Optimisation and Acceleration of Second-order Dynamic Shimming at 7T



Yuhang Shi

Keble College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2017

This thesis is dedicated to
my parents

Acknowledgements

I want to express my sincere gratitude to my supervisors Professor Stuart Clare, Doctor Johanna Vannesjo, and Professor Karla Miller for valuable suggestions and guidance since the inception of my DPhil studies.

I want to thank all of my colleagues and support staff at Wellcome Centre for Integrative Neuroimaging, University of Oxford, for their help during my DPhil studies.

Finally, I want to thank my parents for supporting my studies. Without your encouragement and love, I would not get chance to study at University of Oxford and pursue a DPhil degree.

Contents

Prologue **1**

1 Introduction **4**

1.1 MR physics 5

1.1.1 Magnetic resonance 5

1.1.2 Relaxation 7

1.2 MR image formation 11

1.2.1 Spatial encoding using gradient fields 11

1.2.2 MR pulse sequence: GRE sequence 14

1.2.3 MR hardware 19

1.3 Linking B_0 homogeneity to image quality 21

1.3.1 Susceptibility induced B_0 field inhomogeneity 21

1.3.2 B_0 field inhomogeneity related image artefacts 24

1.4 Shimming and B_0 field mapping 28

1.4.1 Theory of field mapping 29

1.4.2 Field mapping techniques 30

1.4.3 Field map post-processing 32

1.5 Advanced B_0 shimming techniques 34

1.5.1 Active shimming 34

1.5.2 Local shimming 36

1.5.3 Dynamic shimming 39

2 Calibration and implementation of dynamic shimming **41**

2.1 MR scanner and dynamic shimming hardware configuration 42

2.2 Static shim field characterisation 44

2.2.1 Methods 45

2.2.2 Discussion 50

2.3 Eddy current compensation 51

2.3.1 Eddy current model 54
2.3.2 Image-based eddy current measurement 58
2.3.3 SIRF-based eddy current measurement 63
2.3.4 SIRF-based method vs. image-based method 73
2.3.5 Discussion 73
2.4 Preliminary implementation of dynamic shimming 77
2.4.1 Dynamic shimming implementation procedures 78
2.4.2 Experimental validation 85
2.4.3 Simulation with a field map database 88
2.4.4 Discussion 90

3 Dynamic shimming optimisation algorithms 93
3.1 Perspectives on current constraints and regularisation methods 94
3.1.1 Theory of regularisation 96
3.1.2 Regularisation performance evaluation 101
3.2 Regularised second-order dynamic shimming 110
3.2.1 Tikhonov regularisation with the L-curve method 112
3.2.2 Performance evaluation of regularised optimisation algorithms 115
3.2.3 Comparison of LSQlin, TSVD-H and Tikh-L 118
3.2.4 Comparison of different optimisation methods 126
3.2.5 Experimental validation of different shim optimisation methods 129
3.2.6 Discussion 132
3.3 Shim step size optimisation 135
3.3.1 Tikh-based shim step size optimisation 136
3.3.2 Simulation of shim step optimisation methods 140
3.3.3 Discussion 144
3.4 Effects of using various cost functions in dynamic shimming 147
3.4.1 Signal loss and BOLD sensitivity model 149
3.4.2 Implementation of different models in shimming 150
3.4.3 Comparison of different optimisation methods 151
3.4.4 Discussion 157

4 Template-based field map prediction for static and dynamic B_0 shimming 159
4.1 Template-based field map prediction for static B_0 shimming 160
4.1.1 Field map template generation and prediction 161
4.1.2 Simulations of the template-based method 165

4.1.3 Experimental evaluation of the template-based field map prediction . . . 172
4.1.4 Third-order shimming . . . 177
4.1.5 Considerations . . . 177
4.2 Adaptive template-based field map prediction for static and dynamic B_0 shimming . . . 181
4.2.1 Correlation between field map similarity and different factors . 183
4.2.2 Adaptive field map prediction . . . 188
4.2.3 Evaluation of adaptive methods in simulation . . . 193
4.2.4 Discussion . . . 196

5 Conclusion **199**
5.1 Future work . . . 201
5.1.1 Eddy current compensation . . . 202
5.1.2 Shim calculation optimisation using regularisation methods . . 202
5.1.3 The template-based field-map prediction . . . 203
5.1.4 Potential applications of the template-based prediction method 204
5.1.5 Potential applications of dynamic shimming . . . 205

Appendices **207**
A A list of spherical harmonics up to fourth orders . . . 207
B Schematic diagrams of EPI sequences for dynamic shimming . . . 207

Bibliography **207**

Prologue

With the publicity and the gained tremendous benefits of the human brain research, non-invasive imaging modalities such as magnetic resonance imaging (MRI), ultrasound, and positron emission tomography (PET) have been used more and more frequently. These techniques have allowed researchers, clinicians and neuroscientists to cultivate collaborations in diagnosis of neurological disorders and neuroscience research. In recent medical and neuroscience research publications, MRI comes into popular use, as it allows for visualisation of the brain structure and delineating the human brain connections and functions in different brain areas [1]. MRI's important role in the brain research comes from its ability to detect musculoskeletal and nervous activity and to reveal metabolic information.

Much research offers the evidence of the gain in signal when moving to a higher field, as a stronger polarisation of protons is induced by a high B_0 field strength [2, 3, 4]. In an idealised homogeneous field, higher field MR systems allow imaging at higher temporal and spatial resolution. The increased spatial resolution in MR images could potentially reveal complex tissue structures and properties in exquisite detail and improve clinical results (e.g., depiction of pathologies and lesions) [5] [6]. These attributes facilitate early diagnosis of diseases and remedial treatment of patients. Clinical research at higher fields can also help clinicians understand the nature of diseases and the effects of therapy.

Although the benefits of going to higher field strength are evident, MR systems are more demanding on advances in hardware and software. Recent decades have seen many breakthroughs in MR technology (e.g. zero boil-off technology and high-performance gradient coils), which facilitate research and development of imaging techniques at higher field strength [7]. Currently, 1.5T MR systems are still the workhorse for clinical and research practice, while 3T MR systems have demonstrated improved clinical applicability and diagnostic benefits due to better images quality. With the 3T MR techniques becoming mature, there is a push to manufacture scanners with even higher field strength, for instance, 7T and 11.7T. Recently developed solutions to those ultrahigh-field challenges allow 7T imaging to benefit multiple clinical applications [8]. Ultrahigh-field scanners have been demonstrated their clinical superiority in morphological diagnosis, for instance, the detection of pathologies in Alzheimer's disease [9] and multiple sclerosis plaques [10], spectroscopy of brain tumour, and functional diagnosis [11].

There are many ultrahigh-field challenges that have not been addressed [6]. Most MR techniques, especially echo-planar imaging and steady-state free precession, benefit from a relatively homogeneous magnetic field. The field is often distorted by the presence of the subject. In the human brain, the most prominent field heterogeneity occurs around the sinuses and ear cavities, because of the high magnetic susceptibility difference between air and tissue [12]. The heterogeneity of the magnetic field in those regions leads to significant image artefacts, such as signal loss and image distortion. Moreover, it impedes the detection of tissue structures and properties that counteracts the benefits of going to higher field strength.

Improving field homogeneity is an on-going problem since the invention of the MRI scanner. The techniques to homogenise magnetic fields, so called shimming, are used [13, 14, 15]. Most clinical MR scanners are equipped with room-temperature

shim coils that generate magnetic fields to counteract B_0 field inhomogeneity. Most 3T systems are equipped with second-order shim coils, whereas some ultrahigh-field systems have even higher orders available. Typically, MRI studies will acquire a field map as one of the first scans. The field map is used to calculate a current for each shim coil that will result in a total field that is maximally homogeneous.

Numerous shimming techniques have been developed over the years. A more recent advance in shimming is dynamic shimming, which is referred to as the techniques of updating shim settings during scans to optimise field inhomogeneity in small regions of interest (ROI). Dynamic shimming techniques have been shown to be superior to conventional methods [16, 17, 18]. In this thesis, we will introduce and demonstrate the performance of dynamic shimming technique at 7T MRI, and propose novel techniques for optimisation of dynamic shimming.

Chapter 1

Introduction

MRI is a technique to produce images of biological tissues by measuring signal from nuclei through manipulating three types of magnetic fields: (1) the main magnetic field, termed B_0, (2) the radiofrequency field, termed B_1, and (3) the gradient fields. To understand the mechanism behind MRI, **Section 1.1** introduces the background knowledge of the quantum mechanics and basics of nuclear magnetic resonance. **Section 1.2** then dives into key MRI principles and explains how nuclei interact with the three types of magnetic fields to form MR images. In order to get high quality MR images, MRI requires a homogeneous B_0 field, which is the focus of this thesis. **Section 1.3** offers some insights on the role of the B_0 field in MRI and the importance of a homogeneous B_0 field to MRI. Specifically, origins of field inhomogeneity in MR scanners and the influences of field inhomogeneity on MR image quality will be introduced. Following that, **Section 1.4** presents basic techniques to B_0 field optimisation, termed "shimming". Finally, **Section 1.5** provides a literature review of various shimming techniques (e.g., active shimming, local shimming, and dynamic shimming). Familiarity with this knowledge lays the foundation for understanding the following chapters about implementation and optimisation of dynamic shimming techniques.

1.1 MR physics

This section aims to introduce the basic concepts of nuclear magnetic resonance and describe the phenomenon nuclei experience when placed in a B_0 field. We will then describe the behaviour of nuclei under the influence of an external B_1 field and how this makes contributions to MR signal.

1.1.1 Magnetic resonance

In 1946, Felix Bloch and Edward Purcell discovered nuclear magnetic resonance [19, 20]. It describes that the phenomenon of magnetic resonance occurs on an atomic nucleus that contains odd numbers of protons or neutrons, such as hydrogen-1, phosphorus-31, and carbon-13, in the presence of a B_0 field [21]. These nuclei are usually referred to as spins, which can be used as sources of MR signal. In practice, hydrogen nucleus is found abundant in human tissues. Therefore, MR studies usually treat protons as sources of MR signal.

In the natural environment, spins in the sample are oriented in random directions and the sum of their magnetic moment is zero. Under the influence of a B_0 field, the spins will be polarised and their magnetic moment will be oriented in parallel or anti-parallel with the applied external B_0 field they experience. The spins also precess in the transverse plane (xy-plane) about the direction of the B_0 field (z-axis) at a specific frequency, as illustrated in Figure 1.1a. The frequency the spins precess at is proportional to the B_0 field strength, called Larmor frequency. The behaviour of the spins can be characterised as

$$f = \frac{\omega}{2\pi} = \frac{\gamma B_0}{2\pi}, \tag{1.1}$$

where γ denotes the gyromagnetic ratio of H1 protons ($\gamma = 267.5 \times 10^6$ rad/(s $\cdot$ T)). For the H1 nucleus, it precesses at a frequency of 42.57MHz/Tesla. The work in this thesis is based on a 7T scanner and the corresponding resonance frequency is 298MHz. The direction along the applied B_0 field is referred to as the longitudinal direction (z-axis). The directions perpendicular to the longitudinal direction are referred to as the transverse directions (x- and y-axis).

From a quantum mechanics view, the phenomenon of magnetic resonance corresponds to energy exchange when a time-varying radio-frequency (RF) pulse oscillating at the Larmor frequency is applied to the spins. Under the influence of a RF pulse, the spins absorb energy from the RF pulse and are excited to a higher energy state (i.e. opposing the direction of the external magnetic field). This phenomenon is called "excitation".

From an alternative classical view, excitation involves using a B_1 field oscillating at the Larmor frequency in the transverse direction to tip the magnetisation of the spins. During excitation, there is a torque applied to the spins and the direction of magnetisation is tilted away from that of the B_0 field, as illustrated in Figure 1.1b. The angles tilted away from the B_0 field, called flip angle, depend on the B_1 field strength (micro-tesla) and its duration. The magnitude and the phase of the magnetisation in the transverse plane reflect the properties of the sample.

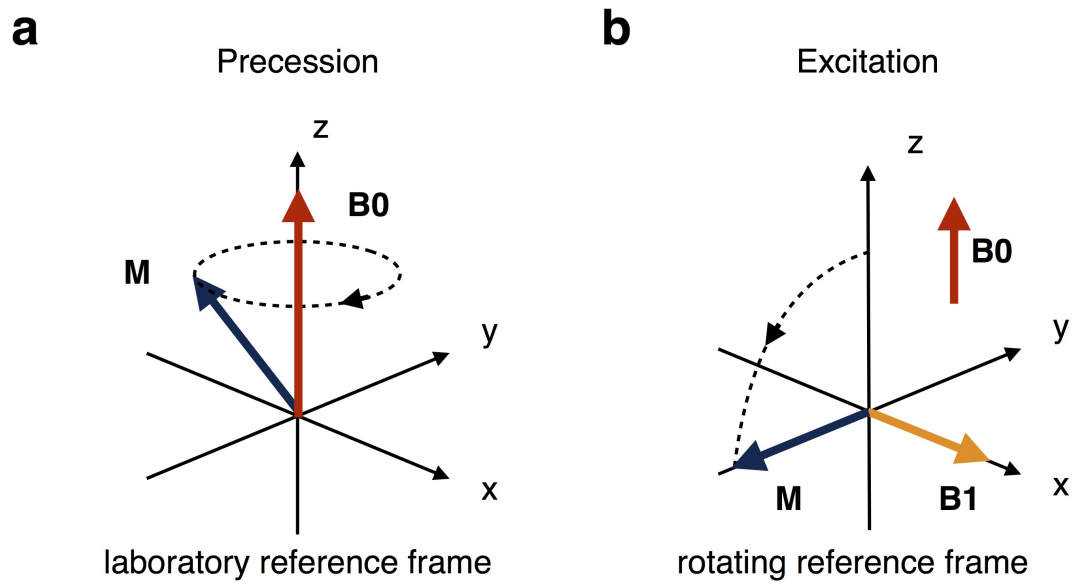


Figure 1.1: Illustration of the precession of magnetic moment of spins in the laboratory reference frame (a) and the rotation of magnetisation of spins in the rotating reference frame after excitation (b).

1.1.2 Relaxation

T1 and T2 relaxation

After excitation, the net longitudinal magnetisation M_z and the transverse magnetisation M_{xy} gradually relax back to the thermodynamic equilibrium state over time. This process is known as relaxation, as illustrated in Figure 1.2a.

The longitudinal magnetisation recovery is mainly due to field fluctuations in the transverse plane. The field fluctuations drive the magnetisation of the spins aligning with the longitudinal axis. In this process, energy exchange occurs between the spins and their environment ("the lattice"). The speed of the longitudinal magnetisation recovery can be characterised by a time constant T1. Thus, this process is usually termed the spin-lattice relaxation or T1 relaxation. The evolution of the longitudinal

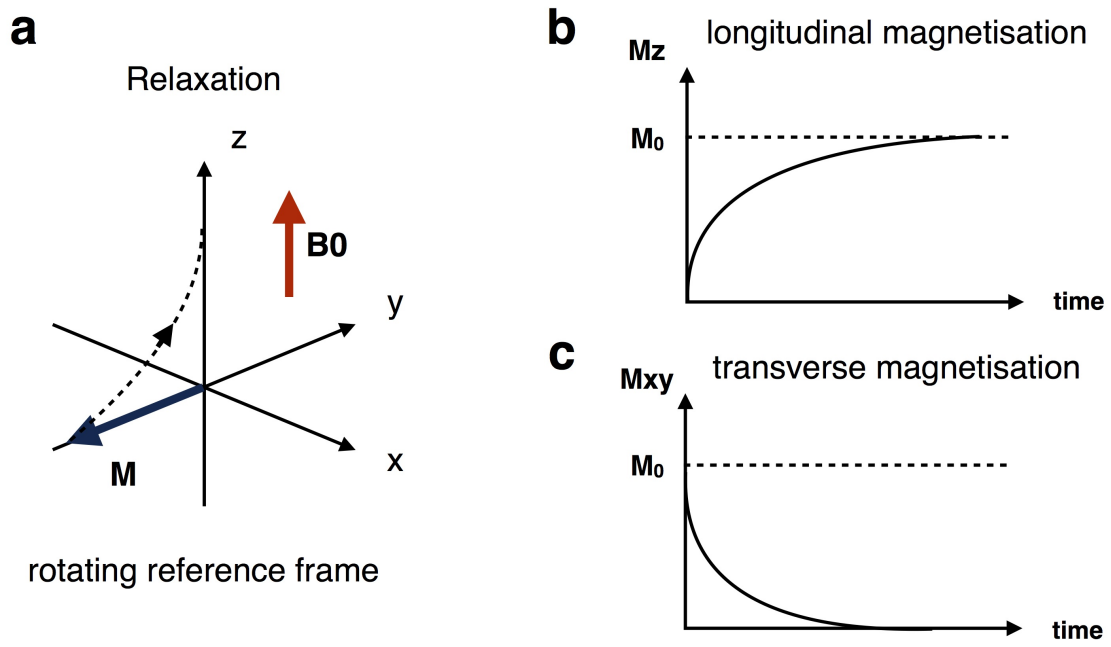


Figure 1.2: Illustration of magnetisation recovery of spins in the rotating reference frame after excitation (a), and the change of the longitudinal (b) and the transverse magnetisation decay (c).

magnetisation M_z over time t after excitation is illustrated in Figure 1.2b. Mathematically, after a 90 degree excitation, the evolution of longitudinal magnetisation can be described as a function of time t:

$$M_z(t) = M_0\left(1 - e^{-t/T1}\right), \tag{1.2}$$

where M_0 indicates the net magnetisation in an equilibrium state in the presence of a B_0 field.

The transverse magnetisation M_{xy} decay is largely attributed to field fluctuations along the longitudinal direction. Thus, spins in the sample experience a wider range of frequencies. As a result, the spins gradually lose phase coherence. This process is described by the term "dephasing". In this process, energy exchange happens between the spins. The speed of the magnetisation decay in the transverse plane can be characterised by a time constant T2. Therefore, this process is usually referred to as the spin-spin relaxation or T2 relaxation. The evolution of the transverse

magnetisation M_{xy} over time t after a 90 degree excitation is illustrated in Figure 1.2c, which can be approximated as:

$$M_{xy}(t) = M_0\left(e^{-t/T2}\right). \tag{1.3}$$

The relaxation time constants T1 and T2 depend on tissues types and vary at different field strengths. Typical values of T1 and T2 time constants for white and grey matter at 3T and 7T are shown in Table 1.1.

	T1 at 3T	T1 at 7T	T2 at 3T	T2 at 7T
White matter	830	1200	80	40
Grey matter	1330	2000	110	60

Table 1.1: Typical values of T1 and T2 relaxation time constants at 3T and 7T. [22, 2, 23]

T2* relaxation

In a homogeneous B_0 field, the transverse magnetisation decays exponentially with a time constant T2. In an inhomogeneous field, the spins experience a wider range of frequencies and the "dephasing" effects are accentuated by local field inhomogeneities. Thus, the transverse magnetisation attenuates at a shorter time constant than the T2 relaxation. The speed of the decay is usually described by a time constant T2*. This process is called T2* relaxation. The T2* relaxation reveals both the microstructure and the tissue function information [24].

Bloch equation

As introduced above, the spins precess at the Larmor frequency and exhibit the phenomenon of magnetic resonance in a B_0 field. A radiofrequency field oscillating

at the Larmor frequency can excite the spins and thus contributes to MR signal. Afterwards, the magnetisation in the longitudinal direction and the transverse plane relaxes with time constants T1 and T2. In the presence of field inhomogeneity, the transverse magnetisation decays at a time constant of T2*.

The behaviour of the macroscopic nuclear magnetisation in the transverse plane and the longitudinal direction can be described by an equation called Bloch equation [19]. It summarises the change of the nuclear magnetisation in the transverse plane and the longitudinal magnetisation, respectively, as:

$$\frac{dM_{xy}(t)}{dt} = \gamma \left(M(t) \times B(t)\right)_{xy} - \frac{M_{xy}(t)}{T_2} \tag{1.4}$$

$$\frac{dM_z(t)}{dt} = \gamma \left(M(t) \times B(t)\right)_z - \frac{M_z(t) - M_0}{T_1}, \tag{1.5}$$

where M_0 indicates the magnetisation of nuclei in the equilibrium state in a magnetic field. The first terms $\gamma \left(M(t) \times B(t)\right)_{xy}$ and $\gamma \left(M(t) \times B(t)\right)_z$ describe the Larmor precession of the nuclear magnetisation in the transverse plane and the longitudinal direction over time t. The second terms $-\frac{M_{xy}(t)}{T_2}$ and $-\frac{M_z(t)-M_0}{T_1}$ describe the relaxation behaviour of the transverse and the longitudinal magnetisation of spins, respectively.

Free induction decay

In MRI, the transverse magnetisation can be measured by Faraday's law of induction in a set of electric coils. The measured signal oscillates at the Larmor frequency and decay exponentially with a time constant T2*. The generated signal is called free induction decay (FID). FID signal $S(t)$ in a volume over a duration of t can be described as

$$S(t) = \int_{volume} M(\mathbf{r}) e^{-t/T2^*} d\mathbf{r}. \tag{1.6}$$

where **r** indicates location information of image volumes. The FID signal contains information from all kinds of sources ranging from the local field inhomogeneity to the chemical shift.

1.2 MR image formation

This section describes the principles for the generation of an MR image from the magnetisation of spins. It mainly covers three aspects: signal encoding using gradient fields (**Section 1.2.1**), MR pulse sequence design (**Section 1.2.2**), and hardware configuration of a typical MR scanner (**Section 1.2.3**).

1.2.1 Spatial encoding using gradient fields

In a homogeneous field, spins in the sample precess at the same frequency. The FID signal acquired after excitation is contributed from the entire sample. The formation of images requires differentiation of the contribution of signal at different locations. As described in the Larmor Equation 1.1, the resonance frequency of the spins depends on local field strength. Thus, modification of the magnetic field that the spins experience allows separating signal contributed from different locations. Lauterbur PC. and Mansfield P. proposed to add a set of additional linearly varying fields orientating in parallel with the B_0 field to spatially encode MR signal [25, 26]. These additional magnetic fields are referred to as gradient fields.

In MRI scanners, signal encoding is accomplished by using three gradient fields orthogonal to each other. To differentiate spins located at different slice positions, a field orientated in parallel with the B_0 field and varying linearly along the axis perpendicular to this plane is applied, termed a slice-selection gradient, as illustrated in

Figure 1.3a. The magnetic field the spins experience varies along this gradient. The field strength at location z is given by

$$B(z) = B_0 + G_z \cdot z, \tag{1.7}$$

where B_0 denotes main magnetic field strength, G_z denotes field strength of the slice-selection gradient. Thus, spins in the plane located at position z recess at a frequency of $\omega(z) = \gamma B(z)$.

In order to only excite spins in the targeted slice with a thickness of Δ z located at position z, a RF pulse tuned to a frequency centred around $\omega(z)$ with a bandwidth of $\omega(\Delta z) = \gamma B(\Delta z)$ is transmitted.

Similar to the mechanism of the slice-selection, in order to spatially encode spins at different locations in a plane, two orthogonal in-plane field gradients (x and y) that are perpendicular to the slice-selection gradient are applied, as illustrated in Figure 1.3b.

Under the influence of the in-plane x-gradient, the accrued phase for the spins at spatial location x at time point t is given by

$$\phi(x,t) = \omega_0 t + \gamma \int_0^t G_x(\tau)\, x d\tau. \tag{1.8}$$

Alternatively, the phase can be represented by

$$\phi(x,t) = \omega_0 t + k_x(t)\, x \tag{1.9}$$

with

$$k_x(t) = \gamma \int_0^t G_x(\tau)\, d\tau, \tag{1.10}$$

where $k_x(t)$ denotes the spatial frequency at time t after excitation during a time course of the x-gradient. After switching off the gradient, spins in the sample start

to precess at the original frequency while preserving previously accrued phase information. The measured MR signal at location x can be expressed as:

$$S(t) = \int M(x)\, e^{-t/T2} e^{-i\phi(x,t)} dx \tag{1.11}$$

$$= \int M(x)\, e^{-t/T2} e^{-i(\omega_0 t\ + k_x(t))} dx, \tag{1.12}$$

where $M(x)$ denotes the magnetisation of the excited spins at location x and $\phi(x,t)$ denotes the phase accumulated at time t after excitation. The measured signal contains a range of spatial frequencies.

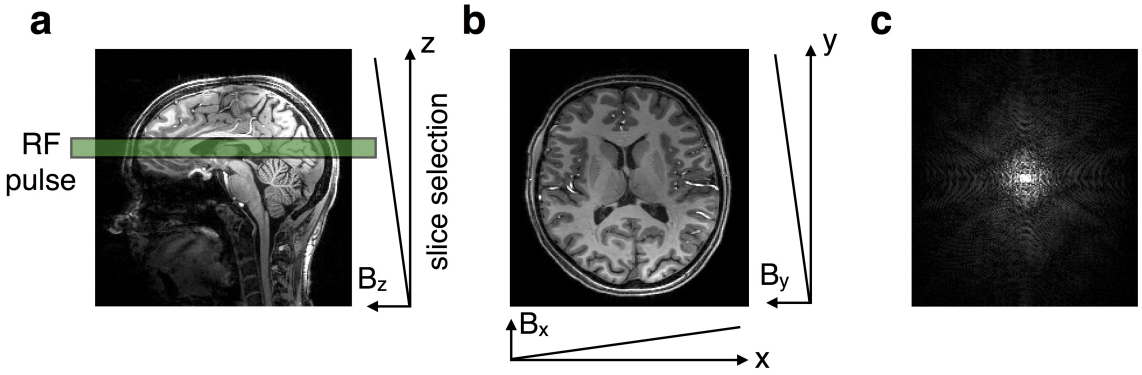


Figure 1.3: Illustration of spatial encoding using the slice selection gradient (a) and the in-plane gradients (b). The image in (c) displays the k-space representation of the brain structural image (b).

The signal measured through the application of three gradients can be expressed as:

$$S(t) = \int\int\int M(x,y,z)\ \ e^{-t/T2} e^{-i(\omega_0 t\ + k_x(t)x + k_y(t)y + k_z(t)z)} dxdydz. \tag{1.13}$$

where $M(x,y,z)$ denotes the magnetisation of the spins at location (x, y, z). The signal equation has a Fourier-like format. The measured signal represents the Fourier transform of the magnetisation of the spins $M(r)$ as a function of spatial frequency. This array is referred to as the k-space, as shown in Figure 1.3c. Each data point in k-space contributes to the information of the entire MR image at a specific spatial

frequency. The k-space data at different locations affect MR images differently. The low frequency components of the k-space data, which are located near the centre of k-space, contain image information about general shapes and contours of the sample. The high-frequency components, located near the periphery of k-space, contain image information about edge details. Sampling data in k-space is a convenient way to acquire MR signal.

1.2.2 MR pulse sequence: GRE sequence

In order to control different components of the MR hardware for image acquisition, pulse sequences, a programmable set of the RF pulses and the imaging gradients, are utilised. To get MR signal from different sources, various sequences can be used (e.g., gradient echo, spin echo, and balanced steady-state free precession). In this subsection, we will introduce the basics of the gradient echo sequence which is of great relevance to this thesis.

Principles of gradient echo

In an inhomogeneous field, after excitation, spins in the sample are dephasing with a time constant T2*. The application of imaging gradients causes spins to relax at a faster rate which crushes FID signal. In order to regain phase coherence, the gradient echo sequence can be used. The gradient echo formation using the readout gradient is illustrated in Figure 1.4.

A typical GRE sequence begins with a RF pulse to tip the magnetisation of the spins into the transverse plane. A dephasing gradient produces different degree of phase

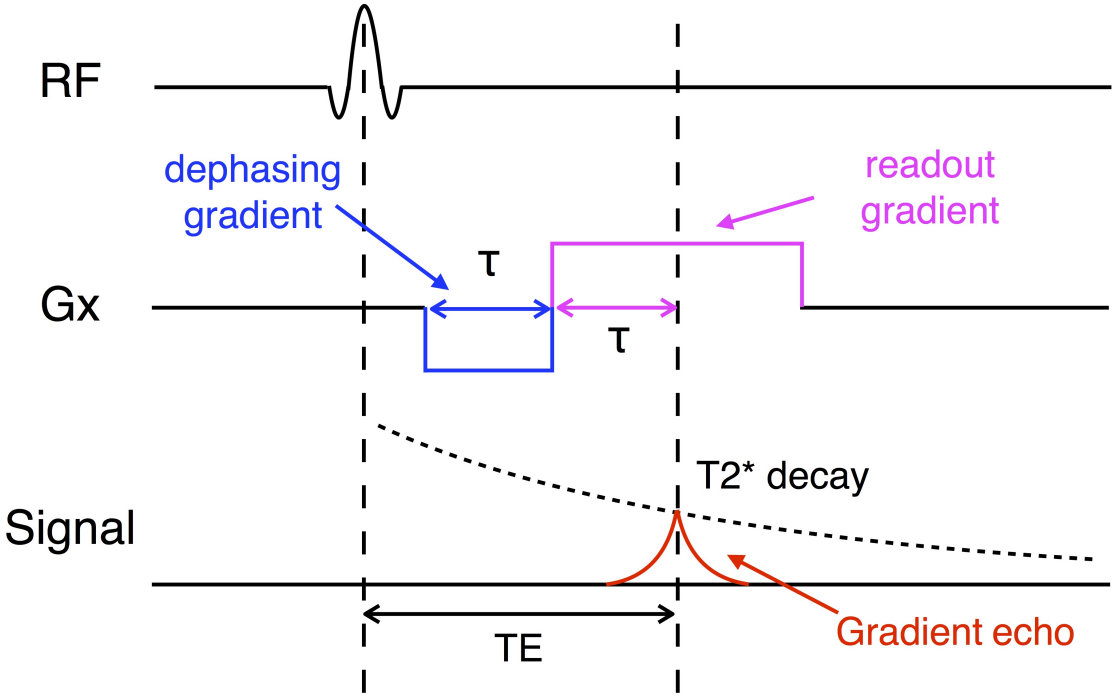


Figure 1.4: Illustration of forming a gradient echo using the readout gradient. The blue pulse represent a dephasing gradient to cause dephasing; the pink pulse represents a readout gradient to unwind dispersed phase; the red peak represents a gradient echo where signal corresponds to T2* relaxation. τ indicates the typical duration of the dephasing and the readout gradients.

dispersion along the applied gradient. To unwind accrued phase induced by the dephasing gradient, a readout gradient with opposite polarity is applied. A gradient echo occurs at the time when the amount of phase shift induced by the dephasing gradient is cancelled out by the rephasing gradient, namely $\phi(x,t) = 0$. At this time point, the spins regain phase coherence and the signal reaches the level determined by T2*. The time elapsed between the RF pulse and the gradient echo is termed echo time (TE).

GRE sequences are sensitive to field variations, as the phase dispersion from field inhomogeneity is not cancelled out at the gradient echo. Therefore, GRE sequences are suitable for measuring field inhomogeneity in the sample. Moreover, it has superiority in detecting field changes induced by physiologic processes, like brain activations. Thus, GRE-based sequences gain popularity in functional MRI. However, the field inhomogeneity can also cause severe image artefacts in GRE images, like signal loss and geometric distortion. Effects of B_0 field inhomogeneity on GRE images will be discussed in detail in **Section 1.3**.

Encoding trajectory

MR images can be reconstructed by applying the inverse Fourier transform to the k-space data. The k-space points at different locations contribute differently to final images. Thus, strategies of sampling k-space data points affect the resulting MR image quality. In MRI, there are various approaches to sample the k-space (e.g., Cartesian, radial, and spiral methods). Among these methods, the Cartesian based method is one of the most widely used acquisition schemes.

In the **conventional Cartesian method**, two in-plane gradients x and y function interactively to sample k-space data. The field gradient applied to acquire a k-space

line after each RF excitation for signal readout is referred to as the frequency encoding gradient.

In order to acquire all the k-space data, k-space lines at different positions are required to be acquired. The gradient applied to encode the position of a k-space line is referred to as the phase encoding gradient. A typical GRE sequence with the conventional Cartesian k-space sampling trajectories and the corresponding pulse sequence timing diagram are illustrated in Figure 1.5. The time interval between each k-space line acquisition is termed repetition time (TR). The total time to fill the whole k-space, termed acquisition time (TA), is given by TA= TR* N, where N is the number of RF excitations. Therefore, the conventional Cartesian acquisition is long in scan time, whereas no distortion occurs in the phase encoding direction.

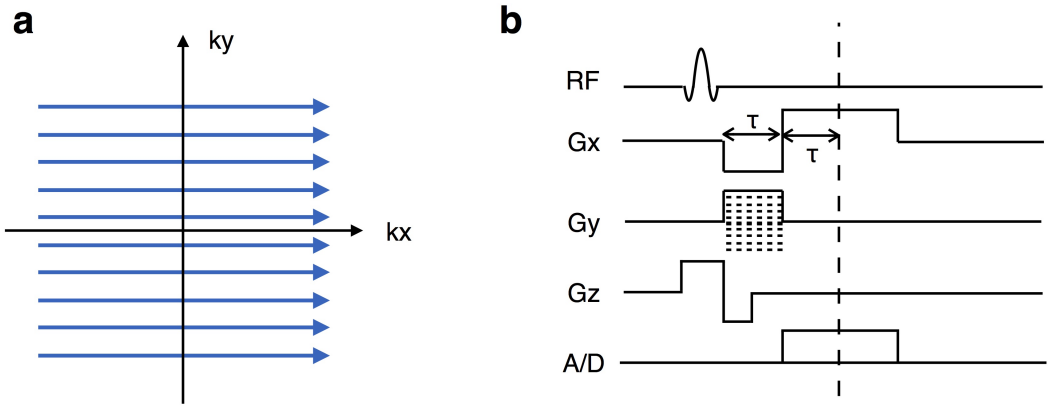


Figure 1.5: Illustration of the k-space trajectories (a) and the pulse sequence diagram (b) of a gradient echo sequence. The horizontal axis represents the timing of sequence activities. RF represents radiofrequency pulse; Gx, Gy, and Gz represent the strength of frequency encoding gradient, phase encoding gradient, and slice selection gradient, respectively. A/D line represents the MR signal receiver.

Echo-planar imaging (EPI) is another type of Cartesian-based acquisition method. It is typically designed to acquire the entire k-space data after a single RF excitation (i.e. single-shot EPI). Figure 1.6 shows a typical single-shot EPI k-space sampling trajectory and the corresponding pulse sequence timing diagram. In the example

single-shot EPI, sampling starts from the bottom of k-space and then sweeps forward and backward by rapidly reversing the frequency-encoding gradient. Applying gradients in this way forms subsequent gradient echoes at the centre of each k-space line. Additionally, each k-space line acquisition in the frequency-encoding direction is followed by upwards blips, small gradients in the phase-encoding direction. The blips step the k-space line acquisition through different positions along the phase-encoding direction. This acquisition method is referred to the blipped Cartesian EPI.

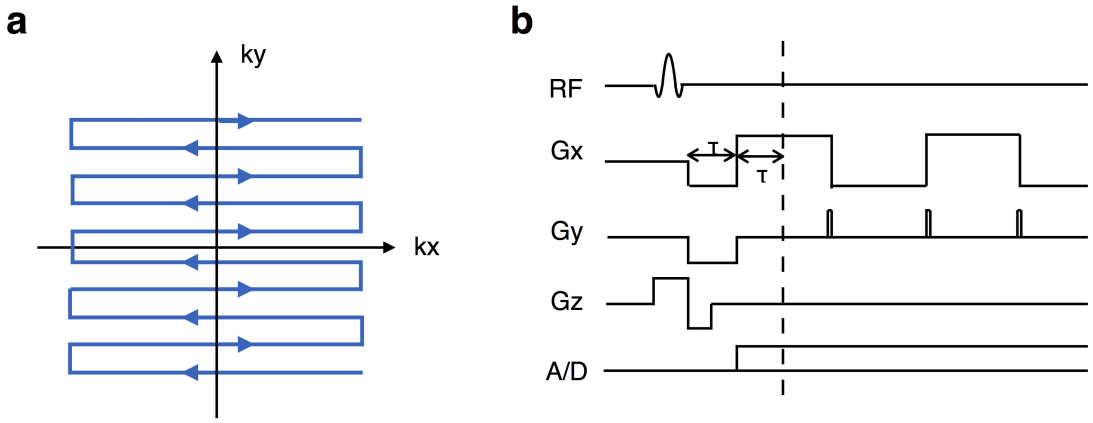


Figure 1.6: Illustration of (a) k-space trajectory and (b) pulse sequence diagram of a EPI sequence. The horizontal axis represents the timing of sequence activities. RF represents radiofrequency pulse; Gx Gy and Gz represent the strength of frequency encoding gradient, phase encoding gradient, and slice selection gradient, respectively. A/D line represents the MR signal receiver.

The single-shot acquisition method has both advantages and disadvantages. On the one hand, the single-shot technique maximises the coverage of k-space acquisition per excitation, and dramatically shortens the image acquisition time. This technique allows MR imaging at a high temporal resolution (TA $\approx$ tens of milliseconds), thereby mitigating motion artefacts [27]. On the other hand, the single-shot EPI has a long phase-encoding duration. The sampling bandwidth of EPI in the phase encoding direction depends on the time interval between subsequent gradient echoes and is relatively low (on the order of milliseconds). Thus, small field inhomogeneity will induce relatively more phase accumulated in the phase-encoding direction, leading to

severe image artefacts (e.g., signal loss and geometric distortion) [28]. To mitigate the artefacts of the single-shot technique, multi-shot EPI has been developed. The multi-shot technique is achieved by equally segmenting the k-space to some subregions to reduce the duration for phase encoding. However, this method requires a longer acquisition duration.

