts are usually more prominent in the slower, phase encode direction. Therefore, motion corrupts MR images acquired at any field strength, but motion has the potential to most severely corrupt images acquired at UHF due to their typically higher resolution. There are many different kinds of motion (continuous versus step-wise, periodic

versus spontaneous, rigid versus elastic, etc.) and realistic motion is not limited to a single type or in a single degree of freedom (DOF) but is always a combination.

The largest and most frequent motion typically occurs in the youngest and oldest subjects, which can ruin the diagnostic quality of images. However, even cooperative subjects can cause involuntary motion, due to breathing, swallowing, blinking or just their cardiac cycle. Sometimes sedation is necessary, but it is expensive, invasive, and not acceptable for research scenarios. Therefore, a technological solution is sought to overcome motion and is currently an active area of research. As summarised by Zaitsev et al. [223] motion artefacts in MRI are a "complex problem with many partial solutions".

8.2 Background of Motion Detection

The effect of motion on k -space is complicated. Additionally, motion also causes changes in B_0 , Rx and Tx sensitivity, and, thus motion degrades the performance of pTx pulses. However, if the subject motion is known, then the initial B_1^+ maps without motion can be updated leading to improved pTx performance [224, 225].

8.2.1 Methods of Motion Detection

Many methods of motion detection have been proposed in the literature and I will briefly introduce a few different methods here.

Optical methods using cameras and markers can be used to monitor patient motion with impressive accuracy and low latency [226]. These markers can be either retro-reflective [227], moiré phase patterns [228] or facial features [229]. Alternatively, active NMR field markers placed on the subject may be tracked, which do not require a line of sight between the marker and sensor [230]. These approaches come at the cost of requiring calibration, additional hardware (which can be intrusive and expensive), and in the case of active markers there is additional complexity involved in implementing into the pulse sequence so as to only be active during the ‘quiet’ periods of a scan.

Perhaps the most commonly used are navigators. Navigators can be implemented into sequences using a sequence building block with a duration of a few hundred milliseconds and therefore enabling periodic motion corrections. One well-known navigator is known as the cloverleaf navigator [231]. Navigators require no additional hardware, but may disrupt sequence timings and affect the magnetisation. To avoid altering the ^1H water signal, Fat navigators (FatNavs) have been used for retrospective motion correction in the brain [232]. FatNavs additionally benefit from the natural sparsity of fat, as it exists largely as sub-cutaneous fat and in the bone-marrow of the skull, leading to the possibility of acquiring highly accelerated acquisitions. Another promising new method for head motion detection is known as a FID navigator which uses the motion information encoded by having multiple receivers surrounding the head [233–235]. A FID is received on every receive channel and motion is estimated from a head model which is generated from simulations. The benefit of this method is FIDs can be acquired very quickly therefore having minimal impact on the sequence timing.

Distributed and Incoherent Sample Orders for Reconstruction Deblurring using Encoding Redundancy (DISORDER) [236] is a self-navigated method which requires no external sensors, navigators or training and can be applied to 3D sequences. DISORDER works by reordering the phase encoding lines of each shot to be uniformly distributed across k -space. This makes the sequence maximally sensitive to motion and redundancy in 3D encoding provides improvement in reconstruction convergence. Alternatively, data-driven methods like Scout Accelerated Motion Estimation and Reduction (SAMER) [237] acquire a quick 3 s scout at the beginning of the scan to correct for motion in all shots.

A common disadvantage of all of these methods is the added complexity through either additional hardware or sequence modifications. Another category of motion detection methods, which drastically reduce this complexity are RF sensors, which I will cover in the next section.

8.2.2 RF Sensor-based Methods of Motion Detection

RF sensor-based methods of motion detection utilise the electrical coupling between a transmitter and a receiver and operate far off-resonance so as to not affect the magnetisation. The benefits of RF sensor-based methods of motion detection are:

- Little-to-no additional hardware is required which can be costly, have difficult calibration and therefore this simplifies implementation.

- Motion estimates are obtained with a high temporal resolution, which are inherently synchronised with data acquisition.
- Can be applied either prospectively or retrospectively.
- Completely independent of sequence therefore no impact of sequence or additional time required and enables the use any sampling trajectory.
- Uses electrical (not MR) signal hence does not affect magnetisation.

However, RF sensor-based signals do not simply yield motion parameters (3 translations and 3 rotations for rigid body motion). If both RF sensor signals and six degrees of freedom motion parameters (e.g. via registration) are acquired at the start of the scan a statistical model can be trained to map between the measured RF sensor signals and the motion parameters. The RF sensor signals can then be used throughout the scan to predict the motion parameters. The two most promising RF sensor-based methods of head motion detection are parallel transmit scattering (pTxS) and pilot tone (PT) and I will introduce both of these methods in the subsequent sections.

8.2.2.1 pTx Scattering

In 1988 Buikman et al. [238] published an RF sensor method of motion detection making use of the body coil in a 2 T MRI. The work of Hess et al. [239–242] built on this by measuring the scattering matrix of a pTx array through utilising the DICOs already present in the scanner, illustrated in Figure 8.1. pTxS requires no additional hardware and obtains scattering matrix measurements every RF pulse. This information is primarily used for monitoring SAR and consists of a digitised and highly sampled (1 MHz) timeseries of forward, $\vec{v}_{fwd}(t)$, and returned voltages, $\vec{v}_{ret}(t)$, from each of the Tx channels, forming an $N_{Tx} \times N_{Tx}$ scattering matrix $\mathbf{S}(t)$.

$$\vec{v}_{ret}(t) = \mathbf{S}(t)\vec{v}_{fwd}(t) \quad (8.1)$$

As a result, there are 72 independent observations per time-point (symmetric 8×8 complex matrix yields 36×2). Hess et al. showed that measurements of $\mathbf{S}(t)$ contain motion information caused by relative changes in tissue conductivity, $\sigma(\vec{x}, t)$, within the region of sensitivity of the coil. This is observed as impedance which changes with respiratory and cardiac motion. Since this method is not dependent on the MR signal or tissue contrast, as it is an electrical signal, there is no increase to the sequence duration. The motion measurements are inherently synchronised with data acquisition and allow for any sampling trajectory.




Figure 8.1: Illustration showing the scattering of an 8 channel transmit system through the body (grey). Due to electrical coupling a forward voltage on channel 1, $\vec{v}_{fwd}^1$, produces a return voltage which is measured on all transmit channels, $\vec{v}_{ret}^1$ to $\vec{v}_{ret}^8$. However, in practice we are transmitting on all channels simultaneously so there is both a forward and returned voltage on all Tx channels.

This can be implemented to existing sequences using the current imaging excitations (unlike navigators which are an MR based method) or using additional small monitoring pulses. However, you can only determine the scattering matrix if you know the contribution from each individual Tx channel- which you do not if you are transmitting the same pulses on each Tx channel simultaneously.

To get around this, it is possible to only use the returned voltages, not the full scattering matrix, from a pTx array. Although, this is inherently susceptible to pre-amplifier drifts [243]. Alternatively, the motion induced change in the scattering coefficients can be monitored [242].

The method proposed by Jaeschke et al. to measure the full scattering matrix involves some small channel-wise RF pulse modifications. Three different strategies to encode the scattering matrix information were explored; time-division multiplexing, frequency-division multiplexing and a hybrid code-division multiplexing [241]. The optimal solution was found to be frequency-division multiplexing, where you add additional monitoring pulses far off-resonance from the main imaging excitations, each channel with a slightly different offset frequency. Frequency-division multiplexing (offset 100 kHz from excitation and with additional 10 kHz individual channel spacings) of the monitoring RF pulses with imaging excitations shows no introduction of artefacts, as shown in Figure 8.2.




Figure 8.2: Phantom data acquired with frequency division multiplexing to measure the full scattering matrix. **a)** Phantom image acquired with monitoring pulses. **b)** Main excitation turned off. Only using the monitoring pulses shows no introduction of artefacts due to excitation outside FOV. **c)** shows the time domain RF signal amplitude measured from the DICOs for a single sinc RF pulse on each transmit channel. The frequency domain in **d)** shows small monitoring pulses are evident, offset from the main excitation frequency, at a much lower amplitude.

Figure reproduced using own data from Jaeschke et al [241].

Hess et al. demonstrated that the accuracy depends on the amount of energy deposited and pTx scattering as a method of motion detection shows good agreement with MRI measurements. They showed 24 ± 17 ms cardiac gating accuracy for free-breathing and 1.3 ± 0.9 mm diaphragm prediction with minimal additional SAR ($\approx 0.01\%$ of SAR limits) as a result of the monitoring pulses.

Although initially developed for cardiac triggering, pTxS has also demonstrated sensitivity to cardiac and respiratory monitoring whilst simultaneously detecting rigid body head motion using only a head coil [244]. Figure 8.3 demonstrates the effect of step-wise head motion on the real and imaginary diagonal terms of the scattering matrix acquired over 10 GRE scans. The scattering matrix elements remain consistent during the scan but change between different head positions.




Figure 8.3: Effect of step-wise head motion on the real and imaginary diagonal terms (S_{11} , S_{22} , etc.) of the scattering matrix acquired over 10 GRE scans. Only diagonal terms are plotted for simplicity.

8.2.2.2 Pilot Tone

PT is a sequence independent solution to motion tracking with no on-subject hardware. PT uses minimal additional hardware, in the form of a RF signal generator, to produce a coherent signal from a continuous wave RF source at a frequency, f_{PT} , inside the scan room, which is received with the MR signal and can be used to track motion. PT gets its name from telecommunications engineering, where a "pilot tone" refers to a signal, usually as a single frequency, that is transmitted as a reference. The PT signal is outside of the frequency range of the MR signal, and outside the FOV of the image but within the receiver frequency band, where it is over-sampled in the frequency encode direction. This enables a high-temporal resolution through changes in the loading of the receive array causing fluctuations in the PT signal. As a result, unlike pTxS, the PT signal appears in the image data. An example of the PT signal in the image data is shown in Figure 8.4 where the signal resides in the $2\times$ oversampled (OS) region.

The PT signal is transmitted at a frequency offset from the scanners frequency, f_{B_0} , given by

$$f_{PT} = f_{B_0} + BW_{\text{pixel}} \cdot N_{RO} \left(\frac{\Delta ISO_{RO}}{FOV_{RO}} + \frac{OS_{\text{factor}} \cdot PT_{pos}}{2} \right) \quad (8.2)$$

where BW_{pixel} is the readout bandwidth per pixel [Hz px^{-1}], N_{RO} is the (non-OS) readout scan matrix size. ΔISO_{RO} is the FOV offset in the RO dimension from isocentre [mm] and FOV_{RO} is the (non-OS) FOV in the readout dimension [mm].




Figure 8.4: Example of pilot tone signal (red arrow) present far into the oversampled (OS) region of GRE image data. **A)** shows image data coil-combined using a basic sum-of-squares approach and **B)** shows coil-combination using receive sensitivity maps using ESPIRiT.

OS_{factor} is the OS factor (typically 2) and PT_{pos} depends on the desired position of the PT signal within the full (OS) FOV. PT_{pos} can be between -1 (bottom of the OS FOV) and 1 (top of the OS FOV).

However, despite the PT signal being outside of the ROI, it is necessary to filter out the PT signal to prevent the leakage artefacts seen in Figure 8.4. To do this, Speier et al. [245] recommended fitting the PT signal in hybrid k -space to a model of the form $Ae^{i2\pi\hat{f}_{PT}t}$ where A is the PT amplitude and $\hat{f}_{PT}$ is the apparent PT frequency after carrier demodulation, $\hat{f}_{PT} = f_{PT} - f_{B_0}$. This modelled PT signal can then be subtracted from the MR data.

8.2.3 Motivation and Aims

There are many methods of motion detection. RF sensor-based methods like pTxS and PT are promising methods that require no (or very little) additional hardware. However, the challenge with RF sensor-based methods are that whilst changes in RF sensor signals correlate well with motion, extracting quantitative motion estimates for each degree of freedom requires training a statistical model.

Previous work by Papp et al. [244] was able to form a general linear model (GLM) using 17 step-wise head positions to predict head motion. However, due to the slow acquisition of these GRE images used for model training only a limited number of poses could be examined. Hence, an EPI sequence was desirable in order to quickly probe and evaluate pTxS and PT over a larger range of head motions. The aim of this chapter was to develop a way to train a GLM using a fast calibration scan for both pTxS and PT.

8.3 Methods

In this chapter I acquired simultaneous pTxS and PT signals during a fast 3D EPI sequence for training a GLM. I compared the accuracy of PT and pTxS in predicting rigid-body head motion with a multivariate linear regression statistical model. Two subjects were scanned in a 7 Tesla MRI (Magnetom VB17, Siemens Healthcare Ltd, Germany) equipped with parallel transmit and imaged using an 8Tx/32Rx head coil (Nova medical, USA).

A video was projected into the bore to promote similar head motion paradigms between participants. The subject was asked to track a moving crosshair with their nose, performing continuous smooth head motions. Subject A was guided to perform circular, up/down, and left/right motion paradigms. Subject B performed similar head motions as subject A but with a larger amplitude and additional roll motion. Motion was performed slowly over a 5 minute scan duration to maintain the accuracy of registration by keeping motion within a single image volumes TR (256 ms) small.

8.3.1 Sequence Implementation

A segmented controlled aliasing in parallel imaging (CAIPI) 3D EPI sequence with a binomial (1-2-1) water excitation RF pulse, to minimise the volume TR, was developed to obtain 1200 volumes ($TR_{vol} = 256 \text{ ms}$) with CAIPI width and slope of 8 and 1, respectively, with a band offset of 2 to match the 3D acceleration factor, as shown in Figure 8.5. Reference lines were acquired in both phase encode directions during which no subject motion occurred for phase correction. $TE = 6 \text{ ms}$, $TR = 16 \text{ ms}$, $FA = 15^\circ$, $BW = 2894 \text{ Hz px}^{-1}$, $FOV = 250 \times 250 \times 125 \text{ mm}^3$, slice thickness = 3.9 mm, scan matrix = $64 \times 49 \times 32$ (including 75% phase PF), total acceleration factor 6.




