

**Understanding and overcoming head
motion in ultra-high field Magnetic
Resonance Imaging with parallel
radio-frequency transmission.**

Alix J. D. Plumley

This thesis is being submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

School of Psychology
Cardiff University — Prifysgol Caerdydd
November 2022

Acknowledgements

I feel very lucky to be able to say that I have really enjoyed the past four years, and I've found a lot of fulfilment in my research. I know for certain that it could have been a much less enjoyable experience without such an encouraging, patient, supportive, kind, and fun supervisor. Thank you Emre.

Thank you Kevin, for being the Dumbledore character - wise, kind and understanding. I really appreciate your support and guidance.

I am grateful to everyone who believed in me enough to give me the opportunity to do this PhD, despite my non-technical background.

Special thanks to the Rabin Ezra Scholarship Trust for your support. The generosity of your family's trust is inspiring.

My wonderful friends at CUBRIC have undoubtedly helped shape me into who I am. For the two years of remote working during the pandemic, I realised how integral the social side of work is to my enjoyment of it. I feel incredibly lucky to be able to work with so many awesome, like-minded and inspiring people (nerds, yay!). I hope I will catch up with lots of you at future ISMRMs and other meetings, wherever we all end up!

Mum, dad and Ruth - thank you for your support and love, and for always being there for me. I couldn't have done this without you. I've missed the warm glow of home over these past few months. I will always appreciate everything you do for me.

And last but not least - thank you Hoa, for supporting me, for cooking me delicious food after these long evenings of write-up, and for keeping our army of houseplant babies happy and watered while I sat, neglecting them, in front of my computer. You inspire me to stay positive and focussed. I'm so happy to spend my life with you.

Thesis Summary

Ultra-high field (UHF) magnetic resonance imaging (MRI) offers more signal-to-noise compared to most clinical systems, but clinical uptake of UHF MRI remains low, partly due to artificial signal contrast variations and a higher risk of undesirable tissue heating at UHF. Parallel RF transmission (pTx) is capable of overcoming both issues, but the implications of patient motion on signal and safety (i.e., specific absorption rate; SAR) when pTx is used are currently not well understood. The work in this thesis aims to better characterise these effects, and presents novel approaches to help overcome them.

The study chapters present investigations into firstly, the effects of motion on signal quality when different RF pulse types are used (Chapter 5); secondly, the effects of pTx coil dimensions on motion-related SAR changes (Chapter 6); and thirdly, the inter-subject variability of motion-related SAR changes in pTx (Chapter 7). Following this, two methods are outlined which aim to reduce the sensitivity of signal to head motion in pTx. These comprise firstly, a method which uses composite B1+ maps for pTx pulse design (Chapter 8), and secondly, a deep learning framework which can estimate B1+ maps following head motion (Chapter 9). Finally, the generalisability of the latter approach across coil models is explored (Chapter 10). Simulated data were used for all of the work presented.

Electromagnetic field distributions were generated using Sim4Life (Zurich MedTech, Zurich, Switzerland). RF pulse design and evaluation was conducted in MATLAB (The MathWorks Inc., Natick, MA).

Findings indicate that systematic differences in the signal and SAR behaviour arise when different RF pulse types and different pTx coils are used, respectively. On the other hand, differences were observed in the SAR sensitivity to motion across different virtual body models, but these were not clearly systematic. These findings indicate that new approaches are needed in order to guarantee good image quality and safety for pTx in cases of subject motion. The two proposed methods reduced the impact of motion on simulated signal profiles.

Contents

I	Background	1
1	MR Signal - basic concepts	2
1.1	The B_0 field and polarisation	2
1.2	Precession and the role of radiofrequency	4
1.3	Gradient fields and k-space	5
1.4	The B_1^+ field	7
1.5	Electric fields and safety	9
1.6	Pulse sequence basics: Relaxation and echoes	11
2	Ultra-High Field MRI	15
2.1	Practical benefits of UHF	16
2.2	Challenges of UHF	19
2.3	Strategies to overcome the non-uniformity at UHF	22
2.4	Conclusions	29
3	Parallel Transmission (pTx)	30
3.1	Basic principles of pTx	31

3.2	Small-tip pTx pulse design	32
3.3	Large-tip pTx pulse design	34
3.4	Parallel transmit implementations	36
3.5	SAR and safety in pTx	40
3.6	Electromagnetic simulations	42
4	Motion in MRI	44
4.1	Motion in MRI	44
4.2	Effects of motion on the signal	46
4.3	Motion at UHF with pTx	48
4.4	Correction techniques and their limitations	49
II	Study Chapters	53
5	Motion and large vs. small tip-angle pTx pulse performance	54
5.1	Introduction	54
5.2	Methods	55
5.3	Results	60
5.4	Discussion	77
6	Specific Absorption Rate and pTx coil design	80
6.1	Overview	80
6.2	Introduction	81

6.3	Methods	82
6.4	Results	85
6.5	Discussion	91
7	Specific Absorption Rate and inter-subject variability	96
7.1	Introduction	96
7.2	Methods	98
7.3	Results	99
7.4	Discussion	105
8	Approaches to reduce motion-sensitivity in pTx:	
	Composite B_1^+ maps	109
8.1	Introduction	109
8.2	Methods	110
8.3	Results	113
8.4	Discussion	127
9	Approaches to reduce motion-sensitivity in pTx:	
	Deep Learning B_1^+ prediction	130
9.1	Introduction	130
9.2	Methods	134
9.3	Results	142
9.4	Discussion	152

10 Network generalisability across coils	157
10.1 Introduction	157
10.2 Methods	159
10.3 Results	161
10.4 Discussion	168
III General Discussion and Concluding Comments	171
11 General Discussion	172
11.1 Motion and image quality in pTx: Concluding comments	173
11.2 Motion and safety in pTx: Concluding comments	174
11.3 The future of deep learning approaches	176
12 Appendix	178

Introduction

Since 1977, when the first whole-body in-vivo images were acquired with nuclear magnetic resonance, Magnetic Resonance Imaging (MRI) has grown to occupy a unique medical imaging niche which is exploited in everyday clinical workflows across the globe. Its excellent soft-tissue contrast and non-ionizing radiation are widely appreciated, and mean that common applications of MRI today include diagnosis and prognosis of pathology, injury assessment, surgical planning, identification of clinically-relevant biomarkers, and characterisation of brain structure, function, and development. Historically, MRI has suffered from low signal-to-noise ratio, meaning image resolution was coarse, scan durations long, and/or high levels of noise were present in images. These issues have been (partly) solved by hardware developments including a move towards systems which use stronger static magnetic fields (B_0).

For the purpose of this thesis, MRI with B_0 strength of 7 Tesla (7T) and above is considered to be ultra-high field (UHF). UHF MRI offers a host of benefits compared to current clinical standard systems (which are typically 1.5 - 3T). Neurological areas including (but not limited to) epilepsy [1–4], multiple sclerosis [5–9], brain tumour [10, 11], cerebrovascular disease [10], traumatic brain injury [10] and degenerative disease [12] have benefited from the increased sensitivity to subtle features and lesions which are often overlooked at lower field strengths. Non-neurological applications have also been explored, such as UHF cardiovascular [13, 14], musculoskeletal (MSK) [15], and whole-body [16] MRI. However so far, this has been largely limited to research contexts, while clinical uptake of UHF MRI remains slow.

One reason is that UHF MRI also suffers from artificial signal contrast variations and a higher risk of undesirable tissue heating [17]. Several partial solutions have been proposed, including the use of parallel RF transmission (pTx) hardware [18–20]. When used correctly, pTx can effectively overcome the signal (artificial contrast) and safety (locally-elevated power deposition) issues inherent to UHF MRI; however pTx hardware presents challenges of its own.

One of the least well-understood aspects of pTx is how patient motion affects performance, both in terms of signal and safety. Since MRI scan durations can be on the order of tens of minutes, this is especially problematic for populations who may have difficulty remaining still, such as paediatric or elderly patients, or patients suffering from tremors. The focus of this thesis is to better characterise these effects, and to describe novel software-based approaches to help overcome them.

The initial sections are dedicated to describing the origin of the signal in MRI, and providing the information necessary to understand how this relates to the problems experienced in UHF MRI. Several candidate solutions and their limitations are briefly discussed, before focusing on pTx in more depth. The issue of motion is discussed, with emphasis on existing motion-correction techniques in MRI, and their limitations. In the study chapters, the effects of motion on pTx are investigated and characterised, with attention paid to different use cases, different pTx coil designs, and different patients. Following this, several chapters are presented, outlining frameworks to address these issues. Finally, conclusions are drawn and recommendations are made for improving pTx performance in cases where subject motion occurs. Throughout the thesis, focus is primarily on signal-related issues, however some attention is also paid to safety-related issues, since these are particularly pertinent for UHF MRI. Although each chapter is designed to be somewhat standalone, the large degree of thematic overlap means that cross-referencing between chapters is generally preferred over excessive repetition.

Part I

Background

Chapter 1

MR Signal - basic concepts

The data acquisition and processing pipeline in MRI begins with generation of MR-visible signal. Image quality depends critically upon the available signal-to-noise ratio (SNR) in an MR experiment. In the majority of cases, where the dominant noise source is the imaged object (rather than the coil), SNR in MRI is linearly proportional to the strength of the static magnetic field (B_0) [17, 21–25]:

$$SNR \propto B_0 \tag{1.1}$$

Today's clinical scanners typically have a B_0 field of 1.5 or 3 Tesla (T). The prevalence of 7T systems is also increasing, especially since the first 7T system gained Food and Drug Administration (FDA) clearance in 2017. Compared to 1.5T, 7T systems yield approximately 4.7 times the SNR, the benefits of which are discussed in Chapter 2. In order to understand the relation described in Equation 1.1, the origin of the MR signal must first be appreciated.

1.1 The B_0 field and polarisation

MR signal relies on two, related properties intrinsic to atomic nuclei. The first is magnetic susceptibility, often denoted with χ . Nonzero magnetic susceptibility indicates that a substance polarises in the presence of an externally-applied magnetic field. In other words,

the substance itself becomes (temporarily or permanently) magnetised. Magnetic susceptibility can be broken down into diamagnetism, paramagnetism, or ferromagnetism. While para- and ferromagnetic materials are attracted by externally-applied magnetic fields (meaning they polarise in the direction of the field), diamagnetic materials oppose externally-applied magnetic fields. The majority of biological tissue is weakly diamagnetic (although tissue can be paramagnetic in locations with high concentrations of metals such as iron). The second atomic property critical to MRI is quantum spin. In essence, an atom with nonzero spin can absorb and emit electromagnetic radiation. This energy exchange is maximally efficient at the atom's so-called resonance condition, where the external electromagnetic frequency matches the intrinsic frequency of the spin. Spin frequency in turn is dependent on the atom, as well as external conditions including magnetic field strength and temperature. For a more detailed overview on quantum spin, see Ref. [26].

Spins are dipole-like in their behaviour. In the presence of a static external magnetic field (known as equilibrium condition), a very small, but critical majority align with the direction of the static field (referred to as the longitudinal, or +Z direction), resulting in a so-called magnetic moment. The magnetic moment describes the summation over a population of spins within a small region of the sample being imaged. For most MRI protocols, hydrogen protons within the water molecules of biological tissue are the nuclei of interest and primary source of signal. Sodium, potassium, oxygen and chlorine are among several other nuclei which are suitable for MRI, though their natural abundance in biological tissue is much lower than hydrogen, meaning that their magnetic moment is smaller. As mentioned in the previous paragraph, the intrinsic spin frequency (and therefore its resonance condition) is atom-dependent, meaning that the resonance condition for water is slightly different from that of fat due to their different chemical compositions - known as chemical shift. These mechanisms allow water-fat separation, spectroscopic analysis, and imaging of metabolic processes using MR techniques.

The net magnetisation vector (M) at equilibrium is:

$$M_0 \propto \phi B_0 \tag{1.2}$$

Here, M_0 is the net magnetisation at equilibrium; ϕ is the proton density of the tissue or

material; and B_0 is the static magnetic field. However M_0 itself is not visible to the electronic receivers used in most MR systems. To generate signal which can be detected by such receivers, the longitudinally-aligned M_0 must be perturbed, or tipped, into the transverse plane. Once tipped, M rotates about its longitudinal axis, and this is known as precession. It is worth noting that individual atomic spins also precess, however they cannot produce a net rotation vector without a strong external magnetic field. The precessing M generates a sinusoidal current in nearby receiver coils, forming the time-domain MR signal. Signal magnitude can be described as:

$$S \propto M_0 \sin(\alpha) \quad (1.3)$$

where S is the complex MR signal, and α is the angle of perturbation from the longitudinal axis – known as the flip-angle. With this in mind, Equation 1.2 illustrates that more signal is available when B_0 is large. It therefore follows that a higher B_0 field translates into higher SNR.

1.2 Precession and the role of radiofrequency

In the previous section, it was explained that M_0 is perturbed in order to generate transverse magnetisation. This is done by applying a radiofrequency (RF) pulse. The magnetic field associated with the RF pulse is known as the B_1 field. RF is transmitted for a finite duration T in a direction perpendicular to B_0 . The net magnetisation rotates about the axis along which RF was applied, according to the left-hand rule. Assuming that the resonance condition is met (discussed next), the time-integrated RF magnitude determines the flip-angle experienced by spins. The flip-angle α can therefore be further described as:

$$\alpha = \gamma \int_0^T B_1(t) \delta t \quad (1.4)$$

where δt is the time between discrete RF samples. According to the principle of resonance, the RF frequency ω_{rf} must closely match that of the spins' frequency for energy exchange to effectively perturb M_0 .

The operating frequency is:

$$\omega = \frac{\gamma}{2\pi} B \quad (1.5)$$

where γ is the gyromagnetic ratio in radians per second per Tesla, and B is the total magnetic field. γ is the gyromagnetic ratio which describes the spin behaviour intrinsic to a given atomic particle, and $\frac{\gamma}{2\pi}$ has units of MHz per Tesla. Ignoring off-resonance effects, this equates to the Larmor frequency ω_0 at equilibrium:

$$\omega_0 = \frac{\gamma}{2\pi} B_0 \quad (1.6)$$

The Larmor frequency of a hydrogen nucleus is 42.58 MHz/Tesla. In principle, the static field is spatially and temporally uniform, meaning that ω_0 is also uniform. Therefore, signal is typically demodulated by the Larmor frequency to simplify calculations required for encoding and acquisition. This is conceptually identical to adopting a frame of reference which is rotating at the Larmor frequency. The so-called rotating frame therefore eliminates precession due to the static field, meaning that remaining frequency terms are due to intentional or unintentional local offsets or deviations in the field. In the rotating frame, the flip-angle experienced by spins is the integral of the demodulated RF pulse envelope.

1.3 Gradient fields and k-space

Three major fields contribute to form the overall magnetic field B in MRI. Firstly, the static B_0 field which was described in the previous section, and can be neglected under the rotating frame. Secondly, the time dependent B_1 field, which was previously introduced, and is described in more detail later in Section 1.4. The third contribution is from spatial gradient fields, which are also time dependent, and are explained in this section.

A gradient coil produces a spatially-varying field (on the order of tens of millitesla per meter), with the spatial variation being in the direction of the gradient. Although non-linear spatial variation can be useful for certain specialist applications [27], linear gradient fields are

desirable for most imaging protocols. While the gradient fields are switched on, ω is also spatially-dependent, with a frequency modulation pattern which is consistent with the gradient field (i.e., linearly-varying in most cases):

$$\omega_G(r) = \frac{\gamma}{2\pi} G \cdot r \quad (1.7)$$

Here, G is the gradient field and ω_G is its associated frequency term at voxel r . In general, three gradient coils are available: one along each of the scanner's physical x, y, and z axes. Gradients are used for spatial encoding in MRI, which can be done in three conceptually distinct ways. Firstly, gradients can be applied during RF pulse playout to modulate the pulse's effective bandwidth. Therefore, only spins within a certain frequency band (which, according to Equation 1.7, equates to a specific spatial region) are affected by the RF pulse. Increasing gradient amplitude narrows the effective bandwidth of any given pulse, since the range of frequency values over a given region becomes larger. Secondly, a gradient can be briefly pulsed (blipped) sometime after excitation. This so-called phase encoding gradient causes a phase dispersion along the gradient direction which persists after the gradient has been switched off. Although it is ultimately the phase that varies in this case, the amount of phase dispersion is determined by differences in frequency caused by the gradient, which accumulate over its (short) duration and are 'locked in' after the gradient pulse. Thirdly, gradients applied continuously during signal readout cause direct frequency modulation in the readout direction. For simple 2-D imaging, each method is generally used to spatially encode a different axis, while for simple 3-D imaging, non-selective RF pulses can be played out, and two different phase-encode directions can be used along with the frequency-encoded readout. However, more advanced spatial encoding methods are also commonly used, such as when imaging with oblique slice orientations, non-linear encoding and/or readout trajectories, and multi-dimensional spatial selection during RF pulse playout.

Prior to image reconstruction, raw MR signal is stored in so-called k-space (or frequency space). K-space can be understood as a finite collection of planar waves with differing frequencies and phases. The central part ($k_x = k_y = 0$) is equivalent to the DC offset, and determines contrast in the image. The periphery of k-space instead contains high-frequency information on image detail such as borders and edges. The radial angle within k-space

contains the phase information. This applies to signal reception and reconstruction, however excitation k-space exists in an analogous way. The way that RF energy is deposited in k-space can affect the uniformity of the measured signal; this is discussed in the context of RF pulse design throughout this thesis. In Cartesian acquisitions, k-space is often described as having a frequency-encoded dimension and one or more phase-encode dimension, which are encoded with the frequency- and phase-encode gradients as described above.

The Nyquist limit dictates that signal must be sampled at twice the highest expected frequency in order to prevent aliasing [28]. Aliasing arises from the mis-allocation of measured signal in k-space and usually manifests as a wrap-around of structures in image space (e.g., the tip of the nose appearing behind the head). The simplest linear image reconstruction technique is to take an inverse Fourier transform of the k-space signal, however, under-sampling in either direction (for example, by skipping certain phase-encode lines) results in aliasing. If signal is under-sampled (for example, to reduce acquisition time), or if non-Cartesian spatial encoding is used, more advanced techniques such as partial or non-uniform Fourier methods [29], compressed sensing [30], or parallel imaging (see Chapter 2) are better suited for image reconstruction.

1.4 The B_1^+ field

The B_1 field in MRI can be decomposed into the transmit field (B_1^+) and the receive field (B_1^-). For the purpose of this thesis, only the B_1^+ field is of interest, and B_1^- is not discussed except where clearly stated. The useful component of the B_1^+ field rotates in the same direction as nuclear precession, and is responsible for tipping the magnetisation when an RF pulse is applied. Most modern scanners typically use quadrature or birdcage RF transmit coils which produce power-efficient B_1^+ fields. These coils contain multiple physically-offset transmit elements which are driven with a fixed phase shift of the equivalent angular offset. This results in a circularly-polarised (CP) B_1 field (which is the B_1^+ field), meaning that the useful components from the coil elements interfere constructively, while other components (which instead only contribute to undesirable tissue heating) cancel. A B_1^+ field is therefore complex-valued, with a magnitude and a phase at each spatial location.

The significance of the B_1^+ field can be appreciated by studying the spatial-domain response of magnetisation to externally-applied B_1 fields [31, 32]:

$$M(r) = i\gamma M_0 S(r) \int_0^T b(t) e^{ir \cdot k(t)} \delta t \quad (1.8)$$

where k is the trajectory traversed in k-space, which is determined by integration of the time-reversed gradient waveforms. This holds for small-tip angle pulses, while large tip-angle pulses require alternative methods for calculating M , discussed later.

The B_1^+ field is temporally varying, with the time evolution determined by $b(t)$: the waveform played out over time by the RF coil. On the other hand, the B_1 sensitivity distribution S of a given transmit element determines the spatial coverage of the associated B_1^+ field. The two are separable, allowing temporal- and spatially-varying aspects of the B_1^+ field to be considered (and manipulated) independently, but the overall B_1^+ field depends on both, since the waveforms act to scale the sensitivities. Although B_1^+ field uniformity has been assumed so far in this chapter, the separability of S and b is especially important in the context of field non-uniformity (i.e., when S varies substantially throughout the imaged object), which will be discussed in Chapter 2.

For neuroimaging at $\leq 3T$, the B_1^+ field is typically produced by transmit coils built into the scanner bore (body coils) or a quadrature/birdcage head coil, and its spatial variation in brain-sized objects is assumed to be negligible (i.e., B_1^+ is relatively uniform across the head). Vendor-supplied RF pulse waveforms and sequences therefore perform relatively well for neuroimaging applications at $\leq 3T$. Spatial mapping of the B_1^+ field (B_1 mapping) is rarely required, except if specialist post-processing methods which incorporate the B_1^+ field are used. This is not the case for torso imaging at 3T or brain imaging at higher field strengths, where signal voids are present (discussed in Chapter 2) and the assumption of field uniformity is violated. In these cases, alternative acquisition approaches which require knowledge of the B_1 sensitivity maps are often needed (further detailed in Chapter 2). Good image quality can be achieved using so-called tailored RF pulses which aim to overcome the B_1^+ non-uniformities present by incorporating B_1 maps in their design [31–34]. Tailored pulses are discussed in the context of parallel transmission in Chapter 3.

1.5 Electric fields and safety

According to Faraday's Law, any time-varying magnetic field produces electric fields (and, according to Maxwell's equations, vice-versa), meaning that the applied B_1 and gradient fields are accompanied by unavoidable concomitant electric fields (E-fields). E-fields are highly relevant in MRI safety, since they can ultimately causes tissue heating through power deposition.

While temperature is not generally measured directly in MRI, it can be estimated, for example, using Pennes' bioheat equation, which describes the temperature (τ) change over time in perfused tissue:

$$\rho_{tis}C_{tis}\frac{\delta\tau}{\delta t} = k_{tis}\nabla^2T - C_bW_b(\tau - \tau_{art}) + P \quad (1.9)$$

where ρ_{tis} is the tissue density, C_{tis} and C_b are the tissue and blood heat capacities, respectively, k_{tis} is the thermal conductivity, W_b is the local perfusion rate, and τ_{art} is the blood temperature (assumed to be constant). However, perfusion cannot be relied upon for thermoregulation in certain cases, such as in subjects with a reduced perfusion response [35]. During an MRI scan, the temperature change experienced in tissue stems primarily from power absorption (P in Equation 1.9). Therefore, the amount of tissue heating can be estimated based on the amount of power absorbed by tissue. The specific absorption rate (SAR) quantifies this in units of Watts per kilogram:

$$SAR = \frac{1}{V} \int_V sar(r)\delta v \quad (1.10)$$

The SAR at each voxel ($sar(r)$) is:

$$sar(r) = \frac{\sigma(r)}{2\rho(r)}|E(r)^2| \quad (1.11)$$

where V is the sample volume over which the SAR is calculated, σ is the tissue conductivity in Siemens/m, E is the electric field in Volts/m, and ρ is the tissue mass density in kg/m³.

Because temperature is not measured directly, SAR is used as a proxy measurement to ensure that temperature limits are not exceeded in MRI (based on known relationships such as the one described in Equation 1.9). The validity of using SAR to estimate temperature is widely debated, as increases in SAR do not always coincide with temperature increases due to complex biological factors such as thermoregulation, perfusion and evaporation, each of which in turn depends on many time- and subject-specific variables [36, 37]. Nevertheless, SAR continues to be the primary safety mechanism for temperature management on MRI systems.

Regulatory guidelines for SAR are set by the Food and Drug Administration (FDA) and International Electric Commission (IEC) - for neuroimaging, SAR limits are designed to prevent temperature in the brain increasing above 38°C. For normal operating mode, whole-head SAR should not exceed 3.2 W/kg. Local SAR limits are also given for the average SAR within a small region; for example, there is an additional limit of 10 W/kg per 10g of tissue when local (surface) coils are used. The values for FDA and IEC pertain to SAR averaged over 10 and 6 minutes, respectively (additionally, the 10-second averaged SAR is also limited to twice these values according to the IEC). The limits applicable to a given MRI scan also depend on the anatomy being imaged; for example, when using local transmit coils, imaging in extremities has a higher local SAR limit (20 W/kg) than that for head and torso imaging (both 10 W/kg) according to the IEC. To calculate local SAR, V in Equation 1.10 can be set to a small amount of tissue (e.g., 10g).

Although it is rare, MRI can lead to burns. Out of 1,548 adverse MRI events reported across a 10-year period, almost two thirds were thermal related [38]. However, burns stemming directly from RF power absorption are rare ($\sim 6\%$ of all burns reported in [39]). It is worth noting that these longitudinal reports were from the early 2000's, when local transmit coils, coil arrays, and 7T systems were less common. It is explained later in Chapters 2 and 3 that SAR is especially relevant when using local transmit arrays, and at UHF due to increased RF power requirements [17]. Therefore, although the primary focus is on B_1^+ and signal, consideration to SAR and safety is also given throughout this thesis.

1.6 Pulse sequence basics: Relaxation and echoes

Immediately following application of an RF pulse, the magnetisation vector begins to relax to its equilibrium value M_0 , which manifests as decaying signal magnitude. (Relaxation happens during the RF pulse application too, but this is usually neglected due to the short pulse duration relative to relaxation). Relaxation occurs via two conceptually distinct mechanisms – namely, loss of the MR-visible transverse component of the net magnetisation vector (M_{xy}), and regrowth of the longitudinal component M_z . In the former case, so-called T2 (or spin-spin) relaxation, which refers to the exchange of energy between adjacent spins, causes M_{xy} to reduce in magnitude due to a loss of intra-voxel phase coherence. In practice, M_{xy} decays more rapidly than T2 relaxation alone due to additional phase incoherence caused by local field non-uniformities. As described in Equation 1.5, precession frequency is a function of the magnetic field strength. The Larmor frequency therefore describes spin behaviour at the nominal field strength (i.e., with zero off-resonance), however spins located in areas of marginally lower or higher magnetic field will lose or gain phase respectively, compared to those at Larmor frequency. Slight differences in frequency between different tissues also contributes to this extra dephasing. This mechanism of phase accrual causes relaxation known as T2', which combines with T2 effects to accelerate transverse decay. The overall transverse relaxation rate (incorporating both T2 and T2') is referred to as T2*.

On the other hand, T1 (or spin-lattice) relaxation relates to the regrowth of M_z towards its equilibrium value (M_0) which occurs as spins gradually re-align with B_0 . T1 relaxation is especially significant in the context of RF pulse sequences, since the magnitude of M_z immediately prior to an excitation pulse determines how much potential signal is available to be tipped into the transverse plane by the subsequent excitation pulse. Signal becomes saturated if insufficient time between RF pulses is left for M_z recovery, meaning that no new measurable signal can be generated.

The evolution of magnetisation in time is well understood and can be described by the rotating frame Bloch equations [40]:

$$\begin{aligned}
 \frac{dM_x}{dt} &= \gamma M_y \left(B_0 - \frac{\omega_0}{\gamma} \right) + \gamma M_z b(t) \sin(\omega_{rf} - \omega_0)t \\
 \frac{dM_y}{dt} &= -\gamma M_x \left(B_0 - \frac{\omega_0}{\gamma} \right) + \gamma M_z b(t) \cos(\omega_{rf} - \omega_0)t \\
 \frac{dM_z}{dt} &= -\gamma M_x b(t) \sin(\omega_{rf} - \omega_0)t - \gamma M_y b(t) \cos(\omega_{rf} - \omega_0)t
 \end{aligned} \tag{1.12}$$

In combination with Equation 1.8 (which describes the initial effect of B_1^+ and gradient waveforms on the magnetisation), the required RF and gradient waveforms (i.e., b and G) can be calculated via RF pulse design to elicit the desired behaviour on the magnetisation vector (discussed in Section 2.3.4). The following section provides a brief overview of RF pulses typically used in MRI, while pulse design is discussed further in Chapter 3.

RF pulses designed to tip magnetisation into the MR-visible transverse plane are referred to as excitation pulses. In the ideal scenario, spins are perturbed by the RF to the desired flip-angle, and there is phase coherence among them. In the case that phase coherence is not achieved (or is lost, e.g., due to T_2^* effects), signal magnitude is decreased due to destructive interference between spins with different phases within the same voxel.

At some point prior to complete signal decay, signal is sampled, forming an MR measurement. This point is known as the echo time (TE). TE can be chosen to provide the desired signal magnitude and contrast in images. Since relaxation mechanisms cause a reduction in signal magnitude, short TEs generally provide more signal than long TEs following a single excitation pulse, as less time is available for relaxation to occur. However, differences between tissues (i.e., contrast in images) are primarily a product of different relaxation rates between tissues, meaning contrast only begins to manifest some time after the RF pulse. There is generally an optimum TE for a given application and imaging setup, which will maximise the contrast between tissues of interest, while retaining sufficient signal.

Figure 1.1 shows example relaxation curves for brain tissue following an excitation pulse. Precise relaxation rates are defined as the amount of time taken for 63% of M_z recovery in the case of T1, or 37% of M_{xy} decay in the case of T2, compared to their values immediately following the RF pulse. T1 relaxation in cerebrospinal fluid (CSF) is long (~ 3000 ms)




Figure 1.1: Example relaxation rates in brain tissue

compared to that in grey matter (~ 1300 ms), which in turn is longer than that in white matter (~ 800 ms) [41, 42]. T2 relaxation is shorter than T1, at ~ 2000 ms, ~ 80 ms, and ~ 110 ms for CSF, grey, and white matter, respectively [41, 42]. T2* relaxation is shorter still, but the tissue hierarchy matches that of T2. These times pertain to 3T, however precise relaxation rates depend on a number of factors including B_0 field strength, with higher B_0 fields generally pertaining to slightly longer T1 relaxation in grey and white matter [43–45].

In addition to TE, the repetition time (TR) is another parameter to consider in the pulse sequence. TR refers to the repetition time within a sequence – that is, the time at which the sequence is repeated to acquire a new section of k-space. With the exception of single-shot encoding techniques, which acquire an entire k-space following a single excitation pulse, many repetitions are usually required, each with different encoding parameters (e.g. a unique phase-encode gradient to fill a line of k-space). The TR has important implications in safety assessments, since power deposition is averaged over time, so a longer TR equates to a lower SAR – discussed in Chapter 2.

While T2 relaxation due to spin-spin interaction is unavoidable, the component of T2* relaxation due to local off-resonance effects can be (partially) reversed. This is achieved by applying a so-called refocussing RF pulse at time $\frac{TE}{2}$ [46]. These pulses generally have larger flip-angles (traditionally 180°), and they recover signal lost due to additional phase terms associated with local off-resonances over the $\frac{TE}{2}$ interval. This can be understood intuitively by considering that a 180° pulse causes phase reversal of spins in the transverse plane. Therefore, spins which have lost phase (with respect to the Larmor frequency) will gain phase following the refocussing pulse, and vice-versa. As spins continue to gain or lose phase due to the same field offsets, following a second time interval of $\frac{TE}{2}$ after the refocussing

pulse, an echo forms in the signal as the phase accrual has been reversed. The principle assumes that spins are stationary and experience the same local field offsets before and after the refocussing pulse, although this is not strictly the case (e.g., due to diffusion or flow in certain regions).

Large flip-angle ($\sim 180^\circ$) pulses can also be applied to invert M_0 prior to an excitation pulse (i.e., inversion pulses). Then, when the excitation pulse is applied, tissue contrast has already begun to manifest due to T1 relaxation, which causes spins to relax towards their original M_z orientation. By applying the excitation pulse at the null point of a certain tissue (i.e. when M_z for the tissue = 0), signal from that tissue can be suppressed (e.g., in fat- or fluid-suppressed sequences).

Pulse types and applications discussed above represent a small portion of the tools and measurements available in MRI. Because different pulse types are more or less well-suited to certain applications, and perform well under different conditions, it is important that pulse design approaches are devised with flexibility in mind. As mentioned, the performance of standard pulse waveforms can break down in the case where B_1^+ is non-uniform and load-dependent, such as at UHF. Furthermore, the design of large flip-angle pulses requires alternative, more complex techniques compared to small tip-angle pulses (discussed in Chapter 3). Throughout this thesis, attention is paid to signal fidelity in the context of parallel transmission at UHF. These regimes are accompanied by their own characteristics and requirements, which will be explained and addressed in the following chapters.

Chapter 2

Ultra-High Field MRI

The primary technical benefit of UHF MRI is the gain in ultimate intrinsic signal-to-noise ratio (SNR), which scales linearly with B_0 field strength [17, 21–24]. Practically, compared to 4T, 7T yields up to double the SNR (depending on the location) [17]. For hip imaging, SNR was reported to increase by 133% at 7T compared to 3T in Ref. [47]. Since it is not possible to eliminate noise, signal must be averaged to form high-quality images. MR signal is always summed spatially, over a small region (voxel) within the sample. When the available SNR is low (for example, due to a low B_0 field), larger voxels are needed (i.e., image resolution is coarse, on the scale of several millimetres). In contrast, sub-millimetre voxel sizes are commonly acquired at 7T.

Similarly, acquisitions are often repeated in time and temporally averaged. Again, to increase SNR more averaging is required, translating into longer scans with more sequence repetitions. Specifically, $SNR \propto \sqrt{t}$, where t is the scan duration. When the native SNR is high (i.e., at UHF), scan time can be held constant or reduced for improved or equivalent image quality, respectively (compared to lower field strengths). The temporal benefits of leveraging high SNR towards shorter scans are simple; patient discomfort and scan costs are reduced, and patient throughput can be increased. This can be supplemented by parallel imaging techniques which allow further acceleration by effectively dividing the acquisition across multiple receiver channels using so-called parallel imaging techniques [48, 49].

In parallel imaging, multiple channels with differing sensitivity profiles are used for signal reception, and the final image is reconstructed by combining the signal received in each channel - either in k-space (e.g. GRAPPA [50] and SMASH [51]) or image space (e.g., SENSE [52]). Therefore, the amount of k-space data acquired by each channel can be reduced by a factor of N (while still adhering to the Nyquist limit in the final image), so scan times can be reduced by the same factor (N being the number of channels). Due to the previously-mentioned dependence of SNR on time, accelerated PI acquisitions have lower SNR than fully sampled acquisitions [52, 53]. Nevertheless, the increased intrinsic SNR available at 7T helps to counter this undesirable effect, meaning that image quality can be retained even with relatively high acceleration factors [53–55]. Parallel imaging, which uses multiple receiver channels, should not be confused with parallel transmission, where multiple channels are used for RF transmission (although similarities exist between the two concepts). Parallel transmission is discussed in Chapter 3.

In addition to SNR, susceptibility contrast and contrast-to-noise ratio (CNR) are also increased at UHF, meaning differences in the signal from different tissues (or between different chemical species) are larger. Increased contrast can be valuable for a wide range of applications, however it can also lead to signal dropout at tissue boundaries.

Some of the practical benefits of UHF are discussed in both clinical and research contexts in the following section, before outlining major challenges associated with the UHF regime. The chapter concludes by describing existing approaches which help to overcome those challenges.

2.1 Practical benefits of UHF

The increased available spatial resolution at 7T allows detection of subtle features and lesions which are often imperceptible at $\leq 3T$ [56]. In Ref. [10], experienced neuroradiologists were given paired UHF (7T) and standard (3T or 1.5T) images of a cohort of 104 patients with a range of pathologies (including epilepsy, traumatic brain injury, multiple sclerosis (MS), cortical infarcts, orbital melanoma, among others). In almost all cases, pathology-related abnormalities such as infarcts, lesions, micro-haemorrhage and

malformation were described as better demarcated and more conspicuous at 7T. In several patients, micro-haemorrhage, malformation, and MS lesions were only identifiable on the 7T images. In one case, post-traumatic encephalomalacia (as determined at 7T) was mis-identified as a cystic lesion at 3T. Even where pathology was also visible at lower field strengths, in many cases, the 7T images allowed a more detailed clinical assessment. For example, while lesions were visible at 3T, the central vein sign within lesions (which distinguishes multiple sclerosis from other conditions [9]) was only visible at 7T. Other examples included better localisation of micro-haemorrhages, and quantification of the proportion of the optic nerve which was affected by orbital choroidal melanoma.

Epilepsy is an area where clinical adoption of 7T is beginning to flourish. Among treatment-resistant epilepsy patients, surgical resection of the seizure onset zone is often necessary. The seizure onset zone (SOZ) must be identified prior to surgery, and it is unfortunately the case that abnormalities often cannot be identified using standard field strength MRI. In these cases, invasive techniques such as installation of stereoelectroencephalography grids are the logical next step. However, if the SOZ can instead be identified using MR images acquired at UHF, invasive treatment pre-surgery can be bypassed. Several studies have been conducted with 7T scans in focal epilepsy patients who had previously received inconclusive MRI results (i.e., no lesions identified) at 1.5 or 3T [1, 3, 4]. For up to 67% of patients, the UHF images revealed previously undetected abnormalities in cortical structure and/or susceptibility. According to a meta-analysis, the overall detection rate of abnormalities at 7T was 65%, compared to just 22% at standard clinical field strengths [2]. Among these cases, the 7T findings helped to identify the SOZ, assist pre-surgical planning and/or electrode placement, and substantially altered the clinical prognosis [1, 4].

Similarly, compared to 3T, 7T images have been found to provide significantly higher sensitivity and specificity in the detection and localisation of multiple sclerosis (MS) lesions in the brain [5, 7, 57–59] and spinal cord [60]. Recently, uptake of 7T in the context of MS was recently recommended for clinical practice [8]. Inter-rater agreement in the categorisation of MS lesions was 0.97 at 7T, compared to 0.12 at 3T, showing that the improved resolution can greatly help reduce clinical ambiguity [5]. Regarding prognosis, high resolution quantitative T1 maps were recently used to identify and differentiate between

demyelinated (active), remyelinated (inactive) and mixed de/re-myelinated lesions associated with MS [6].

Compared to 3T, the increased susceptibility contrast and spatial resolution of quantitative susceptibility maps (QSM) acquired at 7T led to more accurate manual parcellation of small subcortical structures in pre-surgical planning [61]. In line with theory, QSM at 7T was also demonstrated to produce image quality equivalent to that acquired at 3T in half the scan time [62].

In research settings, the blood oxygenation level dependent (BOLD) contrast signal relies on dynamic signal changes caused by deoxyhaemoglobin to infer regions of activation in the brain. The contrast is dependent on field strength [63]; signal changes at 1.5T are around 1-4% of the total signal, increasing to ~4-9% at 7T [64–67]. Smaller voxel sizes available at higher field strengths also allow more precise localisation of the BOLD response [68]. Using 7T functional MRI (fMRI) responses to a range of movements, digit-specific [69] and whole-body [70] somatotopic representations within the cerebellum have been identified. Similarly, detailed retinotopic mapping of the visual cortex is possible at 7T [71].

UHF MRI is especially well-suited to techniques such as chemical exchange saturation transfer (CEST) imaging and magnetic resonance spectroscopy (MRS). These techniques can be used to quantify the prevalence of clinically or functionally relevant molecules other than water with applications in stroke [72], tumour [11], osteoarthritis [73], lymphedema [74], and Alzheimer's disease [12], among others. The increased spectral resolution means that signal peaks associated with each metabolite are more clearly defined compared to lower field strengths. MRS is typically measured using very large voxels (tens of mm^3), leading to uncertainty in the spatial origin of derived spectra. The higher SNR also means that smaller voxels ($\leq 3 \text{ mm}^3$) can be used, allowing more precise localisation of metabolites.

2.2 Challenges of UHF

2.2.1 Signal

The benefits outlined above make clear the motivation to use UHF MRI. There are, however, important disadvantages of the UHF regime. As well as determining the magnitude of M_0 , the strength of B_0 also determines the Larmor frequency, as shown in Equation 1.6. Since wavelength is inversely related to frequency, higher B_0 also implies a shorter operating wavelength. Many of the remaining challenges which currently prevent widespread use of UHF MRI stem from the shorter wavelength which accompanies the higher operating frequency at 7T.

At 3T, the operating frequency and wavelength in tissue are 127 MHz and 26 cm, respectively, while at 7T, they are 298 MHz and 11 cm. Because the wavelength at 3T is longer than the dimensions of a human head, phase variation is minimal, meaning that at $\leq 3T$, the quasi-static regime can be assumed for neuroimaging. Under this assumption, Equation 1.4 need not consider spatial location, since b generates an RF field (B_1^+) which is relatively spatially uniform, meaning that flip-angle (and by extension, signal) is also fairly uniform ($\sim 30\%$ variation) and there are no signal voids across the head [75].

In contrast, at $\geq 7T$, the electromagnetic regime applies since the operating wavelength is similar to the size of the head such that substantial phase variations exist within the imaged volume [17, 76, 77]. The issue can be understood by considering the temporally- and spatially-varying nature of electromagnetic fields (i.e. wave propagation and interference patterns). The shorter wavelength means that spatial variation is greater within the fields produced by each RF source, meaning that a standing wave pattern is observed in head-sized objects at 7T - so-called wavelength effects. This is true even with a single RF source, but the problem can be exacerbated when multiple transmitters are used due to interference effects [78] (which will be discussed in the following paragraphs). Body imaging at 3T encounters similar issues as neuroimaging at 7T since the operating wavelength at 3T is on the order of the size of the human torso. Additionally, at 7T, B_1^+ amplitude is more strongly attenuated in conductive materials, leading to amplitude variation due to there being more zero-crossings (because of the shorter wavelength) across the imaged object

[77]. Local RF transmit coils are used to help overcome this, but B_1^+ amplitude remains variable in space (by an order of magnitude or more) since it decays according to the inverse square law $\frac{1}{d^2}$, where d is the Euclidean distance from the coil.

Therefore, flip-angle (Equation 1.4) and signal (equations 1.3 and 1.8) become spatially-dependent:

$$\alpha(r) = \gamma \int_0^T B_1(t, r) \delta t \quad (2.1)$$

$$S(r) \propto M_0 \sin(\alpha(r)) \quad (2.2)$$

Therefore, for small tip-angle (STA) pulses,

$$M(r) = i\gamma M_0 S(r) \int_0^T b(t) e^{i\gamma \Delta B_0(r)(t-T)} e^{ir \cdot k(t)} \delta t \quad (2.3)$$

with r denoting voxel location and ΔB_0 being the local off-resonance term.

In a perfectly CP field, phase coherence is achieved by applying a phase offset to each RF source. Constructive interference is controlled, with the superposition of different sources summing to a given phase value (e.g., 0°). However, a field can only be truly CP for one spatial location. As described in Chapter 1.6, a relative phase offset must be applied to each channel in order to offset the physical angle between channels (e.g. $+90^\circ$ between 2, orthogonally-oriented quadrature channels, each driven with a sinusoidal current). However, the apparent angle between channels varies depending on the spatial location within the imaged object. This means that by applying 45° offsets in an 8-channel pTx array, CP can be achieved at the centre-most region (the white voxel in fig 2.1). However, other locations perceive the angle between channels differently (i.e., not 45°), resulting in imperfect CP at those locations (the purple and yellow voxels in fig. 2.1). Undesirable destructive interference therefore occurs outside of the centre, manifesting as diminished signal (darkening) in the periphery of the imaged object (although it is often described as central brightening) [79].

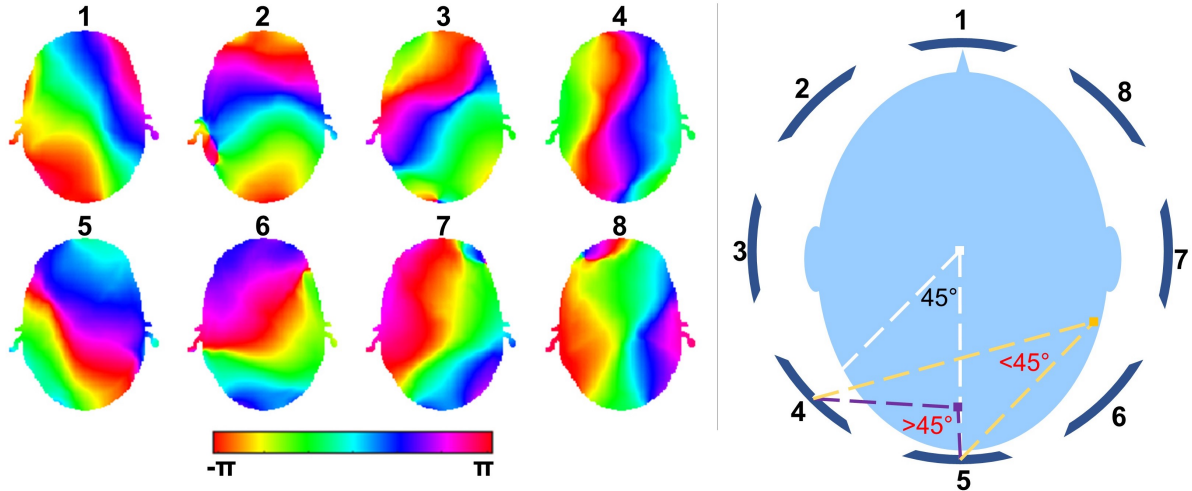


Figure 2.1: **Left:** Simulated B_1^+ phase distributions from an 8-channel pTx array at 7T. In practice, the overall B_1^+ field is a superposition of the fields produced by each channel, leading to areas of constructive interference and areas of destructive interference due to the spatially-varying phase present at 7T. **Right:** Schematic to illustrate the origin of destructive interference in circularly-polarised fields. While CP can be achieved in the centre-most region (white) by applying a 45° offset between transmit channels, the field at other locations (yellow and purple) is not perfectly CP because the perceived angle between transmit channels at these locations is different (i.e., not 45°).

Artifacts also arise from the increased susceptibility contrast at 7T. As described in Chapter 1.6, T_2^* effects arise from local field non-uniformity and energy exchange between neighbouring spins. The field non-uniformity at 7T, along with an increased chemical shift contribute to an increased dynamic range at 7T. Therefore, T_2^* decay is accelerated at tissue boundaries and air-tissue interfaces including the sinuses, causing intra-voxel dephasing and manifesting as signal dropout. Specifically,

$$Artifact \propto \Delta\chi \cdot R \cdot B_0 \cdot TE \quad (2.4)$$

where $\Delta\chi$ is the difference in magnetic susceptibility between tissues/regions and R is the largest voxel dimension. To reduce susceptibility-related artifacts, spin-echo techniques can be used (since these remove the contribution of field non-uniformity to T_2^* for stationary spins), TE can be minimised, or the acquisition bandwidth can be increased [80, 81].

2.2.2 Safety

As well as signal-related effects, there are important safety-related disadvantages of UHF MRI. Compared to 4T, approximately 2.1 times more RF power is required to achieve the same flip-angle at the centre of the head at 7T [17]. SAR scales with the square of the applied RF Voltage, meaning SAR is more than double that at 4T according to these findings.

Additionally, local RF transmit coils are often used as opposed to body transmit coils inbuilt to the scanner bore. For brain imaging with local head coils, this has the advantage that incidental conductive loops are more difficult to form, since the RF field is very small outside of the anatomy being imaged. However, local transmit coils at 7T adhere to the near-field regime, meaning that the dynamic range of the B_1^+ and concomitant E-fields are much higher, with high magnitude regions proximal to the transmit elements. Power deposition in a saline sample was shown to occur deeper inside the sample as a function of increasing operating frequency [77]. It is unclear whether the same phenomenon occurs in biological tissue; the RF skin effect is another phenomenon which suggests the opposite trend (that power deposition occurs closer to the periphery of conductive objects with increasing frequency). Furthermore, when multiple surface coils are used in an array, local hotspots can also form at arbitrary locations within the imaged volume due to constructive interference between E-fields produced by different RF sources (discussed further in Chapter 3). This causes concern, but can be controlled with careful SAR prediction and management.

2.3 Strategies to overcome the non-uniformity at UHF

Like the benefits, the technical difficulties associated with UHF MRI have widespread impact on its utility. However, progress has been made to overcome the major issues described above. In this section, techniques designed for this purpose will be outlined along with their own strengths and pitfalls. For the rest of this thesis, the B_1^+ field will be the primary field of interest, but some techniques which also help to overcome B_0 non-uniformity at UHF are discussed here too.

2.3.1 Inner-volume imaging

One solution to overcome the short wavelength issue is to specify a small region of interest (ROI) for imaging while suppressing signal from non-ROI regions in so-called inner-volume imaging. In an early demonstration, Feinberg et al. applied a 1-D selective excitation pulse before applying a 1-D selective refocussing pulse in an orthogonal direction [82]. Signal was therefore only generated from the echo produced in the intersecting region, where the two planes overlapped. However, spatially-selective refocussing pulses are very SAR-expensive, which is especially problematic at 7T due to the higher RF power requirements previously mentioned. A train of non-selective refocussing pulses were instead used by Mitsouras et al. following application of a 2-D selective pulse to excite a cylinder ROI [83]. Either approach can be used to reduce the signal-generating region such that it effectively becomes smaller than the operating wavelength. This in turn means that phase variation across the ROI is minimised, and interference effects can be neglected (as is the case for lower field strengths).

Inner-volume imaging is especially useful for minimising scan time when only small regions are of interest, or for suppressing signal from problematic regions such as those which are prone to breathing-related motion. However, there is the obvious drawback that the technique is inflexible; large objects such as the whole brain cannot be imaged. Furthermore, when parallel transmit arrays are used for inner-volume imaging and the ROI is off-centre with respect to the array, channels' RF contributions become highly asymmetric (i.e., a few channels contribute most of the RF power), increasing the sensitivity to motion (explained in Chapter 3).

2.3.2 Shimming

An alternative approach is to increase field uniformity via shimming. In the case of the B_0 field, additional coils can be used for the purpose of shimming, while dielectric (high-permittivity) materials can be used in or around the coil to shim the B_1^+ field. These materials are carefully placed around the subject to generate secondary fields which interact with the main field in a desirable way.

Shimming of the B_0 field is necessary for all systems due to manufacturing and environmental imperfections which degrade field uniformity. Therefore, B_0 shim coils are typically integrated into the scanner bore by default, and shim currents are automatically calculated to offset field non-uniformities (measured using B_0 field mapping sequences) during a scan. Recently, multi-channel local shim arrays have been used to further improve B_0 field uniformity, however this requires dedicated hardware [84].

The B_1^+ field can also be shimmed, although the approaches used are different from those used for the B_0 field. One passive technique is to place padding made from high-permittivity materials strategically around the subject to maximise B_1^+ field homogeneity within the imaged region. High-permittivity materials polarise when exposed to externally-applied electric fields. This induces an opposing secondary EM field within the material, which distorts the overall B_1^+ field distribution in the material's proximity by focussing the field around it [85]. Dielectric padding has been used in this way to increase signal in areas such as the frontal and temporal lobes and cerebellum, where signal is typically low due to the previously-mentioned wavelength effects at 7T [86, 87].

Although passive shimming can help homogenise fields in a small region, it often comes at the cost of homogeneity in other regions, leading to the same limitations as inner-volume imaging (i.e., difficulty in imaging larger ROIs). Similarly, although dielectric padding is placed to maximise the B-field homogeneity, it also changes the E-fields -and consequently SAR distributions- in potentially unpredictable and unanticipated ways [88]. Practically, due to the close-fitting nature of head coils used at 7T, limited space is available for dielectric padding. Furthermore, field homogeneity may initially be enhanced with dielectric padding, but it remains dependent on several factors including the load position. Changes in positioning (e.g. due to subject motion) change the optimal padding configuration, causing homogeneity to be lost once again (see Chapter 4).

Dynamic shimming of the B_1^+ field is another approach whereby non-uniformities are explicitly addressed in RF pulse design. This is a form of parallel transmission, which is discussed later in Chapter 3.

2.3.3 Retrospective B_1^+ correction

If the B_1^+ distribution is known, artificial contrast due to its non-uniformity in resulting images can be retrospectively corrected to some extent. In an early demonstration, [89] improved the quality of T1 and proton density measurements retrospectively. The bias introduced by firstly, a global flip-angle miscalibration (i.e., mismatch between the target and actual flip-angle), and secondly, a spatially-dependent flip-angle variation (due to spatially-varying B_1^+ sensitivity) were calculated and removed by acquiring images at several different flip-angles. This meant that effects of the spatially-varying sensitivity could be disentangled from the spatially-varying measures of interest (proton density and T1). The method worked well for B_1^+ variation of $\sim 50\%$ at 1.5T, however at 7T with local transmit coils, large areas of very low magnitude B_1^+ (where signal is indistinguishable from noise) are present due to the near-field regime.

More recently, bias introduced via the B_1 field to measures of myelin have been addressed using retrospective techniques at 3T, with or without a measured B_1^+ map [90, 91]. Although quantitative measures such as these require precision (and therefore even a small bias introduced by B_1^+ needs correcting), the B_1^+ field at 3T is much more uniform than 7T as discussed previously. Furthermore, these techniques assume that the effects of any patient motion are negligible to the B_1^+ distribution, which is reasonable at 3T, but not at 7T, where motion-related field changes are higher (see Chapter 4).

Retrospective techniques are also limited in that they cannot correct for so-called spin history effects (where issues in the previous excitation affect the current excitation). Essentially, while scaling factors may be accurately extracted and applied to correct the final image intensity, contrast which varied due to changes in spin behaviour cannot be corrected.

2.3.4 Pulse design (single transmit)

Another tool which can be leveraged to overcome B_1^+ non-uniformity is RF pulse design. RF pulses can be formulated in the frequency [33, 92] or spatial domain [31, 32, 34]; methods discussed within this thesis will comprise the latter.

For a given target magnetisation profile M , the required RF pulse b must be calculated by essentially inverting the well-understood behaviour of magnetisation described by the Bloch equations (Equations 1.12). This can be achieved using simple linear methods [32], linear approximations [93], iterative optimisation algorithms, or artificial intelligence. The method of choice depends largely on the type of pulse required and the imaging context.

Broadly, linear methods can only be used to design STA pulses, since the relationship between flip-angle and signal becomes non-linear at flip angles of $>30^\circ$ according to Equation 2.2. Although the linear approximation strictly holds to pulses with flip-angles of $\leq 30^\circ$, in practice, pulses up to 90° can be designed using linear methods, with good results.

For small tip-angles, the pulse design problem can be linearised and expressed in matrix form:

$$M = \mathbf{A}b \quad (2.5)$$

The so-called system matrix \mathbf{A} represents the behaviour of the magnetisation in a known system (i.e., with a given B_0 distribution and gradient waveforms). It is an $N_r \times N_t$ matrix incorporating the gyromagnetic ratio γ , the equilibrium magnetisation M_0 , B_0 maps, gradient terms, and channel sensitivity distributions.

The required RF pulse is generally solved by framing Equation 2.5 as a minimisation problem. In its simplest form, a solution can be sought for the following least-squares cost function:

$$b = \arg \min_b \{|M - \mathbf{A}b|^2\} \quad (2.6)$$

This least-squares problem can be solved by taking the psuedoinverse of \mathbf{A} , or using optimisation algorithms such as conjugate gradient [94]. In cases where the phase of M is not critical, the inner terms (both M and $\mathbf{A}b$) can be replaced by their magnitude, transforming the problem into a magnitude-least squares problem. This can lead to improvements in the uniformity of the magnitude profile (at the cost of a degraded phase profile) [95].

To design large tip-angle (LTA; $>90^\circ$) pulses such as those required for inversion or refocussing, consideration of the evolution of magnetisation over time is needed. This is

most accurately achieved by carrying out full Bloch simulations with a candidate RF pulse to generate the elicited magnetisation profile, and updating the candidate RF pulse accordingly. Having said that, linear approximations can be made in some cases, and are generally preferable for efficiency, while slower iterative methods are preferable for flip-angle accuracy but are less clinically feasible.

Rather than overcoming or reducing B_1^+ non-uniformity, an alternative approach is to use techniques which are robust to it, such as adiabatic pulses [96] and B_1^+ -independent rotation (BIR) pulses [97]. In contrast to non-adiabatic pulses which are usually applied at (or very near) the Larmor frequency, these pulses sweep over a wide range of frequencies during their payout. Specifically, the amplitude of adiabatic excitation pulses is typically modulated from 0 to its maximum, while the frequency offset is modulated from its maximum to 0 (where 0 is the Larmor frequency). So long as the sweep is slow with respect to the precession frequency, the net magnetisation follows the effective field (the total magnetic field) in the rotating frame (see Section for a description of the rotating frame). This results in effective perturbation of spins even in areas of high off-resonance (i.e. when ΔB_0 in Equation 2.3 varies substantially across the sample). Nevertheless, adiabatic pulses are very SAR-expensive, and are therefore undesirable at UHF.

More recently, pulse design techniques which compensate for non-uniformities in the B_1^+ distribution have been developed [98–101]. The quality of T2-weighted images was improved by optimising the refocussing pulses of an echo train in a turbo spin echo (TSE, also known as fast spin echo) sequence in Eggenschwiler et al. [101]. The optimisation was achieved using a spatially-resolved extended phase graph approach, which allows direct control to be placed on the signal generated at each echo [102]. TSE is a pulse sequence which is capable of filling k-space very rapidly via a long train of spin echoes (as was done in the inner-volume example mentioned previously [83]). The echoes are typically generated by many nonselective (hard) refocussing pulses which follow an initial excitation pulse.

In Eggenschwiler et al., the standard pulses were replaced by so-called kT-points pulses [99], which distribute energy throughout k-space in a desirable way, reducing the effects of B_1^+ non-uniformity on the acquired signal. This is achieved by splitting the pulse into multiple (N_{kT} ; usually 2-9) sub-pulses (of a few hundred microseconds in duration each), where each sub-pulse is applied at a different k-space location, navigated with gradient blips (see Figure




Figure 2.2: Example single-channel kT -points diagram for a 7 kT -points pulse. Arbitrary units are shown (gradients and RF waveforms were normalised to their respective maximum values). Since kT -points use short, non-selective excitation sub-pulses, their duration can be very short (on the scale of microseconds). Gradient blips are used to rapidly navigate k -space, meaning that power from each RF sub-pulse is deposited in a different k -space location.

2.2). The kT -point locations are determined using subject-specific B_1^+ sensitivity maps. kT -point locations are chosen as the N_{kT} highest energy components in the excitation k -space. In the case of single-channel transmit with small tip-angle (STA) excitation, this could be the Fourier transform of the multiplicative inverse of the B_1^+ map. In the study, optimised kT -points pulses improved image quality substantially, reducing the standard deviation of the signal profile by 34% compared to the standard hard pulse TSE sequence (although SAR was increased) [101]. Compared to adiabatic pulses, kT -points pulses are short and require much less RF power, meaning they are better suited to UHF settings where SAR is a limiting factor. Eggenschwiler et al. demonstrated the feasibility of using kT -points with a readily-available single-channel transmit (sTx) system, which meant that performance

was still limited by many of the issues described above [101]. Dielectric pads were also used to increase signal intensity in the cerebellum, but low-magnitude regions of B_1^+ remained and were omitted from the optimisation as signal fidelity in these regions cannot be optimised in sTx without compromising either signal quality elsewhere, or SAR. kT-points have more commonly been employed with parallel transmit systems (discussed in the next Chapter), which avoid these shortcomings and yield further improvements in signal homogeneity.

2.4 Conclusions

It is clear that, despite many partial solutions, technical problems in the UHF regime remain, hindering its clinical uptake. Although strides have been taken to overcome the B_1^+ non-uniformity using techniques such as inner-volume imaging, shimming, and B_1^+ -insensitive RF pulses, each has its own major drawback. In particular, in the context of imaging patients who are prone to motion, these solutions are inadequate. The optimal shim setup is dependent on the patient's position, while inner-volume imaging (especially when the ROI is off-centre) and lengthy B_1^+ -insensitive pulses are also particularly susceptible to motion. Furthermore, the RF power and SAR requirements of these pulses makes them ill-suited for UHF imaging. Pulses which include measured B_1^+ sensitivities in their design go some way to overcoming the non-uniformities present, but are not capable of generating signal in areas where B_1^+ magnitude is low.