1.2.3 MR hardware

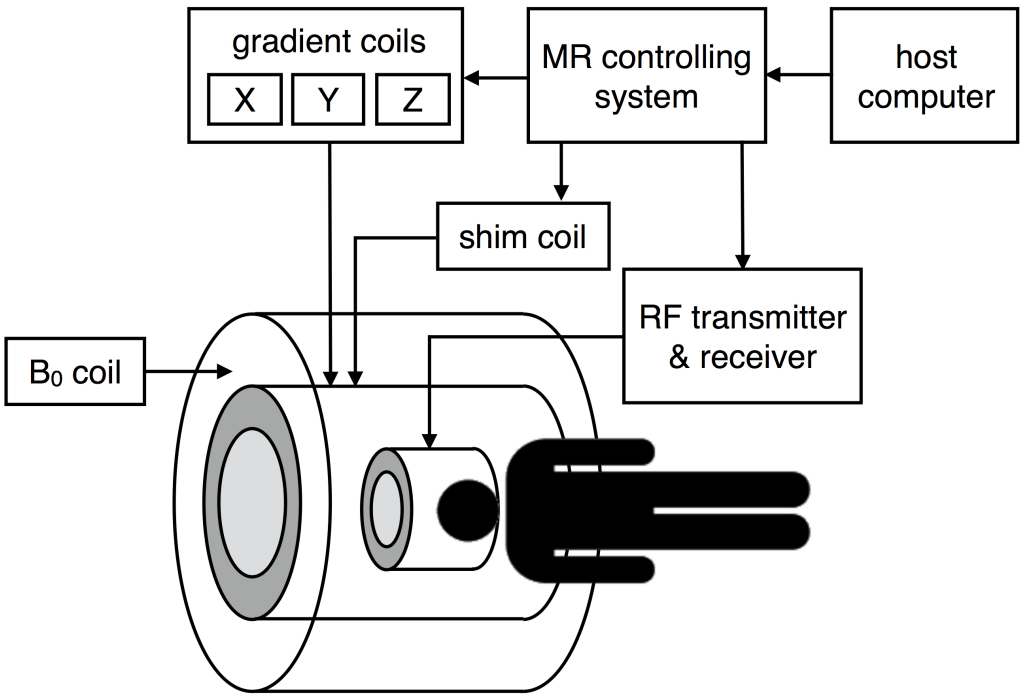


Figure 1.7: Illustration of the hardware configuration of a typical high-field MRI scanner. It mainly consists of a B_0 superconducting magnet, RF transmitter and receiver coils, gradient coils, and shim coils.

A typical MR scanner is mainly comprised of three components: a superconducting magnet, RF transmitter and receiver coils and gradient coils (Figure 1.7)

The cryogenic superconducting B_0 coil is dedicatedly manufactured to generate a strong static B_0 field to polarise the protons. To retain the superconducting effects

of B_0 coils, the magnet is surrounded by the heat shields and the cryostat.

RF transmitter and receiver coils are manufactured and placed near the sample to generate an electromagnetic field perpendicular to the B_0 field for excitation. According to Faraday's law of induction, the change of magnetisation in the transverse plane after excitation will induce currents in conductors nearby, which contributes to MR signal. RF coils are then used to detect the electromagnetic currents, which are amplified and demodulated to form MR signal. There are various types of RF coils (e.g., birdcage coils and cylindrical phased array coils) that are optimised for different imaging applications.

Gradient coils are mounted concentrically in the inner bore of the magnet to generate spatially varying fields to encode spins at different locations. During MR scans, switching of the gradients generates temporally varying magnetic field and thereby induces eddy currents in nearby conducting structures (e.g. the cryostat of the magnet). The eddy currents are undesired and can distort the required output gradient fields. To counteract the eddy current effects, active shield coils are designed and placed close to imaging gradient coils [29].

Another key component in a MRI system is shimming hardware. The knowledge of shimming, a method for compensating B_0 field inhomogeneity, is important to the research described in this thesis, which will be introduced in detail in **Section 1.5**.

B_0 field inhomogeneity affects the quality of MR images. Typical sources of field inhomogeneity can be classified to factors intrinsic and extrinsic to the sample. The extrinsic factors mainly include inaccuracies during the manufacture of the magnet and coils (e.g. the intrinsic physical properties of materials), and some local site issues (e.g., the steel used in site and the magnetic shielding). To homogenise the magnetic field in the bore of an "empty" magnet, passive shimming is used. Passive shimming

is implemented by placing dedicatedly-designed ferromagnetic materials, like steel or iron (shims), at specific positions within the scanner bore to generate correction fields to improve the field homogeneity [30, 31]. With passive shimming, the magnetic field in the bore can reduce the field inhomogeneity to the order of parts per million. The intrinsic factors that affect field inhomogeneity include the susceptibility variations within the subject. This is discussed in detail in **Section 1.3.1**. To correct the field inhomogeneity regardless of its origins, a set of electric-current-driven room temperature shim coils are available in most high-field MR systems to generate an adjustable magnetic field (see **Section 1.5.1**). The shim coils are usually incorporated in the gradient coil structure and wound on a cylindrical form [32]. Conventionally, a shimming field that varies linearly across the sample is achieved by adding offsets to imaging gradient fields. High-order shim fields are set through a digital serial-control-unit prior to image acquisition.

1.3 Linking B_0 homogeneity to image quality

Section 1.2 introduced the principles of MR image formation. This section introduces susceptibility-induced B_0 inhomogeneity (**Section 1.3.1**) and its relevance to image quality (**Section 1.3.2**).

1.3.1 Susceptibility induced B_0 field inhomogeneity

The sample placed in a homogeneous field can induce shift of the B_0 field. The degree of the B_0 field shift depends on the susceptibility of the sample. In vacuum, the magnetic field generated by a MR superconducting magnet is given by

$$B_0 = \mu_0 \cdot H \tag{1.14}$$

where μ_0 indicates the vacuum permeability, H represents is a magnetic field. The susceptibility-induced B_0 field scales with background field strength, which can be expressed by

$$B_{sus} = B_0 + \mu_0 \chi H, \tag{1.15}$$

where χ denotes the susceptibility of tissues. Typical values of the susceptibility of different tissues are listed in Table 3.

Material	χ $(\times 10^{-6})$
Air	0.40
Water	- 9.04
Fat	-7.79
Bone	-8.44
Grey matter	-8.97
White matter	-8.80
Deoxygenated red blood cells	-6.07
Oxygenated red blood cells	-9.41

Table 1.2: Typical susceptibility values of materials [33, 34, 35]

The susceptibility effects have both advantages and disadvantages. On the one hand, the blood-oxygenation-level dependent (BOLD) technique relies on detecting the locations and the degree of field changes resulting from the susceptibility difference in the oxygenated and deoxygenated hemoglobin due to neuronal activity [36]. Specifically, during neuronal activity in the brain, the local magnetic susceptibility of tissues varies with the local cerebral blood flow and oxygenation of the hemoglobin. As shown in Table 1.2, the susceptibility difference between oxygenated and deoxygenated red blood cells amounts up to 2.5 $\times 10^{-6}$. Higher field yields stronger susceptibility-induced field contrast due to brain activities, thereby leading to accentuated T2* dephasing effects [37]. Therefore, BOLD techniques favour high-field systems due to enhanced functional sensitivity [11, 38].

On the other hand, the susceptibility effects cause B_0 field inhomogeneity. As shown in Table 1.2, the susceptibility difference between air and water is approximately 50

times larger than that between grey and white matter. In the human brain, the most prominent macro-scale B_0 field inhomogeneity occurs in the regions where large magnetic susceptibility difference between air and surrounding tissues presents, e.g. the prefrontal cortex and the temporal lobes, as shown in Figure 1.8. According to Equation 1.15, an increase of field strength will increase the inhomogeneity of the magnetic field, leading to exacerbated image artefacts.

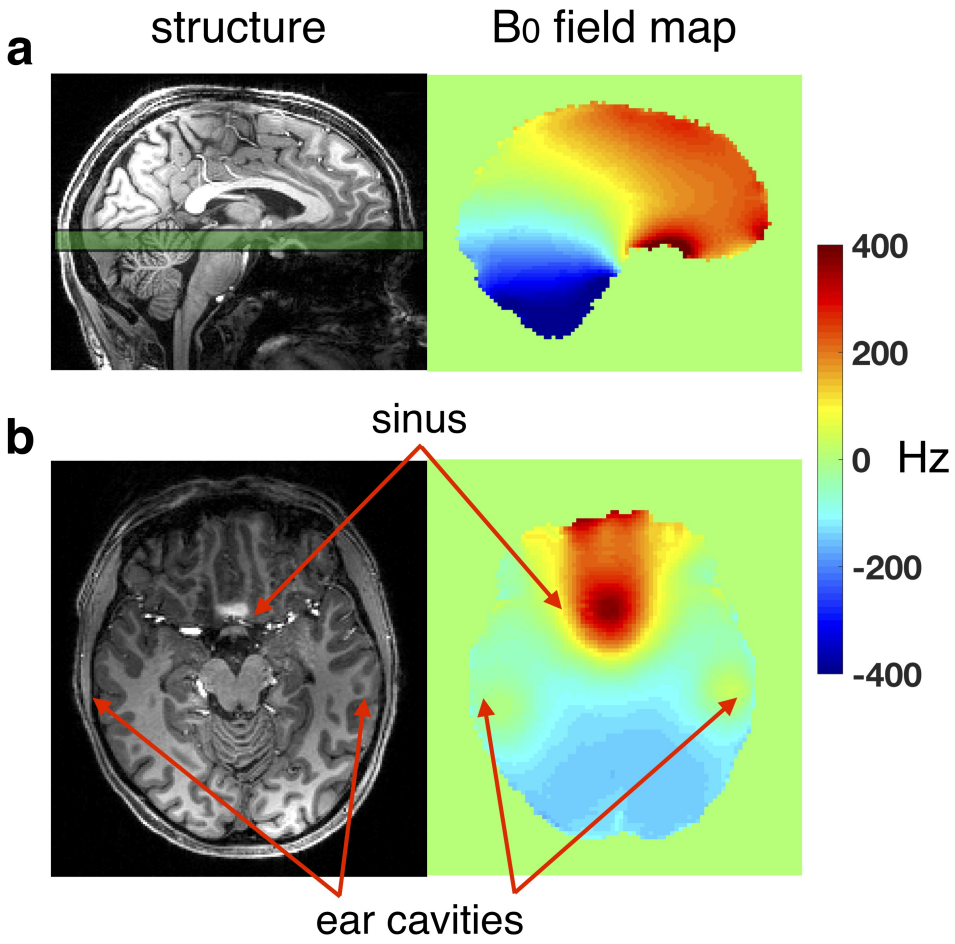


Figure 1.8: The brain structure (left) and the corresponding B_0 field map (right) of an example subject acquired at 7T ($\sim$ 298MHz). The sagittal view and transversal view are shown in (a) and (b) with labels on the regions of large field distortion (e.g., close to sinuses and ear cavities).

1.3.2 B_0 field inhomogeneity related image artefacts

There are two typical image artefacts that are caused by the susceptibility-induced field inhomogeneity, geometric distortion and signal loss [39]. This section introduces models to quantify the two image artefacts based on a B_0 field map. Additionally, some approaches to correct them will be briefly discussed.

Geometric distortion

For 2D imaging, geometric distortion (e.g., deformation and pixel shift) mainly results from the in-plane field inhomogeneity. The degree of distortion in MR images is affected by many factors, such as the type of sequences and the sampling bandwidth. Among different sequences, EPI suffers severely from geometric distortion. Long phase encoding duration in EPI causes predominant geometric distortion in this direction. A popular method to model EPI distortion was proposed by Jezzard P. et al. [40], which estimates image distortion based on an additionally acquired field map. The field-inhomogeneity induced shift of voxels at location (x, y, z) can be described by

$$\Delta y\left(x,y,z\right)=\gamma\Delta B\left(x,y,z\right)/BW_{pe}, \tag{1.16}$$

where $\Delta B\left(x,y,z\right)$ denotes the field inhomogeneity, BW_{pe} denotes the bandwidth of phase encodings.

There are a range of methods that can be used to mitigate geometric distortion artefacts induced by B_0 inhomogeneity. These methods can be classified into two categories: hardware-based and software-based methods.

B_0 shimming is one of the most straightforward hardware-based methods [41]. The advantages and limitations of various B_0 shimming methods will be discussed in detail

in **Section 1.3**. Another important hardware-based method is the parallel imaging techniques [42]. With the implementation of multichannel coils, parallel imaging techniques accelerate phase encoding in EPI which reduce geometric distortion and scan time concurrently, while preserving image resolution and effective echo time. However, the advantages come at a cost of signal-to-noise ratio (SNR) degradation due to k-space under-sampling. Additionally, some image reconstruction related errors also increase with the acceleration factor.

Numerous studies have been performed to investigate software-based techniques to reduce distortion artefacts. Most methods for correcting geometric distortion involve the use of additionally-acquired field maps [43, 44]. Jezzard P. et al. demonstrated the feasibility of correcting image distortion using the geometric distortion model described in Equation1.16. This method has proved effective for fMRI applications, such as locating brain activation and conducting statistic analysis [45, 46]. However, the conventional field map acquisition prolongs the scan time which is an overhead to routine scans. To speed up this process, geometric distortion correction based on EPI-based field map acquisition has been proposed [47]. However, the EPI-based field map itself suffers from geometric distortion and signal loss to some degree, which can affect distortion correction quality. Other methods for distortion correction include algebraic reconstruction approaches and inverse methods for image reconstruction [48, 49].

Signal loss

In the presence of field gradients or field inhomogeneity, spins precess at a wider range of resonance frequencies. The broadening of resonance frequencies accentuates the dephasing effects, resulting in signal loss. The local field gradients along the directions

of the three imaging gradients can be described by

$$G(x,y,z) = \nabla\left(B(x,y,z)\right) \tag{1.17}$$

$$G_{fe} = \frac{\delta\left(B(x,y,z)\right)}{\delta x} \tag{1.18}$$

$$G_{pe} = \frac{\delta\left(B(x,y,z)\right)}{\delta y} \tag{1.19}$$

$$G_{ss} = \frac{\delta\left(B(x,y,z)\right)}{\delta z}, \tag{1.20}$$

where G_{fe}, G_{pe}, and G_{ss} denote field gradients in the frequency encoding, the phase encoding and the slice selective directions, respectively.

Conventionally, signal loss effects can be readily modelled by the sum of the squares of field gradients across a unit volume.

$$I_{loss} \propto \sqrt{G_{fe}^2 + G_{pe}^2 + G_{ss}^2}. \tag{1.21}$$

Besides this general model, an alternative model is available. In EPI, signal loss induced by through-slice gradients is more significant compared to that induced by in-plane gradients [50, 51, 52, 53, 54]. Thus, phase dispersion in voxels at time TE can be approximately modelled with the intra-voxel through-slice field gradients. With the assumption of a Gaussian slice profile, the corresponding phase dispersion is given by

$$\Phi = \gamma\left(G_{ss} \cdot TE\right)\frac{\Delta z}{4\sqrt{ln2}}, \tag{1.22}$$

where Δz denotes the slice thickness. The estimated image intensity under the influence of field gradients is then given by

$$I = I_0 exp\left(-\Phi^2\right), \tag{1.23}$$

where I_0 denotes apparent image intensity in a homogeneous field.

Deichmann R. et al. extends the EPI-based signal intensity model by incorporating the information on the shift of effective echo time induced by error field gradients [52]. Under the assumption of zero-gradient-field, the image intensity and echo time are I_0 and TE_0, respectively. In the presence of error field gradients along the frequency encoding, G_{fe}, and the phase encoding, G_{pe}, directions, the shape of the gradient echo will be distorted and shifted away from the desired echo time. The shift of the related echo time is characterised by a factor of Q,

$$TE_{eff} = \frac{TE_0}{Q}, \tag{1.24}$$

with

$$Q = 1 - \frac{\gamma \cdot \Delta t}{2\pi} \cdot FOV_{pe} \cdot G_{pe}, \tag{1.25}$$

where FOV_{pe} denotes the field of view (FOV) in the phase encoding direction of the EPI sequence, Δt denotes the inter-echo spacing during the EPI readouts. If the echo is shifted outside of the acquisition window, signal will be lost completely, which can be described as

$$I = 0, \qquad if \quad |\gamma \cdot G_{fe} \cdot TE_{eff}| > \frac{\pi}{\Delta x}, \tag{1.26}$$

where Δx indicates the size of voxel in the frequency encoding direction. Meanwhile, TE change also affects the image intensity, which is given by

$$\begin{cases} I = \frac{I_0}{Q} \cdot exp\left[-\frac{TE_{eff}-TE_0}{T_2^*}\right], & if \quad |TE_{eff} - TE_0| < \frac{TA}{2} \\ I = 0, & otherwise. \end{cases} \tag{1.27}$$

where TA indicates image acquisition interval. The use of short TE, though providing high SNR, produces reduced T2* effects. In contrast, long TE causes rapid decay of transverse magnetisation and low signal intensity but yielding high tissue contrast. This forces a trade-off between signal strength and tissue contrast. This limitation affects image quality for BOLD imaging, which will be discussed in **Chapter 3**.

Signal loss in MR images causes loss of information on brain structures and tissue contents. As signal loss is a function of the field inhomogeneity, the most straightforward way to eliminate this artefact is B_0 shimming. Besides these methods, there are a wide range of methods that have been explored to improve the signal intensity. Volz S. et al. suggested that some combinations of images acquired with different echo times could improve the SNR at the cost of increased scan time [55]. Merboldt KD. et al. demonstrated that acquiring thinner slices could reduce field gradient-induced signal loss, though leading to a decrease of the SNR and the spatial coverage [56, 57]. Instead of correcting the image artefacts by optimising imaging protocols, Liu G. et al. proposed a method to improve the SNR by reconstructing the images from k-space data in conjunction with a measured B_0 field map [58]. Other techniques attempted to use multiple refocusing gradients for field inhomogeneity correction and tailoring RF pulses [59, 60].

1.4 Shimming and B_0 field mapping

According to Section 1.3, field inhomogeneity can cause signal loss and geometric distortion. Thus, a homogeneous field is necessary for high quality magnetic resonance imaging. Techniques to compensate for magnetic field inhomogeneity are known as "shimming". This terminology comes from the mechanical shims which is referred to as a piece of material used to adjust position of objects, e.g., lifting the leg of a table to the level surface, or filling gaps between different parts of objects. In MRI, shimming is accomplished by applying an external compensation field that resembles the field inhomogeneity. Conventional shimming methods follow the basic framework [13, 14, 61, 15]:

1. Magnetic field mapping and post-processing

2. Identification of ROIs for shimming

3. Calculation of shims that resemble field inhomogeneity for field inhomogeneity compensation

4. Application of shims to homogenise magnetic fields

Fast and accurate B_0 field mapping is a prerequisite for shimming. In this section, we will review methods to measure field inhomogeneity in the sample and advances in post-processing strategies.

1.4.1 Theory of field mapping

Field inhomogeneity induces additional phase accrued over a prescribed duration. A conventional method of mapping field inhomogeneity is taking the difference between phase scans at different echo times. The measured phase maps at echo time TE_i can be calculated as:

$$\Delta\Phi_i(\mathbf{r}) = \gamma\Delta B(\mathbf{r}) TE_i + \varepsilon_i(\mathbf{r}) + \varepsilon_w(\mathbf{r}), \tag{1.28}$$

where $\Phi_i(\mathbf{r})$ denotes the measured phase images at echo time TE_i, $\varepsilon_i(r)$ denotes systematic errors and $\varepsilon_w(\mathbf{r})$ denotes white noise, i denotes the number of the phase map scans. For a dual-echo field map acquisition, the field inhomogeneity can be obtained by subtracting two phase images acquired at different echo times, which is then calculated as:

$$\Delta B(r) = \frac{\Phi_2(r) - \Phi_1(r)}{\gamma(TE_2 - TE_1)}. \tag{1.29}$$

Typical systematic errors are fairly consistent throughout scans. Their effects can be readily cancelled out by subtracting the phase maps with each other.

The selection of echo time difference, ΔTE, affects the quality of the final estimated field map. On the one hand, long echo time difference produces more field inhomogeneity induced phase accrual. Thus, the noise-induced phase variations have relatively small effects on the measured phase maps. However, the prolonged echo time difference gives rise to more severe phase wrapping problems. On the other hand, measurements become noise-dominated with a short echo time difference. Therefore, it is crucial to choose the appropriate echo time difference with a trade-off between reducing phase variations and improving the phase measurement robustness. The commonly used echo time difference is about a few milliseconds to avoid the ambiguity of the 2π wraps. To circumvent the errors, Dagher J. et al. proposed a 3-echo field mapping method together with a novel echo time optimiser to deliver a robust estimation of the B_0 field distribution [62]. In practice, the low-order shimming is only capable of compensating the large-scale field inhomogeneity. Thus, small noise in measured field maps will not significantly affect the final shimming results. Therefore, the dual-echo field mapping method is generally sufficiently robust to map the field distribution in the sample for shimming purposes.

As field mapping sequences are sensitive to the field inhomogeneity, the measured field maps can suffer from geometric distortion or signal loss in regions of large field inhomogeneity. These errors in the measured field maps will be translated to the final calculated shim settings. In **Chapter 3**, we will investigate methods to mitigate the effects of measurement errors on the final shimming quality.

1.4.2 Field mapping techniques

One conventional B_0 mapping technique is 3-dimensional (3D) volume-based field mapping, which estimates the field based on 3D phase images acquired at different

echo times (TEs). The GRE-based 3D volume field mapping methods provides a robust estimation of a B_0 field map, but often takes a minute or more to acquire. Multi-echo field mapping can improve robustness further, but it comes at the expense of increased scan time [63]. Methods for accelerating the B_0 field mapping have been investigated in the last decade. Single-shot techniques, such as echo-planar imaging (EPI) [64, 47] and spiral-based [65, 66] acquisitions, have been used for field mapping. Although these methods are fast, distortion and signal loss in the acquired EPI or spiral images will be translated into the field maps and will affect the accuracy of shim determination. Alternatively, balanced steady-state free-precession sequences [67] have been demonstrated to provide B_0 field maps with high SNR on a short time scale. However, this performance is degraded if a wide range of frequencies are present in image voxels. In that case, the field estimation will be not accurate in the regions of large field inhomogeneity (e.g., close to sinuses or ear cavities).

In addition to 3D volume-based field mapping methods, projection-based methods are used, which allow rapid shim determination based on the B_0 field distribution along the field map projections (e.g. the fast, automated, shimming technique by mapping along projections) [68, 69, 70, 71]. This method is typically based on the assumption that the shim field can be well described by a sum of spherical harmonics, which is good for local shimming targets, such as spectroscopy. However, it is not applicable for some other shimming applications, such as slice-wise dynamic shimming, because of limited through-slice field map data available for determining shim settings for thin slices.

Instead of measuring the field directly, it has been suggested to computationally model the susceptibility-induced field inhomogeneity [72, 73, 33, 74]. In these methods, the field inhomogeneity is calculated by convolving a 3D dipole field kernel with a predicted susceptibility distribution. However, this relies on additional high-resolution

anatomical images acquired from both CT and MRI to build a subject-specific accurate susceptibility model for B_0 field estimation. The long computational time and high memory cost make this method inconvenient for clinical applications.

Conventional field mapping sequences are conducted prior to any scans and are not capable of capturing time-varying field fluctuations induced by physiological noise that arise from cardiac-respiratory effects and subject motion [75]. Instead of constraining subject motion or physiological behaviour, integration of fast field mapping techniques (e.g., a single-shot magnetic field mapping or field navigators), prior to execution of image acquisition, has demonstrated its feasibility to map dynamic off-resonance field drifts [76]. Roopchansingh V. et al. introduced a single-shot field mapping sequence in conjunction with imaging interleaves to map the dynamic field changes during image acquisition [77].

1.4.3 Field map post-processing

In most cases, the acquired raw field maps cannot be directly used for shim calculation, due to image artefacts, like phase wraps. To recover the signal and enhance the fidelity of field map measurements, image post-processing method are used. This section gives a brief introduction and discussion of two field map post-processing techniques, field unwrapping and denoising.

Phase unwrapping

As signal only takes phase value in the range of 2π, the acquired phase value jumps when it hits the boundaries. Thus, artificial phase wraps occur in the measured phase

maps. The phase accumulation induced by field inhomogeneity at echo time TE can be described as

$$\phi(\mathbf{r}, TE) = \gamma \Delta B(\mathbf{r}) TE + \delta(\mathbf{r}) + 2\pi k_n, \tag{1.30}$$

where k_n represents the number of wraps in the measured phase maps, $\delta(r)$ represents the measurement noise at location $\mathbf{r}$=(x, y, z).

Typical phase unwrapping techniques include region growing and neighbouring filtering. Among these methods, the 3D phase unwrapping method PRELUDE (Phase Region Expanding Labeller for Unwrapping Discrete Estimates) proposed by Jenkinson M. et al. is widely adopted [78]. He used a region-merging algorithm to minimise phase difference across boundaries of different regions. The method is demonstrated to be fast and robust for most phase unwrapping cases. However, this method solves problems by optimising local optimal choice for fast speed at the expense of some accuracy. Figure 1.9 shows two sagittal images of the field maps acquired from a subject with (a) and without (b) phase unwrapping artefacts.

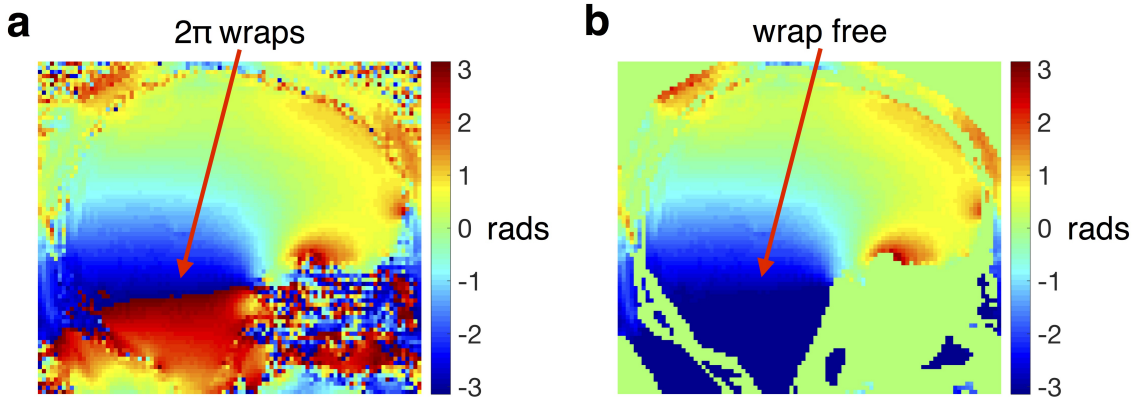


Figure 1.9: Sagittal images of the field maps acquired from a subject with (a) and without (b) phase unwrapping artefacts.

Denoising

The field map estimation can be noisy in the regions of low spin density. Typical image denoising techniques involve using low pass filters to smooth out abrupt phase change induced by noise. However, these methods can also remove the physiology-related phase change and degrade the information delivered by measured field maps. Alternatively, sophisticated methods are available which involve using statistical models to reduce the effects of noise in field map estimation [79].

1.5 Advanced B_0 shimming techniques

The **Section 1.4** presented various techniques available to estimate B_0 field distribution in the brain. Following the description of the shimming procedures indicated previously, this section introduces different shimming techniques to homogenise magnetic fields and gives a discussion on their advantages and limitations.

1.5.1 Active shimming

Active shimming involves using room temperature shim coils to generate a corrective magnetic field. As the magnetic field distribution of source-free systems satisfies the Laplace equations, spherical harmonics, the solution of Laplace equation, are considered as a set of ideal orthogonal basis to describe any source-free magnetic field [32] (Table 1.3). Active shim coils are typically designed to generate a field to resemble a specific spherical harmonic term. The magnetic field can be roughly resembled by a linear combination of spherical harmonic fields generated by different shim coils.

Notation	Order	Degree	Spherical coordinates	Cartesian coordinates
C	0	0	1	1
X	1	-1	$rsin\theta cos\theta$	x
Y	1	1	$rsin\theta sin\theta$	y
Z	1	0	$rcos\theta$	z
Z2	2	0	$(r^2(3cos^2\theta - 1))/2$	$z^2 - (x^2 + y^2)/2$
ZX	2	1	$r^2 sin\theta cos\theta cos\theta$	zx
ZY	2	-1	$r^2 sin\theta cos\theta sin\theta$	zy
X2Y2	2	2	$r^2 sin^2\theta cos^2\theta$	$x^2 - y^2$
XY	2	-2	$r^2 sin^2\theta sin^2\theta$	xy

Table 1.3: List of spherical harmonic orders and the corresponding spatial functions.

Active shimming techniques are marked by the adaptability of generating a corrective field via changing the electric currents in the shim coils. This design allows full automation of the shimming process and offers convenience for researchers and technicians to tailor magnetic fields [71].

A typical method of calculating the shim currents minimises the residual field norm in user-defined ROIs using a linear least-squares approach, which can be described by

$$min\left\{\|\Delta B - S \cdot I\|^2\right\} \tag{1.31}$$

where ΔB denotes a vector containing field inhomogeneity information in the ROIs, I denotes a vector of shim currents in different shim channels, S denotes a $n \times m$ matrix containing the distribution of shim fields generated by shim coils, where n is the number of points in the ROIs and m is the number of shim channels. The shim currents can be readily calculated by using a pseudoinverse matrix:

$$I \ = S^+ \cdot \ \Delta B, \tag{1.32}$$

where S^+ denotes a pseudoinverse of the shim matrix S. Various advanced shim optimisation methods will be presented in **Chapter 3**.

The techniques of using active shim coils to compensate for static field inhomogeneity are commonly termed static shimming. The shim settings are optimised once at the start of the scan based on the field map acquisition and remain fixed throughout the experiment [68]. Shim coils available in most clinical MR scanners can only generate corrective fields up to second-order spherical harmonics, which are low in spatial orders compared to large field inhomogeneity in local structures. Thus, such strong local field distortion is not correctable using low-order shims. To improve the field homogeneity in ROIs, Wilson J et al. proposed to use a software-based shim calculation technique that preferentially homogenises field in specific regions [80]. However, this approach compromises field homogeneity in other regions.

With the trend of going to higher field strength, it has become evident that low-order shimming is no longer capable of providing a sufficiently homogeneous magnetic field over the entire sample. To circumvent the challenges, 3^{rd}- and 4^{th}-order shim coils have been manufactured to produce shim fields with higher spatial orders. The feasibility of using higher-order shims in homogenising magnetic field has been demonstrated in MR spectroscopy and MRI [12, 81]. However, high-order shimming is demanding and expensive on the hardware and the bore space.

1.5.2 Local shimming

Recent research on cognitive functions, e.g., decision-making and working memory, has led to a growing interest in shimming the field inhomogeneity in the prefrontal cortex [82]. As shown in the sagittal view of the example field map in Figure 1.9, strong field inhomogeneity presents in the regions of the prefrontal cortex. Thus, the low-order shims are not capable of producing good field homogeneity in the prefrontal cortex.

To circumvent the limitations of the low-order spherical harmonic shims, Wilson J. et al. proposed to use highly diamagnetic materials in the mouth, referred to as the mouth shims, to compensate for the strong local field distortion on top of spherical harmonic shimming [83, 84, 85, 86, 87]. The schematic of the mouth shim is shown in Figure 1.10. The use of the localised passive shims is a relatively low-cost method compared to using the high-order shims. However, it has not gained popularity in clinical experiments. As the magnetic field inhomogeneity distribution in the brain varies between different subjects, the complicated design of passive dielectric materials and the long preparation time of tailoring subject-specific mouth shims add to the burden of technicians [88].

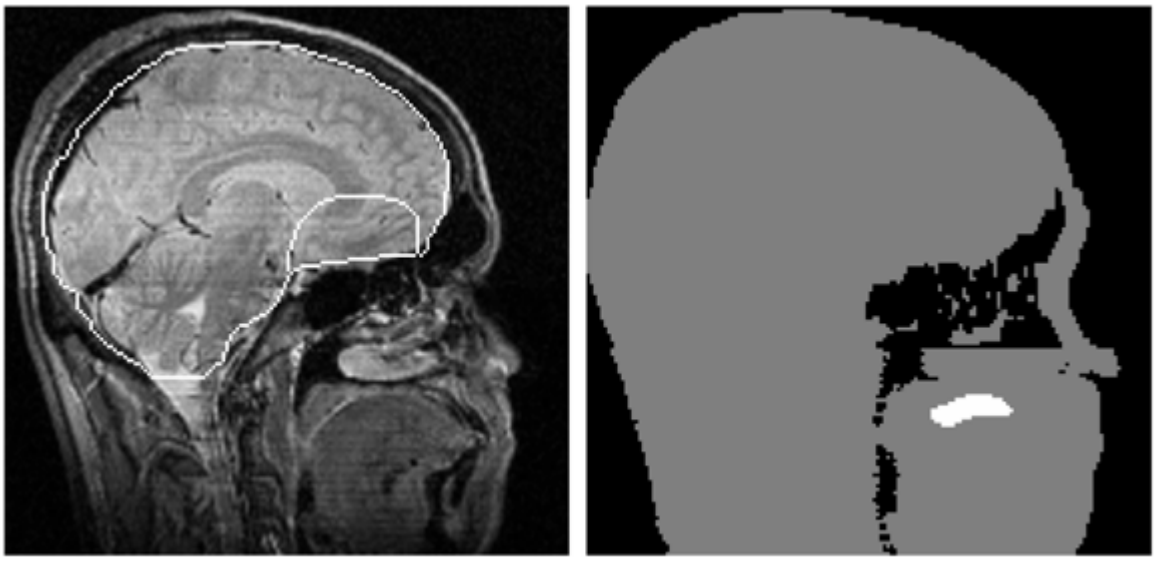


Figure 1.10: The mid-sagittal view of (a) the MR structural image of an example subject's head and (b) the corresponding susceptibility map used for simulating B_0 field in the head, which is comprised of air (black), tissue (grey) and a mouth shim (white). (adopted from Wilson J thesis [89])

An alternative method to correct field inhomogeneity locally is using a set of electric coils to generate a corrective field. Hsu, J. J. et al. demonstrated that using a set of intra-oral electric coils can effectively compensate for field inhomogeneity near sinuses. This technique of using localised shim coils is marked by its adaptability in

providing a corrective magnetic field with high spatial orders by changing the currents in the coils [90]. Based on the studies, Juchem C. et al. attempted to use a set of external electric coils near regions of large field inhomogeneity to achieve localised shimming [91]. Next, he extended the concept of the localised coil technique to the multi-coil technique which involves using a powerful electric coil array to generate required shim fields in a large region with high spatial orders (Figure 1.11). The multi-coil technique has demonstrated its feasibility on the mouse and the human brain [92, 93, 94, 95]. The use of localised shimming techniques enables a correction of subject-specific strong local field inhomogeneity on top of the conventional spherical harmonic shimming. However, the multi-coil techniques need complex hardware design. Moreover, the incorporation of additional coil elements may affect the patient comfort.

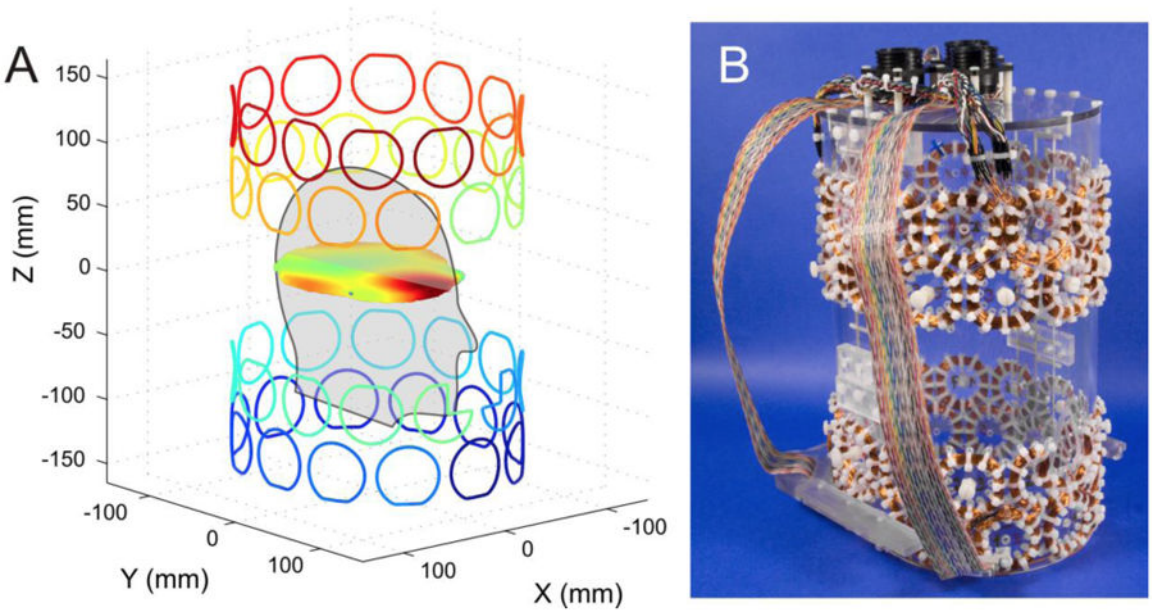


Figure 1.11: Theoretical design (A) and experimental realisation (B) of the multi-coil setup. The multi-coil array is comprised of 48 independently driven, localised electric coils. (adopted from Juchem C et al. [92])

1.5.3 Dynamic shimming

Dynamic shimming extends the concept of localised shimming. This technique is achieved by updating shim settings during MR imaging to perform B_0 shimming in small subvolumes (e.g. slices for multi-slice imaging) [96]. The rationale of dynamic shimming is based on the fact that reduction of the size of ROIs can potentially reduce the spatial frequency of field inhomogeneity. Thus, the low-order shims can better resemble the field inhomogeneity and refine magnetic field more efficiently compared to whole brain static shimming. Dynamic shimming has demonstrated this feasibility for various imaging applications (e.g., the brain, spinal cord [97], breast [98], and MR spectroscopy [99, 100]).

Dynamic shimming technique was first demonstrated back in 1996 and implemented by Blamire A.M. et al. He designed an external hardware console to achieve real-time updating of the gradient fields [16]. In 1997, Morrell G. et al. demonstrated the feasibility of slice-wise first-order dynamic shimming on a wide range of sequences (e.g., GRE and EPI sequences [17]). Following that, Zhao Y. et al. demonstrated the effectiveness of higher-order dynamic shimming in simulation [101]. Then, in 2005, Koch K. et al. experimentally realised second-order dynamic shimming. Additionally, he implemented the "pre-emphasis" method to compensate for eddy-current effects induced by updating second-order shims [18]. In 2010, Juchem C. et al. extended dynamic shimming to include third-order shims and demonstrated its feasibility at 7T [102].

The concept of dynamic shimming has been extended to a wide range of applications. First, dynamic shimming has been demonstrated feasible to achieve real-time correction of time-varying field fluctuations induced by physiological noise and subject motion [103, 104]. Second, Poole M. showed that shimming performed over a

series of cuboidal subvolumes is more efficient in homogenising magnetic fields than the plane volumes in simulation [105]. However, the volume-based imaging technique has not been well developed in realistic settings. Thus, extra effort is demanded to realise this technique. Third, dynamic shimming with multi-coil technique has been demonstrated capable of delivering better field homogeneity compared to spherical harmonic shims in the brain of mouse, rat and human [95, 92, 93, 94].

There are also some limitations that hinder the performance of dynamic shimming.

First, in dynamic shimming, fast shim switching induces strong eddy currents in cold conductors near shim coils. The eddy currents distort desired shim fields and further degrade shimming quality. To solve this problem, eddy current compensation is required. **Chapter 2** presents methods to characterise and compensate for eddy current effects, and then experimentally demonstrates the performance of dynamic shimming with eddy current compensation.