Figure 8.5: CAIPI sampling mask used for data acquisition (colours represent acquired samples). CAIPI width = 8, slope = 1, band offset = 2 (to match the 3D acceleration factor), 16 shots with 17 phase encodes per shot and total acceleration factor $R=6$. Numbers indicate order of samples within a single shot.

8.3.2 Hardware and Sequence Adaptions for RF Sensors

PT and pTxS signals were simultaneously measured. To measure pTxS the binomial water excitation pulse was overlaid with a frequency-division multiplexed scattering matrix measurement (100 kHz offset; 10 kHz spacing; 15% pulse amplitude). The pTxS measurements were recorded by enabling a switch in the scanner's command line interface which exported the digitised DICO data along with the TWIX data for analysis.

PT was implemented by transmitting a pure sine wave (90 kHz off-resonance) using an external RF signal generator (VSG25A, Signal Hound, USA) controlled via VSG25A software version 1.0.6. The signal generator was connected via a coaxial cable through the waveguide in the RF penetration panel to a dipole antenna placed at the end of the patient table. To prevent damage to the RF signal generator, situated in the control room, a pair of cross diodes were soldered inline at the output port. The PT signal was assessed using an oscilloscope and FID measurements were acquired with the scanner to find the nominal output power that should be used during experiments, see Figure 8.6. The nominal output power of -30 dBm was

found to give a clean signal on the order of the k -space data, without saturating the receivers and operating in the linear regime of the preamplifier.




Figure 8.6: Pilot tone spectral density measured from FID measurements on the scanner for different nominal power outputs on the signal generator (SG). A nominal power output of -30 dBm gives a strong peak with amplitude on the order expected for k -space data.

8.3.3 RF Sensor Signal Processing

The timeseries of scattering matrices are obtained by parsing the TWIX data using *mapVBVD* [207] which contains the forward and reflected voltages for every RF pulse. The scattering matrix for each RF pulse is calculated by inverting Equation 8.1 and is unaffected by amplifier drifts. PTxS measurements were averaged across each RF pulse giving 16 pTxS measurements per image volume (one per shot). The real and imaginary components of the (8×8) S-matrix were reshaped as a vector of 1×128 .

The PT signal is found after an inverse fast Fourier transform in the readout direction and extracted using a peak-finding algorithm. Typically, the PT signal is then subtracted from the k -space data prior to image reconstruction using a signal model which assumes a constant PT frequency [245]. However, although a pure sine wave is being transmitted (confirmed via oscilloscope measurements), due to eddy current compensation during the EPI sequence the PT signal cannot be modelled as a single frequency. This means these methods are imperfect at removing PT signal from EPI imaging data [246].

Instead of a pure sine wave signal model, eddy current related phase shifts and $\pm$ readout frequency changes in the PT signal were corrected by multiplying each data point with the conjugate phase from an arbitrarily chosen reference which was chosen to be the complex summation of PT signal across all receive channels. The MR data was then subtracted from this to provide clean PT-free MR data, as shown in Figure 8.7.




Figure 8.7: Example of pilot tone signal (red arrows) present in hybrid k -space EPI data (inverse Fourier transformed along readout) is shown in **A**). Subfigure **B**) shows the hybrid k -space data with pilot tone signal removed and **C**) shows a single readout line. The pilot tone signal could be pushed further into the $2\times$ oversampled region. The pilot tone signal appears as alternating either side of the MR data due to the EPI readout and $\pm$ readout directions.

PT signals were then reordered to account for the segmented CAIPI sampling trajectory and averaged for each EPI shot to obtain 16 PT measurements per image volume (one per shot). The absolute value of the PT signal results in a 1×32 vector for each shot.

A combination of signals was also evaluated by concatenating the pTxS and PT signals resulting in a 1×160 vector. To prevent overfitting during subsequent analysis, the RF signals were normalised, whitened using a SVD, and truncated to reduce their dimensionality to rank 12 which was found to be a suitable threshold for both PT and pTxS and close to the optimal hard rank threshold [247], as shown in Figure 8.8. To allow for an offset in the GLM a "1" was concatenated to the end of the RF sensor signal vector.

A summary of the RF sensor processing pipeline is shown in Figure 8.9.




Figure 8.8: Singular value decomposition of pilot tone (PT) and pTx scattering (pTxS) data. Optimal hard rank threshold calculated from Gavish et al. [247].




Figure 8.9: Illustration of the RF sensor processing pipeline from the raw input data to the output RF signal vectors for pilot tone (PT), pTx scattering (pTxS) and a combination (Comb.) of both methods. OS; oversampled, RF; radiofrequency, SVD; single value decomposition.

8.3.4 Image Reconstruction

After PT removal, the clean PT-free MR data was Nyquist ghost phase corrected using motion-free reference lines. Data was Hann filtered and zero-padded up to $64 \times 64 \times 32$. Image reconstruction was performed using a conjugate gradient SENSE reconstruction in the Berkeley Advanced Reconstruction Toolbox (BART) [248] using coil sensitivities calculated from ESPIRiT [51].

8.3.5 Post-processing

Brain extraction and 3D rigid-body linear (affine) registration (6 DOF) were performed against the first image volume using FSL's [208] brain extraction and FMRIB's linear image registration tools (BET, FLIRT) [249–251] to obtain the ‘actual’ relative head position from the EPI data in the form of transformation matrices. The transformation matrices were converted to 1×6 rotation-translation vectors made from concatenating the Euler angles, from the rotation matrix, with the translations. These were then low-pass filtered to remove any sporadic erroneous measurements from registration. The registered motion estimates for subject A from FLIRT are shown in Figure 8.10.




Figure 8.10: Actual motion parameters for translations (top) and rotations (bottom) for subject A extracted using FSL's FLIRT registration tool. Motion paradigms were still (S), circular clockwise (CW), anti-clockwise (AC), up (U), down (D), left (L) and right (R).

A multivariate linear regression model was calculated in MATLAB (R2021a, MathWorks, USA) and formed by fitting the first 400 EPI rotation-translation vectors (600 for Subject B) to the processed pTxS and PT RF sensor signals. The remaining 800 rotation-translation vectors were used to test the models. The six degrees of motion rotation-translation vectors were collapsed to a single RMS score [mm], E_{RMS} , using an assumed head radius of $R = 80$ mm, as defined by Jenkinson et al. [252] by

$$E_{\text{RMS}}^2 = \frac{1}{5}R^2\text{Trace}(\mathbf{A}^\top \mathbf{A}) + \vec{t}^\top \vec{t} \quad (8.3)$$

where $\mathbf{A}$ is a 3×3 matrix and $\vec{t}$ is a 3×1 vector extracted from an error matrix, $\mathbf{M}$, defined as a 4×4 matrix from the difference between the first transformation matrix, $\mathbf{T}_1$, and the n^{th} transformation matrix, $\mathbf{T}_n$, as

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \vec{t} \\ 000 & 0 \end{bmatrix} = \mathbf{T}_n - \mathbf{T}_1 \quad (8.4)$$

A global RMSE score for PT, pTxS and a combination was calculated by taking the mean of the difference in RMS score across the entire timeseries.

8.4 Results

When the subjects were instructed to move, subjects A/B moved with an average velocity of $2.5 \text{ mm s}^{-1}/4.8 \text{ mm s}^{-1}$ with translations and rotations up to $20 \text{ mm}/30 \text{ mm}$ and $7/10^\circ$. During a stationary 20 s window, the RMSE for subjects A/B was $1.0 \text{ mm}/0.7 \text{ mm}$ for PT and $0.6 \text{ mm}/0.5 \text{ mm}$ for pTxS. During motion, the ‘ground truth’ registration precision for subjects A/B was estimated to be $0.5 \text{ mm}/0.9 \text{ mm}$ for translations and $0.2^\circ/0.3^\circ$ for rotations (determined as the first order derivative/ $\sqrt{2}$). The cumulative energy in the first 12 singular values was 61% for pTxS, 90% for PT, and 61% for combined signals. The RMSE for each degree of freedom is shown in Table 8.1 as well as the average global RMSE which was $1.5 \text{ mm}/2.3 \text{ mm}$ for PT, $1.4 \text{ mm}/2.1 \text{ mm}$ for pTxS and $1.3 \text{ mm}/2.1 \text{ mm}$ for a combination of these for Subjects A/B. The predicted motion traces from PT, pTxS and a combination of both are shown in Figure 8.11 for each degree of freedom and correlate well with the ground truth.

DOF	Subject A			Subject B*		
	PT	pTxS	Comb.	PT	pTxS	Comb.
x [mm]	0.9	1.0	1.0	1.4	1.1	1.1
y [mm]	1.0	0.9	0.9	1.2	1.0	1.0
z [mm]	0.9	0.8	0.7	1.8	1.8	1.8
Yaw [°]	0.3	0.3	0.3	0.5	0.3	0.3
Roll [°]	0.3	0.2	0.2	0.6	0.6	0.6
Pitch [°]	0.3	0.4	0.4	0.4	0.3	0.3
Global RMSE [mm]	1.5	1.4	1.3	2.3	2.1	2.1

Table 8.1: Root-mean-square error for Pilot Tone (PT), pTx Scattering (pTxS) and a combination of PT and pTxS signals for both subjects A and B across the predicted timeseries. *Data for subject B is trained on 600 volumes to include the additional roll motion in training.




Figure 8.11: Actual and predicted translations and rotations from parallel transmit scattering (blue), pilot tone (red) and a combination of both (green). The red shaded area represents the data used for model training.




Figure 8.12: Actual and predicted root-mean-square combined motion from parallel transmit scattering (blue), pilot tone (red) and a combination of both (green). The red shaded area represents the data used for model training. The grey shaded area represents the expected error in 'Actual' due to motion.

8.5 Discussion

I have demonstrated both PT and pTxS show sensitivity to head motion and can be used for monitoring head motion using a linear regression model for estimating quantitative motion parameters. I have also shown the ability to measure both of these methods simultaneously and evaluated a combination of these signals.

I did not find a significant difference between pTxS and PT predicted motion parameters. It was expected that since there are no Tx elements in the z -direction that pTxS would not be able to provide accurate measurements in this dimension, however, this was not the case and pTxS still had good sensitivity to motion in this direction. Both PT and pTxS show an increased global RMSE for subject B which involved larger amplitudes of head motion. This is likely due to changes in coil loading related to rigid motion or departure from the GLM.

Instructing subject motion cannot guarantee that all possible subject poses are examined and how this subject-wise ‘train then apply’ methodology would integrate into a clinical workflow is uncertain. Additionally, the time spent training the GLM (100s) is far too long for clinical use. Therefore, going forward, it would be interesting to try and form a model across a population and explore the accuracy of such a model when applied to an unseen subject. This would remove the necessity for a subject-wise calibration. However, one issue I faced when attempting this was that the pTxS signal amplitude and standard deviation is sensitive to frequency offsets. The issue with having a relative frequency offset when acquiring a 2D sequence for pTxS is that the measurement of the scattering matrix is sensitive to the kilohertz changes in the carrier frequency with each slice position. This is demonstrated in Figure A.13 where small frequency offsets produce linear variation in the scattering matrix. The solution I settled on was using fixed absolute frequencies for the frequency division multiplexed monitoring pulses instead of a relative 100 kHz offset. Alternatively, the frequency effects could be modelled out.

Aside from the frequency dependency, one issue which complicates a universal statistical model for pTxS is that the scattering matrix elements amplitude and standard deviation alters with pulse shape. This is demonstrated in Figure A.14 where truncating the RF pulse prior to calculating the scattering matrix leads to a decrease in both the accuracy and precision of the scattering matrix elements. This means two different pulse shapes, overlaid with frequency division multiplexed monitoring pulses, will output slightly different scattering elements, depending

on pulse shape and amplitude. This is because, as found by Jaeschke et al., the accuracy and precision of the scattering matrix is determined by the energy of the scattering matrix frequency division multiplexed monitoring pulses [241]. Hence, changing the flip angle of the RF pulse results in the mean and standard deviation of the scattering matrix elements changing, as seen in Figure A.15. A universal model may not even be possible due to variation in subject impedances.

Motion predictions using all methods agree with EPI registered motion and where they do not, they tend to agree well with each other, indicating true changes not captured by the rigid-body motion model. For example around 193s in Figure 8.12. In this example, pTxS and PT agree with each other but deviate from the slower "actual" EPI registered motion estimate. Further work is required to explore their sensitivity to smaller amplitudes of motion.

Both RF sensor-based methods demonstrate high temporal resolution which could allow for motion updates per TR. pTxS signals are obtained during transmission whereas PT signals are obtained during reception which occur much more frequently during an MR sequence than transmission. Hence, PT may give more frequent updates, unless additional low-amplitude monitoring pulses are implemented into the sequence.

Further work is also required to interpret non-rigid body system changes and refine the motion model to achieve the accuracy required for in-vivo motion tracking. In some datasets, motion far from the coil (e.g. foot movement) resulted in changes in the PT signal, resulting in sporadic erroneous head motion estimates. This motion did not corrupt pTxS signals since the motion was outside of the Tx sensitivity. Hence, as a result I believe the PT transmitter should be mounted as close to the ROI as possible to avoid this motion corrupting the PT signal. Indeed, recent commercial work implements PT transmitters directly onto coils for cardiac triggering, under the product name "BioMatrix Beat Sensor" [253]. However, adapting coil designs in order to incorporate a PT transmitter directly onto coils could be expensive.