In the following chapter, parallel transmission (pTx) of RF will be described. pTx is a flexible hardware and software approach in which the transmitted RF power is distributed between multiple independent channels. The increased degrees of freedom allow pTx pulses to effectively overcome B_1^+ non-uniformities. Like the other approaches covered in this chapter, pTx has its own associated difficulties. One such area is the response of pTx pulses to patient motion, which is currently not well understood. This is subsequently discussed in Chapter 4, before being addressed in the study chapters.

Chapter 3

Parallel Transmission (pTx)

Following firstly, the broad adoption of parallel signal reception with parallel imaging, and secondly, the move towards high and ultra-high fields (UHF) with their associated B_1^+ uniformity issues, parallel RF transmission (pTx) was introduced. For pTx, transmit coil arrays are used to transmit RF from independently-controlled, spatially-distributed channel elements. Like single-channel (or non-pTx) transmission (sTx) at 7T, local transmit coils are usually used (although the early implementations of pTx were conducted at $\leq 4.7T$). Between 2 and 32 loop or dipole (e.g. stripline) elements are generally used, and they can be spatially distributed in many ways. Spatial distribution patterns influence field distributions directly as well as the degrees of freedom available in pTx pulse design, with the latter depending on the amount of overlap between channels' sensitivities (the so-called geometry factor - see [52]) as well as the number of channels. Optimal pTx coil design depends on the intended application, although maximising B_1^+ uniformity is generally desirable. At Cardiff University an 8-channel Tx 32-channel Rx head coil with loop elements is used (Nova Medical, Wilmington, USA). Therefore, except where specified otherwise, the work presented in the study chapters of this thesis used a simulated model of an 8-channel Tx loop array, with elements distributed at equal 45° azimuth intervals in the axial plane (although details of the Nova coil are not known due to proprietary information, and therefore, the simulated coil bears no resemblance to the Nova coil).

Underlying pTx coil designs is the ability to independently modulate the amplitude and phase of the RF waveform transmitted by each channel. It is a versatile technique which can

be applied to arbitrary imaging scenarios. Although the concept of pTx hardware was already two decades old [103], the technique began to take off in the mid 2000's following early implementations of pTx for coil design [104, 105], MR spectroscopy (MRS) [106], signal optimisation [107, 108], SAR management [34], and -arguably most influentially- pulse acceleration [18, 34] and overcoming B_1^+ non-uniformity at UHF [105]. The latter application forms the focus of this chapter.

Importantly, the number of degrees of freedom available in pulse design increases by a factor of N_c (the number of transmit channels in the array). When pTx arrays are used along with bespoke RF pulse design techniques, substantially improved signal uniformity can be obtained, even across large imaging regions at 7T. This chapter introduces the core concepts of pTx including why it is especially well-suited to overcoming the B_1^+ non-uniformity encountered at UHF. Then, pTx pulse design techniques are discussed, before outlining remaining challenges which must be addressed to increase the feasibility of routine clinical pTx use.

3.1 Basic principles of pTx

As mentioned in Chapter 1.6, the B_1^+ (and therefore flip-angle) distribution for a given imaging setup depends on two, separable aspects related to each individual RF source: its spatial sensitivity distribution, and the RF waveform (i.e., scaling over time). The former decays in space as $\frac{1}{d^2}$, where d is the Euclidean distance from the coil, but it additionally depends on a number of factors related to the imaging setup including field strength, the coil and its load (e.g., patient anatomy and position). It is inherently non-uniform at 7T (see Chapter 2). The RF waveform can be considered as a time-dependent weighting which scales the source's associated sensitivity to produce the overall B_1^+ field. Assuming that channels are decoupled, the RF transmitted from channel c scales only the sensitivity of channel c (although in practice perfect decoupling is difficult to achieve).

The overall B_1^+ field in pTx is produced by the superposition (summation) of all independent RF sources within the array. Therefore,

$$B_1^+(r, t) = \sum_{c=1}^{N_c} B_{1,c}^+(r, t) = \sum_{c=1}^{N_c} S_c(r) b_c(t) \quad (3.1)$$

where the latter expression makes explicit the separability of each channel's spatial sensitivity profile (S_c) and time-dependent RF (b_c). The flip-angle in each voxel location therefore also depends on these same factors. The critical role of RF pulse design becomes obvious in this form; a carefully-designed RF pulse in pTx is capable of reducing the overall B_1^+ non-uniformity by scaling each channel's sensitivity independently, and in a dynamic way. This is the basis of tailored pulse design, for which knowledge of the sensitivity distributions is required (for example via B_1^+ mapping). In other words, tailored pulse design in pTx can be summarised as solving for $b_c(t)$, given known (measured) S_c and a target flip-angle profile.

The kT-points pulses [99] described in Chapter 2 are an example of tailored pulse design. Residual error in that case (where pTx was not used) arose partly from low-magnitude regions of B_1^+ , where the applied RF was unable to produce the nominal flip-angle without sacrificing fidelity elsewhere or exceeding SAR limits [101]. Due to the spatial dispersion and independent control of RF sources, pTx allows more nuanced manipulation of the overall field, and kT-points pulses are capable of yielding highly uniform flip-angle profiles over large 3-D volumes.

While generic vendor-supplied pulses generally perform well for lower field strengths and sTx, pTx requires custom pulses to be designed on a case-by-case basis for optimal performance (with some exceptions, which are discussed later in this chapter). Like in sTx, the RF solution in pTx can be derived using a variety of approaches as described in Chapter 2.

3.2 Small-tip pTx pulse design

Assuming small tip-angles, the pTx pulse design problem described in Equation 2.5 can be modified for pTx by considering each transmit channel c separately, as shown in Equation 3.2 below.

$$M = \sum_{c=1}^{N_c} \mathbf{A}b_c \quad (3.2)$$

In Ref. [32], it was shown that well-defined STA excitation patterns could be achieved using the following cost function, which simply adds a regularisation term and spatial weighting to Equation 2.6:

$$b = \arg \min_b \{ \|M - \mathbf{A}b\|_w^2 + \lambda b^T b \} \quad (3.3)$$

Here, λ is an adjustable Tikhonov regularisation parameter which penalises RF power and T denotes complex transpose. Additionally, w is a spatial weighting factor used to specify ROIs, such that error is minimised within the ROI [32]. The optimal choice of λ was shown to be dependent on the target flip-angle, with small values (<1) yielding better excitation profiles for flip-angles of $\leq 30^\circ$, and large values (e.g., 10 and 100) being preferable for larger flip-angles. This was attributed to the fact that pulses with high RF magnitude are more likely to violate the linear STA assumption when channels are combined. Therefore, when larger tip-angles (which inherently require more power than small tip-angles) are designed, solutions with low RF power are favourable over those with high power.

Peak RF power can be reduced using so-called VERSE (variable rate selective excitation) pulses [109]. The technique dilates RF (and gradient) waveforms in time such that the peak amplitude can be reduced whilst still integrating to the same flip-angle. In Ref. [110], Malik and Hajnal used VERSE pulses for STA pulse design in the context of TSE sequences. The method is used for pulse design in Chapter 5 of this thesis.

To maximise spin echo signal amplitude, refocussing pulses in a TSE train should be applied with a 90° phase offset with respect to the excitation pulse, and elicit a flip-angle of exactly 180° . This is the so-called Carr-Purcell-Meiboom-Gill (CPMG) condition [111]. However the latter is hard to achieve in practice, meaning that echo amplitude is reduced compared to its theoretical maximum. In terms of transverse magnetisation, M_x and M_y can be equivalently expressed as real and imaginary components of M_{xy} respectively (i.e., $M_{xy} = M_x + iM_y$).

In addition to the spatial weighting term w shown in 3.3, Malik and Hajnal's approach also introduces a different weighting factor to real and imaginary signal components [110]:

$$\begin{pmatrix} \text{Im}(m) \circ w_i \\ \text{Re}(m) \circ w_r \end{pmatrix} = \begin{pmatrix} \text{Im}(\mathbf{A}) \circ w_i & \text{Re}(\mathbf{A}) \circ w_i \\ \text{Re}(\mathbf{A}) \circ w_r & -\text{Im}(\mathbf{A}) \circ w_r \end{pmatrix} \begin{pmatrix} \text{Re}(b) \\ \text{Im}(b) \end{pmatrix} \quad (3.4)$$

Assuming an excitation pulse is applied along the rotating frame x axis, magnetisation rotates about that axis to align with the y axis (i.e., $\text{Re}(M_{xy}) = 0$). With the CPMG condition in mind, only the imaginary (M_y) component will be efficiently refocussed. In Ref. [110], the non-CPMG (real / M_x) component is assigned a low weighting relative to the CPMG (imaginary / M_y) component when designing the excitation pulse, thereby penalising error in the CPMG component more than the non-CPMG component:

$$b_s = \arg \min_{b_s} \{ \|M_{sw} - \mathbf{A}_{sw} b_s\|^2 + \lambda \|b_s\|^2 \} \quad (3.5)$$

where $_s$ indicates that real and imaginary components have been separated. By prioritising error reduction in the CPMG component in this way, it is shown that erroneous excitation can be “pushed” into the non-CPMG part and is therefore not refocussed when an imperfect (non-180°) refocussing pulse is applied. This is especially useful in the context of TSE sequences, as error is reduced at every echo of the TSE train [110].

3.3 Large-tip pTx pulse design

For LTA pulse design ($>90^\circ$), simulation of the Bloch equations and/or iterative techniques must be used to characterise the system in place of the Fourier matrix. In a natural progression from the STA method described above, the additive angle method uses a linear approach to build LTA pulses iteratively [93]. Specifically, an initial STA pulse is designed in the above way (Equation 3.3) and a scaling factor is applied to elicit a large target flip-angle. The pulse is evaluated using Bloch simulations. The discrepancy between the resulting and desired profiles is subsequently treated as the target profile for the next iteration (with the profile's phase updated such that solutions from all iterations have a

coherent phase). The calculated “correction” pulses are summed with the waveform from the previous iteration. This is repeated, with the residual error reducing at each iteration until a tolerable error level is reached. A related concept is the linear class of LTA (LCLTA) pulses described in Ref. [92]. The LCLTA approach can be used to design pulses with arbitrary starting conditions, whereas the additive angle method assumes a starting condition of equilibrium. Compared to the initial STA solution, these iterative linear methods improve profile fidelity at the cost of longer pulse design times (50-180 seconds, depending on the excitation trajectory). Nevertheless, residual error remains due to the linear approximations.

Optimal control (OC) approaches have been proposed as a more accurate solution for LTA pulse design [112–115]. In OC, a least-squares cost function similar to that in Equation 3.3 is first solved, usually with the addition of some constraints such as hardware limits and RF power. Bloch simulations are performed with the resulting pulse. Unlike the additive angle method, the derivative of the Bloch equations with respect to small changes in the RF pulse is then calculated, and gradient descent methods are used to determine the pulse updates needed to reduce error. The process is repeated iteratively, leading to long design times (on the scale of minutes to hours [113, 116]) due to the computationally expensive Bloch simulations performed at each iteration. On the other hand, pulses designed in this way yield very low error in their magnetisation profiles due to the analytical approach used.

A “fast optimal control” (fast OC) approach subsequently proposed by Grissom et al. blends these two approaches [116]. Like with full OC, the magnetisation response to small changes in an RF pulse are characterised analytically using Bloch equations, but the relationship is subsequently defined using a linear model which can be used in place of full Bloch simulations for future iterations. To first derive the model, Bloch simulations are formulated in the spin domain, which allows the rotation produced by an RF pulse to be characterised by 2 parameters instead of the usual 3, simplifying the computation (see [117]):

$$\frac{d}{dt} \begin{pmatrix} \beta \\ \alpha^* \end{pmatrix} = \frac{i\gamma}{2} \begin{pmatrix} G(t) \cdot x & b(t)^* \\ b(t) & -G(t) \cdot x \end{pmatrix} \begin{pmatrix} \beta \\ \alpha^* \end{pmatrix} \quad (3.6)$$

where $*$ is the complex conjugate, and α and β are the two Cayley-Klein parameters which describe rotations of angle ϕ about vector $V = (v_x, v_y, v_z)$ via Equations 3.7 below.

$$\begin{aligned}\alpha &= \cos \phi/2 - iv_z \sin \phi/2 \\ \beta &= -i(v_x - iv_y) \sin \phi/2\end{aligned}\tag{3.7}$$

Use of the linear model in place of full Bloch simulations in fast OC was shown to speed up the design process by a factor of ~ 25.6 (depending on the excitation trajectory) compared to full OC in Ref. [116]; design times were reduced from >1 hour to <4 minutes to converge to the same level of error in the resulting profiles. In Chapter 5, fast OC is used for LTA pulse design (along with Malik and Hajnal's STA method described above).

3.4 Parallel transmit implementations

Pulse design approaches can also be broken down into quadrature mode, RF-shimming, and dynamic pTx, which refer to the drive mode of the pTx coil. The three form a spectrum; quadrature mode does not make use of the independent control of each coil, RF-shimming makes partial use, and dynamic pTx mode fully exploits the extra degrees of freedom available with pTx hardware. The quality of the resulting profiles unsurprisingly also form a spectrum, with dynamic pTx pulses yielding the most favourable flip-angle profiles. The pulse design approaches described in this chapter so far were formulated for use with dynamic pTx, however most can also be adapted to design quadrature or RF-shimming pulses.

3.4.1 Quadrature mode

Quadrature, or circularly-polarised (CP) mode is the simplest drive mode and is akin to sTx transmission. Channels' amplitude and phase are fixed, with amplitudes being equal and the phase offset between channels matching the physical offset between channels in the array. In other words, the phase offset between adjacent channels is equal to $2\pi/N_c$ for a cylindrically symmetric array with N_c channels. Ideally, phase coherence is achieved within a small region of the imaged object (as discussed in Chapter 2). Unlike other pTx modes, quadrature therefore does not require detailed knowledge of each channel's B_1^+ sensitivity distribution,

simplifying and accelerating the imaging workflow compared to other pTx approaches. In essence, quadrature pulse design simply applies a global field scaling based on the measured overall B_1^+ sensitivity in either a small ROI, or averaged across the whole region. Local sensitivity variations inherent to the coil and its load therefore cannot be overcome, meaning that destructive interference in peripheral areas leads to the characteristic centre-bright appearance of quadrature images, as previously explained in Chapter 2.

A few studies have compared performance of an 8-channel pTx coil driven in CP mode to a single-channel coil at 7T (both Nova Medical coils). In Ding et al., B_1^+ fields produced by the pTx coil generally had a higher magnitude and were more uniform than those produced by the sTx coil, translating to better image quality [118]. Similarly, Graf et al. showed that while single-voxel MRS spectra derived from the brainstem were statistically indistinguishable, the pTx coil yielded 27% higher SNR than the sTx coil, and higher flip-angles were achieved in the brainstem [119]. These findings demonstrate that pTx hardware offers some benefits even when independent control of channels is not exploited (i.e., using CP mode). Although differences were found, a pTx coil driven in CP mode is considered to behave similarly to standard sTx coils. It is therefore viewed as a safe and well-understood mode of pTx operation – though this is not necessarily true, as will be discussed in Chapters 6 and 7.

3.4.2 RF-shimming and static pTx

Unlike quadrature mode, RF-shimming (or so-called static pTx) allows either the phase or the phase and magnitude assigned to each channel to vary. Therefore, although the same waveform is still transmitted by each channel, its magnitude and phase can be scaled by a complex weight in order to manipulate the overall B_1^+ field in a desirable way.

RF-shimming can be used to maximise power efficiency or B_1^+ uniformity. For the former, a small target ROI is specified wherein channel weights are designed to constructively interfere, meaning that less RF power is needed to achieve the same superposed B_1^+ magnitude in that region. Lower RF power is associated with reduced SAR, and by extension, reduced tissue heating. In Ref. [120], Deniz et al. used maximum-efficiency shimming for hip imaging with a 4-channel pTx array at 7T. A B_1^+ metric (based on subject-specific B_1^+ sensitivity maps)

was defined along with a power metric (based on a subject-specific power correlation matrix calibrated with forward and reflected power measurements). The ratio of the metrics (B_1^+ to power) was used to quantify efficiency. Compared to CP mode, the maximum-efficiency shim was on average 34% more efficient, yielding up to 55% reduction in RF power deposition and more uniform flip-angles within the ROI in simulations and experiments. A similar approach was used for prostate [121] and 2-D brain [122] imaging at 7T.

Improved uniformity is not guaranteed by maximum efficiency shimming. Instead, a shim set can be solved for the purpose of maximising the homogeneity in an ROI. In Ref. [122], Curtis et al. calculated slice-specific RF shim sets to maximise uniformity within one slice. When averaged over the brain, the maximum-uniformity shim yielded a higher normalised mean B_1^+ magnitude (0.8 vs. 0.7) and a smaller standard deviation (13% vs. 18%) compared to CP mode. In Ref. [123], Gras et al. showed that RF-shimming for brain imaging with an 8-channel pTx coil yielded a normalised root-mean-squared-error (nRMSE) in the resulting profile of 20%, compared to 28% with CP mode.

Compared to CP mode, RF-shimming can improve B_1^+ efficiency and/or uniformity when small ROIs such as a single slice or small inner volume are used as the target shim region as was the case in the examples described above. However, improvements are negligible if large shim volumes (such as the whole brain) are specified [122]. This is due to the limited degrees of freedom available for RF-shimming over a large volume which contains high dynamic range in its B_1^+ sensitivity, even when many pTx channels are used. Mao et al. determined that RF-shimming can theoretically yield uniform fields across the whole brain at 7T when 16 Tx channels are available [124], however 16 channels were used in Ref. [122] and non-uniformities were still observed when shims were calculated for the whole brain. This is likely due to the practical methods available for pulse design, which cannot readily attain the same level of performance as numerical or theoretical calculations. Therefore, it is desirable to use dynamic pTx mode when large regions of uniformity are desired.

3.4.3 Dynamic pTx mode

RF-shimming assigns a single weight per channel, whereas dynamic pTx allows the independent channel weights to change throughout the course of an RF pulse. In other

words, each channel transmits a unique RF waveform, introducing many degrees of freedom into pulse design. Further degrees of freedom are available when gradient waveforms are also played out during or between RF pulses to optimally distribute the transmitted power in k-space. These can be leveraged to better control B_1^+ (and/or E-field) distributions over small or large regions. The necessary weights are calculated based on measurements of the B_1^+ complex sensitivity distributions via B_1^+ mapping. In Ref. [123], Gras et al. used tailored kT-point pulses on an 8-channel pTx system, yielding flip-angle profiles with an nRMSE of 7%, compared to 28% using CP mode for neuroimaging.

Channel weights can be updated in a discrete or pseudo-continuous way. When implemented on pTx systems, the kT-points approach previously described in Chapter 2 falls into the former category. Several (typically 2-9) k-space locations are visited via short gradient blips, and a nonselective pTx pulse is played out. 2-D (slice-selective) imaging can be performed using a similar approach known as spokes (or fast-kz) pulses. kT-points vary in all three logical axes, while the k-space locations for spokes pulses vary only in the 2 in-plane dimensions. In both cases, the weights differ between channels, and across kT-points (or spokes), but there are only $N_{kT} \times N_c$ (or $N_{spokes} \times N_c$) parameters to solve for. At the pseudo-continuous end of the spectrum are methods which optimise channel weightings for every timepoint of a several milliseconds-long RF waveform. These pulses are most often designed using approaches based on optimal control [113]. Here, the number of parameters to solve for is much higher - $N_t \times N_c$. High-quality flip-angle profiles with sharp transition bands and non-ROI signal suppression can be obtained for small and large flip-angles [113]. These approaches are also flexible; for example, acceleration factors and arbitrary excitation trajectories and target profiles can be specified in the design. As previously mentioned, the performance and flexibility of these approaches generally comes at a cost of longer pulse design times.

3.4.4 Plug-and-play pTx

A major criticism of pTx -particularly dynamic pTx- is the complexity involved with pTx setup compared to sTx protocols, where pre-defined pulses are used and merely calibrated according to a few standard pre-scan measurements. Therefore, efforts have been made to

increase the ‘plug-and-play’ usability of pTx, while retaining its benefits of for overcoming B_1^+ uniformity. An example is so-called Universal Pulses (UPs) which, in contrast to tailored pulses, are designed offline, using B_1^+ distributions from a database of representative individuals [123]. The approach is based on the assumption that similar B_1^+ distributions are present across human subjects when the same anatomical region is imaged. Therefore, by minimising error using several subjects’ concatenated B_1^+ sensitivity maps, the resulting pulse yields fairly good performance for any subject, including subjects not included in the database. In Ref. [123], UPs yielded profiles with 11% nRMSE; sitting between those produced by tailored pulses (7%) and RF-shimming (20%).

In [125], it was shown that profiles produced by UPs can be improved by using fast online customisation (FOCUS) of a UP. The customisation is based on measured subject-specific B_1^+ maps. Similarly, some generalisability of UPs is traded for improved performance in SmartPulse [126]. In this approach, a neural network was initially trained to categorise individuals, and one Universal-like pulse was designed per category (which performed well for individuals in that category). Following training, categorisation of individuals is used to determine which pulse is most well-suited to a given individual, based on readily-available information from their localiser scan. Like FOCUS, this approach yielded performance sitting between fully tailored pulses and UPs.

Although they represent a compromise between performance and convenience, the effect of subject motion on the performance of these pulse types is unclear. Chapter 4 discusses the effects of motion on pTx in detail.

3.5 SAR and safety in pTx

As discussed in Chapter 2, RF power requirements (and therefore SAR) increase with increasing field strength. An exception is when many (on the order of 128) pTx channels are used for RF transmission. In this case, global SAR reaches a peak at around 5T for head-sized objects, before decreasing at higher field strengths [127]. However, most pTx use cases today have ≤ 16 pTx channels. For the work conducted in this thesis, 8 pTx channels are used, meaning that SAR is generally higher at 7T compared to 3T.

Global SAR can be calculated relatively easily, and can be monitored online using measurements of forward and reflected RF power - differences between the two are attributed to power absorption [19, 128]. In sTx, global SAR is proportional to maximum peak local SAR, meaning both metrics are relatively accessible. However, this is not the case in pTx due to the complex interference patterns. Each channel in a pTx array produces its own, spatially distinct electric field (E-field). Like B_1^+ , the overall E-field in a pTx array is a superposition across channels, meaning that local hotspots can arise unpredictably in locations where channels' E-fields constructively interfere. Therefore, Equation 1.11 in Chapter 1.6 cannot be used, and a more comprehensive approach is required.

The Q-matrix formalism [129] offers a convenient way to account for channel interactions on a voxel-wise basis:

$$Q_{ij}(r) = \frac{1}{2\rho(r)} [J^*_{x,j}(r)E_{x,i}(r) + J^*_{y,j}(r)E_{y,i}(r) + J^*_{z,j}(r)E_{z,i}(r)] \quad (3.8)$$

Here, E is the complex electric field in Volts/m, J is the complex current density in A/m², ρ is the tissue mass density in kg/m³, i and j are channel indices, and $*$ denotes complex conjugate. The i, j th element of each $N_c \times N_c$ Q-matrix therefore contains interactions between channels i and j in one spatial location r . Q-matrices are specific to the coil and subject setup. They are calculated for unit current in each channel, meaning that, once calculated, voxel-wise SAR can be obtained for any arbitrary RF pulse b using the Q-matrices via:

$$SAR(r) = b^T Q(r) b \quad (3.9)$$

By summing over voxels, global SAR can be determined. To calculate local SAR, Q-matrices can be calculated over a small (e.g., 10g) region of tissue rather than voxel-wise. Numerous studies demonstrate that local SAR limits are more likely to be exceeded than global SAR limits at 7T [36, 37, 130], and it is local SAR which causes the most concern regarding hotspots and RF burns. Although the Q-matrix approach means that comprehensive SAR calculations can be conducted, the number of calculations which need to be performed is very high ($N_r \times N_c^2$). For applications where rapid SAR estimates are required, compression

techniques such as virtual observation points (VOPs) can be used instead [131].

Unlike B_1^+ , E-fields cannot be measured using MR techniques, meaning that local SAR cannot be calculated or monitored directly. Instead, E- and J-field information required for the Q-matrix calculations is usually obtained offline using electromagnetic simulations which solve the Maxwell equations for a given coil-subject setup. This can take hours to days per simulation, since the models can be formed of tens of millions of voxels. It also relies on accurate modelling since small modelling discrepancies can alter behaviour substantially. Vendor-supplied SAR monitoring is therefore based on simulated data from a limited set of generic body models. To allow for uncertainty including anatomical variation, positioning and modelling / system imperfections, large safety margins are applied. These limit scanner performance by imposing overly conservative RF power limits resulting in sub-optimal flip-angle and signal profiles, and/or extending scan times due to the necessary increase in TR. The effects of variations in anatomy and pTx coil design on SAR in pTx are discussed later in Chapters 6 and 7 respectively. Recently-proposed methods are emerging for rapid personalised local SAR estimation in pTx based on convolutional neural networks [132, 133], discussed more in Chapter 9.

Although pTx does add complexity, it also adds flexibility. Therefore, although SAR can be higher in pTx, is it not inevitable. Holding B_1^+ uniformity constant, when a 2-port body coil was used in pTx mode, local- and global SAR were reduced to 28% and 17% of its corresponding quadrature mode values, respectively [134]. Indeed, once calculated, SAR can be used directly as a constraint in pulse design (e.g. [115, 123, 135–137]). The flexibility can even be leveraged for applications such as RF hyperthermia and ablation, where small focal areas of very high SAR are intentionally created [138–140]. Conversely, E-field steering techniques can be used to specify locations of low SAR (for example, near a metallic implant) [141, 142].

3.6 Electromagnetic simulations

As mentioned above, local SAR in pTx cannot be observed directly, and therefore electromagnetic (EM) simulations are used to estimate it instead. For the work presented

later in this thesis, Sim4Life (Zurich MedTech, Zurich) software was used to implement EM simulations to generate SAR and B_1^+ distributions. Sim4Life employs the finite difference time domain (FDTD) modelling approach to calculate EM fields produced by realistic MRI hardware. The basis of the approach is to solve Maxwell's time-dependent equations describing EM field dynamics using partial differential equations, discretised in space using a grid of so-called Yee cells, which contain the electric and magnetic field components at that cell's spatial location.

For the simulations conducted in this thesis, the pTx coil models consist of 8-channel loop arrays where the loops were always distributed evenly in a cylindrical formation in the axial plane. Since we are interested in MRI-related EM fields inside of the human head, virtual body models from the Virtual Population (IT'IS Foundation, Zurich, Switzerland) were used in conjunction with the coil models [143]. These body models are whole-body, high-resolution voxelised models comprising approximately 300 different tissue types. Each tissue type is assigned with dielectric properties informed by the literature [144]. Although our interest is in brain imaging, for our simulations, the head, neck and shoulders were also included for improved accuracy as recommended by Wolf et al. [145].

Chapter 4

Motion in MRI

Motion has been problematic to MRI since its inception, and is a widely-discussed topic (see review articles [146] and [147]). In this chapter, the prevalence and cost of motion in clinical settings will first be considered, along with some common motion prevention strategies.

Then, the physical mechanisms underlying motion artefacts will be discussed, before exploring various motion-detection and correction techniques. While periodic motion, deformation (non rigid-body motion) and flow do contribute to the problems experienced in MRI, the primary focus here is on abrupt rigid-body head motion (i.e., rotation and translation), and its effects on the B_1^+ field, especially when pTx coils are used. To this end, the limitations of existing techniques in this context will be highlighted.

4.1 Motion in MRI

Motion occurring during an MRI scan often manifests as blurring, aliasing, signal loss, and/or signal brightening within images. The precise manifestation depends on the system, imaging sequence, signal acquisition strategy, and the type (and timing) of motion. Motion can affect both the internal consistency of measurements taken within the scan, as well as the compliance of any acquired data with assumptions of reconstruction or models fit to the data during post-processing. While the former usually results in visible artefacts, the latter can undermine the validity of conclusions drawn from the data with or without overt

artefacts (e.g., Ref. [148]). In both cases, quality and diagnostic value of data are compromised.

Other, more subtle issues arise when considering who is likely to move. Certain populations have a higher tendency to move (e.g. paediatric or elderly patients, patients suffering from tremors, or stroke patients), which may confound group-level differences in experiments [149–152]. However, even healthy volunteers -who are often assumed to be capable of remaining still- are prone to unintentional motion including physiological motion, task-related motion, and slow drifts [149, 153–156].

A recent review estimated that around 8% of scans were affected by head motion artefacts, with 2% being sufficiently motion-corrupted that a re-scan was necessary [157]. The annual cost to the hospital (Rigshospitalet, Denmark) was estimated to be \$364,242 [157]. The study considered 228 adult and paediatric clinical scans conducted at 3T over 2 years. Similarly, Andre et al. reviewed 192 scans acquired at 1.5 and 3T in a US hospital, estimating a total annual cost of \$434,000 due to motion-related issues [158]. They report a similar figure (8%) for the proportion of images affected by motion for outpatients, but this figure is much higher (29%) for inpatients [158]. Furthermore, a report of 2,384 scans acquired at 1.5T found that 10% contained sufficient motion to mask (or alter the apparent volume of) multiple sclerosis-related lesions [159]. The study highlighted that motion may remain undetected (artefacts were not always visible), yet severely impact the diagnostic value of images.

Reduction of motion-related artefacts can be achieved in several ways. Sedation can be used in cases where motion is difficult to avoid, such as when imaging infants, however this is costly and requires strong justification to outweigh the associated risks inherent to sedation [160, 161]. Moreover, even under general anaesthetic, residual head motion exceeding 1 mm is commonly observed in children [152]. Simpler motion prevention measures are routinely used, such as verbal instructions, and padding specially designed and placed to reduce patient mobility within the scanner.

Modifications can also be made to the scanning protocol to increase robustness to motion. Scans are accelerated (e.g., using parallel imaging techniques) which can reduce the risk of discomfort-induced motion. Modifications can also be made to the acquisition itself;

single-shot sequences acquire an entire 2-D k-space in a very short duration following a single RF excitation, meaning motion is unlikely to affect data consistency within that slice. Whilst gating and triggering (the use of physiological measurements such as cardiac or respiration cycle to determine sequence timings) can be used to minimise the effects of periodic motion, these techniques are inappropriate for abrupt, unpredictable motion, and therefore they will not be discussed.

In practice, these simple prevention measures come at their own cost (for example, low SNR in single-shot acquisitions, and patient discomfort from too much padding), and are insufficient to completely eliminate motion and its downstream effects. They are therefore often combined with motion correction strategies, which will be discussed later in this chapter.

4.2 Effects of motion on the signal

Accurate spatial encoding of the MR signal depends on stability of the object's position due to the reliance on controlled spatial frequency variations (i.e., gradient fields) across the imaging FOV. As previously discussed in Chapter 1.6, the local precession frequency is determined by the total field (i.e., summed B_0 , gradient and B_1 fields) at that location. In the absence of system imperfections (i.e., assuming a homogeneous imaging volume, uniform B_0 and B_1 fields and perfectly linear gradient fields), the only field change experienced by a moving object would be a change in the gradient field term. Incoherent phase accumulation across the imaged object (caused by spatial variation in the magnitude and/or direction of motion) causes signal to be attenuated (incidentally, this is the basis of diffusion-weighted MRI), or mis-assigned in k-space.

When translational motion occurs, a linear phase ramp arises in k-space, while for rotational motion, k-space samples are also rotated. Both cases cause aliasing and/or blurring in images. It is possible to avoid these effects from impacting the acquired signal by updating the FOV in real time to follow the imaged object. This can be achieved by adjusting the central operating frequency and gradient waveforms to prospectively counteract the gradient-related frequency shift experienced by the displaced (translated) object. When

rotational motion occurs, the gradient fields (and therefore k-space) must also be rotated [162–164]. Alternatively, for Cartesian acquisitions, the erroneous phase shift caused by motion can simply be estimated and corrected for retrospectively from the associated shifts and rotations of k-space [165].

For Cartesian acquisitions, signal is typically sampled pseudo-continuously in one (frequency-encode) direction, while the orthogonal phase-encode direction(s) are sampled more distally in time. Therefore, motion has minimal effect on consistency between frequency-encoded samples, and motion artefacts tend to arise in the phase-encode direction. Non-Cartesian acquisitions such as spiral or radial readouts tend to exhibit blurring when motion occurs, which can be harder to correct since frequency- and phase-encode directions are not aligned with the physical x- and y- axes, meaning that the artefacts are encoded in a non-Cartesian way. Finally, the nature of artefacts also depends on the timing of motion with respect to the signal readout - different artefacts can manifest if motion occurs while signal from the centre of k-space is being sampled, compared to at the periphery.

Slow drifts (e.g., sinking into a support pillow), generally do not produce visible artefacts when adjacent k-space locations are acquired sequentially. However, interleaved sequences are prone to artefacts from slow movement, since adjacent k-space locations are sampled distally in time, allowing more time for the drift to occur between neighbouring samples. Furthermore, the effects of slow drifts on E-field distributions (and therefore tissue heating patterns) should not be overlooked. The effects of motion on E-fields and SAR are discussed in more detail in Chapters 6 and 7.

As mentioned these effects can be corrected, however, in practice, system imperfections such as field non-uniformities exist. Linear correction techniques based on central frequency and gradient adjustments are therefore insufficient to eliminate the effects of motion.

Additionally, biological samples (including the human brain) are highly non-uniform in terms of electrical tissue properties, meaning that boundary conditions exist between tissues (and between tissue and air).

In terms of the B_0 field, the air-tissue boundary at the periphery of the head causes strong local field distortions known as the static susceptibility artefact. The artefact usually manifests as geometric distortions near the sinuses and ear canals and undermines data

quality [166]. It can be overcome relatively simply, for example by non-linear registration to an unaffected structural image or by measuring and incorporating a B_0 map in the reconstruction [167]. Alternatively, multiple images can be acquired with opposite phase-encoding directions, meaning the distortions are in opposite directions and can be combined, resulting in a much-reduced distortion in the final image [168, 169].

These techniques are effective at correcting for the static susceptibility effect, however, according to Maxwell's equations, a change in the boundary conditions within a fixed FOV (i.e., due to motion) changes the magnetic field distribution. Specifically, if the angle of a certain boundary changes with respect to the main magnetic field, field changes will occur. This so-called dynamic susceptibility effect means that even if the initial non-uniformities and static portion of the artefact are known and accounted for, they will change following motion. 5° rotations are reported to cause B_0 variations of 30 - 50Hz at 3T, producing ~ 6 mm distortions in EPI images [151]. Dynamic susceptibility effects are difficult to predict and are often overlooked by motion correction techniques.

Although susceptibility effects are often viewed in the context of B_0 , a conceptually similar effect arises within the B_1 field (although it should be noted that the mechanism behind these effects differs, and susceptibility does not affect the B_1 field). While B_0 changes produce geometric distortion, B_1 changes cause over- or under-excitation (i.e., incorrect flip-angle), manifesting as artificial contrast in images due to dynamic changes to the wavelength effects (as explained in Section 2.2.1). One example showed motion-related (receive) B_1 field changes of around 4% due to rotations of 5° at 3T [170].

4.3 Motion at UHF with pTx

These effects are problematic for any MR experiment, however certain system-related factors exacerbate the issues described above. Firstly, when UHF MRI is leveraged for high spatial resolution, sensitivity to motion is inherently increased due to the smaller voxel size. Secondly, the increase in field non-uniformity due to the shorter wavelength at UHF means that a given displacement is likely to cause a greater local field change (and therefore signal attenuation) compared to the equivalent displacement at lower field strengths. Thirdly, the

use of local transmit coils at 7T means that B_1^+ field magnitude decays rapidly in space near the coils, meaning the fields experienced following motion may be much higher or lower in magnitude (depending on the direction of movement). Finally, and of particular interest to this thesis, is the effect of motion on the B_1^+ field produced with pTx coils. Sensitivity of fields to head motion has been shown to be very high when pTx is used [130], however very little attention has been paid to the issue in general. Since each channel in a pTx array produces its own B_1^+ field, motion affects each in a different way. While the individual field changes are problematic in their own right as previously described, the interactions between pTx channels make the effects of motion on the overall (superposed) B_1^+ field highly complicated and unpredictable. There are also important safety implications associated with subject motion at UHF [130], which are discussed in sections 6 and 7.

An early investigation used finite difference time domain (FDTD) simulations to investigate the effect of head motion on B_1^+ and flip-angle in an 8-channel pTx system at 7T [171]. B_1^+ distributions were highly position-dependent, and flip-angle profiles were degraded following large rotations and/or small rotations with highly accelerated (under-sampled) excitation trajectories [171]. The effects of respiratory motion on body imaging at 7T using pTx was investigated in [172]. B_1^+ variation of up to 25% in simulations, and 18% in-vivo were observed due to respiration. Motion sensitivity was channel-dependent, since, for supine imaging, breathing alters the distance between tissue and anterior (but not posterior) channels. These studies focussed on the B_1^+ changes themselves. Following pTx pulse design, the channel weights initially calibrated for uniformity (see Chapter 3) are sub-optimal following head motion. This results in poor fidelity in the flip-angle profile due to under- or over-excitation in regions where B_1^+ changed, and consequently, artificial contrast in images. In Ref. [173], rotational motion of 20° caused normalised root-mean-squared-error (nRMSE) of up to 99% in purely simulated excitation profiles, while moderate head motion caused up to 22% nRMSE in simulated excitation profiles using in-vivo B_1^+ maps.

4.4 Correction techniques and their limitations

Motion correction techniques fall into two categories: retrospective and prospective. Retrospective correction involves manipulating the image reconstruction pipeline to

incorporate motion parameters. Although motion-related artefacts can be effectively removed, retrospective correction is unable to recover signal loss due to motion-related dephasing or interactions of motion with spin history effects. Another important consideration is that retrospective techniques cannot address any motion-related changes in tissue heating patterns due to changes in E-fields, which lead to higher SAR changes in pTx compared to quadrature mode [130]. On the other hand, prospective correction is carried out online at signal acquisition stage by applying real-time updates to the RF and gradient waveforms in response to displacements. In both cases, the motion parameters need to be characterised.

An alternative retrospective approach which does not always require knowledge of motion parameters is to use deep neural networks (NNs). NNs can be trained to remove motion artefacts using corresponding pairs of motion-corrupted and motion-free images [174–176]. MoCoNet uses convolutional NNs for this purpose, and reduced the maximum-observed motion-related error in 3-D T1-weighted images from 5.5% to 2.8% (compared to motion-free ground truth images) [174]. Similarly, motion correction of T2*-weighted images was demonstrated using a conditional generative adversarial network, where motion-related error of 16.4% was reduced to 10.8% [177]. However there is uncertainty surrounding the use of NNs (and other deep- and machine-learning techniques) when pathologies or abnormalities are present (especially if these were not included in the training database). Artefact removal also suffers the same shortcomings as other retrospective techniques as described above.

For techniques which do require knowledge of motion parameters, they can be acquired by incorporating so-called navigators into the sequence. Navigators are low-resolution images or projections, usually collected once per TR. Motion parameters extracted from registering the navigator data can be incorporated in image reconstruction. Navigators require sequence adjustments which often extend the total scan time. Sparse measurements can be taken which shorten the required duration, for example, FatNav uses fat-selective navigator pulses to localise the sparse fat surrounding the brain [178]. Alternatively, gradient fields with non-linear spatial variations can encode location information along multiple directions simultaneously, reducing the number of navigator waveforms needed [179]. Free induction decay navigators (FIDnavs) sample the signal's DC component rapidly (~ 1 ms). Changes in

the signal intensity are used to solve for changes in positioning, given the (known) spatial dependence of receive channels' sensitivity profiles [180, 181]. Wallace et al. developed a forward model framework to characterise changes in FIDnav signal with changes in head position [181]. Following offline calibration of the coil-specific system matrix (which included channels' sensitivities), the 6 rigid-body motion parameters could be accurately solved for using the FIDnav signal. Optical tracking systems can be used with no cost in scan time, and are able to provide motion parameters with high precision and temporal resolution [182–184], although they require specialist hardware. Self-navigation is an alternative group of approaches in which the imaging data itself is used to characterise and correct for motion. PROPELLER is a commonly-used example which sequentially acquires overlapping, rotated rectangular blocks of k-space data [185]. The central k-space portion (which is sampled every time) can be used for motion estimation in a similar way to FID navigators, although field uniformity is assumed.

A few studies have addressed the B_0 dynamic susceptibility artefact. Brackenier et al. demonstrated retrospective correction of motion-related B_0 field changes based on a model relating motion parameters with B_0 field changes [186]. The model is based on linear B_0 coefficient maps for each motion type, which are estimated using a data-driven approach. The recently-proposed FieldMapNet [187] is an alternative, deep-learning based method to retrospectively remove the effect of B_0 non-uniformities dynamically (i.e., on a shot-by-shot or line-by-line basis). Prospective dynamic susceptibility correction was proposed by Ward et al., who applied real-time updates to the gradient waveforms to counteract motion-related B_0 changes at 1.5T and 3T [188]. Because field measurements were 1-D and correction was applied via the gradients, this method was limited to first-order (linear) field correction, which is insufficient at UHF where higher-order terms are present.

In terms of the B_1 field, the 4% signal changes reported in Ref. [170] were overcome by extrapolation and adjustment of each channel's B_1^- map (e.g. by applying a rotation matrix) to match the position of the head. The realigned B_1^- maps were then used in reconstruction. The values within the sensitivities were not altered, meaning dynamic field changes were not considered. This resulted in effective correction at 3T, however a portion of error remained. The increased susceptibility variation at 7T, and interactions between channels in pTx mean that dynamic changes are larger, rendering this approach insufficient.

To avoid this, the dynamic B_1^+ field changes caused by motion should be compensated for prospectively. To the best of our knowledge, there is currently no available technique which fulfils these requirements with pTx coils. Real-time re-design of pTx pulses has been suggested as one potential avenue [189], which could be implemented in conjunction with prospective FOV correction via gradient waveform updates. Re-design of tailored pulses would counteract any motion-related B_1^+ changes, resulting in uniform flip-angle profiles, regardless of position. However, this requires knowledge of the motion-resolved B_1^+ maps in order to determine the new channel weights. A typical pTx B_1 -mapping sequence takes around 30 seconds per channel, so this cannot be measured online. Chapter 9 of this thesis presents an approach which estimates motion-resolved B_1^+ maps for this purpose. An alternative approach to improve the motion-robustness of pTx pulses is presented in Chapter 8. The remainder of this thesis is dedicated to better understanding the effects of motion in pTx, both in terms of signal and safety. Although there is still much work to be done until pTx is routinely used in the clinic, it is hoped that the work presented in the following chapters goes some way towards more widespread use of pTx by highlighting and addressing one of its remaining shortfalls.

Part II

Study Chapters

Chapter 5

Motion and large vs. small tip-angle pTx pulse performance

5.1 Introduction

UHF MRI offers benefits over standard field strengths, but B_1^+ non-uniformities caused by the shorter operating wavelength at 7T remain problematic, as described in Chapter 2. pTx is one solution which can be used flexibly for the transmission of tailored RF pulses, but the design of those pulses relies on knowledge of the B_1^+ distribution present at pulse ployout (see Chapter 3). Instead, B_1^+ maps are usually measured at the start of a scan. As described in Chapter 4, subject motion throughout the scan causes B_1^+ fields to change, resulting in sub-optimal pulse performance due to the mismatch between actual and measured B_1^+ distributions. Flip-angle profiles are therefore degraded, and artificial contrast arises in images, undermining the benefits of UHF MRI.

While limited attention has been paid to the effects of motion on small flip-angle (STA) pTx pulse performance [171, 173], to the best of our knowledge the effects on large flip-angle (LTA) pTx pulse performance has not been investigated. STA findings indicate that pulse design parameters (e.g. the trajectory used to excite k-space, acceleration factor) can affect sensitivity of the pulse to motion [171], raising the possibility that different pulse design algorithms may also produce pulses with differing motion sensitivity.

The primary question addressed here is whether motion-induced error in excitation profiles produced by STA and LTA pTx pulses are quantitatively and/or qualitatively different. This is an important consideration since future approaches for motion correction may also need to be different, or be flexible enough to accommodate different pulse types. To investigate this, we test the performance of STA and LTA pTx pulses with similar k-space trajectories and target profiles, but designed using two different approaches. A secondary interest is the effect of varying parameters within each approach on performance and motion sensitivity. Flip-angle profiles produced by each pulse with and without simulated motion are assessed. If a pulse is inherently more motion-robust, the profile produced under conditions of motion is expected to be more similar to that produced without motion, while differences can be attributed to sensitivity of the pulse to motion.

5.2 Methods

5.2.1 B_1^+ maps and simulations

B_1^+ maps were generated using FDTD simulations (see Section 3.6) in Sim4Life (Zurich MedTech, Zurich) following the protocol described in Ref. [130]. The Ella body model (IT'IS Foundation, Zurich, Switzerland) [143] was simulated using a generic 8-channel pTx coil array tuned to 295 MHz, shown in Figure 5.1. The loops were made from perfect electrical conducting (PEC) material. One RF source, and three capacitors were included per loop. The loop dimensions were 40 mm (width) \times 110 mm (height), and the array radius was 115 mm. The B_1^+ distribution with the head at an initial, centred position was used for subsequent pulse design.

To simulate motion, the array was displaced relative to the head (thereby avoiding any changes to voxelisation within the body model). B_1^+ maps were simulated for the body model in 103 positions, including all 6 rigid-body motion types; left-right (ranging from 0 to 20 mm), anterior-posterior (0 to 10 mm), inferior-superior (0 to 20 mm), pitch (-5° to 20°), roll (0° to 20°) and yaw (-20° to 20°). Off-axis translational movement (e.g. combined rightward and posterior) was simulated for 60 in-plane positions. All positions listed above

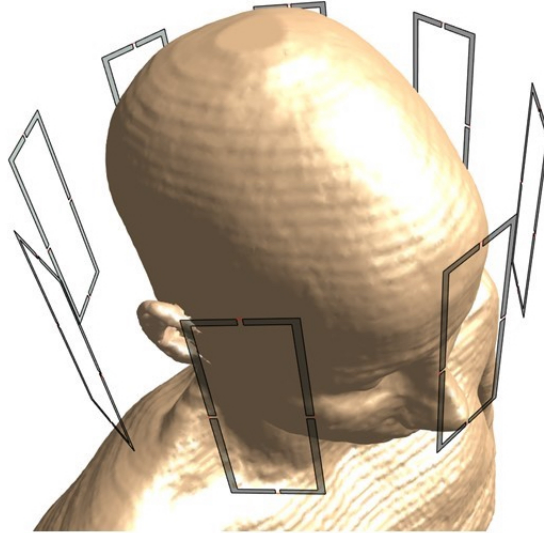


Figure 5.1: *Simulation setup in Sim4Life showing the Ella virtual body model in a central position within the 8-channel pTx array. Simulations included the head, neck and shoulders.*

refer to the effective movement of the head (for example, to simulate the “right 20 mm” position, the coil was translated 20 mm to the left, mimicking a rightward translation of the head).

Simulation results were normalised to an accepted power of 1 W per channel beyond the input port to the coil elements. Although this may arguably remove some of the motion-related variability, the effect of this normalisation step has been previously investigated [130]. Without normalisation, the variability in power delivered to each coil element was found to be small (around 8%) compared to the motion-induced variability in B_1^+ and flip-angle (up to 99% for the same range of motion in Ref. [173]). Therefore, normalisation was carried out in order to override imperfections in coil matching and any positional dependencies. The simulations were manually checked for input impedance and reflection coefficient as well as field smoothness across positions. Fields were masked to exclude background (air) voxels and exported to MATLAB (The MathWorks, Inc.) for further processing.

5.2.2 Pulse design

Two pulse design methods (one STA and one LTA) previously outlined in Chapter 3 were implemented. Malik and Hajnal's phase-relaxed Carr-Purcell-Meiboom-Gill (CPMG) excitation design method [110] was used as the STA design method, and Grissom et al.'s fast optimal control approach [116] was used as the LTA method.

The toolboxes for both pulse design methods are available online and can be freely downloaded (both accessed 20/11/2022):

STA - phase-relaxed CPMG excitation [110]:

https://github.com/mriphysics/phase_relaxed_CPMG_excitation

LTA – fast optimal control (fast OC) [116]:

<https://github.com/ismrmd/ismrmd-paper/tree/master/code/extern/irt/contrib/grissom-mri-rf-large-tip>

To assess differences in pTx pulse sensitivity to motion, three separate comparisons were drawn, in each of which the motion sensitivity of two pulses was compared. The pulses and comparisons are summarised in Figure 5.2. In comparison 1, STA and LTA pulses designed using their respective STA (phase-relaxed CPMG) and LTA (fast OC) methods were compared. In comparison 2, fast OC was used to design both STA and LTA pulses. In comparison 3, STA (90°) pulses designed using both phase-relaxed CPMG and fast OC methods were compared. The second and third comparisons were designed to inform whether any differences observed in comparison 1 were primarily driven by the pulse design method, or the target flip-angle. In addition to these three main comparisons, comparisons 4 and 5 explore the effect of parameter variations within each design method. In comparison 4, motion sensitivity of an echo planar (EP) trajectory was compared with a spiral trajectory for LTA pulses designed using fast OC. Finally, in comparison 5, the effects of CPMG weighting (with values 0, 0.1, 0.2, 0.5 and 1.0) and RF power regularisation (16 values ranging from 0.1 to 1024) on the motion sensitivity of STA pulses designed using the phase-relaxed STA approach were investigated. The roles of CPMG weighting and RF power regularisation in pulse design were previously explained in Section 3.2.

PULSE	DESIGN METHOD	TARGET FLIP-ANGLE	ROI	EXCITATION TRAJECTORY
i	Phase-relaxed CPMG	90°	Rectangle (10 x 5 cm)	Spiral
ii	Fast OC	180°	Rectangle (10 x 5 cm)	Spiral
iii	Fast OC	90°	Rectangle (10 x 5 cm)	Spiral
iv	Fast OC	70°	Whole slice	Single-spoke
v	Fast OC	180°	Whole slice	Single-spoke
vi	Fast OC	180°	Rectangle (10 x 5 cm)	EPI

COMPARISONS	PULSE A	PULSE B
1. STA vs LTA	(i)	(ii)
2. Effect of target flip-angle	(iii)	(ii)
	(iv)	(v)
3. Effect of design method	(i)	(iii)
4. Effect of trajectory	(vi)	(ii)
5. Effect of CPMG & RF power regularisation weight	80 pulses , 90° target, phase-relaxed CPMG method $\mathbf{W}_r = 0/0.1/0.2/0.5/1$; $\lambda = 16$ values between 0.1 and 1024 Small heart-shaped ROI	

Figure 5.2: Pulses used for each comparison. w_r is the weighting assigned to the real (i.e., non-CPMG) signal component, and λ is the Tikhonov RF power regularisation term. The terms are used during pulse design as described in Equation 3.5.

Pulses were designed for a mid-axial slice from the B_1^+ map obtained at the centre position. Figure 5.2 summarises which pulses were used for each pair-wise comparison. Inner-volume imaging was used for most pulses here, and was previously described in Chapter 2. Briefly, inner-volume imaging is 2- or 3-D spatially-selective, and is especially useful for increasing signal uniformity, increasing resolution (if scan time is held constant), or speeding up acquisitions (if resolution is held constant) when imaging small ROIs [82, 83].

Pulse performance under conditions of motion was assessed by evaluating each pulse on the corresponding axial slice of the 103 off-centre B_1^+ maps. Slices were defined in patient coordinates, meaning that even when through-plane motion was simulated, the same slice was used for design and evaluation. In practice, this requires prospective gradient and central frequency adjustment as described in Chapter 4. Flip-angle profiles were generated using a STA Fourier approximation and Bloch simulations in the phase-relaxed CPMG and fast OC cases, respectively.

Normalised root-mean-squared error (nRMSE) was used to quantify pulse performance. First, the RMSE for the centred position (P0) was calculated and normalised by the target profile. This reflected baseline error intrinsic to the pulse design, independent of motion.

Next, to isolate the effects of motion, RMSE was calculated using the difference between central and off-central flip-angle profiles, and normalised to the target flip-angle:

$$nRMSE = 100 \times \frac{\sqrt{\frac{1}{Nr} \sum_{r=1}^{Nr} \|\alpha_{displaced}(r) - \alpha_{centre}(r)\|^2}}{\alpha_{target}} \quad (5.1)$$

where α_{target} is the target flip-angle profile, α_{centre} is the profile achieved at the centre position (i.e., without motion), and $\alpha_{displaced}$ is the profile achieved in off-central position (i.e., following motion). Nr is the number of voxels in the masked region.

For comparison 5, error was calculated separately for non-CPMG (real) and CPMG (imaginary) signal components. The primary interest was the CPMG component because error in the non-CPMG part would not be refocussed in a turbo spin echo sequence, as described in Chapter 3.

5.3 Results

5.3.1 Comparison 1: STA (excitation) vs LTA (inversion) pulses.

Figure 5.3 shows the target and central position (P0) profiles produced by pulses *i* (STA) and *ii* (LTA). Both pulses yielded high-quality excitation or inversion profiles without motion.

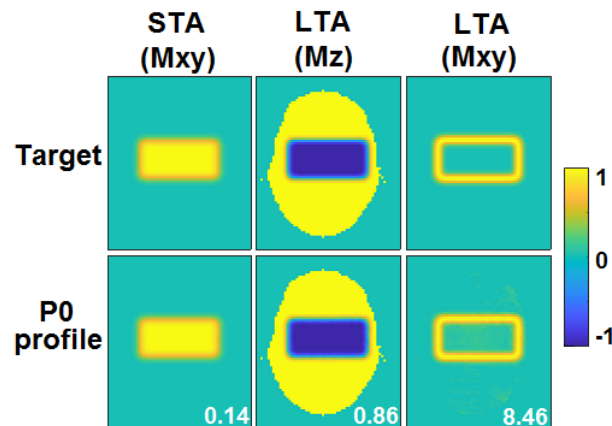


Figure 5.3: Target profiles (top row) and performance (bottom row) of pulses *i* (left column) and *ii* (middle and right columns) at the central position (i.e., without motion). The error (nRMSE [%]) without motion is indicated at the bottom right of each P0 profile. Error is $\leq 1\%$ in the transverse and longitudinal components for the STA pulse and LTA pulse, respectively. The colour bar shows magnetisation in M_{xy} or M_z . The target rectangular ROI shown here (and in all subsequent analyses in this Chapter except where stated otherwise) has dimensions 5×10 cm.

Figures 5.4 and 5.5 demonstrate that the effect of motion on the pulses' performance differed both quantitatively and qualitatively. The LTA pulse was generally more susceptible to motion than the STA pulse. One notable exception was inferior-superior (Z-axis) translation, where motion-related error was lower with the LTA pulse for all positions along this dimension (Figure 5.4, top right). The extent of motion was approximately linearly related to error in the STA regime, while for LTA the relationship appeared quadratic for some movement types (left-right and inferior-superior translations). Figure 5.5 shows magnetisation error maps for a range of positions. The maps show that error in the STA regime was relatively smooth and uniform across the slice, while for LTAs error was less uniform overall and there were small regions of high error. Following small movements,

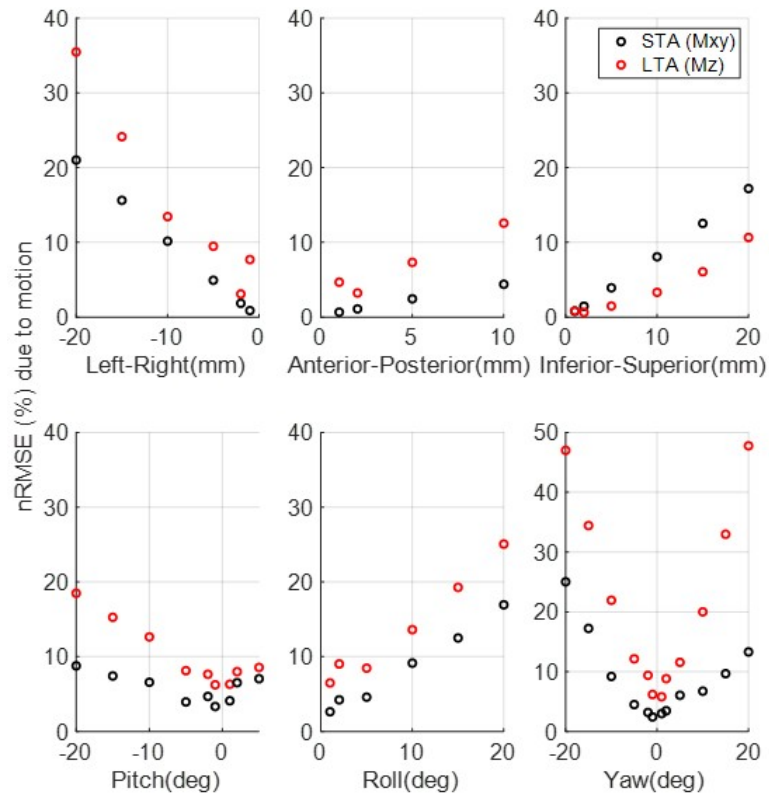


Figure 5.4: Motion-related error for pulse i (STA, black) and ii (LTA, red) designed to excite or invert the same inner-volume target. Error in M_{xy} is shown for the STA pulse, while error in M_z is shown for the LTA pulse. Error is with respect to the respective P0 profile (without motion), normalised to the respective target profile - both shown in Figure 5.3.

profiles produced with LTA pulses were most susceptible to error at locations surrounding the ROI, while the ROI itself was most prone to error when STA pulses were used.

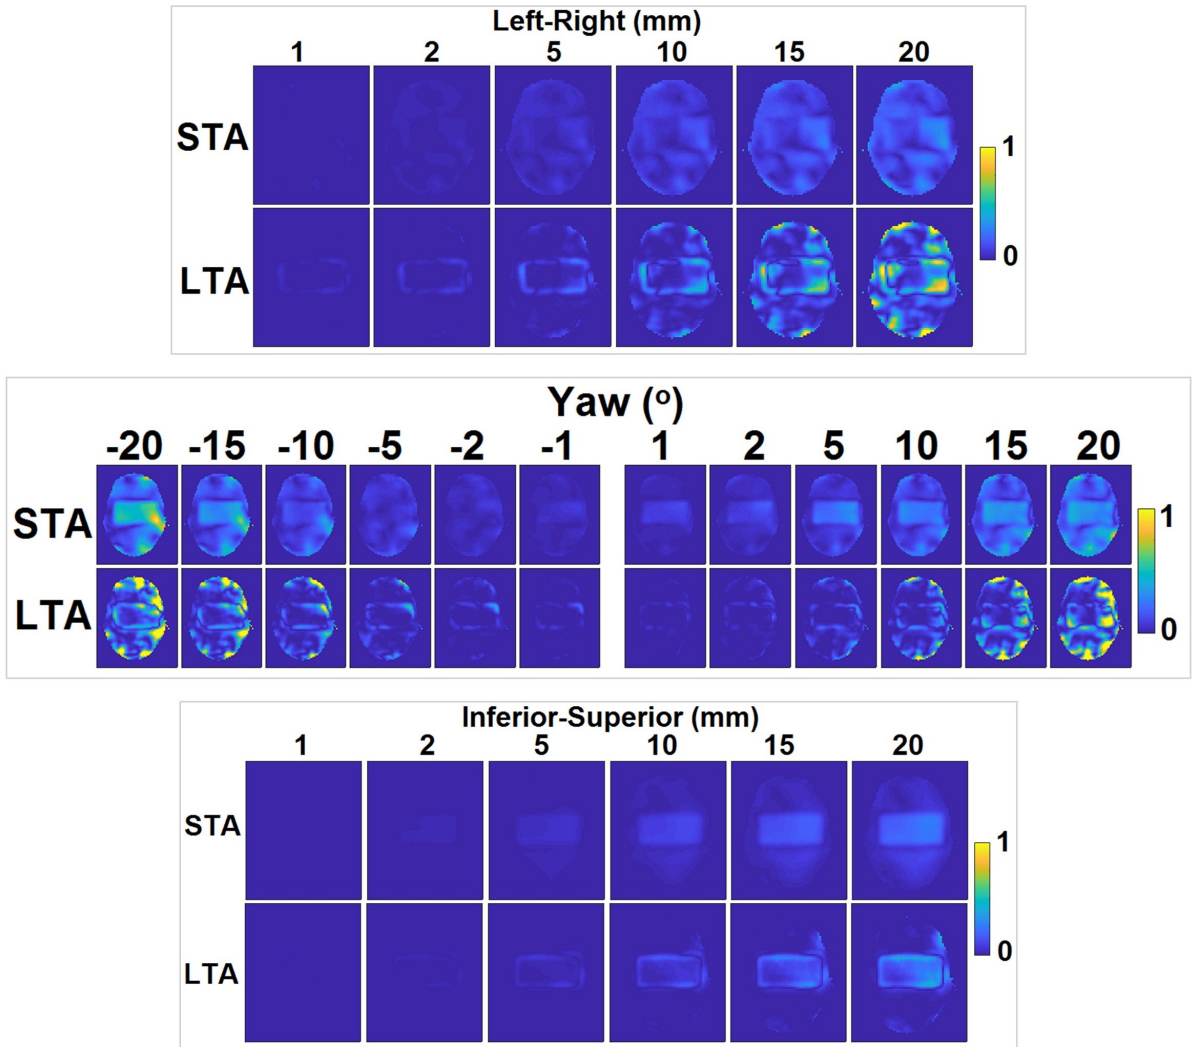


Figure 5.5: Maps showing motion-related error in magnetisation following displacements in rightward, yaw, and superior direction for pulses i (STA) and ii (LTA). An error value of 1 indicates approximately 90° flip-angle error.