Second, shimming in small volumes is intrinsically challenging in comparison to large volumes. Using the field map data with limited spatial coverage for shim calculation is likely to cause the overfitting problem, leading to unstable shim solutions. To mitigate the problem, we investigated a range of regularisation-based methods to improve the robustness of shim calculation in **Chapter 3**.

Finally, dynamic shimming requires high-resolution field maps to provide detailed spatial information on the field distribution for shim calculation. Acquisition of such field maps can take minutes which prolongs the scan time and may affect patient comfort. In order to shorten the preparation time spent on the field map acquisition, **Chapter 4** proposes a fast template-based field map prediction method without measuring the B_0 field to accelerate the field map acquisition stage and evaluates its performance.

Chapter 2

Calibration and implementation of dynamic shimming

High-quality dynamic B_0 shimming requires high-fidelity shim fields. In practice, idealised shim fields can hardly be achieved. The error responses of the shim coils to the input electric currents encompass contributions from not only the imperfection of the shim systems but also errors arising from other sources. Shim field errors can be mainly classified to two categories: static field errors and dynamic field errors. Static field errors are caused by factors, including inaccuracies in the manufacture of the magnet and coils, shim coil non-linearity and concomitant fields. Dynamic field errors are induced by factors, such as eddy current effects, thermal noise, mechanical vibrations and complex amplifier behaviour. Thus, without proper measurement and characterisation of the behaviour of the shim systems and compensation for the unwanted shim fields, desired dynamic shimming results can hardly be achieved [102].

The **aim** of this chapter is to present methods to characterise and calibrate shim fields (used for dynamic shimming) and then evaluate the performance of dynamic shimming in experiment and simulation. This chapter is comprised of four sections. **Section 2.1** introduces the configuration of our dynamic shimming hardware and how different MR components and electronic devices communicate with each other to

perform shim updating. **Section 2.2** presents methods for characterising the static spatial distribution of the magnetic field generated from each shim channel. The focus of **Section 2.3** turns to methods to characterise and compensate for eddy-current-induced fields. We specifically want to compare the performance of eddy current compensation (ECC) using two characterisation methods, an image-based method and a method based on a state-of-art field camera. Finally, **Section 2.4** presents the procedures to implement dynamic shimming and evaluates its performance on one subject in experiment and on 143 subjects in simulation.

2.1 MR scanner and dynamic shimming hardware configuration

Experiments throughout this thesis were performed on a Siemens Magnetom whole-body 7T MRI scanner (Siemens Healthineers, Erlangen, Germany), equipped with active shielded 70 mT/m gradients and 2nd-order spherical harmonic shim coils without active shielding. Imaging data were acquired with a 32-channel head coil (Nova Medical, Wilmington, MA).

The first-order shim updating was achieved using gradients, by setting a slice-wise gradient offset in sequence programming. On standard MR systems, second-order shim settings cannot be updated in real-time. In order to perform dynamic control of all second-order shims, specially designed dynamic shimming hardware is required. In our setup, the dynamic shimming system is comprised of three components: the host computer, the Siemens Transmission-Array system (pTx) and shim amplifiers from Resonance Research Inc. (RRI, Billerica, MA, USA). They communicate with each other to conduct real-time shim updating.

Each RRI shim amplifier has one output on one of the second-order shim coil connectors, and one analogue input port to receive a signal for real-time control of the shim currents (level 3). The input port is connected to a digital signal processor (DSP) with digital-to-analogue conversion. The DSPs receive user-defined digital waveforms of shim pulses.

In order to individually control the digital shim waveforms sent to the amplifiers, the pTx array is used (level 2). Similar to the standard scanner system where one MPCU computer (Measurement and Physiologic Control Unit) controls 3 gradients, called MPCU0, the pTx system assigns two additional MPCUs for second-order shim control where the Z2, ZX, ZY shims are grouped for MPCU1, and the X2-Y2, XY shims for MPCU2.

Two host computers, master and slave, are configured to allow three sequences running in parallel from pTx (level 1) to control the different MPCUs. The master computer conducts sequences to control the RF pulse and the 3 imaging gradients. The slave computer is synchronised to the master and controls shim waveforms in the 5 second-order shim channels. This setup allows full shim waveform control of all second-order shim channels within the same sequence programming framework as for the imaging gradients.

Additionally, the second-order shim currents produced by amplifiers are subject to ECC. The MPCU control allows exponential ECC as for the gradients. Cross-term ECC is only available for shim terms within assigned MPCUs. Taking the Z2 shim for example, cross-term ECC can only be carried out for the ZX and ZY shim terms. The B_0 ECC settings were fed back to MPCU0 to compensate B_0 eddy fields by demodulating transmit and receive frequency. Each shim amplifier is designed to work within current constraints of 10A and the summation of them is limited below 20A.

The RRI shim control schematic is shown in Figure 2.1.

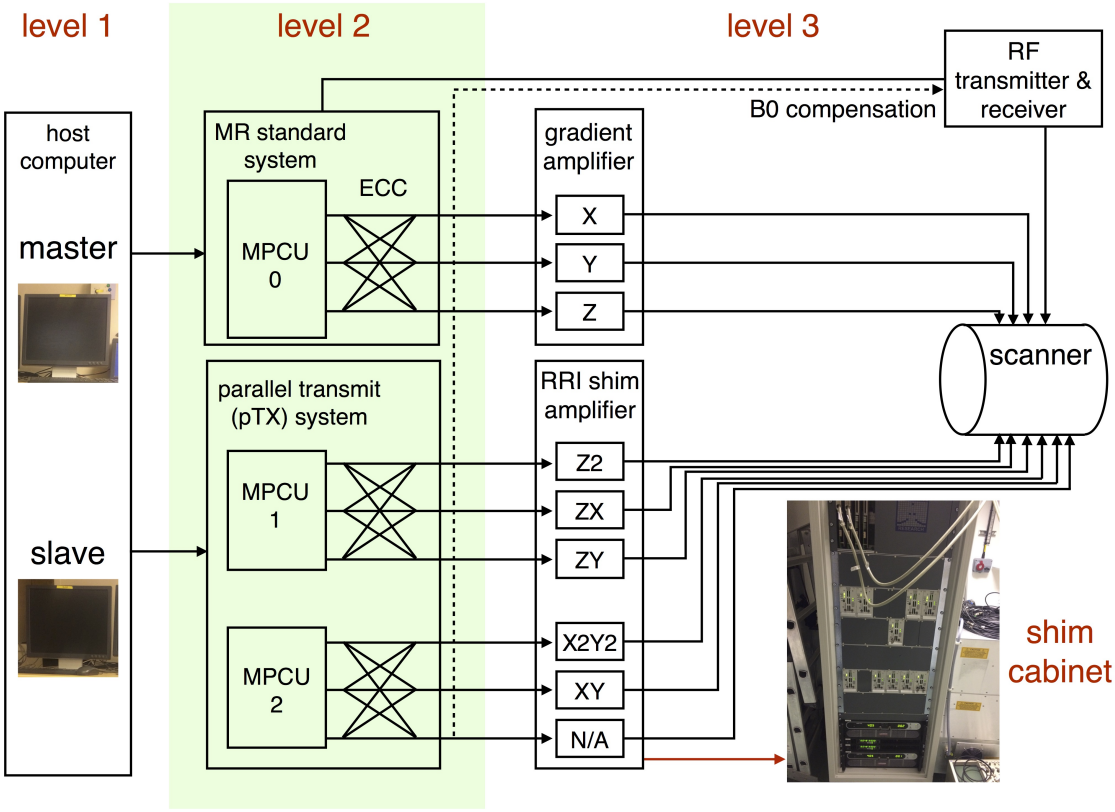


Figure 2.1: Siemens and RRI shim control schematic. MPCU0, 1, and 2 represent three gradient controllers used to control 3 gradient coils, and 5 second-order shim coils through vendor-provided consoles and a RRI shim cabinet separately. B_0 field offsets were achieved by updating the frequency of RF pulses.

2.2 Static shim field characterisation

In this section, we introduce a linear-model-based method to characterise the static field distribution generated by shim coils.

2.2.1 Methods

The method for characterising the spatial field distribution from each shim channel consists of four stages: field measurement, spherical harmonic decomposition, linear regression, and significance test. In the first stage, the spatial distribution of the magnetic field generated by each shim channel is measured with a range of shim strengths. Second, the measured field maps are then decomposed into spherical harmonic functions using least-squares optimisation. Third, the decomposed spherical harmonic data are fit by linear models. Finally, the significance of all decomposed spherical harmonic terms to input shim strengths is tested. The procedure is illustrated in Figure 2.2.

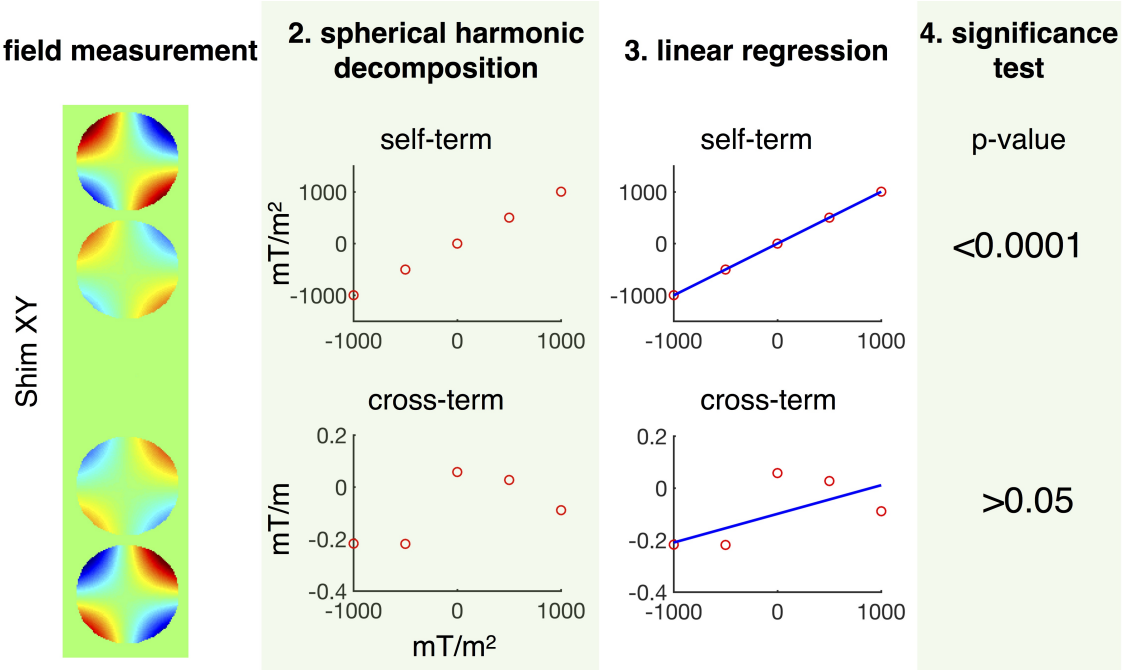


Figure 2.2: Procedures for characterising the spatial field distribution generated from each shim channel. For each shim channel, (1) the spatial distribution of the magnetic field generated by each shim term is measured with a range of shim strengths, and then (2) the data are decomposed into spherical harmonic functions from zeroth to fourth order using least-squares optimisation. (3) The relationships between decomposed spherical harmonic terms and input shim strengths are characterised by linear regression models. (4) significance tests are conducted to evaluate the correlation between spherical harmonic terms and the input shim strengths

Shim field mapping

We resolved the static distribution of the generated magnetic fields from each shim channel by utilising high spatial resolution field maps. The field maps were acquired based on an 18-cm-diameter homogeneous spherical oil phantom placed in the isocentre of the scanner's FOV with an asymmetrical dual gradient-echo sequence (voxel size=2 $\times$ 2 $\times$ 2 mm^3, TE1=4.08 ms, ΔTE=1.02 ms, flip angle=39°, TR=620 ms). The field maps were then unwrapped using PRELUDE algorithm [78]. The ROIs were identified using the magnitude image of the field map acquisition using BET (FMRIB's Brain Extraction Tool) [106]. A calibration field map was acquired with a default shim applied for background field correction at the beginning of the experiment. To characterise the spatial dependence of the magnetic field generated by each shim coil, field maps of the oil phantom were repeatedly measured by adding different shim strengths within a range of shim values ($\pm$500mT/m for first-order shims and $\pm$2000mT/m^2 for second-order shims, with increments of 20% of maximum shim value), while keeping the other shims zero. To mitigate systematic noise, measurements of the shim fields were conducted several times over a period of one month. The measured magnetic fields generated by second-order shims at different shim strengths are shown in Figure 2.3.

Spherical harmonic decomposition

As introduced in **Chapter 1**, spherical harmonics are ideal basis functions to describe the magnetic field. To quantify the magnetic field produced by each shim channel, the field map data were decomposed into spherical harmonic functions from zeroth to fourth order using least-squares optimisation. The corresponding Cartesian representations of spherical harmonic terms are listed in Appendix Table A.1. A matrix

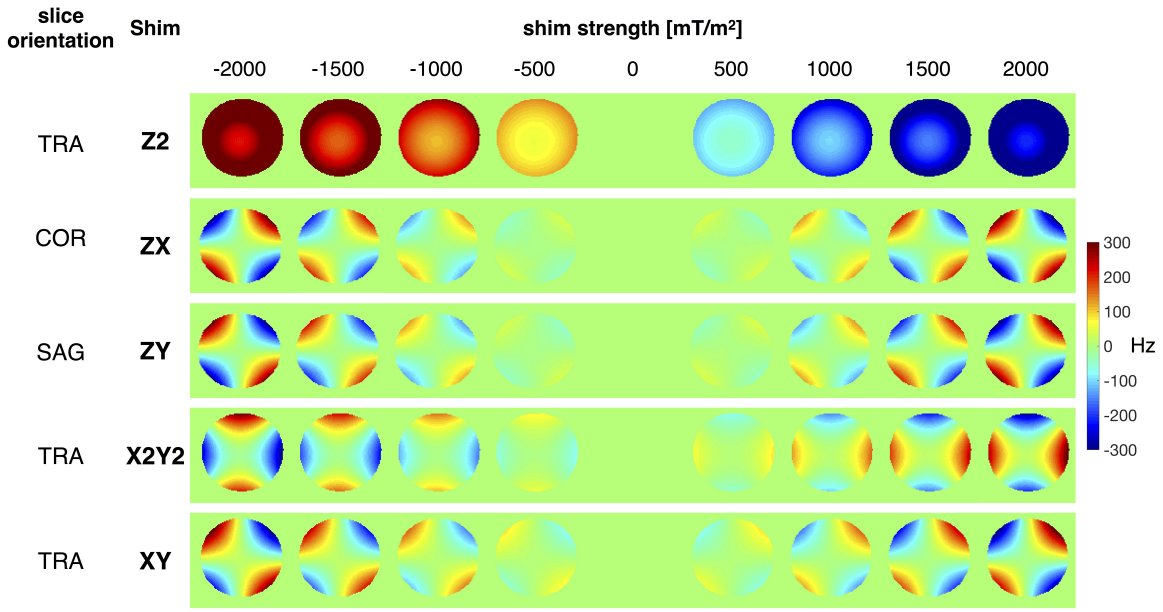


Figure 2.3: The measured shim fields generated by second-order shim channels at different shim strengths for a slice located at the centre of the scanner with different orientation. TRA: transverse plane, COR: coronal plane, SAG: sagittal plane.

of the coefficients of decomposed spherical harmonics S was used to indicate the field distribution generated by shim channels, which is given by

$$\mathbf{S} = \begin{bmatrix} 1 & X_X & X_Y & X_Z & \dots & X_{XY} \\ 1 & Y_X & Y_Y & Y_Z & \dots & Y_{XY} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & XYZ_X & XYZ_Y & XYZ_Z & \dots & XYZ_{XY} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \\ 1 & X4+Y4_X & X4+Y4_Y & X4+Y4_Z & \dots & X4+Y4_{XY} \end{bmatrix}, \tag{2.1}$$

where coefficients in each column represent contributions from decomposed spherical harmonic components of the magnetic fields generated by a specific shim channel, indicated in the subscripts. Conventionally, the spherical harmonic terms that match the shim channels indicated in the subscripts, are referred to as self-term (e.g., X_X which means the X component that corresponds to applying a X shim). The rest of spherical harmonic components are referred to as cross-terms (e.g., $Z2_X$ which means the Z2 component that corresponds to applying a X shim).

Linear regression

In order to model the spatial distribution of generated shim fields at different shim strengths, we fit linear regression lines to the data using the Matlab embedded function 'fitlm' to quantify the dependence of each spherical harmonic term on the input shim strengths.

Figure 2.4 shows the spherical harmonic decomposition of the magnetic fields generated by the XY shim with a range of shim strengths. The results indicate that measured self-terms were highly linearly dependent on the input shim strengths. Some cross-terms show relatively small contributions to the generated shim fields but linearly dependent on the input (e.g., the X, X2-Y2, and Z2XY cross-terms). Other shim terms change nonlinearly with the input shim strengths (e.g., the 3X2-Y2 and X(X2-3Y2) cross-terms).

We modelled shim fields with a superposition of spherical harmonic terms showing high linearity with the input shim amplitudes. F-test was conducted for each shim term and P-value was used as a measure of its statistical significance. The decomposed spherical harmonic components that have relatively small contributions or low correlation ($P < 0.001$) to the input shim strengths were not included in the shim field model (e.g., the X(X2-3Y2) term).

The coefficients of the decomposed spherical harmonic components for all first- and second-order shim channels are listed in Table 2.1. For a direct comparison of the contributions of different spherical harmonic components, their coefficients were scaled to the maximum field offsets within a 10-cm-radius sphere, given in percentage of the input shim strengths, for specific shims. It is evident that highly correlated spherical harmonic terms are self-terms whose coefficients are around 100% of the required input shims. Some cross-term components show strong dependence on the inputs (e.g.,

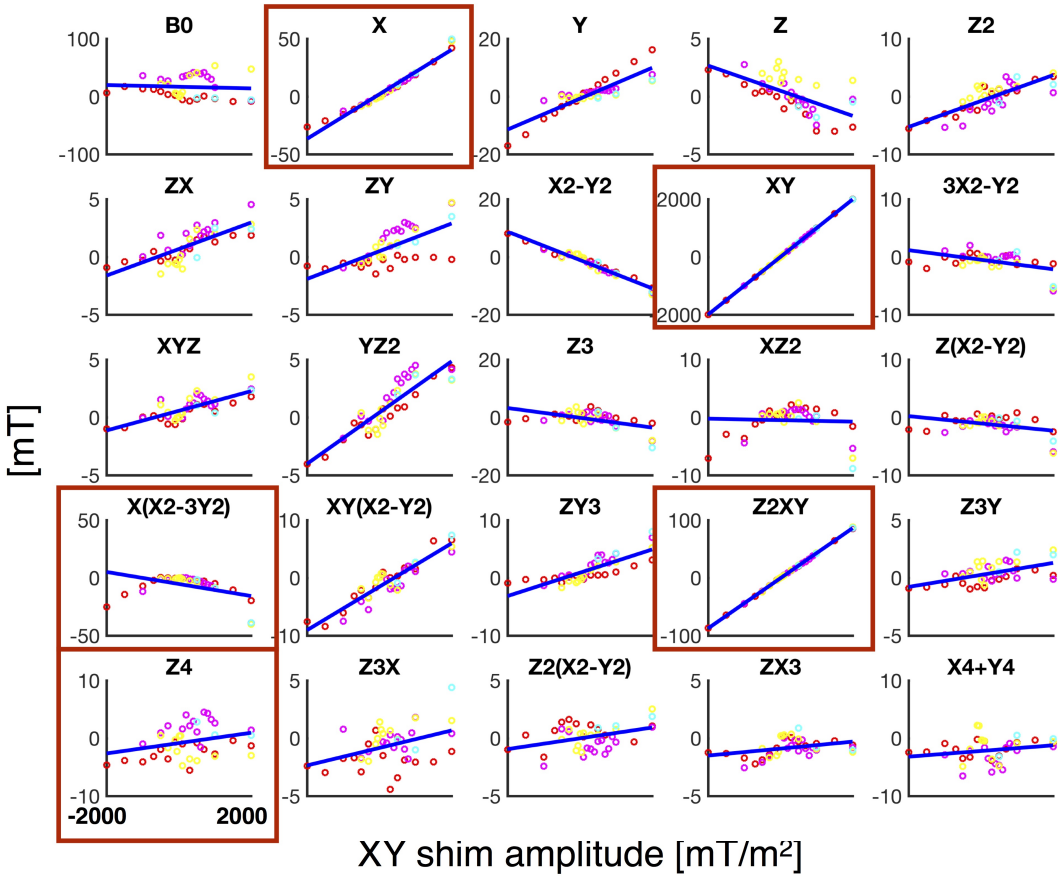


Figure 2.4: The measured self-term and cross-term spherical harmonic components of the XY shim. The lines coloured in red, yellow, pink, and cyan, indicate the measurements at different times over a month. The blue reference line denotes least-square fit results of the measurements.

Shim	X %	Y %	Z %	Z2 %	ZX %	ZY %	X2Y2 %	XY %
X	102.9	0	0	0	-0.9	0	0	1.9
Y	-0.1	100.2	-0.2	-0.9	0.1	-1.3	-1.7	0.5
Z	0.2	0.1	101.4	-1.3	0	1.4	0.2	-0.1
Z2	0	0	0	102.7	0	0.3	-0.3	0.2
ZX	0	0.1	0.2	0	98.8	-0.4	0.1	0.1
ZY	0.1	0	0	0	0	97.6	-0.3	0.1
X2-Y2	0	0	0	-0.2	0.2	0.1	96.4	-0.5
XY	0	0	0	0.2	-0.1	0.2	-2.1	99.7
3X2-Y2	-0.1	0	0	0	0.1	0	0	-0.1
XYZ	0	0.1	0	0	0	0	0	0.1
YZ2	0	0	0	-0.1	0	0	0	0.2
Z3	-0.1	0	-0.2	0	0	0	-0.2	-0.2
XZ2	0	0	0	-2	1.7	0.1	-0.3	0
Z(X2-Y2)	0.1	0.1	0	0	0	-0.7	0.2	-0.1
X(X2-3Y2)	0	0	0	0	0	0	1.6	-0.5
XY(X2-Y2)	0.1	0	0	0	0	0	0	0.4
ZY3	0.2	0	0.2	0	-0.1	0	-0.1	0.2
Z2XY	-0.1	0	0.5	-0.1	0	0.1	-0.2	4.3
Z3Y	0	-0.1	0	-0.2	-0.1	3	0	0.1
Z4	-0.2	0	0	4.3	0	-0.1	-0.2	0.1
Z3X	-0.1	0	0.3	0	3.4	0	-0.2	0.1
Z2(X2-Y2)	0.2	0	0	0.1	0	0	3.9	0
ZX3	0.1	-0.1	0	0	0.3	0	0.1	0
X4+Y4	0	0	0	-0.2	0	0	0.4	0

Table 2.1: Maximum field offsets of the self-term and cross-term within a 10-cm-radius sphere, given in percentage of the input shim strengths, for specific shims. Maximum field offsets of spherical harmonic terms larger than 1% (blue) of the corresponding input shim strengths are highlighted, as well as spherical harmonic term larger than 90% (red).

the Z2 and Z4 cross-terms while switching the Z2 shim). Most cross-terms show field shift below 0.5% of the expected shim strengths.

2.2.2 Discussion

We have presented a shim field characterisation method capable of modelling the magnetic field generated from shim coils using linear models in static operations, based on the field maps of the phantom. As shown in Figure 2.4, most strong spheri-

cal harmonic terms conform to the linear model (e.g., the X and Z2XY cross-terms). Characterising static responses of shim channels in this way does not require any additional equipment. Several shim field calibrations can be conducted to assess the consistency of the shim coil behaviour over time. Second, the method allows the measurement of magnetic field distribution generated by each shim channel at relatively high spatial resolution. Due to the rich amount of information about the field distribution, a more comprehensive characterisation of shim fields can be achieved by including more spherical harmonic terms in the shim model.

The field-map-based field characterisation method has some disadvantages. First, some shim terms that are nonlinearly depending with input shim strengths cannot be captured by the linear models (e.g. the X(X2-3Y2) cross-term). These nonlinear shim field characteristics are attributed to complex system behaviours. To overcome this challenge, nonlinear models can be further investigated for better shim field characterisation. Second, the shim fields generated by the shim coils are sources of field inhomogeneity. For the GRE-based field mapping sequence, high strength shim fields will lead to large image distortion artefacts in the acquired field maps and further degrade the accuracy of shim field measurements. Using high-bandwidth GRE sequence can mitigate this problem at the cost of SNR. In our implementation, the shim fields generated for characterisation are limited in a reasonable shim strengths. Therefore, the field distortion artefact did not much affect the characterisation results.

2.3 Eddy current compensation

Section 2.2 presented methods to measure and characterise actual static field distribution generated from shim channels. Besides static field errors, dynamic shim switching can cause spatially and temporally varying distortion to shim fields. The

distortion resulting from eddy currents generated in the proximate cold conducting structures of the magnet are predominant. Without correction of this artefact, eddy-current-induced fields could reach up to 40% of the amplitude of the switched shims and last a few seconds [107, 108]. Therefore, minimising eddy current effects is one of the most important topics for dynamic shimming

Techniques to counteract the eddy current effects mainly include:

- Active shielding. Using a set of coils to generate compensation magnetic fields to cancel out undesired fields, coming at the cost of bore space and utilisation. [109, 110]
- Sophisticated magnet design. Increasing the relative dimension of the magnet bore against the size of active coils [111]
- Pre-emphasis. Using high pass filters to pre-distort the input waveform to counteract the expected eddy-current-induced field distortion [112]

Among these techniques, the pre-emphasis approach is a relatively inexpensive, space-saving solution to offset eddy current effects and is well demonstrated for dynamic shimming applications [113, 114, 107, 102, 108]. The pre-emphasis method was adopted in this work to perform ECC on second-order shim channels. The following subsection introduces the theory of the pre-emphasis method and the procedures to determine the corresponding pre-emphasis parameters for ECC.

With pre-emphasis, ECC is achieved by overdriving shim currents to compensate for expected distortion in shim fields. Determination of pre-emphasis currents requires the knowledge of the spatial and the temporal field variations induced by eddy currents after shim switching. Many studies have investigated methods to measure eddy-current-induced fields [111]. Among these methods, the image-based method is

conventionally used, which relies on measuring a series of field maps of a phantom over the eddy current settling period [115]. The image-based method is less demanding on hardware and easy to carry out. It has been well demonstrated for dynamic shimming applications.

Besides the image-based method, measuring spatial and temporal field dynamics using nuclear magnetic resonance (NMR) probes has gained much popularity [116, 117]. This technique relies on measuring the FID of samples in NMR probes under different field conditions [118, 119]. Measuring eddy-current-induced fields using NMR probes significantly decreases the scan time and allows relatively accurate description of field dynamics in space and time. Recently, a state-of-art field camera has been developed using high-performance NMR probes, which allows monitoring the spatiotemporal field dynamics in MR scanners at very high temporal resolution [103, 120]. The field camera has superiority in capturing temporal system behaviour like oscillatory responses [121]. Previously it has been demonstrated that ECC using a field camera for pre-emphasis parameter adjustment yielded eddy current settling time about 1ms for rapid shim updating [122] [108].

The **aim** of this section is to compare the performance of ECC using the image-based methods and a field camera, and to evaluate whether eddy current characterisation using field camera produces better ECC results for dynamic shimming than the image-based method. This section starts by introducing the model to describe eddy currents, in **Section 2.3.1**. Then, **Section 2.3.2** provides the theory of the "pre-emphasis" method. Following that, **Section 2.3.3** and **Section 2.3.4** present the procedures and results of ECC using the image-based method and a field camera, respectively. Finally, in **Section 2.3.5**, the ECC performance using the two methods are compared, followed by a discussion of their advantages and limitations.

2.3.1 Eddy current model

According to Faraday's law and Lenz's law, temporally varying magnetic fields induce eddy currents in surrounding conductive materials. The generated eddy currents impede the change of the magnetic fields. Conventionally, eddy currents can be modelled by a set of mutually coupled resistor-inductor electric current loops, as shown in Figure 2.5 [113].

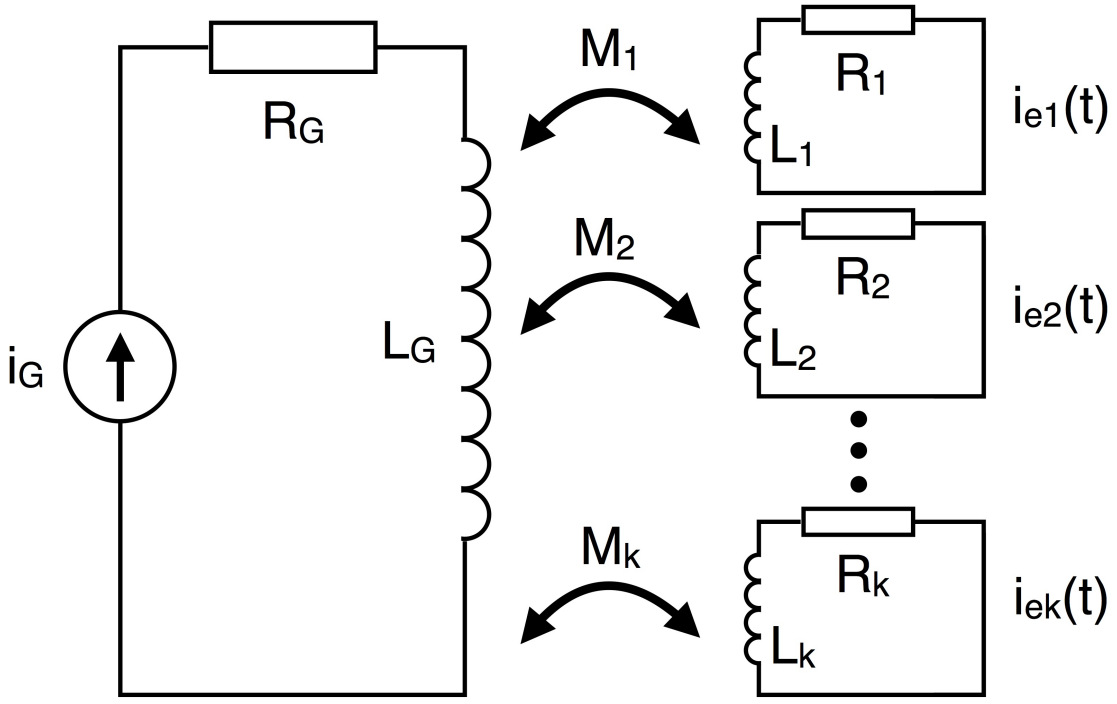


Figure 2.5: R-L circuit representation of the eddy current model. Switching the input shim currents i_G in the shim coils G induces eddy currents i_e in loop k ($k = 1 \cdots n$) that are mutually coupled to the shim coils. L and R denote inductance and resistance in the eddy current loops, and M denotes the mutual inductance between shim coils and eddy current loops.

The overall shim induced fields are the sum of the fields generated by the input currents in the shim coil G, and eddy currents running in separate loops. The shim fields generated by input electric currents in shim coils can be modelled by

$$G(t) = c_G i_G(t) + \sum_{k=1}^{n} c_{ek} i_{ek}(t), \tag{2.2}$$

where k denotes the index of eddy current loops, n denotes the number of eddy current loops, G denotes the shim coil loops. L, R and i denote inductance, resistance and currents in the eddy current loops. c_G denotes a correlation factor between the output shim fields and the input electric currents i_G in shim current loops, c_{ek} denotes the correlation in the eddy current loops k. The first term $c_G i_G(t)$ denotes the shim fields generated solely by input currents i_G in shim coils, and the second term $c_{ek} i_{ek}(t)$ denotes which induced by eddy currents i_{ek},

The relationship between the currents in the eddy current loops i_{ek} and the shim coil i_G can be defined as

$$L_k \frac{di_{ek}(t)}{dt} + R_k i_{ek}(t) + M_k \frac{di_G(t)}{dt} = 0, \tag{2.3}$$

where M denotes the mutual inductance between shim coils and eddy current loops in the surrounding conductors.

Based on Equation 2.2 and 2.3, the corresponding unit responses of the shim channel G can be described by the sum of exponential decay functions

$$G(t) = c_G \left(1 - \sum_{k=1}^{n} A * exp(-B * t)\right), \tag{2.4}$$

where $A = \frac{c_{ek} * M_k}{c_G L_k}$ denotes amplitudes of exponential functions and $B = \frac{R_k}{L_k}$ denotes time constant parameters. Thus, eddy current effects can be modelled using a set of exponential decay functions.

Pre-emphasis

If the outputs of the system are stationary and linearly dependent on the inputs, then the system is a linear time-invariant (LTI) system. Thus, the eddy current model presented in **Section 2.3.1** can be described by a LTI model. The output waveforms

of a shim system can be modelled by convolving the shim input waveforms I with the transfer function of the system S, which is expressed as:

$$O(s) \;=\; S(s) \;*\; I(s), \tag{2.5}$$

where s denotes frequency, $I(s)$ denotes the frequency-domain representation of the input shim pulses on shim channels and $O(s)$ denotes the corresponding output system responses which correspond to the magnetic fields generated by the shim system.

For a trapezoidal input shim pulse, a schematic of the output shim waveforms distorted by eddy currents are illustrated in Figure 2.6a.

The pre-distorted shim waveforms, $I_{distort}(s)$, modulated by the pre-emphasis filters, $H_{PE}\ (s)$, for compensating eddy current effects can be obtained by

$$I_{distort}(s) \;=\; H_{PE}\ (s)\ I_{in}(s), \tag{2.6}$$

where $I_{in}(s)$ denotes the input shim waveforms affected by eddy currents, as shown in Figure 2.6b.

In idealised situations, with pre-emphasis, convolution of the impulse responses of the pre-emphasis filters with the shim system responses should be equal to one, which is expressed as:

$$H_{PE}\ (s)\, S(s) = 1, \tag{2.7}$$

where $H_{PE}\ (s)$ denotes the impulse responses of the pre-emphasis filters and $S(s)$ denotes the impulse responses of a shim system. The pre-emphasised output shim waveforms are shown in Figure 2.6c.

Based on the exponential model of eddy currents as described in Equation 2.4, the impulse responses of the shim system S after Laplace transformation are defined as:

$$S(s) = 1 - \sum_{k=1}^{n} A_k \frac{s}{s + B_k}, \tag{2.8}$$

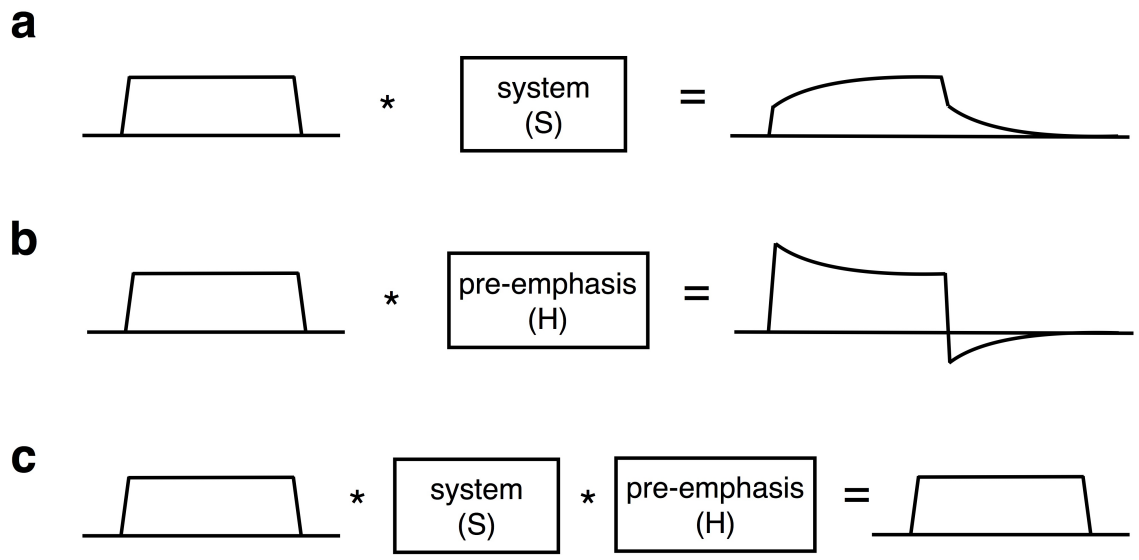


Figure 2.6: A distorted shim waveform due to eddy currents (**a**), a pre-distorted shim waveform (**b**), and a pre-emphasised shim waveform (**c**). S: Transfer functions of a shim system, H: Transfer functions of pre-emphasis filters.

where amplitude $A_k = \frac{c_{ek} * M_k}{c_G L_k}$ and time constant $B_k = \frac{R_k}{L_k}$.

According to Equation 2.7, the impulse responses of pre-emphasis filters are obtained from the inversion of impulse responses $S(s)$. For simplicity, with the assumption of mono-exponential decay of eddy currents, the corresponding impulse responses of shim pre-emphasis filters $H_{PE}\ (s)$ can be obtained by

$$H_{PE}\ (s) = \frac{1}{S(s)} = \frac{1}{1 - A_1 \frac{s}{s+B_1}} \tag{2.9}$$

The step responses of the pre-emphasis filters are given by:

$$h_{PE}\ (t) = 1 + A_{PE1} e^{-B_{PE1} * t}, \tag{2.10}$$

where amplitude $A_{PE1} = \frac{A_1}{1-A_1}$ and time constant $B_{pe1} = \frac{B_1}{1-A_1}$

Following on the works of Jehenson P., the exact pre-emphasis parameters can be readily determined in the inclusion of more exponential terms [113].

2.3.2 Image-based eddy current measurement

This section presents the image-based method to measure and characterise eddy-current-induced fields. This section begins by introducing the pulse sequence design and procedures of eddy field measurement. Following this, the ECC results using the image-based method will be presented.

Pulse sequence design

The image-based method requires two types of sequences.

- A phase mapping sequence running on MPCU0, see **Section 2.1**, to measure eddy-current-induced field change after shim switching. It is comprised of an adjustable echo time delay block, a RF pulse and a k-space acquisition block.
- A shim control sequence running on MPCU1 and 2 to control the waveforms on the second-order shim channels. It consists of a long trapezoidal shim set-up pulse on one of the second-order shim channels and an eddy current settling time block.

In our implementation, a GRE sequence was modified to realise the two functions. The schematic of the eddy-current-induced field measurements is illustrated in Figure 2.7. The shim control sequence generates a long shim pulse on one of second-order shim channels, followed by a long time block for eddy current settling prior to the next measurement. Then, the phase mapping sequence is applied. Each k-space line acquisition is preceded by a pre-defined delay time. The duration of the delay time is equivalent to the total duration of the shim set-up pulse and the pre-defined time delay block. In order to acquire the eddy-current-induced field distribution at

a specific time point, phase map measurements were repeated twice with echo time difference of 1 ms. The echo time difference is relatively small to mitigate effects of eddy-current-induced field variations between measurements.