Although a GLM was used in this work there is no physical justification why this is suitable. Recent work shows linearity in Sim4Life simulations of pTx scattering for translations up to 10 mm but rotations show non-linearity [254]. Hence, alternative methods such as Gaussian Processes may be a better choice than a linear regression model [255]. Alternatively, deep-learning could be employed to learn the complicated relationship between the RF sensor signals and the motion parameters and may be useful for estimating larger non-linear motions.

Recent developments of PT have applied two frequencies (≈ 2.4 GHz) to overcome its limited spatial sensitivity due to the relatively long (compared to the range of motion expected) ^1H Larmor wavelength, called beat pilot tone (BPT) [256, 257]. The two frequencies are separated by the receiver bandwidth, which is essentially the Larmor frequency, ω_L , and results in inter-modulation which is picked up by the receive coils. The two high frequency transmitters cause standing waves to be present with a much shorter wavelength and hence sensitivity to smaller ranges of motion. This means that the transmitter can be located much further away from the subject and there is greater flexibility in what frequencies are emitted, as long as they remain separated by ω_L . BPT has demonstrated much greater sensitivity (approximately $20\times$) to small motion than PT which is only sensitive to bulk motion. Although, at this stage, whether the greater sensitivity offered by BPT is genuinely useful, or whether it just further complicates the interpretation of the PT signal, is unknown.

Alternatively, a hybrid method could be conceptualised where the PT signal is generated by the transmit coil itself. This hybrid method would result in a 8×32 matrix, which is larger than that of both pTxS (8×8) or PT (1×32) and may potentially capture more motion information. Of course, this would require simultaneous transmit and receive which has not yet been demonstrated but may be possible by overriding the T/R switch. This hybrid approach would overcome the problem of sensitivity to exterior motion and reduce costs by not having to adapt coil designs with on-board PT sensors. It would also be applicable to lower field scanners (e.g. 1.5 T or 3 T) unlike pTxS which is only useful for pTx capable UHF systems.

Nevertheless, the challenge with these RF sensor-based methods for quantitative motion estimation remains in the calibration of a suitable statistical model. Recent work by Brackenier et al. [258] demonstrated motion correction, using a similar linear regression model, of an MPRAGE using PT signals calibrated during short (4s) blocks dispersed within the entire scan. This was shown to be a clinically viable solution which is applicable to any sequence making use of sequence dead time or by implementing dedicated calibration blocks.

8.6 Summary

I have shown there is high sensitivity to rigid body head motion in both PT and pTxS signals. Within this work both methods currently achieve a similar degree of accuracy but more investigation is required to determine their sensitivity for smaller amplitudes of motion and whether a universal statistical model is possible.

So long, and thanks for all the fish.

-Douglas Adams

9

Summary

Contents

9.1	Contributions	189
9.2	Future Directions	191
9.3	Closing Remarks	192

9.1 Contributions

This thesis was largely focused on improving B_1^+ mapping which is at the foundation of pTx pulse design, in order to improve pTx performance and particularly in regions outside the brain. SatTFL works well for use at ultra-high field as a 2D sequence but suffers from anisotropic voxels and slice-profile effects. A 3D sequential SatTFL sequence is infeasible for clinical use due to the long acquisition time. This led me to develop a novel 3D B_1^+ mapping sequence based on SatTFL.

In this thesis I have outlined the development of this novel volumetric B_1^+ mapping sequence which I have called "Sandwich". I have demonstrated that Sandwich is fast ($20\times$ faster than a sequential 3D SatTFL) as well as being robust to B_0 and flow with a lower peak B_1^+ than 2D sequences which is ideal for B_1^+ mapping in the body. Most importantly, the 3D Sandwich sequence enables further acceleration that was not possible with the 2D SatTFL sequence. Hence, Sandwich can achieve a significant reduction in acquisition time therefore making tailored pTx design more feasible.

I have demonstrated the practicality of the Sandwich sequence in the brain and the heart and compared it to similar sequences through simulations as well as phantom and in vivo experiments. My simulation framework helped to identify and quantify the limitations of each method under various parameters such as T_1 , B_0 , flow and diffusion.

During my internship with Siemens Healthineers in Erlangen, Germany, I implemented the Sandwich pulse sequence and ICE reconstruction for the Magnetom Terra.X XA60 platform. The Sandwich sequence is hopefully to become the default product sequence for pTx pulse design, replacing the conventional 2D SatTFL. This will help to allow faster and more accurate tailored pTx pulse design.

Through my work on improved B_1^+ mapping I also determined that coil-cycling is essential for the accuracy and precision of relative magnitude maps. Coil-cycling is especially important for 3D relative magnitude maps acquired in the body. I recommend this simple modification is implemented to ensure accurate mapping and therefore improved pTx pulse design.

I also investigated an acquisition strategy to acquire single breath-hold volumetric B_1^+ maps in the heart for pTx. I demonstrated that the acquisition of small scan matrices were adequate for cardiac B_1^+ mapping. Although two Tx modes are likely sufficient for cardiac absolute B_1^+ mapping a third mode would be particularly beneficial for the Sandwich scheme. This work is important for developing ^1H cardiac imaging at UHF.

My initial work with RF sensor-based motion detection methods demonstrated the ability to use a simple linear regression model to predict head motion calibrated using a fast EPI scan. I did not find a significant difference in the training or prediction accuracy between PT and pTxS. I demonstrated the ability to acquire both of these methods simultaneously and saw no clear benefit from combining the signals. However, the accuracy of my "ground truth" calibration method was limited and further investigation is required to evaluate their accuracy to smaller amplitudes of motion where we may see more subtle differences between the methods. I did find some evidence that a linear model may not be suitable for large motion amplitudes.

9.2 Future Directions

At the time of writing, Siemens Healthineers are continuing to develop the Sandwich sequence by incorporating a multi-echo readout in the reference image for simultaneous B_1^+ and B_0 mapping. Other improvements could be made to the Sandwich sequence for example a radial instead of Cartesian acquisition for robustness against respiratory motion or the use of a universal pulse for the imaging excitation.

The largest assumption in this work is the temporal invariance of the transmit field. Currently, B_1^+ field maps are measured once and applied for the remainder of the scanning session. However, previous cardiac work at 3 T has shown respiration leads to B_1^+ field changes between inhalation and expiration ranging from 1% to 14% and this is expected to be worse at higher field strengths [259]. Hence, respiration-resolved cardiac B_1^+ field maps would be useful to design pTx pulses that are either more robust to respiration or a library of different pulses to be applied at certain phases of the respiratory cycle. This could be done by continuously measuring fast coil-cycled relative maps throughout the respiratory cycle and binning them according to the respiratory phase. The respiratory binning could be done without using any bellows or on-subject hardware by utilising RF sensor-based methods as outlined in Chapter 8.

TxLR is an iterative method of reconstruction, which takes a considerable amount of time. Therefore, a low-rank deep learning reconstruction is desirable. Despite aiming to make some progress in this area, I ultimately found this to be challenging as the reconstruction problem is novel. Typically, deep-learning neural networks work with large k -space (or image space) matrices (typically $\geq 256 \times 256$) which generally have a fully sampled centre of k -space. Instead, we are working with a very small (24×24), and highly under-sampled, k -space which could be on the order of what is typically used for the ACS region. I am, however, aware of others actively working on this problem.

RF sensor based methods are an innovative approach to motion detection but they require further development to be fully realised. Many potential methods exist such as PT, BPT, pTxS and a potential hybrid PT/pTxS method but which of these methods is able to provide the most practical and useful motion information is still not known. Hence, further work should focus on evaluating the sensitivity of the RF sensor signals to motion. This could be explored through Sim4Life simulations. Aside from the choice of method, further work needs to be done to refine the motion

model. A linear model is probably inadequate and more complex models should be explored. Deep learning is the most promising way to learn the complex relationship between RF sensor signals and motion and would be helpful to obtain a universal motion model that can be applied across a population. Although, how much and what kind of motion data is required in order to train such a universal model, and if a universal model is even possible, still needs to be determined.

9.3 Closing Remarks

As UHF MRI is increasingly used for clinical scans, the need to be able to adapt to the inter-subject variability only grows, and hence RF pulses should be tailored to the subjects own anatomy and B_1^+ mapping remains relevant. Improvements in B_1^+ mapping methods will only serve to improve the design of pTx pulses. It is therefore difficult for me to believe that UHF MRI will do away with B_1^+ mapping altogether, although, how deep learning will impact scanner calibration is anyone's guess.

The use of the Tx (pTxS) and/or Rx (PT) coils themselves for motion detection is in its infancy and there is a lot of exciting work to be done in this area with potentially huge benefits for motion detection and correction. I am eager to follow along and see RF sensor-based methods maturing as a method of motion detection.

I look forward to seeing how others may pick up from this work, implement their own ideas and put forward new concepts in the field of MRI. But if one day I find myself back in a Siemens 7T scanner and I hear the brief yet distinct sound of the Sandwich sequence, it will certainly bring a smile to my face in the knowledge that I made some *small* contribution to the field.

A

Contents

A.1	T1 Correction Scheme	193
A.2	VB17 Slave Dropout	196
A.3	VB17 Slave Synchronisation	197
A.4	Ineffective Gradient Spoiling	199
A.5	HS8 Versus HS4	201
A.6	Scattering Matrix Sensitivity	203

A.1 T_1 Correction Scheme

An additional modification to the proposed Sandwich scheme was tested in order to further reduce sensitivity to T_1 relaxation using Eq. 4.1. This involved a third readout train, S_2 , which is performed directly after S_1 . Phantom and *in vivo* results for the T_1 Correction scheme, as outlined in Section 4.3.1.2, are shown here.

Initial simulation results into a T_1 correction scheme appeared to reduced the variation in the measured FA due to T_1 which extends the upper bounds of the DR. However, this ultimately comes at the cost of an overall slightly reduced DR due to decreased signal for low FAs. This is evident as an increase in the NRMSE in both the phantom and brain maps. Ultimately, since there were no noticeable T_1 -related artefacts in the Sandwich maps and since the acquisition using the third readout train would not fit entirely within diastole, this T_1 correction modification was not explored further.




Figure A.1: Flip angle maps from a water phantom at a reference voltage of 60 V for various sequences. **i)** segmented satTFL method with $TR = 20$ s **ii)** satTFL **iii)** satTFL acquired segmented and TR reduced to 2 s **iv)** segmented Sandwich scheme with $TR = 1$ s, $\beta = 0.1\alpha$ **v)** segmented T_1 Correction scheme with $TR = 1$ s, $\beta = 0.1\alpha$.




Figure A.2: Flip angle maps in the brain at a reference voltage of 60 V for various sequences. **a)** satTFL **b)** satTFL acquired segmented and TR reduced to 2 s **c)** segmented Sandwich scheme with TR = 1 s, $\beta = 0.1\alpha$ **d)** segmented T₁ Correction scheme with TR = 1 s, $\beta = 0.1\alpha$.

A.2 VB17 Slave Dropout

For general pTx on the Magnetom 7T Step 2.3 VB17 pTx system the master synchronises all the slaves at the start of the sequence. When cardiac gating is enabled, the master has to resynchronise all the slaves each TR. The slaves then need to calculate the sequence timings for the next TR on the fly and wait for the "green light" from the master. If one of the slaves is slightly too slow for whatever reason (e.g. too many *COUTs*), it can lead to instability in the system and potentially a crash which requires a reboot or even a full power cycle.

This was the cause of one of the issues when implementing this sequence on the scanner, due to calculating the encoding trajectory through a new *ReorderInfo* class. Every TR the new *ReorderInfo* class would read in a binary text file containing the undersampling mask for each transmit channel and calculate the encoding trajectory. Although no error was reported by the system, this was taking too long and caused system instability. After this problem was identified, the sequence's *ReorderInfo* class was rewritten to avoid reading in the text file each TR, which improved the stability of the sequence when cardiac gating.

Figure A.3 shows the relative phase between master and slaves for ungated and gated Sandwich sequences. For ungated sequences the phase remains constant throughout the acquisition only changing when the transmit mode changes from CP^+ to CP^{2+} . However, for gated sequences it appears three slave channels begin to have a phase offset that varies randomly. This is due to the slave dropping out and having zero amplitude therefore arbitrary phase for the remainder of the scan. Once this was identified the *ReorderInfo* class was rewritten to minimise the computational time each slave would need to perform the trajectory calculations which fixed the problem.




Figure A.3: VOP data showing the relative phase to the master channel (Tx1; blue) for ungated and gated sequences.

A.3 VB17 Slave Synchronisation

When troubleshooting problems with the accelerated cardiac B_1^+ maps it was noticed that there were additional artefacts introduced when the joint reconstruction was performed with a combination of gated relative maps and ungated Sandwich absolute maps. The green arrow in Figure A.4a) shows a generic shading artefact due to poor CP^+ / CP^{2+} coverage. However, the red arrows in Figure A.4b) show severe B_1^+ errors.




Figure A.4: Artefacts noticeable in pTx maps acquired in a drum phantom when performing a joint reconstruction with an ungated Sandwich absolute maps and a) ungated and b) gated relative maps.

Multiplying the ungated relative maps by the conjugate of the gated relative maps you can see in Figure A.5 a difference in the phase of the relative maps between ungated and gated acquisitions. There is a constant additional phase on the slave channels of approximately 71° .




Figure A.5: Phase offset between ungated and gated relative maps acquired in a drum phantom.

Looking at VOP data for the preparation pulse, as seen in Figure A.6, there was a clear $10\mu\text{s}$ delay between master and slave RF pulses which was the cause of the constant additional phase on the slave channels and the artefacts seen in Figure A.4. I found this was caused by an error in the class which adapted *SeqLoop* to control the halting and re-synchronisation of the slaves when cardiac gating resulting in an unexpected $10\mu\text{s}$ delay between the master and slaves. This delay resulted in the phase difference seen in relative maps.




Figure A.6: VOP data for ungated and gated acquisitions showing a clear $10\,\mu\text{s}$ delay between the master (blue) and slave preparation pulses.