5.3.2 Comparison 2: The effect of target flip-angle.

For comparison 2, pulses were designed using fast OC with small (pulses *iii* and *iv*) and large (pulses *ii* and *v*) target flip-angles. Figure 5.6 shows the target and central profiles for pulses used in comparison 2. Single-spoke whole-slice excitation (70°) and inversion (180°) yielded a higher baseline error than the inner-volume pulses - with nRMSE at the central position being over 30% higher for the former. This was expected, since signal uniformity is harder to achieve when increasing the area over which pulse design optimisation is performed (as described in Section 2.3.1).

Figures 5.7 (left panel) and 5.8 show motion-induced error for pulses *ii* and *iii*. Like in comparison 1, motion-related error in the STA case was lower than that in the LTA case except following inferior-superior motion, for which error was very similar. For both pulses, background suppression remained reasonably good, even when movement was large in the STA case. This is demonstrated most clearly in the inferior-superior 20 mm case, where almost no error was present outside of the ROI.

Compared to STA pulse *i* used in comparison 1, motion-related error for STA pulse *iii* was less uniform across the slice, although it was still more uniform than that for LTA pulse *ii*.

Motion-related error was also consistently higher when LTA pulses were used for single-spoke excitation/inversion (*iv* and *v*), as shown in Figures 5.7 (right panel) and 5.9. The discrepancy between STA and LTA pulse performance was exaggerated when the movement was large, as can be seen from the sharply diverging error across all motion directions in Figure 5.7. This was also the case for the inner-volume spiral pulse, but to a lesser extent. In the out-of-slice (I-S translation, pitch and roll) and anterior-posterior directions, the LTA pulse with large movements produced almost triple the error of the STA pulse.

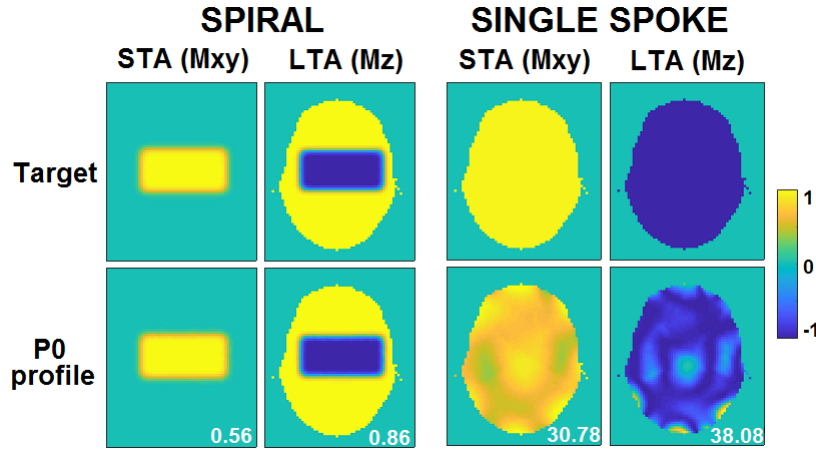


Figure 5.6: Target profiles (top row) and performance at the central position (bottom row) for the pulses used in comparison 2. STA pulses iii (column 1) and iv (column 3) and LTA pulses ii (column 2) and v (column 4) were all designed using fast OC. nRMSE (%) without motion is shown in the bottom right of P0 profile panels. The colour bar shows magnetisation in M_{xy} or M_z .

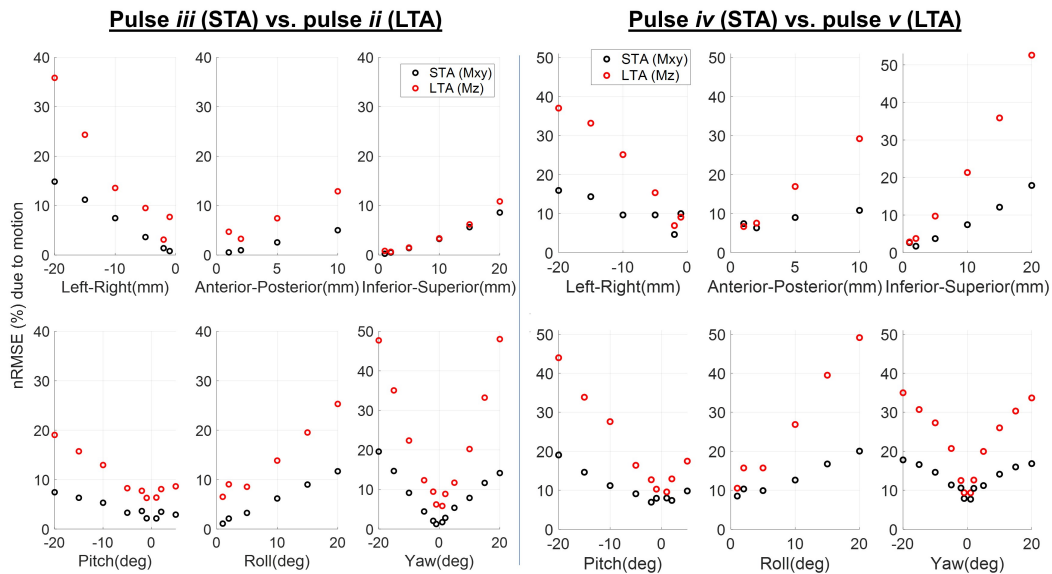


Figure 5.7: Motion-related error for STA (90° , black) and LTA (180° , red) pulses designed using fast OC. The left panel shows performance of spiral inner-volume excitation/inversion, while the right panel shows performance of single-spoke whole-slice excitation/inversion. Error is calculated with respect to the no-motion cases, and normalised to the target profiles shown in Figure 5.6.

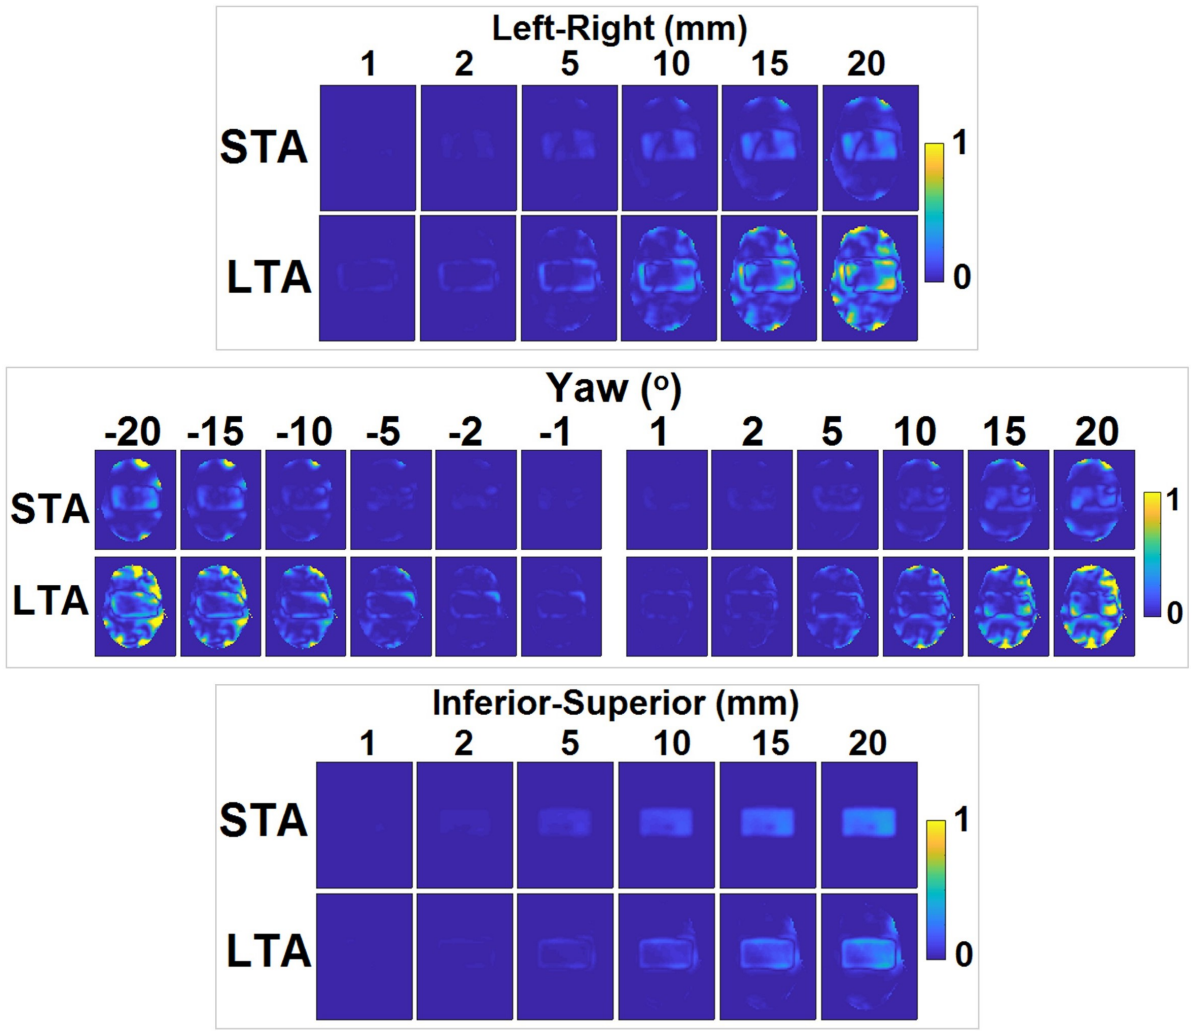


Figure 5.8: Maps showing motion-related error in magnetisation following displacements in rightward, yaw, and superior for pulses iii and ii, both designed using fast OC. An error value of 1 indicates approximately 90° flip-angle error.

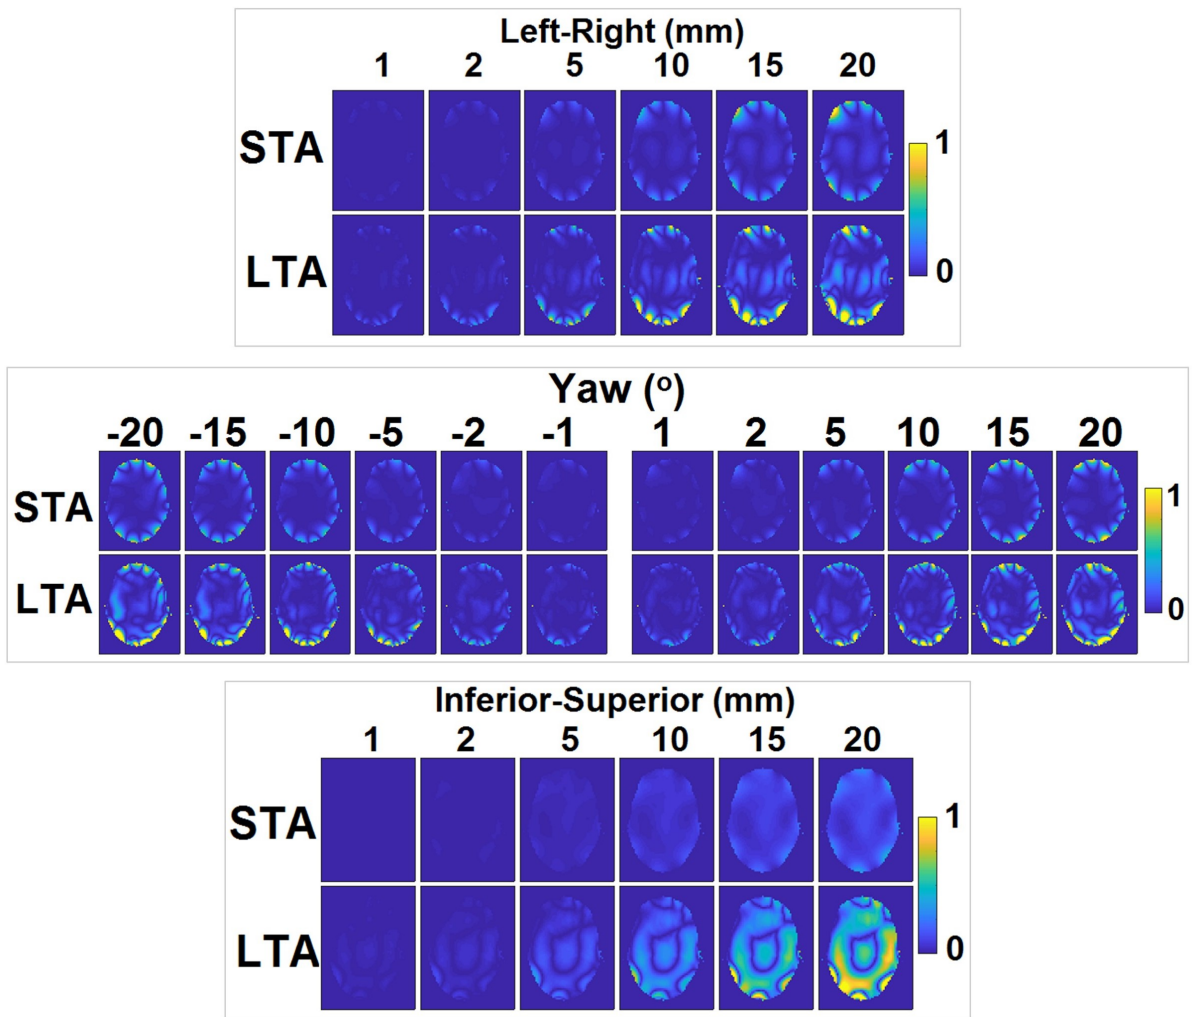


Figure 5.9: Maps showing motion-related error in magnetisation following displacements in rightward, yaw, and superior for pulses iv and v, both designed using fast OC. An error value of 1 indicates approximately 90° flip-angle error.

5.3.3 Comparison 3: The effect of design method.

The third performance comparison was between the 90° excitation pulses designed using the two different design methods (pulses *i* and *iii*). Figure 5.10 shows the target profile and good performance of both pulses in the central position.




Figure 5.10: Target profile (top) and performance at the central position (bottom) for pulses *i* (left) and *iii* (right). $nRMSE$ without motion is shown. Both pulses performed well, so the colour scale (showing M_{xy}) was saturated to emphasise differences between them.

For most movements, both pulses were affected similarly by the same magnitude of movement, as shown in Figure 5.11. Where substantial sensitivity differences were present (L-R and I-S), they were in the opposite direction to those observed in the previous comparisons – that is, the fast OC pulse was less sensitive to motion than the CPMG pulse.

Generally, error distributions between pulses *i* and *iii* were comparable, however as in comparison 1, the fast OC pulse resulted in more spatially-varying error across the slice compared to the phase-relaxed CPMG pulse, shown in Figure 5.12. This was especially noticeable in peripheral locations proximal to the channels where error on the order of 90° was observed with the fast OC pulse. However, spatial error variation remains substantially lower than that observed with the LTA pulses observed in comparisons 1 and 2.

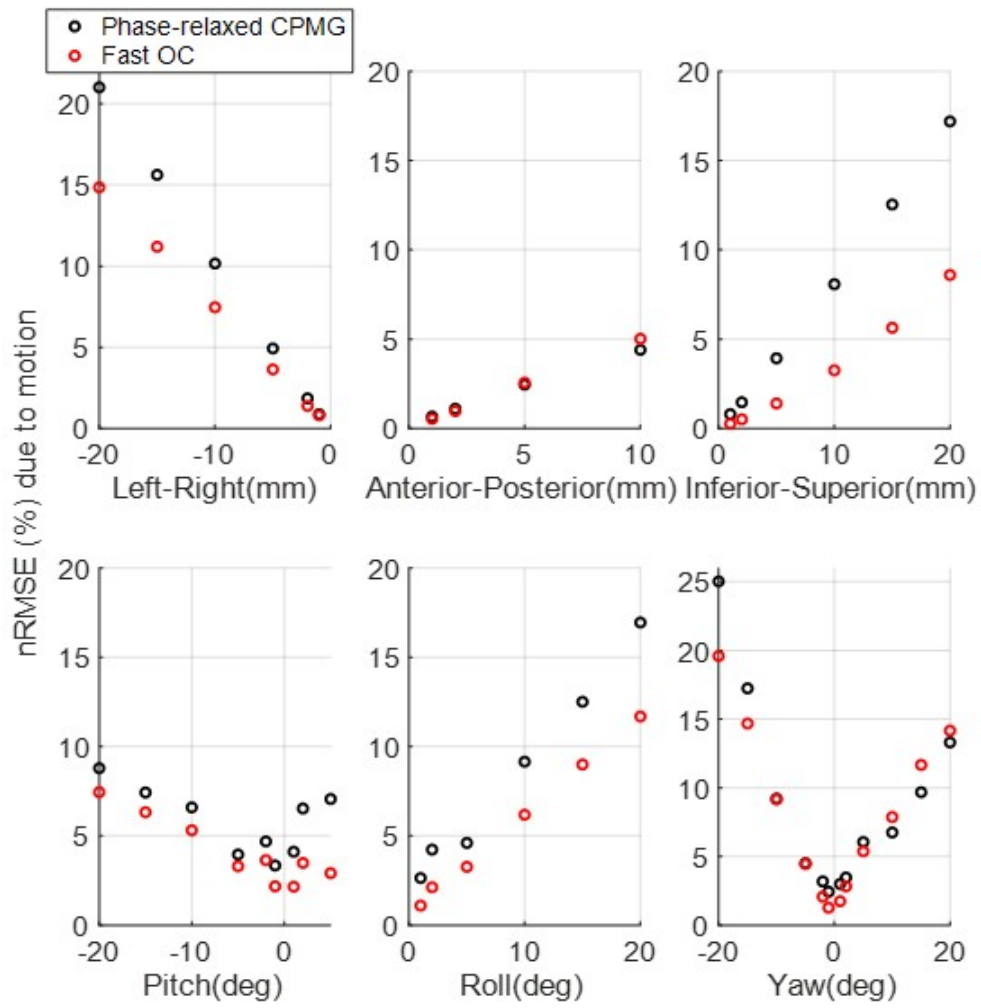


Figure 5.11: Motion-related error for 90° excitation pulses designed using phase-relaxed CPMG (black) and fast OC (red) methods. Error is with respect to the no-motion cases, and normalised to the target profile shown in Figure 5.10.

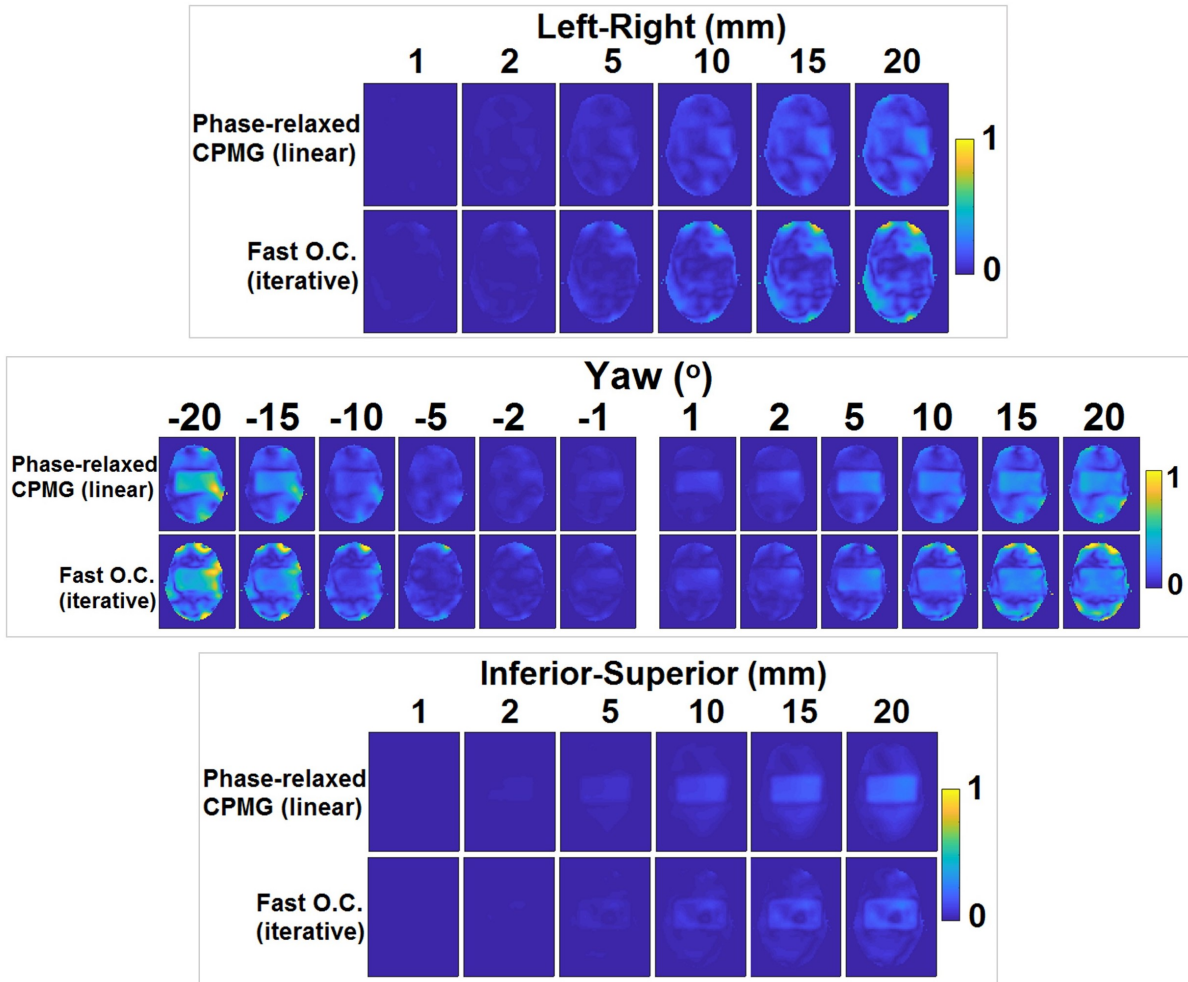


Figure 5.12: Maps showing motion-related error in transverse magnetisation following displacements in rightward, yaw, and superior for pulses i and iii - the two 90° excitation pulses designed using the two different methods. An error value of 1 indicates approximately 90° flip-angle error.

5.3.4 Comparison 4: LTA with fast OC

Figure 5.13 shows baseline performance for the spiral (ii) and EPI (vi) LTA pulses designed using fast OC. The inversion profile produced using a spiral trajectory pulse in a central position was superior to that using an EPI trajectory. This is most apparent within the inverted ROI where a non-trivial amount of M_{xy} was generated by the EP pulse, although small areas of unwanted magnetisation were also apparent outside of the ROI.

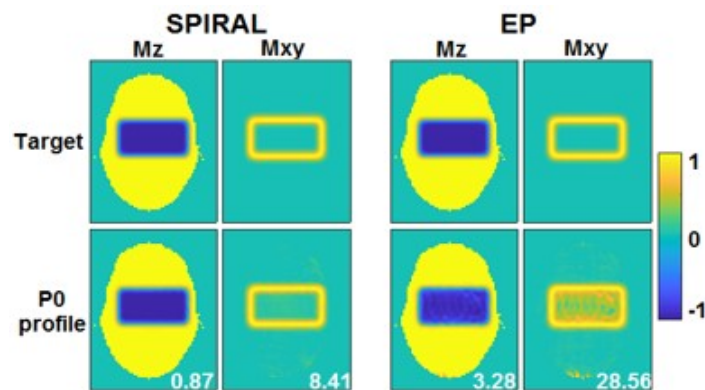


Figure 5.13: Target profiles (top) and performance at the central position (bottom) for pulses used in comparison 4. Two 180° inversion pulses (pulses ii - left, and vi - right) were designed using fast OC, with a spiral and EPI trajectory, respectively. $nRMSE$ (%) without motion is displayed. The colour scale shows M_{xy} or M_z as indicated.

Figure 5.14 shows that, for most positions, the trajectories demonstrated similar sensitivity to motion, with a slightly lower error seen for the EPI trajectory following some displacements. Results for pitch rotation demonstrate the important distinction between motion sensitivity and performance; motion-induced error was consistently higher for the spiral pulse, however when looking at the M_z maps shown in Figure 5.15, it is clear that the profile produced by the spiral pulse -especially within the ROI- is in fact superior up to around 5 mm or 5° of motion.

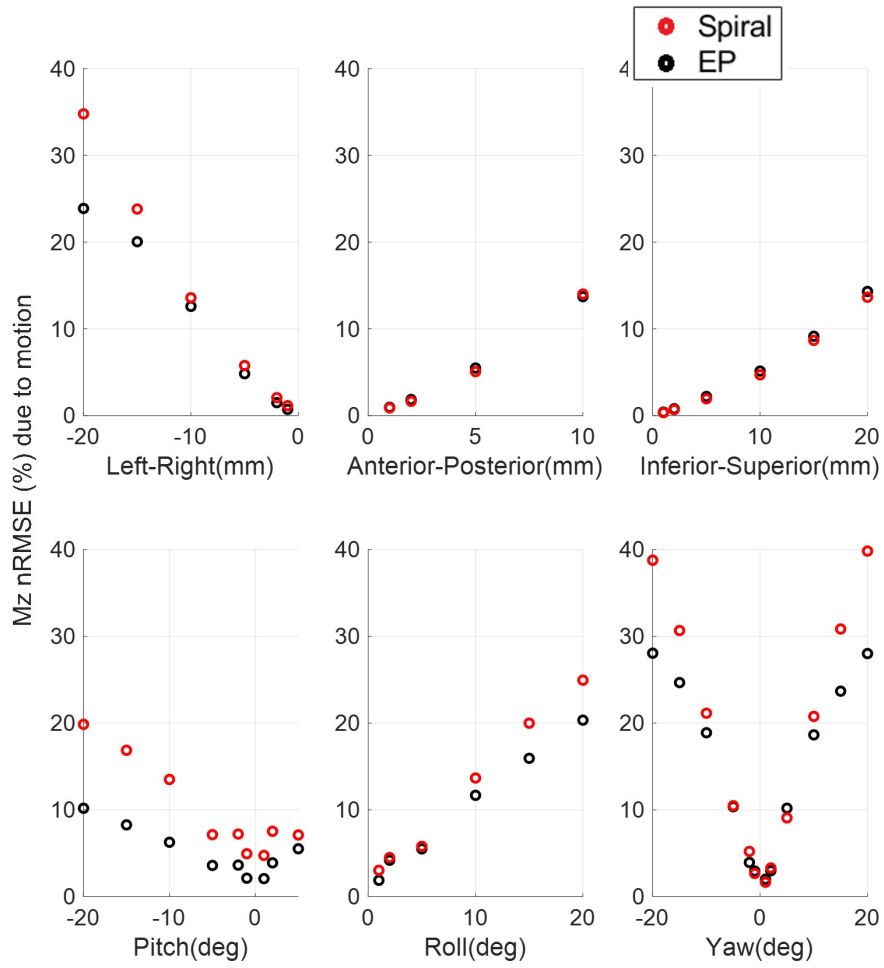


Figure 5.14: Motion-related error for 90° excitation pulses designed using phase-relaxed CPMG (black) and fast OC (red) methods. Error is with respect to the no-motion cases, and normalised to the target profile shown in Figure 5.10.




Figure 5.15: Maps showing inversion profiles following displacements in rightward, yaw, and superior for pulses ii and vi. The colour scale shows M_z .

5.3.5 Comparison 5: Varying weighting values in STA phase-relaxed CPMG pulses

The following results used a mid-range RF power regularisation term ($\lambda = 8$); the effects of changing λ are mentioned later. A w_r of 1 is equivalent to conventional STA pulse design (i.e. CPMG and non-CPMG components are weighted equally) and yielded nRMSE of 5.4% in the signal magnitude (with respect to the target profile), reflecting an acceptable initial M_{xy} excitation.

STA pulses were designed to excite a small heart-shaped ROI. Their performance without motion is shown in Figure 5.16. Subtle regions of excitation in the CPMG/imaginary component outside of the ROI when $w_r = 1$ were present, but these were reduced when a smaller w_r was used instead. As expected and in line with previous work [110], pulses with a low w_r resulted in increased error in the overall signal magnitude following excitation. However, error in the CPMG/imaginary signal component was reduced by around 12% compared to when $w_r = 1$, meaning an overall improvement in the profile when CPMG refocussing pulses are subsequently applied (not shown here, but the principle is explained in Chapter 3).




Figure 5.16: Maps showing excitation profiles without motion. Imaginary (CPMG) and real (non-CPMG) signal components are shown separately, in addition to the magnitude error. The colour scale shows flip-angle.

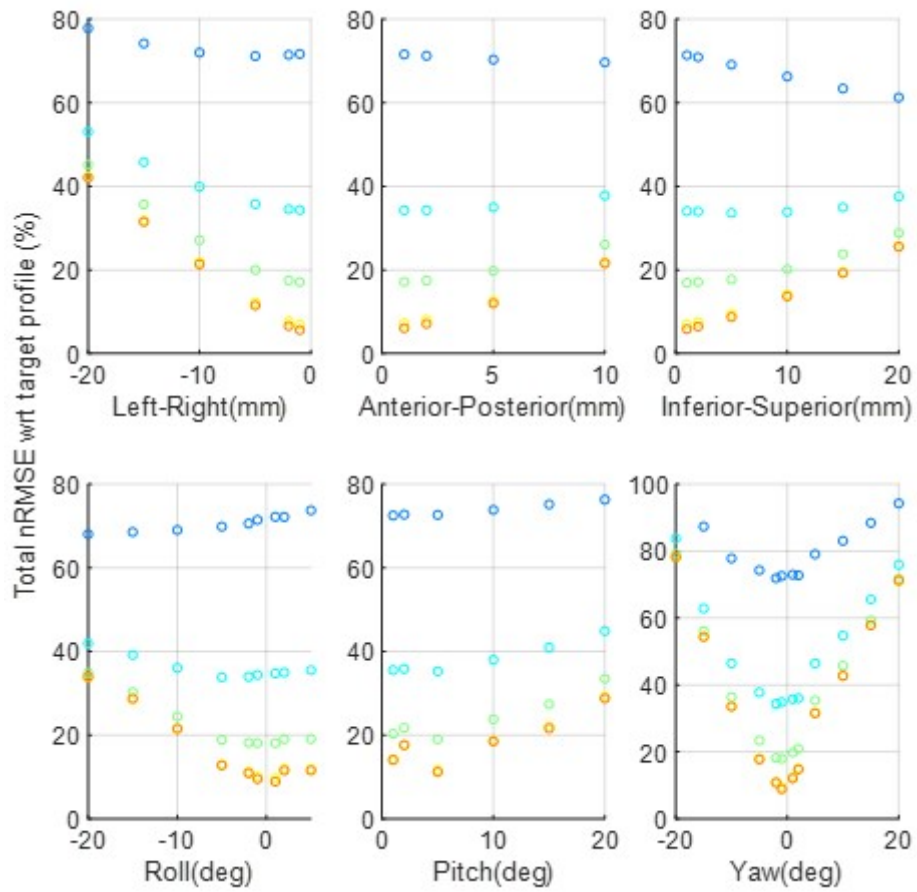


Figure 5.17: Absolute error with respect to the target profile, shown as nRMSE (%). This includes error in both CPMG and non-CPMG signal components. For clarity, the off-axis positions are not shown in the figure. Different colours show different w_r weightings - the colour legend is shown below, in Figure 5.18.

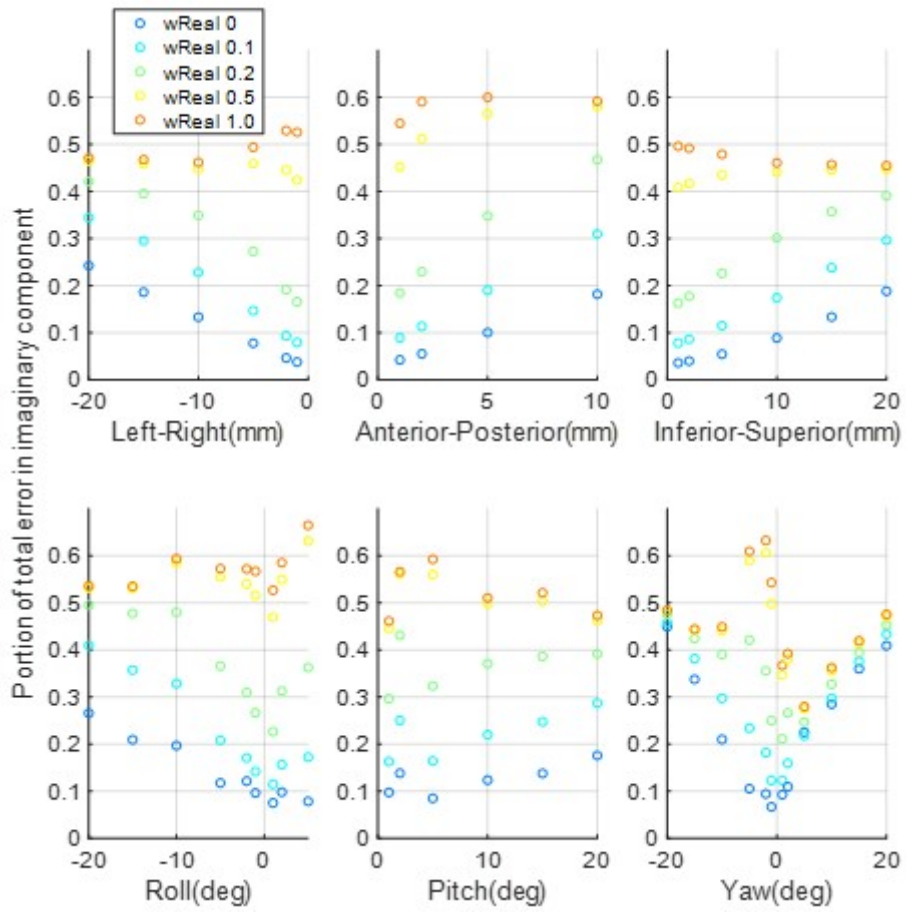


Figure 5.18: The proportion of error shown in Figure 5.17 which arose in the CPMG (imaginary) signal component. For clarity, the off-axis positions are not shown in the figure.

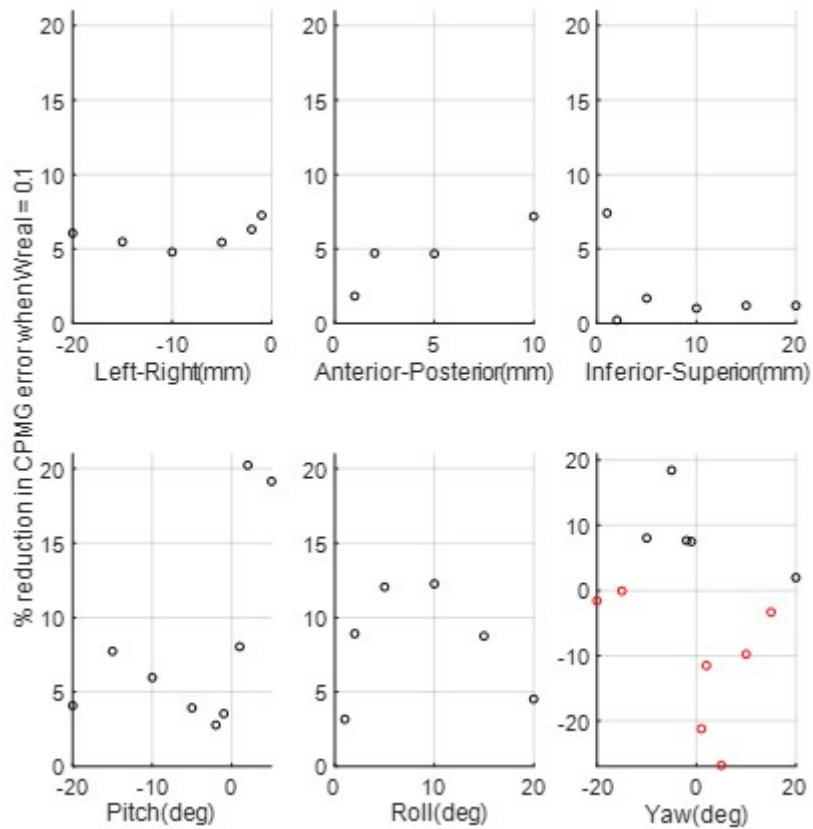


Figure 5.19: Reduction in motion-related error in the CPMG signal component by setting w_r to 0.1, compared to when $w_r = 1$ (negative values indicate that error was lower when $w_r=1$, and are shown in red). For clarity, the off-axis positions are not shown in the figure.

When a low w_r was used, total M_{xy} error (including both CPMG and non-CPMG components) was high, but relatively insensitive to motion ($w_r = 0$, shown in dark blue, is fairly consistent across positions). When a high w_r was used, baseline M_{xy} error was much lower (orange and yellow plots in Figure 5.17), but more sensitive to motion; nRMSE across positions varied by 20-70% depending on the direction and extent of motion.

The effect of w_r on error within the CPMG component (as a proportion of the total error) is shown in Figure 5.18. Unsurprisingly, with a conventional w_r of 1, both real and imaginary parts contributed approximately equally to the total error. When $w_r = 0$, less than a quarter of the total error was comprised of the CPMG part, even following large movements.

By using $w_r = 0.1$ compared to the conventional $w_r = 1$, error in the CPMG signal component was reduced by a mean value of 5.3%, with maximum error reduction (20.2%) following small pitch movement. An exception was seen for some yaw positions, where error

in the CPMG signal component was sometimes higher when w_r was small. When w_r was set to zero, the average error reduction was similar (5.3%), however less positions benefited.

An alternative benefit of setting w_r to 0.1 (compared to 1) is to offer more freedom for movement, assuming that a tolerable error level in the CPMG signal component is specified. Using best fit lines (not shown) between motion magnitude and CPMG error, the movement tolerance provided by setting w_r to 0.1 (compared to 1) is estimated to be on the scale of millimetres.

RF power regularisation (λ) also becomes relevant when choosing an optimal w_r to minimise motion-induced error. When λ is large (>1), CPMG error was reduced by choosing $w_r = 0.1$ for almost all positions as described above. For small λ values (<1), the reductions in CPMG error gained by using low w_r were less pronounced, and were position-dependent.

5.4 Discussion

This study used simulations to investigate whether differences in motion sensitivity exist between LTA and STA pTx pulses at 7T. Our results for STA pulses fall in line with those previously reported, with the relationship between displacement magnitude and error in the flip-angle profile being approximately linear [173]. Here, LTA inversion pTx pulses were affected by motion differently. We found that quantitatively, LTA pulses are more motion-sensitive, with motion-induced error being higher than that caused by the same movement with STA pulses. Additionally, we have shown that qualitative differences exist; error is more spatially varying and less smooth across the imaged object for the LTA pulses in all cases, making global correction approaches especially unsuitable.

The patterns observed in comparisons 2 and 3 helped to clarify whether it is the larger flip-angle itself or the specific methods tested here which primarily drove the difference in performance between STA and LTA pulses. Error observed in comparison 1 followed a similar pattern to that of comparison 2 (where flip-angle was changed, but the design method held constant), suggesting that flip-angle itself is the primary factor underlying the difference in motion sensitivity. Furthermore, in comparison 3 (where a 90° pulse was designed using both methods), the similarity in nRMSE between pulses indicates that

contributions of the design method to overall differences observed are minimal (indeed, we observed the opposite pattern, where the fast OC pulse was less motion-sensitive in terms of nRMSE than the phase-relaxed CPMG pulse in some cases).

The qualitative differences between pulses in comparison 1 are present in comparisons 2 and 3 to some extent, with LTA pulses and / or pulses designed using fast OC demonstrating more spatial variation in error than the STA phase-relaxed CPMG pulse. Although this prevents a conclusive determination of contributions from target flip-angle vs. design method from being drawn, patterns within comparison 3 offer some further clarification. Here, the major difference in error distribution between the two 90° pulses was seen in peripheral regions proximal to the transmit channels, where the fast OC pulse yielded higher error. One possible explanation is a difference in RF power as between the two methods. The phase-relaxed CPMG design incorporates VERSE to limit peak RF power (and therefore, RF power in peripheral regions close to the channel is likely to be reduced), whereas fast OC does not. It is possible that this mechanism contributed to the differences observed between the two pulses, and indicates that the power-limiting effect of VERSE may help to avoid areas of high motion-related error in B_1^+ locations near to channels.

Motion sensitivity of inversion pulses with a spiral vs. EPI excitation trajectory was similar. Although the EPI pulse was less motion-sensitive in terms of quantitative error, the spiral pulse yielded a favourable flip-angle profile (better ROI inversion, and better non-ROI suppression) when the movement was small (≤ 5 mm or 5°). For larger movements, the EP trajectory appeared to be marginally less motion-sensitive due to superior background suppression.

Our results suggest that the phase-relaxed CPMG excitation pulse approach may offer some motion-robustness in the initial excitation within the context of TSE sequences. Error reduction from using a low w_r is around 12% without motion, and between 5 - 20% following motion, although benefits offered by this approach also depend on the level of RF power regularisation.

In comparison 2, the effect of target flip-angle on error following inferior-superior (I-S) translation was negligible. This suggests that the slightly better performance of the fast OC method (as seen in comparison 3) drove the differences observed for this motion type in

comparison 1. In the fast OC design, M_z (which is in line with the direction of I-S motion) is explicitly defined and included in the optimisation, which may explain the better performance. Additionally, for pTx arrays where channels are cylindrically-symmetric about the longitudinal axis (as they are in the simulations here), B_1^+ field dynamics across the diameter of the array (i.e. an axial cross-section) are different from those along the coil's long axis. The latter is more slowly-varying in comparison, meaning that movement in the I-S direction is therefore likely to result in a less perturbed B_1^+ field than a similar magnitude in-plane movement. This combination (M_z -optimised performance, and a small change in B_1^+) therefore results in lower error with the fast OC pulse in this case. This is corroborated by the fact that, when considering error in M_{xy} , the STA pulse outperformed the LTA pulse for all movements including I-S (data not shown).

It should be noted that the transmit coil model used in simulations had an array diameter of 230 mm, which may differ from the 7T pTx coils commonly used in practice such as the 8-channel transmit/32-channel receive Nova Medical pTx coil. Due to proprietary information surrounding the Nova Medical coil, only the inner diameter of the receive array housing is known (205 mm), and the diameter of the transmit coil array is unknown. The results presented here therefore may differ from those obtained using other, realistic coils, and this topic is further investigated in Chapter 6. Nevertheless, we anticipate that the observed differences in behaviour between STA and LTA pulses would remain.

To the best of our knowledge, this is the first demonstration that motion-induced error for STA pTx pulses differs from that of LTA pTx pulses both quantitatively and qualitatively, with LTA pulses generally being more sensitive to motion. Together, the findings presented here suggest that fundamental differences in behaviour and requirements between STA and LTA pTx pulses –and not nuances within each design– underlie the observed differences in motion sensitivity. For this reason, different approaches are likely to be required to address the issue of motion for STA and LTA pTx pulses. One potential avenue would be to introduce motion-robustness into specific protocols and sequences (e.g., by designing a motion-robust LTA pulse). A more flexible approach would be to correct for motion in a dynamic way which allows tailored pulses to be designed using arbitrary methods, while retaining good performance in cases of motion. Chapters 8 and 9 of this thesis present approaches which fall into the latter category.

Chapter 6

Specific Absorption Rate and pTx coil design

6.1 Overview

SAR is sensitive to head motion, especially in pTx due to channel interference and changing proximity of the head to RF transmit elements. This can cause SAR limits to be exceeded following motion when SAR-constrained pulses are designed for one position [130]. In the literature, local SAR has been reported to more-than-triple due to head motion [130].

Previous works examining the impact of motion on SAR have used simulated data from a single body model and a single coil model [130, 190, 191] (although models differed between the studies). Therefore, it is unclear whether (and to what extent) the findings are generalisable across various imaging scenarios.

The following two chapters present studies examining whether (and how) certain factors may affect the sensitivity of SAR to motion in pTx. In this study chapter, the effect of pTx coil geometry on sensitivity of SAR to motion is investigated by simulating 6 differently sized pTx coil models. For each model, the SAR produced by a pTx pulse designed for one central position is evaluated at 19 displaced positions. In Chapter 7, the effects of inter-subject variability are investigated in an analogous way by investigating SAR sensitivity to motion across several different body models.

The work presented in this chapter led to the following conference proceedings:

Plumley, A. et al. Parallel-transmit coil dimensions affect SAR sensitivity to motion at 7T. Proceedings of the International Society for Magnetic Resonance in Medicine, #2315 (2021).

6.2 Introduction

In Chapters 1 and 3, the significance of E-field distributions in pTx were discussed. Briefly, E-fields are directly implicated in SAR and therefore tissue heating patterns. Each pTx channel produces its own E-field, and these interfere, producing local SAR hotspots (small areas of high SAR), especially in locations where channels constructively interfere. If E-field distributions or compressed representations such as virtual observation points (VOPs) are known, an upper-bound on local SAR can be specified during pulse design [192], but crucially, the upper bound cannot accurately be determined when E-fields change. In general, local SAR cannot readily be measured or monitored on a subject-specific basis. Instead, global SAR is constrained during pulse design (e.g., by using RF power limits), and pre-determined SAR models (e.g., VOPs derived from a database of body models) are integrated into vendor software to estimate local SAR. This could be seen as concerning, given the likelihood of local SAR hotspots in pTx (although there are of course widely accepted risks associated with non-pTx MRI, which are arguably equally concerning).

Recent literature has shown that head motion can cause peak local SAR to increase by up to 3.1-fold when pTx pulses are used at 7T [130, 190, 191]. Even in quadrature mode (see Section 3.4.1) which is generally considered safe, peak local SAR could increase by 2.1-fold following axial displacements [130]. Guidelines suggest only constraining whole-head SAR when quadrature mode is used, which could elicit a false sense of security because whole-head SAR increased by just 5% for the same displacement. Therefore, if pulses are calibrated at, or close to, the global SAR limit, motion can cause SAR (especially local SAR) to exceed safety limits.

RF power absorption can be estimated in near real-time using the inbuilt directional couplers to measure forward and reflected power, however this can only be used to estimate global SAR, and not local SAR [19, 128, 193]. More recently, techniques have emerged which can

estimate the theoretical worst-case peak local SAR in real-time [132, 194], although they often overestimate the actual SAR produced by realistic pulses, limiting imaging performance.

Electromagnetic field interference patterns depend on the size, composition, and geometry of the coil [127, 195]. Deniz et al. investigated the effect of increasing the distance between coil and sample (while simultaneously increasing the size or number of pTx channels) [127]. Global- and peak local SAR decreased with increasing coil-to-sample distance, while similarly uniform excitation profiles were achieved. Peak local SAR approximately halved by increasing the coil-to-sample distance from 1 cm to 4 cm, and SAR distributions were noticeably smoother. The benefits were attributed to greater effective degrees of freedom available in pTx pulse design when fields are smoother, since channels contribute RF more equally to all spatial locations.

It is unclear whether similar variations in coil design may affect the response of SAR to motion. Better understanding of the relationship between coil design and motion sensitivity of SAR is therefore needed, and could be leveraged in the design of safe pTx coils. To this end, here we investigate whether the dimensions of pTx coil elements affect the sensitivity of SAR to head motion at 7T. For this, we compare motion-induced SAR changes using simulations with 6 different coil models.

6.3 Methods

The simulation protocol followed that outlined in 5.2.1. Six coil models (A to F) shown in Figure 6.1 were simulated. Each model consisted of 8-channel loop arrays. The loop dimensions of coil C (the base model) were 40 mm (width) x 110 mm (height), and the array radius was 115 mm. Compared to coil C, coils A and B had 25% narrower and shorter loops, respectively, while coils D and E had 25% taller and wider loops, respectively. Coil F had taller, wider loops, and the array radius was 25% larger.

For each coil model, one central, and 19 off-centre positions were simulated using the Ella model (IT'IS, Zurich, Switzerland) [143] totalling 120 datasets. Kopanoglu et al. [130] and Ajanovic et al. [190] found that SAR was most sensitive to displacements within the axial

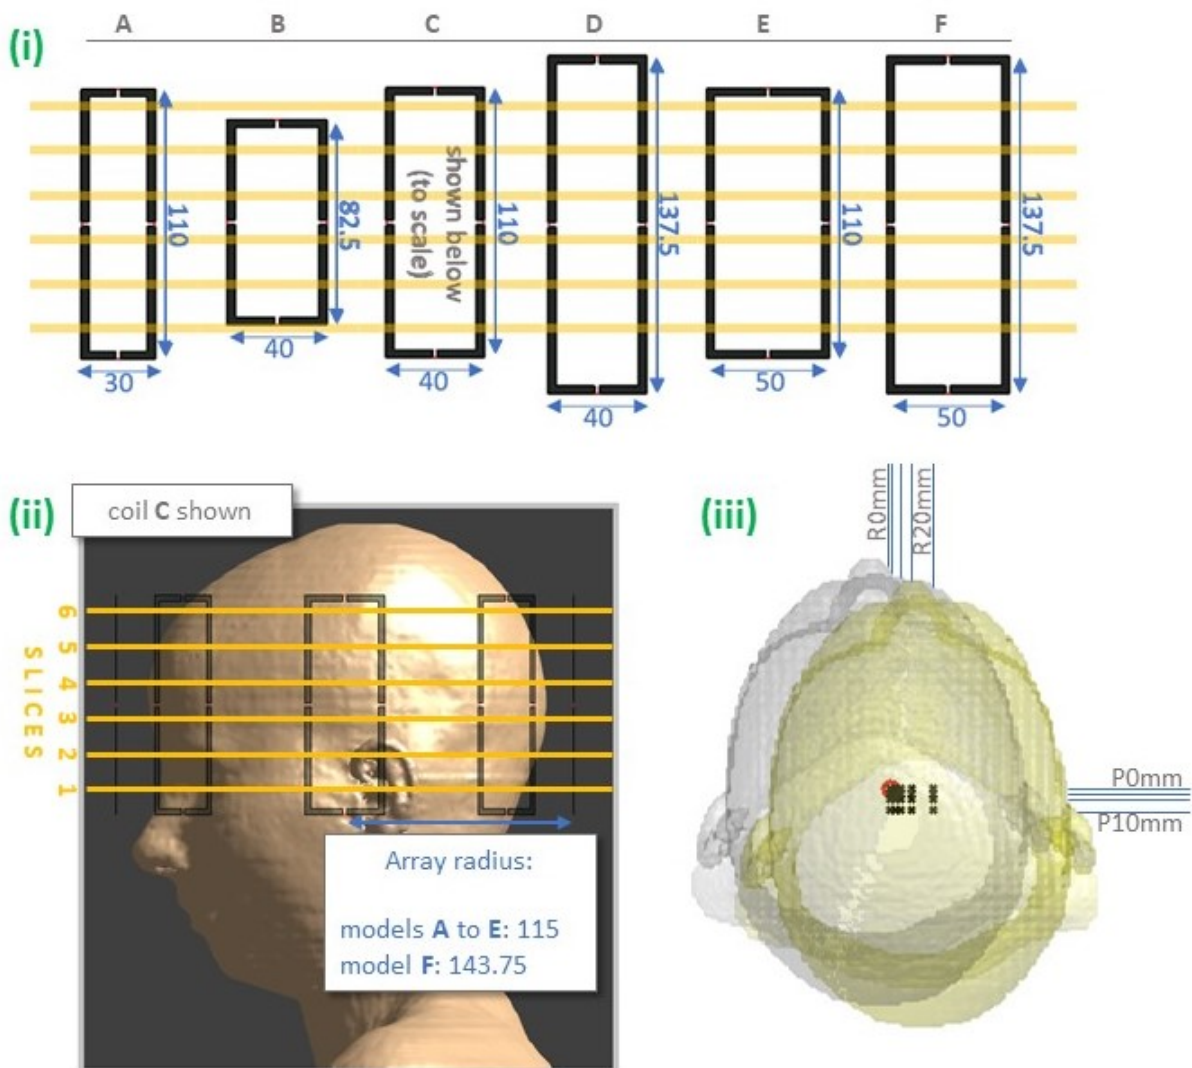


Figure 6.1: (i): Dimensions (in mm) of the 6 pTx coil models. A single loop (channel) is shown for each model. (ii): Sagittal view of Sim4Life setup (coil model C shown). Slice positions -also shown in (i)- used for pulse design are indicated in yellow. Slices remained constant in scanner coordinates across coil models. (iii): Axial view of simulated positions. The 20 positions were combinations of: rightward: 0, 2, 5, 10, 20mm and posterior: 0, 2, 5, and 10mm. The two extremes (0, 0 – grey) and (20, 10 – yellow) are shown as isosurfaces. The origins of all positions are indicated with crosses (with the central position's origin circled in red).

plane, so we focus on these here. Off-centre positions were therefore rightward 0, 2, 5, 10, 20 mm and posterior 0, 2, 5, 10 mm, as well as the grid of all rightward-posterior combinations (Figure 6.1iii). In all cases, positions were defined with respect to the coil array's origin. Following each simulation, B_1^+ , E-field, current density, SAR, and tissue density distributions were extracted.

RF pulses were designed to uniformly excite an axial slice using the spatial domain method [32, 196] with the B_1^+ distributions from the central position for each coil. Quadrature, and pTx (1/2/3/5-spokes; where 1-spoke is RF-shimming) pulses were designed. Separate pTx pulses were designed for the 6 slice locations shown in Figure 6.1, and one quadrature pulse was designed (since the effects of motion on quadrature pulses are independent of slice location), yielding a total of 25 pulses per coil. Pulses were not SAR-constrained, but RF power was penalized (along with magnitude excitation error) in the cost function using Tikhonov regularisation ($\lambda = 0.1$ in Equation 3.3).

To include the effects of pTx channel interactions, Q-matrices were used for SAR calculation [129]. The approach is described in more detail in Chapter 3 (see Equation 3.8). SAR was evaluated in terms of three metrics: whole-head SAR (gSAR), 10g-averaged local SAR (psSAR), and eigenvalue-based SAR (eigSAR). EigSAR is calculated from the eigenvalue of the Q-matrices, illustrating the SAR that would occur if channels within a given array interacted in the most constructive way possible. It therefore excludes the effects of pulse design, producing the worst possible SAR for a given coil-subject setup. Although this is physically possible, it is relatively unlikely to occur when realistic pTx pulses are used. For gSAR, SAR calculated using voxel-wise Q-matrices was summed over the entire head. For the psSAR and eigSAR, 10g averaged SAR was calculated by using Q-matrices averaged over a 10g local tissue region (tissue information was obtained from simulations). Peak (for psSAR and eigSAR) or total (for gSAR) SAR was calculated for each pulse at every position.

The SAR obtained from the centre position (SAR_{centre} ; i.e., the no-motion case, for which the pulse was designed) was used as a baseline. SAR evaluations at all other positions were subsequently normalised by this value to yield a relative SAR value per position, per pulse which quantified SAR changes due to motion. Values are therefore reported as a factor or percentage of SAR_{centre} . Although safety assessments often focus on peak local SAR, this considers SAR in only one spatial location. Another worthwhile consideration is the volume

of tissue experiencing an increase in SAR. For this reason, we also quantified the volume of tissue exposed to higher SAR than $\text{psSAR}_{\text{centre}}$ following motion (we refer to this as 'high-SAR tissue volume' here).

6.4 Results

Among all SAR metrics and pulse types considered, local SAR (psSAR) for RF-shimming was the most motion sensitive metric, especially among smaller coils. Figure 6.2 shows psSAR sensitivity for RF-shimming pulses designed and evaluated on slice 2. For the base coil (C), RF-shimming yielded a worst case psSAR increase of 2.7-fold due to motion. For the smallest coil (A), psSAR sensitivity was much higher - the maximum observed psSAR increase was 3.8-fold. On the other hand, for the largest coil (F) the worst case psSAR increase was much lower (1.3-fold). The local hotspot remained in approximately the same location for all coils' worst cases.

Across all coil models, worst case psSAR increases occurred using pulses designed and evaluated on the most inferior slices (slice 1 or 2). For RF-shimming, slices 2 to 5 showed similar patterns, while slices 1 and 6 did not, and are discussed later. The majority of worst case increases were seen following the largest displacement, but for 14% of pulses, worst case psSAR occurred at intermediate positions (i.e., not the "corner" positions). These cases were using larger coils and/or slices 5 and 6, where psSAR sensitivity was generally lower.

Figure 6.3 shows psSAR for coils' worst cases across all pTx (including RF-shim) pulses. As the number of spokes in pulses increased, SAR sensitivity generally reduced and became more similar across coil models (Figure 6.4), however the three largest coil models' worst cases were multi-spoke pulses. Conversely, coils A - C worst cases were all observed for single-spoke pulses.