In order to obtain the field distribution over a slice, the measurement above was repeated multiple times to achieve adequate k-space coverage. To measure the eddy-current-induced field variations over a period of time, the field map measurements were repeated several times at different time points after shim switching.

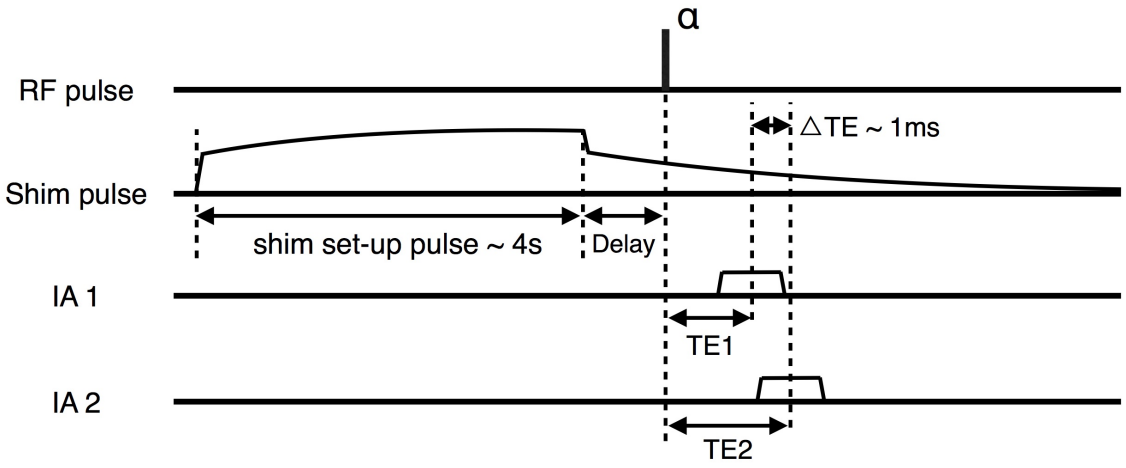


Figure 2.7: Eddy-current-induced field mapping using a modified GRE sequence. SG: pulsed shim gradient; IA: image acquisition (the phase encoding gradients and slice selection gradients are omitted for clarity); ST: pulsed set-up time.

Measurement of eddy-current-induced fields

Eddy-current-induced field measurements were performed on an 18-cm-diameter homogeneous spherical oil phantom positioned in the isocentre of the scanner's FOV. The modified GRE sequence was performed (voxel size=6.25 × 6.25 × 5.6 mm^3, slice=1, TR=10 s, TE1=4 ms, ΔTE=1 ms, FOV=200 × 200 mm^2, total duration $\approx$ 6 hours) to measure eddy-current-induced fields after shim switching at a slew rate of 2500 $mT/(m^2 \cdot ms)$. The duration of the shim set-up pulses was set to 4 s

which is sufficient for second-order eddy current settling prior to the k-space acquisition [102, 108, 107].

Here we aimed to achieve eddy current reduction down to 1% of input shim step about 10 ms. Thus, eddy-current-induced fields decay with time constants > 5 ms were required to be compensated. Field maps of the phantom were repeatedly measured with exponential increase delay time from 5 ms to 2.5 s after shim switching. The nonuniform temporal sampling of eddy fields allow measurements of both the medium- and long-living eddy current effects in a more time efficient way.

In order to correct the field inhomogeneity from the phantom, one additional ground field map was acquired without any preceding shim pulse. The ground field map was then subtracted from subsequent field map measurements. To quantify the degree of field distortion induced by eddy currents, for each shim channel, a baseline field map was acquired with a static current input to the corresponding coil.

Characterisation of eddy-current-induced fields

In order to characterise the eddy-current-induced fields, the measured field evolution was decomposed to spherical harmonic components by least-squares optimisation. In each field map measurement, only a single slice was acquired. Thus, some degenerated spherical harmonic components were removed from the fitting. The reduced shim sets used for fitting the acquired 2D-field maps at different slice acquisition orientation for specific shims are specified in Table 2.2. In addition, masks were created based on the magnitude image of the field map measurements using BET to exclude data in nonphantom regions from fitting [106].

The calculated coefficients of decomposed spherical harmonics for each shim channel

Shim	Slice Orientation	Spherical Harmonic terms for Fitting
Z2	Coronal	Const, X, Z, Z2, ZX
ZX	Coronal	Const, X,Z, Z2, ZX
ZY	Sagittal	Const, Y, Z, Z2, ZY
X2-Y2	Transversal	Const, X, Y, XY, X2-Y2
XY	Transversal	Const, X, Y, XY, X2-Y2

Table 2.2: Slice acquisition orientation and the corresponding nondegenerated spherical harmonics for measuring and characterising 2-D field distribution generated by pulsing specific second-order shims.

were normalised by the corresponding baseline field map. To describe the temporal variation of eddy current effects, relative amplitudes of the eddy-current-induced self-term spherical harmonic over 2.5 s, after switching off the ZX shim, are plotted in Figure 2.8.

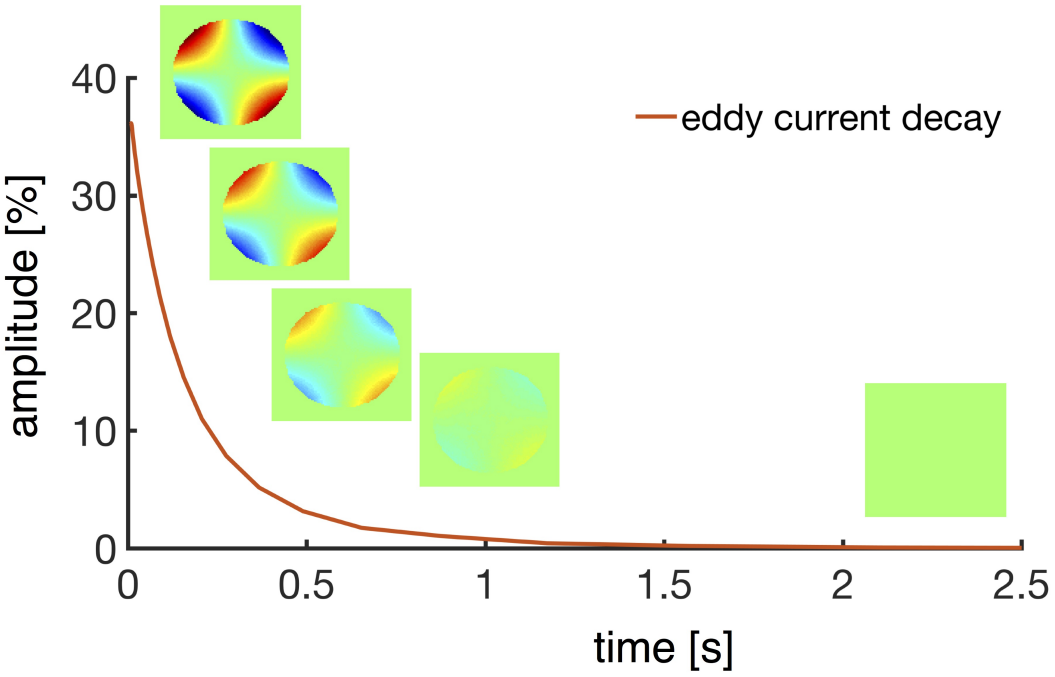


Figure 2.8: Relative amplitudes of the eddy-current-induced self-term spherical harmonic over 2.5 s after switching off the ZX shim.

Based on the exponential models of eddy currents as described in Equation 2.4, the eddy-current-induced fields were characterised by fitting a set of multi-exponential de-

cay functions using the Levenberg-Marquardt least-square algorithm [123]. To mitigate the influence of measurement noise and improve the accuracy of exponential regression, we used an iterative approach for eddy-current-induced field characterisation. In the first step, two or three sets of exponential terms, amplitudes and time constants, were used for initial characterisation. The corresponding pre-emphasis parameters were obtained based on Equation 2.10. The initial pre-emphasis results were used in the subsequent iterations to find an optimum fit by including additional exponential terms (up to 5 terms). These pre-emphasis parameter calibration procedures were repeated iteratively until residual eddy currents reach $< 1\%$ of the pulsed shim field strength.

The current change in the self-term shim channel will also affect the cross-term field responses. Thus, the cross-term ECC was conducted based on the self-term optimisation results. In the current hardware setup, cross-term ECC is only available to shim terms that share a same MPCU (see Figure 2.1). Taking the Z2 shim cross-term ECC for example, a series of 2D field maps in the coronal and the sagittal direction were acquired to measure the contributions of the ZX term. The measured field maps on the ZX cross-term were decomposed to the corresponding spherical harmonic components, as described in Table 2.2. Following that, the corresponding pre-emphasis settings for cross-term ECC were determined. After that, the procedures were repeated for the ZY cross-term ECC. Similarly, the procedures above were repeated for the rest of second-order shims.

Results and discussion

Figure 2.9 plots the measured shim field strength on each second-order channel after shim switching with (red line) and without self-term pre-emphasis (blue line). With-

out self-term pre-emphasis, pulsing the second-order shims induces different degrees of eddy currents in terms of the exponential decay amplitudes and time constants. The eddy-current-induced shim fields range from 20% to 40% of the input shim step and the settling time of eddy-current-induced fields was about 1.5 s long, on average. With self-term pre-emphasis, the strengths of eddy-current-induced fields were reduced to 1% or less of the pulsed shim. The self- and cross-term pre-emphasis parameters for all second-order shims are shown in Table 2.3.

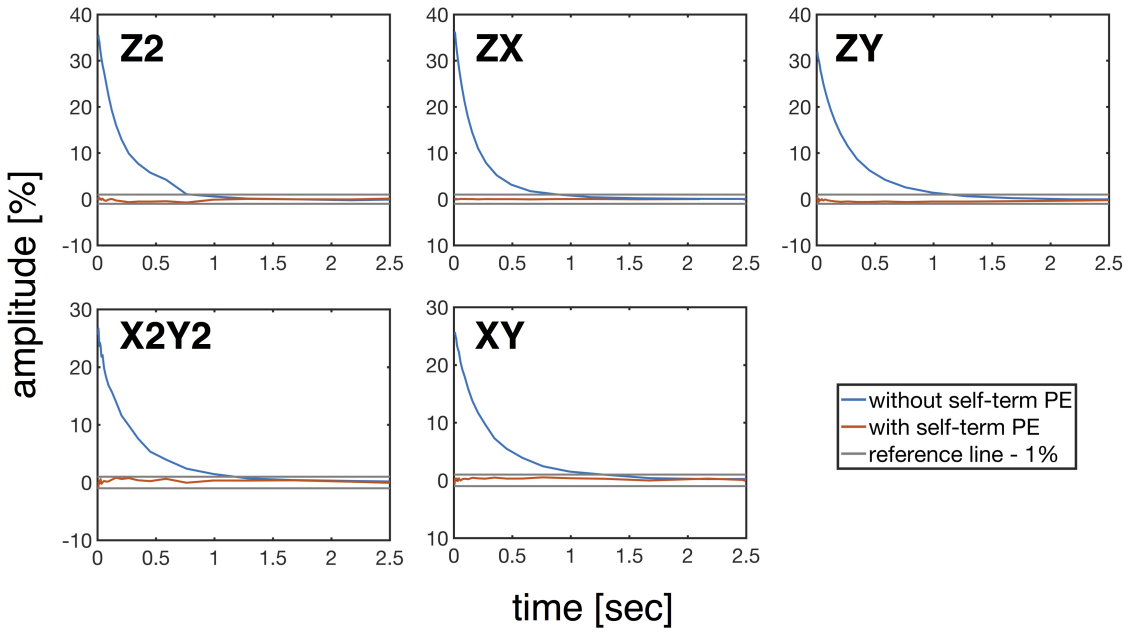


Figure 2.9: The measured self-term step responses after switching the Z2, ZX, ZY, X2-Y2 and XY shims without pre-emphasis (blue) and with the image-based self-term (red) pre-emphasis. Grey lines denotes 1% threshold of pulsed-shim amplitude.

2.3.3 SIRF-based eddy current measurement

This section introduces eddy current characterisation using a field camera. It starts by presenting the procedures of measuring the shim field dynamics using a field camera. Then, the focus will turn to introduce two methods for calculating pre-emphasis

Shim term	A_1 (%)	TC_1 (ms)	A_2 (%)	TC_2 (ms)	A_3 (%)	TC_3 (ms)	A_4 (%)	TC_4 (ms)	A_5 (%)	TC_5 (ms)
$Z2_{Z2}$	2.0	413	50.5	118	2.7	33	-	-	-	-
$B0_{Z2}$	36	1555	76.1	254	49.3	5	-	-	-	-
ZX_{Z2}	-	-	-	-	-	-	-	-	-	-
ZY_{Z2}	-0.25	186	-	-	-	-	-	-	-	-
ZX_{ZX}	7.1	458	31.7	126	-1.3	64	1.5	14	2.2	5
$B0_{ZX}$	-	-	-	-	-	-	-	-	-	-
$Z2_{ZX}$	-	-	-	-	-	-	-	-	-	-
ZY_{ZX}	-	-	-	-	-	-	-	-	-	-
ZY_{ZY}	4.8	553	35.0	135	2.5	34	6.3	4	-	-
$B0_{ZY}$	1.0	180	-2.4	6	-	-	-	-	-	-
$Z2_{ZY}$	-	-	-	-	-	-	-	-	-	-
ZX_{ZY}	-	-	-	-	-	-	-	-	-	-
$X2Y2_{X2Y2}$	4.0	667	-0.2	194	27.3	189	7.3	54	-	-
$B0_{X2Y2}$	-1.7	73	-	-	-	-	-	-	-	-
XY_{X2Y2}	-0.8	110	-	-	-	-	-	-	-	-
XY_{XY}	6.7	527	25.9	163	4.3	35	-0.3	10	-	-
$B0_{XY}$	-0.2,	359	-0.3	201	7.1	5	-	-	-	-
$X2Y2_{XY}$	0.4	3	-	-	-	-	-	-	-	-

Table 2.3: Pre-emphasis parameters for self and cross-term ECC on second-order shim channels using the image-based method. Pre-emphasis amplitudes are defined in percentage of input shim step and time constants are defined in the time scale of millisecond. "-" denotes the pre-emphasis parameters does not provide significant positive contributions to pre-emphasis results.

parameters based on the SIRF measurements. Finally, the self- and cross-term ECC results using the two pre-emphasis parameter determination methods on second-order shims will be compared.

Configuration of the field camera

The dynamic field camera (Skope Magnetic Resonance Technologies, Zurich, Switzerland) is comprised of 16 NMR field probes placed uniformly on the surface of a sphere with a radius of 10 cm to monitor field dynamics at different locations in the scanner bore [117]. Figure 2.10 shows the customised field camera used in this work. Each field probe has its own receiver coils to detect FID signal under the exposure of the applied gradient and shim fields. The signal acquisition system is connected to a set of high-speed analogue-to-digital converters (ADC). The ADCs allow receiving the field variation information at a temporal resolution of 1 μs. In order to conduct field measurements during a prescribed time course, the field camera was synchronised with the timing of the sequences carried out on the second-order shims using transistor-transistor logics (TTL). This design enables quick and accurate measurements of eddy-current-induced system responses.

Shim impulse response functions and field measurement

The field camera can simultaneously measure local magnetic field in the 16 field probes. From another point of view, field camera provides 16 spatially distributed data points that contain information on local field distribution. The field distribution is then quantified by spherical harmonic decomposition of the 16 data points using

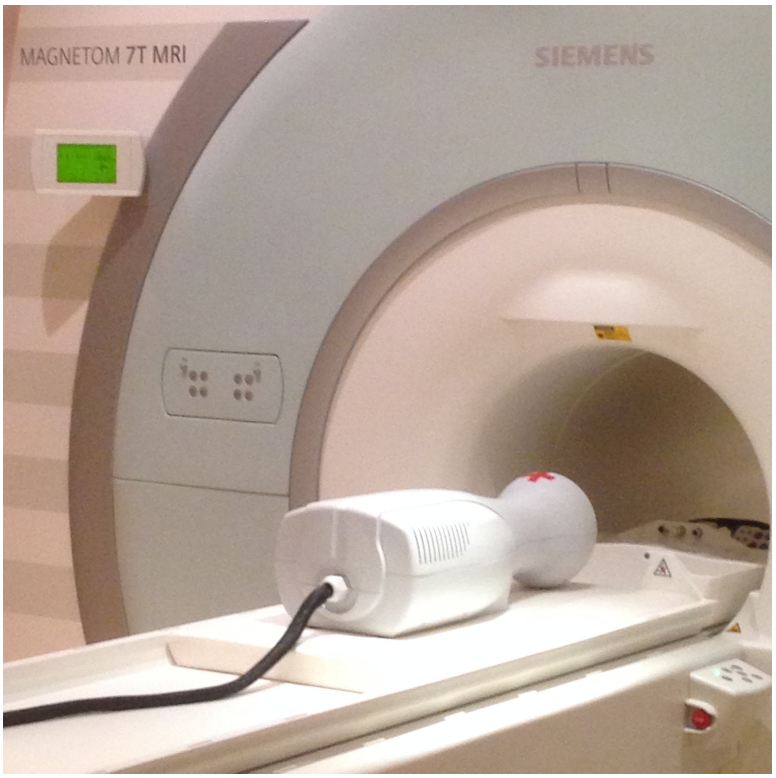


Figure 2.10: The photograph of the field camera used in the experiment

least-squares optimisation. In this way, we can obtain spherical harmonic representation of field distribution up to third orders.

In order to measure field distributions, the first step is to calibrate probe positions. In the implementation, probe positions were measured using local frequency offsets information under the influence of the X, Y, Z gradients. The static fields were retrieved from the calibration for background field correction.

Following the convention used in [120], the characteristics of shim channels can be described by the shim impulse response functions (SIRF). The SIRF is determined by dividing output responses to inputs in Laplace's domain. The output SIRFs of each shim channel were measured using two types of frequency-modulated input shim pulse waveforms:

- A 10 s long linear frequency sweep of 0-10 kHz. It is capable of capturing a wide

band of frequencies responses including both the low-frequency contents (e.g. long-living eddy currents), and the high-frequency contents (e.g. oscillatory vibrations).

- A 10 s long cubic frequency sweep of 0-30 Hz. It provides better characterisation of the low-frequency shim system responses, which mainly cover the medium (time constant 2 ms $\sim$ 100 ms) and long-living eddy currents (time constant $>$ 100 ms).

The field probes used in field camera can measure signal over a duration of 10 ms. The frequency sweeps a longer duration than that field probe can measure. To measure the field evolution spanning the whole time courses of the frequency sweeps, the field probes were excited to measure signal multiple times with a time interval of 100ms. To fill the gaps between each measurement, 10 repetitions of the shim waveforms are conducted with an incremental delay time of 10 ms.

The output SIRFs on Z2 shim term, while sweeping the Z2 shim with the frequency of 0-10 kHz is shown in Figure 2.11a. In the SIRF measurement, it is evident that there is a low-bandwidth peak centred at zero frequency, which corresponds to impulse responses of medium- and long-living eddy currents. The high-frequency components are associated with oscillatory responses observed in a short period after shim switching (Figure 2.11b), which are attributed to system mechanisms, like the amplifier behaviour.

Pre-emphasis parameter determination

The pre-emphasis parameters for self-term and cross-term ECC can be calculated based on the SIRFs measurements of each shim channel. For self-term pre-emphasis,

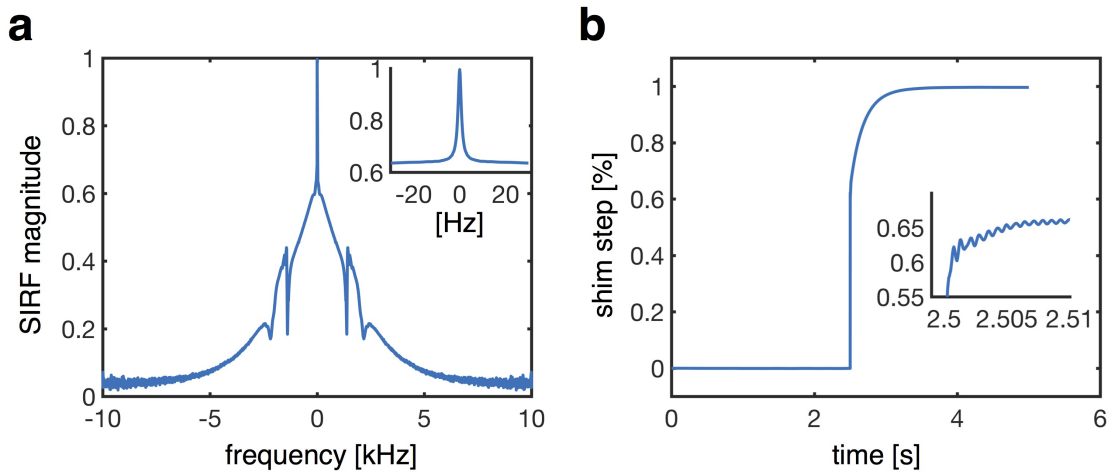


Figure 2.11: The plot (a) shows the impulse responses of the Z2 shim to a 0-10kHz frequency sweeps input waveform without pre-emphasis. The central peak in SIRF spectrum in the zoomed graph (a) corresponds to long-living eddy currents. The plot (b) represents the corresponding step responses of the Z2 shim. The zoomed plot in (b) represents step responses in 10 ms.

we used two strategies to determine the pre-emphasis parameters.

- A time-domain method. Transform SIRFs to time domain with an inverse Laplace transform and fit exponential functions to the transformed SIRFs. The corresponding pre-emphasis parameters were then determined according to Equation 2.10.

- A frequency-domain method [122]. The Fourier Transform of an exponential function is a Lorentzian curve centred at zero frequency. The pre-emphasis parameters were then obtained by fitting the Lorentzian curve to the frequency domain SIRF measurements. The transfer functions of pre-emphasis filters were calculated by the inverse of SIRFs in the frequency domain.

The information on medium- and long-living eddy-current-induced fields is contained in the SIRF measurements within the bandwidth of 30Hz. Thus, pre-emphasis parameter determination was based on the SIRF responses in the 30Hz bandwidth.

Previously studies reported that three sets of pre-emphasis parameters are robust enough for ECC purposes [108, 124]. Thus, we fit the SIRF measurements using three exponential or Lorentzian functions with least-squares optimisation.

For cross-term ECC, we calculated the cross-term pre-emphasis parameters based on the SIRF measurements with self-term ECC. It is because self-term pre-emphasis currents can generate secondary system responses on the cross-terms, on top of input shim currents. This method allows correction cross-terms arising from self-term pre-emphasis currents and input shim currents.

Our shim hardware setup only allows the cross-term ECC between the shim coils that share the same MPCU. Thus, the cross-term compensation was performed between the Z2, ZX, and ZY shims in MPCU1 and the X2Y2, XY shims in MPCU2. In addition to the second-order cross-term correction, the eddy-current-induced B_0 cross-term field is corrected by modifying the frequency waveform generated by the transmitter and receiver amplifiers.

SIRF-based self-term ECC

Figure 2.12 shows the ECC results without and with pre-emphasis using the SIRF-based frequency- and time-domain methods. Without pre-emphasis, as shown in Figure 2.12c, the SIRF measurement the strengths of eddy-current-induced self-term fields after switching the Z2 shim reached up to 35% of the shim step. The eddy current effects then decayed away within 3 seconds.

With self-term pre-emphasis, as shown in Figure 2.12a and 2.12b, the central peak of the SIRFs of the switched shim was eliminated and a plateau response exhibited within the 30Hz bandwidth. In the time domain, as shown in Figure 2.12c and 2.12d,

the difference between the desired and the measured self-term contribution of the Z2 shim was reduced down to 1% of the shim step around 10 ms settling time. Generally speaking, the time-domain and frequency-domain methods showed comparable performance. The frequency-domain method performed slightly better than the time-domain method in compensating long-living eddy-current-induced fields. The ECC results for the rest four second-order shim channels will be introduced in the following section.

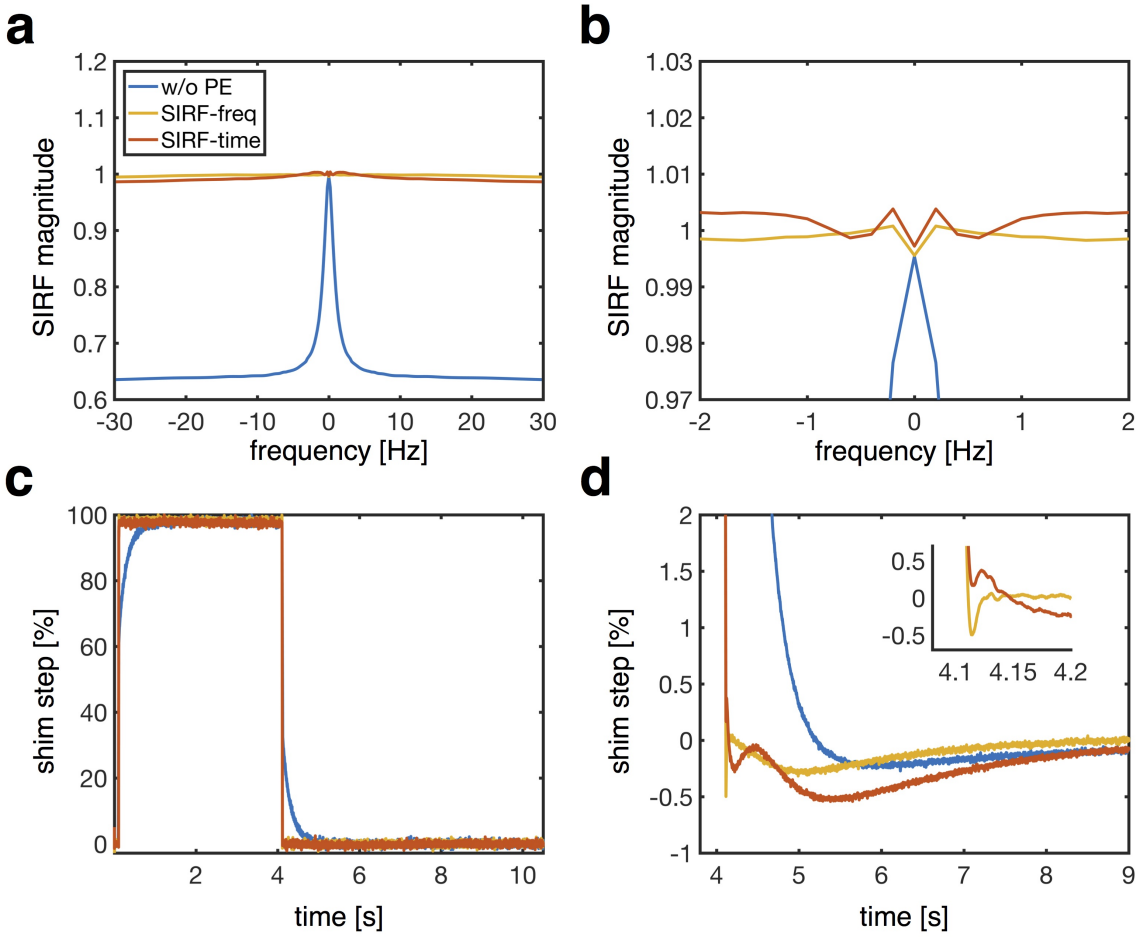


Figure 2.12: SIRFs (**a, b**) and field evolution (**c, d**) on the Z2 shim term, while switching the Z2 shim, without (w/o) pre-emphasis (blue), with self-term pre-emphasis using the frequency-domain (yellow), and the time-domain (red) methods. (**b, d**) show the corresponding zoomed plots of the SIRFs (**a**) and the field evolution (**c**) measurements, respectively.

SIRF-based cross-term ECC

Figure 2.13 shows the measured SIRFs of the Z2 shim in the self-term and the cross-terms B_0, ZX, and XYZ without (w/o) pre-emphasis (1), with the frequency domain self-term pre-emphasis (2), and frequency domain self- and cross-term pre-emphasis (3). Figure 2.13a shows, without pre-emphasis, strong central peak exhibited in the self-term spherical harmonic component. The time-dependent field variations on both the self-terms and the cross-terms settled down within two seconds. Furthermore, the amplitude of eddy-current-induced first- and second-order spherical harmonic cross-term responses were not significant compared to the self-term, whereas the third-order cross-terms show relatively strong eddy currents. The strongest third-order eddy-current-induced shim fields peak at about 0.5% of pulsed shim right after shim switching. Both the self-term and the cross-term pre-emphasis reduced medium-and long-living eddy currents. In this regard, problems of eddy currents can be offset in a large part by compensating for self-term eddy current.

Self-term pre-emphasis affects the cross-term responses. An overshoot of eddy current field is visible after the self-term pre-emphasis. The cross-term pre-emphasis reduces cross-term responses in the compensated shim channels. Taking cross-term pre-emphasis from ZX to Z2 for example, as shown in Figure 2.13b, the maximum amplitude of eddy-current-induced field deviation in the ZX cross-term was reduced from 0.2% to 0%. The respective field variations in the ZY cross-term decreased from 0.1% to 0%. Figure 2.13b suggests that the cross-term pre-emphasis currents gave rise to eddy current effects on the shim terms (e.g. the XYZ cross-term).

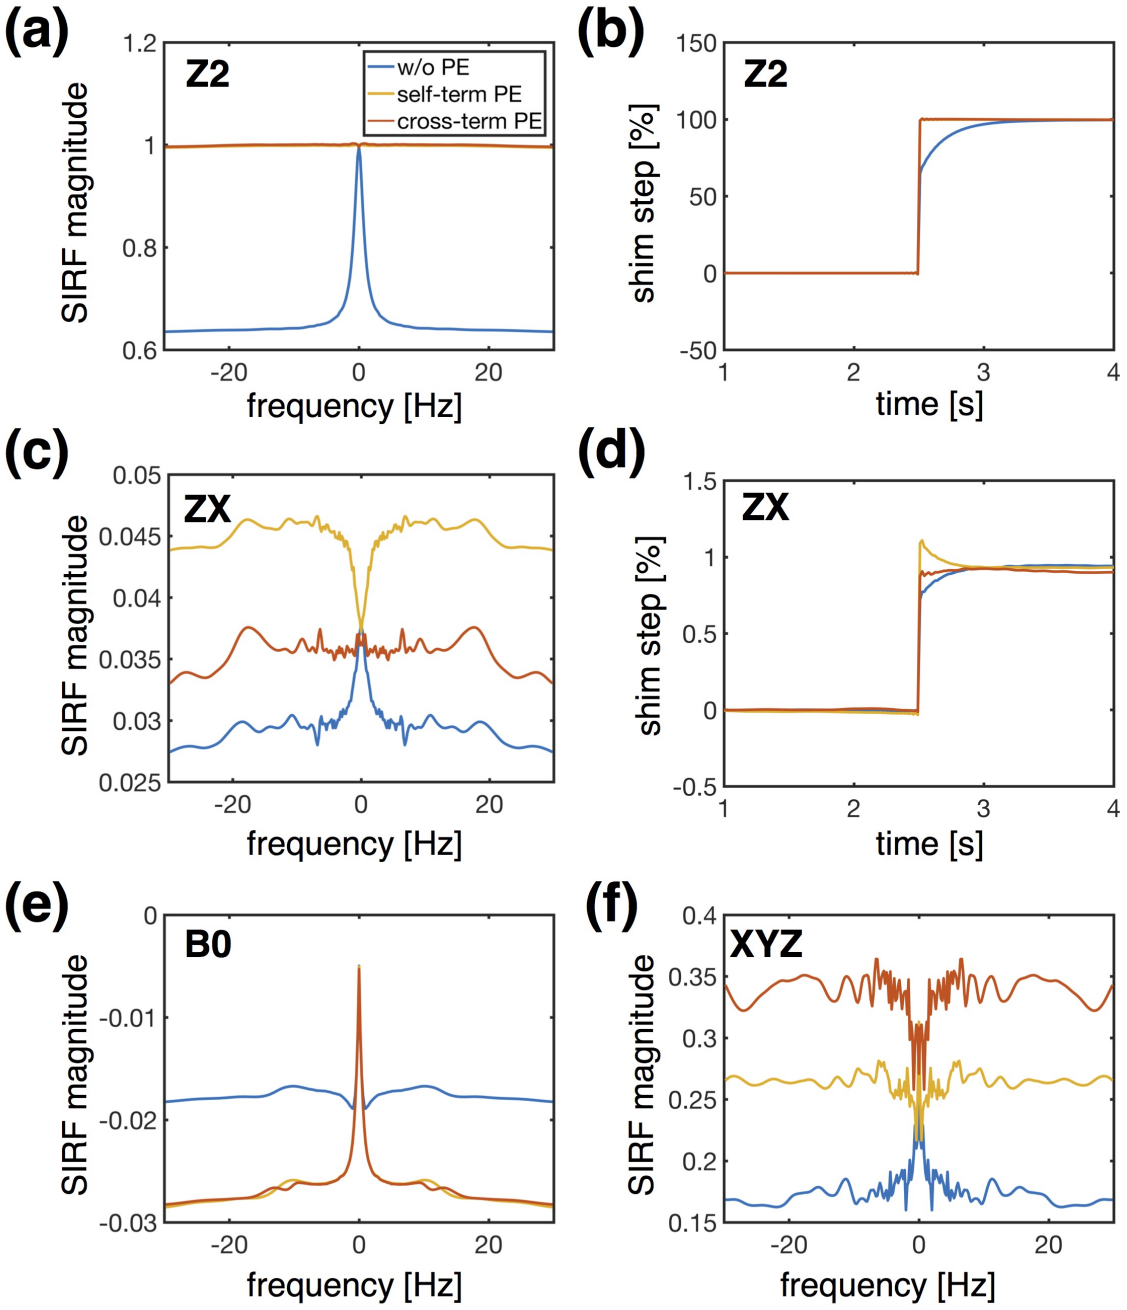


Figure 2.13: **a-d**: SIRFs (**a, c**) and field evolution (**b, d**) on the Z2 (**a, b**) and ZX (**c, d**) shim terms, while switching the Z2 shim, without (w/o) pre-emphasis (blue), with the frequency-domain self-term (yellow) and the self- and cross-term (red) pre-emphasis. The amplitudes of field evolution (**b, d**) are defined as the maximum field offsets on the surface of a 10-cm-radius sphere (normalised by the maximum field offsets on the self-term component). **e-f**: SIRFs of the B_0 (**e**) and XYZ (**f**) cross-terms.

2.3.4 SIRF-based method vs. image-based method

Figure 2.14 compares the performance of self-term pre-emphasis using the image-based and SIRF-based methods. All these methods provided similar results where the self-term eddy currents were reduced down to 1% of the inputs about 10 ms settling time after shim switching. The zoomed plot of the SIRFs in Figure 2.14b shows that the image-based method delivered more residual SIRF magnitude, compared to the SIRF-based methods. Figure 2.14d shows the field evolution after switching the Z2 shim. Pre-emphasis using the SIRF-based frequency-domain method delivered lowest eddy-current-induced field variations. The image-based method delivered 0.2% and 0.4% more maximum self-term field deviation than the SIRF-based time-domain and frequency-domain methods, respectively. In this regard, pre-emphasis using SIRF-based methods is slightly better than the image-based method. In terms of the rest of second-order shim channels, similar pre-emphasis results using the three eddy current characterisation methods were shown in Figure 2.15. The corresponding pre-emphasis parameters are shown in Table 2.4.

2.3.5 Discussion

We have compared the performance of pre-emphasis using the image-based and the SIRF-based methods. The SIRF measurement results shown here demonstrated that the image-based pre-emphasis can deliver comparable ECC results as the SIRF-based methods.

The reasons for the comparable results delivered using the image-based method and the SIRF-based methods are as follows. First, the dominant field errors are induced by medium- and long-living eddy currents, see Table 2.3. Simple field maps of the phan-

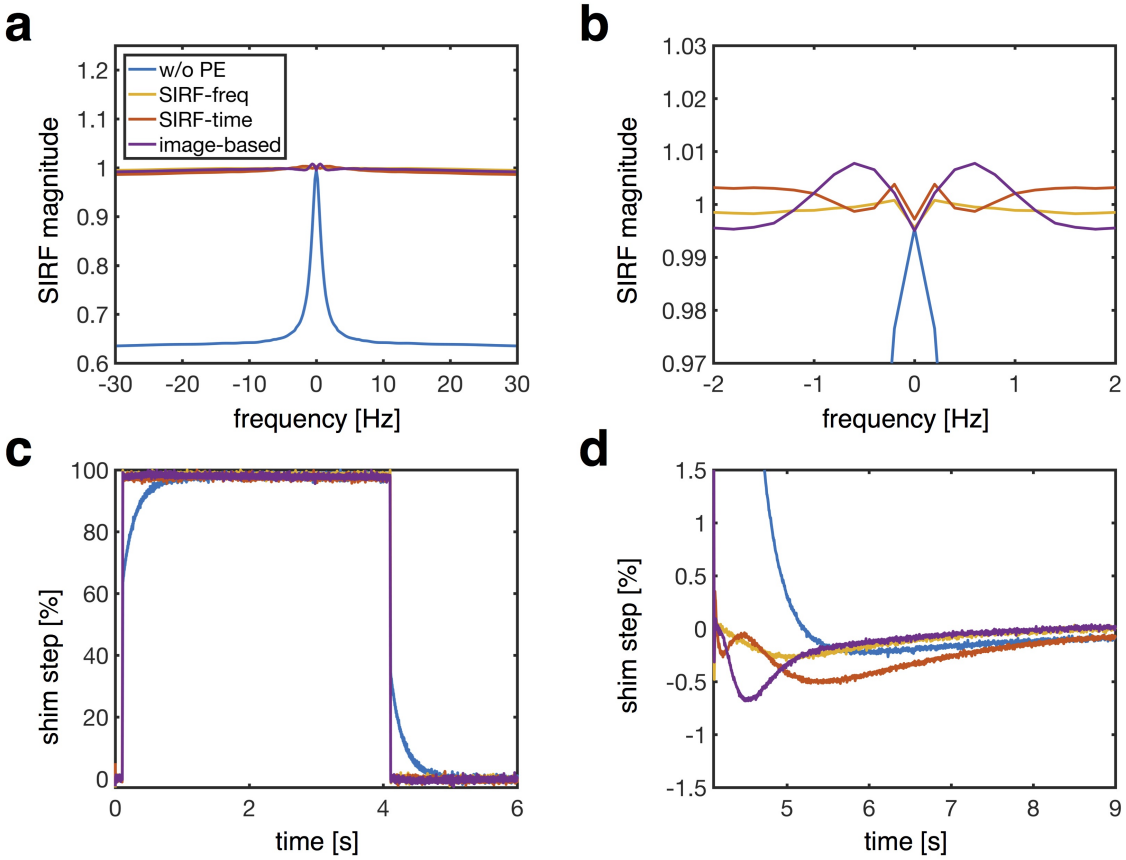


Figure 2.14: SIRFs (**a, b**) and field evolution (**c, d**) on the Z2 shim term, while switching the Z2 shim, without (w/o) pre-emphasis (blue), with self-term pre-emphasis using the frequency-domain (yellow), the time-domain (red), and the image-based (purple) methods. (**b, d**) show the corresponding zoomed plots of the SIRFs (**a**) and the field evolution (**c**) measurements, respectively.