A.4 Ineffective Gradient Spoiling

During the implementation of the Sandwich sequence on XA60 for the Terra.X a few smearing artefacts were initially noticeable in the readout direction of the prepared images and increased with the reference voltage as seen in Figure A.7.




Figure A.7: Artefacts visible in the prepared images of the Sandwich sequence worsening for higher reference voltages.

Initially a variety of parameters were tested to understand their influence on the artefacts. The artefacts in question are clearly visible in Figure A.8.




Figure A.8: Artefacts visible in the prepared images of the Sandwich sequence for various protocols. **A)** Default centric reordering, **B)** Increased gradient spoiling, **C)** 10 dummy TRs, **D)** 10 s TR, **E)** 0° saturation pulse and **F)** Linear reordering.

The artefacts were significantly reduced when no saturation pulse was present or when linear reordering was used. To further investigate the cause of this artefact, the k -space centre line along the readout was plotted for the summed magnitude of receive channels as shown in Figure A.9 for centric versus linear reordering.




Figure A.9: A line through the readout of the **A)** reference and **B)** prepared images for the summed magnitude over all receive channels showing a reduction in an artefact in the prepared image due to a second echo from switching to a linear phase encoding scheme.

A second echo is visible approximately 0.75 ms after the main peak, which is reduced when linear reordering is implemented. It became clear that mixing of the transverse magnetisation from the reference image by the preparation RF pulse

and subsequent refocusing was the cause of the artefact. The second echo has a higher intensity closer to the preparation pulse, decaying through the FLASH train and hence less prevalent for linear versus centric encoding. Hence, after increasing the spoiler of the preparation pulse this artefact went away.

This problem arose since the SatTFL sequence on XA60 implemented the preparation pulse during a typical readout of the FLASH train, which only allowed a short duration for the preparation pulse and associated spoiling. This was not a problem for SatTFL, but since the Sandwich scheme acquires the reference image directly before the prepared image, this became an issue and the preparation pulse needed to be implemented in its own SBB with dedicated spoiler.

A.5 HS8 Versus HS4

During the implementation of the Sandwich sequence on XA60 for the Terra.X, the choice was made to switch to a HS4 pulse, shown in Figure A.10, for the saturation since the HS4 pulse was already present in the Siemens RF pulse library. The motivation for this was to reduce the number of derived classes and sequence files required to compile the sequence making upkeep simpler and less bespoke. The effect of changing to the HS4 pulse is investigated here using the simulations developed in Section 5.




Figure A.10: The magnitude (left) and phase (right) for HS8 (blue) and HS4 (red) preparation pulses.

For the same amplitude integral and duration of $5120\mu\text{s}$ changing to the HS4 instead of the HS8 pulse had an 11% reduction in V_{peak} (35.9 V versus 32.3 V) and a 27.8% reduction in SAR. The reduction in SAR comes from a reduction in the total energy from 4.62 J to 3.62 J with the average 10 s power reducing from 1.40 W to 1.10 W. The change from HS8 to HS4 had no impact on the DR of the Sandwich method.

Figure A.11 shows the longitudinal magnetisation as a function of B_0 and B_1^+ expressed as a dynamic range for the HS8 and HS4 pulses. As expected based on the energy values, the HS8 pulse does have a larger bandwidth than the HS4. However, the region over which the B_1^+ is consistent is approximately constant between the two pulses at around $\pm 1.1\text{ kHz}$. There is, however, a slightly larger ripple in the longitudinal magnetisation within this region for the HS4 which is clearer to see in Figure A.12.




Figure A.11: The longitudinal magnetisation as a function of B_0 and B_1^+ expressed as a dynamic range for **A)** HS8 and **B)** HS4 pulses. The colour represents longitudinal magnetisation with yellow indicating 1 (no excitation) and dark blue indicating -1 (full inversion). Arrows correspond to lines plotted in Figure A.12.




Figure A.12: The longitudinal magnetisation as a function of B_0 plotted for three nominal B_1^+ dynamic ranges.

A.6 Scattering Matrix Sensitivity

The scatter matrix measurements are sensitive to the frequency offsets resulting from slice selection for 2D sequences. This behaviour is demonstrated in Figure A.13. The impact of truncating the RF pulse, prior to calculating the scattering matrix, on the scattering matrix elements is shown in Figure A.14. A decrease in the RF pulse duration leads to a decrease in both the precision and accuracy of the scattering matrix itself, as shown by the variation in mean and increase in the SD. Similarly, the effect of increasing the FA on the measured scattering matrix elements is shown in Figure A.15. Increasing the FA of the RF pulse which contains frequency division multiplexed monitoring pulses (at 15% pulse amplitude) improves the accuracy and precision of the scattering matrix measurement. This complicates the ability to train a model based on scattering matrix measurements from one RF pulse type/amplitude and applying them to data acquired with another.




Figure A.13: Effect of slice selection on the mean and standard deviation of the scattering matrix elements.




Figure A.14: Effect of RF pulse truncation on the mean and standard deviation of the scattering matrix elements.




Figure A.15: Scattering matrix elements concatenated for phantom acquisitions at flip angles from 2° to 8° to show the difference in mean and standard deviation with flip angle.

References

- [1] James L. Kent et al. “Rapid B1+ Mapping using a Sandwiched Train Saturated TurboFLASH Sequence at 7T”. In: *ISMRM Workshop on Ultra-High Field MRI* (2022), p. 46.
- [2] James L. Kent et al. “Rapid B1+ Mapping using a Sandwiched Train Saturated TurboFLASH Sequence at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 30. 2022, p. 0456.
- [3] James L. Kent et al. “Rapid 3D absolute B1+ mapping using a sandwiched train presaturated TurboFLASH sequence at 7 T for the brain and heart”. In: *Magnetic Resonance in Medicine* 89.3 (Mar. 2023), pp. 964–976. DOI: 10.1002/mrm.29497.
- [4] James L. Kent et al. “A Simulation Framework for Comparison of Volumetric B1+ Mapping Methods”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 32. 2024, p. 1602.
- [5] James L. Kent et al. “A Comparison of FLASH-based Volumetric B1+ Mapping Methods in Phantom at 7 Tesla”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 32. 2024, p. 4957.
- [6] James L. Kent et al. “Acceleration Methods to Acquire Fast 3D Multi-Channel Absolute B1+ Maps at 7T for 8-Channel pTx”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 30. 2022, p. 2864.
- [7] James L. Kent et al. “Accelerated Volumetric Multi-Channel pTx B1+ Mapping at 7T for the Brain and Heart”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 31. 2023, p. 4408.
- [8] James L. Kent et al. “Accelerated 3D multi-channel B1+ mapping at 7 T for the brain and heart”. In: *Magnetic Resonance in Medicine* 92.5 (June 2024), pp. 2007–2020. DOI: 10.1002/mrm.30201.
- [9] James L. Kent et al. “Pilot Tone vs pTx Scattering : A Comparison between ‘ RF Sensor ’ Methods for Rigid Body Motion Detection of the Brain at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 31. 2023, p. 5258.
- [10] I. I. Rabi et al. “A New Method of Measuring Nuclear Magnetic Moment”. In: *Physical Review* 53.4 (Feb. 1938), pp. 318–318. DOI: 10.1103/PhysRev.53.318.
- [11] E. M. Purcell, H. C. Torrey, and R. V. Pound. “Resonance Absorption by Nuclear Magnetic Moments in a Solid”. In: *Physical Review* 69.1-2 (Jan. 1946), pp. 37–38. DOI: 10.1103/PhysRev.69.37.
- [12] F. Bloch, W. W. Hansen, and Martin Packard. “Nuclear Induction”. In: *Physical Review* 69.3-4 (Feb. 1946), pp. 127–127. DOI: 10.1103/PhysRev.69.127.
- [13] E. L. Hahn. “Spin Echoes”. In: *Physical Review* 80.4 (Nov. 1950), pp. 580–594. DOI: 10.1103/PhysRev.80.580.
- [14] P. C. Lauterbur. “Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance”. In: *Nature* 242.5394 (Mar. 1973), pp. 190–191. DOI: 10.1038/242190a0.

- [15] John R. Mallard. “Magnetic resonance imaging-the Aberdeen perspective on developments in the early years”. In: *Physics in Medicine and Biology* 51.13 (July 2006), R45–R60. DOI: 10.1088/0031-9155/51/13/R04.
- [16] Peter Mansfield et al. “Human whole body line-scan imaging by NMR”. In: *The British Journal of Radiology* 51.611 (Nov. 1978), pp. 921–922. DOI: 10.1259/0007-1285-51-611-921.
- [17] S Ogawa et al. “Brain magnetic resonance imaging with contrast dependent on blood oxygenation”. In: *Proceedings of the National Academy of Sciences* 87.24 (Dec. 1990), pp. 9868–9872. DOI: 10.1073/pnas.87.24.9868.
- [18] Peter Mansfield and A. A. Maudsley. “Planar spin imaging by NMR”. In: *Journal of Physics C: Solid State Physics* 9.15 (1976). DOI: 10.1088/0022-3719/9/15/004.
- [19] Peter Mansfield. “Multi-planar image formation using NMR spin echoes”. In: *Journal of Physics C: Solid State Physics* 10.3 (1977). DOI: 10.1088/0022-3719/10/3/004.
- [20] Raymond Damadian. “Tumor detection by nuclear magnetic resonance”. In: *Science* 171.3976 (1971), pp. 1151–1153. DOI: 10.1126/science.171.3976.1151.
- [21] Peter Mansfield. *The Long Road to Stockholm: The Story of MRI - An Autobiography*. Oxford University Press, 2013, pp. 1–241.
- [22] Qiuyun Fan et al. “Mapping the human connectome using diffusion MRI at 300 mT/m gradient strength: Methodological advances and scientific impact”. In: *NeuroImage* 254. February (2022), p. 118958. DOI: 10.1016/j.neuroimage.2022.118958.
- [23] Susie Y. Huang et al. “Connectome 2.0: Developing the next-generation ultra-high gradient strength human MRI scanner for bridging studies of the micro-, meso- and macro-connectome”. In: *NeuroImage* 243. April (Nov. 2021), p. 118530. DOI: 10.1016/j.neuroimage.2021.118530.
- [24] Pierre-François Van de Moortele et al. “B1 destructive interferences and spatial phase patterns at 7 T with a head transceiver array coil”. In: *Magnetic Resonance in Medicine* 54.6 (2005), pp. 1503–1518. DOI: 10.1002/mrm.20708.
- [25] Qing X. Yang et al. “Analysis of wave behavior in lossy dielectric samples at high field”. In: *Magnetic Resonance in Medicine* 47.5 (2002), pp. 982–989. DOI: 10.1002/mrm.10137.
- [26] Christopher M. Collins and Zhangwei Wang. “Calculation of radiofrequency electromagnetic fields and their effects in MRI of human subjects”. In: *Magnetic Resonance in Medicine* 65.5 (May 2011), pp. 1470–1482. DOI: 10.1002/mrm.22845.
- [27] Axel Haase et al. “FLASH imaging. Rapid NMR imaging using low flip-angle pulses”. In: *Journal of Magnetic Resonance* 67.2 (Apr. 1986), pp. 258–266. DOI: 10.1016/0022-2364(86)90433-6.
- [28] Jens Frahm, Axel Haase, and Dieter Matthaei. “Rapid NMR imaging of dynamic processes using the FLASH technique”. In: *Magnetic Resonance in Medicine* 3.2 (Apr. 1986), pp. 321–327. DOI: 10.1002/mrm.1910030217.

- [29] Axel Haase. “Snapshot FLASH MRI. Applications to T1, T2, and Chemical-Shift imaging”. In: *Magnetic Resonance in Medicine* 13.1 (1990), pp. 77–89. DOI: 10.1002/mrm.1910130109.
- [30] Stefan Blüml et al. “Spin-lattice relaxation time measurement by means of a TurboFLASH technique”. In: *Magnetic Resonance in Medicine* 30.3 (Sept. 1993), pp. 289–295. DOI: 10.1002/mrm.1910300304.
- [31] Geoffrey J.M. Parker et al. “Improving image quality and T1 measurements using saturation recovery turboFLASH with an approximate K-space normalisation filter”. In: *Magnetic Resonance Imaging* 18.2 (2000), pp. 157–167. DOI: 10.1016/S0730-725X(99)00124-1.
- [32] Iman Khodarahmi, Mahesh Keerthivasan, and Jan Fritz. “Turbo-spin echo based B1+ mapping in the presence of metallic hardware”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 30. 2022, p. 0100. DOI: 10.58530/2022/0100.
- [33] Alberto Tannús and Michael Garwood. “Adiabatic pulses”. In: *NMR in Biomedicine* 10.8 (Dec. 1997), pp. 423–434. DOI: 10.1002/(SICI)1099-1492(199712)10:8<423::AID-NBM488>3.0.CO;2-X.
- [34] David G. Norris. “Adiabatic radiofrequency pulse forms in biomedical nuclear magnetic resonance”. In: *Concepts in Magnetic Resonance* 14.2 (2002), pp. 89–101. DOI: 10.1002/cmr.10007.
- [35] Robin A. De Graaf and Klaas Nicolay. “Adiabatic rf pulses: Applications to in vivo NMR”. In: *Concepts in Magnetic Resonance* 9.4 (1997), pp. 247–268. DOI: 10.1002/(sici)1099-0534(1997)9:4<247::aid-cmr4>3.3.co;2-a.
- [36] Yasvir A. Tesiram. “Implementation equations for HSn RF pulses”. In: *Journal of Magnetic Resonance* 204.2 (June 2010), pp. 333–339. DOI: 10.1016/j.jmr.2010.02.022.
- [37] Michael Garwood and Lance DelaBarre. “The return of the frequency sweep: Designing adiabatic pulses for contemporary NMR”. In: *Journal of Magnetic Resonance* 153.2 (2001), pp. 155–177. DOI: 10.1006/jmre.2001.2340.
- [38] Matt A. Bernstein, Kevin F. King, and Xiaohong Joe Zhou. *Handbook of MRI Pulse Sequences*. Elsevier, 2004, p. 1040. DOI: 10.1016/B978-0-12-092861-3.X5000-6.
- [39] Y. Zur, Michael L. Wood, and L. J. Neuringer. “Spoiling of transverse magnetization in steady-state sequences”. In: *Magnetic Resonance in Medicine* 21.2 (Oct. 1991), pp. 251–263. DOI: 10.1002/mrm.1910210210.
- [40] Adrian P. Crawley, Michael L. Wood, and R. Mark Henkelman. “Elimination of transverse coherences in FLASH MRI”. In: *Magnetic Resonance in Medicine* 8.3 (Nov. 1988), pp. 248–260. DOI: 10.1002/mrm.1910080303.
- [41] R. R. Ernst and W. A. Anderson. “Application of Fourier Transform Spectroscopy to Magnetic Resonance”. In: *Review of Scientific Instruments* 37.1 (Jan. 1966), pp. 93–102. DOI: 10.1063/1.1719961.
- [42] Matthias Weigel. “Extended phase graphs: Dephasing, RF pulses, and echoes - pure and simple”. In: *Journal of Magnetic Resonance Imaging* 41.2 (Feb. 2015), pp. 266–295. DOI: 10.1002/jmri.24619.