PsSAR sensitivity and high-SAR tissue volume was lowest for the taller coil models, where psSAR increased by a maximum of 96% and 59% for coils D and F, respectively (see Figure 6.5). By comparison, coil E (wider, but equal in length to coil C) yielded a worst case psSAR increase of 2.5-fold, and SAR increases occurred over a larger tissue volume (176 cm^3 compared to 53 cm^3 and 73 cm^3 for coils D and F respectively). If psSAR is ignored,

RF-shimming, slice 2

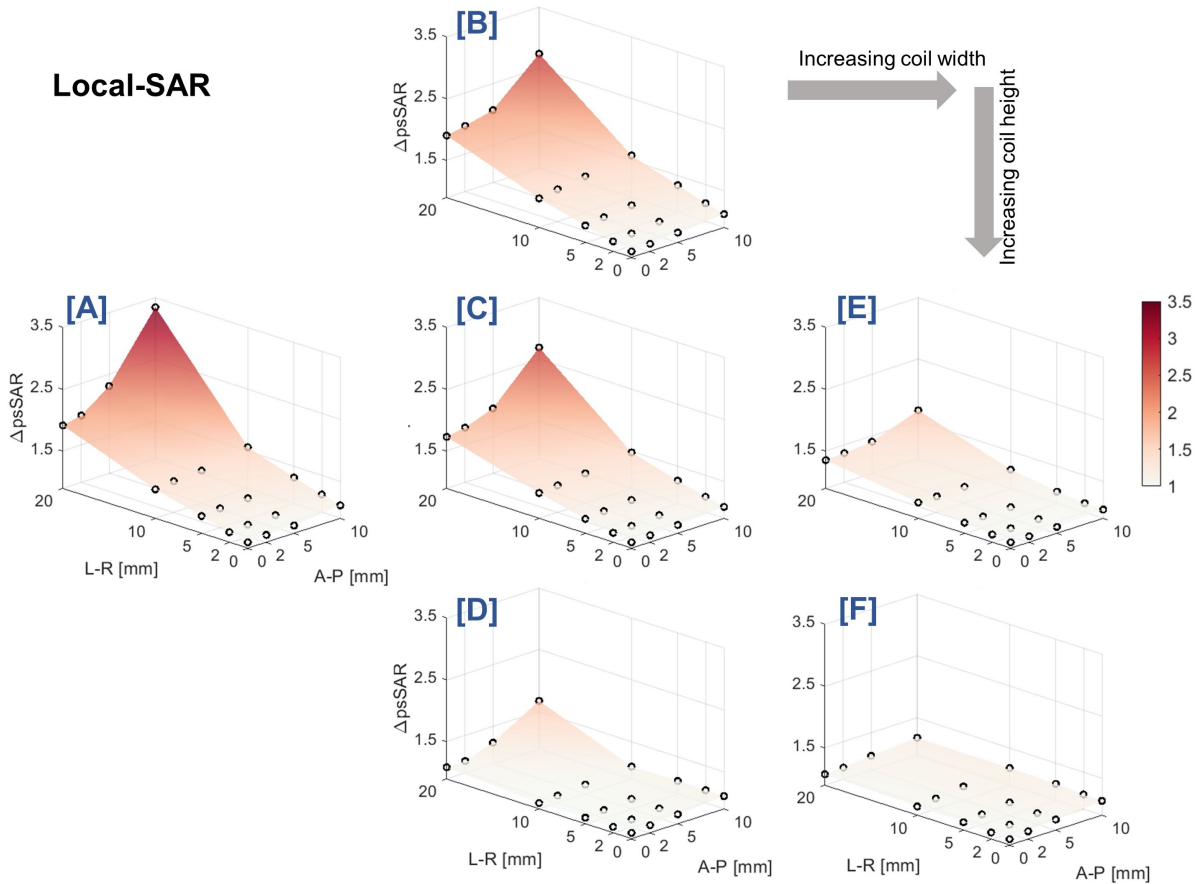


Figure 6.2: Motion-sensitivity of peak local SAR for RF-shim pulses designed using slice 2 at the central position for each coil model A to F. Slices 2 - 5 showed similar trends. Y-axes refer to SAR at each evaluated position (left-right [L-R] and anterior-posterior [A-P] displacements), as a factor of SAR at the central position (i.e. without motion). SAR_{centre} was comparable across coil models (data not shown).

the worst case in terms of high-SAR tissue volume was observed using coil A (306 cm^3) followed by coils B (257 cm^3) and C (194 cm^3). Note that Figure 6.5 does not use this definition, and instead shows high-SAR tissue volume for the worst case as defined by the evaluation which yielded the highest psSAR.

A different pattern emerged for quadrature pulses, which were generally less sensitive to motion than pTx pulses. For quadrature, psSAR was most sensitive using coil C where it almost doubled, while psSAR for coils A and E increased by 18% and 5%, respectively (Figure 6.6). High-SAR tissue volume exhibited a similar trend; 53 cm^3 and 49 cm^3 was exposed to increased SAR for coils C and B respectively, while $<3 \text{ cm}^3$ was for coils A and

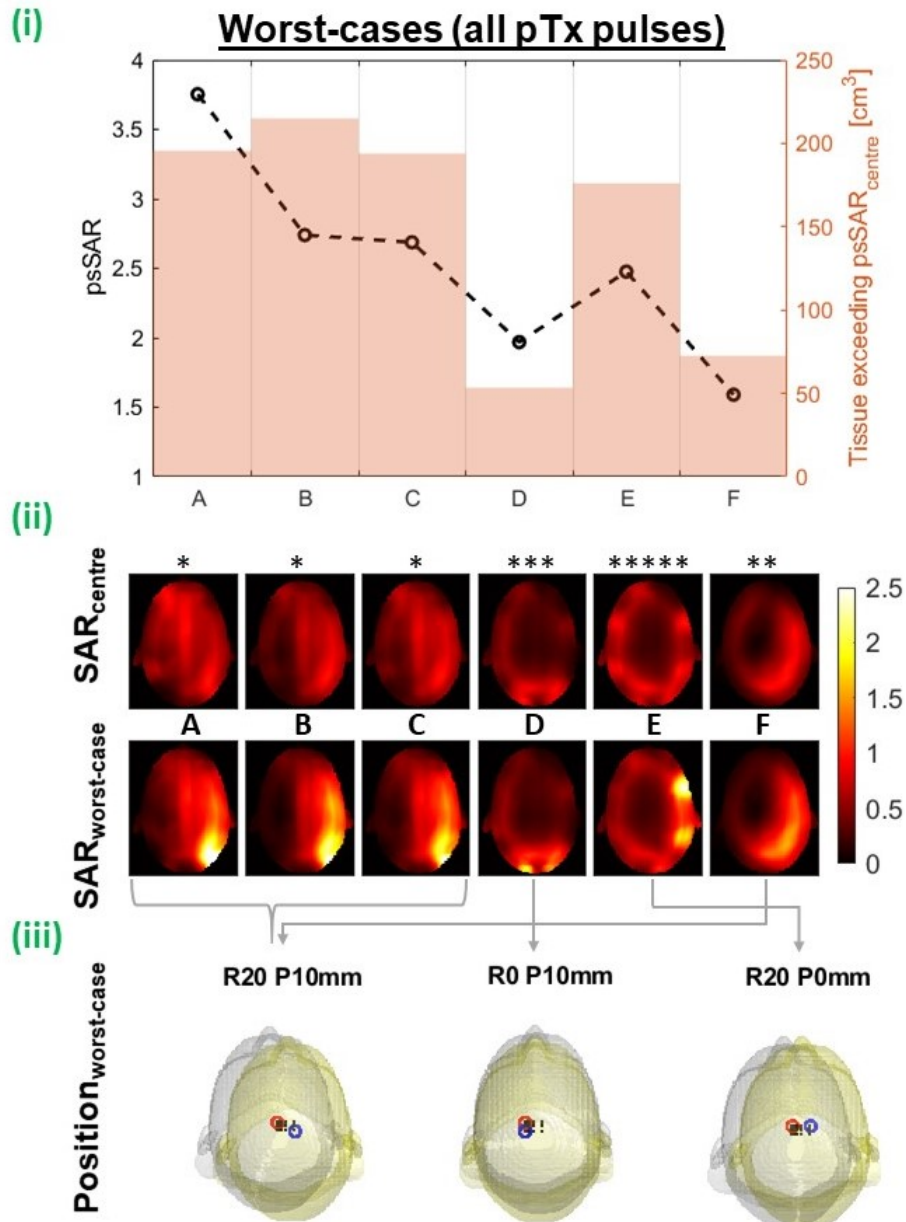


Figure 6.3: Worst-case local SAR following motion for pTx pulses. **(i):** psSAR for worst cases as a factor of $psSAR_{centre}$ (the same pulse, evaluated without motion). Tissue volume exposed to higher SAR than $psSAR_{centre}$ is shown as orange bars. **(ii):** SAR profiles for corresponding worst cases (bottom row) compared to their SAR_{centre} (top row). Colour bar shows psSAR as a factor of $psSAR_{centre}$ (normalised per coil). The number of spokes in each coil's worst case pulse is indicated with asterisks. **(iii):** The positions at which these worst case psSAR were observed. All worst cases were observed at slice 1 or 2.

F. Unlike pTx where the hotspot location was relatively stable, in all quadrature worst cases, the local-hotspot shifted from the front to back of the head.

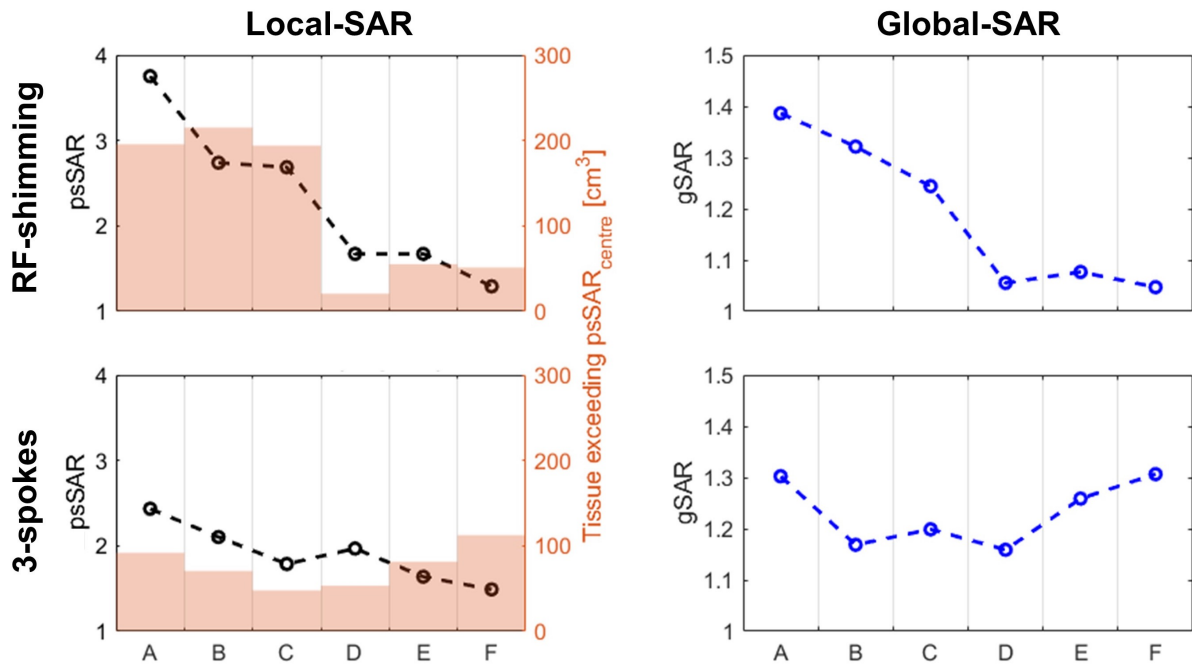


Figure 6.4: Relative local ($psSAR$; left) and global ($gSAR$; right) SAR for 1-spoke (top) and 3-spoke (bottom) pTx pulses for coil models A to F. The volume of tissue exposed to $>psSAR_{centre}$ is also shown. Worst-cases are defined by the largest increase in $psSAR$ or $gSAR$ for left and right figures respectively (i.e. the worst case $psSAR$ pulse was not necessarily the worst case $gSAR$ pulse).

Whole-head SAR ($gSAR$) was substantially less sensitive to motion than $psSAR$; a representative RF-shimming pulse is shown in Figure 6.7. Among pTx pulses, the maximum observed increase was 49%, which again occurred using the smallest coil (A). Unlike $psSAR$, coils' worst cases were relatively similar; coil C had the lowest worst case $gSAR$ increase, which was 24%. When only RF-shimming pulses were considered, relative $gSAR$ followed an almost monotonically decreasing trend with increasing coil size. However, coils B and C were less motion-sensitive than the largest coil models for multi-spoke pulses (Figure 6.4), resulting in a U-shaped relationship between worst case $gSAR$ and coil size when all pTx pulses were considered. For quadrature pulses, $gSAR$ was relatively stable, increasing by a maximum of 5% (coil A) and minimum of 2% (coil E).

In terms of $eigSAR$, motion caused a 43% increase for coil C. However, like the other metrics, $eigSAR$ was more motion sensitive among smaller coil models, increasing by up to 72% for coil A. $eigSAR$ was generally more sensitive to anterior-posterior than left-right

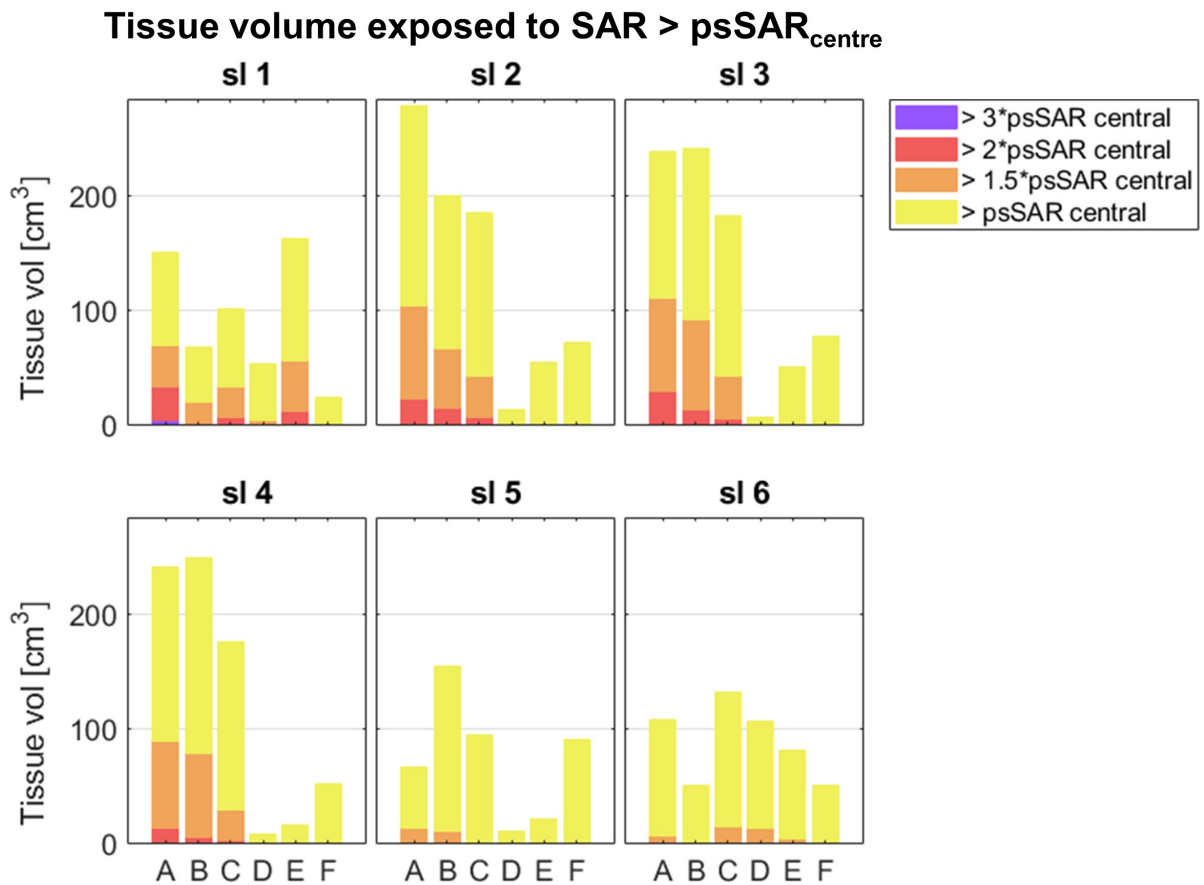


Figure 6.5: High-SAR tissue volume: the amount of tissue exposed to higher SAR than $psSAR_{centre}$ by >100% (all colours), >150% (orange + red + purple), >200% (red + purple) and >300% (purple) for worst case $psSAR$ pTx pulses separated by slice (slices 1 to 6 are shown in separate subplots).

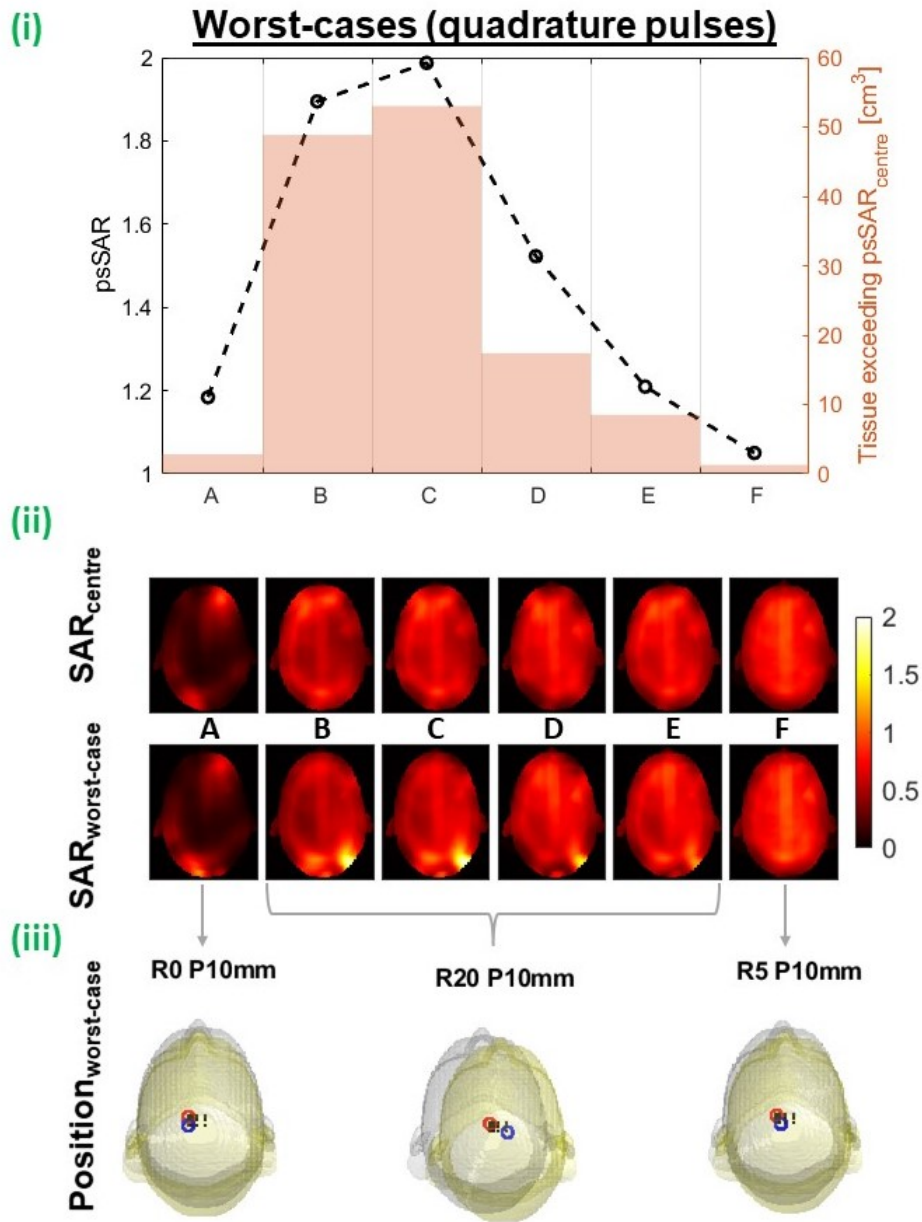


Figure 6.6: Worst-case local SAR following motion for quadrature pulses. **(i):** $psSAR$ for worst cases as a factor of $psSAR_{centre}$ (same pulse without motion). Tissue volume exposed to higher SAR than $psSAR_{centre}$ shown in orange. **(ii):** SAR profiles for corresponding worst cases (bottom row) compared to their SAR_{centre} (top row). Colour bar shows $psSAR$ as a factor of $psSAR_{centre}$ (normalised per coil). **(iii):** The positions at which these worst case $psSAR$ were observed. Relative $psSAR$ across positions for quadrature pulses is independent of slice location.

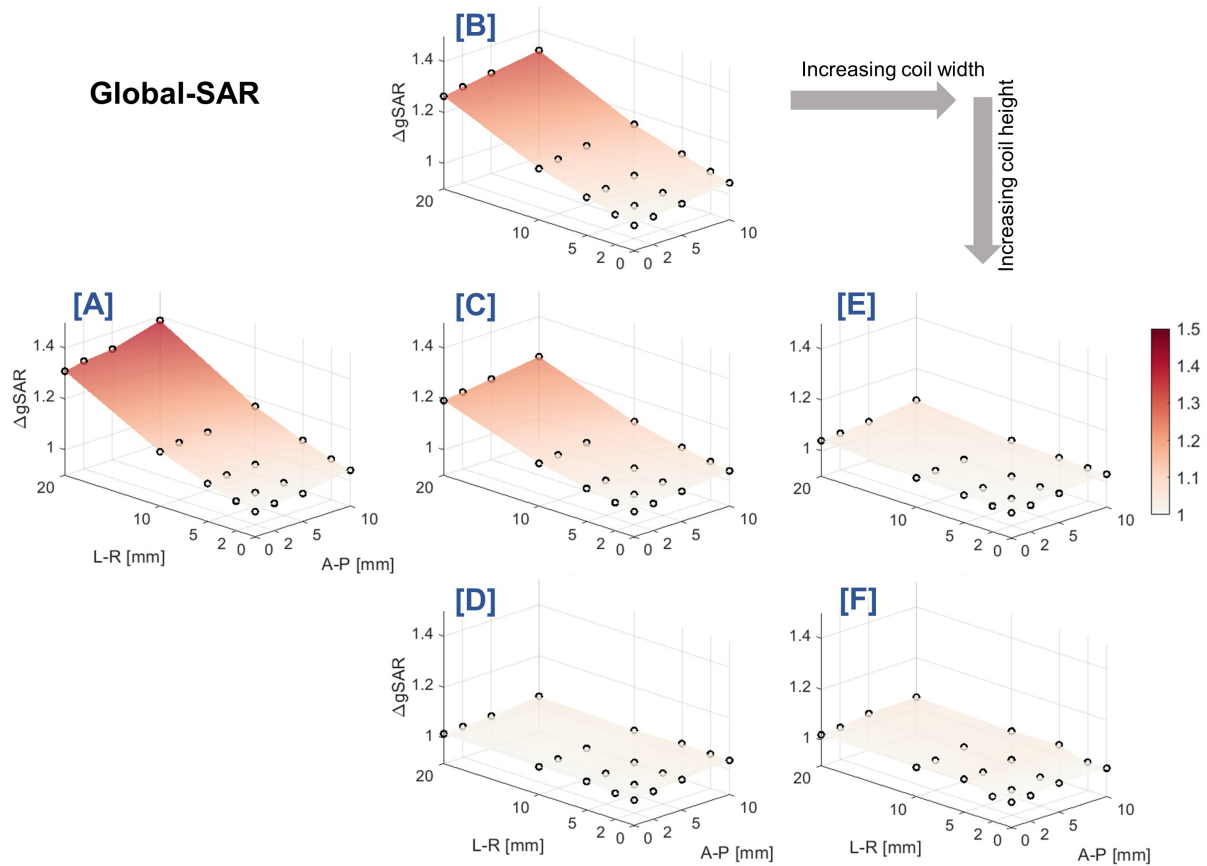


Figure 6.7: Motion-sensitivity of whole-head SAR for RF-shim pulses designed using slice 2 at the central position for each coil model A to F. Slices 2 - 5 showed similar trends. Y-axes refer to SAR at each evaluated position (left-right [L-R] and anterior-posterior [A-P] displacements), as a factor of SAR at the central position (i.e. without motion). SAR_{centre} was comparable across coil models (data not shown).

displacements. The lowest eigSAR was observed at the $x=-5$, $y=2$ mm position for all coils (except coil E, where it was $x=-5$, $y=5$ mm). Coil E yielded the lowest worst case eigSAR increase (6%). The largest coil had the second-highest worst case eigSAR increase (46%), however $eigSAR_{centre}$ was much lower (around 25% of the next closest $eigSAR_{centre}$).

6.5 Discussion

In this study we investigated the relationship between SAR sensitivity to motion and pTx channel dimensions. Coil C has similar dimensions to that in Kopanoglu et al. [130] and we found similar SAR sensitivity to motion as expected. For RF-shimming, we observed 2.7-fold

increase in peak local SAR (psSAR) and 24% in whole-head SAR (gSAR), while Kopanoglu et al. reported 2.4-fold and 33%, respectively. Eigenvalue-based SAR (eigSAR) for coil C (43%) was also in accordance with the 42% reported previously [130]. We found that, for pTx pulses, pTx coils with narrower and/or shorter loops were substantially more motion sensitive (psSAR = 3.8-fold for the smallest coil's worst case) than those with wider, taller loops and/or a larger array radius (psSAR = 1.3-fold for the largest coil's worst case). For local SAR, increasing the coil's height yielded marginally lower motion-sensitivity compared to increasing its width (although both options reduced motion-sensitivity compared to the base coil). The overall worst case pTx pulses for gSAR and psSAR increase were not the same (slice 70 vs slice 60), however they were both seen using coil A (the smallest coil).

When quadrature pulses were used, non-monotonic SAR behaviour was observed (with respect to loop height and width), with the mid-sized coil C exhibiting the most motion-sensitivity. We attribute this to the tendency of local hotspots to change location following motion. Even without motion, coil A experienced small, intense local-hotspots, meaning change in psSAR was low when the hotspot shifted to a different location, but remained at a similar intensity (due to the high baseline). For large coils, fields were smoother throughout the head, leading to less intense hotspots, even after motion. Coil C did not exhibit either behaviour, hence the largest relative increase. Although, in the small-coil case, the change in psSAR was low, the change in hotspot location represents an alternative safety concern, especially if SAR monitoring based on one psSAR location is used, or if pulses have been calibrated to avoid heating in certain areas (e.g. near implants) [141, 142]. Similarly, techniques which intentionally change the SAR hotspot location over the course of a pulse sequence (to avoid heat build-up in one hotspot location) may be compromised following unanticipated interactions with motion-related changes to the hotspot location [197, 198].

For slices 2-5, baseline local and whole-head SAR (SAR_{centre} , data not shown) were comparable across coil models, whereas motion-sensitivity was systematically higher among smaller coils for these slices. $psSAR_{\text{centre}}$ and $gSAR_{\text{centre}}$ were higher when pulses were designed using slices 1 and 6. Interestingly, for coil B, slices 1 and 6 fell beyond the loops' vertical extent, and motion-sensitivity was notably lower than other slices (though $psSAR_{\text{centre}}$ was around threefold). Similarly, SAR was less motion sensitive in slice 6 (most

superior) for all coils. Although it could be argued that the low sensitivity of these slices was at least partly due to the higher SAR_{centre} used to calculate SAR change, slice 1 (which also had a high SAR_{centre}) yielded among the largest SAR increases when motion occurred. We attribute this to the highly non-uniform fields seen in inferior slices near the cerebellum, which are likely to experience more field changes for a given displacement compared to more homogeneous slice locations. Nonetheless, high SAR_{centre} is less concerning than motion-sensitivity of SAR, as SAR_{centre} can be constrained during pulse design.

Deniz et al. showed that SAR was generally lower when the coil-to-sample distance was large [127]. In this scenario, B_1^+ and E-fields at all spatial locations receive contributions from all coil elements since channel sensitivities are smoother and contain more overlapping regions. In contrast, voxels near the coils only receive contributions from the nearest coil when the coil-to-sample distance is very small. In the latter case, pTx pulse design cannot leverage degrees of freedom in areas close to the coils, since only the most proximal channel is available to deliver RF. Deniz et al. hypothesised that the extra degrees of freedom in pulse design were utilised for SAR reduction [127]. Although we did not constrain SAR here, a similar mechanism may explain the lower SAR sensitivity to motion seen in the most inferior and superior slice locations. These slices are physically smaller than central slices (and therefore further from the coil, where fields from different channels overlap more). Similarly, when the array radius is larger (coil F), the size of the head is smaller relative to the array. The overlapping fields in these scenarios therefore allow a pTx pulse to be designed which is more symmetric, in terms of channel contributions. In previous work, it was found that highly asymmetric pulses designed to excite off-central target ROIs were substantially more susceptible to motion than symmetric pulses designed to excite central ROIs [199]. Similarly, we generally observed less motion-sensitivity among symmetric (quadrature) pulses here. This supports the notion that more symmetry among channel weightings (which is made possible when fields are smooth and overlapping - more likely when the array radius is large) may decrease sensitivity to motion.

EigSAR was lowest at a similar off-centre position across coil models ($x, y = -5, 2$ position for coils A to D and coil F, and $x, y = -5, 5$ for coil E), showing that interference patterns were similar, and the centre position corresponded well across coil models. However, for larger shifts (especially in y), eigSAR for the smallest coil was more motion sensitive than

the larger coils. While the largest coil model (F) unexpectedly yielded the second-highest worst case eigSAR, $\text{eigSAR}_{\text{centre}}$ was much lower. Generally, eigSAR was more sensitive to anterior-posterior than left-right displacements, with this discrepancy being greater for smaller coils. We believe this is due to the elliptical shape of the head compared to the circular array (as an A-P translation results in a closer proximity to coil elements than the equivalent L-R translation). EigSAR is independent of pulse design, so the higher motion-sensitivity observed among smaller coils reflects field interference patterns intrinsic to the coil and subject setup.

One limitation of the study here is the inability to disentangle the reasons for apparent low SAR sensitivity when the baseline SAR which occurs without motion is high. Nevertheless, it is likely to be what occurs in practice, since we used realistic pulses with the same target excitation pattern for evaluations at all positions.

Deniz et al. found that the relationship between coil-to-sample distance and *RF power* did not follow that between coil-to-sample distance and *global SAR* [127] (which was attributed to higher power dissipation in the larger / more numerous coil elements when the coil-to-sample distance was large). In this study, we penalized RF power via Tikhonov regularisation in the pulse design. It is possible that a different pattern may have emerged if we had instead constrained global SAR.

This investigation focused on altering the loop channel dimensions, while SAR behaviour for diverse pTx coil designs (e.g., stripline designs) may be very different. The optimal pTx coil design also depends on the imaged region. Wu et al. compared several coil designs including single- and multi-rung stripline arrays and a 2-port birdcage body coil [134]. The 2-rung stripline array was found to yield the most favourable flip-angle and SAR performance for head and spine imaging (where it reduced local- and global SAR by approximately 8-fold compared to quadrature mode), but differences between coils were less well-defined when imaging over a larger region covering the pelvis. The patterns in SAR behaviour observed here may also differ for different anatomical regions.

6.5.1 Conclusion

While our results agree with those previously reported for the similar-sized coil, we generally observed lower SAR motion-sensitivity in larger (wider or taller loops, and/or larger array radius) coil models, and much higher sensitivity for the smallest coils. Maximum observed local SAR increases due to motion were 3.8-fold and 1.6-fold for the smallest and largest coil models respectively. This seems to be somewhat independent of pulse design, since eigenvalue-based SAR (which depends only on the transmit *fields*) followed this pattern, however pTx pulses with small coils elicit the most concerning local SAR increases.

Chapter 7

Specific Absorption Rate and inter-subject variability

The work presented in this chapter led to the following conference proceedings:

Plumley, A. et al. Inter-subject differences in SAR sensitivity to motion with parallel-transmit at 7T. Proceedings of the International Society for Magnetic Resonance in Medicine, #0457 (2022).

7.1 Introduction

In the previous chapter, the sensitivity of SAR to motion across different coil models was explored. It was found that pTx coils with different sized channel elements exhibit different SAR behaviour when motion occurs. In addition to the coil itself, E-field and SAR distributions also depend critically on the position, orientation, size, shape, and composition of the coil's load due to dielectric effects. RF propagation in a tissue depends on the tissue's electrical properties (conductivity and permittivity) [200]. Conductivity in brain tissue varies by almost 5-fold (CSF: ~ 2 , white matter: ~ 0.4 , grey matter: ~ 0.7 Siemens/meter), and relative permittivity by around 1.6-fold (CSF: ~ 73 , white matter: ~ 44 , grey matter: ~ 60) [144]. Air inside the sinus and ear canal cavities has zero conductivity and a relative permittivity of 1. As well as propagation directly from the RF sources, reflection and

refraction at boundaries between different materials also impact field distributions, especially where the electrical properties at the boundary are very different (for example, air-tissue interfaces). These factors therefore determine the load-dependent RF energy deposition pattern within the imaged object, which in turn yields a SAR distribution. In MRI, these factors reduce to inter-subject variability of anatomy and positioning within the scanner.

Unsurprisingly, SAR therefore varies between subjects, even without motion [201–204]. Using simplified torso models consisting of fat, skin and muscle, Ipek et al. found peak local (10g-averaged) SAR variation of more than two-fold among four subjects with realistic RF-shim pulses designed for uniform B_1^+ in the prostate [201]. Peak local SAR was 2.8 W/kg vs. 5.7 W/kg for the least sensitive and most sensitive models, respectively. Using a similar approach but including an extra tissue type (cortical bone) in models of 23 subjects, Meliado et al. found peak local SAR values ranging from 2.6 to 4.6 W/kg [202]. For head imaging, De Greef et al. found local SAR varied by almost two-fold across 6 subject models with RF-shimming pulses (10.3 vs. 19.2 W/kg) [203]. Peak local SAR was lower for quadrature-mode excitation, however still varied by a factor of approximately two across subjects (4.6 vs. 8.1 W/kg). Similarly, whole-head SAR varied between 2 and 4 W/kg. Despite this substantial variation, standard vendor-supplied SAR estimations are typically calibrated using generic models which can be unrepresentative of all populations. To account for the uncertainty and variability, large conservative safety margins are applied, which limit imaging performance.

Although this is currently the most common approach, alternative SAR management methods are available. Electrical properties needed for SAR calculation can be obtained from B_1^+ maps using neural networks [205]. However, to calculate SAR distributions, EM simulations must typically be conducted on these models, which can take on the order of hours and so are infeasible for clinical pipelines. Virtual observation points (VOPs) provide a compressed representation of local SAR, allowing rapid calculations which can be used for SAR-constrained pulse design. Another approach aims to reduce the overestimation of the generic safety margins by using probabilistic SAR estimates based on a subject-specific observation in relation to observations obtained over a 23-subject database [206]. Alternatively, deep neural networks can be used to estimate SAR distributions directly from a subject's measured B_1^+ maps [132]. A comprehensive approach was also recently proposed

where patient-specific body models can be generated rapidly from structural scans using deep neural networks. Subject-specific SAR calculations can be conducted on the models for pTx in less than 45 seconds using a custom accelerated EM solver [133]. These emerging methods effectively address the inter-subject variability of SAR, but they do not account for patient motion occurring throughout the scan. The same issue applies to Universal Pulses (described in Chapter 3) which can incorporate SAR constraints in their design [123, 125].

As described in Chapters 4 and 6, motion causes substantial changes to E-field and SAR distributions at 7T, but the limited works on this topic have each used only a single body model [130, 190, 191]. Therefore, it is unknown whether the sensitivity of SAR to motion depends substantially on the subject or population being imaged (as SAR distributions themselves do), or whether these findings are generalisable. Here, we investigate whether (and to what degree) SAR sensitivity to motion depends on the subject and their anatomy by comparing the sensitivity of SAR to motion across six virtual human body models.

7.2 Methods

Two child models (Billie and Dizzy) and four adult models (Duke, Ella, Fats and Glenn) of the Virtual Population [143] were simulated using the same simulation approach as described in Section 5.2.1. The coil's loops were 110 mm (height) × 40 mm (width), and the array radius was 115 mm. Each body model was simulated at the same 30 positions on the axial plane as shown in Figure 6.1iii. While relative positions were identical across models, the centre positions (R0, P0) across models were aligned by eye with the centre of the model occupying roughly the centre of the coil. The Duke, Fats and Glenn models were scaled to 90% of their original sizes, to avoid overlapping the coil and body models.

Like in Chapter 6, quadrature-mode and pTx (1, 2, 3 and 5-spokes) excitation pulses were designed for uniform in-plane excitation in an axial slice using a spatial domain small tip-angle method [32, 75]. Pulses were designed using the B_1^+ maps from the centred position for each body model. Separate pTx pulses were designed for 7 slice locations spanning a 10.8 cm axial region (figure 6.1). One quadrature pulse per model was also designed, yielding 29 pulses in total per model.




Figure 7.1: *The six Virtual Population body models used, each shown in their respective centre position. Models were downloaded from IT'IS (Zurich MedTech, Zurich) [143]*

SAR was evaluated using the same approach as that described in Chapter 6. Briefly, whole-head (gSAR), 10-g average (psSAR) and eigenvalue-based SAR (eigSAR) were calculated using Q-matrices [129]. In all cases, SAR is reported relative to that obtained by the same pulse without motion (SAR_{centre}). The volume of tissue experiencing higher SAR than SAR_{centre} was also investigated. Our analysis focussed on similarities or differences in sensitivity of these metrics to motion across the 6 body models.

7.3 Results

7.3.1 Quadrature mode

In quadrature mode, peak local SAR (psSAR) more-than doubled for the Ella model, however sensitivity varied substantially across models, as shown in Figure 7.2 (left panel). Worst-cases for Glenn, Fats, Dizzy, Billie and Duke models were 1.9, 1.8, 1.6, 1.3 and 1.2, respectively. All models' worst cases occurred following the largest displacement (at the R20, P10 position). Correspondingly, the peak SAR hotspot changed to a location in the right posterior part of the head for all models (data not shown).

For Billie and Duke models' worst cases, 13 - 14 cm³ was exposed to higher SAR than the peak SAR value without motion. For Dizzy, Fats and Glenn, the volume of high-SAR tissue was 32 - 37 cm³, while for Ella it was higher (53 cm³, data not shown).

Figure 7.3 (left panel) shows relative gSAR for quadrature pulses, which was fairly insensitive to motion, yielding worst-case increases between 1% (Billie) and 5% (Fats). For quadrature pulses, a higher SAR_{centre} (not shown) coincided with a lower SAR sensitivity to motion.

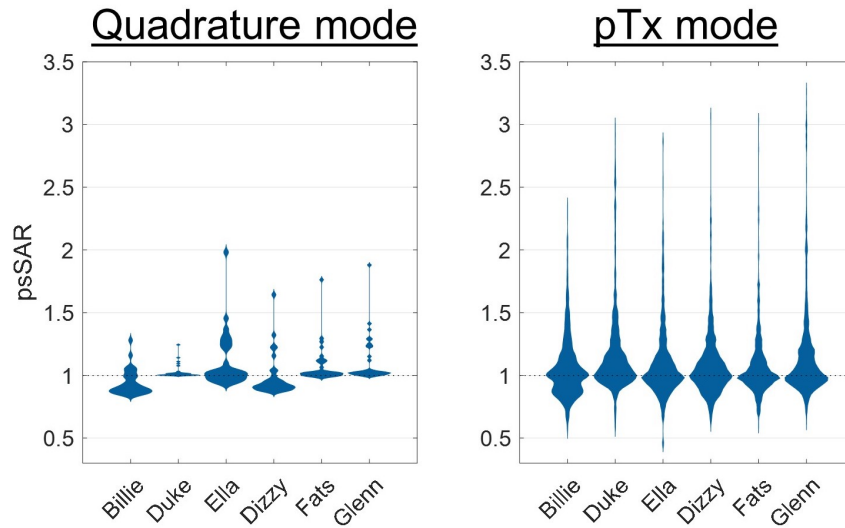


Figure 7.2: Relative peak local SAR ($psSAR$) when using quadrature (left column) and pTx (right column) modes. Relative SAR for quadrature mode is independent of slice and number of spokes, so only one quadrature pulse was designed per model, meaning quadrature violin plots contain just 30 evaluations per model (one per position). pTx plots contain all pTx pulse evaluations (840 in total) per body model. In each case, SAR is shown as a factor of that produced by the same pulse without motion.

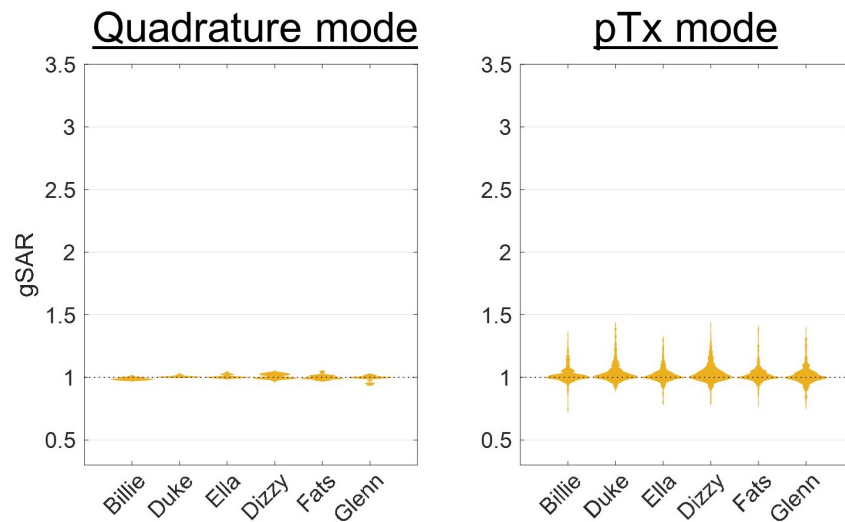


Figure 7.3: Relative whole-head ($gSAR$) when using quadrature (left column) and pTx (right column) modes. Relative SAR for quadrature mode is independent of slice and number of spokes, so only one quadrature pulse was designed per model, meaning quadrature violin plots contain just 30 evaluations per model (one per position). pTx plots contain all pTx pulse evaluations (840 in total) per body model. In each case, SAR is shown as a factor of that of that produced by the same pulse without motion.

7.3.2 pTx mode

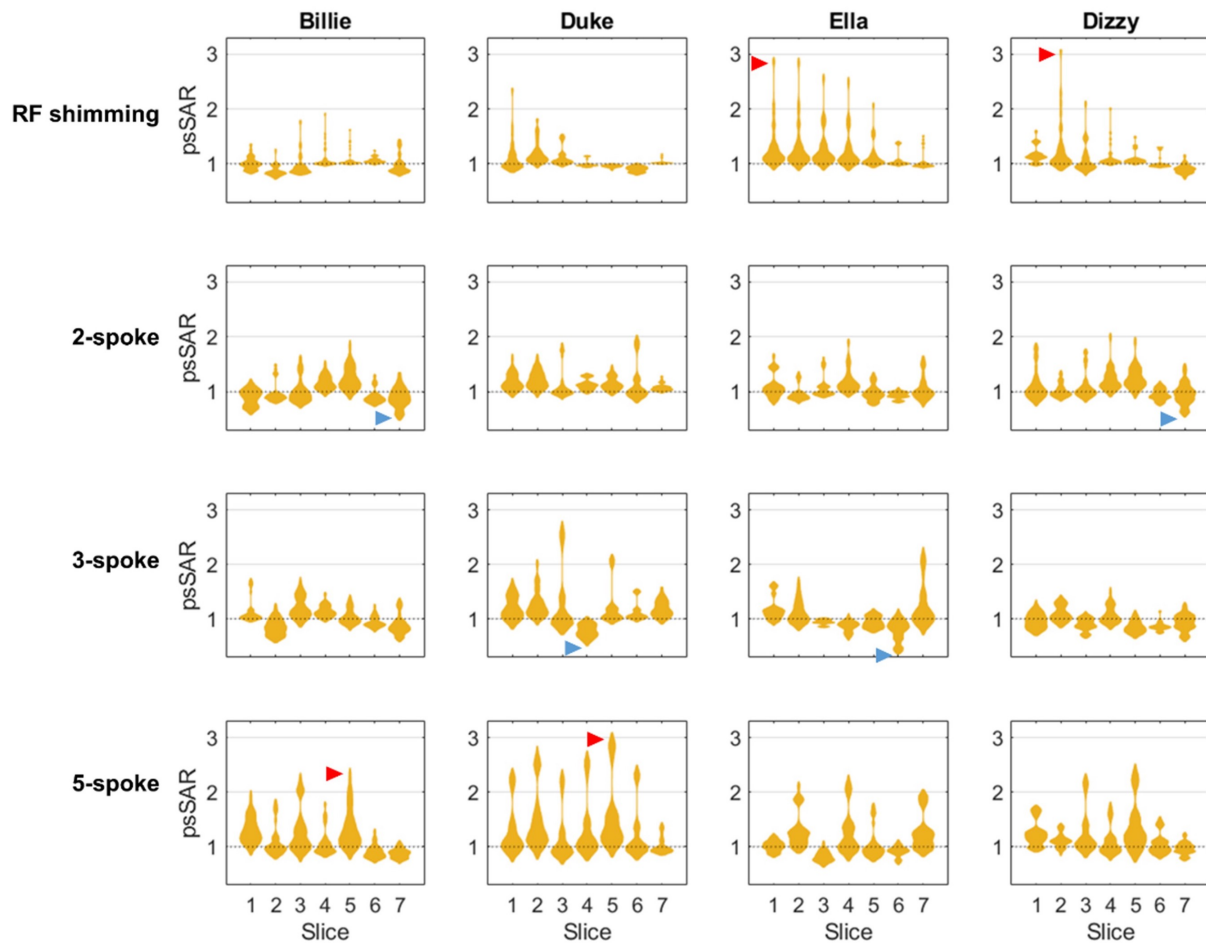


Figure 7.4: Relative local SAR ($psSAR$) sensitivity of pTx pulses to motion, split by slice location and number of spokes in the pTx pulse (RF-shimming = 1 spoke). Each independent violin plot contains the 30 evaluations (i.e., all positions) for that pulse. SAR is shown as a factor of that of the same pulse without motion. Each body model's lowest and highest observed relative $psSAR$ are indicated with blue and red arrows, respectively. The Fats and Glenn models are omitted for clarity, but they showed similar variability to the models shown.

In pTx mode, $psSAR$ was highly sensitive to motion across all body models. Resulting $psSAR$ was highly variable across pulses and positions, as seen by the large distributions in Figure 7.2 (right panel). Having said that, $psSAR$ increased by at least 2.4-fold in all models' worst cases. Figure 7.4 shows SAR distributions for all evaluations, separated by pulse (only 4 body models are shown for clarity). Interestingly, worst- (and best-) cases (indicated with red and blue arrows, respectively) did not occur for the same slice or number of spokes across models. Worst cases for Billie (2.4-fold) and Duke (3-fold) models arose

with 5-spokes pulses at slice 5. For Ella (2.9-fold) and Dizzy (3.1-fold), worst cases were with 1-spoke pulses at inferior slice locations. Worst cases in the Fats (3.1-fold) and Glenn (3.3-fold) models were obtained using 3-spokes pulses at mid-superior slice locations (Fats: slice 4; Glenn; slice 6).




Figure 7.5: local SAR profiles for each model's worst-case SAR increase. SAR_{centre} shows SAR using the same pulse without motion; models' worst-case profiles were normalised to psSAR from the respective motion-free profile. The crosses indicate the maximum psSAR hotspot. Slice, number of spokes, and position at which worst-cases occurred are indicated below profiles (R=rightward, P=posterior).

Figure 7.5 shows local SAR profiles for each models' worst case. Except for Duke and Glenn, whose worst cases occurred at the (R20, P0) position, all models' worst cases occurred following the largest displacement (R20, P10). The psSAR hotspot substantially changed location in all models' worst-cases, except Ella and Duke. Unlike quadrature mode, the location of the hotspot following displacement varied across models. For the models whose worst cases were at the (R20, P10) position, the hotspot moved to the right posterior part of the head, while for Duke and Glenn, it was at the right anterior part.

Like for quadrature pulses, gSAR was less sensitive than psSAR, and worst-cases were between 1.3 and 1.4-fold (shown in Figure 7.3, right panel). Like psSAR, worst cases arose for different pulses and positions across models. Duke, Billie and Dizzy yielded worst-case

gSAR with 5-spokes pulses at mid-axial slices and at the (R20, P10) position, while 2-spokes in the most superior slice and (R20, P0) position yielded highest gSAR for Ella. For Glenn and Fats, gSAR suffered most with the 3-spokes pulse in mid-superior slices at the (R20, P0) position.

Figure 7.6 shows the tissue volumes exposed to higher SAR than $psSAR_{centre}$ for pTx pulses, separated by body model and slice location. Compared to quadrature pulses, pTx pulses also yielded larger high-SAR tissue volumes. Considering models' respective worst cases, Duke yielded the overall largest high-SAR tissue volume (295 cm^3), but at least 177 cm^3 of tissue was exposed to increased SAR for all models' worst-cases. All models also experienced substantial volumes ($10 - 40\text{ cm}^3$) which more-than doubled $psSAR_{centre}$ (except Billie, where up to 6.4 cm^3 of tissue experienced double $psSAR_{centre}$). High-SAR tissue volumes depended more on slice location than body model, with pulses designed for mid-axial slices yielding the largest high-SAR tissue volumes following motion.

Similarly, in terms of both gSAR and psSAR, pTx pulses designed for the inferior-most and superior-most slices consistently yielded higher SAR_{centre} (data not shown) but were generally less sensitive to motion compared to mid-axial slices, whose SAR_{centre} tended to be lower, but relative SAR following motion was highly variable. These intra-model differences dominated over inter-model differences.




Figure 7.6: Tissue volume exposed to higher SAR than psSARcentre by >100% (all colours), >150% (orange + red) and >200% (red) for worst-case pTx pulses separated by slice.

7.3.3 Eigenvalue-based worst-case peak local SAR




Figure 7.7: Relative eigenvalue-based SAR (eigSAR) sensitivity to motion distribution across all 30 positions for each body model. EigSAR was calculated as the maximum eigenvalue of the 10g-averaged Q -matrices and is therefore pulse-independent.

Figure 7.7 shows that the range in eigSAR evaluations was approximately 0.9 to 1.4 for all body models except for Glenn, where worst-case was 1.8-fold at the (R20, P10) position, and Duke, where eigSAR more-than doubled at the (R20, P0) position. Billie, Dizzy and Fats models yielded worst-case eigSAR at the (R0, P10) position.

7.4 Discussion

In light of previous work reporting the effects of patient motion on SAR at 7T [130, 190, 191], this study compared the SAR response to head motion across six different subject models to investigate whether the previous findings are likely to be generalisable or population-specific. Our worst-case psSAR (2.9-fold) and eigSAR (43%) results for the Ella model are in line with previous work (3.1-fold and 42% respectively in Kopanoglu et al. [130]) although our worst-case gSAR increase (32%) was lower than the 60% reported in Kopanoglu et al. [130]. We attribute the discrepancy in gSAR to differences in pulse design (e.g., different Tikhonov regularisation).

As expected from the literature, we found psSAR to be more motion-sensitive than gSAR. Worst-cases for psSAR ranged between 2.4 and 3.3-fold that without motion across the 6 body models we tested here. Local SAR sensitivity varied across body models in quadrature mode, but pTx mode yielded the most SAR variability and the largest SAR increase for all models, in line with previous findings [130, 203]. Similarities in the low gSAR sensitivity across body models may mask these important psSAR differences, especially in quadrature mode, where gSAR did not exceed 5% of that without motion, but psSAR more-than doubled in some cases. Having said that, gSAR was also moderately sensitive to motion with pTx pulses (increasing by 30-40%). Also, for every model, the pulse yielding the worst-case psSAR increase was different to the pulse yielding worst-case gSAR increase, highlighting the need to monitor both metrics. The sensitivity of eigSAR (which is pulse-independent and depends only on the fields) varied somewhat across models, but was also much less sensitive than psSAR, supporting the use of pulse-specific SAR assessments over eigenvalue-based approaches.

De Greef et al. found that, based on eigenvalue analysis of SAR, the RF-shim sets yielding each subject's worst cases were similar (i.e., their worst cases arose from a similar pulse) [203]. In contrast, we found that, for realistic pulses, each pulse of the same type (same slice, same number of spokes) responded very differently across the different subjects. Correspondingly, worst-case increases across the different models were elicited with very different pulses. The pulses we tested here were realistic pulses designed for uniform excitation, while De Greef et al. obtained the worst-case single-spoke shim set from all possible shim sets using the eigenvalue-based approach, which may explain the difference in conclusions drawn (as our pulse database was much more constrained). Our results suggest that a single "safe" pulse which yields good flip-angle performance and low SAR sensitivity to motion across all subjects is unlikely to exist.

Although the hotspot location following motion differed across models for pTx pulses, its location was congruent with the worst-case position observed for each model. Where the worst case was observed at the most rightward and posterior position, the hotspot appeared at this location in the head, while it was further anterior in the models whose worst cases were in more anterior positions, presumably due to close tissue-channel proximity. Changes to the hotspot location (which occurred in some, but not all body models) can be

considered an advantage, in the sense that the RF power absorption appears to be better distributed in space and time (avoiding the build-up of intense local heat over time), however there is no guarantee that the previous hotspot location experiences reduced RF power deposition when a new, more intense hotspot forms (i.e., the old hotspot location may continue to build heat too). Furthermore, if VOPs or similar local SAR monitoring approaches are calibrated based on SAR distributions from one position, SAR changes in new, unanticipated hotspot locations may go unnoticed.

For the majority of our evaluations, SAR distributions did not change substantially compared to the no-motion case. However, there were large tails in the distributions where local SAR (and whole-head SAR in the case of pTx pulses) increased substantially, indicating that SAR management based on probabilistic methods (e.g., Ref. [206]) may unintentionally permit very high SAR values in some cases where motion occurs.

We applied the same displacements to all body models and considered each displacement to be equally likely, however the amount of motion is likely to be population-dependent (e.g., patients with tremors may be more prone to motion, and children may be more prone to large movements). Therefore, SAR variability may become population-dependent when the amount of motion that is likely to occur in practice is also considered.

Differences in sensitivity between models may also be influenced by slight differences in the reference position (R_0 , P_0) across models. For example, Duke's nose was very close to the anterior-most coil for some positions, meaning SAR response to a given displacement may be different to other, differently placed models. This highlights the distinct but related issue of initial patient positioning in SAR evaluations [207]. Similarly, the findings from this study should be considered in parallel to the findings reported in Chapter 6 in the sense that using different coil models may affect the inter-subject variation observed.

Overall, this study demonstrates that motion-sensitivity of SAR (especially local SAR) at 7T is somewhat variable across subjects, with models' worst-cases occurring with different pulse types (different number of spokes), and at different slice locations. In agreement with Ipek et al. [201], our results therefore indicate that a single model is insufficient for local SAR estimation due to the variability in SAR distributions across different subjects, and the difference in SAR response of each model to each pulse type. Despite the variation

observed, all models experienced an over-doubling of local SAR due to motion in their respective worst-cases. Additionally, although there were differences, they were not clearly systematic (e.g., between child and adult models), demonstrating that as a safety concern, SAR sensitivity to motion at 7T should be considered generalisable, and is not likely to be limited to certain populations. Even in quadrature mode, which is generally considered safe, local SAR doubled or almost doubled in half of the models tested here, while whole-head SAR for quadrature pulses severely underestimated this change, increasing by a maximum of 5%. Therefore, tailored approaches with real-time, subject-specific local SAR measurements are needed to guarantee safe and efficient performance with pTx.

Chapter 8

Approaches to reduce motion-sensitivity in pTx: Composite B_1^+ maps

The work presented in this chapter led to the following conference proceedings:

Plumley, A. et al. Large tip-angle, motion robust pulse design for parallel transmission at 7T using average B1 distributions. Proceedings of the International Society for Magnetic Resonance in Medicine, #3700 (2020).

8.1 Introduction

For optimal tailored pulse performance in pTx, the B_1^+ distribution present when the pulse plays out should match that which was used in the pulse design process. This is not the case when patient motion causes B_1^+ field changes during the scan, as explained in previous chapters. It is therefore desirable to reduce the positional dependence of tailored pulses.

In Chapter 5, it was shown that there are differences in motion-sensitivity between small tip-angle (STA) and large tip-angle (LTA) parallel-transmit (pTx) pulses. The error profiles observed were also qualitatively different, being spatially smoother for STA pulses compared

to LTA pulses. To fully exploit the versatility of MRI, a range of RF pulse types and sequences are needed, often even within the same protocol. Therefore, motion correction should be approached in a flexible way such that it can be applied to arbitrary sequences including small and large flip-angle pulses.

Here, we introduce motion-robustness to pTx pulses by designing them using composite (averaged) B_1^+ maps, which cover a larger area of the field of view (FOV) than the head. Average maps are calculated by weighted-averaging B_1^+ maps derived from the head in multiple positions. Hence, they do not reflect the precise B_1^+ distribution of the head in any one position. Rather, they include and combine (through voxel-wise averaging) information from several head positions. When these average maps are used in pTx pulse design, the resulting pulses are therefore less tailored to head position compared to standard tailored pulses (which typically use B_1^+ maps measured at a single position only).

Versatility of the approach is demonstrated through simulations using non-selective 3-D STA pulses and spatially-selective 2-D LTA pulses. To assess performance, pulses calculated using the averaged B_1^+ distributions ($\text{pulse}_{\text{average}}$) are evaluated at different head positions. The resulting flip-angle profiles are compared with those produced by conventional tailored pulses which are designed using the B_1^+ maps measured at a single position only ($\text{pulse}_{\text{centre}}$).

8.2 Methods

8.2.1 B_1^+ maps and simulations

The same simulated data as that used in Chapter 5 was used for this study. Simulation details were reported in Section 5.2.1. To recap, B_1^+ distributions from a single virtual body model were generated at 104 different head positions inside an 8-channel pTx array at 7T, covering all 3 translation (left-right, anterior-posterior and inferior-superior) and all 3 rotation (pitch, roll and yaw) directions.

8.2.2 B_1^+ averaging

Averaged B_1^+ distributions were defined as:

$$B_1^+ \text{average}(c, r) = \frac{\sum_{i=1}^{Np(r)} w_i B_1^+(i, c, r)}{\sum_{i=1}^{Np(r)} w_i}$$

where $Np(r)$ is the number of simulated positions with non-zero B_1^+ in voxel r , $B_1^+(c, r)$ is the c 'th channel's complex B_1^+ sensitivity at voxel r , and w_i is a position-specific weighting factor used to optimally weight contributions from different head positions.

For most patients, small movements are generally assumed to be more likely to occur than large movements. Additionally, it is anticipated that averaging will trade some accuracy in pulse performance for generalisability across positions. For these reasons, we establish two different weighting regimes which prioritise small or large displacements. In the small-displacement regime, displacements of 0, 1, 2, 5, 10, 15 and 20 mm/ $^\circ$ were assigned weights of 8, 7, 7, 5, 2, 1 and 1, respectively (i.e., contributions to $B1_{\text{average}}$ from the central position were weighted eight-fold compared to those from the R20 mm position in all voxels where the positions overlap). In the large-displacement regime, the same displacements were assigned weights of 0, 1, 2, 5, 7, 8 and 8, respectively.

Simulated B_1^+ maps were automatically co-registered upon exportation from Sim4Life (since the coil, and not the body, was moved when simulating motion). To create the average maps (and thereby gain B_1^+ information over an area larger than the head), maps from each position therefore had to be "de-registered" into scanner co-ordinates using interpolation (for rotations), or zero-padding (for translations). Due to the different padding methods used, the dimensions of the translations average map differed from those of the rotations average map. Placing the two average maps into a common space (i.e., one average map covering both translations and rotations) was therefore not done. As a result, rotation and translation displacements were treated separately in all cases (i.e., the average map containing translated positions did not include any B_1^+ information from rotated positions, and vice-versa). We also did not include any off-axis positions in the translation average maps, but the designed pulses were evaluated on all translated B_1^+ maps (including those from off-axis positions).

In addition to the translation and rotation average maps, direction-specific average maps were also generated for displacements along (or about) a single axis. For example, a direction-specific rightward average map was created using seven positions following different magnitude rightward motion. Where positive and negative displacements of the same type were present in the simulated data, separate maps were created (e.g., positive and negative yaw had separate direction-specific average maps). Findings for these direction-specific pulses are reported separately. Their purpose was to evaluate the importance of the direction of motion in the context of the accuracy-generalisability trade-off.

8.2.3 RF pulse design and evaluation of performance

For 2-D inversion, 180° spiral trajectory pulses were designed on a mid-axial slice using a fast optimal control algorithm [116] (previously used in Chapter 5, and described in Section 3.3). In the case of through-plane motion (inferior-superior, pitch and roll), the slice for which the pulse was designed was different to that on which it was evaluated, simulating the effects of motion without prospective through-plane correction. On the other hand, gradient and central frequency adjustment were included in evaluations at positions following in-plane motion. Motion susceptibility of pTx pulses is known to be slice-dependent [130, 173], so for one motion type (yaw), separate pulses were also designed and evaluated at five additional axial slice locations (each separated by 20 mm).