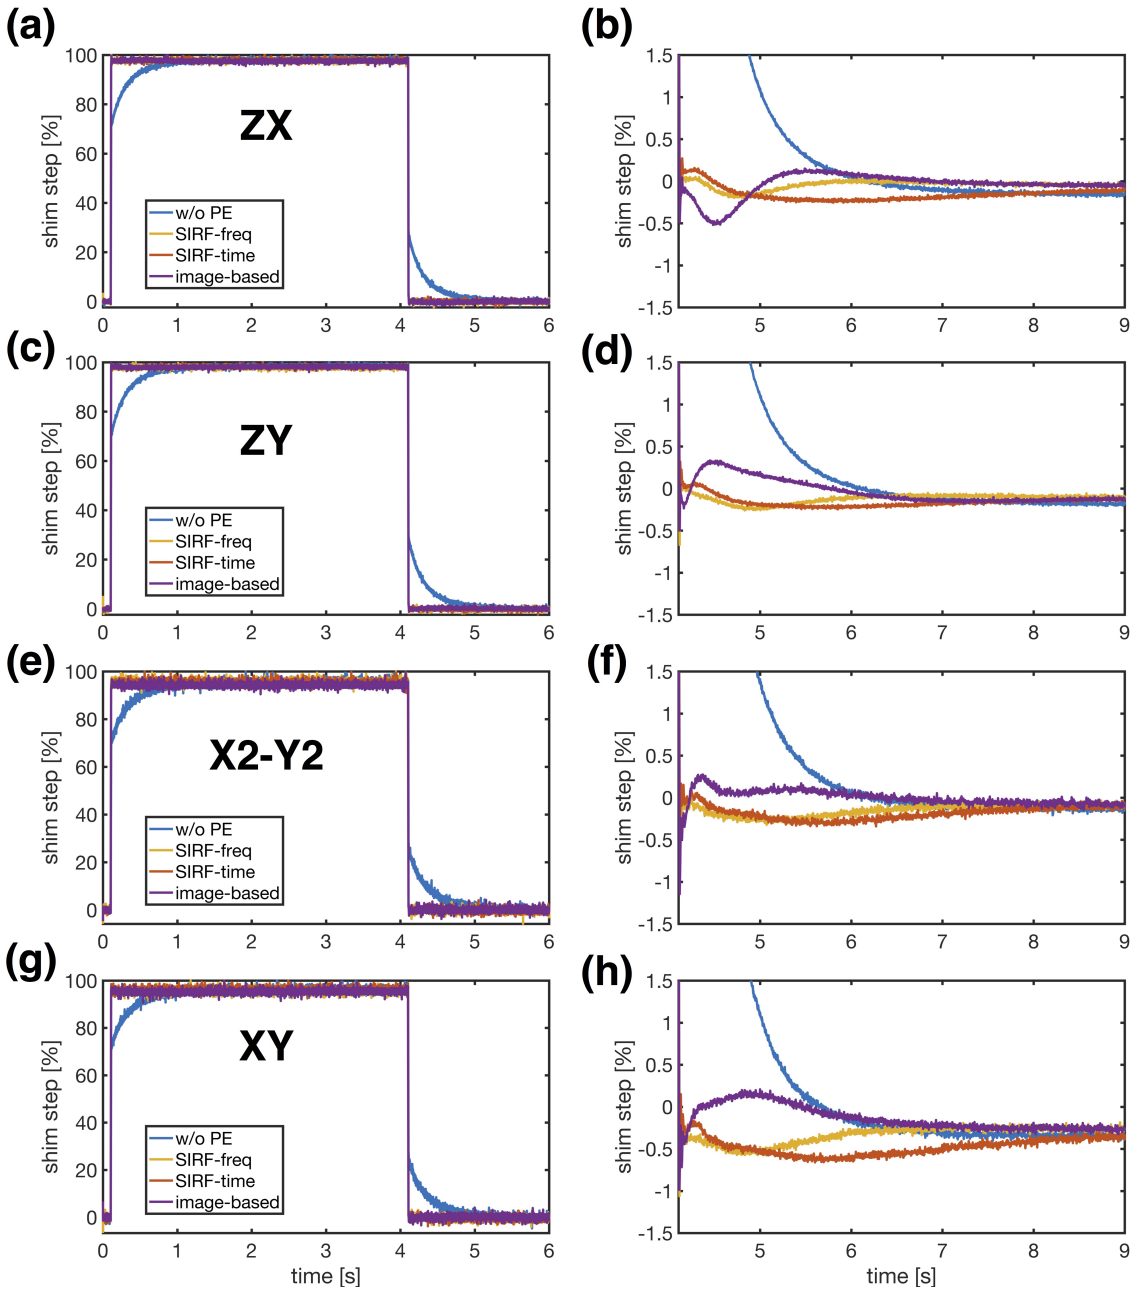


Figure 2.15: Field evolution on the ZX (**a**), ZY (**b**), X2Y2 (**c**), and XY (**d**) shim terms, while switching the corresponding shims without pre-emphasis (blue), with self-term pre-emphasis using the frequency-domain (yellow), the time-domain (red), and the image-based (purple) methods, respectively. (**b, d, f, h**) show the corresponding zoomed plots of the field evolution measurements (**a, c, e, f**), respectively.

Shim term	ECC method	A_1 (%)	TC_1 (ms)	A_2 (%)	TC_2 (ms)	A_3 (%)	TC_3 (ms)	A_4 (%)	TC_4 (ms)	A_5 (%)	TC_5 (ms)
$Z2$	SIRF-freq	2.8	413	51	118	2.7	33	-	-	-	-
	SIRF-time	0.9	2047	54.1	122	-	-	-	-	-	-
	Image-based	12.5	250	32.2	149	-111.7	125	39.7	112	83.3	110
ZX	SIRF-freq	7	458	31.7	126	1.5	14	-	-	-	-
	SIRF-time	1.2	1643	8.7	298	29.1	121	-	-	-	-
	Image-based	24.2	279	37.8	165	-83.9	161	61.2	126	-	-
ZY	SIRF-freq	5.3	553	2.5	34	34.6	135	-	-	-	-
	SIRF-time	1.4	1231	9.2	285	30.9	119	-	-	-	-
	Image-based	3	1107	-31.6	569	59.2	448	-26.8	353	38.2	132
$X2Y2$	SIRF-freq	4	667	27.3	189	7.3	54	-	-	-	-
	SIRF-time	1.8	1187	26	216	10.1	72	-	-	-	-
	Image-based	16.9	305	21	110	1.1	30	-	-	-	-
XY	SIRF-freq	6.7	527	25.9	163	4.3	35	-	-	-	-
	SIRF-time	1.6	1757	23.8	227	10.6	79	-	-	-	-
	Image-based	23.5	244	3.3	124	9.1	81	0.7	33	-	-

Table 2.4: Pre-emphasis parameters for self-term ECC only on second-order shim channels using the SIRF-freq, the SIRF-time, and the image-based methods (calibration results before gradient coil replacement). Pre-emphasis amplitudes are defined in percentage of input shim step and time constants are defined in the time scale of millisecond. "-" denotes the pre-emphasis parameters does not provide significant positive contributions to pre-emphasis results.

tom, are sufficient to capture the dominant time-varying field errors for ECC. Moreover, the image-based method has the advantage of providing relatively higher in-slice spatial resolution compared to the SIRF measurements. This attribute allows accurate characterisation of the contributions from self- and cross-term responses. Second, hardware constraints affect the performance of pre-emphasis using the SIRF measurements. Besides the medium- and long-living eddy currents, the SIRF measurements provides additional field variation information on system responses of short-living eddy currents and cross-term system responses up to third orders. However, there are limited number of pre-emphasis parameters available to interpret complex shim system responses. Some system responses in a very short timeframe measured with SIRF methods are not able to be compensated. Additionally, in terms of cross-term compensation, in the current system setup, cross-term ECC are only allowed to be performed on shims within shared MPCUs (see **Section 2.1**). Therefore, the third-order eddy-current-induced system responses, see Figure 2.13, were not be able to

compensated. Thus, the SIRF measurements only yield slightly better performance than the image-based method. To take advantages of SIRF measurements, inclusion of higher-order cross-term compensation and more pre-emphasis parameters is necessary.

There are some challenges that hinder the performance of the image-based method. First, the GRE sequence is sensitive to the field inhomogeneity. The time-varying eddy-current-induced fields will introduce different degree of distortion to the field maps measured at different echo times and further lead to errors in the eddy current characterisation. Second, the image-based method has limited temporal resolution. Some time-varying field features are lost in the image-based measurement. Thus, the image-based method is not capable of achieving better temporal responses compared to the SIRF methods. Third, the image-based method is time-consuming. The time used for compensating self-term eddy currents on a single shim channel is ≈6 h. The cross-term ECC requires repeating the field measurements for the other two orthogonal planes. Thus, pre-emphasis using the image-based method requires long scan time for the full characterisation of eddy-current-induced shim fields. Even though, in practice, if a field camera is not available, the image-based method is considered as a good alternative for eddy-current-induced field characterisation on any system.

2.4 Preliminary implementation of dynamic shimming

Section 2.2 and **Section 2.3** presents methods for calibrating shim fields and ECC. The demonstration of ECC effects on second-order shims is the first milestone for

dynamic shimming. This section builds on the hardware calibration results and investigates the performance of slice-wise second-order dynamic shimming. **Section 2.4.1** presents explicit details on procedures for the implementation of dynamic shimming. **Section 2.4.2** experimentally evaluates the performance of slice-wise dynamic shimming. Simulation of dynamic shimming on field maps acquired from 143 subjects is presented in **Section 2.4.3**. **Section 2.4.4** identifies the strengths and challenges of dynamic shimming.

2.4.1 Dynamic shimming implementation procedures

The implementation of the field-map-based slice-wise dynamic shimming consists of four main stages: sequence development (1), B_0 field distribution measurements (2), the ROI determination (3) and slice-wise shim calculation (4). The procedures are detailed as follows.

Sequence development

A GRE based field mapping sequence and an EPI sequence were modified to enable real-time shim corrections. The sequences running on the Master MPCU (MPCU0) conducted the standard image acquisition and the first-order shim corrections. The sequences running on the Slave MPCUs (MPCU1 and MPCU2) were synchronised to the sequences used on the Master and generated waveforms to control second-order shims. The schematic diagrams of the modified GRE-based field mapping sequences running on MPCU0 and MPCU1 are shown in Figure 2.16. The corresponding schematic diagrams of the modified EPI sequences are shown in Appendix B.1.

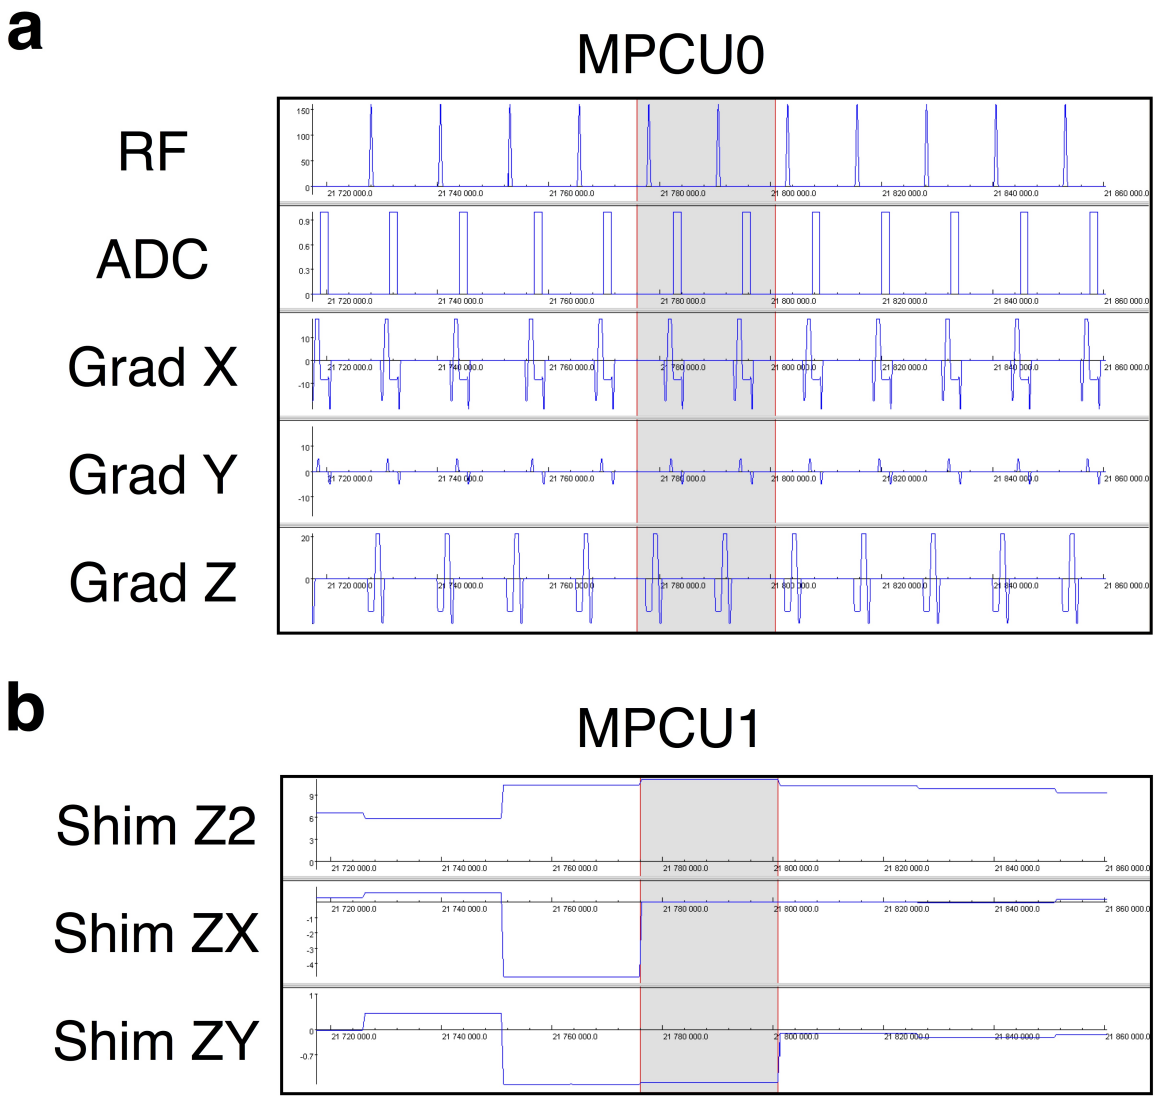


Figure 2.16: Schematic diagrams of the modified GRE-based field mapping sequences running on MPCU0 (a) and MPCU1 (b). The grey boxes indicate a GRE field mapping acquisition event block. It is visible that the shim waveforms running on MPCU1 is synchronised with the imaging protocol running on MPCU0

Each image acquisition block was preceded by a settling period for eddy currents as a consequence of shim switching. According to ECC results as described in **Section 2.3**, eddy currents were reduced down to 1% of the amplitude of the shim change after an initial settling time of 10 ms. In order to mitigate the eddy current effects to a great extent, prolonged settling time was desired. Therefore, the shim settling time was set to 20 ms throughout the experiment.

Field distribution measurements

A high spatial resolution field map was obtained for a subject to get adequate information on the field inhomogeneity for optimal slice-wise shim settings determination. In our implementation, high spatial resolution field maps were acquired with a modified asymmetrical dual gradient-echo sequence (voxel size=$2 \times 2 \times 2$ mm^3, TE1=4.08 ms, ΔTE=1.02 ms, slices=80, flip angle=39°, TR=640 ms, FOV=220×220 mm) with a default shim applied at the beginning of experiment to resolve the field distribution across the brain. Additionally, the modified field mapping sequence was customised to output a configuration file that contains the protocol set-up parameters, such as the rotation matrix of user-defined image volume, image position and orientation, FOV, slice thickness, and shimming volume. Following that, the acquired field maps were unwrapped using PRELUDE [78].

It is worth noting that the selection of echo time difference, ΔTE, affects the resulting field map measurements. Long ΔTE produces more field inhomogeneity induced phase change information. However, it leads to more phase wraps which pose challenges to phase unwrapping. Short ΔTE provides phase images that are largely free of phase wraps. However, images could be contaminated by additive noise. In the experiment, the echo time difference was set to 1.02 ms which is a trade-off between

the two scenarios and ensures that signal contributed from protons of water and fat go in phase at 7T [125].

ROI determination

The target of the experiment is to obtain optimal localised shimming for EPI. In practice, the field mapping calibration scans often have different protocol settings from the EPI acquisition [126]. In that case, we used an interactive slice-wise shim calculation method to obtain shim settings for dynamic shimming of the EPI based on a calibration field map. The procedures are as follows:

- Create a brain mask by using BET on the magnitude image from the field map acquisition and apply the mask on the corresponding field map to exclude non-brain tissue [106].
- Load a modified EPI sequence which is customised to output a configuration file containing the EPI protocol parameters. The generated configuration file then pass both the field mapping and the EPI protocol set-up files to a custom program.
- Reformat the acquired field map and transfer it to the space of EPI image based on their protocol files.
- Calculate shim settings based on the slice-by-slice target volume in the EPI.

Slice-wise shim calculation

Optimal shim currents can be obtained by fitting the magnetic fields generated by shim coils to the measured field map data in the ROI, as described in Chapter 1. In practice, many factors affect the performance of shim setting calculation for dynamic shimming. In the initial implementation, we specifically consider the two issues: (1) lack of field variation information in determining through-slice shims; (2) limited range of shim currents.

As discussed in **Chapter 1**, field variations in the through-slice plane are a major cause of signal loss. With the assumption that field variations are relatively smooth over the brain, neighbouring slices confer through-plane field variations information. To determine the through-slice shims, two neighbouring slices were included in shim calculations. However, this operation compromises in-slice field inhomogeneity information to some degree.

The input shim currents in the gradient/shim coils are subject to hardware constraints [127]. To address this problem, we extend the cost function of least-squares optimisation, as introduced in **Chapter 1 Equation 1.31**, by including a constraint term, which is described as

$$min\left\{\|\Delta B_{meas} - S \cdot I\|^2\right\} \ subject\ to\ \left(I_i\right)^2 \leq L_i^2 \tag{2.11}$$

where S indicates the shim model developed in **Section 2.2**, I indicates the currents required for shimming, ΔB_{meas} indicates the field inhomogeneity in the ROIs, L_i is the shim current limit, where i is the number of the shim channels. The results as shown in Table 2.4 suggest that the pre-emphasised shim currents can amount up to 150% of the inputs. Thus, the total input currents in shim coils may exceed the amplifier limitations. In order to prevent this from happening, the allowable second-

order shim strengths of manipulation were restricted within 3A, i.e. 30% of the shim hardware constraints (Table 2.5).

Shim	Z2	ZX	ZY	X2Y2	XY
function	$z^2 - (y^2 + x^2))/2$	$z * x$	$z * y$	$x^2 - y^2$	$x * y$
Sensitivity ($mT/m^2/mA$)	0.717±1.9e-3	0.731±1.8e-3	0.710±1.0e-3	0.355±5.0e-4	0.359±6.0e-4
Maximum (mT/m^2)	2151	2193	2130	1065	1077

Table 2.5: Shim coil sensitivities and maximum attainable shim strengths for second-order shims driven by RRI shim amplifiers.

The shim currents were calculated by fitting the shim model to the measured field map using the Matlab-embedded constrained least-squares method with "interior-point" optimisation, here termed the "LSQlin".

Central frequency offsets of each slice were calibrated aiming at minimising the in-plane field offsets. The calibrated centre frequency offsets in conjunction with the calculated shim settings were saved to a shim configuration files and loaded to sequences for shim controls. This step is looked at in more detail in **Chapter 3**.

The procedures of implementing dynamic shimming are summarised in Figure 2.17. The procedures begin with a stage of pre-scan where a localiser is acquired to get brain position information, and central frequency and transmission voltages are calibrated for the scanned subject. At the second stage, the protocol settings of targeted EPI sequences are recorded and initial calibration field map was acquired. At the following stage, optimal slice-wise shim settings are calculated and optimised for efficient hardware implementation. At the last stage, final calculated slice-wise shim settings are applied during multi-slice field map and EPI acquisitions.

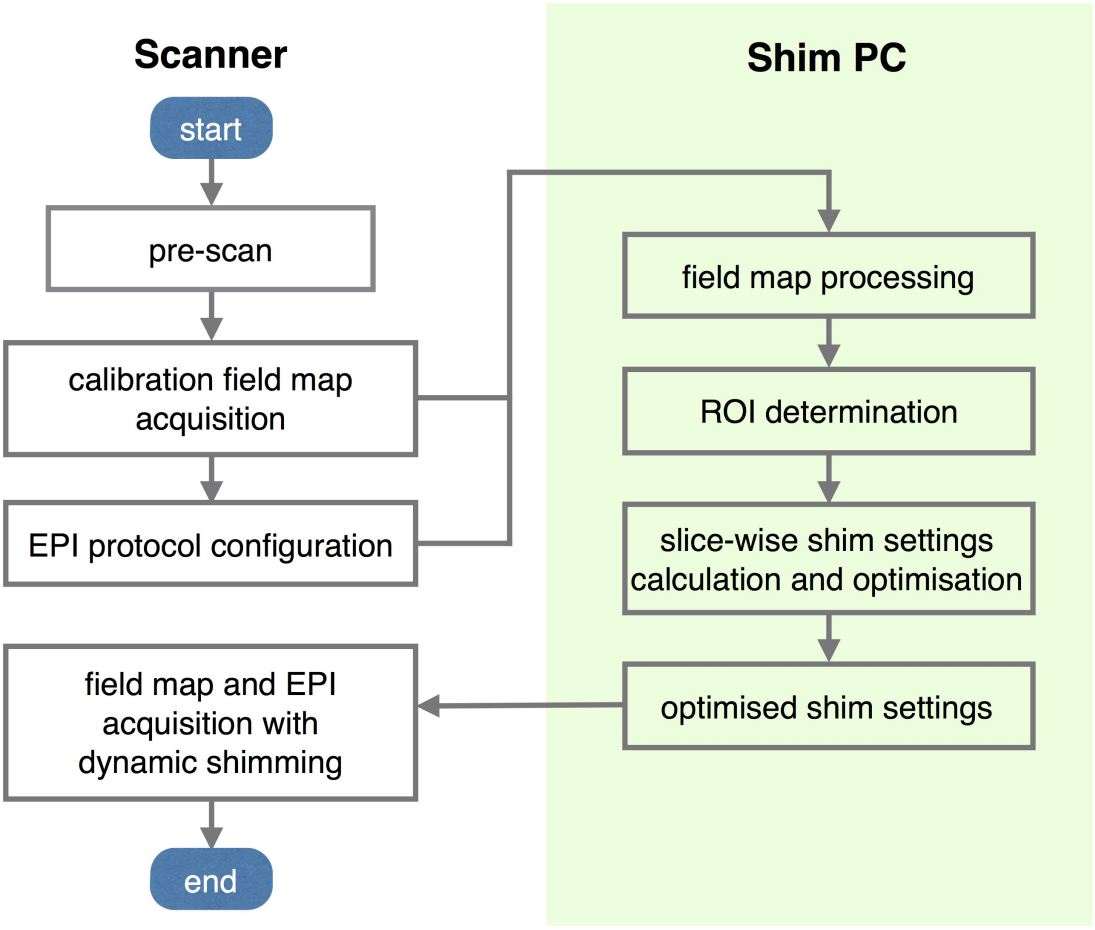


Figure 2.17: Illustration of steps to implement dynamic shimming.

2.4.2 Experimental validation

To demonstrate the performance of slice-wise dynamic shimming in realistic settings, experiments were conducted on one healthy volunteer. Shimming performance was quantitatively evaluated by the residual field standard deviation (STD) of a field map (voxel size=2 × 2 × 2 mm^3, TE1=4.08 ms, ΔTE=1.02 ms, slices=80, flip angle=39°, TR=640 ms, FOV=220 × 220 mm^2, acquisition time=4:24 min) acquired with the calculated slice-wise shim settings, and qualitatively evaluated by signal loss and distortion artefacts in the EPI images (voxel size=2 × 2 × 2 mm^3, slices=80, TR=4610 ms, TE=25 ms, FOV=220 × 220 mm^2, GRAPPA factor=2, echo spacing=0.72 ms). It is noted that the EPI volume was set the same as the B_0 field map. Such settings offer some convenience in the slice-wise shim settings calculation. Moreover, using the same resolution and slice locations in image acquisition allows direct visualisation and comparison of shimming effects in the EPI images. The field map post-processing and shim setting calculation was carried out in Matlab (R2017a) and FSL software [128].

Figure 2.18 shows the slice-wise first-order and second-order shims from the bottom to the top of the brain calculated based on a measured field map acquired from one subject. The shim settings vary smoothly across different slices for most regions of the brain. For slices in the inferior part of the brain, the variations in shim settings are relatively large due to the complexity of the brain structure. In addition, a steep increase of Z and Z2 shim strengths was observed in the most inferior and superior parts of the brain where the size of the ROIs is small which results from the severe shim interaction problems.

Figure 2.19 shows the standard deviation of the residual fields across different slices on the subject using different shimming methods. Field inhomogeneity was reduced substantially using both static shimming and second-order DS. With static shimming,

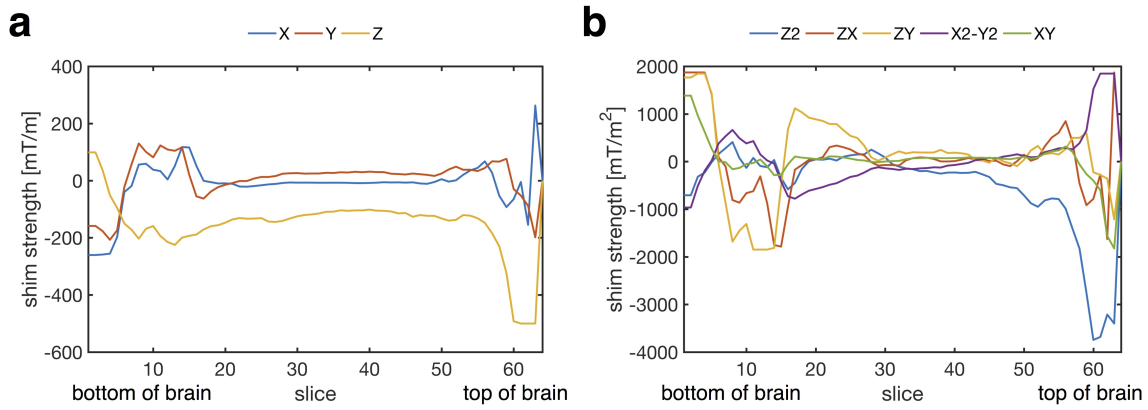


Figure 2.18: Slice-wise first-order (a) and second-order (b) shims from the bottom to the top of the brain calculated based on a measured field map acquired from one subject. The x-, y-, and z-axis corresponds to the left-right, anterior-posterior, and superior-inferior directions of the subject coordinates, respectively.

the improved overall shimming quality is traded at the expense of the deterioration of field homogeneity for slices located in the lower brain region. With second-order DS, the experimental results agreed with the simulation. The second-order DS provided better shimming quality than static shimming for all the slices. However, the high residual field standard deviation for slices close to the sinuses and ear cavities were still preserved after second-order DS. It is because the compensation fields generated by the second-order shims are relatively low in spatial order compared to the native field distribution of the brain.

Figure 2.20 shows the transversal views of B_0 field maps and EPI images acquired on the subject using different shimming methods. The EPI images acquired using dynamic shimming show gain in signal (e.g. the frontal lobe region at slice 20) and reduced distortion (e.g., the brain geometry at slice 10 and 50) in comparison to static shimming. In the B_0 field maps, in the regions of strong field distortion, the field deviation could exceed 200Hz even with slice-wise dynamic shimming. Additionally, it is visible from the slice 30, dynamic shimming reduced the field distortion in the inferior frontal lobe of the brain at the expense of increased field inhomogeneity in

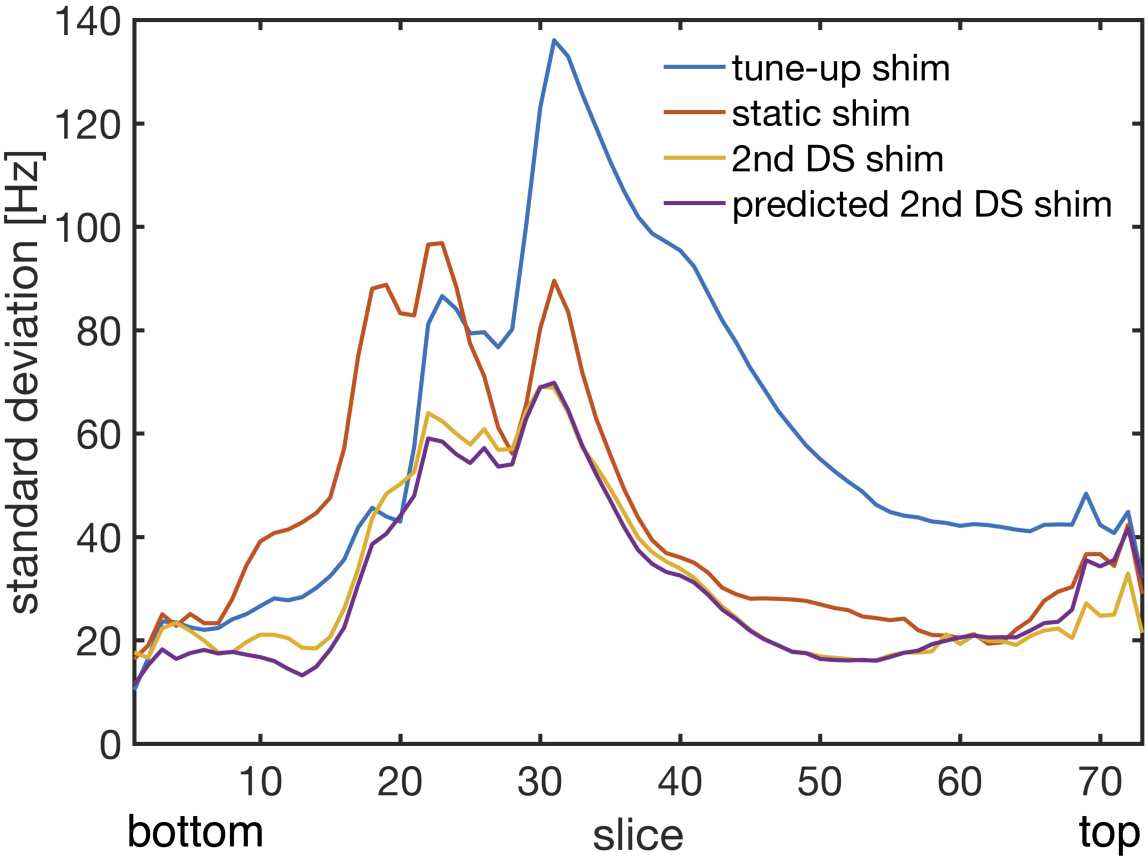


Figure 2.19: Comparison of simulated shimming performance using the tune-up shim (blue line), the static shim (red line), and the second-order dynamic shim (yellow line). In comparison, the predicted the second-order dynamic shim results (purple line) are plotted for reference. The x-axis shows the slice number from bottom to top. The y-axis indicates the standard deviation of residuals after shimming

the rest of regions, due to the limited spatial orders of the shim fields.

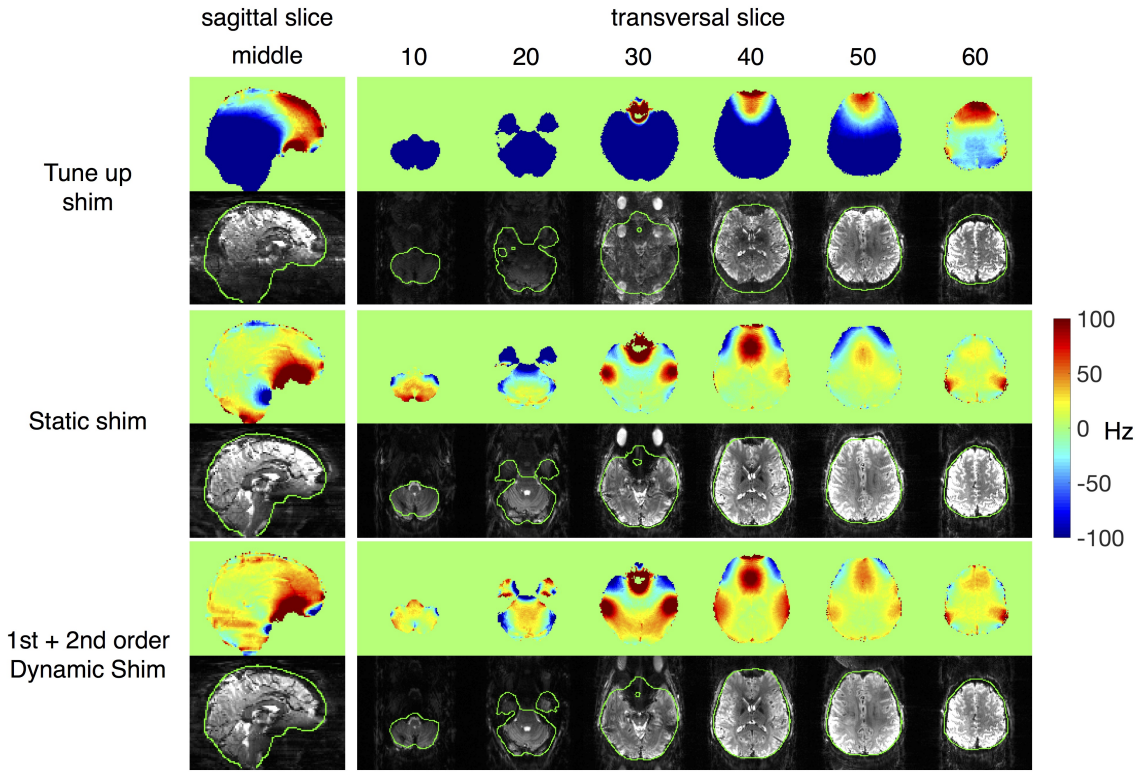


Figure 2.20: Field maps and EPI images acquired using different shimming methods in one subject. The columns from left to right indicate mid-sagittal images and transversal images at slice number 10, 20, 30, 40, 50 and 60. The rows from top to bottom indicate the shimming results using the tune-up shim (1), the static shim (2), and the second-order dynamic shim (3).

2.4.3 Simulation with a field map database

Field map database generation

Field maps used for providing group statistics of shimming results were acquired on the 7T scanner. A total of 176 healthy subjects (age range 20-66) were scanned as part of separate studies carried out by different researchers on the scanner. The scans were acquired in compliance with approved ethics protocols, which also allowed for

anonymous data to be shared for the purpose of this study. In positioning the subjects in the scanner, no special instruction in terms of head position or rotation angle was given. For each subject, the vendor-supplied routine shimming procedure was applied before any further scans. The shim settings were recorded in Digital Imaging and Communications in Medicine files (DICOM), enabling later removal of the baseline field. All study protocols acquired a B_0 field map for the purpose of removing image distortion in postprocessing, and a T1-weighted structural scan. The field maps were acquired with a 2-dimensional dual gradient-echo sequence (voxel size=2 $\times$ 2 $\times$ 2 mm^3, TE1=4.08 ms, ΔTE=1.02 ms, flip angle=39°, TR=620 ms, slices=64). The T1-weighted structural image was acquired with a 3D inversion recovery gradient echo sequence (TR=2200 ms, TE=2.82 ms, inversion time=1050 ms, voxel size=1 $\times$ 1 $\times$ 1 mm^3, GRAPPA (GeneRalized Autocalibrating Partial Parallel Acquisition) factor=4). To ensure that all field maps had sufficient brain coverage, field map scans with less than 90% brain content coverage, as calculated by the ratio of the brain contained in the field map to that in the T1-weighted image, were rejected.

Field map post-processing was carried out using Matlab (R2017a) and FSL software [128]. For each brain in the database, the field maps were unwrapped using PRELUDE [78]. The baseline shim field was removed using the recorded shim settings from the DICOM file and premeasured spatial field distributions of each individual shim coil, as described above. Each T1-weighted structural image and the magnitude image from the field map acquisition were brain extracted using BET [106].

Simulation results

The performance of slice-wise second-order dynamic shimming was evaluated in comparison to static shimming and the tune-up shim using the field map database de-

scribed above.

Figure 2.21 shows standard deviation of the residual field on all of the slices for each subject in the 143 field map database with different shimming methods. Static shimming reduced the residual field standard deviation from 160Hz to 50Hz and dynamic shimming further reduced it down to 41Hz. The simulation results confirmed the superiority of dynamic shimming over static shimming.