- [43] Jürgen Hennig. “Echoes—how to generate, recognize, use or avoid them in MR-imaging sequences. Part I: Fundamental and not so fundamental properties of spin echoes”. In: *Concepts in Magnetic Resonance* 3.3 (July 1991), pp. 125–143. DOI: 10.1002/cmr.1820030302.
- [44] Jürgen Hennig. “Echoes—how to generate, recognize, use or avoid them in MR-imaging sequences. Part II: Echoes in imaging sequences”. In: *Concepts in Magnetic Resonance* 3.4 (Oct. 1991), pp. 179–192. DOI: 10.1002/cmr.1820030402.
- [45] Aaron T. Hess and Matthew D. Robson. “Hexagonal gradient scheme with RF spoiling improves spoiling performance for high-flip-angle fast gradient echo imaging”. In: *Magnetic Resonance in Medicine* 77.3 (2017), pp. 1231–1237. DOI: 10.1002/mrm.26213.
- [46] Shaihan J. Malik et al. “Spatially Resolved Extended Phase Graphs: Modeling and Design of Multipulse Sequences with Parallel Transmission”. In: *Magnetic Resonance in Medicine* 68.5 (Nov. 2012), pp. 1481–1494. DOI: 10.1002/mrm.24153.
- [47] Matthias Weigel et al. “Extended phase graphs with anisotropic diffusion”. In: *Journal of Magnetic Resonance* 205.2 (Aug. 2010), pp. 276–285. DOI: 10.1016/j.jmr.2010.05.011.
- [48] Jonathan Endres et al. “Phase distribution graphs for fast, differentiable, and spatially encoded Bloch simulations of arbitrary MRI sequences”. In: *Magnetic Resonance in Medicine* 92.3 (Sept. 2024), pp. 1189–1204. DOI: 10.1002/mrm.30055.
- [49] Deyun Wei, Qiwen Ran, and Yuanmin Li. “Generalized sampling expansion for bandlimited signals associated with the fractional fourier transform”. In: *IEEE Signal Processing Letters* 17.6 (2010), pp. 595–598. DOI: 10.1109/LSP.2010.2048642.
- [50] Klaas P. Pruessmann et al. “SENSE: Sensitivity encoding for fast MRI”. In: *Magnetic Resonance in Medicine* 42.5 (1999), pp. 952–962. DOI: 10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S.
- [51] Martin Uecker et al. “ESPIRiT - An eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA”. In: *Magnetic Resonance in Medicine* 71.3 (2014), pp. 990–1001. DOI: 10.1002/mrm.24751.
- [52] Mark A. Griswold et al. “Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)”. In: *Magnetic Resonance in Medicine* 47.6 (2002), pp. 1202–1210. DOI: 10.1002/mrm.10171.
- [53] Mark A. Griswold et al. “Partially parallel imaging with localized sensitivities (PILS)”. In: *Magnetic Resonance in Medicine* 44.4 (Oct. 2000), pp. 602–609. DOI: 10.1002/1522-2594(200010)44:4<602::AID-MRM14>3.0.CO;2-5.
- [54] Anagha Deshmane et al. “Parallel MR imaging”. In: *Journal of Magnetic Resonance Imaging* 36.1 (2012), pp. 55–72. DOI: 10.1002/jmri.23639.
- [55] David S. Taubman. “JPEG2000: Image Compression Fundamentals, Standards and Practice”. In: *Journal of Electronic Imaging* 11.2 (Apr. 2002), p. 286. DOI: 10.1117/1.1469618.

- [56] Michael Lustig, David L. Donoho, and John M. Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic Resonance in Medicine* 58.6 (2007), pp. 1182–1195. DOI: 10.1002/mrm.21391.
- [57] P. B. Roemer et al. “The NMR phased array”. In: *Magnetic Resonance in Medicine* 16.2 (1990), pp. 192–225. DOI: 10.1002/mrm.1910160203.
- [58] David O. Walsh, Arthur F. Gmitro, and Michael W. Marcellin. “Adaptive reconstruction of phased array MR imagery”. In: *Magnetic Resonance in Medicine* 43.5 (2000), pp. 682–690. DOI: 10.1002/(SICI)1522-2594(200005)43:5<682::AID-MRM10>3.0.CO;2-G.
- [59] Vladimir Jellus and A.R. Stephan Kannengiesser. “Adaptive Coil Combination Using a Body Coil Scan as Phase Reference”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 22 (2014), p. 4406.
- [60] Souheil J Inati, Michael S. Hansen, and Peter Kellman. “A Solution to the Phase Problem in Adaptive Coil Combination”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 21. 2013, p. 2672.
- [61] Souheil J. Inati, Michael S. Hansen, and Peter Kellman. “A Fast Optimal Method for Coil Sensitivity Estimation and Adaptive Coil Combination for Complex Images”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 22. 2014, p. 4407.
- [62] Tao Zhang et al. “Coil Compression for Accelerated Imaging with Cartesian Sampling”. In: *Magnetic Resonance in Medicine* 69 (2013), pp. 571–582. DOI: 10.1002/mrm.24267.
- [63] Martin Blaimer et al. “Virtual coil concept for improved parallel MRI employing conjugate symmetric signals”. In: *Magnetic Resonance in Medicine* 61.1 (2009), pp. 93–102. DOI: 10.1002/mrm.21652.
- [64] Martin Buehrer et al. “Array compression for MRI with large coil arrays”. In: *Magnetic Resonance in Medicine* 57.6 (2007), pp. 1131–1139. DOI: 10.1002/mrm.21237.
- [65] Zhi Pei Liang, E. Mark Haacke, and Cecil W. Thomas. “High-resolution inversion of finite Fourier transform data through a localised polynomial approximation”. In: *Inverse Problems* 5.5 (1989), pp. 831–847. DOI: 10.1088/0266-5611/5/5/011.
- [66] I. Dologlou, D. Van Ormondt, and G. Carayannis. “MRI scan time reduction through non-uniform sampling and SVD-based estimation”. In: *Signal Processing* 55.2 (1996), pp. 207–219. DOI: 10.1016/S0165-1684(96)00131-4.
- [67] Mathews Jacob, Merry P Mani, and Jong Chul Ye. “Structured Low-Rank Algorithms: Theory, Magnetic Resonance Applications, and Links to Machine Learning”. In: *IEEE Signal Processing Magazine* 37.1 (Jan. 2020), pp. 54–68. DOI: 10.1109/MSP.2019.2950432.
- [68] Francesco Padormo et al. “PRIMO: Precise radiofrequency inference from multiple observations”. In: *Magnetic Resonance in Medicine* 74.2 (Aug. 2015), pp. 372–383. DOI: 10.1002/mrm.25404.
- [69] Aaron T. Hess, Iulius Dragonu, and Mark Chiew. “Accelerated calibrationless parallel transmit mapping using joint transmit and receive low-rank tensor completion”. In: *Magnetic Resonance in Medicine* 86.5 (Nov. 2021), pp. 2454–2467. DOI: 10.1002/mrm.28880.

- [70] Stephen Boyd et al. “Distributed optimization and statistical learning via the alternating direction method of multipliers”. In: *Foundations and Trends in Machine Learning* 3.1 (2010), pp. 1–122. DOI: 10.1561/22000000016.
- [71] Andrew D. Scott, Jennifer Keegan, and David N. Firmin. “Motion in cardiovascular MR imaging”. In: *Radiology* 250.2 (2009), pp. 331–351. DOI: 10.1148/radiol.2502071998.
- [72] Hajime Sakuma et al. “Effect of breath holding on blood flow measurement using fast velocity encoded cine MRI”. In: *Magnetic Resonance in Medicine* 45.2 (2001), pp. 346–348. DOI: 10.1002/1522-2594(200102)45:2<346::AID-MRM1044>3.0.CO;2-I.
- [73] Agnes E. Holland, James W. Goldfarb, and Robert R. Edelman. “Diaphragmatic and cardiac motion during suspended breathing: Preliminary experience and implications for breath-hold MR imaging”. In: *Radiology* 209.2 (1998), pp. 483–489. DOI: 10.1148/radiology.209.2.9807578.
- [74] Robert W. Brown et al. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. 2nd. Wiley & Sons Ltd, 2014, pp. 1–1008. DOI: 10.1002/9781118633953.
- [75] Kamil Ugurbil. “Imaging at ultrahigh magnetic fields: History, challenges, and solutions”. In: *NeuroImage* 168 (2018), pp. 7–32. DOI: 10.1016/j.neuroimage.2017.07.007.
- [76] Essa Yacoub et al. “Imaging brain function in humans at 7 Tesla”. In: *Magnetic Resonance in Medicine* 45.4 (2001), pp. 588–594. DOI: 10.1002/mrm.1080.
- [77] Thomas F. Budinger and Mark D. Bird. “MRI and MRS of the human brain at magnetic fields of 14 T to 20 T: Technical feasibility, safety, and neuroscience horizons”. In: *NeuroImage* 168.January 2017 (2018), pp. 509–531. DOI: 10.1016/j.neuroimage.2017.01.067.
- [78] Ogan Ocali and Ergin Atalar. “Ultimate intrinsic signal-to-noise ratio in MRI”. In: *Magnetic Resonance in Medicine* 39.3 (1998), pp. 462–473. DOI: 10.1002/mrm.1910390317.
- [79] Caroline Le Ster et al. “Magnetic field strength dependent SNR gain at the center of a spherical phantom and up to 11.7T”. In: *Magnetic Resonance in Medicine* 88.5 (2022), pp. 2131–2138. DOI: 10.1002/mrm.29391.
- [80] Mark E. Ladd et al. “Pros and cons of ultra-high-field MRI/MRS for human application”. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 109 (Dec. 2018), pp. 1–50. DOI: 10.1016/j.pnmrs.2018.06.001.
- [81] Siegfried Trattnig et al. “Ultrahigh-field MRI: where it really makes a difference”. In: *Die Radiologie* (Aug. 2023), pp. 1–8. DOI: 10.1007/s00117-023-01184-x.
- [82] Giske Opheim et al. “7T Epilepsy Task Force Consensus Recommendations on the Use of 7T MRI in Clinical Practice”. In: *Neurology* 96.7 (Feb. 2021), pp. 327–341. DOI: 10.1212/WNL.00000000000011413.
- [83] Carlijn J.A. Tenbergen, Gregory J. Metzger, and Tom W.J. Scheenen. “Ultra-high-field MR in Prostate cancer: Feasibility and Potential”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 35.4 (2022), pp. 631–644. DOI: 10.1007/s10334-022-01013-7.

- [84] Matthijs H. S. de Buck. “Imaging Cerebrovascular Health using 7T MRI”. PhD thesis. University of Oxford, 2023.
- [85] Qing X. Yang et al. “Manipulation of image intensity distribution at 7.0 T: Passive RF shimming and focusing with dielectric materials”. In: *Journal of Magnetic Resonance Imaging* 24.1 (July 2006), pp. 197–202. DOI: 10.1002/jmri.20603.
- [86] Andrew G. Webb. “Dielectric materials in magnetic resonance”. In: *Concepts in Magnetic Resonance Part A: Bridging Education and Research* 38 A.4 (2011), pp. 148–184. DOI: 10.1002/cmr.a.20219.
- [87] Paul S. Jacobs et al. “In vivo B1+ enhancement of calf MRI at 7 T via optimized flexible metasurfaces”. In: *Magnetic Resonance in Medicine* 92.3 (Sept. 2024), pp. 1277–1289. DOI: 10.1002/mrm.30060.
- [88] R. Scott Staewen et al. “3-D FLASH Imaging Using a Single Surface Coil and a New Adiabatic Pulse, BIR-4”. In: *Investigative Radiology* 25.5 (May 1990), pp. 559–567. DOI: 10.1097/00004424-199005000-00015.
- [89] David C. Alsop, Thomas J. Connick, and Gabor Mizsei. “A spiral volume coil for improved rf field homogeneity at high static magnetic field strength”. In: *Magnetic Resonance in Medicine* 40.1 (1998), pp. 49–54. DOI: 10.1002/mrm.1910400107.
- [90] J. Thomas Vaughan et al. “Detunable transverse electromagnetic (TEM) volume coil for high-field NMR”. In: *Magnetic Resonance in Medicine* 47.5 (2002), pp. 990–1000. DOI: 10.1002/mrm.10141.
- [91] Victor Han and Chunlei Liu. “Multiphoton magnetic resonance in imaging: A classical description and implementation”. In: *Magnetic Resonance in Medicine* 84.3 (2020), pp. 1184–1197. DOI: 10.1002/mrm.28186.
- [92] John M. Drago et al. “Multiphoton parallel transmission (MP-pTx): Pulse design methods and numerical validation”. In: *Magnetic Resonance in Medicine* 92.4 (Oct. 2024), pp. 1376–1391. DOI: 10.1002/mrm.30116.
- [93] Harry H. Pennes. “Applied physiology”. In: *Journal of Applied Physiology* 1.2 (1948), pp. 93–122. DOI: 10.5005/jp/books/12678{_}10.
- [94] International Electrotechnical Commission (IEC) Committee. *IEC 60601-2-33. Medical Electrical Equipment - Part 2-33: Particular Requirements for the Basic Safety and Essential Performance of Magnetic Resonance Equipment for Medical Diagnosis*. Geneva, Switzerland, 2016.
- [95] David I. Hoult and D. Phil. “Sensitivity and power deposition in a high-field imaging experiment”. In: *Journal of Magnetic Resonance Imaging* 12.1 (2000), pp. 46–67. DOI: 10.1002/1522-2586(200007)12:1<46::AID-JMRI6>3.0.CO;2-D.
- [96] Tamer S. Ibrahim et al. “Application of finite difference time domain method for the design of birdcage RF head coils using multi-port excitations”. In: *Magnetic Resonance Imaging* 18.6 (2000), pp. 733–742. DOI: 10.1016/S0730-725X(00)00143-0.
- [97] Ulrich Katscher et al. “Transmit SENSE”. In: *Magnetic Resonance in Medicine* 49.1 (2003), pp. 144–150. DOI: 10.1002/mrm.10353.
- [98] Yudong Zhu. “Parallel Excitation with an Array of Transmit Coils”. In: *Magnetic Resonance in Medicine* 51.4 (2004), pp. 775–784. DOI: 10.1002/mrm.20011.