For 3-D excitation, non-selective 15 kT-points STA pulses were designed to excite a 36 mm axial slab using a spatial domain approach [32, 196] as described in Section 3.2. The *dzSpokes* tool from the ISMRM RF Pulse Design Challenge was adapted for 3-D excitation pulse design. The toolboxes are openly available here:

<https://github.com/wgrissom/ISMRM-RF-Pulse-Design-Challenge> (accessed 23/11/22).

2-D inversion and 3-D excitation pulses were designed using, firstly, the B_1^+ map from the central head position ($B_{1_{\text{centre}}}$). This represents the conventional approach for tailored pulse design. Secondly, pulses were designed using the averaged B_1^+ distribution ($B_{1_{\text{average}}}$), which represents the proposed approach. Target excitation and inversion profiles were uniform inversion or excitation across the entire region over which B_1^+ information was available. Therefore, the target profile for pulse_{average} was always larger than that of pulse_{centre}, since

the area covered by $B_{1\text{average}}$ was always larger than $B_{1\text{centre}}$. B_1^+ maps from the simulations were used for pulse evaluation. Excitation and inversion profiles produced by each pulse were calculated using non-uniform Fast Fourier Transform and Bloch simulations, respectively. Pulse performance was assessed using normalised root-mean-squared-error (nRMSE) in the resulting profiles (calculated according to Equation 5.1 in Chapter 5).

8.3 Results

In figures, the weighting regime which yielded the lowest nRMSE value is shown, and the weighting regime is indicated in the figures. The effects of weighting regime were not always in the anticipated direction, and are discussed later.

8.3.1 STA 3-D kT-points excitation

Three kT-points excitation pulses ($\text{pulse}_{\text{centre}}$, $\text{pulse}_{\text{average}}$ for translations and $\text{pulse}_{\text{average}}$ for rotations) were evaluated on B_1^+ maps from all positions. The maximum-observed nRMSE for the average pulses was around 5% lower than that of $\text{pulse}_{\text{centre}}$ (Figure 8.1). The average pulses yielded lower nRMSE at 61-73% of the evaluated positions (depending on the weighting regime used - this is discussed later in Section 8.3.3). Mean and maximum error reduction was 2.2% and 5%, respectively. For the remaining positions, the average pulses yielded higher error than $\text{pulse}_{\text{centre}}$, and error was increased by a similar amount (mean: 2.3% and maximum: 4.9%). As shown in Figure 8.1, the positions which did not benefit from the proposed approach were mostly roll and yaw rotations. Error was relatively insensitive to inferior-superior translation (<8%) using either approach.

Excitation profiles following rotation motion are shown in Figure 8.2. Figure 8.4 shows equivalent results following translation motion. For rotations, $\text{pulse}_{\text{average}}$ produced smoother profiles than $\text{pulse}_{\text{centre}}$, including (unexpectedly) at the central position. Improvements in uniformity could also be appreciated through the slab's cross-section (a roll example is shown in the top right of Figure 8.2). The most visible improvements in uniformity were seen for pitch rotation and axial-plane translations. Similar improvements were also present

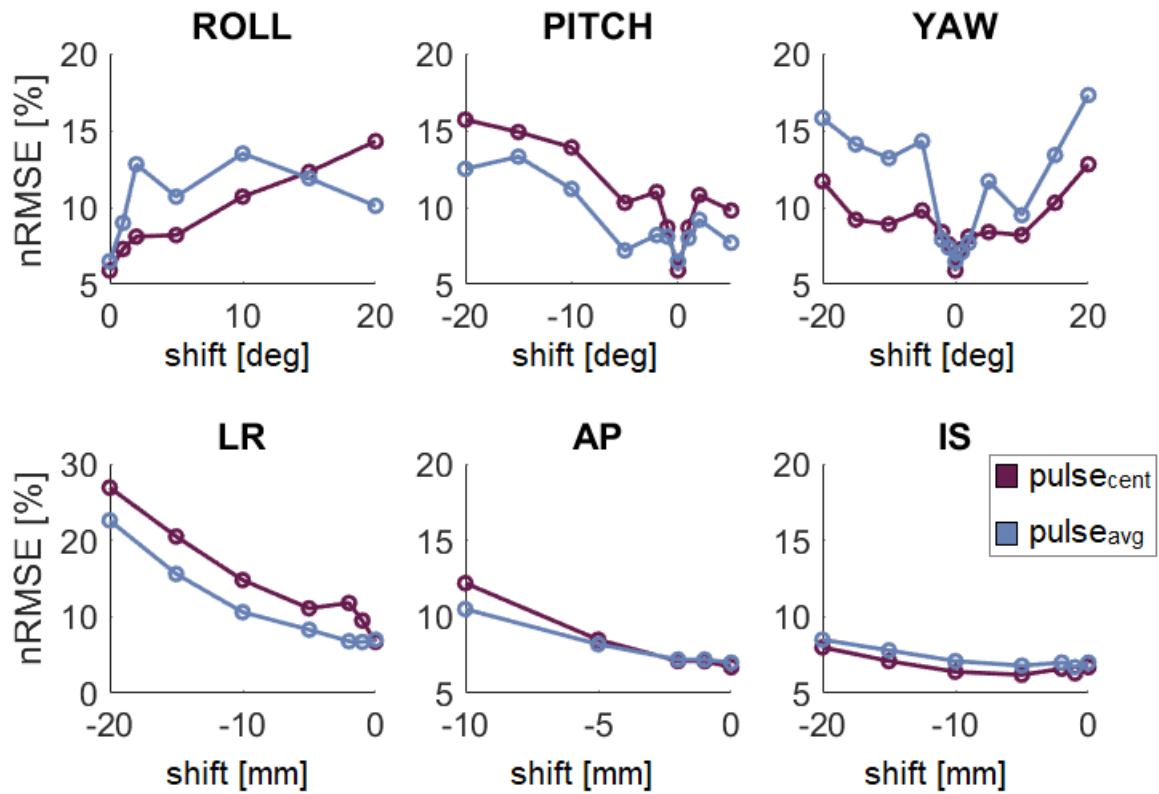


Figure 8.1: *nRMSE for STA kT -points pulses following all rotations and on-axis translations (off-axis translations are omitted for clarity). $nRMSE$ for both conventional ($pulse_{cent}$) and proposed approaches ($pulse_{avg}$) are shown. The $pulse_{average}$ results shown here comprise results from two different weighting regimes (small- and large-displacement), as indicated on the corresponding profile plots (Figures 8.2 and 8.4a).*




Figure 8.2: Excitation profiles for a mid-axial slice under conditions of rotation motion. Rows 1,3 and 5 show profiles generated by $pulse_{centre}$. Rows 2,4 and 6 show profiles produced by $pulse_{average}$. The proposed approach included two pulses which prioritise small (<10 mm, light blue outline) or large (≥ 10 mm, brown outline) displacements. The colour bar indicates signal magnitude (M_{xy}) as proportion of the target. The dark blue box indicates evaluation at the central position (i.e., without motion). An example cross section of the 3-D slab is shown in the upper right (bright green outline).

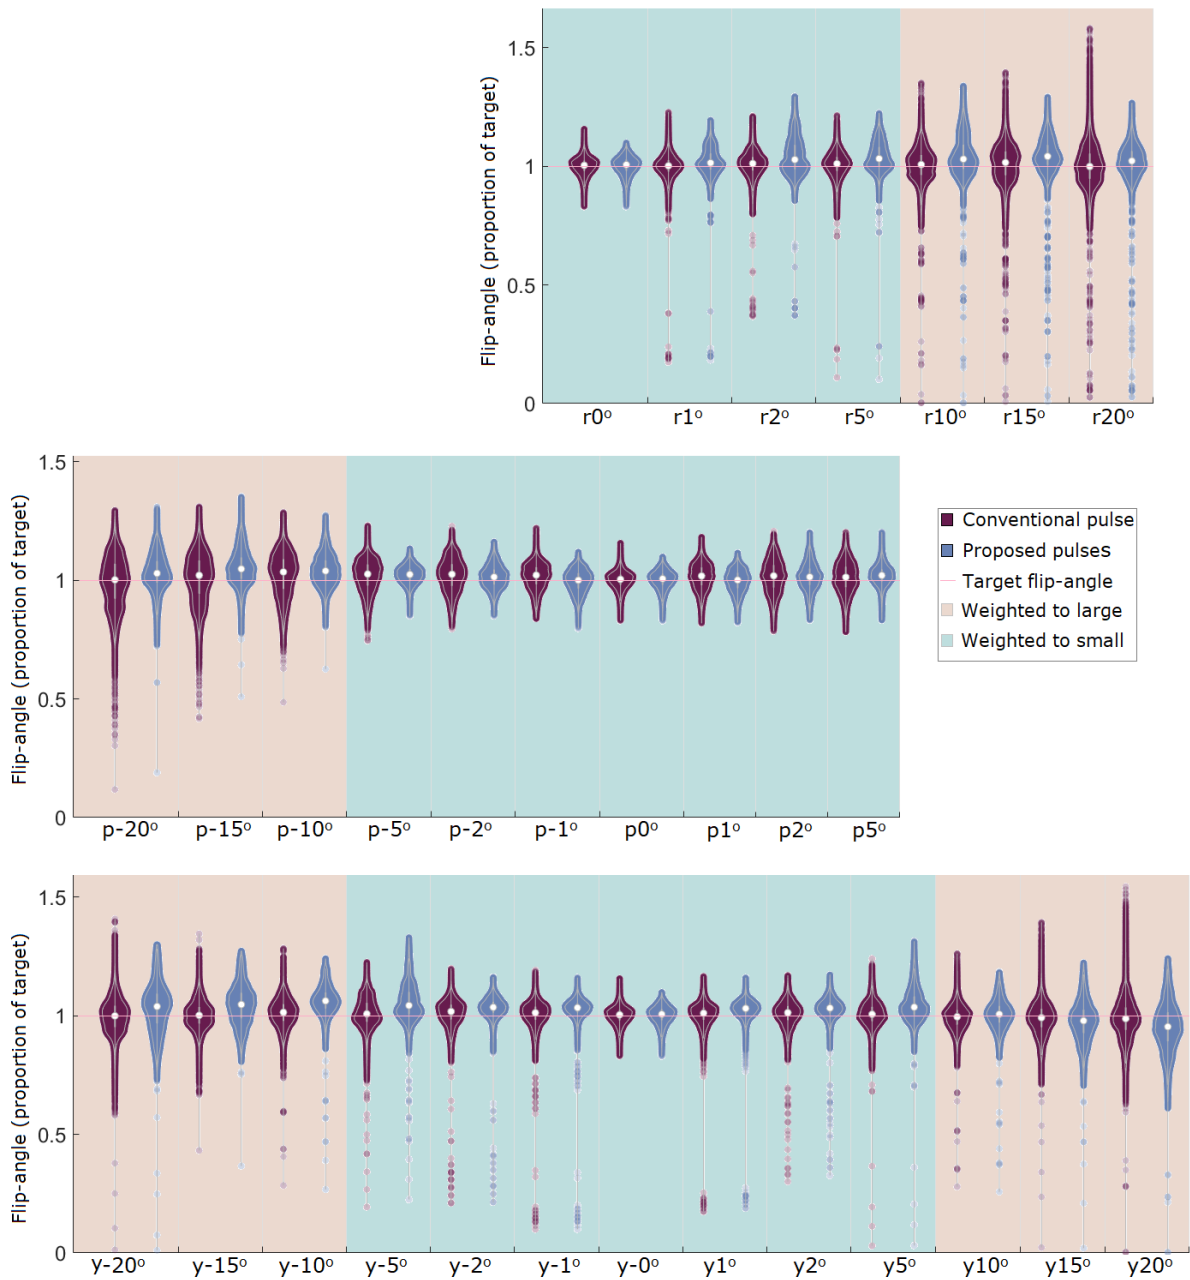


Figure 8.3: Flip-angle distributions throughout the entire 3-D volume for STA pulse_{centre} (dark purple distributions) and STA pulse_{average} (light purple distributions) following rotation motion. Distributions on a pale brown background were derived from the large-displacement weighting regime, while those on a pale blue background were derived from the small-displacement weighting regime.

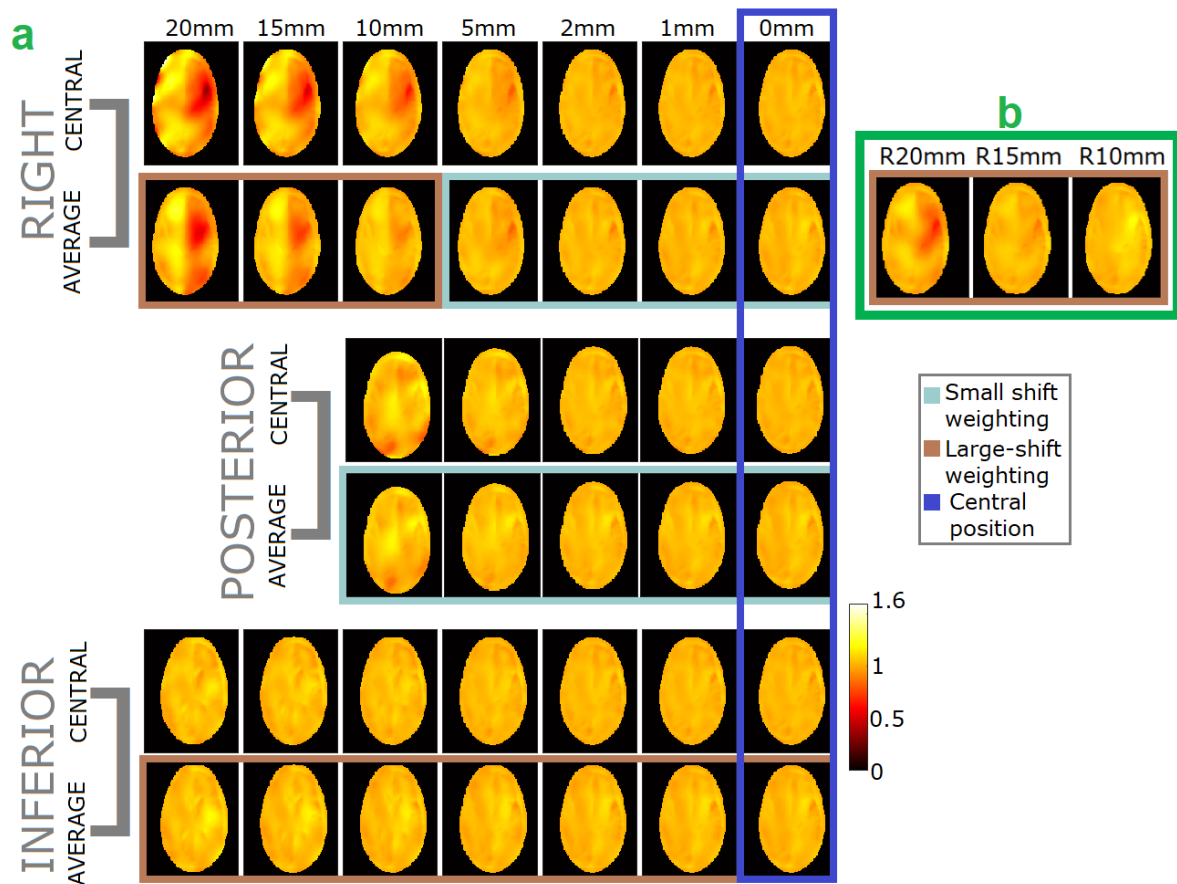


Figure 8.4: **a.** Excitation profiles for a mid-axial slice following translation motion. Rows 1,3 and 5 show profiles produced by $pulse_{centre}$, while rows 2,4, and 6 show profiles produced by $pulse_{average}$. The proposed approach included two pulses which prioritise performance when motion is small (light blue outline) or large (brown outline). The colour bar indicates signal magnitude (M_{xy}) as proportion of the target signal. The dark blue box indicates evaluation at the central position (i.e., without motion). **b.** (green box) Profiles for large rightward motion when $pulse_{average}$ is direction-specific (discussed in Section 8.3.4).

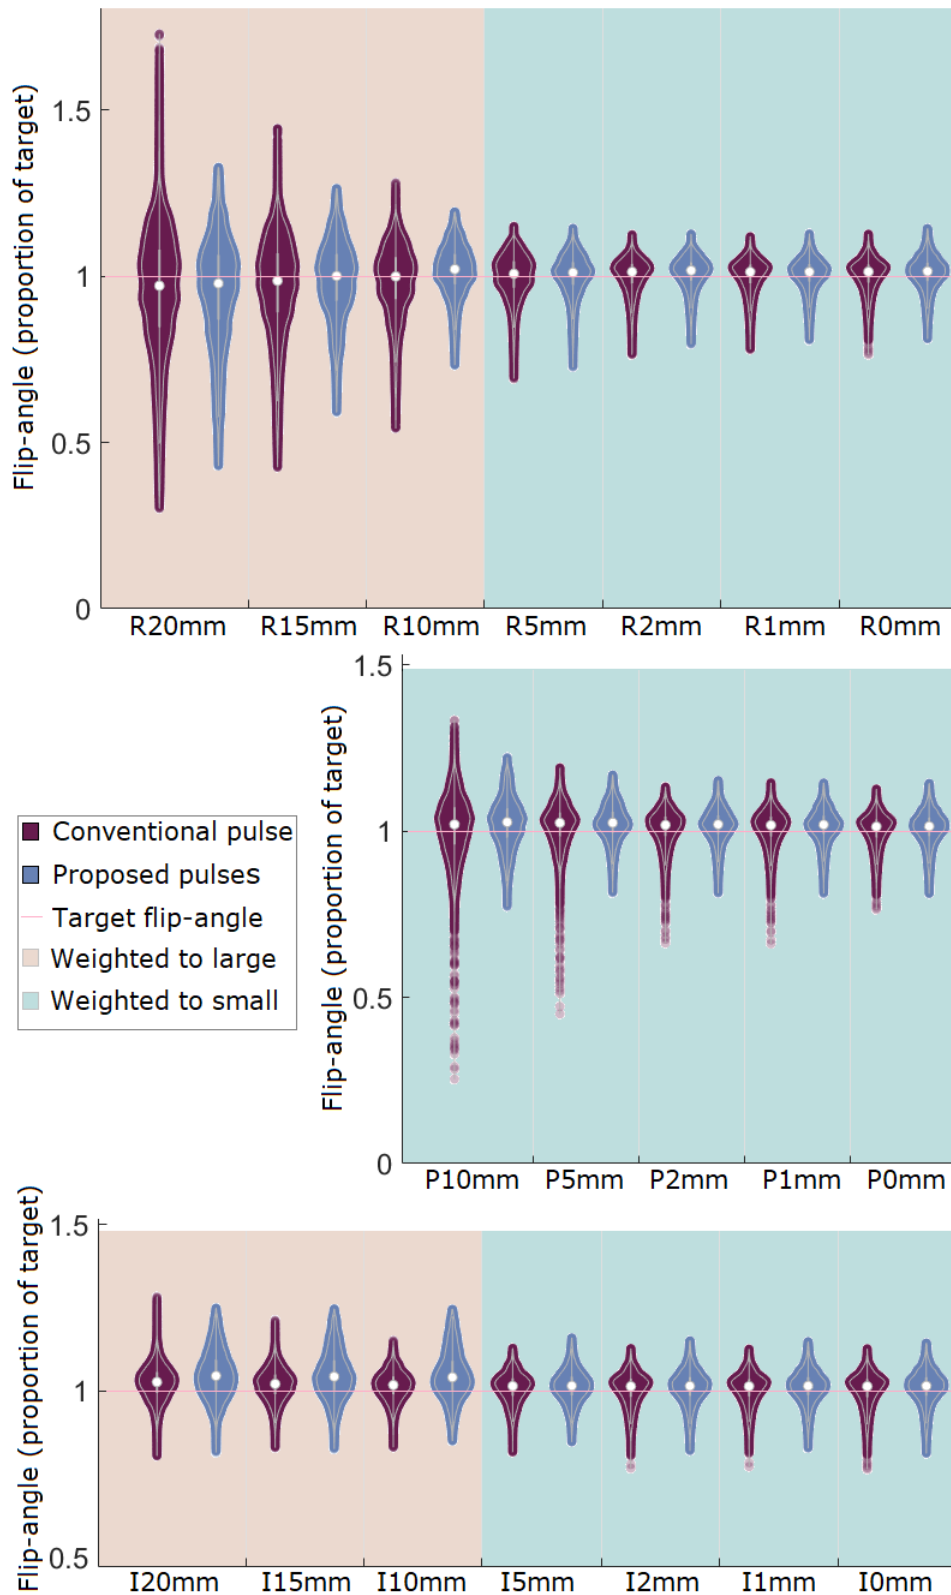


Figure 8.5: Flip-angle distributions throughout the entire 3-D volume for STA pulse_{centre} (dark purple distributions) and STA pulse_{average} (light purple distributions) following translation motion. Distributions on a pale brown background were derived from the large-displacement weighting regime, while those on a pale blue background were derived from the small-displacement weighting regime.

for off-axis positions. There was very little difference in performance between $\text{pulse}_{\text{centre}}$ and $\text{pulse}_{\text{average}}$ following inferior-superior translation.

Figures 8.3 and 8.5 show flip-angle distributions throughout the entire slab following rotation and translation motion, respectively. Generally, $\text{pulse}_{\text{average}}$ yielded distributions which were closer to the target flip-angle, and had a smaller range in flip-angle values. Unexpectedly, for rotations, $\text{pulse}_{\text{average}}$ outperformed $\text{pulse}_{\text{centre}}$ even without motion. A possible explanation is given later, in the Discussion section.

The proposed approach also reduced motion-induced phase change, even in cases where it did not provide benefit to the signal magnitude. Figure 8.6 shows phase error in terms of motion-related change in M_x and M_y signal components in each voxel, shown on the x and y axes respectively. Points along the diagonals therefore indicate that signal phase has not changed following motion. Pulses $_{\text{average}}$ yielded distributions which were closer to the diagonal than $\text{pulse}_{\text{centre}}$, indicating a reduction in motion-related phase error. Other motion types showed a similar pattern to the yaw example shown.




Figure 8.6: Voxel-wise motion-induced error for six yaw positions using the (phase-relaxed) kT -points $\text{pulse}_{\text{centre}}$ (top row) and $\text{pulse}_{\text{average}}$ (bottom row). X and Y axes show the motion-induced change in M_x and M_y signal components, respectively. Results across other movement directions showed a similar pattern to the yaw example shown here.

8.3.2 LTA 2-D spatially-selective inversion

Three LTA spiral-trajectory inversion pulses ($\text{pulse}_{\text{centre}}$, $\text{pulse}_{\text{average}}$ for rotations, and $\text{pulse}_{\text{average}}$ for translations) were evaluated using 2-D slices of B_1^+ maps from all positions. Compared to $\text{pulse}_{\text{centre}}$, $\text{pulse}_{\text{average}}$ yielded lower nRMSE for 63% of on- or about-axis positions, with a mean and maximum error reduction of 1.3% and 10.7%, respectively. Error following off-axis translations was reduced similarly, and is discussed later in Section 8.3.3. As shown in Figure 8.7, for the LTA case, roll and pitch rotations and large left-right displacements appeared to benefit the most from $\text{pulse}_{\text{average}}$. $\text{pulse}_{\text{average}}$ yielded slightly higher error in the remaining 37% of positions (mean and maximum increase were 0.4% and 1.3%, respectively).

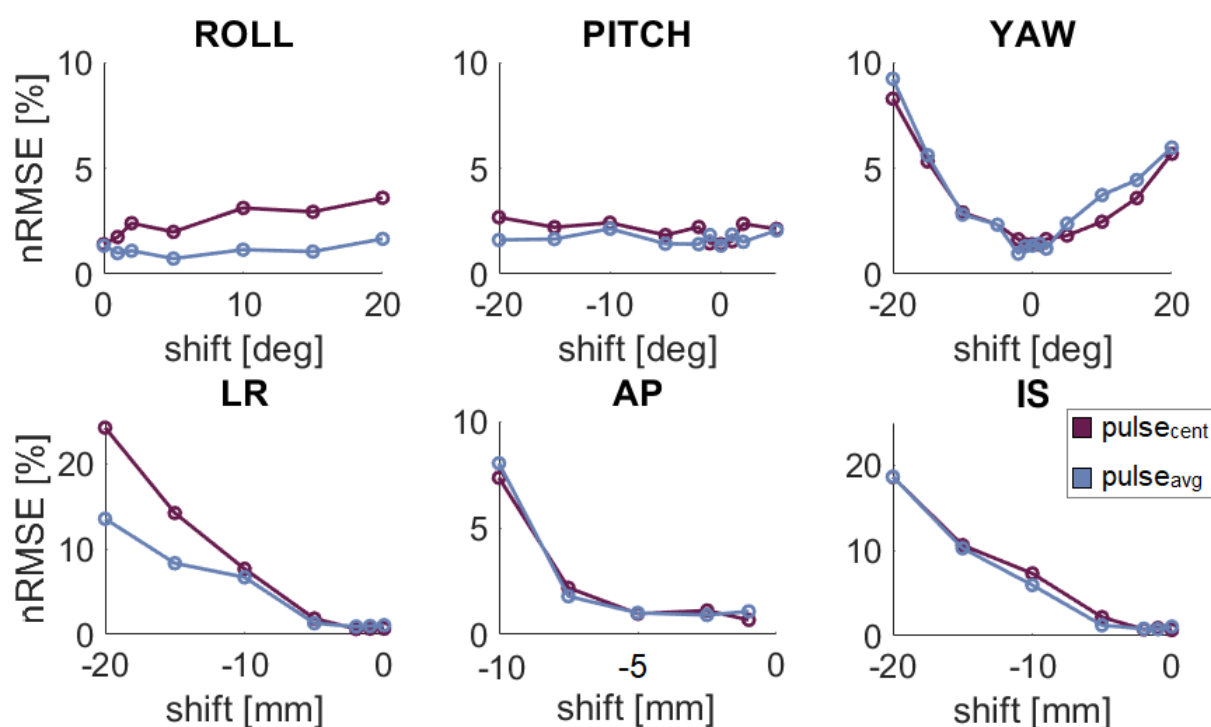


Figure 8.7: nRMSE for all rotations, and on-axis translations using conventional ($\text{pulse}_{\text{centre}}$) and proposed approaches ($\text{pulse}_{\text{average}}$) for LTA inversion. $\text{pulse}_{\text{average}}$ comprises two weighting regimes (small- and large-displacement), as indicated on the corresponding profile plots (Figs. 8.8 and 8.9). In all LTA cases, nRMSE was calculated within a smaller masked region shown in Fig. 8.9a

However, it is important to note that nRMSE was calculated from a region which excluded the outermost voxels. Due to spatial-selectivity of the LTA pulses used here, the most

obvious benefits of $\text{pulse}_{\text{average}}$ were in locations where no B_1^+ distribution information was available during B_1^+ mapping using the standard approach (because the head did not occupy those locations prior to motion). This can be seen following in-plane motion (right, posterior and yaw) shown in Figures 8.8 and 8.9. Large peripheral regions into which the head moved suffer under-inversion using $\text{pulse}_{\text{centre}}$, and this is mitigated when using $\text{pulse}_{\text{average}}$. The nRMSE shown in Figure 8.7 excludes these areas by calculating nRMSE within a smaller masked region (see Fig. 8.9a). Therefore, nRMSE reflects the subtler effects on uniformity of the main body of the inversion profile, and does not include the obvious benefits to the peripheral regions. This was done to make the error reduction more directly comparable to the non-selective STA pulses. Practical benefits of the proposed approach are substantially higher in spatially-selective contexts, as can be visibly appreciated.




Figure 8.8: Inversion profiles produced by spatially-selective LTA $\text{pulse}_{\text{centre}}$ (odd-number rows) and $\text{pulse}_{\text{average}}$ (even-number rows) following rotation motion. The colour scale indicates longitudinal magnetisation (M_z). Target $M_z = -1$

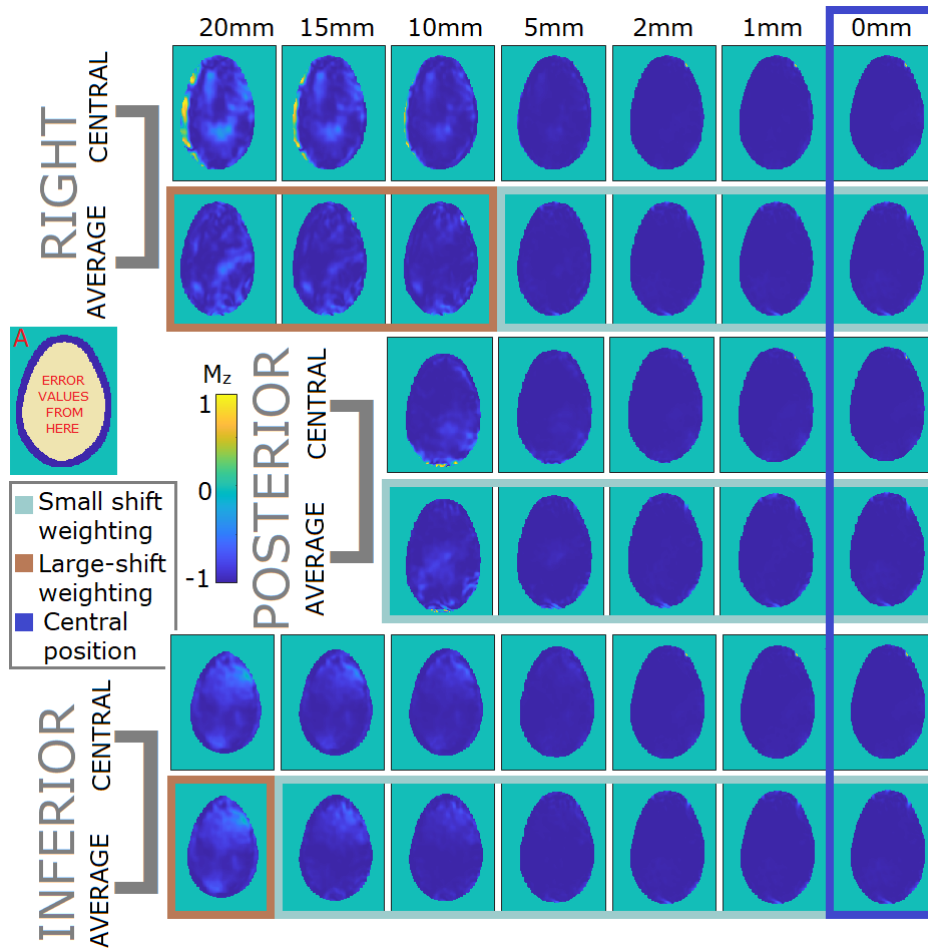


Figure 8.9: Inversion profiles produced by spatially-selective LTA pulse_{centre} (odd-number rows) and pulse_{average} (even-number rows) following translation motion. Colour scale indicates longitudinal magnetisation (M_z). Target $M_z = -1$ **A.** (left) The masked region used to calculate nRMSE for LTA pulses.

8.3.3 Weighting regime

In the small-displacement weighting regime, STA kT-points $\text{pulse}_{\text{average}}$ reduced error in 73% of positions, by a mean and maximum of 1.9% and 5.3%, respectively. For the remaining positions, error was increased by a mean and maximum of 1.6% and 6.1%, respectively. In the large-displacement weighting regime, error was more polarised: compared to $\text{pulse}_{\text{centre}}$, $\text{pulse}_{\text{average}}$ reduced motion-related error for 61% of positions by mean and maximum of 5.5% and 11.6%, while error was increased for the other 39% of positions by a mean and maximum of 5.3% and 9.5%, respectively.

Performance of the large- and small-displacement weighting regimes with LTA pulses following translation motion is shown in Figure 8.10. In the small-displacement weighting regime, $\text{pulse}_{\text{average}}$ reduced motion-related error compared to $\text{pulse}_{\text{centre}}$ in 78% of cases (mean, max = 1.5%, 4.6%). By comparison, large-displacement weighting yielded a lower error in only 35% of cases (mean, maximum = 5.1%, and 9.9%). This indicates that large-displacement weighting had a similar polarising effect to that seen for the STA pulses. For translations, the positions which benefited from the large-displacement weighting were those following large translations (as expected).

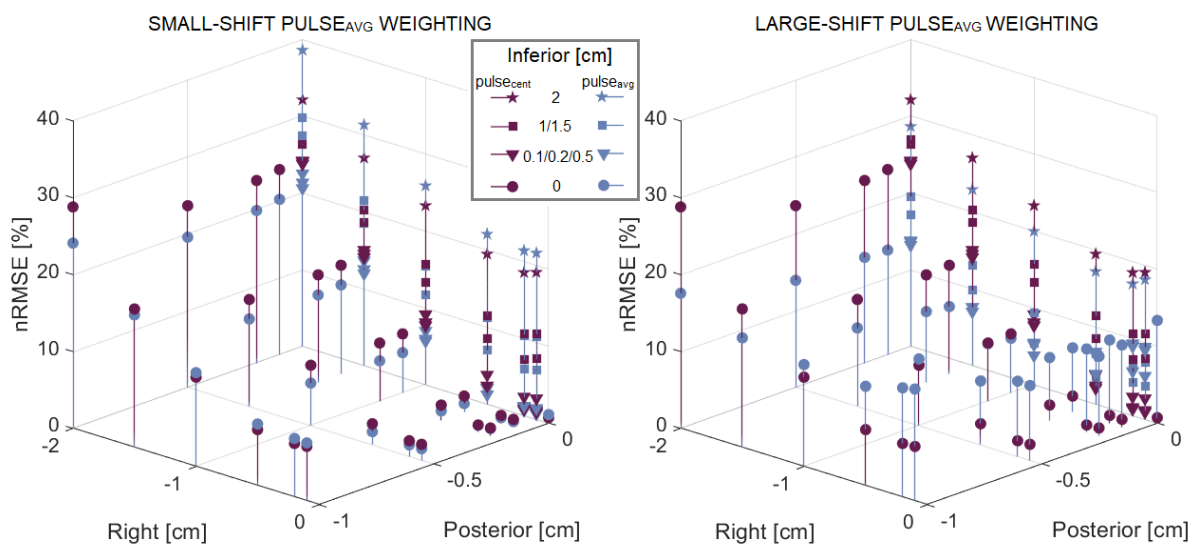


Figure 8.10: $nRMSE$ for LTA $\text{pulse}_{\text{centre}}$ (dark purple) and $\text{pulse}_{\text{average}}$ (light purple) for all on- and off-axis translations. The left and right figures show error obtained with pulses designed using small- and large-displacement B_1^+ weighting regimes, respectively.

In contrast, for the LTA rotations pulse, large-displacement pulse_{average} outperformed both pulse_{centre} and small-displacement pulse_{average} at all evaluated rotation positions. Even in the absence of motion, error produced by large-displacement pulse_{average} (1.3%) was lower than that of small-displacement pulse_{average} (3.1%), and was comparable to that of pulse_{centre} (1.4%). Similarly, the large-displacement kT-points pulse_{average} unexpectedly outperformed its small-displacement counterpart for all inferior translations (although motion-related error was very low for all pulses following this motion type, as discussed previously).

8.3.4 Direction-specificity of pulse_{average}

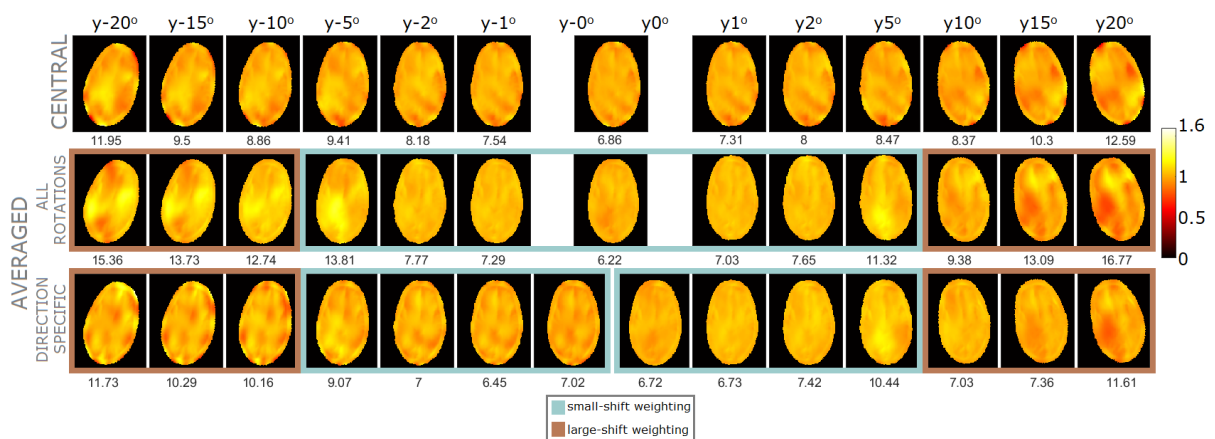


Figure 8.11: Comparison between multi-directional and direction-specific averaging in the B_1^+ maps used in the design of 3-D excitation pulses. Rows 1, 2 and 3 show profiles obtained using pulses designed for the central, multi-direction averaged, and direction-specific averaged B_1^+ maps, respectively. (Rows 1 and 2 are identical to rows 5 and 6 in Figure 8.2).

To explore improvements offered by using direction-specific average pulses, we designed extra pulses using average B_1^+ maps comprising a range of positions along a single motion direction. Flip-angle uniformity was improved by using direction-specific averaging, compared to both pulse_{centre} and the multi-direction average pulse following motion, with a smaller range in flip-angle throughout 3-D volumes. This was expected, as generalisability is reduced (and therefore performance should improve) by including less positions in the average maps. Motion-related error also reduced in cases which did not benefit much from the multi-direction pulses_{average}, such as yaw, shown in Figure 8.11. Figure 8.4b shows performance of the rightward-specific excitation pulse at three evaluated positions. Finally,

the performance of LTA direction-specific, multi-direction and centre-only pulses are shown for yaw in Figure 8.12.

8.3.5 Slice location

To investigate the effects of slice location in the 2-D LTA case, several inversion pulses were designed at the five slice locations shown in Figure 8.12. For each slice, the four pulses were: $\text{pulse}_{\text{centre}}$, $\text{pulse}_{\text{average}}$ (multi-direction), $\text{pulse}_{\text{average}}$ (positive yaw) and $\text{pulse}_{\text{average}}$ (negative yaw). All average pulses in this analysis were designed using the small-displacement weighting regime. Across $\text{pulses}_{\text{centre}}$, error was generally higher for superior slices, however there was no clear slice-dependence in terms of the benefit provided by $\text{pulse}_{\text{average}}$.

Negative yaw caused more motion-related error, as seen by the steeper increase in error across pulses and slices, compared to positive yaw. However, error caused by negative yaw motion was reliably reduced by the negative yaw $\text{pulse}_{\text{average}}$, whereas the positive yaw $\text{pulse}_{\text{average}}$ did not reliably reduce error for positive yaw motion.

Small areas of incomplete inversion were observed for some pulses (e.g. positive yaw pulse, slice 100). This is likely to stem from a pulse design issue, as it was also reported in the original pulse design paper [116].

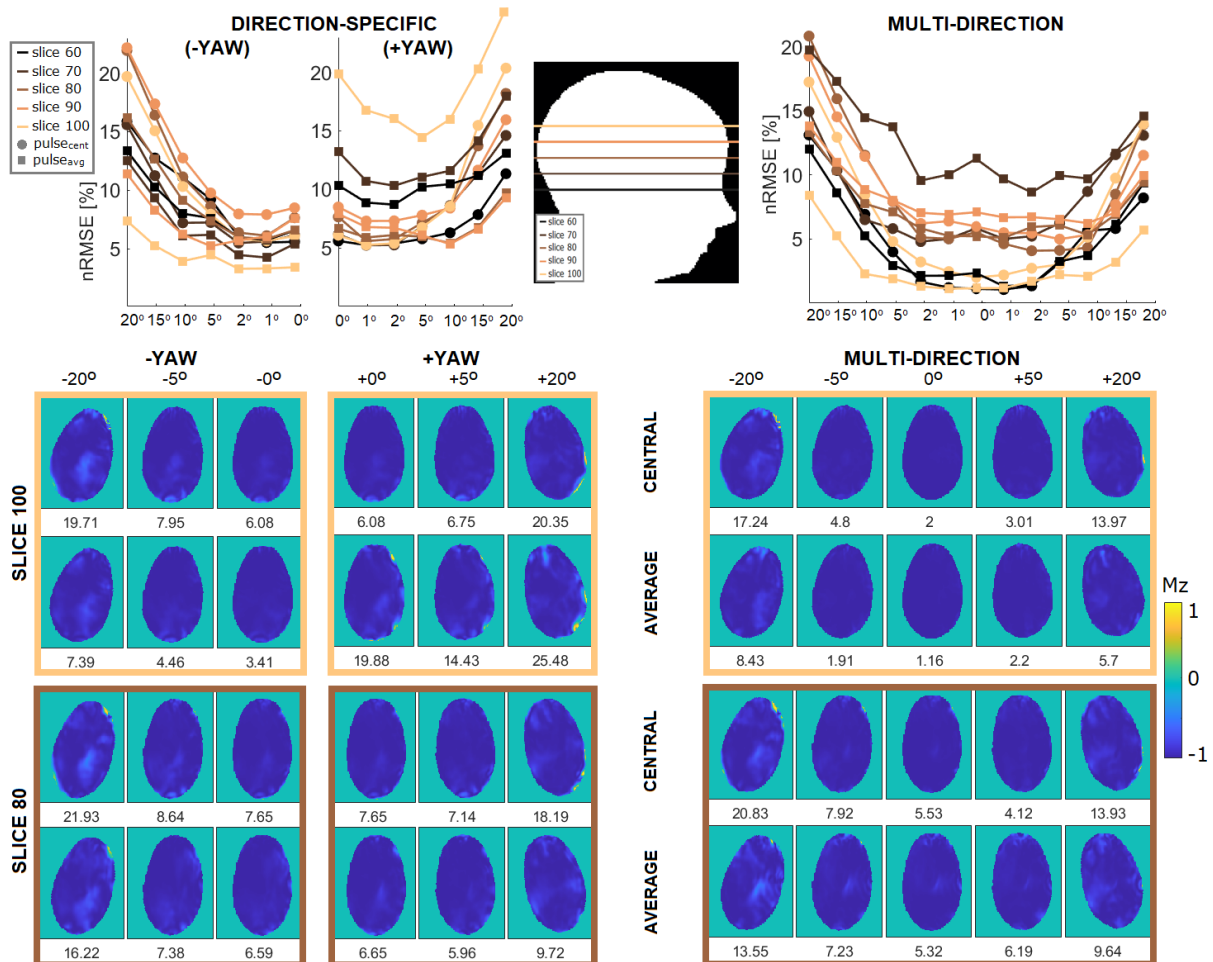


Figure 8.12: Error and inversion profiles for 5 slice locations following yaw rotation. **Top-left:** nRMSE for direction-specific (separate negative yaw and positive yaw) pulses_{average}. **Top-middle:** Axial slice locations. All slices were separated by 18 mm. Four pulses (positive yaw pulse_{average}, negative yaw pulse_{average}, multi-direction rotation pulse_{average}, and pulse_{centre}) were designed for each slice. **Top-right:** nRMSE for multi-direction pulse_{average}. **Bottom-left:** Inversion profiles for direction-specific pulses_{average} shown at slices 80 and 100. **Bottom-right:** Inversion profiles for multi-direction pulses_{average} shown at slices 80 and 100.

8.4 Discussion

This study presented a way to reduce the impact of head motion on flip-angle profiles for pTx. Since STA and LTA pulses respond to motion differently (see Chapter 5), versatility of the method was a priority. This was achieved by reducing the degree of positional information contained in B_1^+ maps used for pulse design. pTx pulses were designed using averaged B_1^+ distributions, which combine B_1^+ information from a single subject at multiple head positions. In line with previous research, we found that conventional pulse performance was most sensitive to axial-plane displacements [130, 173]. The average B_1^+ map approach outlined here reduced motion-related nRMSE by up to 5% for STA excitation and 10.7% for LTA inversion. The approach was most effective for in-plane translations and pitch rotations, while for yaw rotations, direction-specific averaging was required for error reduction (thereby trading generalisability for accuracy).

Our qualitative results highlight that B_1^+ averaging affects signal in two conceptually distinct ways. Firstly, through expansion of the B_1^+ map periphery and secondly, through smoothing the maps' interior regions. Since B_1^+ cannot be mapped outside of the head in practice, the former mechanism always provides a benefit compared to the conventional approach when motion occurs (as no B_1^+ information is available in those locations). This is especially appreciable in spatially-selective contexts such as the LTA pulse tested here (although the mechanism was excluded from nRMSE calculations for LTA pulses through masking, as previously explained). On the other hand, the latter (smoothing) effect of averaging effectively reduces the degree of precision tailoring of the designed pTx pulse. This mechanism, unlike the former, does not guarantee improvement over the conventional approach. Our results support the variable effect of this mechanism on performance, but suggest that over 50% of cases do benefit in both non-selective and spatially-selective cases.

With this in mind, it may be preferable to use only the outer regions of $B_{1_{\text{average}}}$ (i.e., the peripheral areas which are not occupied by the head in the central position). These areas could be superimposed onto a standard B_1^+ map containing the B_1^+ information from the central position. According to our results, pulses designed in this way would yield reduced error in peripheral locations, while minimally affecting flip-angle at the centre. A similar concept was employed by Wallace et al. [181], where B_1 maps were extrapolated using radial

basis functions into areas of the FOV beyond the measurable region occupied by the head.

As expected, the proposed approach required that some flip-angle accuracy was traded for generalisability across positions. Compared to tailored pulses, Universal Pulses (see Section 3.4.4) typically increase flip-angle nRMSE by around 5-7% [123], which is comparable to the 1-3% error increase we observed at the central position. In either case, this trade-off may be justified in imaging contexts which rely on relative signal change over time (e.g., functional MRI), where reducing sensitivity to motion may be more valuable than a high-performing (i.e., low nRMSE), but motion-sensitive pulse.

The method proposed here is conceptually compatible with a Universal Pulse-like approach [123]. Positional variance could be directly included in Universal Pulse design (by including B_1^+ maps from various positions -as well as various individuals- in their design). This approach has recently been demonstrated to reduce the effects of respiration on flip-angle uniformity [208]. However, combining the two different approaches offers the benefit that weaknesses within each method are not inherently augmented, as is the case if the same method was applied to generalise across both individuals and positions.

In most cases, small-displacement weighted pulses_{average} (where small displacements were prioritised in the averaging process) led to small but fairly consistent error reduction across positions, whereas large-displacement weighted pulses_{average} had a more polarising effect. The latter reduced (but also increased) error by a larger amount depending on the type and amount of motion. Only two weighted-average regimes were tested here, however optimisation of the weightings may improve the benefits offered by the proposed approach.

Weighted-averaging B_1^+ maps only ever extends the area over which B_1^+ information is known, bringing the mapped information closer to the pTx channels, where B_1^+ magnitude is highest. If RF power is limited during pulse design (e.g. to manage global SAR), averaging therefore effectively serves to reduce the RF power available at channels for which channel-tissue distance is decreased (due to the extended B_1^+ information in those locations). This is likely to have unpredictable consequences on local SAR behaviour, which depends on channel weightings and tissue-channel distance in pTx [130].

It was anticipated that translation and rotation pulses_{centre} would be similar in design and performance, however their kT-points excitation trajectories differed substantially. Importantly, STA rotations pulse_{centre} also yielded a less uniform profile than translations pulse_{centre} even without motion, and this may have confounded motion-related changes in performance between the two types of movement.

To use this approach in practice, B_1^+ mapping of the subject at several different head positions would be required in order to create the average maps, which could be rather time-consuming. Instead, subject-specific B_1^+ maps for several different head positions could be predicted (following a single B_1^+ mapping sequence) using a deep learning method described in Chapter 9 of this thesis.

8.4.1 Conclusions

We have presented a versatile method which can reduce motion-related error in pTx pulses at 7T. The method uses average B_1^+ maps, calculated by weighted-averaging B_1^+ maps from the head in several different positions. The approach is flexible and can be modified to prioritise generalisability or accuracy. Importantly, the concept can be applied with a variety of pulse designs including STA, LTA, spatially-selective and non-selective pulses. SAR response to the proposed pulses was not investigated here and represents an important future consideration.

Chapter 9

Approaches to reduce motion-sensitivity in pTx: Deep Learning B_1^+ prediction

The work presented in this chapter led to the following publications:

Plumley, A. et al. Rigid motion-resolved B_1^+ prediction using deep learning for real-time parallel-transmission pulse design. *Magnetic Resonance in Medicine* 87(5), p.2254-2270 (2022).

Plumley, A. et al. Motion-resolved B_1^+ prediction using deep learning for real-time pTx pulse-design. Proceedings of the International Society for Magnetic Resonance in Medicine, #0223 (2021).

9.1 Introduction

As discussed throughout this thesis, pTx with tailored RF pulses offers an effective solution to help overcome the non-uniform B_1^+ fields intrinsic to MRI at 7T (see Chapter 3), but subject motion alters electromagnetic field distributions during the scan (Chapter 4). Motion therefore degrades performance of tailored pTx pulses due to their reliance on knowledge of

the B_1^+ field. Retrospective motion correction can be used to reduce the appearance of motion artefacts, but dynamic field changes cannot be corrected without motion-resolved field maps. Additionally, E-fields (which relate directly to SAR) are also affected by motion, and these changes clearly cannot be addressed retrospectively due to the safety implications.

It is therefore desirable to overcome the motion dependence of tailored pTx pulse performance, and to do so using prospective techniques. In Chapter 8, a method to increase the motion-robustness of pTx pulses was described. The method used composite B_1^+ maps for pulse design, however flip-angle profiles were compromised compared to tailored pulses as some accuracy in performance is traded for generalisability in this approach. Instead, real-time pTx pulse design has been proposed as a solution [189], in which channels' complex coefficients are pseudo-continuously updated to counteract motion-induced B_1^+ sensitivity changes. Multi-spokes pTx pulses can be designed in less than 0.5 seconds [189] whereas 2-D spatially selective spiral pulses can be estimated in about 9 ms using deep neural networks [209]. With motion detection (e.g., [181, 184, 210, 211]), channel updates could be determined by instantaneous head position, retaining flip-angle uniformity in cases of arbitrary and / or extreme motion. However, the required updates to channel coefficients depend on the motion-related field changes. Because real-time (i.e., motion-resolved) B_1^+ maps are not measurable, this requires that the relationship between head position and B_1^+ distribution is characterised. The complex interference patterns between channels in pTx make this a non-trivial problem.

9.1.1 Convolutional Neural Networks

Deep convolutional neural networks (CNNs) are well-suited to learning complex non-linear transformations. Encoder-decoder generative networks can be used for so-called image-to-image translation, where networks are trained to “translate” an input image from one initial domain into a corresponding output image from the target domain. Training uses corresponding pairs of input and ground-truth images, which are from the initial and target domains, respectively. The approach is versatile and has been demonstrated for applications including style transfer (for example, converting photographs into sketches or maps), day-to-night, black-and-white to colour, or spatial labels to images [212].

A commonly-used network for image translation is the so-called U-Net [213]. Consecutive layers of spatial convolution underlie the learned transformation. Specifically, a number of small (e.g., 5×5) 2-D or 3-D spatial filters are convolved over the input image in a step-wise fashion. Each filter is assigned a weight (initially random), which ultimately serves to reflect how well the contents of the spatial filter represents the contents of a region in the image. To introduce non-linearity to the network, weights are passed through a non-linear function known as the activation function. One example is the rectified linear unit (ReLU), which is a linear function for positive values, but zero for negative values, thereby forcing all negative weights to be zero. This process is repeated throughout a series of layers, in which the filters and weights from one layer act as input to the next layer.

The size of the input at each stage decreases, either due to using a convolutional stride greater than one, or using pooling techniques (where values over a small region are pooled to a single weight according to the maximum or mean across the region). Simultaneously, the number of filters at each convolutional layer increases. The image representation thereby becomes compressed. At this stage, the spatial filters can be likened to basis functions which comprise the image in a (learned) weighted combination. At the decoding stage, the process is reversed, yielding larger inputs and fewer filters at each successive layer, until an image is output from the final layer. To avoid the loss of image detail, so-called skip-connections can be used, which are connections joining non-adjacent corresponding convolution and de-convolution layers, thereby avoiding the “bottleneck” (the narrowest part of the network, where information is maximally compressed).

Throughout training, network weights are updated using backpropagation of errors [214, 215]. In this process, the derivative of the network loss (i.e., the cost function) is calculated with respect to the weights in the network, and the weights are updated accordingly using optimisation algorithms such as Adam [216]. The process is repeated across the entire training data set a set number of times (epochs), although training can be stopped at any time if the loss calculated over a separate validation data set begins to increase (as this can indicate overtraining / overfitting). This is known as early stopping. Another form of regularisation is the use of dropout layers, where a proportion of the weights are temporarily ignored during training. This has the effect of reducing inter-dependence among weights (which can cause network instability), and encouraging sparsity among weights.

A closely related implementation to U-Nets is conditional generative adversarial networks (cGANs) [212, 217]. In the generative model architecture described above, the cost function usually consists of L1 or L2 loss between predicted and target images. Resulting outputs can therefore be blurry because these loss types have an averaging effect. To reduce blurring effects, cGANs include an additional loss term provided by a second CNN known as the discriminator, which is trained alongside the generator. The discriminator is trained to distinguish between generator-output and ground-truth images. By maximising this adversarial loss function while simultaneously minimising the generative loss, the generator produces more realistic outputs which are capable of “fooling” the discriminator.

Deep CNNs have recently been used for image-to-image translation problems within MRI. Mandija et al. trained cGANs to predict electrical property (EP; conductivity and permittivity) maps from input images consisting of B_1^+ fields and transceive phase [205]. The method produced high-quality EP maps which were comparable to existing MR electrical properties tomography (MR-EPT) methods. EP distributions are implicated in pathology, and are required for individual SAR calculations (see Equation 1.11). Similarly deep learning has been applied to quantitative susceptibility mapping (QSM), which can show myelin, iron and calcium deposition in pathology. Examples include QSMNet [218] and DeepQSM [219], both of which used U-Net-based architectures. In Meliado et al., cGANs were trained to predict tailored SAR distributions from measured B_1^+ maps [132].

Deep CNNs have also been used to estimate (non-pTx) B_1^+ distributions. Wu et al. predicted high-quality maps from reconstructed T1-weighted images, removing the need for B_1^+ mapping while still allowing retrospective correction of B_1^+ related artifacts in quantitative MRI [220]. This approach was limited to post-processing; prediction quality deteriorated when B_1^+ was predicted directly from under-sampled images. Abbasi-Rad et al. used a CNN to reconstruct B_1^+ from a localiser scan for the purpose of SAR reduction through pulse scaling based on slice-wise B_1^+ magnitude; however, B_1^+ prediction quality was dependent on head position [221].

In this work, we train a system of cGANs [217] to predict pTx B_1^+ distributions following simulated head motion, given the initial B_1^+ maps at the centered position as input. If used in conjunction with motion detection, this would constitute motion-resolved B_1^+ map estimation, and therefore permit real-time tailored pulse design [189]. B_1^+ map prediction

quality is assessed by comparison with ground-truth (simulation output) B_1^+ maps following motion. Furthermore, flip-angle distributions of multi-spokes pTx pulses designed using network-predicted B_1^+ maps are compared with those produced by tailored pulses designed using the initial subject-specific B_1^+ maps alone. Finally, we also observe peak 10-g averaged local SAR for both pulses following motion.

9.2 Methods

9.2.1 Simulations and data

Dizzy, Billie, Duke, and Ella (Figure 9.1A–D) of the Virtual Population [143] (IT'IS, Zurich, Switzerland) were simulated with a generic 8-channel pTx coil in Sim4Life (ZMT, Zurich, Switzerland) using the same simulation framework as described in Section 5.2.1. Each model was simulated at one central, and 32 off-center positions. Off-center positions included rightward 2, 4, 5, 10, and 20 mm, posterior 2, 4, 5, and 10 mm, and all possible combinations thereof. These 29 positions are hereafter referred to as the R-P grid (Figure 9.1E). In addition, yaw 5°, 10°, and 15° positions were also simulated (Figure 9.1F). The Duke model was scaled to 90% of the original size, as the body and coil models intersected at some positions when the model was full-sized.

Channels' 3-D B_1^+ , electric field, current density, and SAR distributions were exported to MATLAB (The MathWorks, Natick, MA). To incorporate interactions between channels for local SAR evaluations, 10g-averaged Q-matrices were calculated as defined in Equation 3.8 [47, 129].

B_1^+ maps from 51 slices spanning a mid-axial slab with a thickness of 9 cm from the Duke and Ella body models (Figure 9.1C,D) were prepared for network training by interpolating to 256×256 in-plane resolution. The same preprocessing was applied to the Billie and Dizzy data, but at only six slice locations (Figure 9.1A,B). Magnitude and phase data were separated and normalised between 0 and 1, where 1 corresponds to the maximum magnitude across all channels, slices, and body models, and to 2π for phase. Random offsets were applied to phase maps so that the phase wrap boundary did not occur at the same location

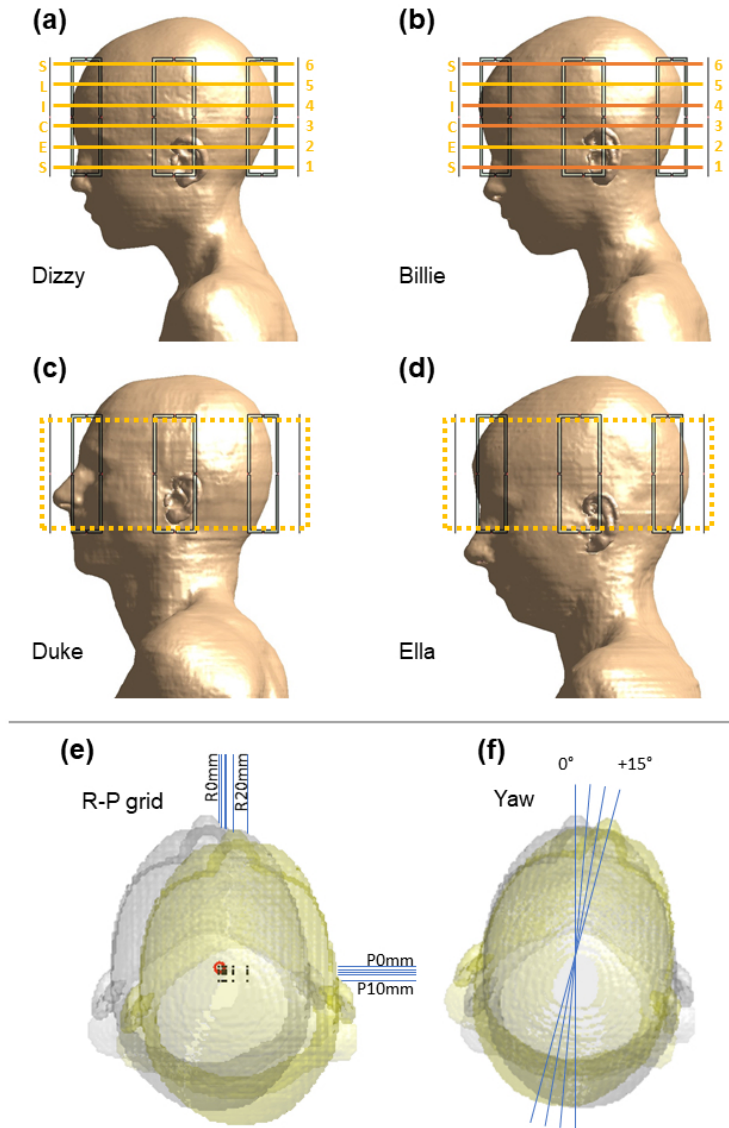


Figure 9.1: *Simulation model setup. (A–D) The four body models used in Sim4Life simulations. Ella and Duke (C, D) were used to generate training data for networks, Billie (b) was used for network validation and testing, and Dizzy (a) was used for testing only. Testing (including pulse design) was conducted at the six indicated slice locations. Validation slices were offset by about 4 mm from these, but within the same axial range. Slices in orange were also used for specific absorption rate (SAR) evaluations. All axial slices within the dashed slab were used for training. (E) Positions simulated for the R-P grid. The origin of the central position is indicated with a red circle, whereas all other positions' origins are indicated with black dots. Axial displacements were all possible combinations of rightward (R) 0, 2, 4, 5, 10, and 20 mm and posterior (P) 0, 2, 4, 5, and 10 mm. (F) Yaw rotations were 5°, 10°, and 15°. The head at the central position (gray isosurface) and most extreme displaced position (yellow isosurface) are shown.*

across slices. B_1^+ maps were input to networks as individual axial slices with size $256 \times 256 \times 8$, where the third dimension is channels. Corresponding B_1^+ map slices before (input) and after (ground truth) a given displacement formed the networks' input-target pairs. Note that inputs are not necessarily at the centered position (explained later).