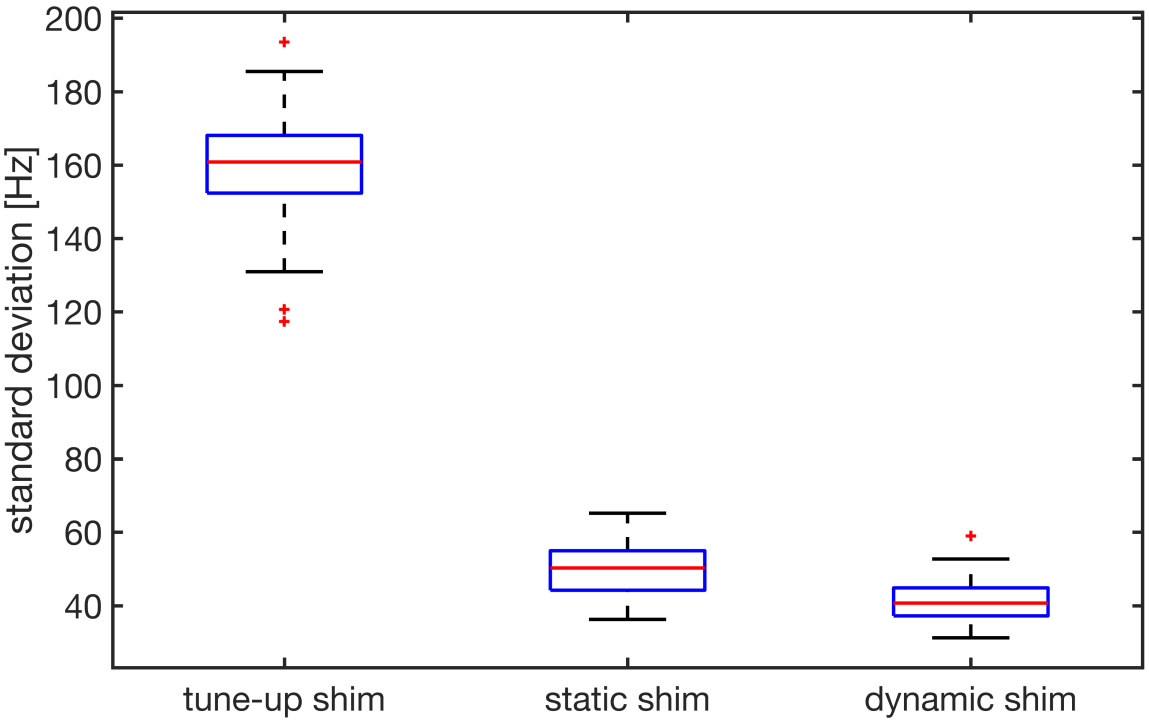


Figure 2.21: Simulated standard deviation of the residual field using the tune-up shim, the static shim, and the dynamic shim on all of the slices for each subject in the 143 field map database.

2.4.4 Discussion

Both simulations and experiments demonstrated that slice-wise dynamic shimming is capable of delivering better field homogeneity than static shimming. However, the experimental dynamic shimming results did not match the simulation results in the

inferior part of the brain. The slice-wise dynamic shimming was performed based on the assumption that there is no subject movement between field map calibration scan and dynamic shimming experiments. If there is movement between scans, the pre-calculated shim settings will not shim the targeted regions. This is especially problematic for slices associated with strong local field distortions, as head rotation and shift will significantly change the field inhomogeneity in the target regions. To address this issue, the optic tracker or quick localiser scan can be used to detect the movement. Incorporating the information in shim calculation could potentially yield better shimming quality.

The errors in dynamic shimming can also be attributed to inaccuracy in the shim calculation. With the reduction in shimming volumes, the effects of noise become more prominent. The calculated shimming solutions are becoming more sensitive to measurement and instrument errors. Thus, small changes in the measured field can lead to relatively large deviation in solutions. Furthermore, small ROIs result in serious adverse interaction between shim terms. In the case of shimming a thin off-centre slice, to correct for field variations along Z direction, fitting any shim terms containing Z components, like the Z2 shim, may lead to overfitting. According to Figure 2.18a and 2.18b, strong Z and Z2 shim interactions are observed at the most inferior and superior parts of the brain. The use of strong shims provides little improvement in shimming quality. Also, limited image intensity and spatial resolution might affect the results of shim calculation and induce a dramatic change of shim currents. **Chapter 3** aims to study different regularisation algorithms and apply them to optimise shim settings for second-order dynamic shimming.

Limitations of the shim hardware affect the performance of dynamic shimming. Shimming slices of severe field distortion sometimes requires high shim currents. Excessive current in shim channels causes heat problems and might exceed hardware constraints

of currents. Additionally, in the situation of high temporal resolution imaging, short eddy current settling period is demanded. In the case of fast switching of strong shim values, even with proper ECC, 1% shim field errors can lead to significant field distortions. As eddy currents last a few seconds, the resultant eddy-current-induced field will build up based on eddy field result from previous shim switching. Methods to mitigate the shim change over different slices are of vital importance. Regularisation algorithms can be used to reduce the shim strength on each channel and shape shim changes, while preserving satisfactory shimming quality.

With dynamic shimming, the field inhomogeneity located in slices near the sinuses is still dominant, see Figure 2.19. It is because second-order spherical harmonics are relatively low in spatial orders compared to the field distortion in these regions. As shown in Figure 2.20 slice 30, the reduction in standard deviation comes at the cost of field homogeneity in other regions. However, field inhomogeneity in the regions close to the sinuses are not significantly reduced and still cause signal loss in the corresponding EPI images. Therefore, the results suggest that the standard deviation optimisation does not necessarily deliver the best image quality. Some field homogeneity measures directly related to signal of interest can be used, for example, signal intensity and BOLD sensitivity. Targeting to optimise different homogeneity measures involves using nonlinear optimisation algorithms. In **Chapter 3**, various homogeneity measures will be explored.

Dynamic shimming requires a high-resolution field map to determine slice-wise shim settings. Acquiring a high-resolution field map, though delivering more information about the field distribution, takes a long scan time. Therefore, it is not convenient for some clinical studies. Methods to accelerate the process could be useful in the dynamic shimming process. In **Chapter 4**, a template-based method to accelerate field acquisition stage for dynamic shimming is proposed and evaluated.

Chapter 3

Dynamic shimming optimisation algorithms

The conventional linear least-squares based shim calculation method, as introduced in **Chapter 2**, assumes idealised performance of the shim coils which have high shim strength, high slew rates (a measure of how quickly shim can be switched), high linearity, and low heating effects. In practice, these coil characteristics are difficult to achieve. Errors in measurements can affect the solutions of the shim calculation and degrade the shimming results. In order to improve the robustness of shim current calculation to related errors, advanced optimisation algorithms are required.

This chapter aims to investigate various shim setting optimisation methods. Previous studies have demonstrated that incorporating regularisation in least-squares optimisation could improve the robustness and the stability of the shimming calculation [66]. **Section 3.1** explores the performance of slice-wise shimming using three typical regularisation methods. **Section 3.2** presents a novel and fully automated regularised shim optimisation technique accounting for hardware current constraints, as described by Equation 2.11. The proposed method is based on Tikhonov regularisation (Tikh) whereby the L-curve method is utilised to search for an optimal regularisation parameter. The performance of slice-wise dynamic shimming using the

proposed method was evaluated in comparison to some conventional methods on the field map database, introduced in **Chapter 2 Section 2.4.3**, in simulation and one subject in experiment.

Eddy currents induced by shim switching affect the performance of dynamic shimming. As described in **Chapter 2**, even with good ECC, the residual eddy currents can be significant if the size of shim steps, change of shim amplitudes between contiguous slices, on a second-order channel is large. As, eddy currents scale with the size of shim steps. A simple method to mitigate eddy current effects is to reduce the size of shim steps. **Section 3.3** presents two shim step size optimisation methods based on the Tikh and the constrained LSQlin methods to mitigate the large shim step sizes. The performance of the two methods is compared in simulation.

Besides minimising the standard deviation of the residual field, there are some other cost functions (e.g., minimisation of the gradients of the local field, or maximising image intensity and BOLD sensitivity) that make calculated shim settings better suited to various clinical applications [129]. **Section 3.4** compares the performance of dynamic shimming with different cost functions in simulation. The strengths and limitations of using different cost functions in slice-wise shimming are discussed.

3.1 Perspectives on current constraints and regularisation methods

Dynamic shimming has been shown to improve the B_0 field homogeneity in the brain in **Chapter 2 Section 2.4**. However, as the shim volumes get smaller, robust calculation of optimal shim settings gets harder. Partly, this is because shim currents required to homogenise the B_0 field in small volumes are more likely to run into

current constraints of the shim amplifiers. Smaller volumes also make the optimisation problem increasingly ill-conditioned, and thereby more susceptible to measurement noise.

Conventional shim optimisation techniques arrive at permissible shim settings with a constrained linear least-squares approach (LSQlin) [127]. The LSQlin searches for optimal shim settings through some iterative methods, for example the interior point and the trust-region reflective methods. With unsuitable solution search algorithms and limited number of iteration, the iterative methods are not guaranteed to converge and yield non-optimal shimming settings.

As introduced in Chapter 1, signal loss is largely related to field variations in the through-slice direction. The through-plane field variations usually have higher spatial orders than quadratic terms. To compensate the through-slice field inhomogeneity, high-order shims can be used. However, due to lack of information and underlying model on the field distribution in through-slice direction, fitting high-order terms may result in overfitting due to noise in the measured field maps, as illustrated in Figure 3.1. Additionally, in practice, shim terms are partly linearly dependent. Thus, using high-order shims sometimes makes limited improvement in field homogeneity.

To address the overfitting problem, a straightforward approach is leaving out some high-order terms which have relatively small contribution to field homogeneity, for example, the through-plane Z2 shim which has strong interaction with the Z shim [63, 18, 130]. Alternatively, to improve the robustness of fitting and mitigate the impact of measurement errors, regularisation can be used [131, 132, 133]. A study has demonstrated that least-squares optimisation using truncated singular value decomposition (TSVD) could improve the robustness of static shimming and concurrently address the hardware constraint problem in shim current calculation, as described by Equation 2.11 [66]. Besides the TSVD, other regularisation methods, like the Tikh and the

damped singular value decomposition (DSVD), can be used to optimise least-squares problems [134].

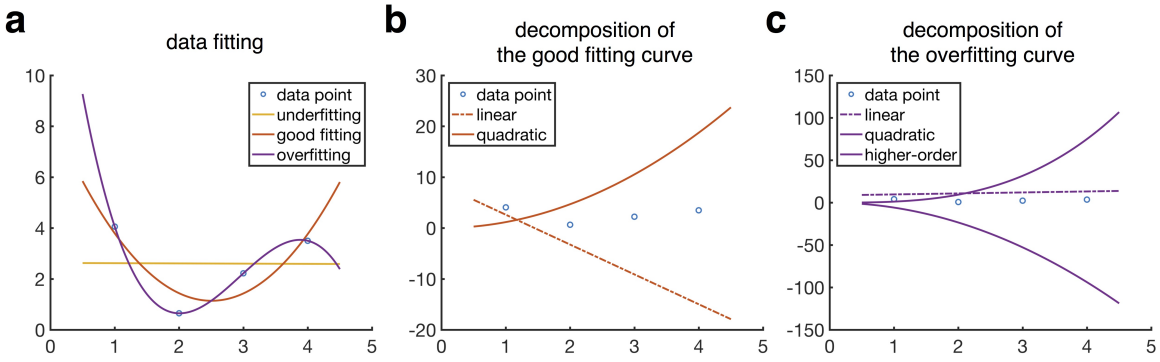


Figure 3.1: (a) Illustration of underfitting (yellow line), good fitting (red line), and overfitting (purple line); (b) decomposition of components used for good fitting; (c) decomposition of components used when the overfitting problem occurs.

This section investigates the performance of slice-wise shimming using different regularisation methods. **Section 3.1.1** is to introduce the theory of regularisation. In **Section 3.1.2**, their performance is comparatively studied based on the 143 field map database, as described in **Chapter 2, Section 2.4.3**, in simulation.

3.1.1 Theory of regularisation

This subsection starts by introducing the concepts of the singular value decomposition (SVD). Being acquainted with the theory of the SVD and its role in shimming helps to understand the concepts of regularised shimming methods. Following that, we will briefly introduce three widely used regularisation methods, the TSVD, the DSVD and the Tikh.

Singular values and condition number

The SVD is a useful mathematical tool in linear algebra that enables a matrix factorisation. The SVD of a shim matrix S is given by

$$S = U\Sigma V^T = \sum_{i=1}^{n} \sigma_i u_i v_i^T, \tag{3.1}$$

where n is the rank of the matrix S, an orthogonal matrix $U = [u_1,\ u_2, u_3,\ \cdots u_m] \in R^{m\times m}$, an orthogonal matrix $V = [v_1,\ v_2, v_3,\ \cdots v_n] \in R^{n\times n}$, a diagonal matrix $\Sigma = diag\,(\sigma_1,\ \sigma_2, \sigma_3,\ \cdots \sigma_n) \in R^{m\times n}$, σ_n is the singular value with $\sigma_1 > \sigma_2 > \sigma_3,\ \cdots \sigma_n > 0$

In practice, the measured field inhomogeneity is modelled by incorporating a noise term

$$\Delta B_{meas} = \Delta B_{exact} + \delta b, \tag{3.2}$$

where ΔB_{exact} denotes the intrinsic field inhomogeneity and δb denotes errors in the measured field map. Therefore, the shim current solutions, as introduced in **Chapter 2**, in terms of the SVD, can be extended to

$$I = S^+ \cdot \Delta B_{meas} = \sum_{i=1}^{n} \frac{v_i u_i^T}{\sigma_i} \left(\Delta B_{exact} + \delta b\right), \tag{3.3}$$

where S^+ denotes the pseudoinverse of the matrix S, which can be defined as a function of the reciprocal of singular values σ_i. The corresponding shim currents calculated by the pseudoinverse is inversely proportional to the singular values. Therefore, with small σ_i, S^+ has a large condition number in respect of inversion. The noise in the measured field map δb is amplified by S^+ and propagates to the final current solutions, futher yielding errors in the compensation field. The sensitivity of shim current solutions to measurement errors is called the conditioning of the least-squares problem. Conventionally, a condition number is used to indicate the conditioning of

the shim determination problem, which is defined by the ratio of the largest singular value to the smallest one, which is expressed by

$$cond\,(S) = \frac{\sigma_1}{\sigma_n}, \tag{3.4}$$

where σ_1 denotes the largest singular value of the shim matrix and σ_n denotes the corresponding smallest singular value. If the shim matrix has large condition number, the calculated current solutions are relatively more sensitive to the related noise or errors. This scenario is commonly referred to as ill-conditioned.

Regularisation effects

The conditioning of the shim matrix is related to the characteristics of singular values decay. A better conditioning of the least-squares problem can be achieved by increasing small singular values. The methods of introducing additional filters ψ on singular values to solve the ill-conditioned problem are called regularisation [131].

With regularisation, the calculated shim solution norm, an indication of the amount of current required for shimming, can be derived as

$$\|I_\lambda\|^2 = \sum_{i=1}^{n} \left(\psi^\lambda \frac{u_i^T \Delta B_{meas}}{\sigma_i} \right)^2. \tag{3.5}$$

The corresponding residual field norm, an indication of field inhomogeneity in the ROIs, can be expressed as

$$\|\Delta B_{meas} - S \cdot I_\lambda\|^2 = \sum_{i=1}^{n} \left(\left(1 - \psi^\lambda\right) u_i^T \Delta B_{meas} \right)^2. \tag{3.6}$$

The solution norm decreases with the regularisation, and the residual field norm increases. It will be discussed in the next section in different shimming cases.

The difference in shim currents between the noise-free current solution I_{exact} and the regularised ones I_λ is given by

$$I_{exact} - I_\lambda = \Delta I_\lambda + \Delta I_{pert}. \tag{3.7}$$

The first term ΔI_λ denotes the contribution of regularisation errors to current solutions, which is given by

$$\Delta I_\lambda = V\left(\mathbb{I} - \Phi\right) V^T I_{exact}, \tag{3.8}$$

where Φ represents a regularisation filter vector, $\mathbb{I}$ represents an identity matrix.

The second term ΔI_{pert} is the contribution of perturbation errors to current solutions, which is given by

$$\Delta I_{pert} = V\Phi\Sigma^{-1}U^T\delta b. \tag{3.9}$$

The regularisation results illustrate that incorporating regularisation filters in the singular value vectors can mitigate the perturbation errors at the expense of introducing some regularisation errors in the solutions. From the point of view of shimming, small regularisation parameters have subtle effects on the calculated shim currents and noise. Large regularisation parameters, though largely mitigating noise effects, cause overregularisation of shimming results - degradation of field homogeneity.

TSVD, DSVD and Tikh regularisation

Each singular value contains different amounts of information about the field distribution. Thus, the regularisation effects on solutions depend on the regularisation filter functions. The filter functions of the TSVD, the DSVD and the Tikh are shown as follows:

- The TSVD filter is designed to directly truncating small singular values in the shim matrix. The filter function is defined by

$$\psi_j^k = \begin{cases} 1, 1 \leq j \leq k \\ 0, k < j \leq l \end{cases} \tag{3.10}$$

where k denotes a truncation coefficient which indicates the number of retained singular values, l denotes the total number of singular values.

- The DSVD regularises the least-squares problem by incorporating an additional constant λ ($\lambda > 0$) to all of singular values to damp the rate of decay of the singular values. The filter function is given by

$$\psi_i^\lambda = \frac{\sigma_i}{\sigma_i + \lambda}. \tag{3.11}$$

This regularisation yields continuous regularisation results compared to the discrete characteristics provided by the TSVD.

- The Tikh method introduces quadratic weightings on singular values. The filter function in terms of the SVD is given by

$$\psi_i^\lambda = \frac{\sigma_i^2}{\sigma_i^2 + \lambda^2}. \tag{3.12}$$

Incorporating the Tikh filter in least-squares optimisation formulates an extra solution norm term into the cost function of the least-squares. Thus, least-squares optimisation using the Tikh method allows optimising the solution norm and the residual field norm simultaneously. The standard form of the Tikh-based least-squares will be explained in more detailed in **Section 3.21**.

3.1.2 Regularisation performance evaluation

Evaluation on example slices and a whole brain

To get a whole picture of the regularisation performance, simulations of the least-squares shim optimisation with the TSVD, the DSVD and the Tikh were performed in different shimming cases for a second-order shim system using a regularisation tool package [131]. Preliminary analysis was carried out to compare the regularisation effects in three different ROIs: (i) the whole brain, (ii) a slice located at the level of the frontal cortex (relatively large volume coverage) and (iii) a slice at the level of the brainstem (relatively small volume coverage). The regularisation results are quantified based on three measures:

1. The condition number of the shim matrix - an indication of the conditioning of least-square problems
2. The solution norm $\|I_\lambda\|^2$ - an indication of shim current strengths required for shimming
3. The residual B_0 field norm $\|\Delta B_{meas} - SI_\lambda\|^2$ - an indication of field inhomogeneity

The corresponding regularised shimming results are illustrated in Figure 3.2. With the increase of the regularisation parameters, both the Tikh and the DSVD method produced smooth changing regularised solutions, whereas the TSVD delivered discrete solutions due to the limited number of singular values available for truncation. As illustrated in Figure 3.2(a-c) and 3.2(g-i), shimming on the small ROIs was more ill-conditioned compared to the larger image volumes and more likely produced large

shim values. The condition number and the solution norm were drawn toward smaller values with the increase of the regularisation parameters. Figure 3.2(d-f) show that an increase of regularisation factors worsen the residual field norm.

The effects of the regularisation on the shimming results depend on the choice of regularisation factors, as shown in Figure 3.2(d-i). With small degree of regularisation, the increase of regularisation parameters caused a sharp decrease in the solution norm, without degrading much residual field norm. Follow that, with the increase of regularisation parameters, the current norm was reduced to a very low value at a relatively large cost of the residual field norm. Taking the Tikh on a slice located in the middle of the brain for example, with a regularisation parameter of 0.1, the solution norm and the condition number were reduced to 31% and 11% of the unregularised solution, respectively. However, the residual norm only increased from 7.4Hz to 7.7Hz. With further increase in the degree of the regularisation, the residual norm was degraded to 65Hz, corresponding to the results of no shim applied.

The calculated shim strengths applied on each channel for the slice located in the middle of the brain are plotted in Figure 3.3. The results have intuitive interpretation of the observations on the solution norm, as shown in Figure 3.2(g-i), that the increase in regularisation parameters delivered smoothly decreasing shim values on each shim term, compared to the discrete nature of the TSVD.

Evaluation of regularisation performance

In shimming, it is inefficient to apply large currents in shim coils to correct field inhomogeneity. The reasons are as follows:

- Large shim currents are likely to hit hardware constraints of currents. For

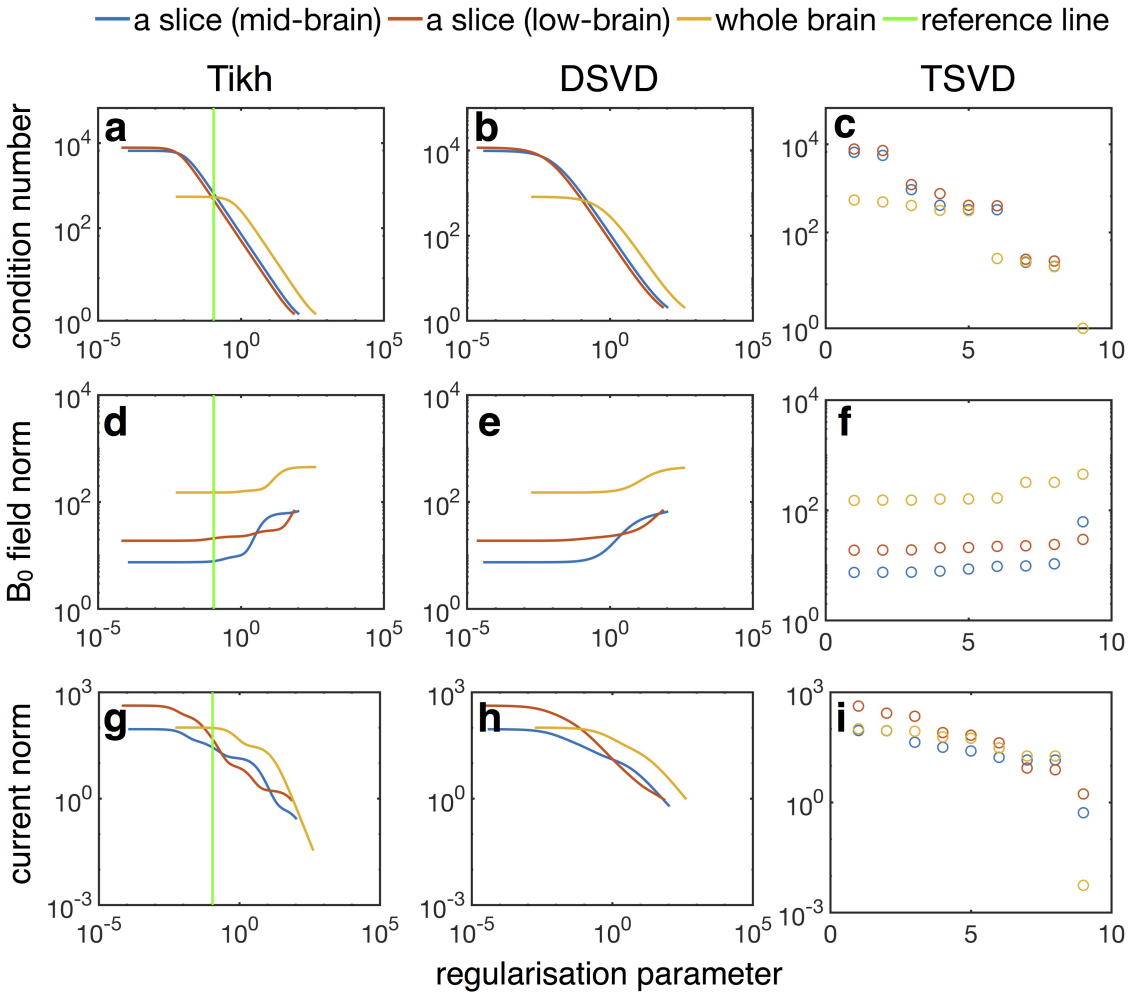


Figure 3.2: Comparison of regularisation effects using the TSVD, the DSVD and the Tikh methods for different ROIs, a whole brain (yellow line), an example slice in the middle part of the brain (blue line) and lower part of the brain (red line), respectively. (**a-c**), (**d-f**), (**g-i**) indicate the condition number of the shim matrices, residual B_0 field norm, and solution norm optimised with the three methods. The reference line (green) indicates the shimming results using the Tikh method, which provided a 7.4Hz residual field standard deviation for a slice located in the middle part of the brain

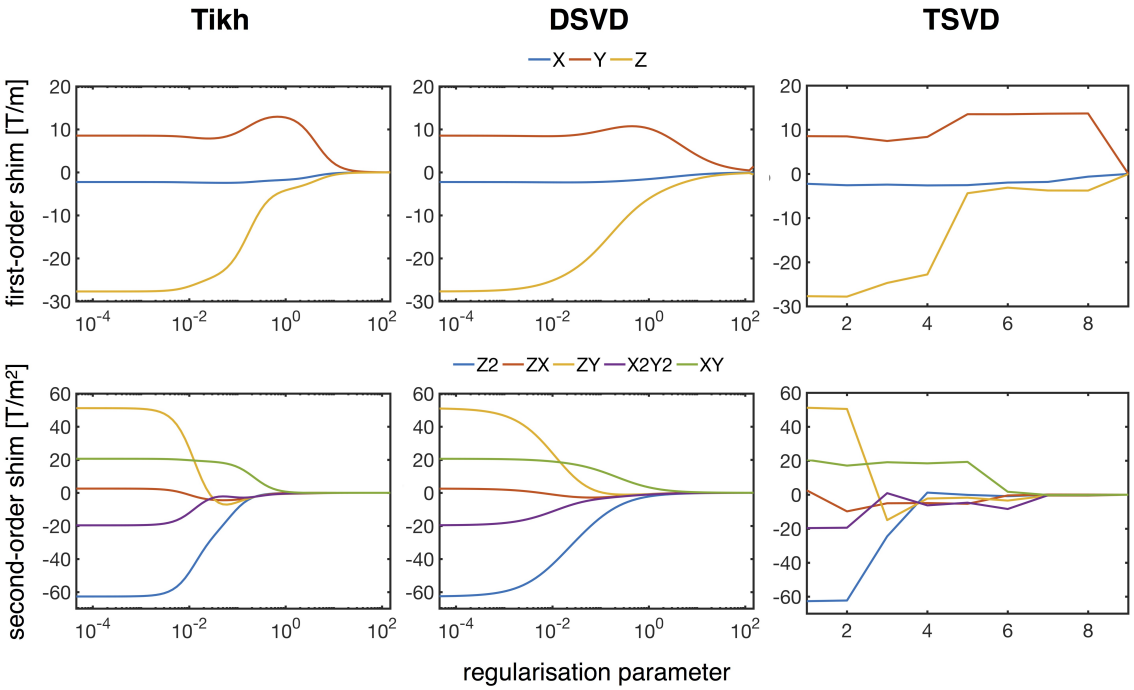


Figure 3.3: The amplitudes of first- (first row) and second-order (second row) shims using the Tikh, the TSVD, and the DSVD methods with different regularisation levels.

example, in our shim control hardware, the maximum allowable currents that can be used on the shim amplifiers are 10A and the summation of them are 20A. Driving high strength currents in all shim amplifiers can potentially damage the equipment. Moreover, inputting high currents in shim coils could cause heating problems.

- Using high shim currents does not certainly yield much better shimming results. As illustrated by regularisation results in Figure 3.2, 69% reduction in the shim solution norm only yielded a degradation of the residual field standard deviation of 0.3Hz, which is considered small for shimming applications.

- Errors in the field characterisation will be amplified when high shim strengths are used. According to the shim field characterisation results as shown in Chapter 2 Figure 2.4, the measured shim fields varied between different measure-

ments. The inaccuracy in shim field characterisation (e.g. limited orders of spherical harmonics), leads to errors in the shim field model. The errors scale with the strength of required currents.

- For dynamic shimming applications, rapid switching high shim currents leads to strong eddy current effects. The use of low shim currents could potentially reduce eddy currents.

- From the SVD perspectives, as noted in Equation 3.5, the high shim solution norm is associated with small singular values. Increasing singular values can reduce shim currents and lead to smaller condition numbers. Thus, small shim solution norm is also an indication of good conditioning of the shim calculation problems.

For all the reasons above, there is a desire to investigate methods to reduce shim current strengths in shim channels. As demonstrated in Figure 3.3, incorporating regularisation methods in the shim current calculation can reduce shim currents. However, the reduction of the shim currents comes at the cost of the field homogeneity. Thus, if the residual field standard deviation after shimming with different regularisation methods is the same, the regularisation methods that deliver lower solution norm is considered having better performance.

To intuitively reflect how the regularised shim solutions are correlated with the residual B_0 field norm with different degrees of regularisation, a log-log plot is used to give insight into the regularisation performance. The parametric curve usually has a characteristic shape of "L". Thus, the log-log plot is usually referred to as the L-curve plot, which is defined as:

$$L := \{(log\|I_\lambda\|,\ log\|\Delta B_{meas} - S \cdot I_\lambda\|)\} \tag{3.13}$$

where the first term, $\|I_\lambda\|$, is the norm of shim currents, the second term, $\|\Delta B_{meas} - S \cdot I_\lambda\|$, is the norm of the residual field.

Figure 3.4 shows the L-curve plot of the shimming results using the TSVD (yellow dot), the DSVD (red line) and the Tikh (blue line) with varying regularisation parameters for a slice located at the middle of the brain. The leftmost points where the solutions of the three methods converge correspond to the solution produced by the unconstrained pesudoinverse. The L-curve plot provides an intuitive interpretation that the regularisation errors dominant in the flat region (overregularisation), while the solution is dominated by perturbation errors (sensitivity to noise) in the vertical region.

To compare the performance of different regularisation methods, we choose the TSVD regularisation results as baseline. Then, we find the solutions of the DSVD and the Tikh yielding the same residual norm delivered by the TSVD regularisation by adjusting the regularisation levels. Therefore, the difference in the shimming performance can be attributed solely to the difference in the solution norm. For instance, the dashed black line indicates the residual norm of 8.5Hz which was produced by the TSVD with a truncation of 4 singular values. With the same residual norm, the TSVD, the DSVD, and the Tikh produced solution norm of 30, 25, and 18, respectively.

To compare the performance of slice-wise shimming using the TSVD, the DSVD, and the Tikh in different conditions, simulations were performed over different regularisation levels in the TSVD and carried out on field maps acquired from 143 subjects, as introduced in **Chapter 2, Section 2.3**. The difference in the solution norm between the three methods is plotted against the different regularisation levels in the TSVD - the number of singular values truncated (Figure 3.5). The boxplots coloured with blue, red, and yellow shows difference in solution norm between the TSVD and the

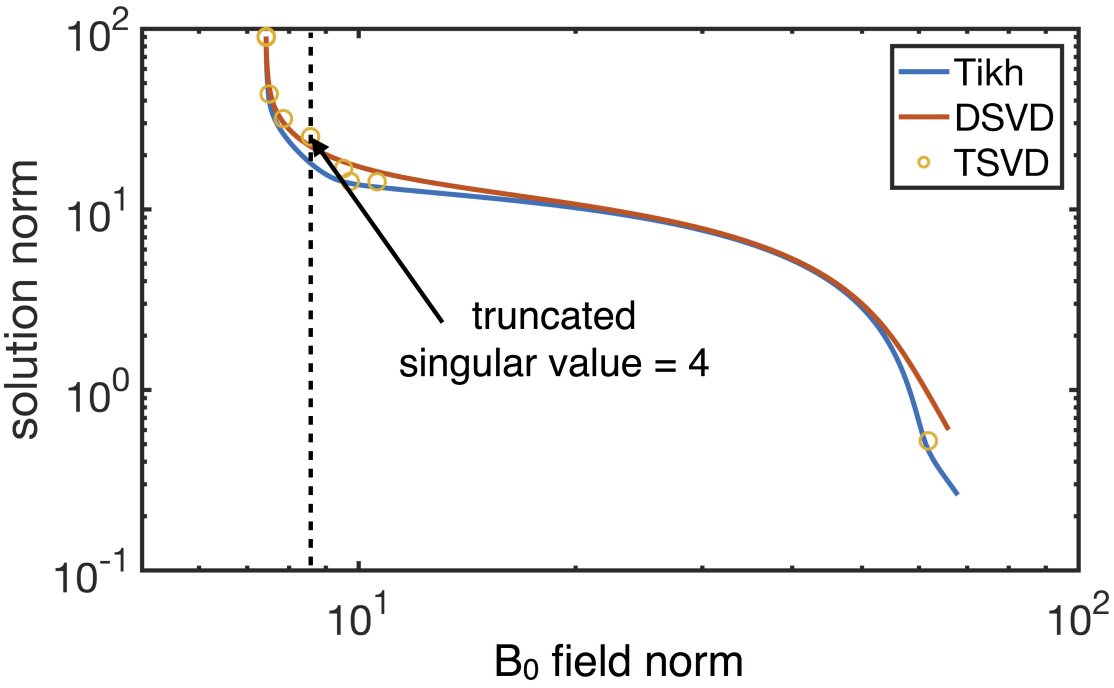


Figure 3.4: Comparison of regularisation effects of the Tikh (blue line), the DSVD (red line) and the TSVD (yellow circles) with varying regularisation levels on an example slice of a subject. In the TSVD, the circle marked by the dashed black line corresponds to the number of truncated singular values equals to 4. The x-axis represents the residual B_0 field norm and the y-axis represents the shim solution norm.

DSVD (blue boxplot), the TSVD and the Tikh (red boxplot), and the DSVD and the Tikh (yellow boxplot), respectively. The Tikh method consistently yielded lowest shim values compared to the TSVD and the DSVD for almost all of the cases. The DSVD and the TSVD delivered comparable regularisation results. The Tikh method produced lower shim solution norm at little cost of B_0 field homogeneity, compared to the DSVD and the TSVD methods.

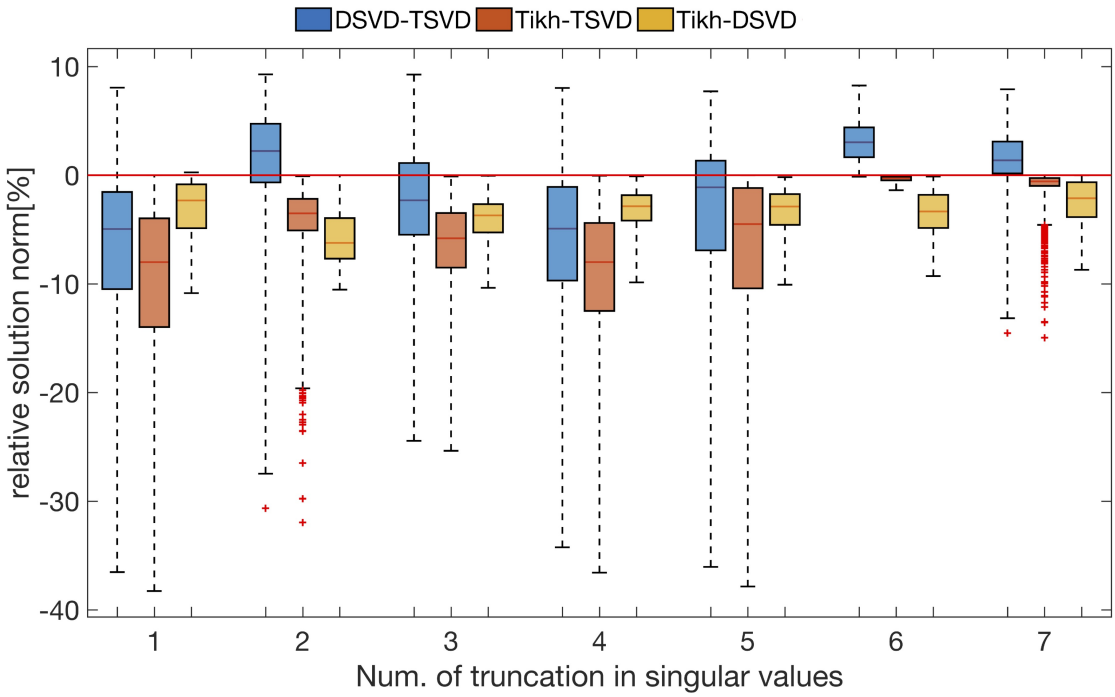


Figure 3.5: Comparison of regularisation effects of the TSVD, the DSVD and the Tikh methods on all of the slices for each subject in the 143 field map database in a second-order shim system. The boxplots coloured with blue, red, and yellow show difference in relative solution norm between (**a**) the TSVD and the DSVD, (**b**) the TSVD and the Tikh and (**c**) the DSVD and the Tikh, respectively.

As demonstrated above, the Tikh method consistently delivers lower solution norm compared to the DSVD and the TSVD at the same residual norm. It suggests that the use of the Tikh method can potentially improve the current use efficiency for shimming applications. Additionally, the Tikh and the DSVD methods have many advantages over the TSVD method. First, in the TSVD, simple truncation of the singular values

will lead to the loss of information about the field variations. Additionally, this will cause round-off errors and which often lead to over-regularisation. The DSVD and the Tikh use smooth filter factors to amplify small singular values. This design overcomes the sudden change of the filter factors in the TSVD. In that case, the field information is preserved and the round-off errors are mitigated. Second, the DSVD and the Tikh can yield a range of regularised solutions. The two methods provide more degrees of freedom on the regularisation which allows more flexible control on the shimming results. It facilitates the use of optimal regularisation parameter searching algorithms, like the L-curve method. However, the TSVD has a limited number of regularisation parameters. The corresponding truncation number is usually chosen empirically, for instance truncation of singular value below a threshold [133].

TSVD Regularisation on high-order shim systems

The evaluation of regularisation methods so far was performed with second-order slice-wise shimming. Therefore, truncation of singular values causes round-off errors, often leading to over-regularisation and non-optimal residual field inhomogeneity. However, a study suggested the TSVD and the Tikh deliver comparable results for higher-order static shimming [133]. To assess the reliability of this argument, simulations were performed on the first-, second-, and third-order slice-wise shimming with the TSVD method on the field maps acquired from 143 subjects.

Figure 3.6 indicates that, with the increase of the shim orders, the change of the average residual field norm between adjacent regularisation levels was reduced. It partly confirms the argument suggested by the literature. High-order shim systems contain more shim terms. Thus, the corresponding shim matrix contain more singular values. Truncation of small singular values will not significantly over-regularise the

shim determination. In low-order shim systems where limited singular values are available, the truncation of singular values yield more regularisation on shimming results. Thus, the Tikh and the DSVD methods are preferable in low-order shimming cases.