- [99] Francesco Padormo et al. “Parallel transmission for ultrahigh-field imaging”. In: *NMR in Biomedicine* 29.9 (Sept. 2016), pp. 1145–1161. DOI: 10.1002/nbm.3313.
- [100] Johannes T. Schneider et al. “Inner-volume imaging in vivo using three-dimensional parallel spatially selective excitation”. In: *Magnetic Resonance in Medicine* 69.5 (2013), pp. 1367–1378. DOI: 10.1002/mrm.24381.
- [101] Sydney Nicole Williams, Paul McElhinney, and Shajan Gunamony. “Ultra-high field MRI: parallel-transmit arrays and RF pulse design”. In: *Physics in Medicine & Biology* 68.2 (Jan. 2023), 02TR02. DOI: 10.1088/1361-6560/aca4b7.
- [102] Yudong Zhu et al. “System and SAR characterization in parallel RF transmission”. In: *Magnetic Resonance in Medicine* 67.5 (2012), pp. 1367–1378. DOI: 10.1002/mrm.23126.
- [103] Gabriele Eichfelder and Matthias Gebhardt. “Local specific absorption rate control for parallel transmission by virtual observation points”. In: *Magnetic Resonance in Medicine* 66.5 (2011), pp. 1468–1476. DOI: 10.1002/mrm.22927.
- [104] Ingmar Graesslin et al. “A specific absorption rate prediction concept for parallel transmission MR”. In: *Magnetic Resonance in Medicine* 68.5 (2012), pp. 1664–1674. DOI: 10.1002/mrm.24138.
- [105] Bastien Guerin, Elfar Adalsteinsson, and Lawrence L. Wald. “Local SAR reduction in multi-slice pTx via “SAR hopping” between excitations”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 20 (2012), p. 642.
- [106] Stephan Orzada et al. “RF excitation using time interleaved acquisition of modes (TIAMO) to address B1 inhomogeneity in high-field MRI”. In: *Magnetic Resonance in Medicine* 64.2 (Aug. 2010), pp. 327–333. DOI: 10.1002/mrm.22527.
- [107] Sascha Brunheim et al. “Fast and accurate multi-channel mapping based on the TIAMO technique for 7T UHF body MRI”. In: *Magnetic Resonance in Medicine* 79.5 (May 2018), pp. 2652–2664. DOI: 10.1002/mrm.26925.
- [108] William A. Grissom et al. “Spatial domain method for the design of RF pulses in multicoil parallel excitation”. In: *Magnetic Resonance in Medicine* 56.3 (2006), pp. 620–629. DOI: 10.1002/mrm.20978.
- [109] John M. Pauly, Dwight G. Nishimura, and Albert Macovski. “A k-space analysis of small-tip-angle excitation”. In: *Journal of Magnetic Resonance* 81.1 (Jan. 1989), pp. 43–56. DOI: 10.1016/0022-2364(89)90265-5.
- [110] John M. Pauly et al. “Parameter Relations for the Shinnar-Le Roux Selective Excitation Pulse Design Algorithm”. In: *IEEE Transactions on Medical Imaging* 10.1 (1991), pp. 53–65. DOI: 10.1109/42.75611.
- [111] Frank Ong et al. “iSLR : An Improved Shinnar-Le-Roux Frequency Selective Pulse Design Algorithm with Reduced Energy and More Accurate Phase Profiles”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 29. 2021, p. 0784.
- [112] William A. Grissom et al. “Additive angle method for fast large-tip-angle RF pulse design in parallel excitation”. In: *Magnetic Resonance in Medicine* 59.4 (2008), pp. 779–787. DOI: 10.1002/mrm.21510.
- [113] Dan Xu et al. “Designing multichannel, multidimensional, arbitrary flip angle RF pulses using an optimal control approach”. In: *Magnetic Resonance in Medicine* 59.3 (2008), pp. 547–560. DOI: 10.1002/mrm.21485.

- [114] Hai Zheng et al. “Improved large tip angle parallel transmission pulse design through a perturbation analysis of the bloch equation”. In: *Magnetic Resonance in Medicine* 66.3 (2011), pp. 687–696. DOI: 10.1002/mrm.22827.
- [115] A. Hoyos-Idrobo et al. “On variant strategies to solve the magnitude least squares optimization problem in parallel transmission pulse design and under strict sar and power constraints”. In: *IEEE Transactions on Medical Imaging* 33.3 (2014), pp. 739–748. DOI: 10.1109/TMI.2013.2295465.
- [116] Martijn Anton Cloos et al. “kT-points: Short three-dimensional tailored RF pulses for flip-angle homogenization over an extended volume”. In: *Magnetic Resonance in Medicine* 67.1 (Jan. 2012), pp. 72–80. DOI: 10.1002/mrm.22978.
- [117] Kawin Setsompop et al. “Slice-selective RF pulses for in vivo B1+ inhomogeneity mitigation at 7 tesla using parallel RF excitation with a 16-element coil”. In: *Magnetic Resonance in Medicine* 60.6 (Dec. 2008), pp. 1422–1432. DOI: 10.1002/mrm.21739.
- [118] Suwit Saekho et al. “Small tip angle three-dimensional tailored radiofrequency slab-select pulse for reduced B1 inhomogeneity at 3 T”. In: *Magnetic Resonance in Medicine* 53.2 (2005), pp. 479–484. DOI: 10.1002/mrm.20358.
- [119] Suwit Saekho et al. “Fast-kz three-dimensional tailored radiofrequency pulse for reduced B1 inhomogeneity”. In: *Magnetic Resonance in Medicine* 55.4 (Apr. 2006), pp. 719–724. DOI: 10.1002/mrm.20840.
- [120] Adam C. Zelinski et al. “Fast slice-selective radio-frequency excitation pulses for mitigating B1+ inhomogeneity in the human brain at 7 Tesla”. In: *Magnetic Resonance in Medicine* 59.6 (2008), pp. 1355–1364. DOI: 10.1002/mrm.21585.
- [121] Shaihan J. Malik et al. “Tailored excitation in 3D with spiral nonselective (SPINS) RF pulses”. In: *Magnetic Resonance in Medicine* 67.5 (2012), pp. 1303–1315. DOI: 10.1002/mrm.23118.
- [122] Vincent Gras et al. “Universal pulses: A new concept for calibration-free parallel transmission”. In: *Magnetic Resonance in Medicine* 77.2 (Feb. 2017), pp. 635–643. DOI: 10.1002/mrm.26148.
- [123] Simon Schmidt, Xiaoxuan He, and Gregory J Metzger. “Universal modes: Calibration-free time-interleaved acquisition of modes”. In: *Magnetic Resonance in Medicine* 92.1 (July 2024), pp. 43–56. DOI: 10.1002/mrm.30032.
- [124] Christoph Stefan Aigner et al. “Calibration-free pTx of the human heart at 7T via 3D universal pulses”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 29. 2021, p. 0777.
- [125] Christoph Stefan Aigner et al. “Calibration-free pTx of the human heart at 7T via 3D universal pulses”. In: *Magnetic Resonance in Medicine* 87.1 (Jan. 2022), pp. 70–84. DOI: 10.1002/mrm.28952.
- [126] Jürgen Herrler et al. “Fast online-customized (FOCUS) parallel transmission pulses: A combination of universal pulses and individual optimization”. In: *Magnetic Resonance in Medicine* 85.6 (2021), pp. 3140–3153. DOI: 10.1002/mrm.28643.
- [127] Jürgen Herrler et al. “The effects of RF coils and SAR supervision strategies for clinically applicable nonselective parallel-transmit pulses at 7 T”. In: *Magnetic Resonance in Medicine* 89.5 (2023), pp. 1888–1900. DOI: 10.1002/mrm.29569.

- [128] Wyger M Brink et al. “Ventricular B1+ perturbation at 7 T – real effect or measurement artifact?” In: *NMR in Biomedicine* April (2014), pp. 617–620. DOI: 10.1002/nbm.3112.
- [129] Ulrich Katscher et al. “Determination of Electric Conductivity and Local SAR Via B1 Mapping”. In: *IEEE Transactions on Medical Imaging* 28.9 (2009), pp. 1365–1374. DOI: 10.1109/TMI.2009.2015757.
- [130] Pierre-François Van De Moortele et al. “Calibration Tools for RF Shim at Very High Field with Multiple Element RF Coils : from Ultra Fast Local Relative Phase to Absolute Magnitude B1+ Mapping”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 15. 2007, p. 1676.
- [131] Francesco Padormo et al. “Large dynamic range relative B1+ mapping”. In: *Magnetic Resonance in Medicine* 76.2 (2016), pp. 490–499. DOI: 10.1002/mrm.25884.
- [132] Pierre-François Van de Moortele and Kamil Ugurbil. “Very Fast Multi Channel B1 Calibration at High Field in the Small Flip Angle Regime”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 17. 2009, p. 367.
- [133] Francesco Padormo, Shaihan J. Malik, and Joseph V. Hajnal. “Relative B1+ mapping directly from k-space for rapid Multi-Transmit calibration”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 20. 2012, p. 604.
- [134] Felix Krueger et al. “Rapid estimation of 2D relative B1+-maps from localizers in the human heart at 7T using deep learning”. In: *Magnetic Resonance in Medicine* 89.3 (Mar. 2023), pp. 1002–1015. DOI: 10.1002/mrm.29510.
- [135] Glen R. Morrell and Matthias C. Schabel. “An analysis of the accuracy of magnetic resonance flip angle measurement methods”. In: *Physics in Medicine and Biology* 55.20 (Oct. 2010), pp. 6157–6174. DOI: 10.1088/0031-9155/55/20/008.
- [136] Rolf Pohmann and Klaus Scheffler. “A theoretical and experimental comparison of different techniques for B 1 mapping at very high fields”. In: *NMR in Biomedicine* 26.3 (Mar. 2013), pp. 265–275. DOI: 10.1002/nbm.2844.
- [137] E.K. Insko and L. Bolinger. “Mapping of the Radiofrequency Field”. In: *Journal of Magnetic Resonance* 103.1 (June 1993), pp. 82–85. DOI: 10.1006/jmra.1993.1133.
- [138] Rudolf Stollberger and Paul Wach. “Imaging of the active B1 field in vivo”. In: *Magnetic Resonance in Medicine* 35.2 (1996), pp. 245–251. DOI: 10.1002/mrm.1910350217.
- [139] Charles H. Cunningham, John M. Pauly, and Krishna S. Nayak. “Saturated double-angle method for rapid B1+ mapping”. In: *Magnetic Resonance in Medicine* 55.6 (June 2006), pp. 1326–1333. DOI: 10.1002/mrm.20896.
- [140] Dingxin Wang, Sven Zuehlsdorff, and Andrew C. Larson. “Rapid 3D radiofrequency field mapping using catalyzed double-angle method”. In: *NMR in Biomedicine* 22.8 (Oct. 2009), pp. 882–890. DOI: 10.1002/nbm.1403.
- [141] Shaihan J. Malik, Gavin D. Kenny, and Joseph V. Hajnal. “Slice profile correction for transmit sensitivity mapping using actual flip angle imaging”. In: *Magnetic Resonance in Medicine* 65.5 (2011), pp. 1393–1399. DOI: 10.1002/mrm.22739.