9.2.2 Neural networks and network training

Models were implemented in TensorFlow 2.3 [222] using Python 3.7. Network architecture is summarised in Figure 9.2. Except where specified, network hyperparameters were the same as those used in the Pix2Pix cGAN described in Ref. [212]. The generators were U-Nets with eight convolutional (encoding) and eight de-convolutional (decoding) layers linking the input and output (predicted) B_1^+ maps, each followed by ReLU activation layers. Filters were 4×4 for magnitude and 8×8 for phase. Although comprehensive hyperparameter optimisation was beyond the scope of this project, during initial testing it was found that phase networks benefited from the large receptive field of 8×8 filters. Conversely, magnitude networks generated smoother maps when more filters were used. To avoid increasing the number of trainable parameters, filter size for magnitude was reduced. The number of filters (initially 128 for magnitude and 64 for phase) increased to a maximum of 1,024 (512 for phase) for the middle layers, and stride size was 2. Filters were split into eight groups to facilitate simultaneous processing of all pTx channels. For phase, batch normalisation was applied at all layers except the first convolution layer. For magnitude, removing batch normalisation resulted in a smoother training curve and higher-quality estimated maps. Skip connections joined each convolution layer to the symmetric de-convolution layer. The network was regularised using dropout layers following each of the first three de-convolution layers (dropout rate = 0.5).

The input B_1^+ maps, concatenated with either ground-truth or generator-predicted B_1^+ maps, served as input to the discriminators, which consisted of five convolution layers. The discriminators used leaky ReLU activation layers ($\alpha = 0.3$) as recommended in Radford et al. [223]. Filter size was the same as that for the generators, and convolution stride was 2 except for the final two layers, where it was 1. A single 2-D distribution of probability (entropy) values was output.

The overall cGAN loss can be expressed as:

$$loss = \arg \min_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G) \quad (9.1)$$

where G denotes the generator, D is the discriminator, and λ (set to 100) is a scaling parameter acting on the L1-norm between generator-predicted and ground-truth maps.

The first term is binary cross-entropy loss which can be further described as:

$$\mathcal{L}_{cGAN}(G, D) = \mathbb{E}_{B1_{gt}, B1_{predicted}} [\log D(B1_{gt}, B1_{predicted})] + \mathbb{E}_{B1_{gt}, B1_{initial}} [\log(1 - D(B1_{gt}, G(B1_{gt}, B1_{initial})))] \quad (9.2)$$

where $B1_{gt}$ are the ground-truth displaced $B1^+$ maps; $B1_{predicted}$ are the generator-predicted displaced $B1^+$ maps; $B1_{initial}$ are the pre-displacement $B1^+$ maps (network input); and \mathbb{E} is the expected value.

The effect of head motion on $B1^+$ depends on the displacement type (i.e., direction and magnitude) [171, 173]. Because data-driven approaches assume that all input-target pairs share a common underlying mapping, separate networks were trained for different displacement types (e.g., rightward vs posterior). Head motion was discretised into large (5 mm) and small (2 mm) displacements in rightward (R) and posterior (P) directions to cover the R-P grid. Additional networks were trained for 5° yaw rotation. Separate networks were trained for magnitude and phase, yielding a total of 10 networks.

The Duke and Ella data were used for training. All available relative displacements were included. For example, to train the R5 mm network, (input)–(target) pairs included (R0, P0 mm)–(R5, P0 mm); (R5, P0 mm)–(R10, P0 mm); (R5, P2 mm)–(R10, P2 mm). This yielded a training data set of 1,020 unique slices for rightward networks, and 1,224 for each posterior network. The yaw network training database was smaller (306 slices).

The Adam optimiser [216] was used to train models for 60 epochs. Learning rate proved to be critical during initial testing, so learning rates within the range .00005 to .001 were tested. The default value of .0002 converged most effectively and was therefore used for all

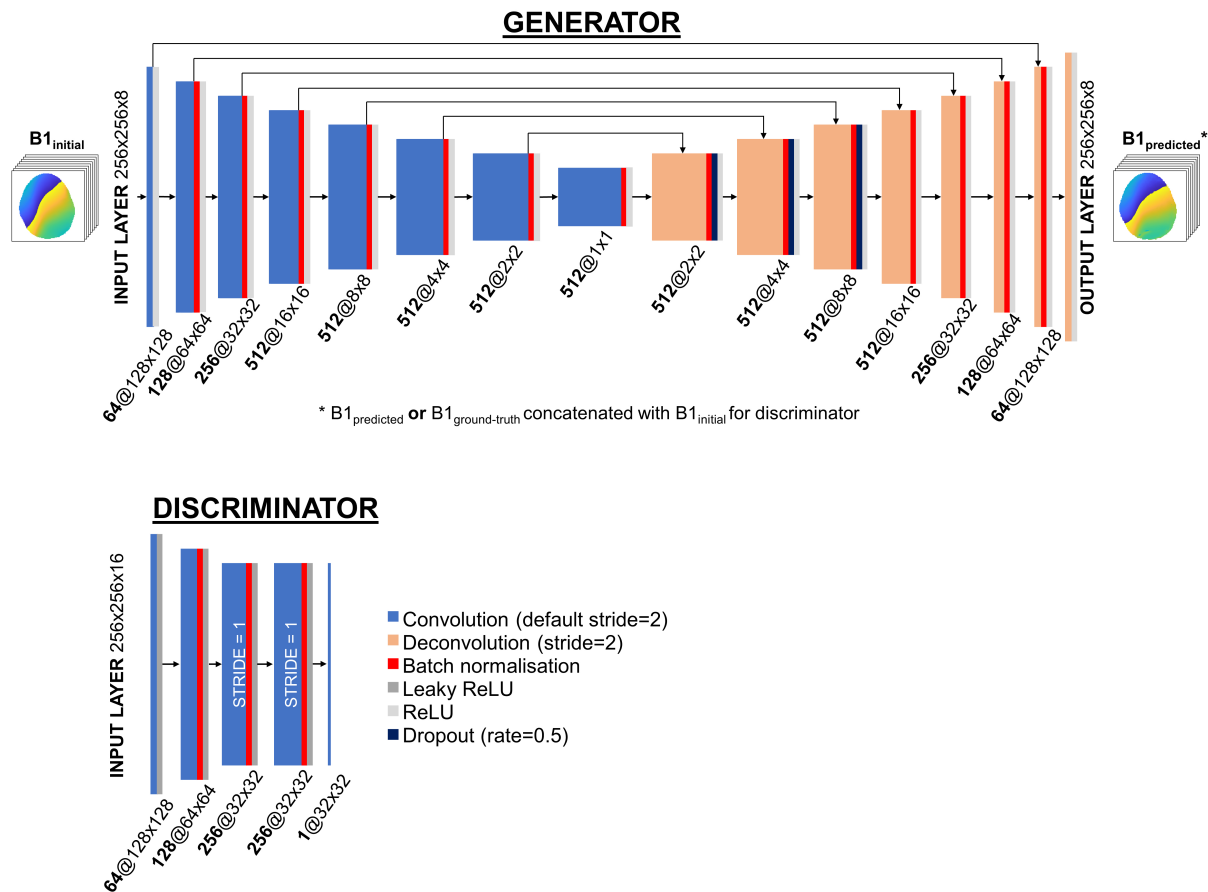


Figure 9.2: *cGAN architecture. Generators were U-Nets with eight convolution and eight deconvolution layers, each with ReLU activation. Discriminators consisted of five convolutional layers with ReLU activation. Square matrix size and number of filters (initially 64 for phase networks) are indicated beneath the layers. Convolution stride was 2 except where specified. Skip connections are shown with arrows. Dropout was applied at indicated layers (dark blue). Batch normalisation (red) was used for phase networks, but not for magnitude networks. Filters for phase networks were 8×8 . Magnitude networks used double the number of filters, with filter size = 4×4 .*

networks. Network weights were saved at the epoch yielding the lowest total error across the validation data set (the Billie data) as a form of early stopping. Networks took approximately 16 hours to train with a batch size of 1 using a standard PC with NVIDIA GeForce GTX 1050 Ti.

9.2.3 Network evaluation and cascading

Networks were tested using the Billie and Dizzy data at six slice locations (Figure 9.1A). For Billie, different slices were used compared with those used for early stopping (and Dizzy was not involved in the training process at all). Like training, testing was conducted for all available examples of each displacement, yielding test data sets of between 6 and 72 slices. In addition to the positions listed in Section 2.1, Billie and Dizzy models were simulated at three combined yaw-rightward positions to test performance for motion involving both rotation and translation.

Because networks were only trained for five displacements but evaluated at 35 positions, networks were cascaded where necessary. Starting with the center position's B_1^+ maps as input, generators were run sequentially, with the output of one generator used as input to the next, until the desired evaluation position was reached. For example, R5 mm, R5 mm, and P2 mm networks were cascaded for evaluation at the (R10, P2 mm) position. Finally, the Billie model was also simulated at inferior 5, 10, and 15 mm to investigate error for through-plane motion.

Predicted B_1^+ maps were exported to MATLAB (The MathWorks, Natick, MA). Voxels with $<1\%$ of the maximum magnitude were smoothed with a Gaussian kernel (standard deviation = 5). Corresponding magnitude and phase network outputs were subsequently combined to form complex predicted maps ($B_{1_{\text{predicted}}}$).

$B_{1_{\text{predicted}}}$ quality was assessed through voxel-wise correlation (using MATLAB's *corrcoef* function) and prediction error between predicted and ground-truth maps at each position. These values were compared with error and correlation following head motion (i.e., between the two simulated maps). Calculations were performed on the tissue-masked region, with the

outermost voxels excluded to avoid artificial amplification of error due to partially filled voxels.

Prediction error for magnitude and phase distributions were assessed through nRMSE and L1 norm, respectively:

$$nRMSE_{||prediction||} = 100 \times \frac{\sqrt{\frac{1}{N_r} \sum_{r=1}^{N_r} ||B1_{gt}(r) - B1_{predicted}(r)||^2}}{\frac{1}{N_r} \sum_{r=1}^{N_r} ||B1_{gt}(r)||} \quad (9.3)$$

$$L1_{\angle prediction} = \frac{1}{N_r} \sum_{r=1}^{N_r} ||\angle \exp^{i(B1_{gt}(r) - B1_{predicted}(r))}|| \quad (9.4)$$

where j is $\sqrt{-1}$; and N_r is the number of voxels in a slice, indexed by r . Motion-induced error was calculated analogously, but substituting $B1_{initial}$ for $B1_{predicted}$ in the equations above.

9.2.4 Pulse design and analysis

Outputs from the R-P grid positions were further processed to assess whether predicted maps were of sufficient quality to be used for tailored pTx pulse design. Five-spoke excitation pulses were designed using a small tip-angle spatial domain method (see Section 3.2) [32, 75, 196] and two approaches were compared in terms of their performance following motion within the R-P grid. A schematic of the process is shown in Figure 9.3. First, a conventional tailored pulse ($pulse_{initial}$) was designed using the subject-specific B_1^+ maps at the initial position ($B1_{initial}$). A uniform magnitude target excitation profile (target flip-angle = 70°) was specified for $pulse_{initial}$. Pulse coefficients were optimised iteratively to minimise magnitude error, whereas the profile's phase was relaxed. The resultant complex profile was used as the target profile for a second pulse ($pulse_{re-designed}$), which was designed without phase relaxation (because magnitude and phase distributions need to be consistent across positions to ensure data consistency following motion which occurs mid-acquisition). $pulse_{re-designed}$ was designed using the network-output $B1_{predicted}$ (the proposed approach). Pulses were approximately 7.7 ms long.

Both pulses ($\text{pulse}_{\text{initial}}$ and $\text{pulse}_{\text{re-designed}}$) were subsequently evaluated using the ground-truth B_1^+ maps at the displaced position ($B_{1,\text{gt}}$) to quantify motion-induced effects on the conventional approach, and improvement provided by the proposed method. Their flip-angle distributions were compared with that of $\text{pulse}_{\text{initial}}$ without motion in terms of nRMSE expressed as percent target flip-angle (and calculated according to Equation 5.1 in Chapter 5).

Peak local SAR (psSAR) of both pulses was also evaluated using the 10-g averaged Q-matrices at each position. Because psSAR sensitivity to motion has been reported to exhibit no clear slice dependence [130], SAR was evaluated at four target imaging slices (out of the six used for pulse design) (Figure 9.1A).

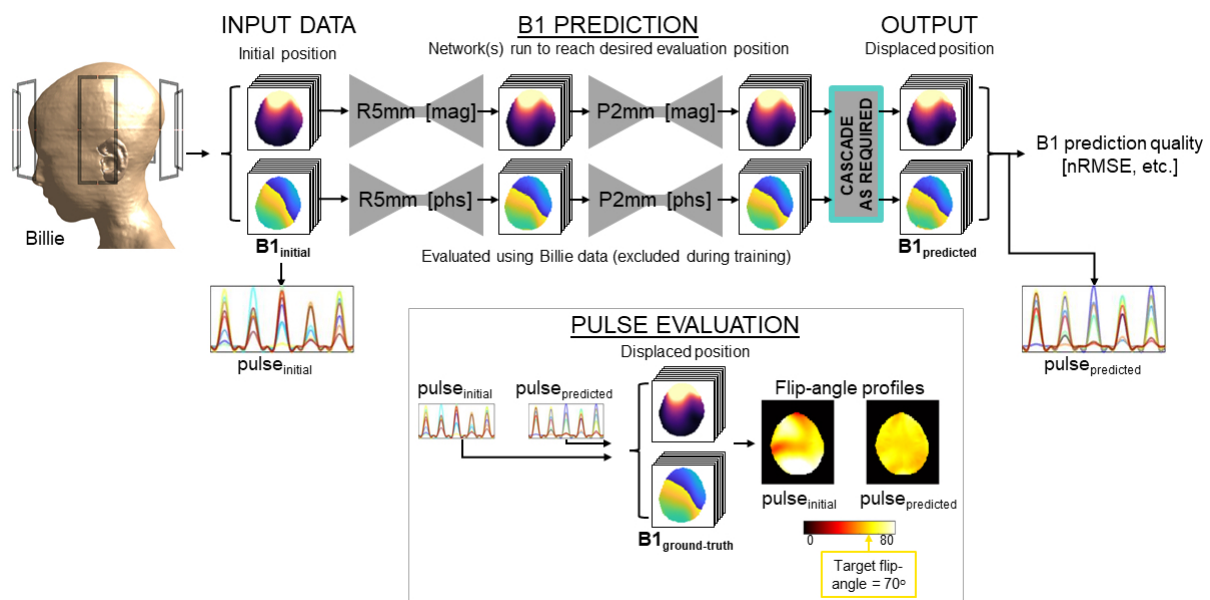


Figure 9.3: Outline of the testing workflow. Simulated B_1^+ maps from the center position are input to the first trained generator. Generators were trained for small displacements but can be run sequentially (cascaded) until the desired off-central (displaced) position is reached in evaluations. Prediction quality is assessed by normalised RMS error (nRMSE) and voxel-wise correlation with respect to the ground-truth (simulation output) displaced B_1^+ map. In addition, pulses designed using the initial B_1^+ map are compared with those designed using predicted maps, in terms of their excitation profiles following head motion.

9.3 Results

9.3.1 B_1^+ prediction quality

B_1^+ maps were predicted by networks in approximately 14 ms using the same PC as used for training. Example B_1^+ maps, motion-induced error, and prediction error are shown for a 5 mm displacement in Figure 9.4. Motion-induced error (averaged across channels) for this example was 15.1% (magnitude) and 4.9° (phase), whereas mean prediction error was 3.2% (magnitude) and 3.5° (phase).

Figure 9.5 shows a summary of error and correlation coefficient for magnitude and phase at each evaluated displacement (averaged across Dizzy and Billie models, slices, channels, and initial positions). Positional dependence of prediction quality was minimal compared with motion-related error, as seen by the reduced gradient with respect to displacement norm in all cases. Billie and Dizzy models yielded very similar prediction quality (see Figures 12.1 and 12.2 in the Appendix).

Mean motion-induced magnitude error scaled linearly with displacement magnitude at about 3% per millimeter (or 3.2% per degree of rotation), compared with 0.36% per millimeter (0.27% per degree) for prediction error. Prediction error was lower than motion-related error in 99.8% of translation, and 90% of rotation evaluations. Figure 9.6A shows B_1^+ nRMSE for magnitude for all slices and channels for 10 example displacements. Motion caused a worst-case magnitude error of 117% following a displacement of R20, P10 mm, whereas maximum prediction error was 33% (at the y15°, R4 mm position). Worst-case prediction error from the R-P grid was lower (20% at the R20, P10 mm position).

Example magnitude correlations are shown in Figure 9.7A. The lowest observed correlation coefficient between $B1_{\text{initial}}$ and $B1_{\text{gt}}$ magnitudes was 0.79 following a y15°, R4 mm displacement. Correlation between $B1_{\text{predicted}}$ and $B1_{\text{gt}}$ did not fall below 0.96.

Motion-induced error and correlation were observed to be slice-dependent and channel-dependent (i.e., the error depended on the displacement relative to each channel's location, as expected). The $B1_{\text{predicted}}$ quality was similar across channels, as demonstrated by the strong correlation across all channels in Figure 9.7A. However, prediction error was

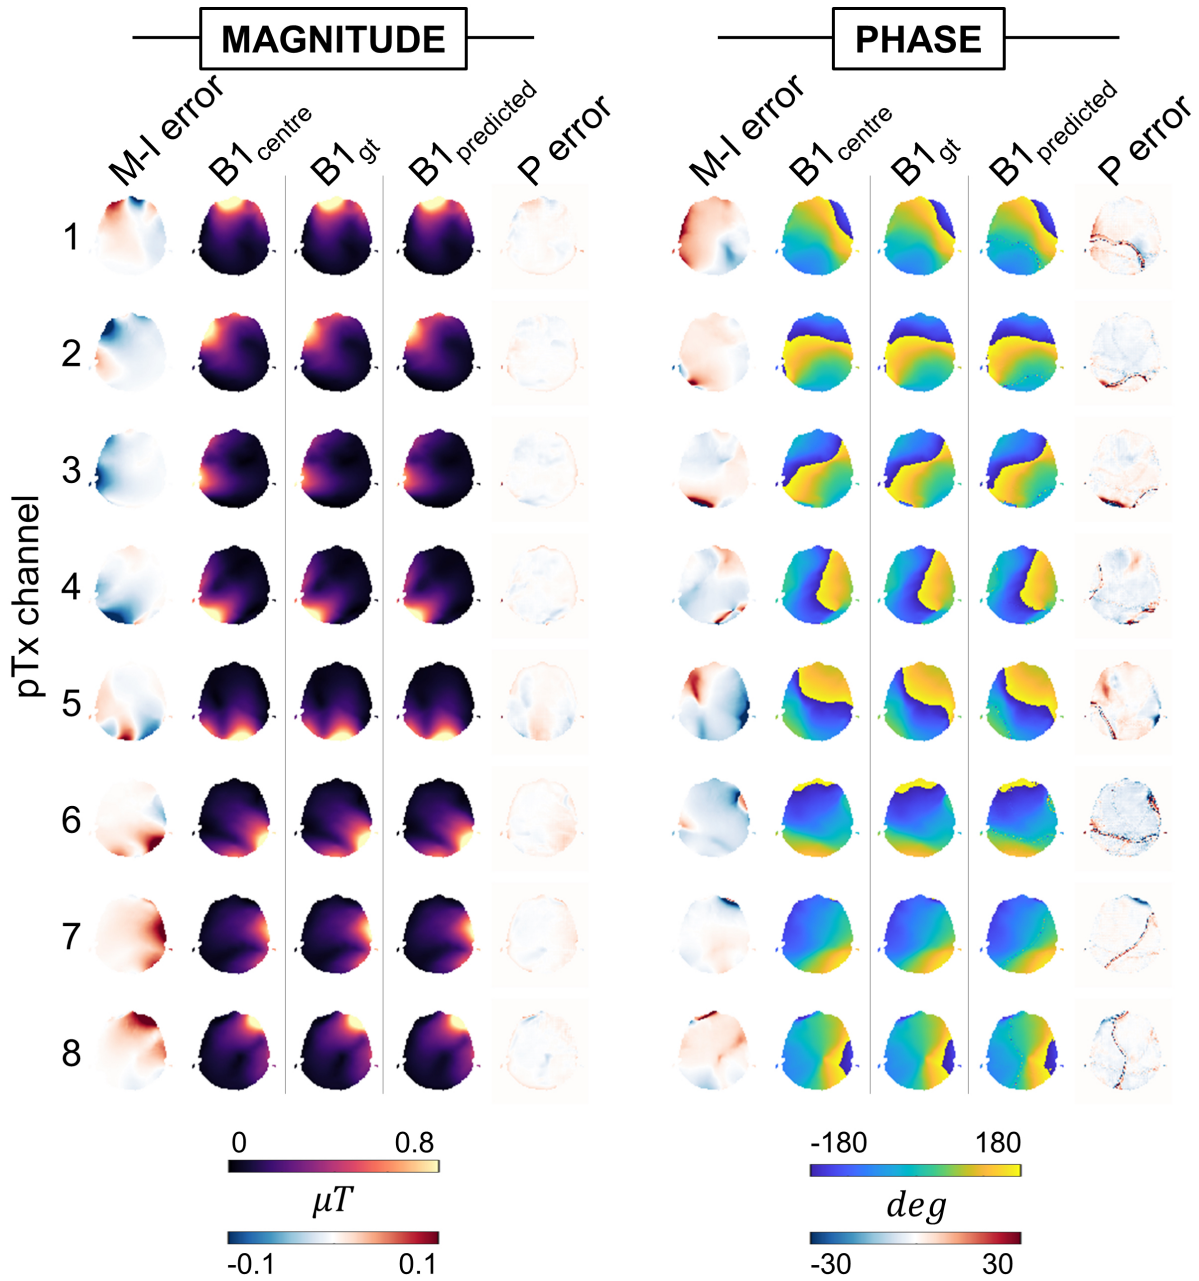


Figure 9.4: Example magnitude and phase B_1^+ maps and error following a rightward displacement of 5 mm (slice location = 2). Motion-induced (M-I) error shows difference between simulation-output B_1^+ at the centered and displaced positions ($B1_{initial}$ and $B1_{gt}$, respectively). Prediction (P) error shows the difference between simulation-output $B1_{gt}$ and generator-predicted B_1^+ ($B1_{predicted}$). Motion-induced error (averaged across channels) for this example was 15.1% (magnitude) and 4.9° (phase), whereas mean prediction error was 3.2% (magnitude) and 3.5° (phase).

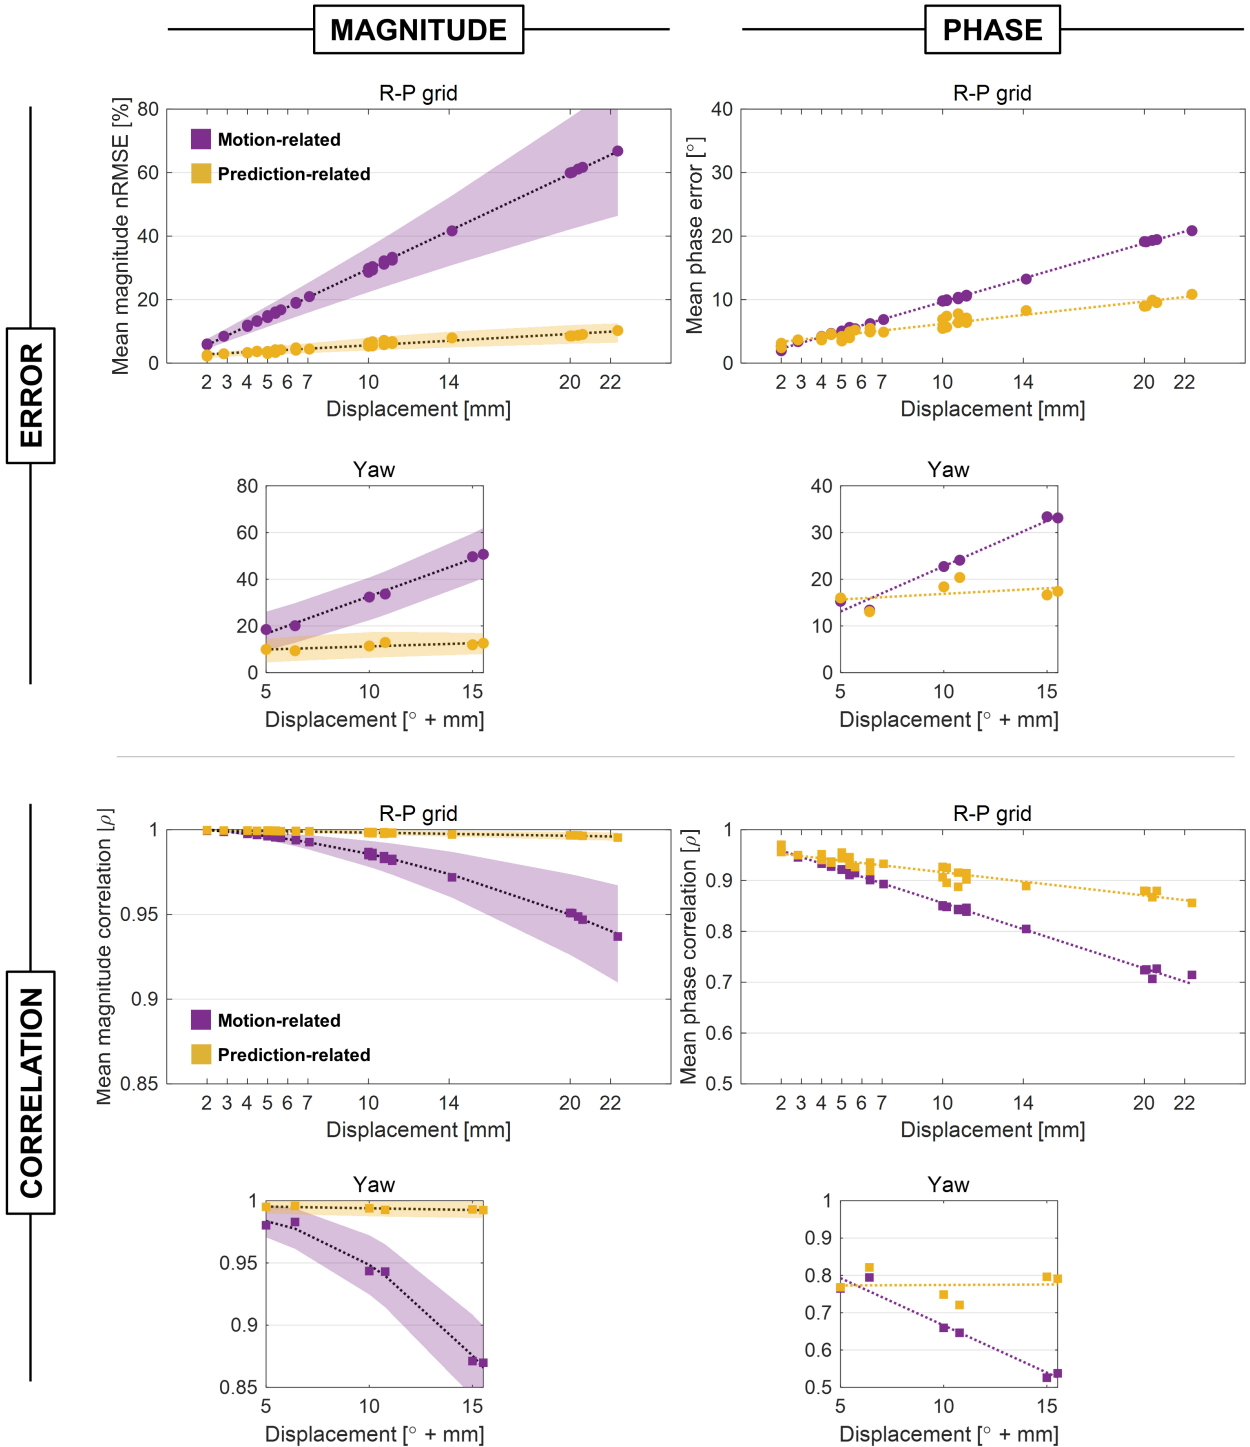


Figure 9.5: Error and correlation coefficient (ρ), averaged over Dizzy and Billie body models, channels, slices, and initial positions for each evaluated displacement. Translations (the R-P grid) are shown in large panels, while rotations (yaw) and combined rotation-translations (yaw plus a 4 mm translation) are shown in smaller panels (for the purpose of the x-axis, the amount of yaw rotation is treated as magnitude displacement; e.g., yaw 5° plus 4mm translation is shown at x=6). Standard deviation is shown as shaded regions for magnitude but omitted for phase for clarity.

B₁ ERROR

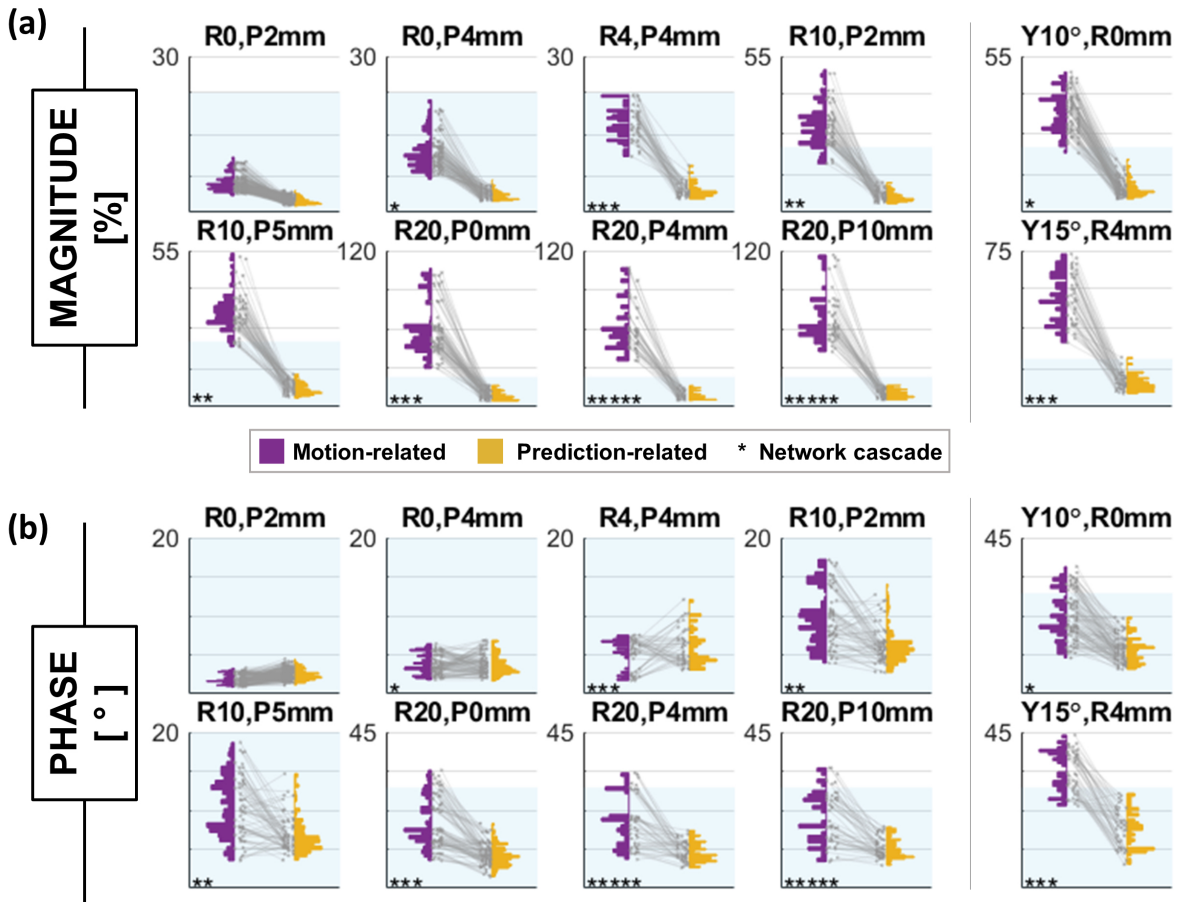


Figure 9.6: B_1^+ error ($nRMSE$ for magnitude **(a)**, L_1 norm for phase **(b)**) for all evaluations with the Billie model following 10 example displacements. Motion-related error is shown in purple, while error for predicted maps is in yellow. Asterisks indicate the number of network cascades required for evaluation. The blue-shaded region shows the maximum observed prediction error across all 35 displacements for the Billie model (consistent across panels).

B₁ CORRELATIONS

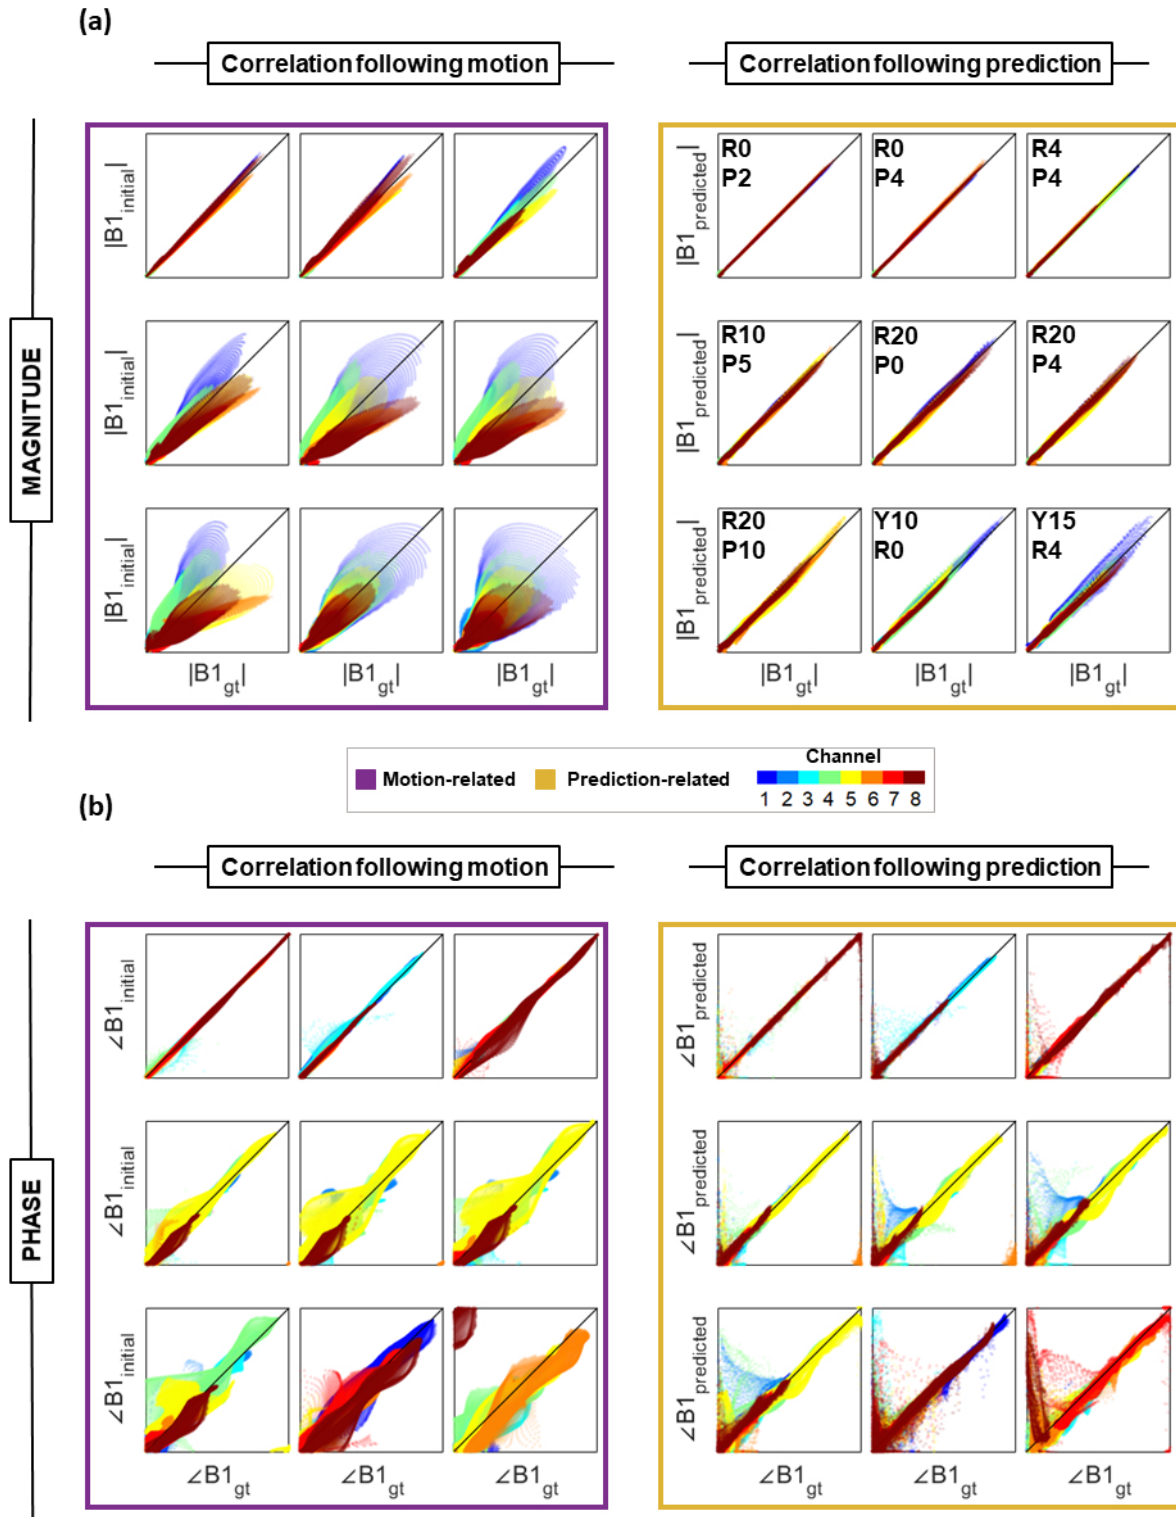


Figure 9.7: Example voxel-wise correlations between $B1_{initial}$ and $B1_{gt}$ (left) and $B1_{predicted}$ and $B1_{gt}$ (right) for nine example displacements. The pTx channels are indicated by color. The x and y axes range between 0 and $3 \mu T$ for magnitude (a), and 0 and 2π for phase (b).

somewhat slice-dependent, with the most inferior slice locations yielding the highest prediction errors (slice information not shown).

Phase error for 10 example displacements are shown in Figure 9.6B. For phase, maximum observed prediction error (57°) was similar to maximum motion-related error (55°). These worst cases arose in the Dizzy model; for the Billie model, maximum prediction error (29.4°) was lower than that caused by motion (44.2°). Furthermore, prediction error was lower than motion-induced error for phase in 68% and 66% of translation and rotation evaluations, respectively (across both models).

Yaw rotation caused substantially higher error than axial translations; for the R-P grid, maximum prediction-related and motion-related errors were 19.8° and 34.7° , respectively. Mean phase prediction error was less position-dependent than motion-related error, with axial translations causing error of approximately 0.9° per millimeter displacement, compared with 0.4° per millimeter in predicted maps. For displacements including rotation, analogous gradients were 2° and 0.2° per degree of yaw, respectively.

Mean phase correlation coefficient between predicted and ground-truth maps was higher than (or very similar to) that between initial and ground-truth maps for all displacements. Phase correlation examples are shown in Figure 9.7B. Correlation coefficient between $B1_{\text{predicted}}$ and $B1_{\text{gt}}$ exceeded that between $B1_{\text{predicted}}$ and $B1_{\text{initial}}$ in 69% of cases.

9.3.2 Parallel-transmit pulse performance

Subsequent analyses were conducted using the Billie model with the R-P grid data only. Five-spokes pTx pulses designed using $B1_{\text{initial}}$ ($\text{pulse}_{\text{initial}}$) yielded uniform flip-angle profiles (mean nRMSE $\sim 1\%$) without motion. However as expected, uniformity was lost following axial translation.

Figure 9.8 shows that flip-angle nRMSE for $\text{pulse}_{\text{initial}}$ was strongly position-dependent, reaching a maximum of 76% following a displacement of R20, P5 mm. Conversely, pulses re-designed using $B1_{\text{predicted}}$ ($\text{pulse}_{\text{re-designed}}$) produced much improved flip-angle profiles when evaluated at the displaced position, yielding nRMSE of 14% for the same displacement. Maximum nRMSE for $\text{pulse}_{\text{re-designed}}$ was 17% (at the R2, P10 mm position),

whereas this error value was exceeded by $\text{pulse}_{\text{initial}}$ (i.e., without any correction) after displacements of just ≥ 5 mm. The largest errors occurred in inferior slice locations for both pulses (slice information not shown). Maximum motion-related error in the excitation profile's phase (110.4°) was reduced by 7.8° when using $\text{pulse}_{\text{re-designed}}$.

Figure 9.8b shows flip-angle profiles for both pulses following several example displacements. Figure 12.3A in the Appendix also shows flip-angle nRMSE for nine example displacements. It should be noted that flip-angle uniformity for $\text{pulse}_{\text{re-designed}}$ could be further improved by including phase relaxation in the design (as was done for $\text{pulse}_{\text{initial}}$); however, this would permit excitation phase to vary throughout the scan, causing reconstruction inconsistencies.

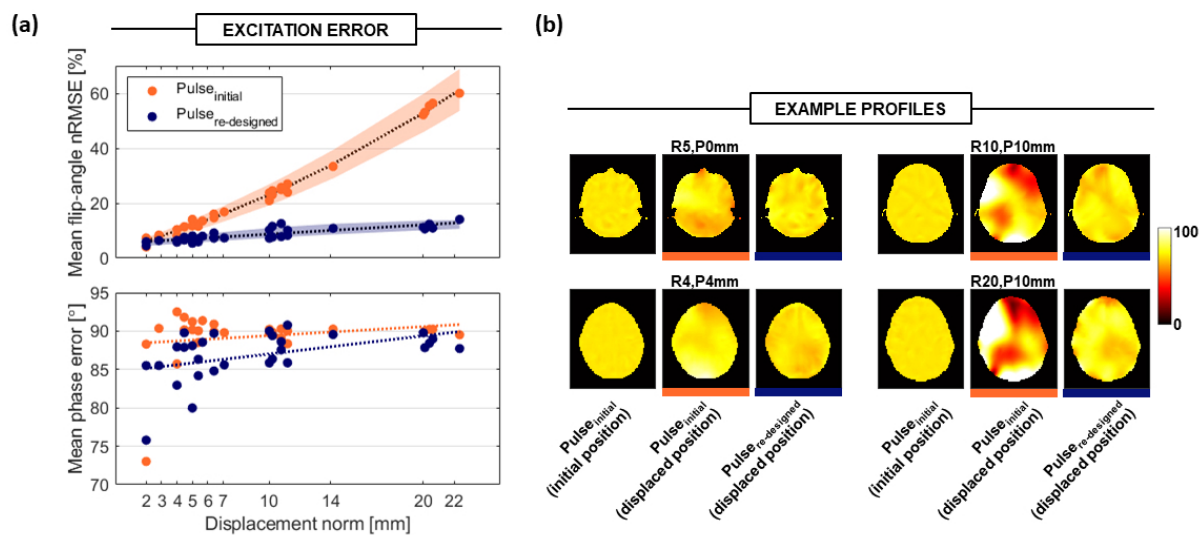


Figure 9.8: Excitation profile results for five-spoke pTx pulses following head motion. **(a)** Mean flip-angle nRMSE (above) and phase RMSE (below) for excitation profiles, averaged over slices and initial positions for each evaluated displacement. Excitation pulses were five-spoke pTx pulses designed using either the initial position ($\text{pulse}_{\text{initial}}$) or predicted ($\text{pulse}_{\text{re-designed}}$) B_1^+ maps. Standard deviation is shown as shaded regions for magnitude but is omitted for phase for clarity, as values were similar. **(b)** Example flip-angle profiles produced by $\text{pulse}_{\text{initial}}$ at the initial position, by $\text{pulse}_{\text{initial}}$ at the displaced position, and by $\text{pulse}_{\text{re-designed}}$ at the displaced position.

9.3.3 Cascading

The $B1_{\text{predicted}}$ quality remained high when networks were cascaded multiple times; however, there was a weak linear relationship between prediction error and motion magnitude. To investigate the impact of cascading on prediction quality, we ran secondary analyses for displacements of R0, P10 mm, R-2, P10 mm, and R-5, P10 mm using only the P2 mm network for the posterior component. Running the 2 mm network five times (i.e., four cascades) led to approximate average increases in magnitude and phase error of 1.2% and 1.2°, respectively, compared with running the 5 mm network twice (one cascade). There was also reduced flip-angle uniformity compared with using the 5 mm network. Nevertheless, Figure 9.9 shows that motion-induced error was appreciably reduced using either approach.

9.3.4 Specific absorption rate

In addition to flip-angle, SAR was also evaluated for the R-P grid positions. Following motion, psSAR produced by pulse_{re-designed} was lower than that of pulse_{initial} in 89% of cases. For pulse_{initial}, motion caused psSAR to increase (relative to that without motion; psSAR_{center}) in 72% of evaluations. When pulse_{re-designed} was used instead, psSAR increased relative to the centered case in only 16% of the cases. The psSAR for nine example displacements is shown in Figure 12.3B in the Appendix.

Figure 9.10A shows worst-case psSAR for each evaluated position (relative to psSAR_{center}). Figure 9.10B shows the same information, separated by slice location. Motion-related SAR change was similar across slices, whereas benefits offered by pulse_{re-designed} were most consistently seen in slices 1 and 6 (furthest from mid-axial locations). pulse_{re-designed} yielded lower psSAR than pulse_{initial} following large displacements, but sometimes resulted in higher psSAR for small displacements, especially in slice 4 (mid-axial). In the worst observed case, pulse_{re-designed} yielded a 3.1-fold increase in psSAR (slice 4 at the posterior 5-mm position), whereas pulse_{initial} increased by a maximum of 3.3-fold (following the largest axial displacement). It should be noted that SAR was not used as a constraint in either pulse's design.

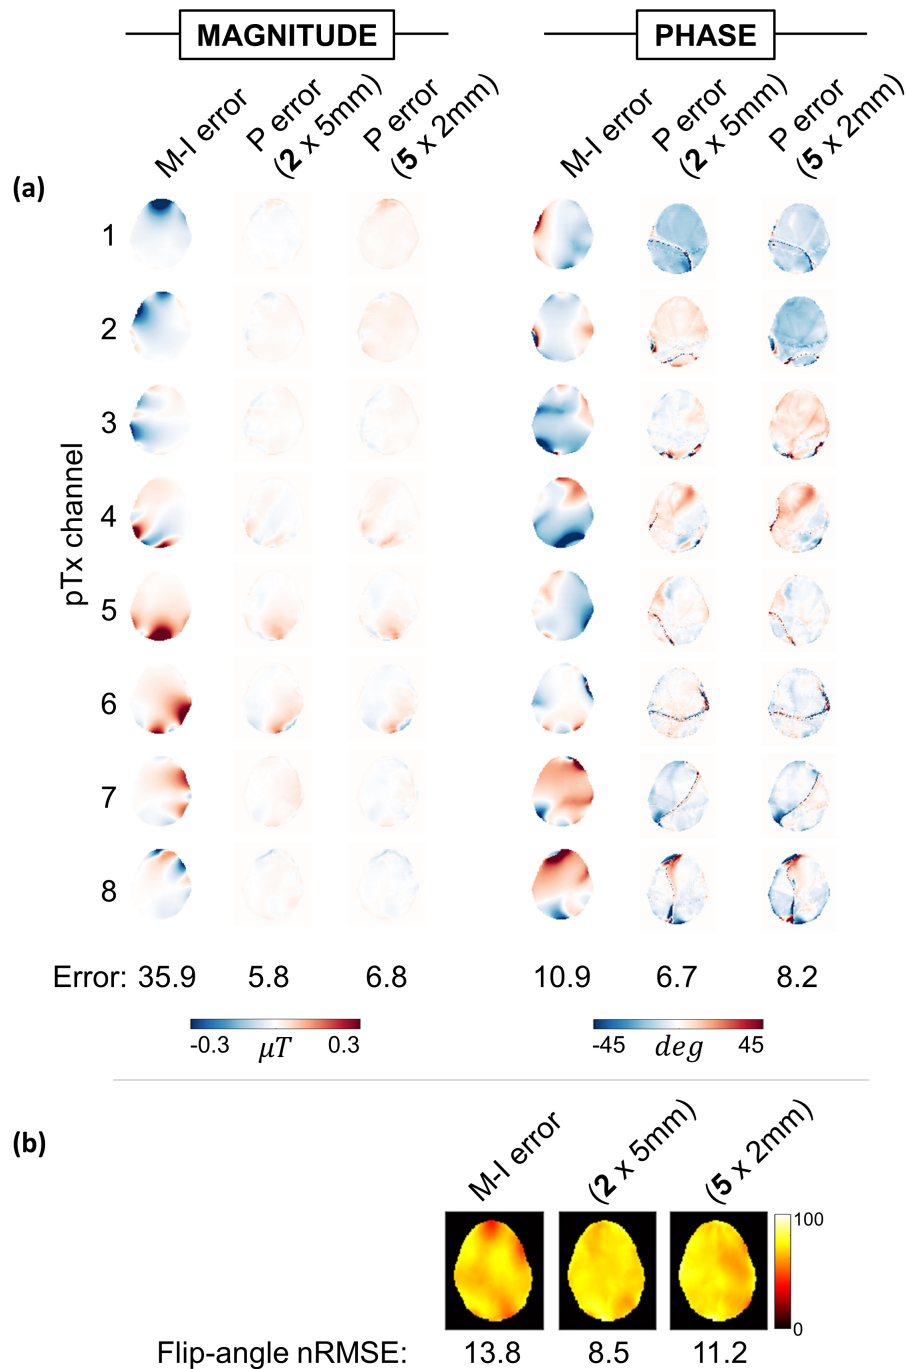


Figure 9.9: Effect of cascading the P2 mm network four times compared with cascading the P5 mm network once for evaluation at the R5, P10 mm position (along with the R5 mm network for the rightward component). **(a)** Example motion-induced (M-I) error and prediction (P) error for both cascade approaches for magnitude (left) and phase (right). Error shown below maps is nRMSE (%) for magnitude and L1 norm (°) for phase, both averaged over channels. **(b)** Comparison of flip-angle profiles and nRMSE for pulses designed using initial (left) and predicted maps using both cascade regimes (center and middle). Target flip-angle is 70°.

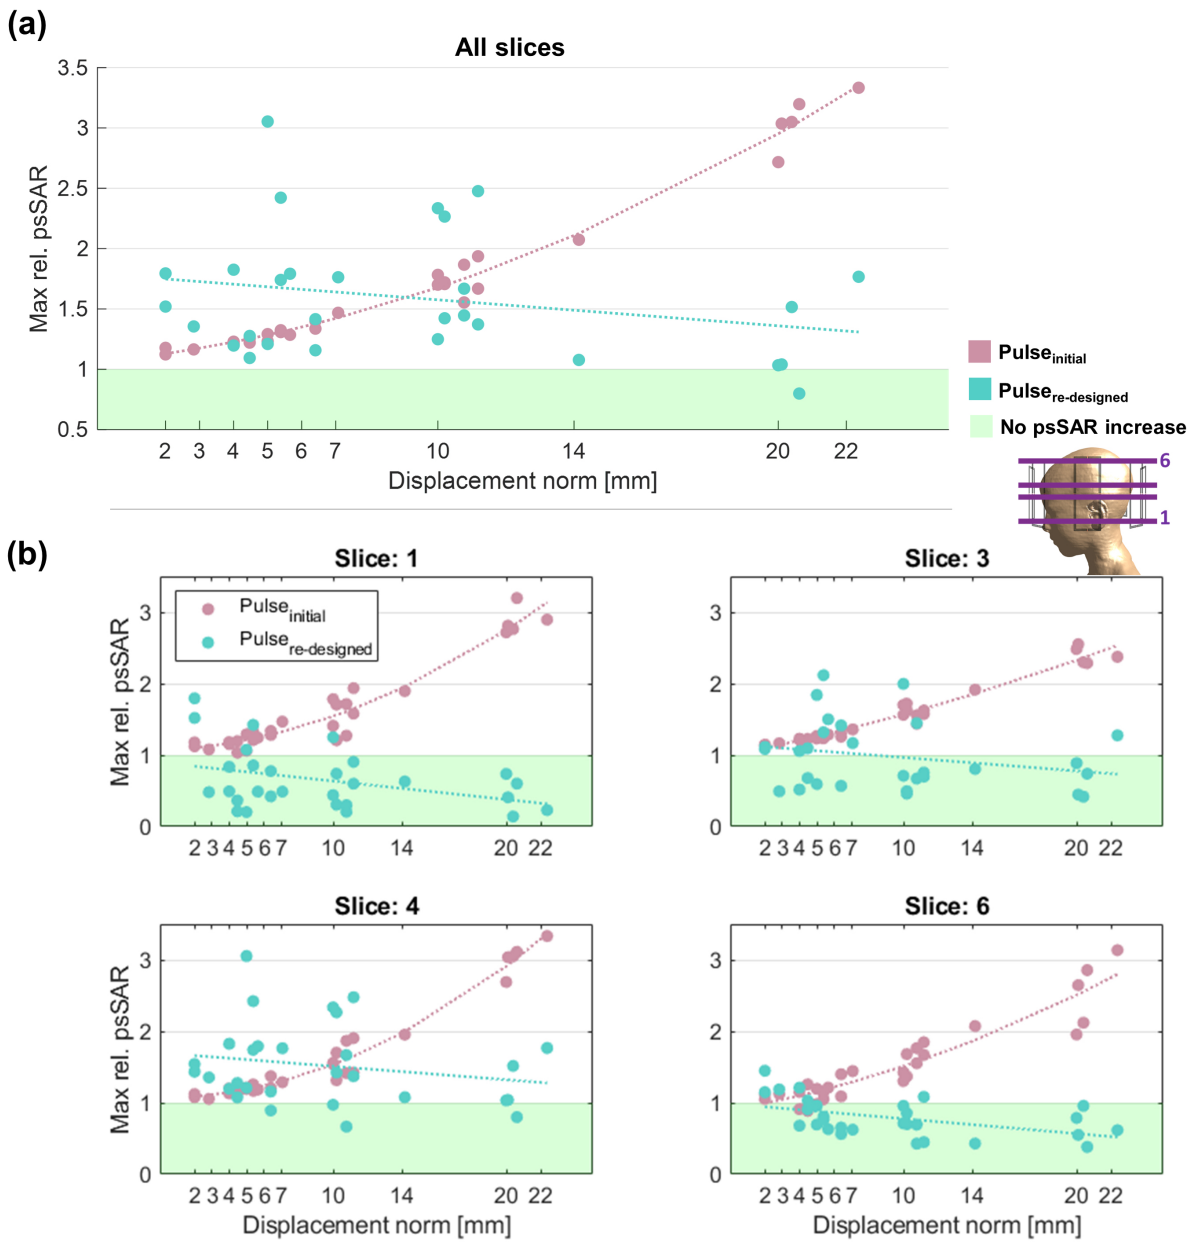


Figure 9.10: Peak 10g-averaged local SAR ($psSAR$) for $pulse_{initial}$ and $pulse_{re-designed}$ following motion. **(a)** Worst-case $psSAR$ for both pulses across all observed slice locations. **(b)** Slice-wise worst-case $psSAR$ for both pulses. Slice locations are indicated in the inset on the right. Vertical axes show relative $psSAR$, calculated as $psSAR$ as a factor of that without motion using $pulse_{initial}$ ($psSAR_{center}$). The green-shaded region shows $psSAR$ at or below $psSAR_{center}$ (i.e., indicating that $psSAR$ did not increase following motion). Neither pulse included SAR constraint in the design.

9.4 Discussion

To the best of our knowledge, this study is the first to demonstrate motion-resolved B_1^+ map prediction in pTx. We successfully trained a system of deep neural networks to estimate B_1^+ sensitivity distributions following simulated in-plane head motion. Predicted B_1^+ maps were of sufficiently high quality to be used for tailored pTx pulse design, and resulting pulses outperformed conventionally designed tailored pulses under conditions of head motion.

Across the four magnitude networks used for the R-P grid positions, prediction error was $4.5\% \pm 1.5\%$ (mean \pm SD) of the ground-truth B_1^+ magnitude (tested individually without cascading and error calculated according to the definition in Abbasi-Rad et al. [221]). This is in line with expectations based on previous work, in which magnitude B_1^+ maps were predicted with mean error of $9.5 \pm 7.1\%$ [221]. The experimental and 3-D nature of data in Abbasi-Rad et al. may have caused the slightly higher error. We also observe similar $B_{1_{\text{predicted}}}$ correlation coefficients to those reported in Wu et al. (~ 0.99) [220], which was retained even when our networks were cascaded multiple times. Yaw rotation caused higher error than translational motion, and prediction error was generally higher than that of translations. We attribute this to the much smaller training database used to train yaw networks.

It was previously reported that excitation profile nRMSE increases by approximately 2.4% per millimeter of axial displacement in simulations [173], and our results (2.7% per millimeter) are in good agreement. A 12% – 22% flip-angle error was observed in-vivo following motion in the range of about 5 – 16 mm in Kopanoglu et al. [173], and we observed an error of about 11% – 35% for a similar range of motion. Error was reduced to approximately 8% – 10% using the proposed approach in our study. For larger movements, the benefit of pulse re-design using predicted maps was even greater.

There was some slice dependence for both $B_{1_{\text{predicted}}}$ quality and flip-angle error, with the highest errors observed for inferior slice locations. This is in line with previous research, in which higher motion sensitivity was observed for inferior slices passing through the temporal lobes and the cerebellum [173]. Inferior slices yield lower field magnitude, and B_1^+ prediction quality was lower in regions with very low field magnitude, which may explain the residual

slice dependence in predicted maps. An alternative explanation is that there were fewer inferior slices in the training data set compared with mid-axial slices, which could have resulted in better training for mid-axial slices.

Considering that the R5 mm training data included positions up to just 10 mm along the rightward axis, it is noteworthy that the R5 mm network was able to extrapolate beyond this by successfully cascading four times to estimate the fields at the R20 mm position. We attribute this to the global normalisation, conducted over all data sets. Results for large displacements could likely be further improved by including extreme positions in the training data set.

Magnitude networks consistently outperformed those of phase. Prediction error exceeded motion-induced error for phase in about a third of evaluations. Improvements to excitation phase were evident but modest. It has been previously acknowledged that phase changes due to motion are difficult to model, and other work on motion-related field changes [181], including B_1^+ prediction [221], often neglect phase altogether. Most of the phase-prediction error occurred at phase-wrapping boundaries. This was somewhat reduced by applying random phase offsets to pairs of input B_1^+ maps, but it was not eliminated. In terms of pulse performance, these small, local errors are likely to have less impact than the global changes caused by motion, which could potentially accumulate when channels are superposed. Incidentally, when we instead trained networks using unwrapped phase data, the error in $B_{1,\text{predicted}}$ was globally higher and yielded inferior results compared with wrapped data. In contrast, the local prediction errors seen with wrapped data are not structured. We believe that the increase in error seen for unwrapped data was due to the increased dynamic range of unwrapped data, meaning that relative changes due to motion were smaller following normalisation. Similarly, when magnitude and phase networks were substituted for real and imaginary networks, B_1^+ prediction quality reduced.