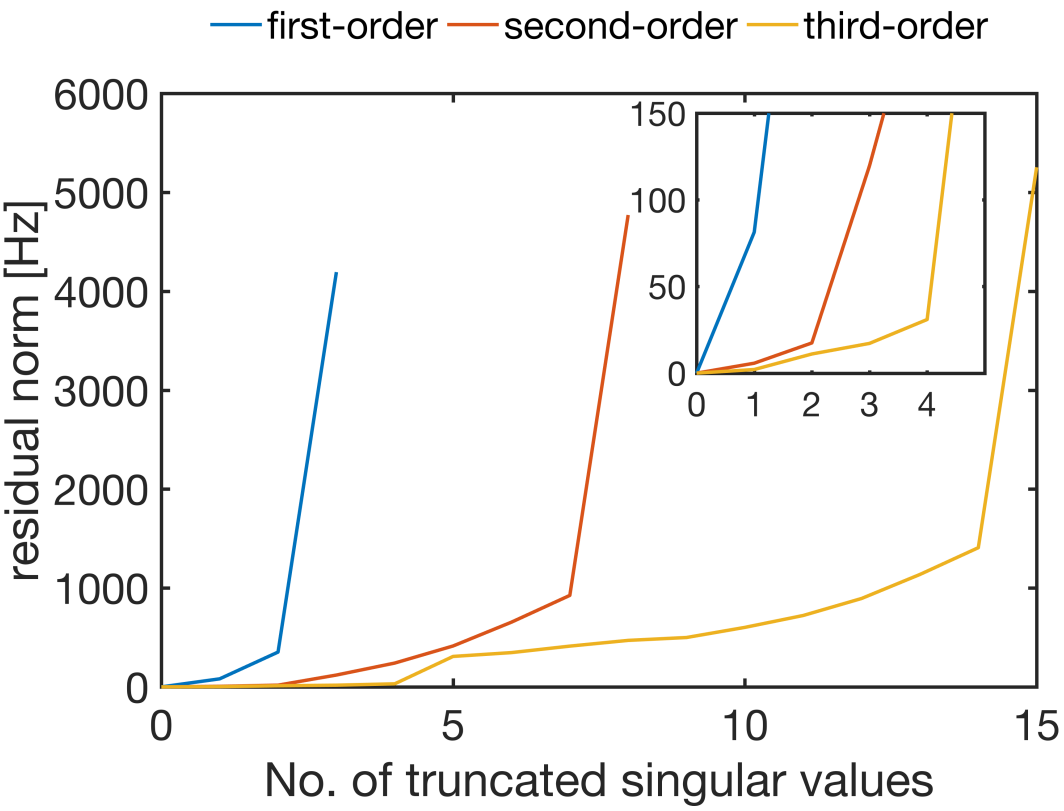


Figure 3.6: Change of the average residual norm with different degree of the TSVD regularisation (number of truncated singular values) for first-, second- and third-order shim systems on all of the slices for each subject in the 143 field map database. The zoomed plot represents regularisation effects with small numbers of the singular value truncation.

3.2 Regularised second-order dynamic shimming

Section 3.1 introduced the TSVD, the DSVD, and the Tikh methods and compared their performance in different shimming cases. We have demonstrated that the regu-

larisation methods could improve the conditioning of the least-squares problems and reduce the current use at the cost of some field homogeneity. Among these methods, the Tikh method can produce lower solution norm than the DSVD and the TSVD when achieving the same field homogeneity.

With regularisation, current strength on each shim channel and their norm attenuates with the increase of the degree regularisation. Thus, we can readily resolve the hardware constraint problems by applying additional penalty on the regularisation parameters, and decouple the need of iterations in searching for current solutions, compared to the iterative LSQlin methods.

One challenge of the regularised least-squares method is choosing regularisation parameters to balance the need for regularisation while minimising the residual field inhomogeneity for slice-wise shimming cases. Using small regularisation yields solutions approximate to the unconstrained least-squares, while large regularisation parameters cause severe degradation of the shimming quality. As imaging conditions vary from slices to slice, applying a universal regularisation parameters for all the shimming cases produces unstable regularisation effects, which is not ideal for the slice-wise shim calculation.

In this section, we propose a fully automated shim determination technique based on the Tikh method in conjunction with the L-curve method (Tikh-L) for optimal regularisation parameter search. The performance of Tikh-L was evaluated in comparison to a wide range of constrained least-squares methods (e.g. the LSQlin and the TSVD) on 143 subjects in simulation and one subject in experiment.

3.2.1 Tikhonov regularisation with the L-curve method

Hardware constraints

As introduced in **Chapter 1** Equation 1.31, the cost of least-squares optimisation is the residual field norm. Inclusion of the Tikh method in least-squares optimisation yields an extra weighting on solution norm, which is written as

$$min\{\|\Delta B_{meas} - S^{\lambda} \cdot I_{\lambda}\|^2 + \|\lambda I_{\lambda}\|^2\} \tag{3.14}$$

where the first term indicates the regularised residual field norm and the second term corresponds to the weighted contribution of the solution norm to the cost function. With regularisation parameter $\lambda = 0$, the formula is reduced to the standard least-squares function which delivers solutions that minimise the residual field norm. With the increase of the regularisation parameters, the function will give more and more weight on minimising the solution norm term. Particularly, the cost function of least-squares optimisation with the Tikh method has an equivalent formulation as the hardware current constraint problem (**Chapter 2** Equation 2.11), where the currents are subject to $\|\lambda I_{\lambda}\|^2 \leq L^2$. Thus, we can readily solve the problem of current constraints by adjusting the regularisation parameters.

Trade-off between solution norm and residual field homogeneity

In the regularised least-squares, if the regularisation parameter is too large, the calculated shim currents would produce much lower shim currents, while severely degrading the residual field homogeneity. Thus, it is important to choose proper regularisation parameters to balance the current use and the residual field homogeneity after shimming. A common method to find the optimal regularisation parameters to make the trade-off between the two quantities is the L-curve method [135]. The optimal solution is defined at the corner of the "L" curve where the curvature of the L-curve

reaches the maximum. The curvature function is defined by

$$\kappa(\lambda) = \frac{\varsigma' \eta'' - \varsigma'' \eta'}{\{(\varsigma')^2 + (\eta')^2\}^{3/2}} \tag{3.15}$$

where η' and η'' correspond to the first and the second derivatives of the logarithm of the shim solution norm with respect to the regularisation parameters λ and similarly ς' and ς'' denotes the first and the second derivatives of the logarithm of the residual field norm.

Implementation

The procedures of using the Tikh-based regularised least-squares and the L-curve method to find the optimal shim current solutions are shown step by step in Figure 3.7a, and illustrated in the plot of the shim solution norm versus the residual field norm on an example slice at the level of the middle of the brain in Figure 3.7b.

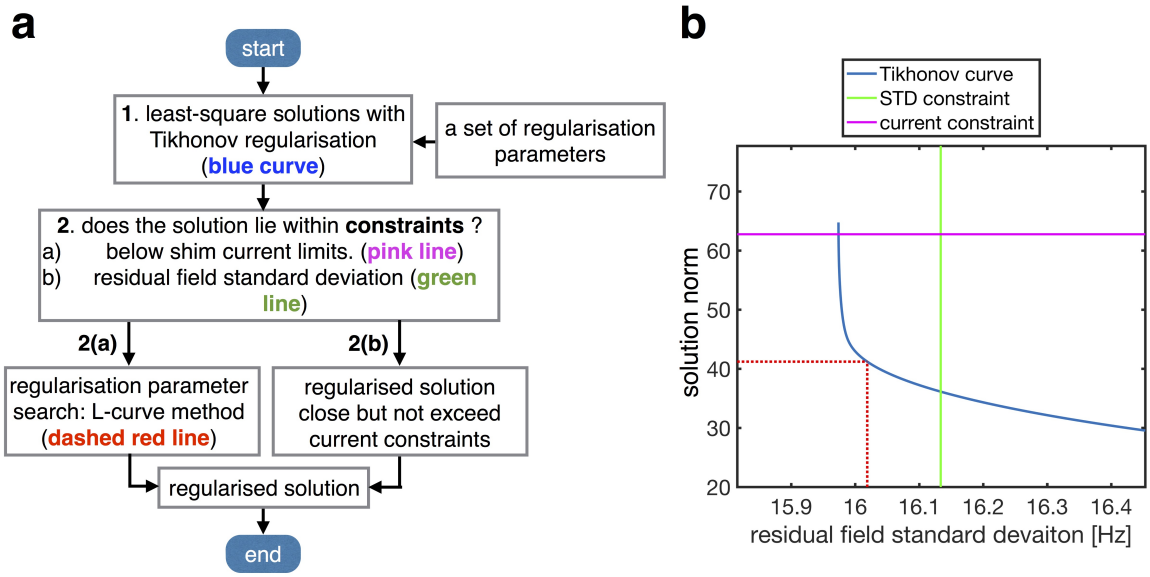


Figure 3.7: **a**: Flow chart of shim determination procedure of the proposed Tikh-L method. **b**: Schematic example of determining optimal regularisation parameter using the Tikh-L method. The optimal shim solution (dashed red line) is determined from a range of regularised solutions (blue line with a characteristic L shape) using the L-curve method along with dual constraints on the solution norm (solid or dashed pink line) and residual B_0 standard deviation (green line).

The Tikh-L method mainly consists of three stages:

1. Calculate multiple sets of shim current solutions using least-squares optimisation with the Tikh method. The calculated shim currents correspond to the blue curve in the plot in Figure 3.7b. The left end point of the curve corresponds to the result with no regularisation ($\lambda = 0$), which is equivalent to the unconstrained method.

2. Find the solutions that lie within the dual boundaries of (1) maximum shim current limits and (2) standard deviation constraints (not greater than 1% of the unconstrained method solution) to ensure that shim currents are within hardware constraints and avoid over-regularisation. The solid pink lines represent the maximum boundary shim currents of different current constraints with 3A on each channel. The green line represents the maximum boundary residual norm of residual field.

   The effects of B_0 field inhomogeneity on MR image quality vary in different imaging conditions, such as imaging sequence use, acquisition timing and so on. In different imaging scenarios, some sequences used are highly sensitive to B_0 field inhomogeneity, such as EPI, while some are less so, for instance MPRAGE. To give some indication of the impact of shim quality, we simulated field inhomogeneity induced EPI distortion [40]. 1% standard deviation difference corresponds to actual residual B_0 standard deviation degradation below 1 Hz. For the EPI acquisition setup used in the text, 1 Hz difference corresponds to $\sim$ 0.04 pixel shift in EPI images which is still relatively small compared to noise during image acquisitions. Therefore, we consider the 1 Hz standard deviation criterion is likely to be acceptable for many shimming applications.

   (a) If solutions exist within the boundaries, use the L-curve method to determine the optimal trade-off between the shim solution norm (an indication

of shim current amplitude) and residual field inhomogeneity. The dashed red line represents the optimal solution determined by the L-curve method.

(b) If no solutions exist fulfilling both of the boundary conditions, pick the solution yielding lowest residual standard deviation within the shim current constraints.

3.2.2 Performance evaluation of regularised optimisation algorithms

The performance of the proposed Tikh-L method for slice-wise second-order dynamic shimming was evaluated in comparison to a range of shimming algorithms, as follows

1. The unconstrained linear least-squares approach using the pseudoinverse. Shim currents were calculated by fitting the first- and second-order shim fields to the measured field map using the unconstrained method. This optimisation minimises standard deviation of field inhomogeneity in the ROIs. As the optimal shim settings solutions that deliver best field homogeneity does not necessarily lie within shim constraints, the unconstrained method was considered as a benchmark for the ideally achievable shimming quality.

2. A Matlab-embedded constrained least-square method with current constraints using two optimisation methods: interior-point algorithm (LSQlin) and trust-region-reflective (LSQlin-R).

3. Regularised least-squares with hardware constraints of currents based on the TSVD (TSVD-H), the DSVD (DSVD-H) and the Tikh (Tikh-H) methods. These methods were implemented in two stages. (1) Calculate multiple sets

of shim current solutions with a range of regularisation parameters and (2) pick the solution yielding lowest residual standard deviation within the shim current constraints.

4. Regularise shim matrix using two fixed and empirical regularisation parameters, 0.01 (Tikh-F1) and 0.02 (Tikh-F2), respectively. If calculated shims run into hardware constraints, we adopt the Tikh-H method.

To assess the performance of slice-wise shimming using aforementioned shim determination methods, simulations of algorithms were performed on all of the slices for each subject in the 143 field map database, as introduced in Chapter 2. The ROIs of the brain content and their masks for each slice were determined using BET [106]. Those slices can be roughly classified to three categories according to their field characteristics:

- Slices with large volume coverage and strong field variations, which are mainly located in the mid part of the brain (the frontal cortex);
- Slices associated with slow field variations, located in the superior brain (the cerebrum);
- Slices with small volume coverage and strong field distortions, located at the brain's lower-level structures (the brainstem and cerebellum).

Assessment of slice-wise shimming quality in different parts of the brain allows evaluating the performance of aforementioned algorithms in different imaging conditions.

In the dynamic shimming hardware setup, the shim amplifiers are designed to work within current constraints of 10A and the summation of them is limited below 20A. The pre-emphasis method overdrives the input currents, which can reach 150% of the

current step applied, to compensate for expected eddy-current-induced fields. Thus, in the simulation and experiment setup, the current limits of the second-order shims were set to 3A to prevent output currents hitting the current limits of the amplifiers. The performance of the algorithms is quantified by the following measures.

Residual field standard deviation

The standard deviation of the residual field in the region defined by the brain mask was used to evaluate the field homogeneity after shimming. The unconstrained least-squares method achieves the theoretically smallest residual field standard deviation. With the hardware limits of currents, the shimming performance is constrained. The performance of shimming using different methods was evaluated by how close the obtained residual standard deviation is to the unconstrained shimming results.

Condition number

The condition number is an indication of the sensitivity of shim current solutions to the measurement noise. It reveals the properties of the shim matrix, which is defined by Equation 3.4. In the case of severely ill-conditioned shim matrix, where condition number is large, noise in the measured field maps or shim fields causes significant change of the shim solutions. Small condition number corresponds to a high degree of robustness of the shim solutions to the noise. Therefore, the condition number of the shim matrix was calculated as a measure of the robustness of the algorithms.

Shim currents

As stated in **Section 3.12**, it is inefficient to use high shim currents for shimming applications. Large shim solution norm can lead to unstable corrective fields, problems of coil heating, and hardware current constraints. Additionally, fast switching high shim currents could cause strong eddy current effects. Thus, the norm of the

shim solutions was calculated for the assessment of the algorithm performance.

Optimisation efficiency

Shim optimisation methods trade some residual field homogeneity for reduced current use to meet the current limits or improve the conditioning of the least-squares problem. As it is inefficient to use high amplitude shim currents for shimming, the optimised results with large reduction in the solution norm and low degradation in the residual field standard deviation are preferred. Thus, a measure of how much reduction in the solution norm resulting from 1 Hz degradation in the standard deviation is of interest. In this work, the performance of different shim optimisation methods were also evaluated by the ratio of the average reduction in the solution norm to the average increase in the residual field standard deviation, termed optimisation efficiency. This measure allows direct comparison of various optimisation algorithms accounting for both the residual field standard deviation and the solution norm.

Computational speed

The computational time spent in determining slice-wise shim currents is an overhead to shimming procedures. The speed of calculating shim settings was considered as another measure of the algorithm performance. To evaluate the computational speed of proposed algorithms, assessments were carried out on a 2.6 GHz Intel Core i7 CPU with 16GB of RAM.

3.2.3 Comparison of LSQlin, TSVD-H and Tikh-L

The performance of slice-wise shimming using the LSQlin, the TSVD-H and the Tikh-L was evaluated on all of the slices for each subject in the 143 field map database.

The scatter plots in Figure 3.8 show the relationships between the solution norm and the residual field standard deviation for the same slices using these three methods in two scenarios: shim determination without (Figure 3.8a) and with (Figure 3.8b) hitting current constraints.

In scenario (a), slices without hitting current constraints, the solution of the LSQlin and the TSVD-H is unconstrained and deliver shim results equivalent to the unconstrained method. In the histograms next to the horizontal and the vertical axis, the LSQlin and the TSVD shimming results are overlaid. The Tikh-L provided a reduction in the solution norm without degrading the standard deviation of residuals more than 1.1Hz compared to the unconstrained method.

In scenario (b), slices hitting current constraints, it can be observed that the TSVD-H produced larger increase in mean standard deviation of residuals, 4.5Hz, compared to 0.7Hz for the LSQlin. The result intuitively shows poor performance of the TSVD-H. The Tikh-L reduced 52% solution norm on average, compared to the unconstrained method, while delivering mean residual standard deviation increase of merely 1.4Hz.

The detailed interpretation of the simulation results on slices hitting current constraints in the lower, the middle and the superior parts of the brain are depicted in Figure 3.9. The scatter plots in Figure 3.9(a-c) show the relationships between the standard deviation of residuals and the solution norm produced by the LSQlin, the TSVD-H, and the Tikh-L in these three regions, separately. The boxplots in Figure 3.9d and Figure 3.9e quantify the solution norm and the standard deviation of the corresponding residuals in three brain regions, respectively.

Figure 3.9b shows that in the middle of the brain only a few slices suffered from the current constraints. These results can be partially explained by the fact that slices located near the middle of the brain have large ROI size. Thus, their shim

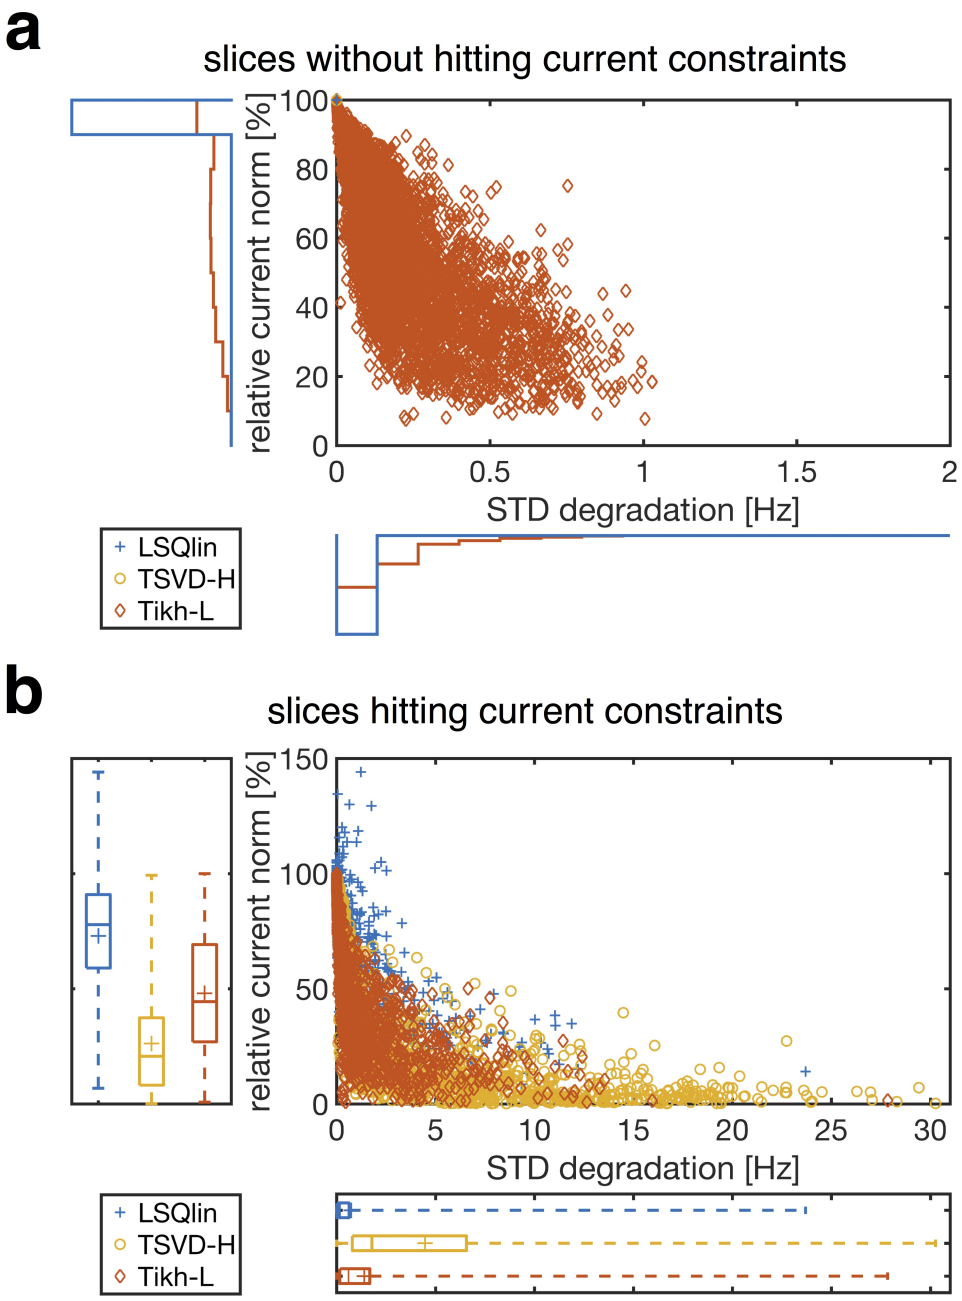


Figure 3.8: Comparison of simulated shimming performance with the LSQlin (yellow), the TSVD-H (red), and the Tikh-L (blue) method on all of the slices for each subject in the 143 field map database. **a**: The scatter plot of the relative solution norm reduction, along with the relative increase in the residual field standard deviation compared to the unconstrained method, on slices that do not hit current constraints. **b**: The scatter plot of results on slices that do run into current limits. The corresponding distributions are shown separately in the histograms and boxplots.

matrices are relatively well-conditioned. As illustrated in Figure 3.9a and 3.9c, the current constraint problems were more likely to occur on slices in the inferior and superior parts of the brain where the ROI sizes are small. Figure 3.9a shows that the TSVD-H was evidently poorly performed for slices in the inferior part of brain in comparison to the LSQlin and Tikh-L. The TSVD-H traded a 4 Hz mean standard deviation degradation, compared to the unconstrained method, for a 71% reduction of the solution norm. The Tikh-L delivered on average 50% reduction of the solution norm, while only degrading the residual field standard deviation by 1 Hz, compared to the unconstrained method. The simulation results demonstrated the superiority of the Tikh-L in effectively reducing the solution norm, while preserving a good level of field homogeneity, compared to the TSVD-H method.

Performance under different current constraints

In dynamic shimming, high performance shim amplifiers are desired due to a wide output current range and high precision amplification. However, high-performance shim amplifiers are not guaranteed in different shim systems. Low performance shim amplifiers have limited current operation range which limits the best shimming quality that can be obtained. The reproducibility of high shimming quality under different hardware constraints of currents is a criterion to assess the performance of the shim current optimisation algorithms. Thus, the performance of the LSQlin, the TSVD-H, and the Tikh-L method under the shim current limits of 1A, 2A, 3A, and 4A on second-order shim channels was evaluated on all of the slices for each subject in the 143 field map database.

Figure 3.10a and 3.10c illustrate, for slices hitting current constraints, a decrease of current limit caused a degradation of residual field standard deviation and an reduc-

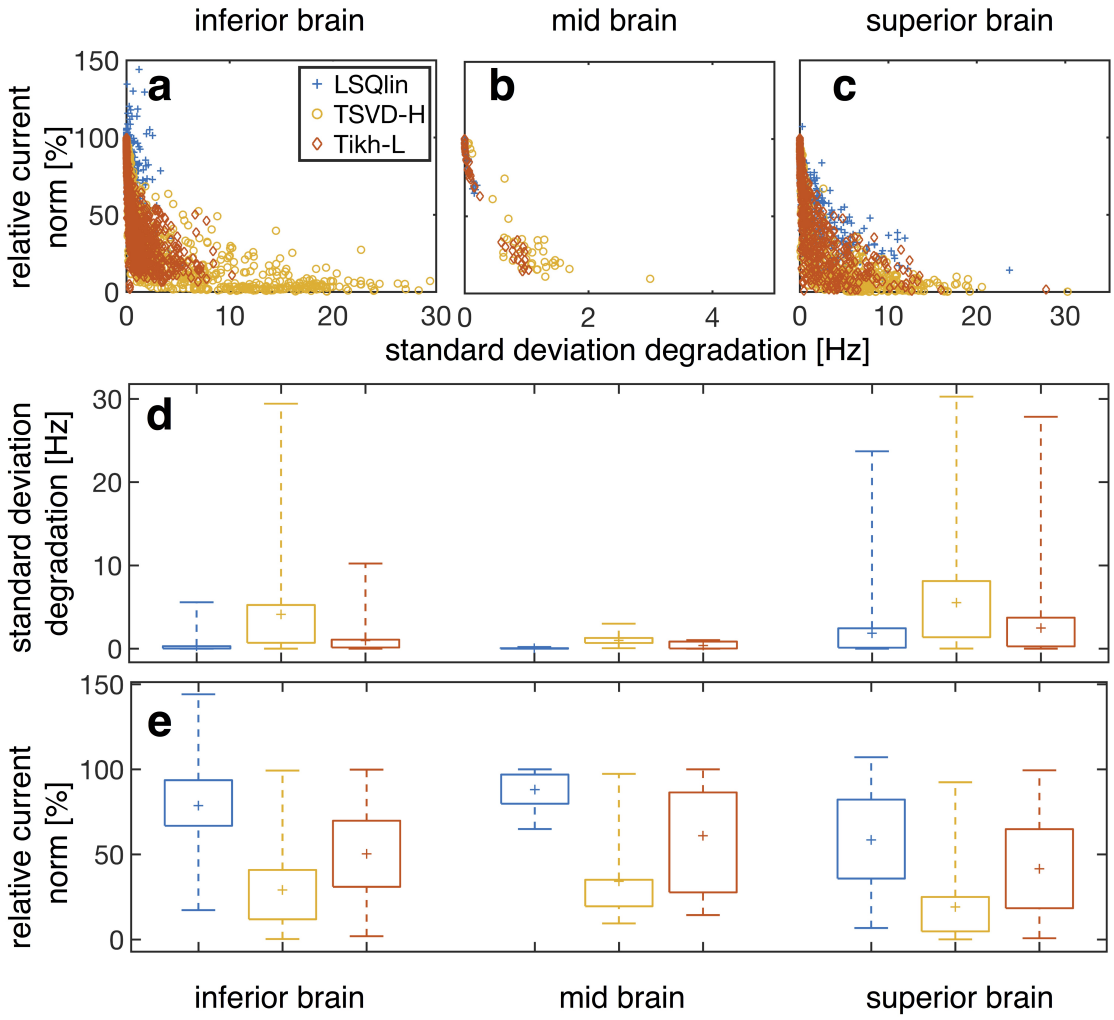


Figure 3.9: Comparison of simulated shimming performance with the LSQlin (blue), the TSVD-H (yellow) and the Tikh-L (red) on slices hitting current constraints, in the inferior, mid and superior brain regions on all of the slices for each subject in the 143 field map database. The scatter plots (**a-c**) indicate the relative solution norm, along with the relative standard deviation degradation of the residuals compared to the unconstrained method. Boxplots (**d, e**) show the corresponding distributions of scatter plots in **a-c**.

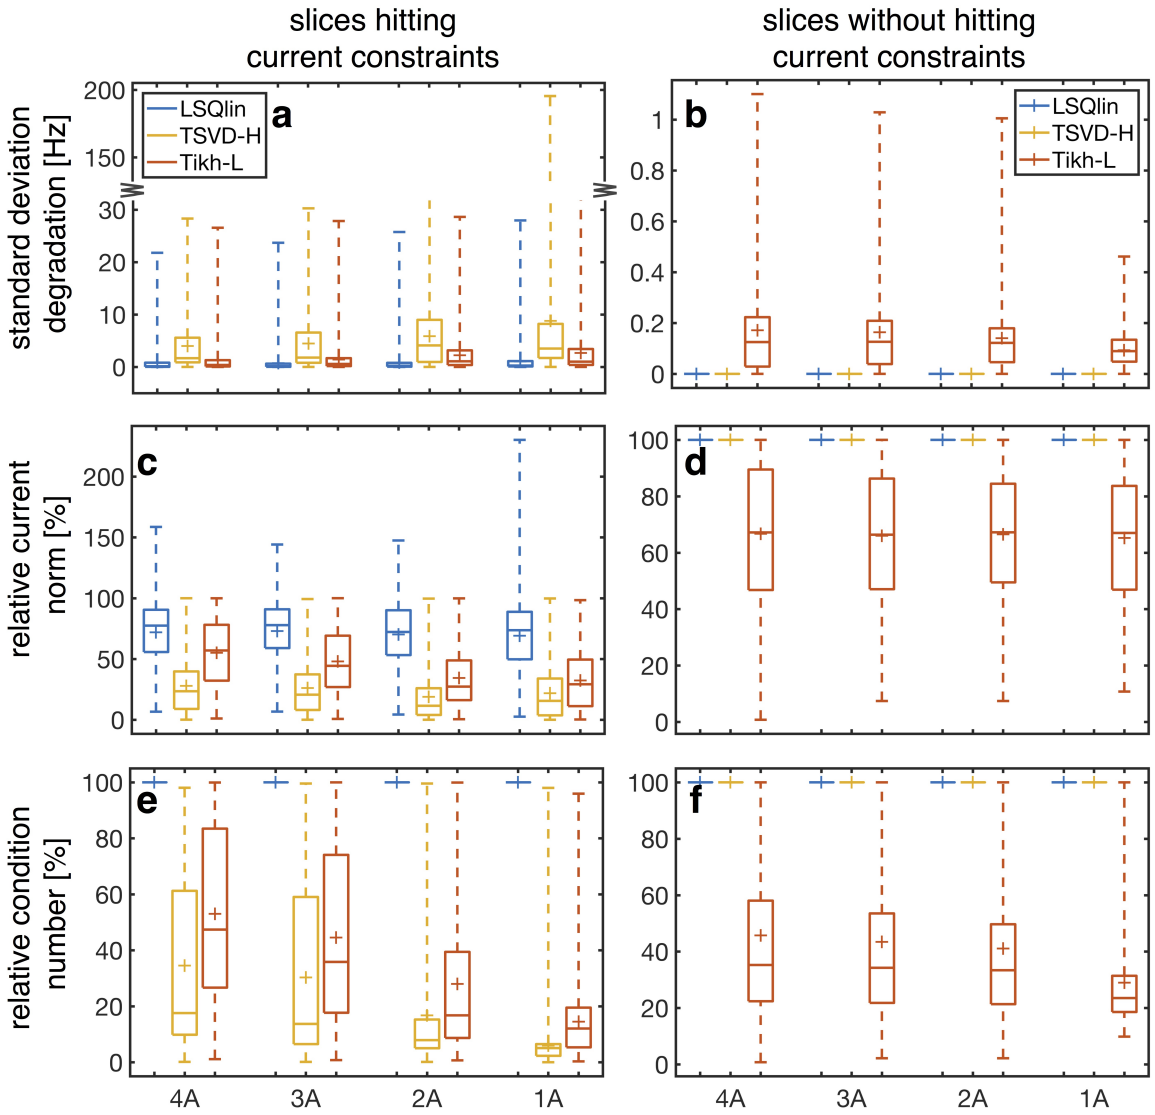


Figure 3.10: Evaluation of the performance of the LSQlin, the TSVD-H, and the Tikh-L on slices with (**a, c, e**) and without (**b, d, f**) hitting current constraints from 1A to 4A on second-order shim channels. The boxplots indicate relative residual field standard deviation degradation (**a, b**), solution norm (**c, d**), and condition number (**e, f**) compared to the unconstrained method, respectively.

tion of the solution norm. With the 4A current constraints, the LSQlin, the TSVD-H, and the Tikh-L methods traded on average 0.9Hz, 4Hz, and 1.3Hz degradation of the residual field standard deviation compared to the unconstrained method, for 28%, 72%, and 45% reduction of the solution norm, respectively. In contrast, with the 1A current constraints, the three methods produced 1Hz, 8.8Hz, and 2.7 Hz mean degradation of standard deviation, respectively. The results demonstrated that the reduction in current limits did not significantly affect the shimming performance of the LSQlin and the Tikh-L compared to the TSVD-H.

As shown in Figure 3.10b and 3.10d, for slices without running into current constraints, the Tikh-L produced a reduction of the residual field standard deviation and the solution norm with the reduction of current limits. The improvement in the field homogeneity is due to the reduced number of slices without running into current constraints. This observations also imply that slices having low residual field standard deviation were less likely to run into current constraints..

Figure 3.10e and 3.10f compare the condition number optimised using the three methods. The LSQlin optimisation did not alleviate the degree of conditioning. The Tikh-L produced much reduction in condition number in both scenarios. However, the TSVD method only reduced the condition number for slices running into current constraints.

Figure 3.11 compares the optimisation efficiency of slice-wise dynamic shimming using the LSQlin, the TSVD-H and the Tikh-L. In general, both the Tikh-L and LSQlin exhibited much higher optimisation efficiency than the TSVD-H at all of the tested current constraints. Moreover, the Tikh-L outperformed LSQlin at current constraint of 2-4 A.

A monotonic reduction of optimisation efficiency was observed for the Tikh-L and

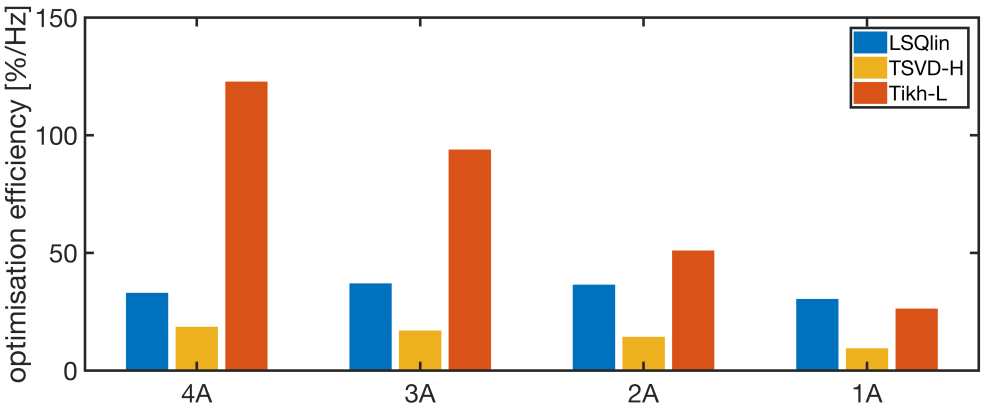


Figure 3.11: Comparison of optimisation efficiency of LSQlin (blue), the TSVD-H (yellow), and Tikh-L (red) methods at current constraints from 1A to 4A on second-order shim channels

the TSVD-H. It can be partially explained by the fact that, with the decreasing of the current limits, more and more slices hit the current constraints.the hardware constraints force the regularisation parameters to large values which lead to over-regularisation of the least-squares problem. Thus, slices running into those current constraints exhibited lower optimisation efficiency.

Computational speed

A comparison of the average computational time of the LSQlin, the TSVD-H, the Tikh-H, and the Tikh-L methods in calculating the shim settings with different ROI sizes is shown in Figure 3.12. In general, the computation time of these methods increased approximately linearly with the size of ROIs. The TSVD-H, the Tikh-H and the Tikh-L exhibited faster computational speed in comparison to the LSQlin. It can be explained by the fact that the LSQlin requires several iterations in order to get converged solutions, whereas the regularisation-based least-squares obtain current solutions in non-iterative and analytical ways which are more time efficient. The TSVD-H and the Tikh-H showed a comparable computational time. Incorporating

the L-curve method in determining the optimal regularisation parameter to the Tikh-H added a little more computational time.

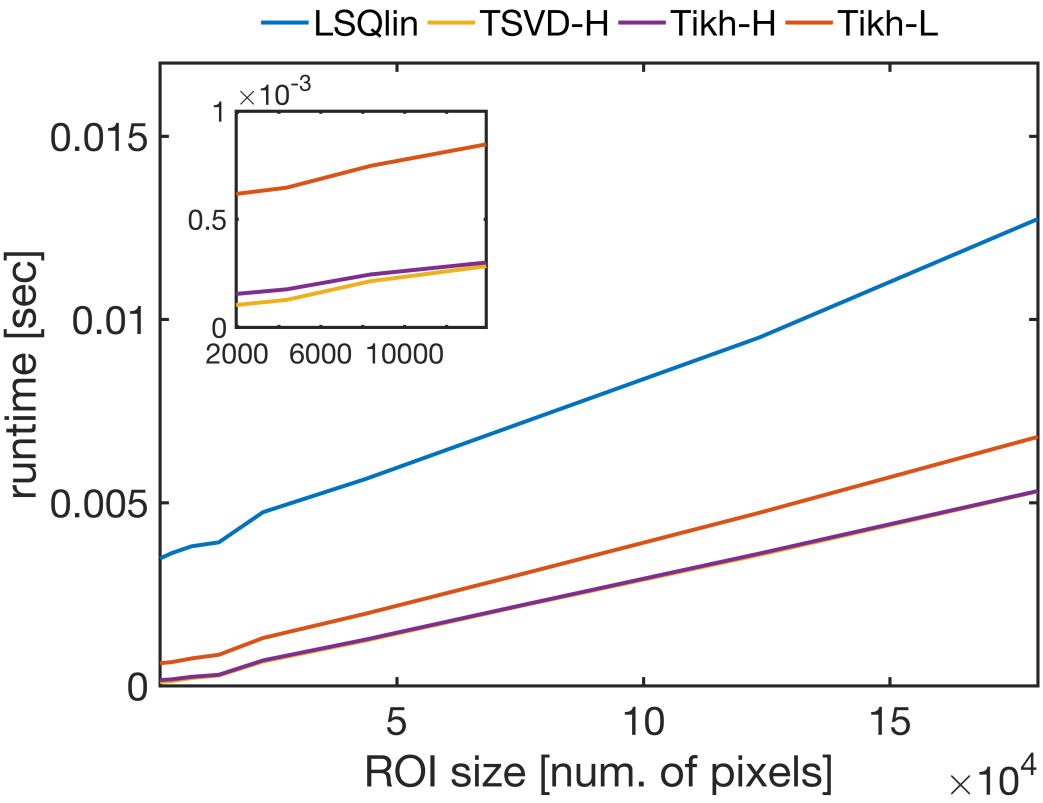


Figure 3.12: Comparison of average runtime of the LSQlin (1), the TSVD-H (2), the Tikh-H (3), and the Tikh-L (4) methods, against the size of ROIs.

3.2.4 Comparison of different optimisation methods

Simulations so far have compared the performance of the Tikh-L, the LSQlin and the TSVD-H. Next, we would like to compare the Tikh-L with a range of methods.

Figure 3.13 shows the comparison of the optimisation algorithms, the LSQlin, the LSQlin-R, the TSVD-H, the DSVD-H, the Tikh-H, the Tikh-F1, the Tikh-F2 and the Tikh-L, on slices with (a, c, e) and without (b, d, f) running into current constraints.