- [142] Vasily L. Yarnykh. “Actual flip-angle imaging in the pulsed steady state: A method for rapid three-dimensional mapping of the transmitted radiofrequency field”. In: *Magnetic Resonance in Medicine* 57.1 (Jan. 2007), pp. 192–200. DOI: 10.1002/mrm.21120.
- [143] Vasily L. Yarnykh. “Optimal radiofrequency and gradient spoiling for improved accuracy of T1 and B1 measurements using fast steady-state techniques”. In: *Magnetic Resonance in Medicine* 63.6 (2010), pp. 1610–1626. DOI: 10.1002/mrm.22394.
- [144] Marco Andrea Zampini et al. “A preparation pulse for fast steady state approach in Actual Flip angle Imaging”. In: *Medical Physics* 51.1 (2024), pp. 306–318. DOI: 10.1002/mp.16624.
- [145] Kay Nehrke. “On the steady-state properties of actual flip angle imaging (AFI)”. In: *Magnetic Resonance in Medicine* 61.1 (Jan. 2009), pp. 84–92. DOI: 10.1002/mrm.21592.
- [146] Alexey A. Samsonov and Vasily L. Yarnykh. “Accurate actual flip angle imaging (AFI) in the presence of fat”. In: *Magnetic Resonance in Medicine* 91.6 (June 2024), pp. 2345–2357. DOI: 10.1002/mrm.30000.
- [147] Sebastian Dietrich et al. “3D Radial Phase Encoded Flip Angle Imaging at Ultra-High Field Strength”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 27. 2019, p. 4560.
- [148] Lin Wei and Kwon Song Hee. “Improved signal spoiling in fast radial gradient-echo imaging: Applied to accurate t1 mapping and flip angle correction”. In: *Magnetic Resonance in Medicine* 62.5 (2009), pp. 1185–1194. DOI: 10.1002/mrm.22089.
- [149] Armin M. Nagel et al. “Sodium MRI using a density-adapted 3D radial acquisition technique”. In: *Magnetic Resonance in Medicine* 62.6 (2009), pp. 1565–1573. DOI: 10.1002/mrm.22157.
- [150] Marta B. Maggioni, Martin Krämer, and Jürgen R. Reichenbach. “Optimized gradient spoiling of UTE VFA-AFI sequences for robust T1 estimation with B1-field correction”. In: *Magnetic Resonance Imaging* 82 (2021), pp. 1–8. DOI: 10.1016/j.mri.2021.06.011.
- [151] Nico Egger et al. “Accelerated B1+ mapping and robust parallel transmit pulse design for heart and prostate imaging at 7 T”. In: *Magnetic Resonance in Medicine* 92.5 (June 2024), pp. 1933–1951. DOI: 10.1002/mrm.30185.
- [152] Kay Nehrke and Peter Börnert. “DREAM-a novel approach for robust, ultrafast, multislice B1 mapping”. In: *Magnetic Resonance in Medicine* 68.5 (Nov. 2012), pp. 1517–1526. DOI: 10.1002/mrm.24158.
- [153] Kay Nehrke et al. “Volumetric B1+ Mapping of the Brain at 7T using DREAM”. In: *Magnetic Resonance in Medicine* 71.1 (Jan. 2014), pp. 246–256. DOI: 10.1002/mrm.24667.
- [154] Jens Frahm et al. “Stimulated echo imaging”. In: *Journal of Magnetic Resonance* 64.1 (Aug. 1985), pp. 81–93. DOI: 10.1016/0022-2364(85)90033-2.

- [155] Jens Frahm et al. “Rapid NMR imaging using stimulated echoes”. In: *Journal of Magnetic Resonance* 65.1 (Oct. 1985), pp. 130–135. DOI: 10.1016/0022-2364(85)90381-6.
- [156] Kay Nehrke and Peter Börnert. “iDREAM: T1 -bias compensated DREAM B1+ mapping using adiabatic inversion pulses”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 30 (2022), p. 3314.
- [157] Gunther Helms, Hampus Olsson, and Mads Andersen. “Removing bias and increasing dynamic range in DREAM flip angle mapping at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 27. 2019, p. 4500.
- [158] Hampus Olsson, Mads Andersen, and Gunther Helms. “Reducing bias in DREAM flip angle mapping in human brain at 7T by multiple preparation flip angles”. In: *Magnetic Resonance Imaging* 72 (Oct. 2020), pp. 71–77. DOI: 10.1016/j.mri.2020.07.002.
- [159] Peter Börnert, Kay Nehrke, and Wang Jinnan. “Magnetization prepared DREAM for fast flow-robust B1+ mapping”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 22. 2014, p. 4338.
- [160] Stefan E. Fischer et al. “Limitations of stimulated echo acquisition mode (STEAM) techniques in cardiac applications”. In: *Magnetic Resonance in Medicine* 34.1 (1995), pp. 80–91. DOI: 10.1002/mrm.1910340113.
- [161] Teresa Rincón-Domínguez et al. “Accelerated multi-snapshot free-breathing B1+ mapping based on the dual refocusing echo acquisition mode technique (DREAM): An alternative to measure RF nonuniformity for cardiac MRI”. In: *Journal of Magnetic Resonance Imaging* 49.2 (Feb. 2019), pp. 499–507. DOI: 10.1002/jmri.26234.
- [162] Kay Nehrke et al. “Fast B1+ Mapping for Cardiac MR using a Black Blood DREAM Sequence”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 21 (2013), p. 4271.
- [163] Philipp Ehse et al. “Whole-brain B1-mapping using three-dimensional DREAM”. In: *Magnetic Resonance in Medicine* 82.3 (Apr. 2019), pp. 924–934. DOI: 10.1002/mrm.27773.
- [164] Marten Veldmann et al. “Spiral 3DREAM sequence for fast whole-brain B1 Mapping”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 30. 2022, p. 1612.
- [165] Svenja Niesen et al. “Spiral 3DREAM sequence for fast whole-brain B1 mapping”. In: *Magnetic Resonance in Medicine* 93.1 (Jan. 2025), pp. 321–329. DOI: 10.1002/mrm.30282.
- [166] Florent Eggenschwiler et al. “SA2RAGE: A new sequence for fast B1+-mapping”. In: *Magnetic Resonance in Medicine* 67.6 (2012), pp. 1609–1619. DOI: 10.1002/mrm.23145.
- [167] Riccardo Metere et al. “Fast Measurement of B1+ using the NO2RAGE pulse sequence”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 2018.
- [168] Gabriele Bonanno et al. “Whole-brain 3D B1+ mapping in under 30 seconds: compressed-sensing accelerated SA2RAGE”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 31. 2023, p. 1423.

- [169] Hampus Olsson et al. “MP3RAGE: Simultaneous mapping of T1 and B1+ in human brain at 7T”. In: *Magnetic Resonance in Medicine* 87.6 (2022), pp. 2637–2649. DOI: 10.1002/mrm.29151.
- [170] Uwe Klose. “Mapping of the Radio Frequency Magnetic Field With A MR Snapshot Flash Technique”. In: *Medical Physics* 19.4 (1992), pp. 1099–1104. DOI: 10.1118/1.596828.
- [171] Sohae Chung et al. “Rapid B1+ mapping using a preconditioning RF pulse with TurboFLASH readout”. In: *Magnetic Resonance in Medicine* 64.2 (Aug. 2010), pp. 439–446. DOI: 10.1002/mrm.22423.
- [172] Dario Bosch et al. “Rapid pre-saturated TFL transmit field mapping with an optimized 3D centric single-shot readout”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2020, p. 3703.
- [173] Jan Sedlacik et al. “Calibration of saturation prepared turbo FLASH B1+ maps by actual flip angle imaging at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 30. 2022, p. 2867. DOI: <https://doi.org/10.58530/2022/2867>.
- [174] Michael Dubiner et al. “A Machine-Learning Approach for Saturation-Prepared Turbo FLASH B1+ Maps Calibration”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 31. 2023, p. 0953.
- [175] Dario Bosch et al. “Optimized ultrahigh field parallel transmission workflow using rapid presaturated TurboFLASH transmit field mapping with a three-dimensional centric single-shot readout”. In: *Magnetic Resonance in Medicine* December 2021 (2022), pp. 322–330. DOI: 10.1002/mrm.29459.
- [176] Dan Ma et al. “Magnetic resonance fingerprinting”. In: *Nature* 495.7440 (2013), pp. 187–192. DOI: 10.1038/nature11971.
- [177] Guido Buonincontri et al. “Spiral MR fingerprinting at 7 T with simultaneous B1 estimation”. In: *Magnetic Resonance Imaging* 41 (Sept. 2017), pp. 1–6. DOI: 10.1016/j.mri.2017.04.003.
- [178] Jiaen Liu et al. “Electrical Properties Tomography Based on B1 Maps in MRI: Principles, Applications, and Challenges”. In: *IEEE Transactions on Biomedical Engineering* 64.11 (2017), pp. 2515–2530. DOI: 10.1109/TBME.2017.2725140.
- [179] Max Lutz et al. “Low power free-breathing absolute B1+ mapping in the human body at 7T using magnetic resonance fingerprinting”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 2022, p. 0386. DOI: 10.58530/2022/0386.
- [180] Max Lutz et al. “MRF-based channel-wise absolute B 1 + mapping at low RF power in the human abdomen at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 31. 2023, p. 1420.
- [181] Max Lutz et al. “B1-MRF: Large dynamic range MRF-based absolute B1+ mapping in the human body at 7T”. In: *Magnetic Resonance in Medicine* 92 (2024), pp. 2473–2490. DOI: 10.1002/mrm.30242.
- [182] Glen R. Morrell. “A phase-sensitive method of flip angle mapping”. In: *Magnetic Resonance in Medicine* 60.4 (2008), pp. 889–894. DOI: 10.1002/mrm.21729.
- [183] Yulin V. Chang. “Rapid B1 mapping using orthogonal, equal-amplitude radio-frequency pulses”. In: *Magnetic Resonance in Medicine* 67.3 (Mar. 2012), pp. 718–723. DOI: 10.1002/mrm.23051.

- [184] Jie Wen, Alexander L. Sukstanskii, and Dmitriy A. Yablonskiy. “Phase-sensitive B1 mapping: Effects of relaxation and RF spoiling”. In: *Magnetic Resonance in Medicine* 80.1 (July 2018), pp. 101–111. DOI: 10.1002/mrm.27009.
- [185] F. Bloch and A. Siegert. “Magnetic Resonance for Nonrotating Fields”. In: *Physical Review* 57.6 (Mar. 1940), pp. 522–527. DOI: 10.1103/PhysRev.57.522.
- [186] Laura I. Sacolick et al. “B1 mapping by Bloch-Siegert shift”. In: *Magnetic Resonance in Medicine* 63.5 (2010), pp. 1315–1322. DOI: 10.1002/mrm.22357.
- [187] Mohammad-Mehdi Khalighi, Laura I. Sacolick, and Brian K. Rutt. “Signal to Noise Ratio Analysis of Bloch-Siegert B1+ Mapping”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 18. 2010, p. 2842.
- [188] Mohammad-Mehdi Khalighi et al. “Fast and Robust B1 Mapping at 7T by the Bloch-Siegert Method”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 18. 2010, p. 2831.
- [189] Patrik Wyss et al. “Interferometric Bloch-Siegert B1+ mapping at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 21. 2013, p. 2594.
- [190] Xiaoying Cai, Wu Dan, and Jiangyang Zhang. “Three dimensional fast spin echo Bloch-Siegert B1 mapping with navigator based phase correction at 11.7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 21. 2013, p. 4268.
- [191] Kay Nehrke and Peter Boernert. “Fast B1 Mapping using a STEAM-based Bloch-Siegert Preparation Pulse”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 19. 2011, p. 4411.
- [192] Kalina V. Jordanova, Dwight G. Nishimura, and Adam B. Kerr. “B1 estimation using adiabatic refocusing: BEAR”. In: *Magnetic Resonance in Medicine* 72.5 (2014), pp. 1302–1310. DOI: 10.1002/mrm.25049.
- [193] Jinghua Wang et al. “Factors influencing flip angle mapping in MRI: RF pulse shape, slice-select gradients, off-resonance excitation, and B0 inhomogeneities”. In: *Magnetic Resonance in Medicine* 56.2 (2006), pp. 463–468. DOI: 10.1002/mrm.20947.
- [194] Gabriela Belsley et al. “The effect of and correction for through-slice dephasing on 2D gradient-echo double angle B1+ mapping”. In: *Magnetic Resonance in Medicine* 91.4 (Apr. 2024), pp. 1598–1607. DOI: 10.1002/mrm.29966.
- [195] Alexis Amadon et al. “Slice profile corrections in the XFL (magnetization-prepared turbo-FLASH) B1-mapping sequence”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 23. 2015, p. 2377.
- [196] Matthijs H. S. de Buck et al. “Head-and-neck multichannel B1+ mapping and RF shimming of the carotid arteries using a 7T parallel-transmit head coil”. In: *Magnetic Resonance in Medicine* 91.1 (Jan. 2024), pp. 190–204. DOI: 10.1002/mrm.29845.
- [197] Hans-Peter Fautz et al. “B1 mapping of coil arrays for parallel transmission”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 16. 2008, p. 1247.
- [198] Shaihan J. Malik, David J. Larkman, and Joseph V. Hajnal. “Optimal linear combinations of array elements for B1 mapping”. In: *Magnetic Resonance in Medicine* 62.4 (2009), pp. 902–909. DOI: 10.1002/mrm.22068.