One limitation of this study is that it deals with simulated data only. This is because of the lack of models for the RF coils at the institution (due to proprietary information). However, in DeepQSM [219], training data were solely synthetic; the ground truth consisted of overlapping cubes and spheres with known susceptibility values. This was convolved with a forward dipole kernel to create the corresponding input. Networks were able to resolve high-quality susceptibility maps for human brains, despite only being trained on simple

geometric shapes. Similarly, Meliado et al. demonstrated successful in-vivo validation following training with synthetic data [132]. We believe that the method proposed here would also be generalisable if a realistic RF coil model (i.e., a model of the coil to be used) is used for simulations. Moreover, using simulated training data avoids the requirement for choreographed in-vivo head movement to be replicated precisely across several subjects to create the training data sets, which would be practically infeasible. Aside from the initial (measured) input B_1^+ map, networks output maps in patient coordinates, making online registration unnecessary. However, motion tracking is needed to determine which network(s) are required, and for online corrections to gradient waveforms to update the imaging volume, as in Zaitsev et al. [164].

The minimum displacement we consider here is 2 mm. Smaller movements could remain problematic for quantitative MR protocols that rely on signal changes on the scale of 1% – 4% of the total signal [170]. Finer discretisation is possible through simulations of smaller displacements. The most appropriate motion discretisation will depend on the user’s primary aims. While using finer discretisation and cascading more (i.e., 5×2 mm vs 2×5 mm) did result in slightly higher error here, motion-induced error was still largely ameliorated.

In this study we considered only positive displacements (e.g., rightward, but not leftward). Because the trained networks are direction-specific, the same networks cannot handle negative displacements. This is due to the nature of inverting a cAGN, which is not possible, as reversing convolutional layers results in a highly underdetermined problem. However, the training of networks for positive displacements does not limit generalisability, and additional networks can similarly be trained for negative displacements. In fact, by training a network for positive and negative displacements for each of the 6 degrees of freedom of motion, all rigid-body motion could theoretically be covered.

We focused on in-plane motion (axial translations and yaw rotation), as motion-induced error was observed to be more spatially varying, which was also in line with findings presented in Chapters 5 and 8. Although some motion-induced B_1^+ error was observed in peripheral slices following inferior translation, the error was relatively low, and importantly, more global (i.e., spatially smoothly varying) within the slice (see Figure 12.4 in the Appendix). We attribute this to the fact that relative tissue-channel distances remain constant for through-slice

translation. This means that simpler correction methods (e.g., slice-dependent pulse scaling) could feasibly be used to counteract the B_1^+ effects of through-plane motion.

In contrast to B_1^+ , through-plane motion was shown to be more disruptive to B_0 than within-plane motion [186, 224]. Although B_0 off-resonance can be incorporated in tailored pulse design, this cannot currently be updated in real time with pTx, as it increases the degrees of freedom to be optimised in the pulse, pushing re-design times beyond practically feasible TR values [189]. Instead, motion-related effects on B_0 can be corrected retrospectively, such as using data-driven coefficients to link motion with field changes [186]. Alternatively, real-time B_0 shimming may be possible with multi-coil shim arrays by predicting B_0 field changes due to motion in a manner similar to the method proposed here.

The SAR observations reported here are incidental. The focal point of this study was to develop a method to accurately estimate B_1^+ maps following motion. Using SAR as a pulse design constraint would trade flip-angle homogeneity for reduced SAR, thereby overshadowing B_1^+ quality. Hence, SAR was not used as a pulse design constraint. Although motion-related SAR increase was generally lower for pulses re-designed with predicted B_1^+ maps, it was higher for a minority of cases. Motion sensitivity of SAR in pTx has previously been reported to be similar across axial slices [130], and we also did not observe clear slice dependence for SAR motion sensitivity. However, we did observe that inferior-most and superior-most slices benefited from the proposed approach more consistently than mid-axial slices. Nevertheless, the overall improvement offered by pulse_{re-designed}, especially for larger displacements, is promising for future development of this approach. Neural networks have previously been used to predict B_1^+ maps for the purpose of SAR reduction [221]. This was achieved through slice-wise pulse scaling based on a predicted 3-D B_1^+ magnitude. Here, the entire 9 cm axial slab could be predicted within approximately 0.8 seconds, permitting concatenation into pseudo-3-D B_1^+ maps that could feasibly be used for pulse scaling or similar SAR management. Pulse scaling based on B_1^+ cannot guarantee SAR compliance, as B_1^+ does not necessarily reflect electric-field distributions. Compared to B_1^+ , E-fields are more sensitive to different anatomies; Alon et al. found the change in E-field to be 10 times that of B_1^+ for the same anatomical variation [225]. A similar finding could be expected in cases of motion. However, SAR compliance could be ensured if (3-D) electric fields were also predicted.

9.4.1 Conclusions

We have demonstrated a framework for a deep-learning approach for motion-resolved B_1^+ estimation in pTx. Estimated maps can be used for real-time tailored pulse re-design, yielding homogeneous flip-angle profiles in cases of head motion. If real-time pulse re-design is not possible, an alternative application of the method proposed here is to predict tailored B_1^+ maps at several positions to be used in the composite B_1^+ map approach described in Chapter 8 of this thesis. Importantly, networks can be run sequentially to predict B_1^+ maps following arbitrary displacements comprising multiple directions. Here, error was reduced for 35 displacements using networks trained for just five displacements. Our findings represent one potential avenue toward user-friendly, optimised pTx at 7T.

Chapter 10

Network generalisability across coils

The work presented in this chapter led to the following conference proceedings:

Plumley, A. et al. Inter-coil applicability of trained neural networks for B_1^+ prediction in parallel-transmit. Proceedings of the International Society for Magnetic Resonance in Medicine, #2869 (2022).

10.1 Introduction

In Chapter 9, a novel approach for prospective correction of motion-related B_1^+ changes in pTx was described. Conditional generative adversarial networks (cGANs) were trained to predict B_1^+ distributions following head motion, given an initial B_1^+ map (without motion) as input. Predicted maps were used for the purpose of tailored pTx pulse re-design. The fidelity of the flip-angle profiles produced by standard pulses (which are designed using the initial B_1^+ map alone) deteriorated when motion occurred, while re-designed pulses produced relatively uniform profiles across all evaluated positions.

As mentioned in Chapter 9, cGANs and U-Nets (which are essentially cGANs without the additional discriminator network used during training) have been used for various applications within MRI. Mandija et al. reconstructed conductivity and permittivity maps using B_1^+ maps as input to cGANs [205]. Local SAR distributions in the pelvis were

estimated using cGANs with B_1^+ maps as input by Meliado et al. [202]. Yoon et al. [218] and Bollmann et al. [219] reconstructed quantitative susceptibility maps (QSM) using U-Nets with local field maps (derived from phase images) as input. For the latter, training samples consisted of small patches. Patch-based approaches tend to generalise well across space and subjects, since they are forced to rely on local features in the input, which could arise in any combination when performing network inference. In Bollmann et al., networks trained using patches of purely simulated data containing simple geometric shapes were able to generalise to in-vivo images without modification, which was attributed to the network learning the dipole inversion underlying susceptibility contrast [219]. Nevertheless, when the EM regime changes (for example, at a different field strength or using a different coil model), the relationship between input and target also changes. It is therefore unclear whether networks such as these are regime specific. For their B_1^+ -to-SAR cGAN, Meliado et al. highlight that the learned SAR model is likely to be coil-specific [202].

It is well known that field distributions depend on the coil design [127, 195], and it was further shown in Chapter 6 that the response of SAR distributions to motion was also coil dependent, with motion-induced SAR changes generally being higher for coil models with smaller loop elements and / or a smaller array radius. With this in mind, it is unclear whether our networks trained for B_1^+ prediction using data from one pTx coil model can generalise well to other coil models, or whether coil-specific networks are required for accurate B_1^+ prediction. In the former case, adoption of the proposed method is relatively simple, since only pre-trained networks are required for implementation. On the other hand, if the latter is true, implementation of the method may instead require large coil-specific training datasets, involving EM simulations for network training.

In this chapter, networks are trained for motion-resolved B_1^+ prediction using data simulated from two virtual body models with one pTx coil model, as was done in Chapter 9. To test inter-coil generalisability, the trained networks are subsequently evaluated without further modification using a third body model simulated with the same coil, as well as five pTx coils with differing dimensions. We compare performance across coil models by assessing error in network-output predicted maps and compare this to motion-related error for each coil. To assess the practical benefits provided by networks, like in Chapter 9, excitation pulses are designed using the initial B_1^+ map (standard approach; no motion correction) and their

flip-angle performance is compared with pulses redesigned using predicted maps (proposed approach) when both pulses are evaluated on the ground-truth (post-motion) B_1^+ maps.

10.2 Methods

Billie, Duke, and Ella of the Virtual Population were simulated using Sim4Life (ZMT, Zurich) at 30 different in-plane positions (rightwards (R) 0, 2, 4, 5, 10 and 20 mm, and posterior (P) 0, 2, 4, 5, and 10 mm) as previously described. All three body models were simulated with an 8-channel pTx coil with loop dimensions 40 mm x 110 mm, and array radius of 115 mm. Consistent with Chapter 6, this coil is referred to as coil model C. Data from Billie and Duke with coil C were used to train networks, while the data from Ella were used for validation and testing. The Ella model was additionally simulated with 5 different coil models which ranged in loop height and width (and / or array radius) by $\pm 25\%$ of coil C (these were the same as the coils used in Chapter 6). Training data consisted of 51 axial slices (covering a mid-axial slab) per model per position, while test data from Ella (with all coil models) was acquired at 6 slice locations spanning the same axial slab (Figure 10.1). Different slices were used for validation vs. testing, and validation was conducted with coil C only (while testing was done with all coils).

cGANs were trained for B_1^+ prediction following motion occurring within the axial plane. Separate networks were trained for rightward (R) and posterior (P) displacements of 2 and 5 mm. A total of 8 networks (4 magnitude, 4 phase) were trained. As in Chapter 9, networks were cascaded for evaluation at positions reflecting larger displacements (explained in Section 9.2.3). Networks took ~ 6 hours to train using a DGX-1 NVIDIA GPU. Other network details followed those described in Chapter 9.

The difference was calculated between firstly, the centre and displaced simulated maps (to quantify the impact of motion), and secondly, the network-predicted and displaced simulated maps (to quantify prediction quality). Note that the 'centre' map in this context was not necessarily the R0, P0 position; rather, it was the position occupied before a given relative displacement, as previously explained in Chapter 9. Performance metrics (nRMSE for magnitude and L1 loss for phase) were subsequently calculated from these differences using

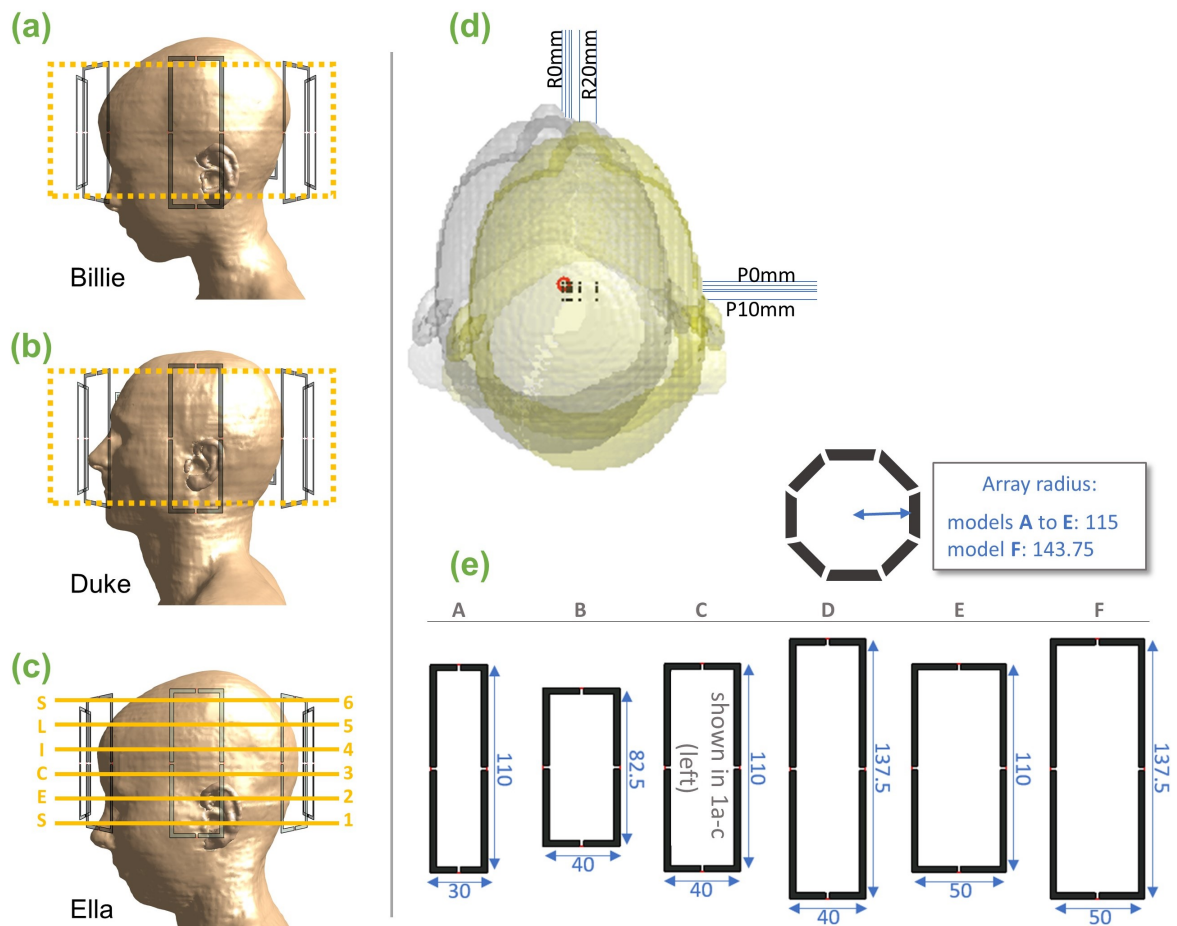


Figure 10.1: Slices from a slab (shown in yellow) from Billie (a) and Duke (b) body models were used to train networks with one coil model (model C). Networks were tested using the Ella model (c) at the 6 indicated slice locations. Testing was conducted with the differently sized coil models shown in (e). (d) The 30 positions simulated for each model. Positions covered an axial grid at 0, 2, 5, 10, 20mm right and 0, 2, 5, 10mm posterior. (e) Coil models A to F. One loop from each pTx array is shown (dimensions shown in mm).

equations 9.3 and 9.4. Motion- and prediction-related errors were compared to assess the benefits of B_1^+ prediction. Our primary interest was prediction quality across coil models compared to that for coil C (which networks were trained on). Pair-wise one-way ANOVAs between coils were conducted to assist this comparison (with corrections for multiple comparisons).

5-spokes pTx pulses were also designed and evaluated using the approach described in Section 9.2.4, with $\text{pulse}_{\text{initial}}$ and $\text{pulse}_{\text{re-designed}}$ being pulses designed using the initial and predicted B_1^+ maps, respectively. Flip-angle nRMSE was calculated to compare their performance and quantify practical benefit provided by the approach for each coil model.

10.3 Results

Figure 10.2 shows B_1^+ maps for 2 example displacements. Motion-related error was qualitatively similar across coils A to E, while prediction error appeared more variable. In both cases, the largest magnitude errors generally occurred adjacent to the nearest coil element, where field magnitude is highest. On the other hand, phase errors were highest for the phase-wrap boundary and at locations where phase is 0.

The distribution of error across all evaluations (and all voxels therein) for each coil are shown in Figure 10.3. For magnitude, mean motion-related nRMSE decreased monotonically with increasing coil size (coil A = 31.6%; coil F = 18.2%). Conversely, a U-shaped pattern emerged for prediction nRMSE, with coil C (the coil used to train networks) yielding significantly lower prediction nRMSE (mean = 6.4%) than all other coils (between 8.4 and 9.7%). Similarly, the maximum-observed prediction nRMSE was just 33% for coil C, compared to 63% and 56% for coils F and A, respectively.

Despite these systematic differences in performance, magnitude prediction nRMSE was significantly lower than motion-related nRMSE for all coils, which is evident from the downward shift in all distributions from motion-related to prediction nRMSE.

nRMSE expresses mean error as a proportion of the field magnitude, however if this normalisation step is omitted, worst-case prediction error for coil F appears to be much lower

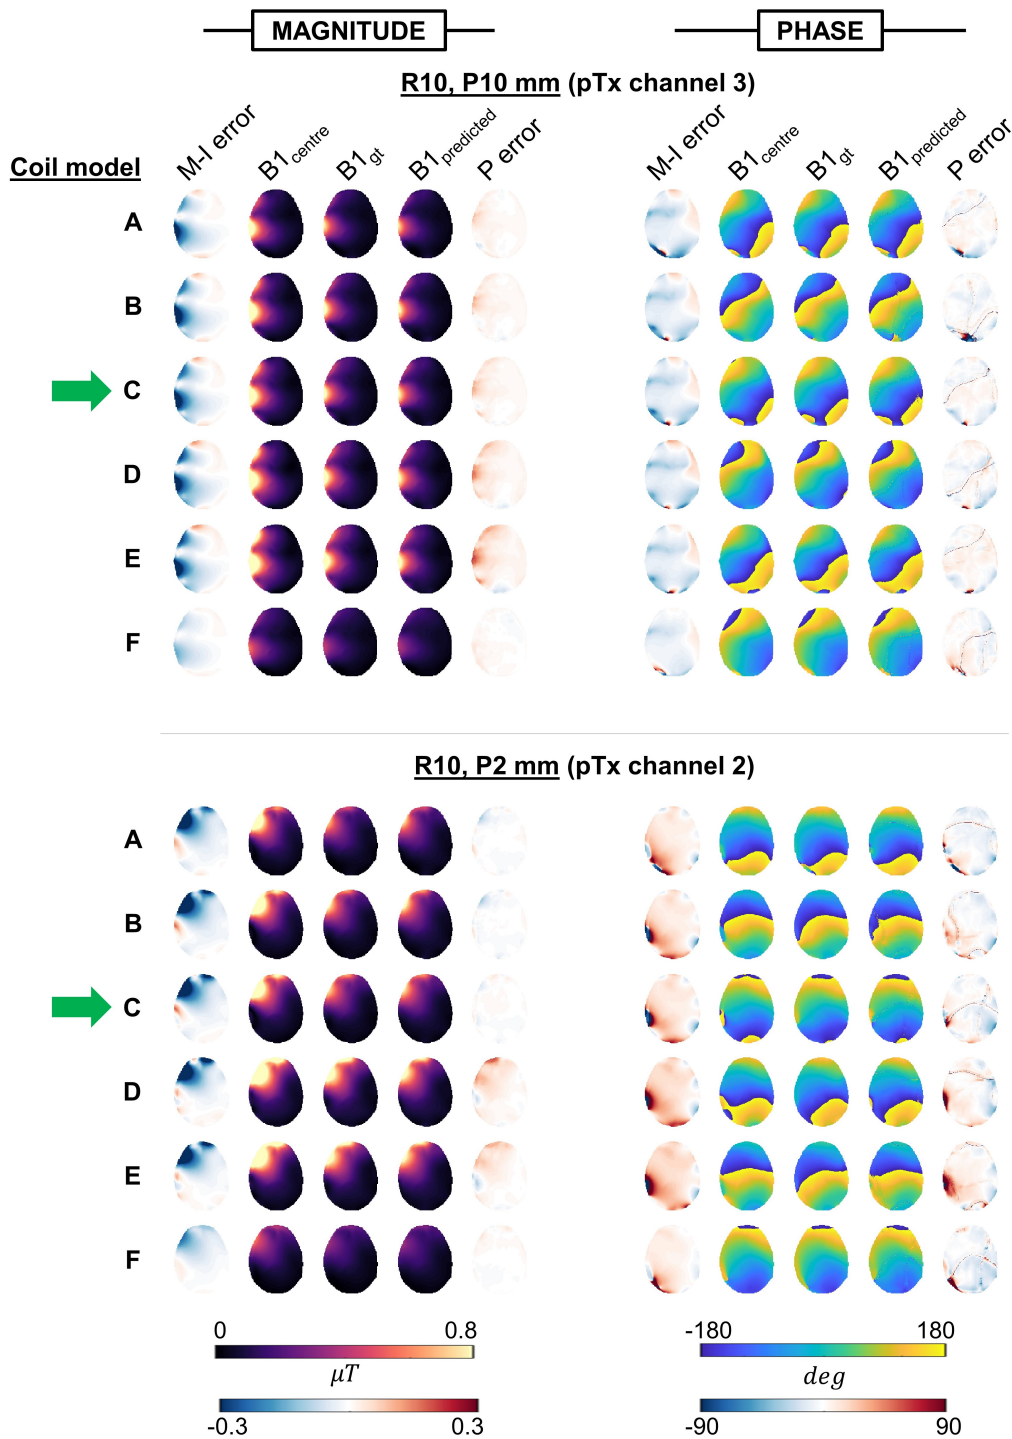


Figure 10.2: B_1^+ magnitude and phase maps for 2 example displacements (a single pTx channel is shown). B_1^+ centre is the simulated map without motion (network inputs). B_1^+ gt is the ground-truth simulated map at the displaced position. B_1^+ predicted is the network-output map at the displaced position. Motion-induced (M-I) error shows difference between B_1^+ initial and B_1^+ gt. Prediction (P) error shows difference between predicted and ground truth maps. Networks were trained using coil C (green arrow).

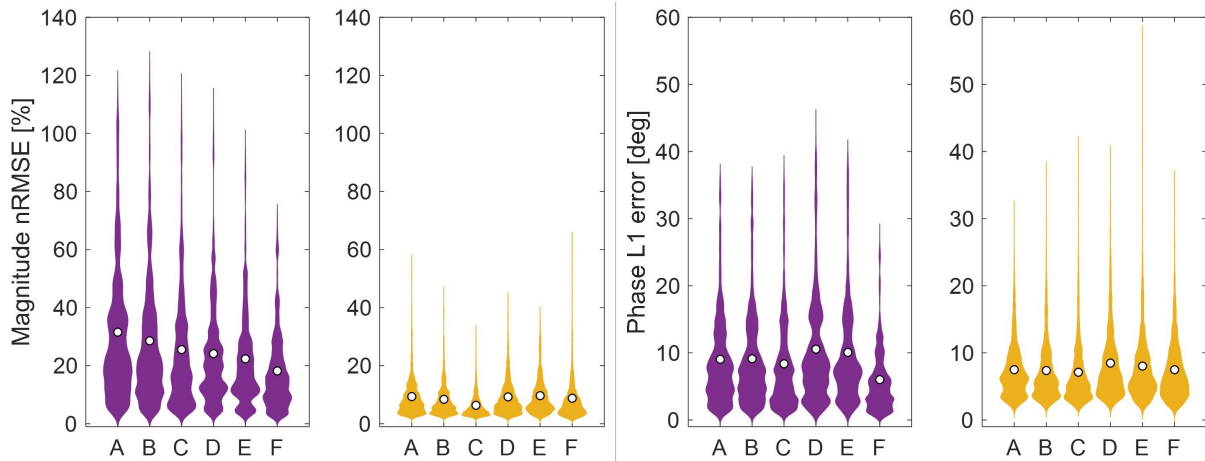


Figure 10.3: Magnitude (left) and phase (right) error distributions including all evaluations (all slices, channels and positions) with each coil model. Motion- and prediction-related errors are shown in purple and yellow, respectively.

than the other coils (data not shown). There was still a general trend that motion-induced error decreased monotonically from 1.7 to 0.4 μT with increasing coil size (the short-looped coil B had the highest-observed magnitude error caused by motion). Also, in keeping with nRMSE, mean prediction error was significantly lower for the training coil C than all other coils ($6 \times 10^{-3} \mu\text{T}$, compared to between 7×10^{-3} and $14 \times 10^{-3} \mu\text{T}$ for the other coils).

Figure 10.4 shows the average error as a function of displacement magnitude for each coil. Positional sensitivity for magnitude nRMSE was highest for coil A (3.7% per mm), and monotonically decreased with increasing coil size to 2.1% per mm for coil F. Conversely, positional sensitivity of prediction nRMSE remained below 0.9% per mm for all coils but was lowest for coil C (0.49% per mm). In terms of phase, coils' sensitivity to motion increased monotonically with increasing coil size (from 1.02° to 1.2° per mm), except for the large-arrayed coil F which was less sensitive (0.7° per mm). As expected, prediction error for phase was less dependent on position than motion-related error in all cases (between 0.5 and 0.7° per mm), although for all coils, error in predicted maps sometimes exceeded motion-related error when the movement was small.

Coil C experienced significantly lower motion-related phase error (mean L1 loss for coil C = 8.4°) than all other coils (9.0 to 10.6°) except for coil F which was much less sensitive to motion (6°). Importantly, coil C also yielded significantly lower prediction phase error (7.1°) than all coils (between 7.5 and 8.5°) except for coil B which was higher than coil C, but not

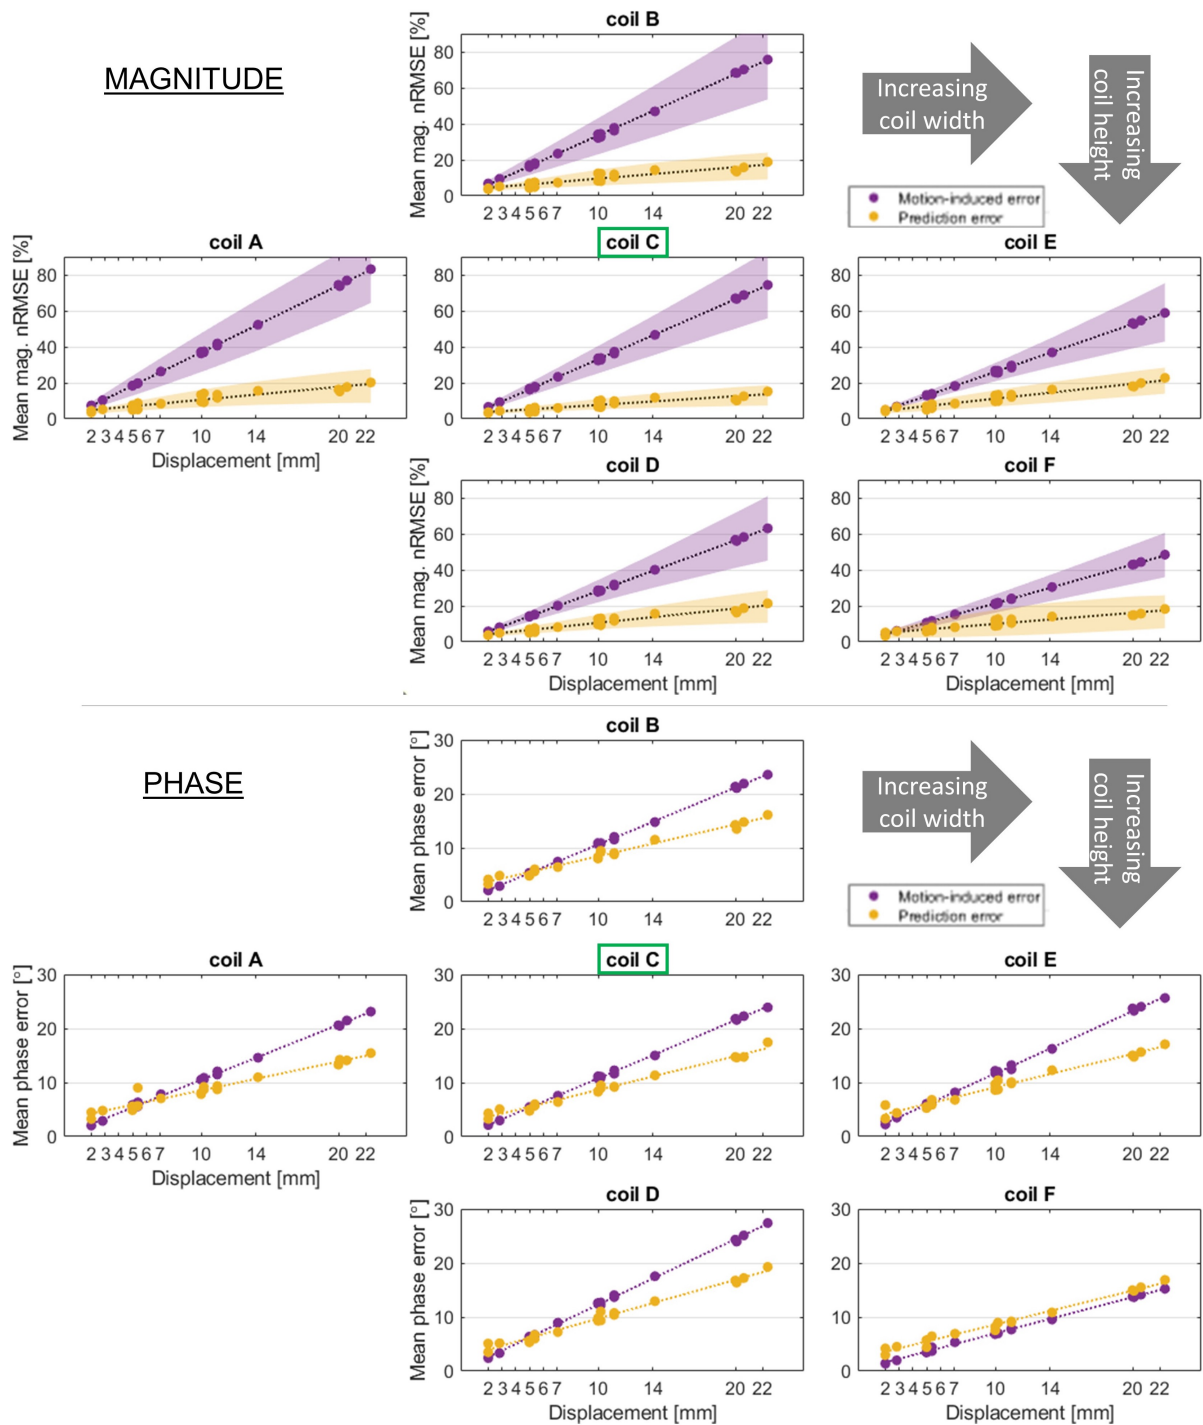


Figure 10.4: Mean magnitude and phase error, averaged across all slices and channels for each displacement. Motion- and prediction-related errors are shown in purple and yellow, respectively. Different coil models (with varying loop size and / or array radius) are shown in subplots. Networks were trained for coil C. Magnitude nRMSE is shown as a percentage of the mean magnitude without motion. Standard deviation is shown for magnitude (shaded regions) but omitted from phase for clarity.

significantly so (7.4°). Like magnitude, Figure 10.3 shows that the downward shift in error distributions between motion-related and predicted L1 loss is appreciable in all cases except the largest coil model. Prediction error was significantly lower than motion-related error for coils A to E, while it was significantly higher for coil F.

The percentage of evaluations in which prediction-related error was exceeded by motion-related error for each coil was also investigated. Performance of phase networks was less consistent than those of magnitude for all coils, but in particular, phase prediction error exceeded motion-related error in around 70% of evaluations for coil F. Coils A to E (which had smaller/larger loops, but the same array radius as the coil used for network training) benefitted in at least 91% of magnitude, and at least 55% of phase evaluations. Coil F also benefitted less consistently for magnitude ($\sim 65\%$ of evaluations) compared to the other coils.

Figures 10.5 and 10.6 show flip-angle results for pulses designed using centre ($\text{pulse}_{\text{initial}}$) vs. predicted ($\text{pulse}_{\text{re-designed}}$) B_1^+ maps. For all coils, the mean flip-angle nRMSE for $\text{pulse}_{\text{re-designed}}$ (between 10 - 13%) was significantly lower than that for $\text{pulse}_{\text{initial}}$ (between 18 - 20%). Coil C yielded the lowest mean and maximum flip-angle nRMSE, and in the case of coils D and E that difference was statistically significant. In terms of worst-cases, coils A, C and F benefitted from the proposed method the most, with their highest-observed nRMSE for $\text{pulse}_{\text{re-designed}}$ being less than half those observed using $\text{pulse}_{\text{initial}}$. Positional sensitivity of predictions was lower than motion-related error, however, like B_1^+ phase, prediction error was higher than motion-related error for some smaller movements.

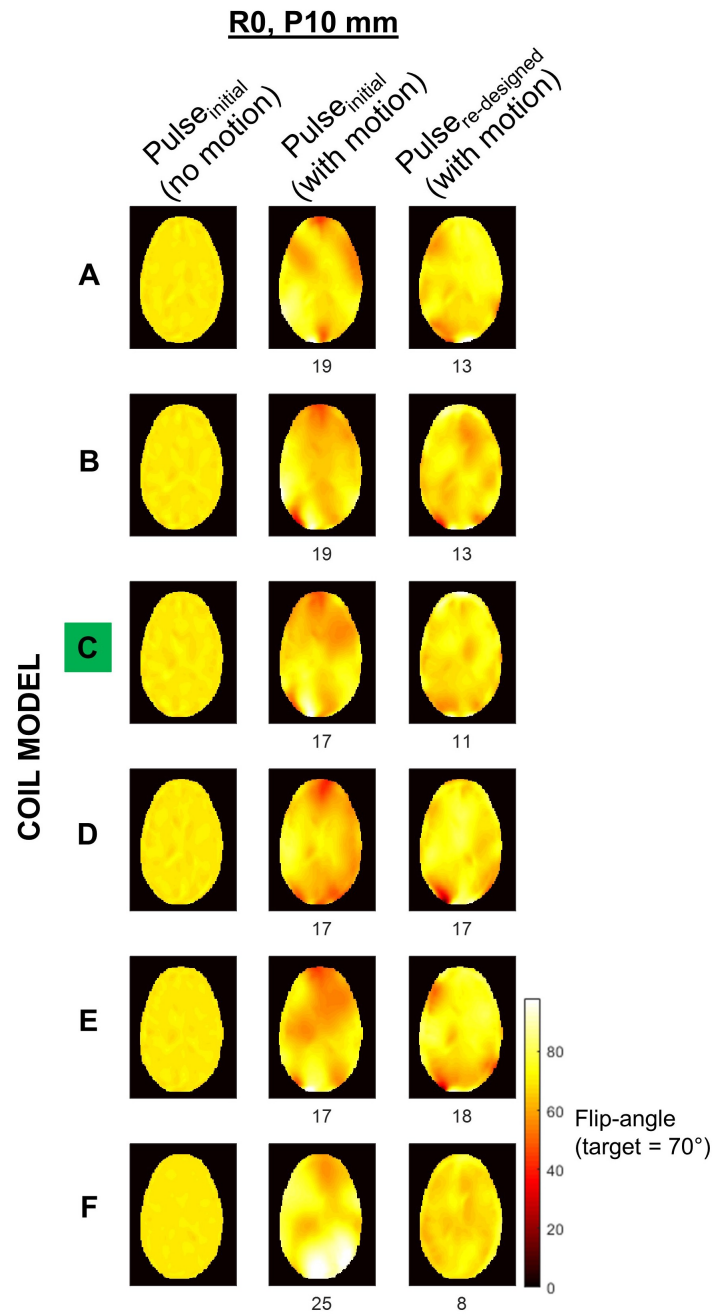


Figure 10.5: Example flip-angle profiles. Pulse_{initial} was designed conventionally, using the B_1^+ map without motion. Its performance following 10 mm posterior displacement occurs is shown in the middle column. Pulse_{re-designed} was designed using predicted maps and evaluated on the same displaced map. nRMSE (%) for both pulses following motion is shown below profiles.

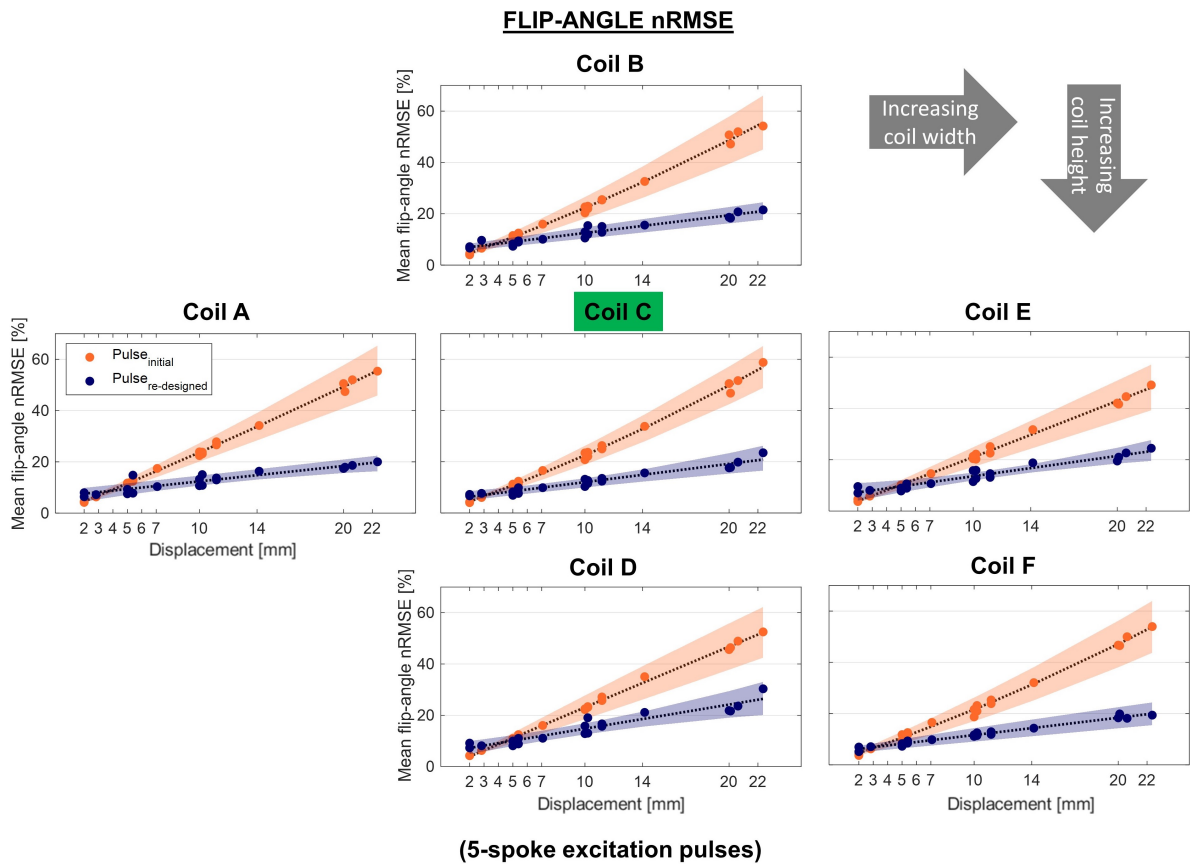


Figure 10.6: Flip-angle nRMSE for both pulses, averaged over all slices for each displacement. Orange and purple colours show performance of pulse_{initial} and pulse_{re-designed}, respectively. nRMSE is shown as a percentage of target flip-angle (70°).

10.4 Discussion

In this study, the inter-coil generalisability of neural networks trained for B_1^+ prediction was tested. We trained networks using simulated data from one pTx coil model, and performed network inference with five additional, differently sized coil models. In terms of B_1^+ magnitude, the error caused by motion was significantly reduced in predicted maps for all coils, although the coil used for network training yielded significantly lower prediction error than the other coils. For B_1^+ phase, all coils except for the largest coil model yielded significantly lower prediction than motion-related error. Like magnitude, network predictions for phase were significantly more accurate for the training coil compared to all other coils (except for one, similarly-sized coil, where the difference was non-significant). For all coils, B_1^+ error reduction translated into improved flip-angle uniformity following motion when pTx pulses were re-designed using predicted maps.

Motion-related B_1^+ error for coil C was similar to that in Chapter 9 as expected (maximum magnitude nMRSE = 117% in Ch. 9 vs. 116% here; maximum phase L1 loss = 35° vs. 38°). However, prediction error was generally slightly higher here. Maximum-observed B_1^+ magnitude nRMSE was 33% here, compared to 20% for the same displacements in Chapter 9, and phase L1 loss was 41° compared to 20°. Like in Chapter 9, phase predictions were generally less accurate than magnitude predictions for all coils, with the worst phase performance seen for small movements (<5mm) where motion-related error was sometimes lower than prediction error. We believe that the use of one less body model in training (Billie and Duke here, vs. Ella, Duke and Dizzy in Chapter 9) underlies the poorer prediction quality here. Nevertheless, the majority of evaluations yielded substantially lower B_1^+ and flip-angle error when the proposed motion correction method was used.

Motion sensitivity of B_1^+ magnitude generally decreased as the coil size increased. Phase error showed the opposite dependency, increasing with larger coils (except for coil F, where phase error was low). However, in both cases, the mid-sized coil C (which was the coil used for network training) yielded the most accurate predicted maps. Although performance for coil C was significantly better, prediction error was significantly lower than motion-related error across all coils (except phase for coil F), indicating some degree of network generalisability. Motion-related flip-angle nRMSE was also significantly reduced for all coils

when pulses re-designed using predicted maps were used, showing that networks are sufficiently generalisable to bring practical benefit to a range of coil sizes. An exception (in all coils) was in cases where the movement was very small, which we attribute to the poorer B_1^+ phase prediction quality at these positions.

Motion-related B_1^+ magnitude error in the coil with the larger array radius (coil F) was notably lower than the smaller coil models, in keeping with the SAR findings reported in Chapter 6. The difference between predicted and ground-truth maps was also low, however, prediction nRMSE appeared high for coil F due to the normalisation step. The magnitude of B_1^+ fields produced by coil F is inherently lower, meaning that nRMSE is inflated compared to coils with a higher magnitude B_1^+ field. However, nRMSE reflects differences arising due to different power requirements between coils. Simulated fields were normalised to 1 W input power for our study, whereas in reality, coils which have intrinsically lower field magnitudes (e.g., coil F) require more RF power to elicit the same flip-angle compared to smaller coils with higher-magnitude fields. Therefore, nRMSE is more practically useful than the B_1^+ difference alone. This is in line with the pTx pulse evaluations presented here, which inherently incorporate different power requirements during pulse design - in terms of flip-angle nRMSE, coil F was among the most motion-sensitive models.

As we saw in Chapter 9, networks were able to accurately predict fields following displacements from -for example- 0 to 5 mm positions as well as from 5 to 10 mm positions. This is despite non-linearity in the spatial decay of fields, meaning that changes in magnitude for a given relative displacement will differ, depending on the starting position. We believe that the (normalised) magnitude of the input data implicitly informs the network as to what the starting position was. In other words, with the spatial decay of B_1^+ in mind, the magnitude of the input informs where on that curve the input sample is from, and therefore, what the motion-induced change will be. Since fields inside the head are smoother when the coil-sample distance is large (e.g., with a large array radius), the curve is likely to differ substantially, which may explain the poorer performance for coil F here.

Lv et al. [226] used transfer learning to improve reconstruction quality of under-sampled images using GANs. Networks were initially trained for parallel reconstruction using a large, generic database of 4,700 brain images. Transfer learning was subsequently conducted on the pre-trained networks using small 'tune-up' datasets (the smallest of which was 216

images) for specialist reconstructions, including brain images containing tumours, and other anatomies (e.g., knee). Transfer learning yielded superior reconstruction quality compared to both the initial pre-trained model alone, and a model trained using only the small specialist dataset. A similar approach may improve B_1^+ prediction quality for other coils by using a small, coil-specific training dataset to tune pre-trained networks [227]. In light of our results, where performance was worst for the coil structure least similar to the training coil, it is likely that transfer learning would be the minimum requirement for use with coil structures which substantially differ in their design (e.g., stripline vs. loop arrays).

10.4.1 Conclusions

B_1^+ prediction networks for prospective motion correction in pTx were generalisable across coil models to some extent, however the accuracy of predicted maps reduced as the coil became less similar to that used for network training. In particular, we saw the poorest performance for the largest coil model (coil F) which –unlike the other models– had a larger array radius compared to coil C. Therefore, our findings support the notion that networks are somewhat generalisable, but the coil which was included in training benefits the most from the proposed approach, as evidenced by the U-shaped relationship between coil size and error which we observed. Whether the network’s correction capability is sufficient, or whether transfer learning is required to improve performance on a coil-by-coil basis will depend on the application, as well as how different the coil model is from that used for network training.

Part III

General Discussion and Concluding Comments

Chapter 11

General Discussion

As emphasised throughout this thesis, parallel transmission (pTx) of tailored RF pulses is an effective way to overcome B_1^+ non-uniformity in the head at 7T [18, 20, 32, 128]. The performance of tailored pulses is compromised when the measured B_1^+ distributions do not match those present at the time of pulse playout, for example due to a change in coil loading caused by patient motion. Approaches such as Universal Pulses [123], SmartPulse [126] and FOCUS pulses [125] (discussed in Section 3.4.4) go some way to addressing differences in coil loading caused by different anatomy, however, their performance remains susceptible to motion because the RF pulse is fixed throughout the scan, meaning that motion-related changes to the B_1^+ fields are overlooked. Similarly, regarding safety, motion-related changes in SAR due to changes in E-fields have, historically, not been well characterised. Additionally, these plug-and-play approaches inevitably yield inferior flip-angle profiles compared to fully tailored pulses. Motion correction techniques exist, but are either retrospective (and therefore do not address the safety concerns), or poorly suited to pTx (where real-time updates to RF waveforms are more complex to implement compared to sTx).

Recently, research has emerged which has begun to quantify the effects of motion on image quality [171, 173] and SAR [130, 173, 190, 191, 207]. However, the generalisability of these findings across different pTx imaging contexts remained unclear. This thesis has presented findings from investigations into differences in motion-sensitivity between large- and small-tip angle pTx pulses; the impact of pTx coil geometry on motion-related SAR changes; and inter-subject differences in motion-related SAR changes. Furthermore, it has

presented two versatile methods which could be used to reduce the impact of motion on image quality in pTx.

The following paragraphs summarise the main findings of this thesis in the context of the current state of the field. Some limitations of the work are addressed, and future research directions are discussed.

11.1 Motion and image quality in pTx:

Concluding comments

The B_1^+ changes caused by motion are especially spatially complex in pTx due to the interaction of independent channels' B_1^+ fields, and the differing effects of a given displacement on each. The artificial contrast which manifests in images depends on acquisition and reconstruction parameters, however, in cases where the movement is small, it may go unnoticed. While this is arguably less concerning for qualitative image quality, quantitative measures will be severely impacted by (even very small) B_1^+ field changes. Currently, in clinical settings, advanced motion correction strategies are rarely used, with scans being repeated instead [157, 159]. This has clear negative effects on efficiency, cost, and patient comfort. As quantitative MR-based measures begin to creep into clinical workflows, motion correction will be required.

The approaches presented in this thesis represent ways to overcome the positional-dependency of the B_1^+ field on pTx performance. Motion also affects image quality in other ways, including by altering the B_0 field and effective field of view. In Chapter 4, dynamic field effects were discussed in the context of B_0 field changes. Previous work addressing the effects of motion on the B_0 field found that roll (ear towards shoulder) and pitch (nodding) rotations have the largest impact on B_0 [186], while others have reported in-plane displacements cause the largest changes in B_1^+ (in line with the results presented in this thesis) [173]. The discrepancy is expected because wave propagation depends on the angle between the field and tissue boundary, and B_0 and B_1^+ fields are orthogonal. While the largest contributor (B_0 or B_1^+ field effects) to error and / or artefacts may therefore vary depending on the motion type, the B_1^+ approaches outlined could feasibly

be incorporated into a battery of correction methods including those designed for dynamic B_0 correction [186].

11.2 Motion and safety in pTx:

Concluding comments

Although neither correction approach presented in this thesis actively constrains or manages SAR, Chapters 6 and 7 contribute much to our understanding of motion-related SAR behaviour in pTx. We found that local SAR produced by pTx arrays with large loop elements (or a large array radius) was less motion-sensitive than that produced by arrays with small loop elements. This knowledge aligns with other research investigating SAR behaviour of pTx coils without motion [127], and may be useful for future hardware design. In the investigation into inter-subject variability, models' worst-cases were observed for different pulses and different slice locations, and we did not observe population-specific patterns in the sensitivity of SAR to motion. It is therefore likely that the previously-reported findings on motion and SAR in pTx are applicable to all individuals [130, 190, 191].

The complexity of dynamic field interactions between multiple transmitting elements means that there is currently no simple or foolproof method to establish if, when, and where local SAR limits are being exceeded throughout a pTx scan. This, along with the findings presented in Chapters 6 and 7 present a worrying picture for the safety of pTx, however there are important considerations which yield a more balanced view.

Firstly, the worst-cases in our studies (as well as those previously reported [130]) were usually (although not always) observed following the largest displacements, which were on the scale of centimetres. With sufficient well-placed padding, this scale of movement is rare with most compliant, healthy volunteers. Movement on the scale of a few millimeters is instead more common [152, 228, 229]. Nevertheless, clinical populations including those suffering from motor diseases and elderly patients are prone to large movements. Children also tend to move more, and in different ways compared to adult subjects, including more axial motion [229]. Furthermore, as previously mentioned, slow drifts occur in almost all subjects, and can often go undetected. So, while -statistically speaking- safety in pTx is unlikely to be

compromised due to typically-observed motion patterns, safety certainly cannot be guaranteed without further development of SAR monitoring and management in pTx.

Secondly, the metric of SAR does not reflect tissue heating directly, and their spatial correspondence has been contested in the literature (see Ref. [230], where good correspondence was found, compared to Refs. [36, 37], where it was not). SAR limits also include a temporal component, meaning that TR can be increased to reduce the likelihood of temperature rise, although this is detrimental to the time-efficiency of clinical routines. Finally, as mentioned in Chapters 6 and 7, the tendency for SAR hotspots to change location under conditions of motion can be viewed as positive, from the point of view that energy is more spatially distributed [197, 198]. However this view may be problematic, firstly, because it assumes that SAR at the previous hotspot location reduces - which is not guaranteed, and secondly, because the timing and location of the hotspot change is difficult to predict. The latter point is especially true for pTx, where a major contributor to motion-related SAR changes is the complex channel interference patterns. This is in addition to the changing tissue-channel proximity, which yields somewhat more predictable changes to the location of peak local SAR, and is assumed to drive motion-related SAR changes in quadrature transmission [127, 130, 207].

In general, our results showed greater SAR variability than those reported in the literature. While previous studies suggest that a factor of ~ 2 variation exists for local SAR among different subjects [201–203], and different subject positions [190, 193] the variation we observed due to motion was typically higher (>3 -fold), more in line with Ref. [130]. It is well-known that SAR distributions depend on the coil design. In Chapter 6, it is demonstrated that SAR sensitivity to motion also depends on the coil, with taller or wider loops, or a larger array radius yielding reduced SAR sensitivity than smaller coil designs. It is likely that the local SAR variation we observed would be lower for a coil model with a larger radius, in line with our results in Chapter 6. However, although local SAR may be reduced with larger coils, more RF power is usually required with increasing sample-to-coil distance [127], meaning that global SAR is likely to increase compared to the values we observed (although it is unclear whether the variation across positions would also increase). Additionally, close-fitting coils may be used for applications such as neonatal imaging, where motion would be a particular concern. In all, the SAR behaviour we observed across coils

and across subjects suggests that further investigation and development of methods for real-time SAR management is required on a subject-by-subject basis in order to ensure safety for all patients in pTx.

11.3 The future of deep learning approaches

The deep learning method presented in Chapter 9 used simulated data only.

Multi-compartment, anthropomorphic phantoms exist and could theoretically be used to create training data instead. The drawback of this approach is twofold - a database of realistic phantoms would be required to increase inter-subject generalisability, and the B_1^+ fields produced in such phantoms do not correspond well with those observed in human subjects [231, 232] (B_1^+ values over double those seen in a human volunteer in Ref. [231]). Considering the financial and temporal investments required to create intricate, heterogeneous phantoms, these drawbacks are rather prohibitive.

Compared to phantom data, in-vivo data would yield a training dataset more representative of the networks' intended use (i.e., human subjects), however this brings its own challenges. For network training, the mapping between positional changes and field changes needs to be well-defined in the training dataset. While this is easily achieved in simulations, precise choreographed motion is difficult to achieve in-vivo - even more so when the exact same movement is required to be replicated by several subjects. An alternative approach would be to train one network for arbitrary movements, allowing subjects to move freely inside the scanner while creating the in-vivo training data, and retrospectively extracting motion parameters which could be input to the network. A network architecture known as StarGAN [233] is a strong candidate for this. StarGAN is a generative network which performs one of a range of different transformations on an input image. The transformation required at each evaluation instance is specified to the network as a separate input. StarGAN could feasibly take motion parameters as an input to determine which transformation to apply to the input image (the B_1^+ field). Future research in the field could benefit from adopting this approach.

Very recently, a pipeline was demonstrated which is capable of predicting pTx B_1^+ and SAR distributions on a subject-by-subject basis using only a 9-second long localiser scan as input

[133]. The key here was an intermediate deep learning segmentation step where tissue properties (density, permittivity and conductivity) were inferred directly from the localiser. Following that step, a fast EM solver was used to calculate subject-specific B_1^+ and SAR distributions. Another attractive possibility for future research would be to also incorporate subject position into this approach.

Chapter 12

Appendix

This appendix shows the supporting information from the work presented in Chapter 9, which is also available in the online published version at <https://onlinelibrary.wiley.com/doi/full/10.1002/mrm.29132> (accessed 25/05/2023).

BILLIE

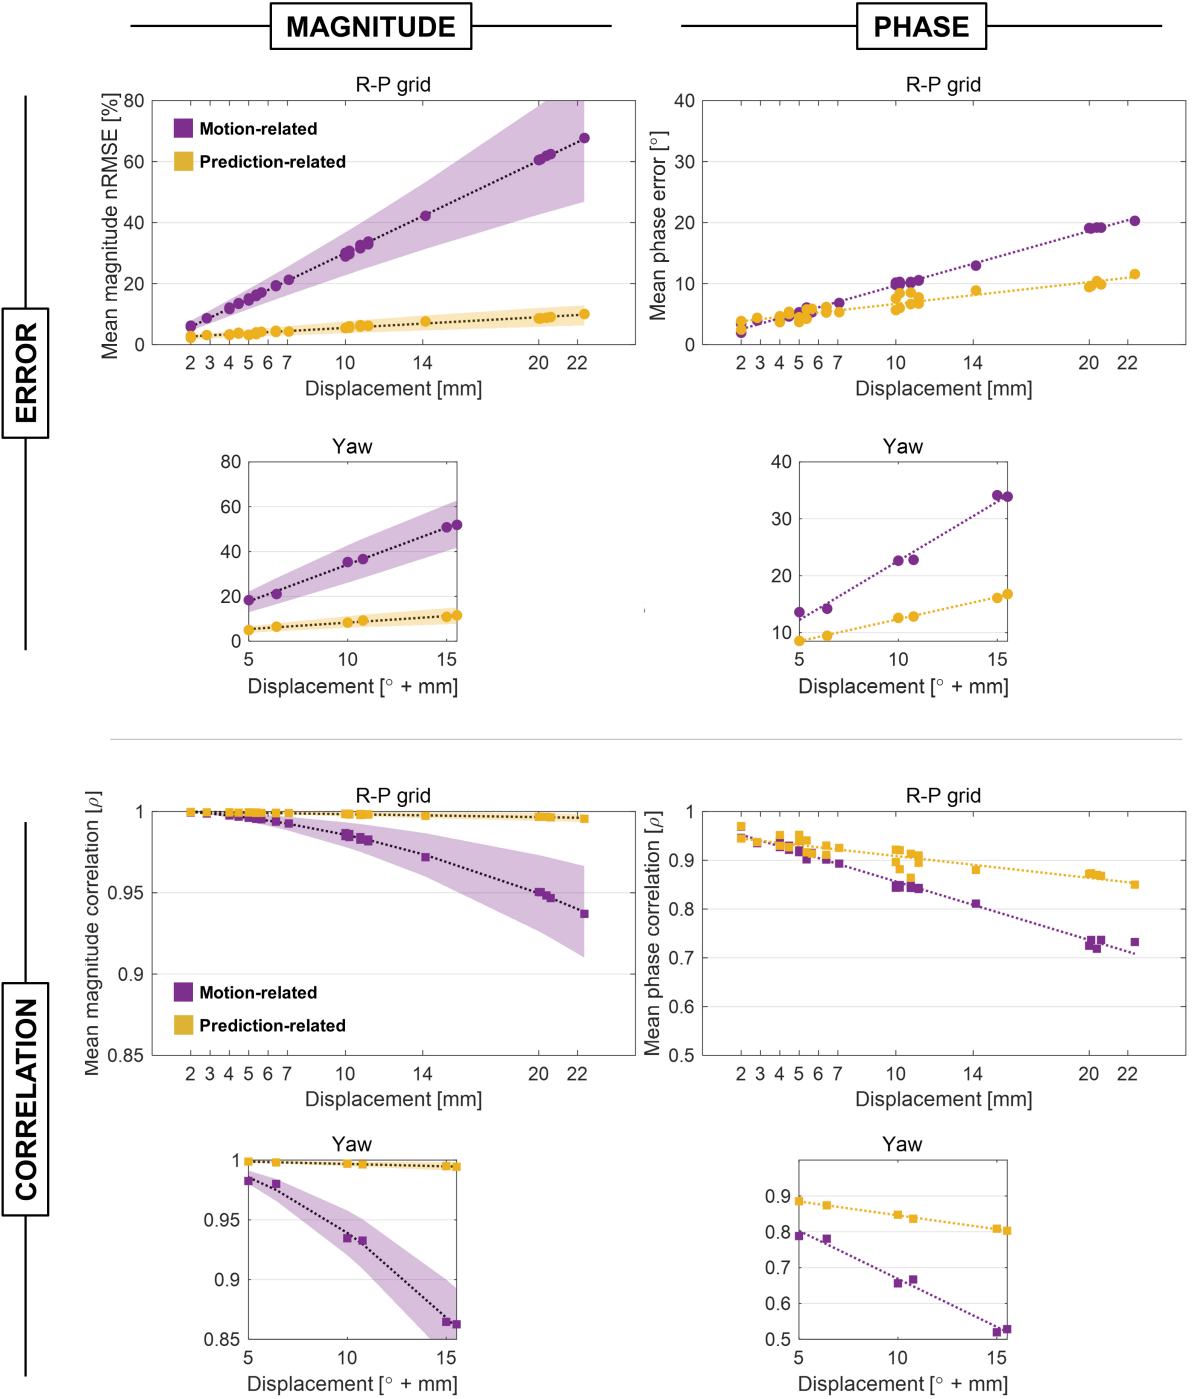


Figure 12.1: B_1^+ prediction quality (error and correlation) for the Billie model only. (Figure 5 displays the same evaluations but averaged over the Billie and Dizzy body models.) Data from the Billie model was used for validation during the training process, while the Dizzy model was not involved at all during training.