- [199] David Otto Brunner and Klaas P. Pruessmann. “B1+ interferometry for the calibration of RF transmitter arrays”. In: *Magnetic Resonance in Medicine* 61.6 (June 2009), pp. 1480–1488. DOI: 10.1002/mrm.21893.
- [200] Desmond H.Y. Tse et al. “Encoding methods for B1+ mapping in parallel transmit systems at ultrahigh field”. In: *Journal of Magnetic Resonance* 245 (Aug. 2014), pp. 125–132. DOI: 10.1016/j.jmr.2014.06.006.
- [201] Kay Nehrke and Peter Börnert. “Eigenmode Analysis of Transmit Coil Array for Tailored B1 Mapping”. In: *Magnetic Resonance in Medicine* 63 (2010), pp. 754–764. DOI: 10.1002/mrm.22239.
- [202] Bobby A Runderkamp et al. “Whole-liver flip-angle shimming at 7 T using parallel-transmit kT-point pulses and Fourier phase-encoded DREAM B1+ mapping”. In: *Magnetic Resonance in Medicine* 91.1 (Jan. 2024), pp. 75–90. DOI: 10.1002/mrm.29819.
- [203] Sascha Brunheim et al. “Erratum to: Fast and accurate multi-channel B1+ mapping based on the TIAMO technique for 7 Tesla UHF body MRI (Magn Reson Med 2018;79:2652–2664)”. In: *Magnetic Resonance in Medicine* (Dec. 2023), p. 1. DOI: 10.1002/mrm.29868.
- [204] David I. Hoult and R.E Richards. “The signal-to-noise ratio of the nuclear magnetic resonance experiment”. In: *Journal of Magnetic Resonance* 24.1 (Oct. 1976), pp. 71–85. DOI: 10.1016/0022-2364(76)90233-X.
- [205] Sebastian Weingärtner et al. “Motion-robust cardiac B1+ mapping at 3T using interleaved Bloch-Siegert shifts”. In: *Magnetic Resonance in Medicine* 78.2 (2017), pp. 670–677. DOI: 10.1002/mrm.26395.
- [206] Yuehui Tao et al. “B0-Independent quantitative measurement of low B1 field for human cardiac MRI at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 21. 2013, p. 2596.
- [207] Philipp Ehse. *mapVBVD*. URL: <https://github.com/pehses/mapVBVD>.
- [208] Mark Jenkinson et al. “FSL”. In: *NeuroImage* 62.2 (Aug. 2012), pp. 782–790. DOI: 10.1016/j.neuroimage.2011.09.015.
- [209] Swati D. Rane and John C. Gore. “Measurement of T1 of human arterial and venous blood at 7T”. In: *Magnetic Resonance Imaging* 31.3 (2013), pp. 477–479. DOI: 10.1016/j.mri.2012.08.008.
- [210] Dan Zhu, Michael Schär, and Qin Qin. “Ultrafast B1 mapping with RF-prepared 3D FLASH acquisition: Correcting the bias due to T1-induced k-space filtering effect”. In: *Magnetic Resonance in Medicine* 88.2 (Aug. 2022), pp. 757–769. DOI: 10.1002/mrm.29247.
- [211] Trevor Wade and Brian K. Rutt. “Comparison of Current B1-Mapping Techniques”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 15. 2007, p. 0354.
- [212] Daniel J. Park et al. “Statistical Analysis of B1 Mapping Techniques”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 19. 2011, p. 4424.
- [213] Daniel J. Park et al. “A statistical analysis of the Bloch-Siegert B1 mapping technique”. In: *Physics in Medicine and Biology* 58.16 (2013), pp. 5673–5691. DOI: 10.1088/0031-9155/58/16/5673.

- [214] Matthijs H. S. de Buck, Aaron T. Hess, and Peter Jezzard. *Head-and-neck multi-channel 7T pTx B1+ database (10 subjects) for a Nova Medical 8Tx/32Rx coil*. 2022. DOI: 10.5287/bodleian:ZB6Gk8QzN.
- [215] Shaihan J. Malik, Rui Pedro A.G. Teixeira, and Joseph V. Hajnal. “Extended phase graph formalism for systems with magnetization transfer and exchange”. In: *Magnetic Resonance in Medicine* 80.2 (2018), pp. 767–779. DOI: 10.1002/mrm.27040.
- [216] Alexander J.E. Raaijmakers et al. “The fractionated dipole antenna: A new antenna for body imaging at 7 Tesla”. In: *Magnetic Resonance in Medicine* 75.3 (2016), pp. 1366–1374. DOI: 10.1002/mrm.25596.
- [217] Sebastian Schmitter et al. “Cardiac imaging at 7 tesla: Single- and two-spoke radiofrequency pulse design with 16-channel parallel excitation”. In: *Magnetic Resonance in Medicine* 70.5 (Nov. 2013), pp. 1210–1219. DOI: 10.1002/mrm.24935.
- [218] Christoph Stefan Aigner, Sebastian Dietrich, and Sebastian Schmitter. “Three-dimensional static and dynamic parallel transmission of the human heart at 7 T”. In: *NMR in Biomedicine* 34.3 (2021), pp. 1–15. DOI: 10.1002/nbm.4450.
- [219] Sebastian Dietrich et al. “3D Free-breathing multichannel absolute B1+ Mapping in the human body at 7T”. In: *Magnetic Resonance in Medicine* 85.5 (2021), pp. 2552–2567. DOI: 10.1002/mrm.28602.
- [220] Sven H. F. Jaeschke and Aaron T. Hess. “Cartesian prospective, dual-gated, cardio-respiratory encoding for 3D cardiac CINE compared to spiral profile retrospective encoding.” In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2020, p. 2173. DOI: 10.1002/mrm.27884.5..
- [221] Ali Aghaeifar and Klaus Scheffler. “RF Power Amplifier Drift Compensation for Reliable B1 Mapping Using Bloch-Siegert Technique”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 32. 2024, p. 1252.
- [222] Ali Aghaeifar et al. “Intra-scan RF power amplifier drift correction”. In: *Magnetic Resonance in Medicine* 92.2 (Aug. 2024), pp. 645–659. DOI: 10.1002/mrm.30078.
- [223] Maxim Zaitsev, Julian Maclaren, and Michael Herbst. “Motion artifacts in MRI: A complex problem with many partial solutions”. In: *Journal of Magnetic Resonance Imaging* 42.4 (Oct. 2015), pp. 887–901. DOI: 10.1002/jmri.24850.
- [224] Alix Plumley et al. “Rigid motion-resolved B1+ prediction using deep learning for real-time parallel-transmission pulse design”. In: *Magnetic Resonance in Medicine* 87.5 (2022), pp. 2254–2270. DOI: 10.1002/mrm.29132.
- [225] Emre Kopanoglu and Katherine Anna Blanter. “Redesigning Parallel-transmit Pulses in Runtime to Correct for the Effect of Patient Motion on B1+-maps for Ultrahigh-field MRI”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 32. 2024, p. 4093.
- [226] Jessica Schulz et al. “An embedded optical tracking system for motion-corrected magnetic resonance imaging at 7T”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 25.6 (2012), pp. 443–453. DOI: 10.1007/s10334-012-0320-0.

- [227] M. Zaitsev et al. “Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system”. In: *NeuroImage* 31.3 (2006), pp. 1038–1050. DOI: 10.1016/j.neuroimage.2006.01.039.
- [228] Daniel Stucht et al. “Highest Resolution In Vivo Human Brain MRI Using Prospective Motion Correction”. In: *PLOS ONE* 10.7 (July 2015). Ed. by Jan Kassubek, pp. 1–17. DOI: 10.1371/journal.pone.0133921.
- [229] Lei Qin et al. “Prospective head-movement correction for high-resolution MRI using an in-bore optical tracking system”. In: *Magnetic Resonance in Medicine* 62.4 (2009), pp. 924–934. DOI: 10.1002/mrm.22076.
- [230] Laura Bortolotti and Richard Bowtell. “Simulating the use of active magnetic markers for motion correction using NMR field probes on a 7T scanner”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 29. 2021, p. 1365.
- [231] André J.W. Van Der Kouwe, Thomas Benner, and Anders M. Dale. “Real-time rigid body motion correction and shimming using cloverleaf navigators”. In: *Magnetic Resonance in Medicine* 56.5 (2006), pp. 1019–1032. DOI: 10.1002/mrm.21038.
- [232] Daniel Gallichan, José P. Marques, and Rolf Gruetter. “Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T”. In: *Magnetic Resonance in Medicine* 75.3 (2016), pp. 1030–1039. DOI: 10.1002/mrm.25670.
- [233] Tess E. Wallace et al. “Head motion measurement and correction using FID navigators”. In: *Magnetic Resonance in Medicine* 81.1 (Jan. 2019), pp. 258–274. DOI: 10.1002/mrm.27381.
- [234] Tess E. Wallace et al. “Rapid measurement and correction of spatiotemporal B0 field changes using FID navigators and a multi-channel reference image”. In: *Magnetic Resonance in Medicine* 83.2 (2020), pp. 575–589. DOI: 10.1002/mrm.27957.
- [235] Tess E. Wallace et al. “Dynamic distortion correction for functional MRI using FID navigators”. In: *Magnetic Resonance in Medicine* 85.3 (2021), pp. 1294–1307. DOI: 10.1002/mrm.28505.
- [236] Lucilio Cordero-Grande et al. “Motion-corrected MRI with DISORDER: Distributed and incoherent sample orders for reconstruction deblurring using encoding redundancy”. In: *Magnetic Resonance in Medicine* 84.2 (2020), pp. 713–726. DOI: 10.1002/mrm.28157.
- [237] Daniel Polak et al. “Scout accelerated motion estimation and reduction (SAMER)”. In: *Magnetic Resonance in Medicine* 87.1 (2022), pp. 163–178. DOI: 10.1002/mrm.28971.
- [238] D. Buikman, T. Helzel, and P. Röschmann. “The rf coil as a sensitive motion detector for magnetic resonance imaging”. In: *Magnetic Resonance Imaging* 6.3 (1988), pp. 281–289. DOI: 10.1016/0730-725X(88)90403-1.
- [239] Sven H. F. Jaeschke, Matthew D. Robson, and Aaron T. Hess. “Cardiac gating using scattering of an 8-channel parallel transmit coil at 7T”. In: *Magnetic Resonance in Medicine* 80.2 (2018), pp. 633–640. DOI: 10.1002/mrm.27038.

- [240] Sven H. F. Jaeschke. “Motion Estimation Using Radiofrequency Scattering of a Parallel Transmit Coil for Cardiac Magnetic Resonance Imaging at 7T”. PhD thesis. 2019, pp. 1–177.
- [241] Sven H. F. Jaeschke, Matthew D. Robson, and Aaron T. Hess. “Scattering matrix imaging pulse design for real-time respiration and cardiac motion monitoring”. In: *Magnetic Resonance in Medicine* 82.6 (2019), pp. 2169–2177. DOI: 10.1002/mrm.27884.
- [242] Aaron T. Hess et al. “Diaphragm position can be accurately estimated from the scattering of a parallel transmit RF coil at 7 T”. In: *Magnetic Resonance in Medicine* 79.4 (2018), pp. 2164–2169. DOI: 10.1002/mrm.26866.
- [243] Hans Hoogduin et al. “Motion detection using reflected signals from an eight channel parallel transmit head coil at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 28. 2020, p. 1280.
- [244] Daniel Papp et al. “Simultaneous detection of cardiac, respiratory and rigid body head motion using the scattering of a parallel transmit RF coil at 7T”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 26. 2018, p. 1168.
- [245] Peter Speier, Matthias Fenchel, and R. Rehner. “PT-Nav: a novel respiratory navigation method for continuous acquisitions based on modulation of a pilot tone in the MR-receiver”. In: *Proc. ESMRMB*. Vol. 28. 128. 2015, pp. 97–98. DOI: 10.1007/s10334-015-0487-2.
- [246] Juliane Ludwig et al. “Pilot tone-based motion correction for prospective respiratory compensated cardiac cine MRI”. In: *Magnetic Resonance in Medicine* 85.5 (2021), pp. 2403–2416. DOI: 10.1002/mrm.28580.
- [247] Matan Gavish and David L. Donoho. “The Optimal Hard Threshold for Singular Values is $4/\sqrt{3}$ ”. In: *IEEE Transactions on Information Theory* 60.8 (Aug. 2014), pp. 5040–5053. DOI: 10.1109/TIT.2014.2323359.
- [248] Martin Uecker. *BART Toolbox for Computational Magnetic Resonance Imaging*. DOI: 10.5281/zenodo.592960.
- [249] Stephen M. Smith et al. “Advances in functional and structural MR image analysis and implementation as FSL”. In: *NeuroImage* 23.SUPPL. 1 (2004), pp. 208–219. DOI: 10.1016/j.neuroimage.2004.07.051.
- [250] Mark W. Woolrich et al. “Bayesian analysis of neuroimaging data in FSL”. In: *NeuroImage* 45.1 (Mar. 2009), S173–S186. DOI: 10.1016/j.neuroimage.2008.10.055.
- [251] Mark Jenkinson et al. “Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images”. In: *NeuroImage* 17.2 (2002), pp. 825–841. DOI: 10.1006/nimg.2002.1132.
- [252] Mark Jenkinson and Stephen Smith. “A global optimisation method for robust affine registration of brain images”. In: *Medical Image Analysis* 5.2 (2001), pp. 143–156. DOI: 10.1016/S1361-8415(01)00036-6.
- [253] Peter Speier and Mario Bacher. “Skip the Electrodes, But Not A Beat: The Engineering Behind the Beat Sensor”. In: *MAGNETOM Flash* 84 (2023), pp. 106–117.

- [254] Hugh Simmons, Boris Mailhe, and Aaron T. Hess. “Relationship between Head Motion and Scattering: in-Silico Evaluation of Linearity”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 32. 2024, p. 2658.
- [255] Niek R.F. Huttinga et al. “Three-dimensional rigid head motion correction using the Beat Pilot Tone and Gaussian Processes”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 31. 2023, p. 1019.
- [256] Suma Anand and Michael Lustig. “Beat Pilot Tone : Exploiting Preamplifier Intermodulation of UHF / SHF RF for Improved Motion Sensitivity over Pilot Tone Navigators”. In: *Proc. Intl. Soc. Mag. Reson. Med.* 2021.
- [257] Suma Anand and Michael Lustig. “Beat Pilot Tone (BPT): Simultaneous MRI and RF motion sensing at arbitrary frequencies”. In: *Magnetic Resonance in Medicine* 92.4 (2024), pp. 1768–1787. DOI: 10.1002/mrm.30150.
- [258] Yannick Brackenier et al. “Sequence-agnostic motion-correction leveraging efficiently calibrated Pilot Tone signals”. In: *Magnetic Resonance in Medicine* 92.5 (Nov. 2024), pp. 1881–1897. DOI: 10.1002/mrm.30161.
- [259] Francesco Padormo. “Advanced Methods for Mapping the Radiofrequency Magnetic Fields in MRI”. PhD thesis. 2012, pp. 1–243.