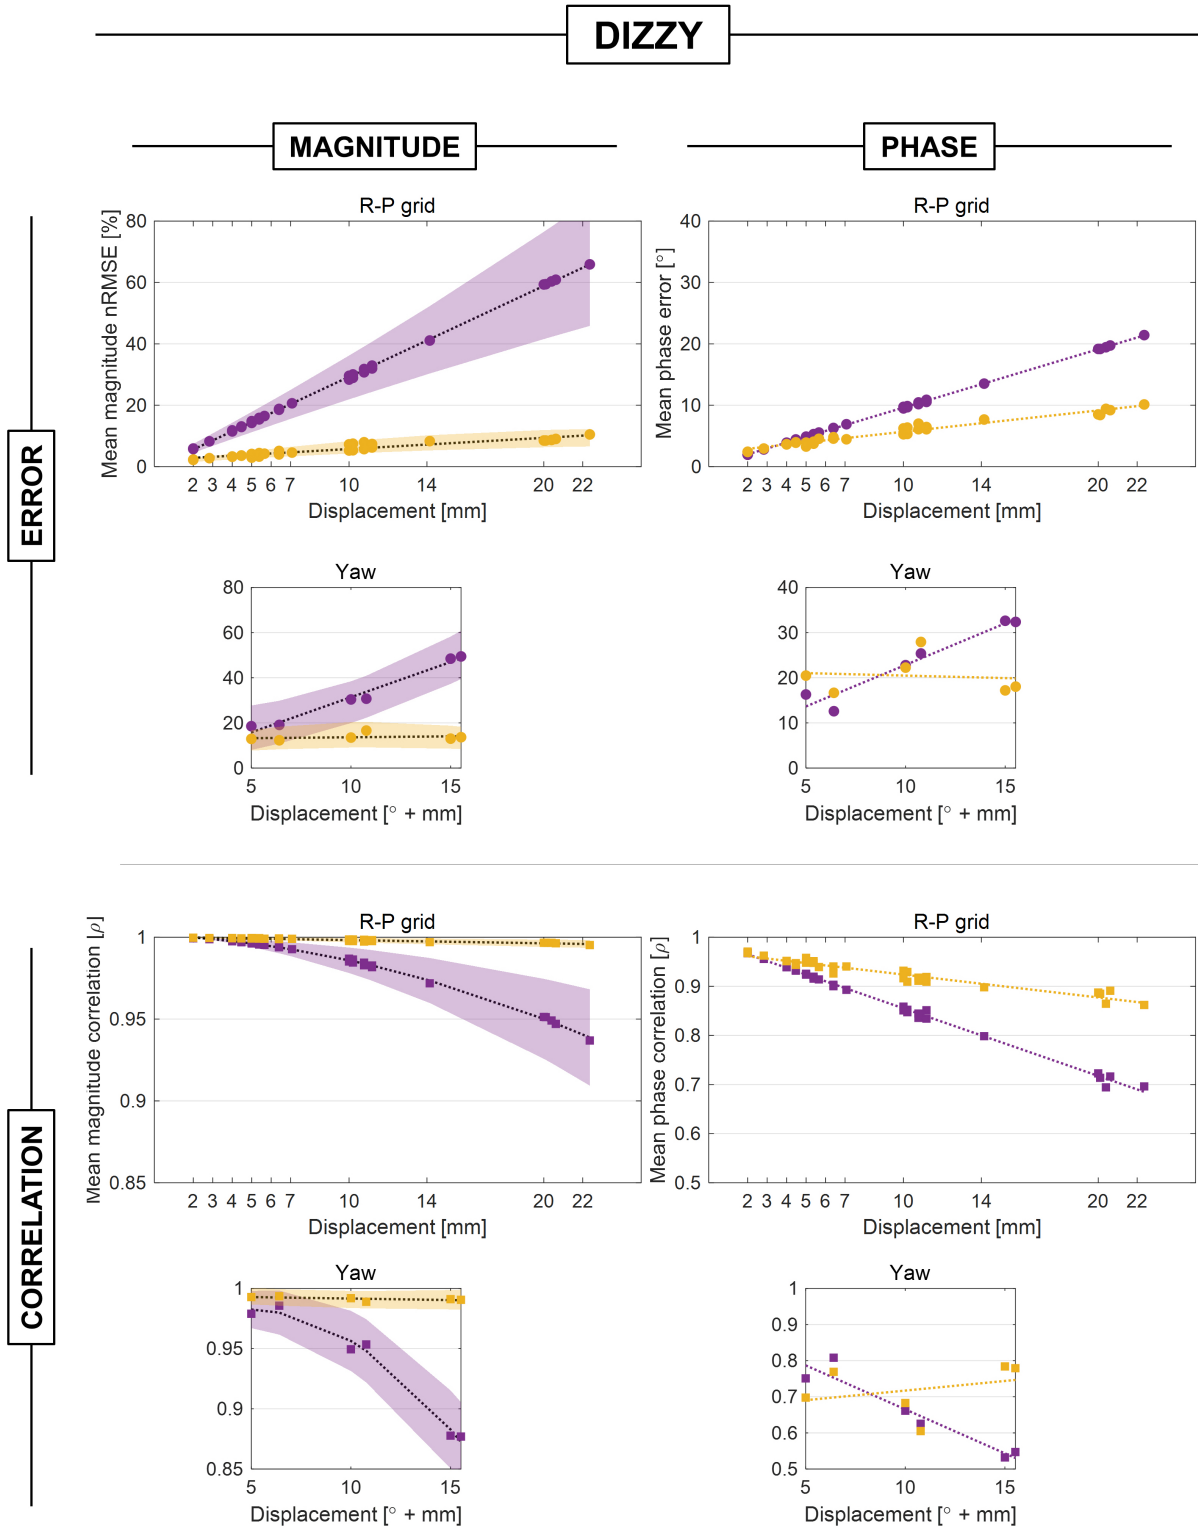


Figure 12.2: B_1^+ prediction quality (error and correlation) for the Dizzy model only. (Figure 5 displays the same evaluations but averaged over the Billie and Dizzy body models.) Data from the Billie model was used for validation during the training process, while the Dizzy model was not involved at all during training.

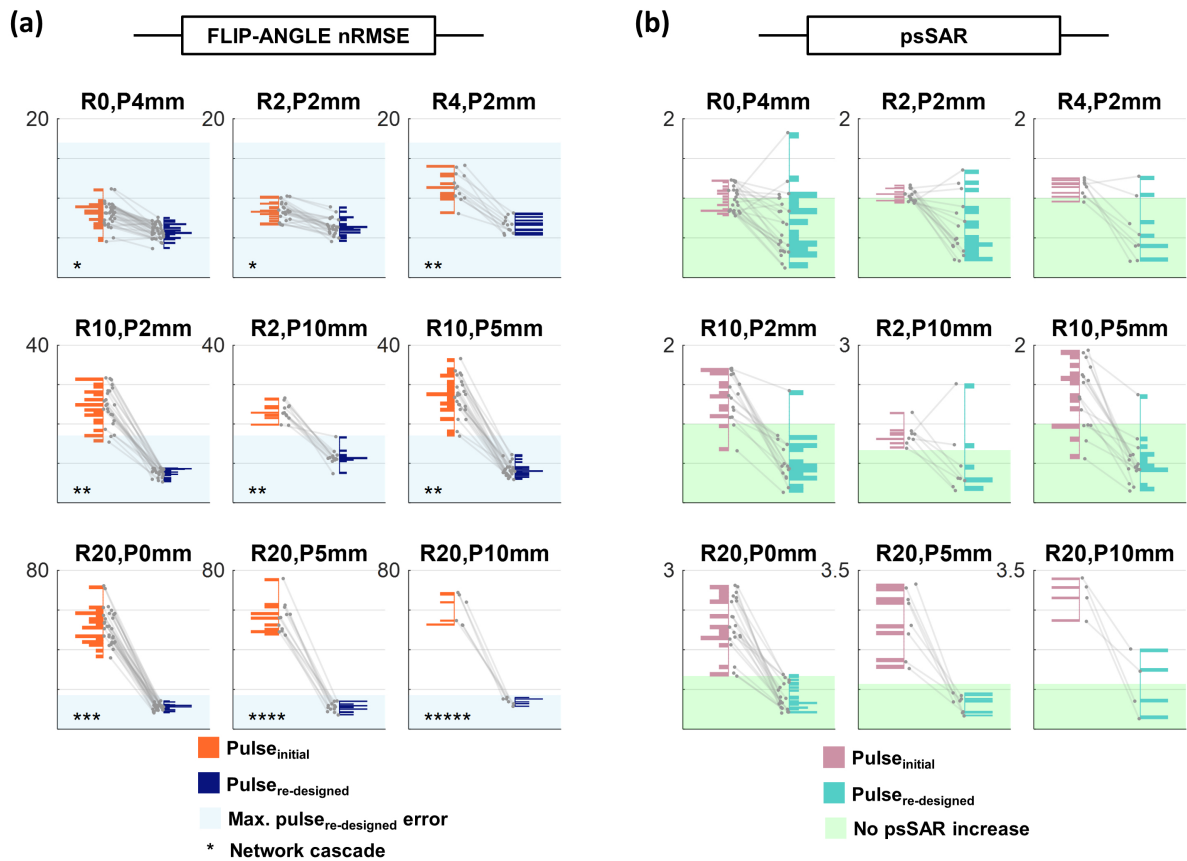


Figure 12.3: (a) Raincloud plots showing flip-angle nRMSE (% of target flip-angle) for 9 example displacements using $pulse_{initial}$ (orange) and $pulse_{re-designed}$ (dark blue). Asterisks indicate the number of network cascades required for evaluation. The light blue shaded region shows maximum prediction nRMSE across all 35 evaluated displacements (consistent across panels). (b) All SAR evaluations for the same 9 example displacements. Vertical axes show relative psSAR, calculated as psSAR as a factor of that without motion using $pulse_{initial}$ ($psSAR_{centre}$). The green shaded region shows psSAR at or below $psSAR_{centre}$ (i.e., indicating that psSAR did not increase following motion). Neither pulse included SAR-constraint in the design.

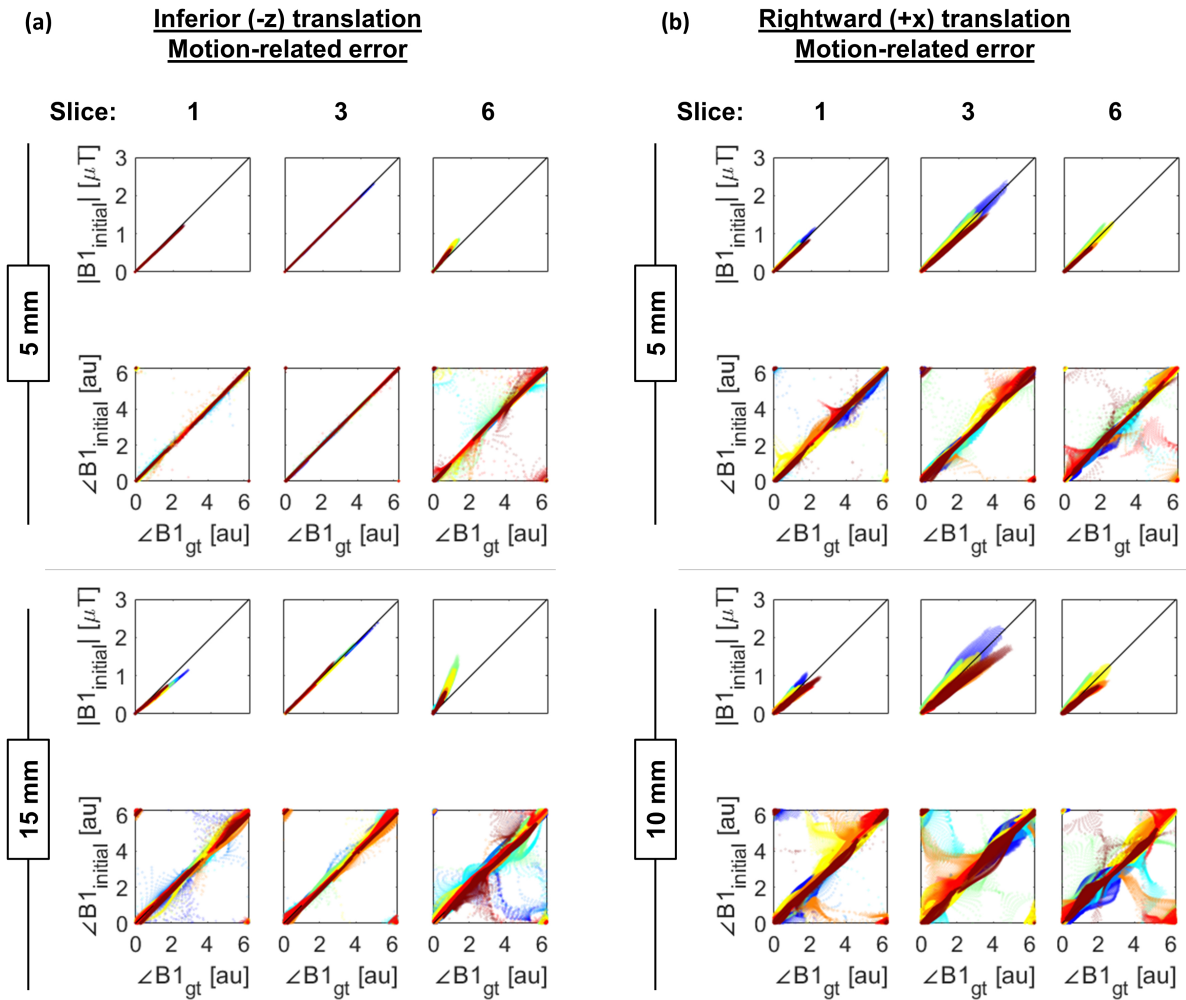


Figure 12.4: Voxel-wise correlations between initial ($B1_{initial}$) and ground-truth ($B1_{gt}$) B_1^+ maps, to compare motion-related error following through-plane (a) and within-plane (b) displacements. Results are shown for two positions, and three slice locations (1 = least superior, 6 = most superior). Note that the large inferior displacement (15 mm) is greater than the large rightward displacement (10 mm), and yet error remains lower due to the through-plane nature of the movement.

Bibliography

1. Wang, I. *et al.* Value of 7T MRI and post-processing in patients with nonlesional 3T MRI undergoing epilepsy presurgical evaluation. *Epilepsia* **61**, 2509–2520 (2020).
2. Park, J. E., Cheong, E.-N., Jung, D. E., Shim, W. H. & Lee, J. S. Utility of 7 Tesla Magnetic Resonance Imaging in Patients With Epilepsy: A Systematic Review and Meta-Analysis. *Frontiers in Neurology* **12** (2021).
3. De Ciantis, A. *et al.* 7T MRI in focal epilepsy with unrevealing conventional field strength imaging. *Epilepsia* **57**, 445–454 (2016).
4. Feldman, R. E. *et al.* 7T MRI in epilepsy patients with previously normal clinical MRI exams compared against healthy controls. *PLOS ONE* **14**, e0213642 (2019).
5. Nielsen, A. S. *et al.* Focal cortical lesion detection in multiple sclerosis: 3 tesla DIR versus 7 tesla FLASH-T2*. *Journal of Magnetic Resonance Imaging* **35**, 537–542 (2012).
6. Kolb, H. *et al.* 7T MRI Differentiates Remyelinated from Demyelinated Multiple Sclerosis Lesions. *Annals of Neurology* **90**, 612–626 (2021).
7. Maranzano, J. *et al.* Comparison of Multiple Sclerosis Cortical Lesion Types Detected by Multicontrast 3T and 7T MRI. *American Journal of Neuroradiology* **40**, 1162–1169 (2019).
8. Louapre, C. & Beigneux, Y. 7 Tesla MRI will soon be helpful to guide clinical practice in multiple sclerosis centres – Yes. *Multiple Sclerosis Journal* **27**, 360–362 (2021).
9. Suh, C. H., Kim, S. J., Jung, S. C., Choi, C. G. & Kim, H. S. The “Central Vein Sign” on T2*-weighted Images as a Diagnostic Tool in Multiple Sclerosis: A Systematic Review and Meta-analysis using Individual Patient Data. *Scientific Reports* **9**, 18188 (2019).

10. Obusez, E. C. *et al.* 7T MR of intracranial pathology: Preliminary observations and comparisons to 3T and 1.5T. *NeuroImage* **168**, 459–476 (2018).
11. Zaiss, M. *et al.* Downfield-NOE-suppressed amide-CEST-MRI at 7 Tesla provides a unique contrast in human glioblastoma. *Magnetic Resonance in Medicine* **77**, 196–208 (2017).
12. Wells, J. A. *et al.* In vivo imaging of tau pathology using multi-parametric quantitative MRI. *NeuroImage* **111**, 369–378 (2015).
13. Stäb, D. *et al.* Cardiac Magnetic Resonance Imaging at 7 Tesla. *Journal of Visualized Experiments* (2019).
14. Reiter, T. *et al.* On the way to routine cardiac MRI at 7 Tesla - a pilot study on consecutive 84 examinations. *PLOS ONE* **16**, e0252797 (2021).
15. Aringhieri, G., Zampa, V. & Tosetti, M. Musculoskeletal MRI at 7 T: do we need more or is it more than enough? *European Radiology Experimental* **4** (2020).
16. Vaughan, J. T. *et al.* Whole-body imaging at 7T: Preliminary results. *Magnetic Resonance in Medicine* **61**, 244–248 (2009).
17. Vaughan, J. *et al.* 7T vs. 4T: RF power, homogeneity, and signal-to-noise comparison in head images. *Magnetic Resonance in Medicine* **46**, 24–30 (2001).
18. Katscher, U. & Börnert, P. Parallel RF transmission in MRI. *NMR in Biomedicine* **19**, 393–400 (2006).
19. Zhu, Y., Alon, L., Deniz, C. M., Brown, R. & Sodickson, D. K. System and SAR characterization in parallel RF transmission. *Magnetic Resonance in Medicine* **67**, 1367–1378 (2012).
20. Deniz, C. M. Parallel Transmission for Ultrahigh Field MRI. *Topics in Magnetic Resonance Imaging* **28**, 159–171 (2019).
21. Ocali, O. & Atalar, E. Ultimate intrinsic signal-to-noise ratio in MRI. *Magnetic Resonance in Medicine* **39**, 462–473 (1998).
22. Kopanoglu, E., Erturk, V. B. & Atalar, E. Analytic expressions for the ultimate intrinsic signal-to-noise ratio and ultimate intrinsic specific absorption rate in MRI. *Magnetic Resonance in Medicine* **66**, 846–858 (2011).

23. Guérin, B. *et al.* The ultimate signal-to-noise ratio in realistic body models. *Magnetic Resonance in Medicine* **78**, 1969–1980 (2017).
24. Le Ster, C. *et al.* Magnetic field strength dependent SNR gain at the center of a spherical phantom and up to 11.7T. English. *Magnetic Resonance in Medicine* **88**, 2131–2138 (2022).
25. Edelstein, W., Glover, G., Hardy, C. & Redington, R. The intrinsic signal-to-noise ratio in NMR imaging. *Magnetic Resonance in Medicine* **3**, 604–618 (1986).
26. Farnell, D. J. & Parkinson, J. B. *An Introduction to Quantum Spin Systems* (Springer, 2010).
27. Galiana, G. *et al.* The role of nonlinear gradients in parallel imaging: A k-space based analysis. *Concepts in Magnetic Resonance Part A* **40**, 253–267 (2012).
28. Nyquist, H. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers* **47**, 617–644 (1928).
29. McGibney, G., Smith, M., Nichols, S. & Crawley, A. Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI. *Magnetic Resonance in Medicine* **30**, 51–59 (1993).
30. Lustig, M., Donoho, D. & Pauly, J. M. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine* **58**, 1182–1195 (2007).
31. Yip, C.-Y., Fessler, J. A. & Noll, D. C. Iterative RF pulse design for multidimensional, small-tip-angle selective excitation. *Magnetic Resonance in Medicine* **54**, 908–917 (2005).
32. Grissom, W. *et al.* Spatial domain method for the design of RF pulses in multicoil parallel excitation. *Magnetic Resonance in Medicine* **56**, 620–629 (2006).
33. Katscher, U., Börnert, P., Leussler, C. & Van Den Brink, J. S. Transmit SENSE. *Magnetic Resonance in Medicine* **49**, 144–150 (2003).
34. Zhu, Y. Parallel excitation with an array of transmit coils. *Magnetic Resonance in Medicine* **51**, 775–784 (2004).

35. Van Lier, A. L., Kotte, A. N., Raaymakers, B. W., Lagendijk, J. J. & van den Berg, C. A. Radiofrequency heating induced by 7T head MRI: Thermal assessment using discrete vasculature or Pennes' bioheat equation. *Journal of Magnetic Resonance Imaging* **35**, 795–803 (2012).
36. Collins, C. M. *et al.* Temperature and SAR calculations for a human head within volume and surface coils at 64 and 300 MHz. *Journal of Magnetic Resonance Imaging* **19**, 650–656 (2004).
37. Wang, Z. *et al.* SAR and temperature: Simulations and comparison to regulatory limits for MRI. *Journal of Magnetic Resonance Imaging* **26**, 437–441 (2007).
38. Delfino, J. G., Krainak, D. M., Flesher, S. A. & Miller, D. L. MRI-related FDA adverse event reports: A 10-yr review. *Medical Physics* **46**, 5562–5571 (2019).
39. Weil, K. M. *et al.* A review of thermal MR injuries. *Radiologic technology* **81** (2010).
40. Bloch, F. Nuclear induction. *Physical Review* **70**, 460 (1946).
41. Wansapura, J. P., Holland, S. K., Dunn, R. S. & Ball Jr., W. S. NMR relaxation times in the human brain at 3.0 tesla. *Journal of Magnetic Resonance Imaging* **9**, 531–538 (1999).
42. Hopkins, A. L., Yeung, H. N. & Bratton, C. B. Multiple field strength in vivo T1 and T2 for cerebrospinal fluid protons. *Magnetic Resonance in Medicine*. **3**, 303–311 (1986).
43. Oros-Peusquens, A. M., Laurila, M. & Shah, N. J. Magnetic field dependence of the distribution of NMR relaxation times in the living human brain. *Magnetic Resonance Materials in Physics, Biology and Medicine* **21**, 131–147 (2008).
44. Rooney, W. D. *et al.* Magnetic field and tissue dependencies of human brain longitudinal $^1\text{H}_2\text{O}$ relaxation in vivo. *Magnetic Resonance in Medicine* **57**, 308–318 (2007).
45. Wright, P. J. *et al.* Water proton T1 measurements in brain tissue at 7, 3, and 1.5T using IR-EPI, IR-TSE, and MPRAGE: results and optimization. *Magnetic Resonance Materials in Physics, Biology and Medicine* **21**, 121–130 (2008).
46. Hahn, E. L. Spin echoes. *Physical Review* **80**, 580 (1950).

47. Deniz, C. M., Alon, L., Brown, R., Sodickson, D. K. & Zhu, Y. Specific absorption rate benefits of including measured electric field interactions in parallel excitation pulse design. *Magnetic Resonance in Medicine* **67**, 164–174 (2012).
48. Hyde, J. S. *et al.* Parallel image acquisition from noninteracting local coils. *Journal of Magnetic Resonance* **70**, 512–517 (1986).
49. Roemer, P. B., Edelstein, W. A., Hayes, C. E., Souza, S. P. & Mueller, O. M. The NMR phased array. *Magnetic Resonance in Medicine* **16**, 192–225 (1990).
50. Griswold, M. A. *et al.* Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine* **47**, 1202–1210 (2002).
51. Sodickson, D. K. & Manning, W. J. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine* **38**, 591–603 (1997).
52. Pruessmann, K. P., Weiger, M., Scheidegger, M. B. & Boesiger, P. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine* **42**, 952–962 (1999).
53. Ohliger, M. A., Grant, A. K. & Sodickson, D. K. Ultimate intrinsic signal-to-noise ratio for parallel MRI: Electromagnetic field considerations. *Magnetic Resonance in Medicine* **50**, 1018–1030 (2003).
54. Wiesinger, F. *et al.* Parallel imaging performance as a function of field strength? An experimental investigation using electrodynamic scaling. *Magnetic Resonance in Medicine* **52**, 953–964 (2004).
55. Wiesinger, F., Boesiger, P. & Pruessmann, K. P. Electrodynamics and ultimate SNR in parallel MR imaging. *Magnetic Resonance in Medicine* **52**, 376–390 (2004).
56. Tractnig, S. *et al.* Key clinical benefits of neuroimaging at 7 T. *NeuroImage* **168**, 477–489 (2018).
57. Tallantyre, E. C. *et al.* 3 Tesla and 7 Tesla MRI of multiple sclerosis cortical lesions. *Journal of Magnetic Resonance Imaging* **32**, 971–977 (2010).
58. Tallantyre, E. C. *et al.* Ultra-high-field imaging distinguishes MS lesions from asymptomatic white matter lesions. *Neurology* **76**, 534–539 (2011).
59. Sati, P. Diagnosis of multiple sclerosis through the lens of ultra-high-field MRI. *Journal of Magnetic Resonance* **291**, 101–109 (2018).

60. Dula, A. N. *et al.* Magnetic resonance imaging of the cervical spinal cord in multiple sclerosis at 7T. *Multiple Sclerosis Journal* **22**, 320–328 (2016).
61. Isaacs, B. R. *et al.* 3 versus 7 Tesla magnetic resonance imaging for parcellations of subcortical brain structures in clinical settings. *PLOS ONE* **15**, e0236208 (2020).
62. Spincemaille, P. *et al.* Quantitative Susceptibility Mapping: MRI at 7T versus 3T. *Journal of Neuroimaging* **30**, 65–75 (2020).
63. Okada, T., Yamada, H., Ito, H., Yonekura, Y. & Sadato, N. Magnetic field strength increase yields significantly greater contrast-to-noise ratio increase: Measured using BOLD contrast in the primary visual area1. *Academic Radiology* **12**, 142–147 (2005).
64. Ogawa, S. *et al.* Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. *Biophysical Journal* **64**, 803–812 (1993).
65. Van der Zwaag, W. *et al.* fMRI at 1.5, 3 and 7 T: characterising BOLD signal changes. *Neuroimage* **47**, 1425–1434 (2009).
66. Gati, J. S., Menon, R. S., Uğurbil, K. & Rutt, B. K. Experimental determination of the BOLD field strength dependence in vessels and tissue. *Magnetic resonance in medicine* **38**, 296–302 (1997).
67. Yacoub, E. *et al.* Imaging brain function in humans at 7 Tesla. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **45**, 588–594 (2001).
68. Uğurbil, K. Imaging at ultrahigh magnetic fields: History, challenges, and solutions. *NeuroImage* **168**, 7–32 (2018).
69. Van Der Zwaag, W. *et al.* Digit somatotopy in the human cerebellum: A 7T fMRI study. *NeuroImage* **67**, 354–362 (2013).
70. Boillat, Y., Bazin, P.-L. & van der Zwaag, W. Whole-body somatotopic maps in the cerebellum revealed with 7T fMRI. *NeuroImage* **211**, 116624 (2020).
71. Hoffmann, M. B., Stadler, J., Kanowski, M. & Speck, O. Retinotopic mapping of the human visual cortex at a magnetic field strength of 7 T. *Clinical neurophysiology* **120**, 108–116 (2009).

72. Wang, M. *et al.* Simultaneous detection and separation of hyperacute intracerebral hemorrhage and cerebral ischemia using amide proton transfer MRI. *Magnetic Resonance in Medicine* **74**, 42–50 (2015).
73. Singh, A. *et al.* Chemical exchange saturation transfer magnetic resonance imaging of human knee cartilage at 3 T and 7 T. *Magnetic Resonance in Medicine* **68**, 588–594 (2012).
74. Donahue, M. J. *et al.* Assessment of lymphatic impairment and interstitial protein accumulation in patients with breast cancer treatment-related lymphedema using CEST MRI. *Magnetic Resonance in Medicine* **75**, 345–355 (2016).
75. Kopanoglu, E. & Constable, R. T. Radiofrequency pulse design using nonlinear gradient magnetic fields. *Magnetic Resonance in Medicine* **74**, 826–839 (2015).
76. Yang, Q. X. *et al.* Analysis of wave behavior in lossy dielectric samples at high field. *Magnetic Resonance in Medicine* **47**, 982–989 (2002).
77. Hoult, D. I. Sensitivity and power deposition in a high-field imaging experiment. *Journal of Magnetic Resonance Imaging* **12**, 46–67 (2000).
78. Collins, C. M., Liu, W., Schreiber, W., Yang, Q. X. & Smith, M. B. Central brightening due to constructive interference with, without, and despite dielectric resonance. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **21**, 192–196 (2005).
79. Van De Moortele, P.-F. *et al.* B1 destructive interferences and spatial phase patterns at 7 T with a head transceiver array coil. *Magnetic Resonance in Medicine* **54**, 1503–1518 (2005).
80. Huang, S. Y. *et al.* Body MR imaging: artifacts, k-space, and solutions. *Radiographics* **35**, 1439 (2015).
81. Port, J. D. & Pomper, M. G. Quantification and minimization of magnetic susceptibility artifacts on GRE images. *Journal of Computer Assisted Tomography* **24**, 958–964 (2000).
82. Feinberg, D. A. *et al.* Inner volume MR imaging: technical concepts and their application. *Radiology* **156**, 743–747 (1985).

83. Mitsouras, D., Mulkern, R. V. & Rybicki, F. J. Strategies for inner volume 3D fast spin echo magnetic resonance imaging using nonselective refocusing radio frequency pulses a. *Medical physics* **33**, 173–186 (2006).
84. Vinding, M., Lund, T., Stockmann, J. & Guerin, B. INSTANT (INtegrated Shimming and Tip-Angle NormalizaTion): 3D flip-angle mitigation using joint optimization of RF and shim array currents. *Proc. Intl. Soc. Mag. Reson. Med.*, 28, 0612 (2020).
85. Webb, A. Dielectric materials in magnetic resonance. *Concepts in Magnetic Resonance Part A* **38**, 148–184 (2011).
86. Teeuwisse, W. M., Brink, W. M. & Webb, A. G. Quantitative assessment of the effects of high-permittivity pads in 7 Tesla MRI of the brain. *Magnetic Resonance in Medicine* **67**, 1285–1293 (2012).
87. Vaidya, M. V. *et al.* Improved detection of fMRI activation in the cerebellum at 7T with dielectric pads extending the imaging region of a commercial head coil. *Journal of Magnetic Resonance Imaging* **48**, 431–440 (2018).
88. Garcia, M. M. *et al.* Investigating the influence of dielectric pads in 7T magnetic resonance imaging—simulated and experimental assessment. *Current Directions in Biomedical Engineering* **6**, 24–27 (2020).
89. Venkatesan, R., Lin, W. & Haacke, E. M. Accurate determination of spin-density and T1 in the presence of RF-field inhomogeneities and flip-angle miscalibration. *Magnetic Resonance in Medicine* **40**, 592–602 (1998).
90. Glasser, M. F. *et al.* Empirical Transmit Field Bias Correction of T1w/T2w Myelin Maps. *bioRxiv* (2022).
91. Emmenegger, T. M. *et al.* The Influence of Radio-Frequency Transmit Field Inhomogeneities on the Accuracy of G-ratio Weighted Imaging. *Frontiers in Neuroscience*, 770 (2021).
92. Pauly, J., Nishimura, D. & Macovski, A. A k-space analysis of small-tip-angle excitation. *Journal of Magnetic Resonance* **81**, 43–56 (1989).
93. Grissom, W. A., Yip, C.-Y., Wright, S. M., Fessler, J. A. & Noll, D. C. Additive angle method for fast large-tip-angle RF pulse design in parallel excitation. *Magnetic Resonance in Medicine* **59**, 779–787 (2008).

94. Sutton, B. P., Noll, D. C. & Fessler, J. A. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE transactions on Medical Imaging* **22**, 178–188 (2003).
95. Setsompop, K., Wald, L., Alagappan, V., Gagoski, B. & Adalsteinsson, E. Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels. *Magnetic Resonance in Medicine* **59**, 908–915 (2008).
96. Tannús, A. & Garwood, M. Adiabatic pulses. *NMR in Biomedicine* **10**, 423–434 (1997).
97. Uğurbil, K., Garwood, M., Rath, A. R. & Bendall, M. R. Amplitude and frequency/phase-modulated refocusing pulses that induce plane rotations even in the presence of inhomogeneous B1 fields. *Journal of Magnetic Resonance* **78**, 472–497 (1988).
98. Yip, C.-Y., Fessler, J. A. & Noll, D. C. Advanced three-dimensional tailored RF pulse for signal recovery in T2*-weighted functional magnetic resonance imaging. *Magnetic Resonance in Medicine* **56**, 1050–1059 (2006).
99. Cloos, M. A. *et al.* kT-points: Short three-dimensional tailored RF pulses for flip-angle homogenization over an extended volume. *Magnetic Resonance in Medicine* **67**, 72–80 (2012).
100. Eggenschwiler, F., O'Brien, K. R., Gruetter, R. & Marques, J. P. Improving T2-weighted imaging at high field through the use of kT-points. *Magnetic Resonance in Medicine* **71**, 1478–1488 (2014).
101. Eggenschwiler, F., O'Brien, K. R., Gallichan, D., Gruetter, R. & Marques, J. P. 3D T2-weighted imaging at 7T using dynamic kT-points on single-transmit MRI systems. *Magnetic Resonance Materials in Physics, Biology and Medicine* **29**, 347–358 (2016).
102. Malik, S. J., Padormo, F., Price, A. N. & Hajnal, J. V. Spatially resolved extended phase graphs: modeling and design of multipulse sequences with parallel transmission. *Magnetic Resonance in Medicine* **68**, 1481–1494 (2012).
103. Sotgiu, A. & Hyde, J. S. High-order coils as transmitters for NMR imaging. *Magnetic Resonance in Medicine* **3**, 55–62 (1986).

104. Ibrahim, T. S., Lee, R., Baertlein, B. A., Kangarlu, A. & Robitaille, P.-M. L. Application of finite difference time domain method for the design of birdcage RF head coils using multi-port excitations. *Magnetic Resonance Imaging* **18**, 733–742 (2000).
105. Adriany, G. *et al.* Transmit and receive transmission line arrays for 7 Tesla parallel imaging. *Magnetic Resonance in Medicine* **53**, 434–445 (2005).
106. Hetherington, H. P., Avdievich, N. I., Kuznetsov, A. M. & Pan, J. W. RF shimming for spectroscopic localization in the human brain at 7 T. *Magnetic Resonance in Medicine* **63**, 9–19 (2010).
107. Ullmann, P. *et al.* Experimental analysis of parallel excitation using dedicated coil setups and simultaneous RF transmission on multiple channels. *Magnetic Resonance in Medicine* **54**, 994–1001 (2005).
108. Zhu, Y. & Giaquinto, R. Improving flip angle uniformity with parallel excitation. *Proc Int Soc Magn Reson Med* **13**, 2752 (2005).
109. Lee, D., Lustig, M., Grissom, W. A. & Pauly, J. M. Time-optimal design for multidimensional and parallel transmit variable-rate selective excitation. *Magnetic Resonance in Medicine* **61**, 1471–1479 (2009).
110. Malik, S. J. & Hajnal, J. V. Phase relaxed localized excitation pulses for inner volume fast spin echo imaging. *Magnetic Resonance in Medicine* **76**, 848–861 (2016).
111. Meiboom, S. & Gill, D. Modified spin-echo method for measuring nuclear relaxation times. *Review of Scientific Instruments* **29**, 688–691 (1958).
112. Conolly, S., Nishimura, D. & Macovski, A. Optimal control solutions to the magnetic resonance selective excitation problem. *IEEE transactions on Medical Imaging* **5**, 106–115 (1986).
113. Xu, D., King, K. F., Zhu, Y., McKinnon, G. C. & Liang, Z.-P. Designing multichannel, multidimensional, arbitrary flip angle RF pulses using an optimal control approach. *Magnetic Resonance in Medicine* **59**, 547–560 (2008).
114. Cao, Z., Donahue, M. J., Ma, J. & Grissom, W. A. Joint design of large-tip-angle parallel RF pulses and blipped gradient trajectories. *Magnetic Resonance in Medicine* **75**, 1198–1208 (2016).

115. Vinding, M. S., Guérin, B., Vosegaard, T. & Nielsen, N. C. Local SAR, global SAR, and power-constrained large-flip-angle pulses with optimal control and virtual observation points. *Magnetic Resonance in Medicine* **77**, 374–384 (2017).
116. Grissom, W. A., Xu, D., Kerr, A. B., Fessler, J. A. & Noll, D. C. Fast large-tip-angle multidimensional and parallel RF pulse design in MRI. *IEEE transactions on Medical Imaging* **28**, 1548–1559 (2009).
117. Pauly, J., Le Roux, P., Nishimura, D. & Macovski, A. Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm (NMR imaging). *IEEE transactions on Medical Imaging* **10**, 53–65 (1991).
118. Ding, B., Rua, C. & Rodgers, C. Comparison of Nova 1Tx vs 8Tx head coils for routine 7T neuroimaging. *Proc. Intl. Soc. Mag. Reson. Med.*, 30, 2094 (2022).
119. Graf, C., Ding, B., Rua, C. & Rodgers, C. Brainstem 7T1H-MRS: comparison between single-channel (1Tx32Rx) and parallel transmit (8Tx32Rx) coils. *Proc. Intl. Soc. Mag. Reson. Med.*, 30, 2095 (2022).
120. Deniz, C. M. *et al.* Maximum efficiency radiofrequency shimming: Theory and initial application for hip imaging at 7 tesla. *Magnetic Resonance in Medicine* **69**, 1379–1388 (2013).
121. Metzger, G. J. *et al.* LocalB1+ shimming for prostate imaging with transceiver arrays at 7T based on subject-dependent transmit phase measurements. *Magnetic Resonance in Medicine* **59**, 396–409 (2008).
122. Curtis, A. T., Gilbert, K. M., Klassen, L. M., Gati, J. S. & Menon, R. S. Slice-by-slice B1+ shimming at 7T. *Magnetic Resonance in Medicine* **68**, 1109–1116 (2012).
123. Gras, V., Vignaud, A., Amadon, A., Le Bihan, D. & Boulant, N. Universal pulses: A new concept for calibration-free parallel transmission. *Magnetic Resonance in Medicine* **77**, 635–643 (2017).
124. Mao, W., Smith, M. B. & Collins, C. M. Exploring the limits of RF shimming for high-field MRI of the human head. *Magnetic Resonance in Medicine* **56**, 918–922 (2006).

125. Herrler, J. *et al.* Fast online-customized (FOCUS) parallel transmission pulses: A combination of universal pulses and individual optimization. *Magnetic Resonance in Medicine* **85**, 3140–3153 (2021).
126. Tomi-Tricot, R. *et al.* SmartPulse, a machine learning approach for calibration-free dynamic RF shimming: Preliminary study in a clinical environment. *Magnetic Resonance in Medicine* **82**, 2016–2031 (2019).
127. Deniz, C. M., Vaidya, M. V., Sodickson, D. K. & Lattanzi, R. Radiofrequency energy deposition and radiofrequency power requirements in parallel transmission with increasing distance from the coil to the sample. *Magnetic Resonance in Medicine* **75**, 423–432 (2016).
128. Padormo, F., Beqiri, A., Malik, S. J. & Hajnal, J. V. RF system calibration for global Q matrix determination. *Magnetic Resonance Imaging* **34**, 690–693 (2016).
129. Graesslin, I. *et al.* A specific absorption rate prediction concept for parallel transmission MR. *Magnetic Resonance in Medicine* **68**, 1664–1674 (2012).
130. Kopanoglu, E., Deniz, C. M., Erturk, M. A. & Wise, R. G. Specific absorption rate implications of within-scan patient head motion for ultra-high field MRI. *Magnetic Resonance in Medicine* (2020).
131. Eichfelder, G. & Gebhardt, M. Local specific absorption rate control for parallel transmission by virtual observation points. *Magnetic Resonance in Medicine* **66**, 1468–1476 (2011).
132. Meliadó, E. *et al.* A deep learning method for image-based subject-specific local SAR assessment. *Magnetic Resonance in Medicine* **83**, 695–711 (2020).
133. Brink, W., Staring, M., Remis, R. & Webb, A. Fast Subject-Specific Local SAR and B1+ Prediction for PTx at 7T using only an Initial Localizer Scan. *Proc. Intl. Soc. Mag. Reson. Med.*, *30*, 0387 (2022).
134. Wu, X. *et al.* Comparison of RF body coils for MRI at 3T: a simulation study using parallel transmission on various anatomical targets. *NMR in Biomedicine* **28**, 1332–1344 (2015).

135. Hoyos-Idrobo, A., Weiss, P., Massire, A., Amadon, A. & Boulant, N. On variant strategies to solve the magnitude least squares optimization problem in parallel transmission pulse design and under strict SAR and power constraints. *IEEE transactions on Medical Imaging* **33**, 739–748 (2013).
136. Pendse, M., Stara, R., Mehdi Khalighi, M. & Rutt, B. IMPULSE: A scalable algorithm for design of minimum specific absorption rate parallel transmit RF pulses. *Magnetic Resonance in Medicine* **81**, 2808–2822 (2019).
137. Gras, V., Luong, M., Amadon, A. & Boulant, N. Joint design of kT-points trajectories and RF pulses under explicit SAR and power constraints in the large flip angle regime. *Journal of Magnetic Resonance* **261**, 181–189 (2015).
138. Winter, L. *et al.* Design and evaluation of a hybrid radiofrequency applicator for magnetic resonance imaging and RF induced hyperthermia: electromagnetic field simulations up to 14.0 Tesla and proof-of-concept at 7.0 Tesla. *PloS one* **8**, e61661 (2013).
139. Guérin, B. *et al.* Computation of ultimate SAR amplification factors for radiofrequency hyperthermia in non-uniform body models: impact of frequency and tumour location. *International Journal of Hyperthermia* **34**, 87–100 (2018).
140. Ertan, K., De Bever, J., Pendse, M., Decuzzi, P. & Rutt, B. Simulation of Focused RF Heating using a High-Channel-Count RF Array and a Maximum SAR Algorithm. *Proc. Intl. Soc. Mag. Reson. Med.*, 28, 4055 (2020).
141. Eryaman, Y., Akin, B. & Atalar, E. Reduction of implant RF heating through modification of transmit coil electric field. *Magnetic Resonance in Medicine* **65**, 1305–1313 (2011).
142. Eryaman, Y. *et al.* Parallel transmit pulse design for patients with deep brain stimulation implants. *Magnetic Resonance in Medicine* **73**, 1896–1903 (2015).
143. Gosselin, M.-C. *et al.* Development of a new generation of high-resolution anatomical models for medical device evaluation: the Virtual Population 3.0. *Physics in Medicine & Biology* **59**, 5287–5303 (2014).
144. Hasgall, P., Neufeld, E., Gosselin, M.-C., Klingenbock, A. & Kuster, N. IT'IS database for thermal and electromagnetic parameters of biological tissues, Version 2.4 (2013).

145. Wolf, S., Diehl, D., Gebhardt, M., Mallow, J. & Speck, O. SAR simulations for high-field MRI: How much detail, effort, and accuracy is needed? *Magnetic Resonance in Medicine* **69**, 1157–1168 (2013).
146. Zaitsev, M., Maclaren, J. & Herbst, M. Motion artifacts in MRI: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging* **42**, 887–901 (2015).
147. Godenschweger, F. *et al.* Motion correction in MRI of the brain. *Physics in Medicine & Biology* **61**, 32–56 (2016).
148. Reuter, M. *et al.* Head motion during MRI acquisition reduces gray matter volume and thickness estimates. *NeuroImage* **107**, 107–115 (2015).
149. Seto, E. *et al.* Quantifying head motion associated with motor tasks used in fMRI. *NeuroImage* **14**, 284–297 (2001).
150. Tisdall, M. D. *et al.* Prospective motion correction with volumetric navigators (vNavs) reduces the bias and variance in brain morphometry induced by subject motion. *NeuroImage* **127**, 11–22 (2016).
151. Graham, M. S., Drobnjak, I., Jenkinson, M. & Zhang, H. Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI. *PloS one* **12**, e0185647 (2017).
152. Eichhorn, H. *et al.* Characterisation of Children’s Head Motion for Magnetic Resonance Imaging With and Without General Anaesthesia. *Frontiers in Radiology* **1**, 789632 (2021).
153. Mattern, H., Sciarra, A., Lüsebrink, F., Acosta-Cabronero, J. & Speck, O. Prospective motion correction improves high-resolution quantitative susceptibility mapping at 7T. *Magnetic Resonance in Medicine* **81**, 1605–1619 (2019).
154. Gretsch, F., Mattern, H., Gallichan, D. & Speck, O. Fat navigators and Moiré phase tracking comparison for motion estimation and retrospective correction. *Magnetic Resonance in Medicine* **83**, 83–93 (2020).
155. Sciarra, A. *et al.* Quantitative evaluation of prospective motion correction in healthy subjects at 7T MRI. *Magnetic Resonance in Medicine* **87**, 646–657 (2022).

156. Laustsen, M., Andersen, M., Xue, R., Madsen, K. H. & Hanson, L. G. Tracking of rigid head motion during MRI using an EEG system. *Magnetic Resonance in Medicine* (2022).
157. Slipsager, J. M. *et al.* Quantifying the Financial Savings of Motion Correction in Brain MRI: A Model-Based Estimate of the Costs Arising From Patient Head Motion and Potential Savings From Implementation of Motion Correction. *Journal of Magnetic Resonance Imaging* **52**, 731–738 (2020).
158. Andre, J. *et al.* Toward Quantifying the Prevalence, Severity, and Cost Associated With Patient Motion During Clinical MR Examinations. *Journal of the American College of Radiology : JACR* **12** (2015).
159. Gedamu, E. L. & Gedamu, A. Subject movement during multislice interleaved MR acquisitions: Prevalence and potential effect on MRI-derived brain pathology measurements and multicenter clinical trials of therapeutics for multiple sclerosis. *Journal of Magnetic Resonance Imaging* **36**, 332–343 (2012).
160. Vanderby, S. A., Babyn, P. S., Carter, M. W., Jewell, S. M. & McKeever, P. D. Effect of anesthesia and sedation on pediatric MR imaging patient flow. *Radiology* **256**, 229–237 (2010).
161. Runge, S. B., Christensen, N. L., Jensen, K. & Jensen, I. E. Children centered care: Minimizing the need for anesthesia with a multi-faceted concept for MRI in children aged 4–6. *European journal of radiology* **107**, 183–187 (2018).
162. Lee, C. C. *et al.* Real-time adaptive motion correction in functional MRI. *Magnetic Resonance in Medicine* **36**, 436–444 (1996).
163. Lee, C. C. *et al.* A prospective approach to correct for inter-image head rotation in fMRI. *Magnetic Resonance in Medicine* **39**, 234–243 (1998).
164. Zaitsev, M., Dold, C., Sakas, G., Hennig, J. & Speck, O. Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. *NeuroImage* **31**, 1038–1050 (2006).
165. Bookwalter, C. A., Griswold, M. A. & Duerk, J. L. Multiple overlapping k-space junctions for investigating translating objects (MOJITO). *IEEE transactions on Medical Imaging* **29**, 339–349 (2009).

166. Irfanoglu, M. O., Walker, L., Sarlls, J., Marenco, S. & Pierpaoli, C. Effects of image distortions originating from susceptibility variations and concomitant fields on diffusion MRI tractography results. *NeuroImage* **61**, 275–288 (2012).
167. Jezzard, P. & Balaban, R. S. Correction for geometric distortion in echo planar images from B₀ field variations. *Magnetic Resonance in Medicine* **34**, 65–73 (1995).
168. Chang, H. & Fitzpatrick, J. M. A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE transactions on Medical Imaging* **11**, 319–329 (1992).
169. Gallichan, D., Andersson, J. L., Jenkinson, M., Robson, M. D. & Miller, K. L. Reducing distortions in diffusion-weighted echo planar imaging with a dual-echo blip-reversed sequence. *Magnetic Resonance in Medicine* **64**, 382–390 (2010).
170. Faraji-Dana, Z., Tam, F., Chen, J. J. & Graham, S. J. A robust method for suppressing motion-induced coil sensitivity variations during prospective correction of head motion in fMRI. *Magnetic Resonance Imaging* **34**, 1206–1219 (2016).
171. Bammer, R. *et al.* Impact of Motion on Parallel Transmission. *Proc. Intl. Soc. Mag. Reson. Med.*, *19*, 4590 (2010).
172. Doran, E. *et al.* Effects of Respiration on B₁₊ and SAR in Whole-Body Imaging at 7 Tesla. *Proc. Intl. Soc. Mag. Reson. Med.*, *26*, 4179 (2018).
173. Kopanoglu, E., Plumley, A., Erturk, M. A., Deniz, C. & Wise, R. Implications of within-scan patient head motion on B₁₊ homogeneity and Specific Absorption Rate at 7T. *Proc. Intl. Soc. Mag. Reson. Med.*, *27* (2019).
174. Pawar, K., Chen, Z., Shah, N. J. & Egan, G. F. MoCoNet: motion correction in 3D MPRAGE images using a convolutional neural network approach. *arXiv:1807.10831* (2018).
175. Ghaffari, M., Pawar, K. & Oliver, R. Brain MRI motion artifact reduction using 3D conditional generative adversarial networks on simulated motion. *IEEE*, 1–7 (2021).
176. Lee, S., Jung, S., Jung, K.-J. & Kim, D.-H. Deep Learning in MR Motion Correction: a Brief Review and a New Motion Simulation Tool (view2Dmotion). *Investigative Magnetic Resonance Imaging* **24**, 196–206 (2020).

177. Johnson, P. M. & Drangova, M. Conditional generative adversarial network for 3D rigid-body motion correction in MRI. *Magnetic Resonance in Medicine* **82**, 901–910 (2019).
178. Gallichan, D., Marques, J. P. & Gruetter, R. Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magnetic Resonance in Medicine* **75**, 1030–1039 (2016).
179. Kopanoglu, E., Wang, H., Galiana, G. & Constable, R. T. Motion Tracking using Nonlinear Gradient Fields: Experimental Verification and Oblique Slices. *Proc. Intl. Soc. Mag. Reson. Med.*, *24*, 4261 (2016).
180. Kober, T., Marques, J. P., Gruetter, R. & Krueger, G. Head motion detection using FID navigators. *Magnetic Resonance in Medicine* **66**, 135–143 (2011).
181. Wallace, T. E., Afacan, O., Waszak, M., Kober, T. & Warfield, S. K. Head motion measurement and correction using FID navigators. *Magnetic Resonance in Medicine* **81**, 258–274 (2019).
182. Qin, L. *et al.* Prospective head-movement correction for high-resolution MRI using an in-bore optical tracking system. *Magnetic Resonance in Medicine* **62**, 924–934 (2009).
183. Stucht, D. *et al.* Highest resolution in vivo human brain MRI using prospective motion correction. *PloS one* **10**, e0133921 (2015).
184. Digiacomo, P. *et al.* A within-coil optical prospective motion-correction system for brain imaging at 7T. *Magnetic Resonance in Medicine*, 1661–1671 (2020).
185. Pipe, J. G. Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging. *Magnetic Resonance in Medicine* **42**, 963–969 (1999).
186. Brackenier, Y. *et al.* Data-driven motion-corrected brain MRI incorporating pose-dependent B0 fields. *Magnetic Resonance in Medicine*, 817–831 (2022).
187. Haskell, M. W., Lahiri, A., Nielsen, J.-F., Fessler, J. A. & Noll, D. C. FieldMapNet MRI: Learning-based mapping from single echo time BOLD fMRI data to fieldmaps with model-based image reconstruction. *Proc. Intl. Soc. Mag. Reson. Med.*, *30*, 0235 (2022).
188. Ward, H. A., Riederer, S. J. & Jack, C. R. Real-time autoshimming for echo planar timecourse imaging. *Magnetic Resonance in Medicine* **48**, 771–780 (2002).

189. Kopanoglu, E. Near real-time parallel-transmit pulse design. *Proc. Intl. Soc. Mag. Reson. Med.*, 26, 3392 (2018).
190. Ajanovic, A., Hajnal, J., Tomi-Tricot, R. & Malik, S. Motion and Pose Variability of SAR Estimation with Parallel Transmission at 7T. *Proc. Intl. Soc. Mag. Reson. Med.*, 29, 2487 (2021).
191. Schön, N. *et al.* Impact of respiration on B1+ field and SAR distribution at 7T using a novel EM simulation setup. *Proc. Intl. Soc. Mag. Reson. Med.*, 28, 1120 (2020).
192. Lee, J., Gebhardt, M., Wald, L. L. & Adalsteinsson, E. Local SAR in parallel transmission pulse design. *Magnetic Resonance in Medicine* **67**, 1566–1578 (2012).
193. Graesslin, I. *et al.* Real-Time Global and Local SAR Monitoring for Parallel Transmission Systems. *Proc. Intl. Soc. Mag. Reson. Med.*, 17, 0302 (2009).
194. Meliaddò, E. F., Sbrizzi, A., van den Berg, C. A., Luijten, P. R. & Raaijmakers, A. J. Real-time assessment of potential peak local specific absorption rate value without phase monitoring: Trigonometric maximization method for worst-case local specific absorption rate determination. *Magnetic Resonance in Medicine* **85**, 3420–3433 (2021).
195. Katscher, U., Röhrs, J. & Börnert, P. Basic considerations on the impact of the coil array on the performance of Transmit SENSE. *Magnetic Resonance Materials in Physics, Biology and Medicine* **18**, 81–88 (2005).
196. Grissom, W. A. *et al.* Advancing RF pulse design using an open-competition format: Report from the 2015 ISMRM challenge. *Magnetic Resonance in Medicine*, 1352–1361 (2016).
197. Graesslin, I. *et al.* Local SAR constrained hotspot reduction by temporal averaging. *Proc. Intl. Soc. Mag. Reson. Med.*, 18, 4932 (2010).
198. Guerin, B., Adalsteinsson, E. & Wald, L. L. Local SAR reduction in multi-slice pTx via “SAR hopping” between excitations. *Proc. Intl. Soc. Mag. Reson. Med.*, 20, 642 (2012).
199. Plumley, A. *Motion in a parallel-transmit array coil: impact on flip angle as a function of RF pulse sequence and design parameters at 7 Tesla (unpublished Master’s Thesis)* MA thesis (Cardiff University, 2018).

200. Winarno, N., Firdaus, R. & Afifah, R. The effect of conductivity and permittivity on propagation and attenuation of waves using FDTD. *Materials Physics and Mechanics* **42**, 617–624 (2019).
201. Ipek, Ö., Raaijmakers, A. J., Lagendijk, J. J., Luijten, P. R. & Van Den Berg, C. A. T. Intersubject local SAR variation for 7T prostate MR imaging with an eight-channel single-side adapted dipole antenna array. *Magnetic Resonance in Medicine* **71**, 1559–1567 (2014).
202. Meliadó, E. F., van den Berg, C. A., Luijten, P. R. & Raaijmakers, A. J. Intersubject specific absorption rate variability analysis through construction of 23 realistic body models for prostate imaging at 7T. *Magnetic Resonance in Medicine* **81**, 2106–2119 (2019).
203. De Greef, M., Ipek, O., Raaijmakers, A. J. E., Crezee, J. & Van Den Berg, C. A. T. Specific absorption rate intersubject variability in 7T parallel transmit MRI of the head. *Magnetic Resonance in Medicine* **69**, 1476–1485 (2013).
204. Petzold, J., Ittermann, B. & Seifert, F. On the complexity of pTx systems in SAR assessment. *Proc. Intl. Soc. Mag. Reson. Med.*, *30*, 2555 (2022).
205. Mandija, S., Meliadó, E. F., Huttinga, N. R. F., Luijten, P. R. & Van Den Berg, C. A. T. Opening a new window on MR-based Electrical Properties Tomography with deep learning. *Scientific Reports* **9**, 8895 (2019).
206. Meliadó, E. F. *et al.* Conditional safety margins for less conservative peak local SAR assessment: A probabilistic approach. *Magnetic Resonance in Medicine*, 3379–3395 (2020).
207. Kopanoglu, E. Actual patient position versus safety models—specific absorption rate implications of initial head position at ultrahigh field MRI. *NMR in Biomedicine* (2022).
208. Aigner, C. S., Dietrich, S. & Schmitter, S. Respiration induced B₁₊ changes and their impact on universal and tailored 3D kT-point parallel transmission pulses for 7T cardiac imaging. *Magnetic Resonance in Medicine* **87**, 2862–2871 (2022).
209. Vinding, M. S., Aigner, C. S., Schmitter, S. & Lund, T. E. DeepControl: 2DRF pulses facilitating inhomogeneity and B₀ off-resonance compensation in vivo at 7 T. *Magnetic Resonance in Medicine* **85**, 3308–3317 (2021).

210. Afacan, O., Wallace, T. E. & Warfield, S. K. Retrospective correction of head motion using measurements from an electromagnetic tracker. *Magnetic Resonance in Medicine* **83**, 427–437 (2020).
211. Bortolotti, L., Mougin, O. & Bowtell, R. Measurement of head motion using a field camera in a 7T scanner. *Proc. Intl. Soc. Mag. Reson. Med.*, 28, 0464 (2020).
212. Isola, P., Zhu, J., Zhou, T. & Efros, A. A. Image-to-Image Translation with Conditional Adversarial Networks. *CoRR* (2016).
213. Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *CoRR* (2015).
214. Rumelhart, D. E., Hinton, G. E., McClelland, J. L., *et al.* A general framework for parallel distributed processing. *Parallel distributed processing: Explorations in the microstructure of cognition* **1**, 26 (1986).
215. McClelland, J. L., Rumelhart, D. E. & Hinton, G. E. The appeal of parallel distributed processing. *MIT Press, Cambridge MA* **3**, 44 (1986).
216. Kingma, D. & Lei Ba, J. Adam: A method for stochastic optimization. *Proceedings of the International Conference on Learning Representations* (2015).
217. Goodfellow, I. *et al.* Generative adversarial nets. *Advances in neural information processing systems* **27** (2014).
218. Yoon, J. *et al.* Quantitative susceptibility mapping using deep neural network: QSMnet. *NeuroImage* **179**, 199–206 (2018).
219. Bollmann, S. *et al.* DeepQSM - using deep learning to solve the dipole inversion for quantitative susceptibility mapping. *NeuroImage* **195**, 373–383 (2019).
220. Wu, Y., Ma, Y., Du, J. & Xing, L. Accelerating quantitative MR imaging with the incorporation of B1 compensation using deep learning. *Magnetic Resonance Imaging* **72**, 78–86 (2020).
221. Abbasi-Rad, S. *et al.* Improving FLAIR SAR efficiency at 7T by adaptive tailoring of adiabatic pulse power through deep learning estimation. *Magnetic Resonance in Medicine*, 2462–2476 (2020).
222. Martin Abadi *et al.* *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems* Software available from tensorflow.org. 2015.

223. Radford, A., Metz, L. & Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015).
224. Liu, J., de Zwart, J. A., van Gelderen, P., Murphy-Boesch, J. & Duyn, J. H. Effect of head motion on MRI B 0 field distribution. *Magnetic Resonance in Medicine* **80**, 2538–2548 (2018).
225. Alon, L., Deniz, C. M., Sodickson, D. K. & Zhu, Y. Do constraints on B1+ also constrain E and SAR in high field MR? *Proc. Intl. Soc. Mag. Reson. Med.*, 19, 0491 (2011).
226. Lv, J. *et al.* Transfer learning enhanced generative adversarial networks for multi-channel MRI reconstruction. *Computers in Biology and Medicine* **134**, 104504 (2021).
227. Frégier, Y. & Gouray, J.-B. Mind2Mind: transfer learning for GANs. *Springer*, 851–859 (2021).
228. Maclaren, J., Herbst, M., Speck, O. & Zaitsev, M. Prospective motion correction in brain imaging: A review. *Magnetic Resonance in Medicine* **69**, 621–636 (2013).
229. Cusack, R. *et al.* Differences in the spatial and temporal patterns of head motion during MRI of adults and infants. *bioRxiv*, 114447 (2017).
230. Steensma, B. R. *et al.* SAR and temperature distributions in a database of realistic human models for 7 T cardiac imaging. *NMR in Biomedicine* **34**, e4525 (2021).
231. Wood, S. *et al.* Design and fabrication of a realistic anthropomorphic heterogeneous head phantom for MR purposes. *PLOS ONE* **12**, e0192794 (2017).
232. Brink, W. M., Wu, Z. & Webb, A. G. A simple head-sized phantom for realistic static and radiofrequency characterization at high fields. *Magnetic Resonance in Medicine* **80**, 1738–1745 (2018).
233. Choi, Y. *et al.* Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8789–8797 (2018).