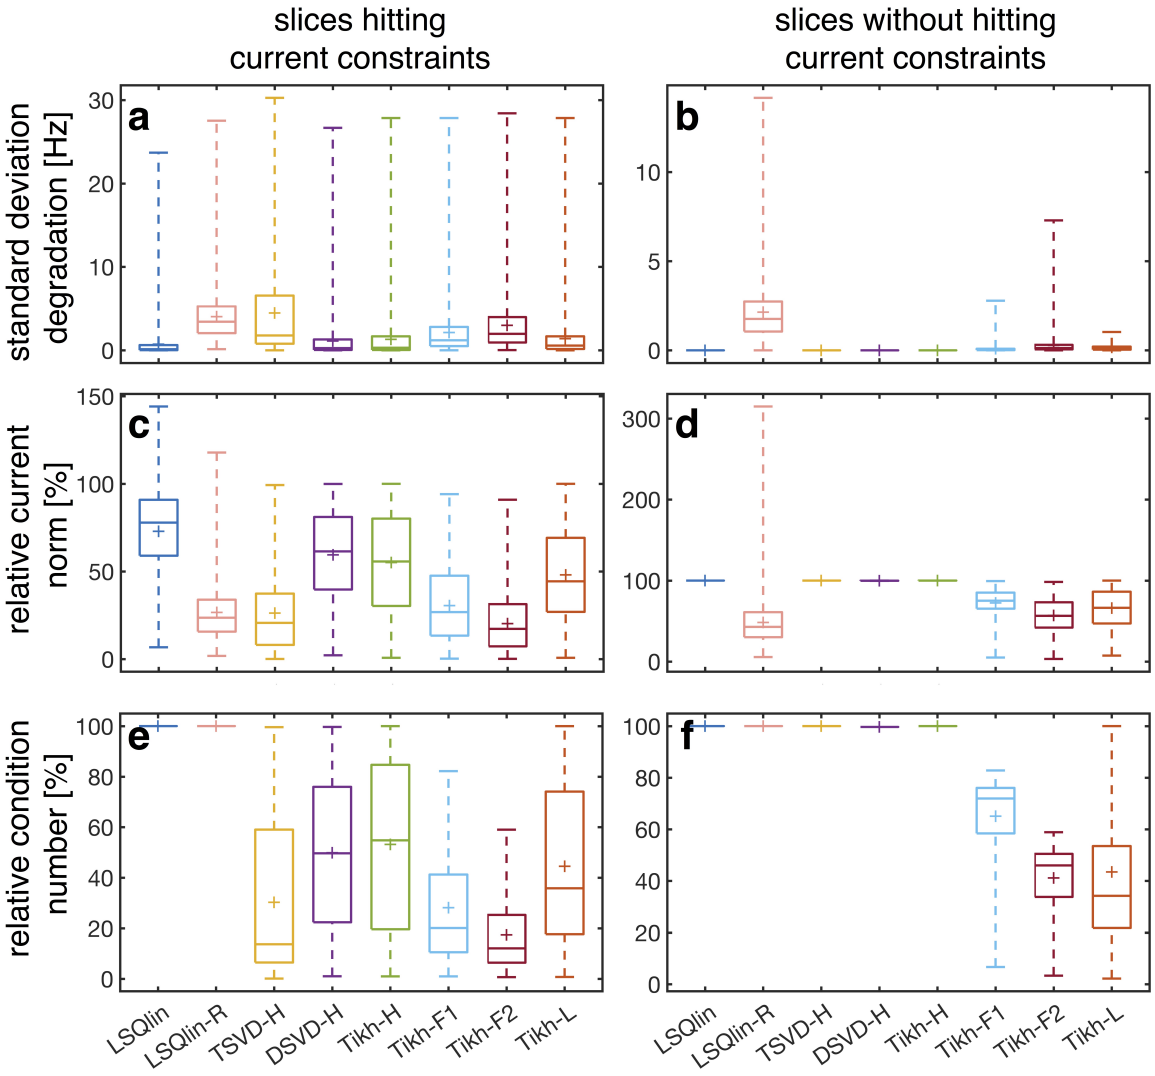


Figure 3.13: Comparison of simulated shimming performance of the LSQlin, the LSQlin-R, the TSVD-H, the DSVD-H, the Tikh-H, the Tikh-F1, the Tikh-F2 and the Tikh-L, on slices with **(a, d, e)** and without (**b, d, f**) hitting current constraints. The boxplots indicate relative residual field standard deviation degradation (**a, b**), solution norm (**c, d**), and condition number (**e, f**) compared to the unconstrained method, respectively.

For slices hitting current constraints, the LSQlin produced on average 0.7 Hz standard deviation degradation of the residuals, compared to the unconstrained method, which was lowest among all the tested methods. However, the LSQlin did not reduce much solution norm, compared to other methods, and it did not improve the conditioning of the least-squares problem. The TSVD-H yielded on average 4.5 Hz standard deviation degradation of the residuals and produced 74% solution norm reduction. In contrast, the Tikh-F2 yielded more solution norm reduction, 80%, smaller standard deviation degradation, 3 Hz, and lower condition number, 17%, compared to the TSVD-H. The results confirmed the findings that the Tikh method outperformed the TSVD method in shimming. The Tikh-L produced on average 1.4 Hz standard deviation degradation of the residuals which was smaller than the Tikh-F1 and Tikh-F2. This result implies that using fixed regularisation causes slight over-regularisation to the least-squares problem.

In order to conclusively compare the performance of these methods, the average optimisation efficiency of slice-wise shimming using each method was calculated for all of slices on subjects in the database, regardless of whether or not hitting current constraints, as shown in Figure 3.14. The Tikh-H and the DSVD-H yielded on average 34 %/Hz and 37 %/Hz optimisation efficiency, compared to 17 %/Hz for the TSVD-H. Introducing small regularisation parameters in the Tikh method improved the optimisation efficiency. The Tikh-F1 produced on average 73 %/Hz optimisation efficiency. However, further increase in regularisation parameters did not deliver higher optimisation efficiency. Incorporation of the L-curve method increased the average optimisation efficiency to 93 %/Hz, which outperformed the rest of shim optimisation methods.

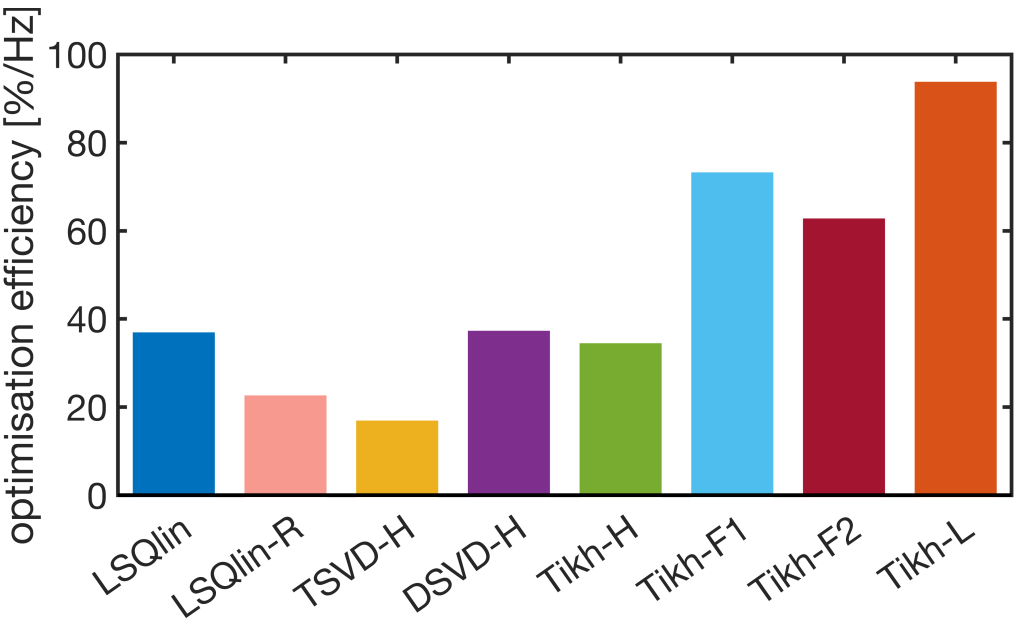


Figure 3.14: Comparison of optimisation efficiency of the LSQlin, the LSQlin-R, the TSVD-H, the DSVD-H, the Tikh-H, the Tikh-F1, the Tikh-F2 and the Tikh-L.

3.2.5 Experimental validation of different shim optimisation methods

In order to test the validity of the simulation results, slice-wise dynamic shimming experiments were performed based on the hardware and procedures described in **Chapter 2**. The performance of the LSQlin, the TSVD-H, and the Tikh-L was tested. The solution norm, condition number and standard deviation of the residuals were simulated based on an experimentally acquired calibration field map (voxel size=2 × 2 × 2 mm^3, TE1=4.08 ms, ΔTE=1.02 ms, slices=75, flip angle=39°, FOV=220 × 220 mm^2). Echo-planar images (voxel size=2 × 2 × 2 mm^3, slices=75, TR=3000 ms, TE=28 ms, FOV=220 × 220 mm^2, GRAPPA factor=2, echo-spacing=0.72 ms) were used to evaluate geometric distortion and signal dropout.

Figure 3.15 compares shimming results on solution norm (a), condition number (b), and residual B_0 standard deviation (c) using the LSQlin, the TSVD-H and the Tikh-L methods across all the slices of the field map acquired from the scanned subject.

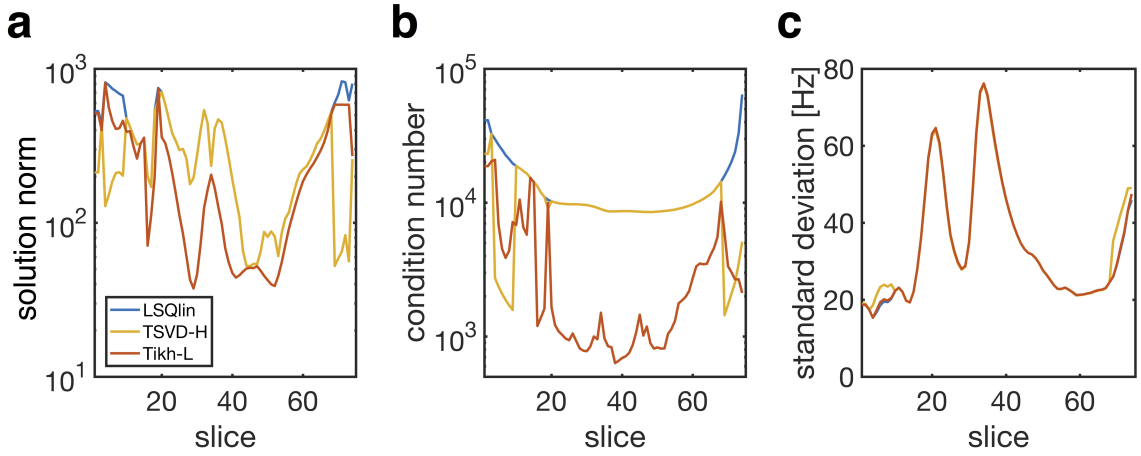


Figure 3.15: Comparison of solution norm (an indication of required shim current amplitude) (**a**), condition number of the shim matrices (**b**) and residual B_0 field standard deviation (**c**) using the LSQlin (yellow), the TSVD-H (red) and the Tikh-L (blue) algorithms across different slices based on a field map experimentally acquired from one subject not included in the database.

Shimming for slices containing fewer brain voxels was more ill-conditioned compared to other regions (Figure 3.15c). The LSQlin yielded solutions with high norm in the inferior part of brain which was likely attributed to the strong local field distortions. The Tikh-L dramatically reduced solution norm and condition number, but without significantly degrading the residual field homogeneity. For instance, at slice 17, the level of frontal cortex, the Tikh-L provided a 34% decrease of the solution norm while degrading the residual B_0 field standard deviation by only 0.2 Hz, on average, compared to the LSQlin. For slices hitting current limits, the TSVD-H yielded an evident reduction in the solution norm and the condition number at the cost of a large field homogeneity.

The corresponding calculated B_0 offset, first- and second-order shim settings for this subject with the LSQlin, the TSVD-H, and the Tikh-L are shown in Figure 3.16. The calculated shim settings reflect the simulation results on the solution norm in Figure 3.15. Shim calculation was more likely to run into shim constraints for those slices associated with B_0 variations of high spatial dependence and the severely ill-

conditioned shim matrices, for instance, slices located in the inferior and superior brain regions. The Tikh-L method delivered relatively more stable and lower amplitude shim settings, especially for shim terms having strong spatial field dependence in the through-slice direction (e.g., ZX and ZY), across the slices.

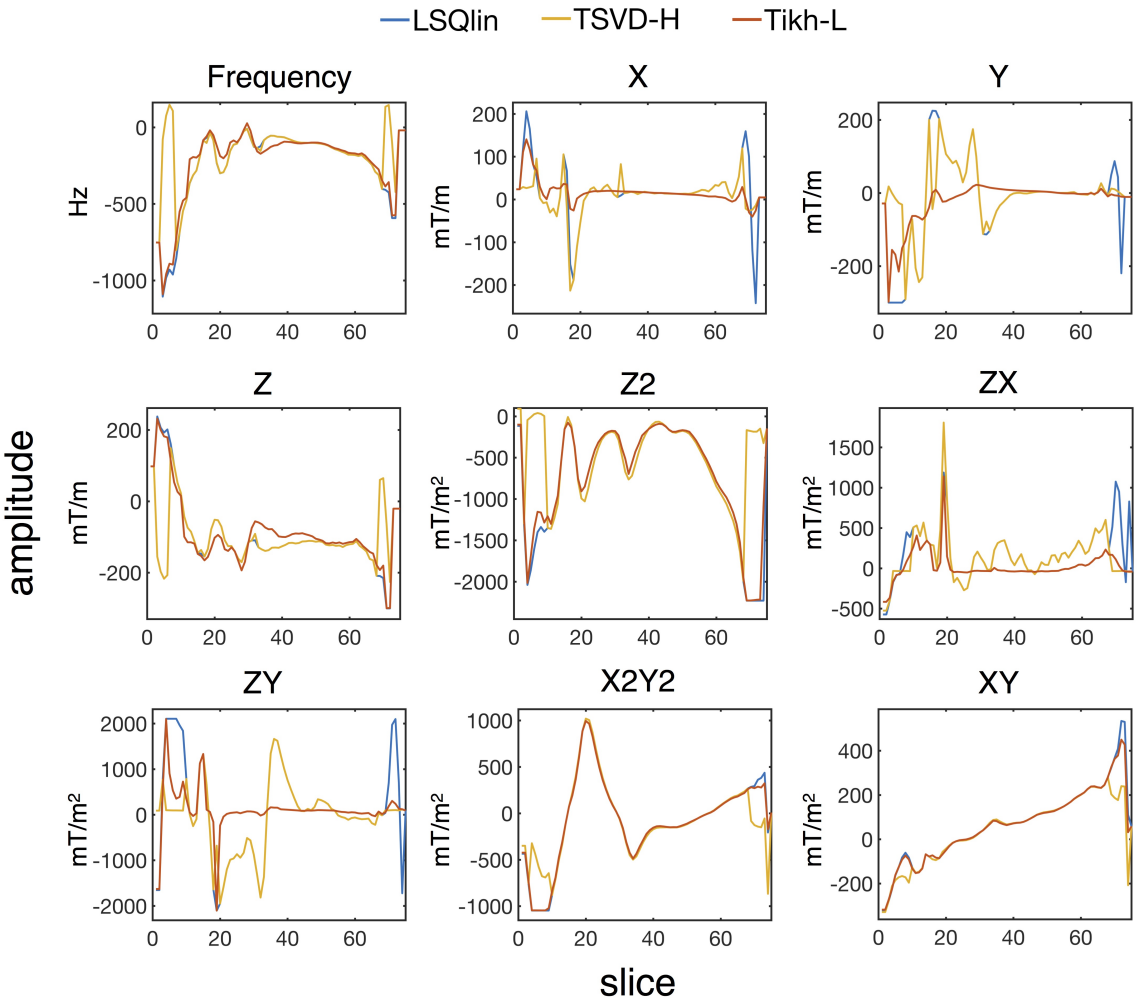


Figure 3.16: Comparison of strengths of the B_0 offset, first- and second-order shim terms determined by the LSQlin, the TSVD-H, and the Tikh-L across different slices based on a field map experimentally acquired from one subject not included in the database.

Figure 3.17 shows sagittal and transversal views of EPI images acquired on a subject using different shimming methods. To visualise the image distortion artefacts, green outlines generated based on the field map acquisition were overlaid to the EPI images.

In EPI images, a substantial reduction of the distortion artefacts was observed. The shimming quality of slice-wise dynamic shimming using the proposed Tikh-L method was comparable to the LSQlin. Moreover, around 50% reduction on average in the shim currents is visible across most of the slices. For slices hitting current constraints, the dramatic current reduction was observed using the TSVD-H and the Tikh-L. However, the TSVD-H yielded substantial signal loss at slices located at the brainstem due to large shim step sizes, as shown in Figure 3.16a. The Tikh-L did not visibly cause signal loss.

3.2.6 Discussion

The feasibility of the Tikh-L approach was demonstrated for slice-wise dynamic shimming at 7T. The technique simultaneously addresses fixed hardware constraints and improves shim current efficiency and the conditioning of the least-squares problem by tunable regularisation. Simulations and experiments indicate that the Tikh-L method exhibited great effectiveness in reducing shim currents, while delivering the shimming quality comparable to that achieved by the LSQlin. In the simulation, the LSQlin achieved the lowest residual B_0 standard deviation numerically. However, this comes at a cost of high shim current amplitude and does not translate to better shimming in practice due to measurement and instrumental noise. Previously it was reported that the LSQlin is not able to deliver better solutions than the TSVD-H using higher-order shims. This problem was not observed here. This can by partly explained by the fact that, in low-order shim systems,the LSQlin does not suffer from high-order shimming problem, like shim interaction.

The regularisation-based methods trade some standard deviation for much improved conditioning of the least-square problems, which potentially reduces the effects of

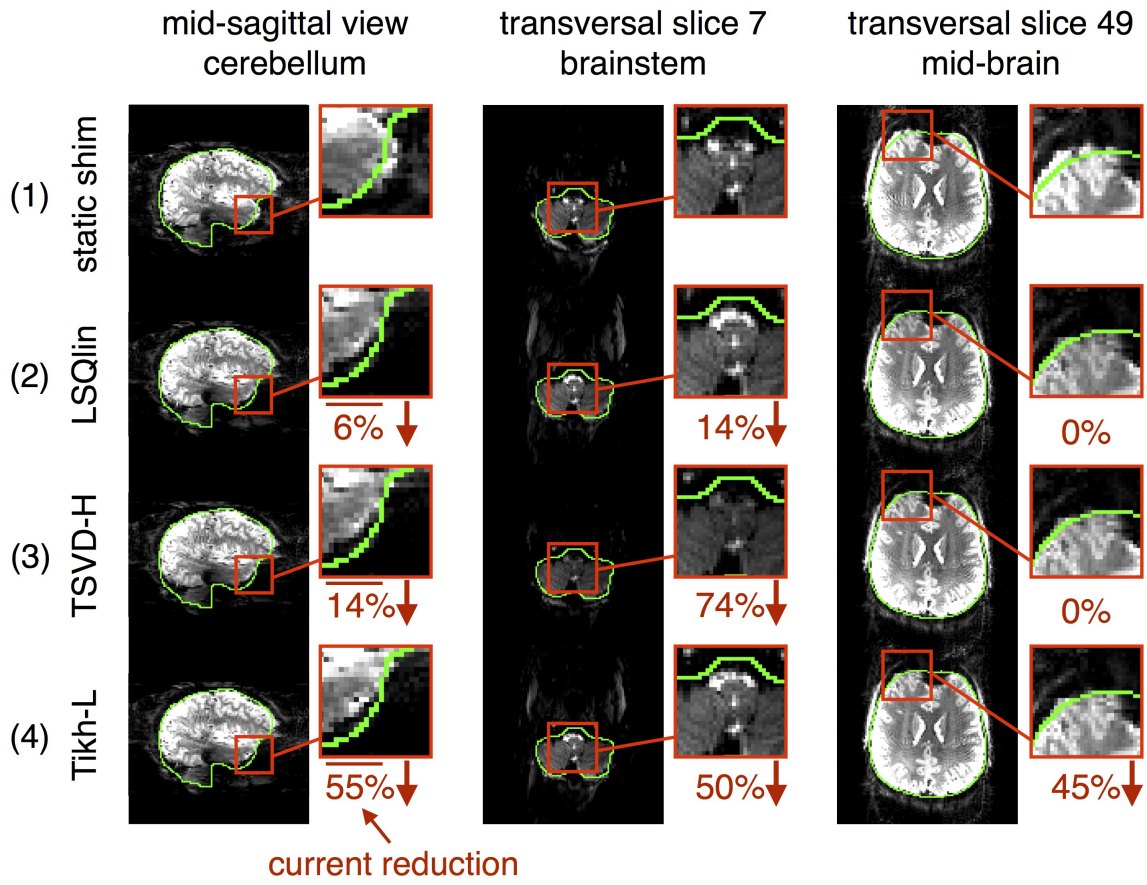


Figure 3.17: EPI images acquired using the different shim calculation methods in one subject not included in the database. The columns from left to right contain a sagittal image through the temporal lobe, and transversal images at the brain stem and the mid-brain regions. The rows from top to bottom show the EPI images acquired with the static shim (1), the dynamic shim with the LSQlin (2), the TSVD-H (3) and the Tikh-L (4). The numbers located next to the images show the reduction in solution norm produced by the corresponding methods compared to the unconstrained method.

noise. The Tikh-L was demonstrated to outperform the TSVD-H method in optimising the current use and the shimming quality, and exhibited superior robustness in the shim systems with different current constraints. In the TSVD, the level of regularisation corresponds to the number of singular values truncated. For low-order shim systems, the shim matrix has limited singular values. Simple truncation of singular value leads to abrupt change of shim settings. At the high regularisation levels, the shimming performance using the TSVD is likely to be severely degraded. Therefore, the TSVD-H was sensitive to constraints applied compared to other methods, which may be not ideal for the second-order slice-wise shim current regularisation. A benefit of the Tikh method over the TSVD for low-order shim systems is that it provides more degrees of freedom of the regularisation and allows more flexible control on shimming results, according to findings in **Section 3.1.3**. For high-order shim systems, inclusion of more shim terms in the shim determination could potentially mitigate the effects of over-regularisation using the TSVD method.

An advantage of the L-curve method is its capability of providing near-optimal regularisation that balances the regularisation and the perturbation errors in different imaging conditions. Additional measurement noise may change image conditions and yield slightly different near-optimal solutions. To avoid over-regularisation, the standard deviation constraint can be used here to constrain regularisation results. Further to this, incorporating extra constraints (e.g., the conditioning of the shim matrices and the current limits on different shim channels) can facilitate the selection of optimal regularisation parameters. Besides the L-curve method, some other regularisation parameter search algorithms (e.g. Generalised Cross-Validation [136]) can be used for determining the near optimal regularisation levels. A comparison of the L-curve method and the Generalised Cross-Validation methods would be interesting but beyond the scope of this section.

The Tikh-L yielded not only a reduction of solution norm but also a relatively smoother shim change across different slices than other methods. As discussed in **Chapter 2**, eddy currents scale with the size of shim steps. Even with proper ECC, some residual higher-order eddy field still exhibit and cause spatial and temporal field variations. Shim settings devoid of abrupt shim step is likely to have low eddy current effects. The TSVD-H induces abrupt shim changes and leads to strong eddy currents. This can partially explain why the most significant signal loss was found for slices in the inferior part of the brain (see slice 7 in Figure 3.17).

The regularisation can be used to address the shim degeneracy problem, due to lack of information on the field distribution along the axis associated with degeneracy terms. Modification of the regularisation matrix allows independent control of the regularisation effects on specific shim terms. Applying different weights on specific shim terms could potentially mitigate the shim interaction effects associated with related shim terms.

3.3 Shim step size optimisation

As introduced in **Chapter 2 Section 2**, with pre-emphasis, the shim-induced eddy current effects are largely reduced and the eddy-current-induced fields only deviate from desired values by 1% of the amplitude of shim step. Abrupt shim steps could still drive nonnegligible eddy current effects. This error in generated shim fields will be translated to the shimming results.

Numerous studies have been conducted to investigate methods to further mitigate the eddy current effects in MR images after the hardware calibration, for instance, post-processing [137, 40], tailoring RF pulses, [138, 139], and manipulating gradient

waveforms [140]. However, none of these techniques was designed to mitigate eddy current effects on higher-order shims for dynamic shimming application. Previous studies on reducing shim-induced eddy current effects is using eddy currents that built up from previous pulses to cancel out the expected distortion induced by the subsequent shim switches [141]. However, this method relies on precise calibration of eddy-current-induced fields which is inconvenient for clinical applications.

The aim of this section is to investigate methods to mitigate the eddy current effects on top of the hardware calibration. Eddy current effects are related to the size of the shim step. One way to reduce the shim-induced eddy current fields is to constrain the size of shim steps. Therefore, shim settings free of abrupt shim step between contiguous slices are ideal for dynamic shimming. The **Section 3.2** has demonstrated the superiority of the Tikh-based methods in reducing the amplitude of shim currents by manipulating the regularisation levels. Base on this investigation, we propose a novel optimisation method to constrain the size of shim steps by reducing abrupt shim amplitudes using the Tikh method (Tikh-S), in **Section 3.3.1**. The performance of the Tikh-S was evaluated in comparison with an alternative method of directly constraining shim amplitudes using the LSQlin method (**Section 3.3.2**). The advantages and limitations of the two methods are discussed in **Section 3.3.3**.

3.3.1 Tikh-based shim step size optimisation

The proposed Tikh-S relies on determining the change in shim amplitudes between contiguous slice pairs, and using that information to optimise related shim settings for the shim step size reduction. The method is comprised of two phases: abrupt shim step size determination and shim step size optimisation. The first stage is to identify slice pairs which exhibit changes in shim amplitude exceeding the step size threshold.

In the second stage, the Tikh optimisation was adopted to reduce the amplitudes of the shim settings which cause abrupt shim steps. The method is illustrated step by step in Figure 3.18.

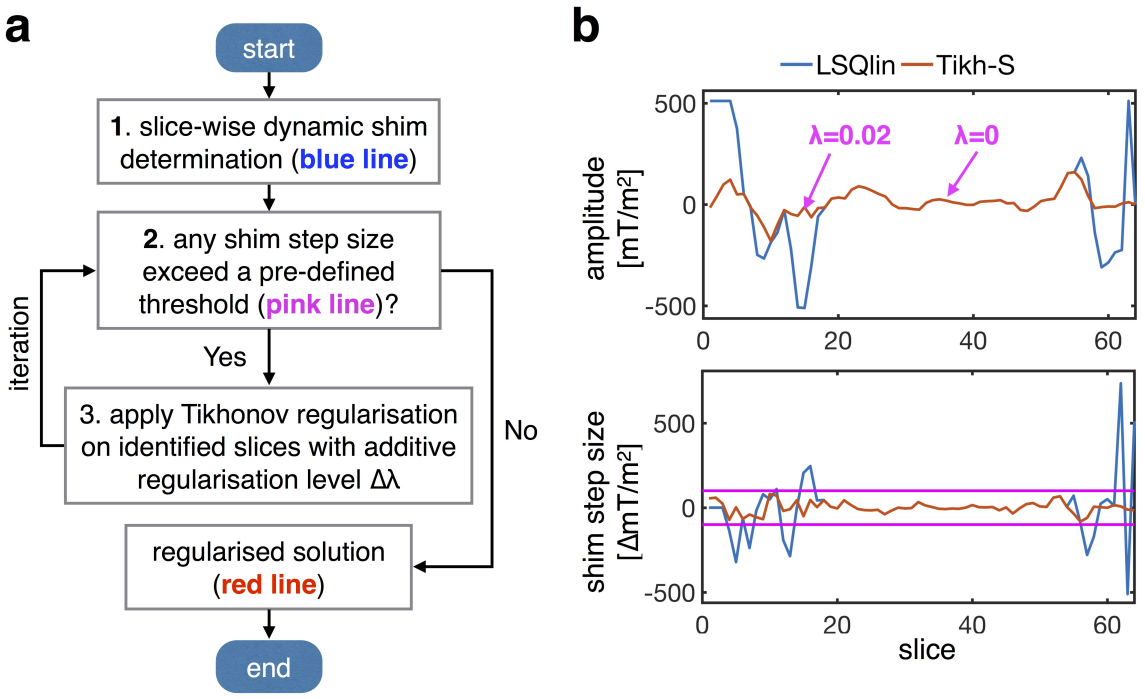


Figure 3.18: **a**: Flow chart of the shim step size optimisation procedure of the proposed Tikh-S method. **b**: Schematic example of amplitude and shim step size of shim settings optimised by the Tikh-S method based on which are calculated using the LSQlin.

Abrupt shim step identification

The identification of abrupt shim steps for specific shim settings was comprised of the following steps

1. Define the constraints on the size of shim steps for each shim term. In the scanner, gradient coils are shielded and well pre-emphasised, whereas eddy currents

induced by second-order shims are relatively significant even with ECC. Therefore we set no shim step size constraints on gradients and strong constraints of Δ 100mT/m^2 on all second-order shims.

2. Calculate the size of shim steps of the target shim settings between all the contiguous slices on each shim term. In practice, the shim settings between the contiguous slices depend on the setups of the MR protocols. Moreover, some sequences require the shim settings to be applied cyclically. Therefore the shim change between the last slice and the very beginning slice were taken into the optimisation.

3. Identify neighbouring slice pairs where the largest shim step size among all shim terms exceeded the threshold as indicated in Step 1.

Shim step size optimisation

In each identified slice pair, one of the shim terms shows the largest shim step size, corresponding to $S_{k=max}$. The strengths of this shim term $S_{k=max}$ on one slice is larger than the other. Reducing the strengths on the shim term $S_{k=max}$ at the slice, where the larger absolute value presents, can mitigate the size of shim steps between slice pairs. Based on this finding, we adopted the following steps to optimise shim shim size:

1. Apply the Tikh method in shim calculation for the slices where the absolute amplitudes on the shim term $S_{k=max}$ are larger than the other.

2. Add extra penalty $\Delta\lambda$ to the regularisation parameters to increase the degree of regularisation on slices required to be optimised. Whilst, slice pairs that meet shim step size constraints remained unchanged.

3. Iterate from Step 2 to 3 to seek for the near optimal regularisation parameters that constrain the size shim steps on all of the shim terms until the specified criteria on shim step size constraints were met.

Evaluation

Besides using the Tikh-S method, the LSQlin method can be used for shim step optimisation (LSQlin-S). The LSQlin-S follows the basic framework of the Tikh-S method. Differently, in Stage 2, the current reduction was achieved by setting lower shim current constraints. The performance of slice-wise dynamic shimming using the Tikh-S and the LSQlin-S methods was evaluated based on the field map database, as introduced in **Chapter 2 Section 2.4.3**.

To quantify the performance of the shim step size optimisation methods, three criteria are used:

1. The number of slices affected by optimisation. In shim step optimisation, constraining the largest shim amplitude affect the calculated shim settings on all shim terms. Thus, new large shim steps may appear on the neighbouring slices. Therefore, optimisation methods that have minimum influences on the original shim settings are preferred.

2. Residual field standard deviation. Assessment of the standard deviation of the residual fields allows estimating the shimming quality using optimisation methods.

3. Computational time. The shim step size optimisation is an overhead to the routine slice-wise shim calculation procedures. Shorter computational time is

preferred for clinical scans.

The order of slice acquisition affects the size shim steps across slices and the shim step size optimisation results. There are two conventionally used slice acquisition schemes:

- The ascending/descending slice acquisition where slices are acquired sequentially.
- The interleaved slice acquisition where slices at odd indices are collected first then even indices, in order to avoid cross talk between consecutive slices.

The performance of the two methods was assessed with the sequential and the interleaved acquisition schemes, respectively.

3.3.2 Simulation of shim step optimisation methods

Simulation on an example subject

Figure 3.19 shows the Z2 shim amplitudes determined from the LSQlin and the size of shim steps on this shim term, together with the outcomes of using the Tikh-S and the LSQlin-S optimisation methods. For the majority of slices, field changes are intrinsically smooth, thus the size of shim steps is relatively small. Large shim step size was observed in the regions associated with small ROIs and strong field distortions, like the most inferior and superior parts of the brain. Additionally, large shim step size occurred at slices associated with the transition regions from the pulsed

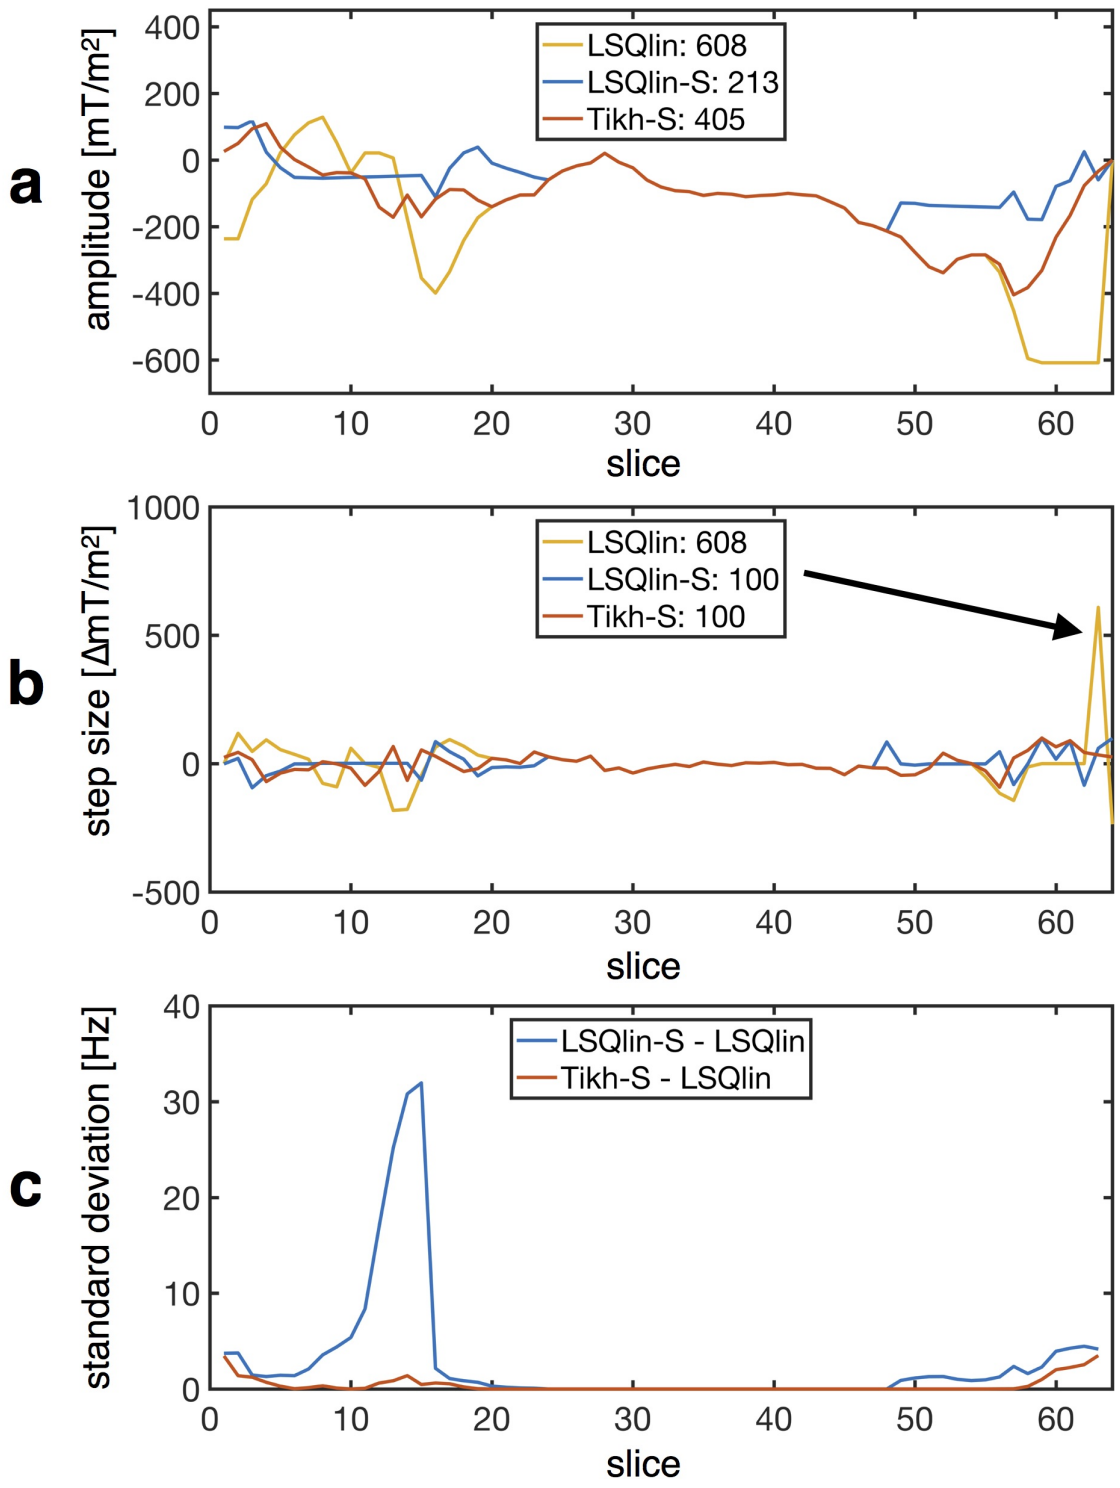


Figure 3.19: Comparison of the amplitude of shim settings, the size of shim steps and the residual B_0 standard deviation for the Z2 shim resulting from the LSQlin (yellow), the LSQlin-S (blue), and the Tikh-S (red) using a sequential slice acquisition method.

shim to zeros. Taking the shim settings for the Z2 shim term for example, the peak shim step size produced by the LSQlin reached up to 608 mT/m^2.

After optimisation with the Tikh-S, the maximum shim step size was reduced to $100mT/m^2$ and only a few slices show standard deviation degradation. Compared to the Tikh-S, the LSQlin-S produced lower maximum shim step size across all the shim terms, but affected shim settings of more slices and significantly degraded the corresponding residual field standard deviation, especially for slice 17. The computational time for optimising the size of shim steps across all the slices of this example subject is 0.73 s with the Tikh-S. However the LSQlin took 25 s for this optimisation.

In the case of the interleaved slice acquisition, as shown in Figure 3.20, more shim step sizes that exceeded the required thresholds were observed. Moreover, the size of shim steps between every other slice was larger than that of adjacent slices. In the interleaved acquisition, though the Tikh-S showed better shimming performance than the LSQlin-S, both methods delivered higher residual field standard deviation compared to the sequential mode.

Simulation on a field map database

Figure 3.21 compares the maximum shim step size across second-order shim terms (a, b), the residual field standard deviation (c, d) and the number of slices affected (e, f) resulting from the LSQlin-S and Tikh-S methods with the sequential slice acquisition mode. As illustrated in Figure 3.21a and 3.21b, both the Tikh-S and LSQlin-S can effectively reduce the peak shim step size of given shim settings down to the user-defined threshold of $100mT/m^2$ on each shim term. As shown in Figure 3.21c and 3.21d, the shim step size optimisation methods trade some standard deviation for reduced shim step sizes. The Tikh-S increased the mean standard deviation by

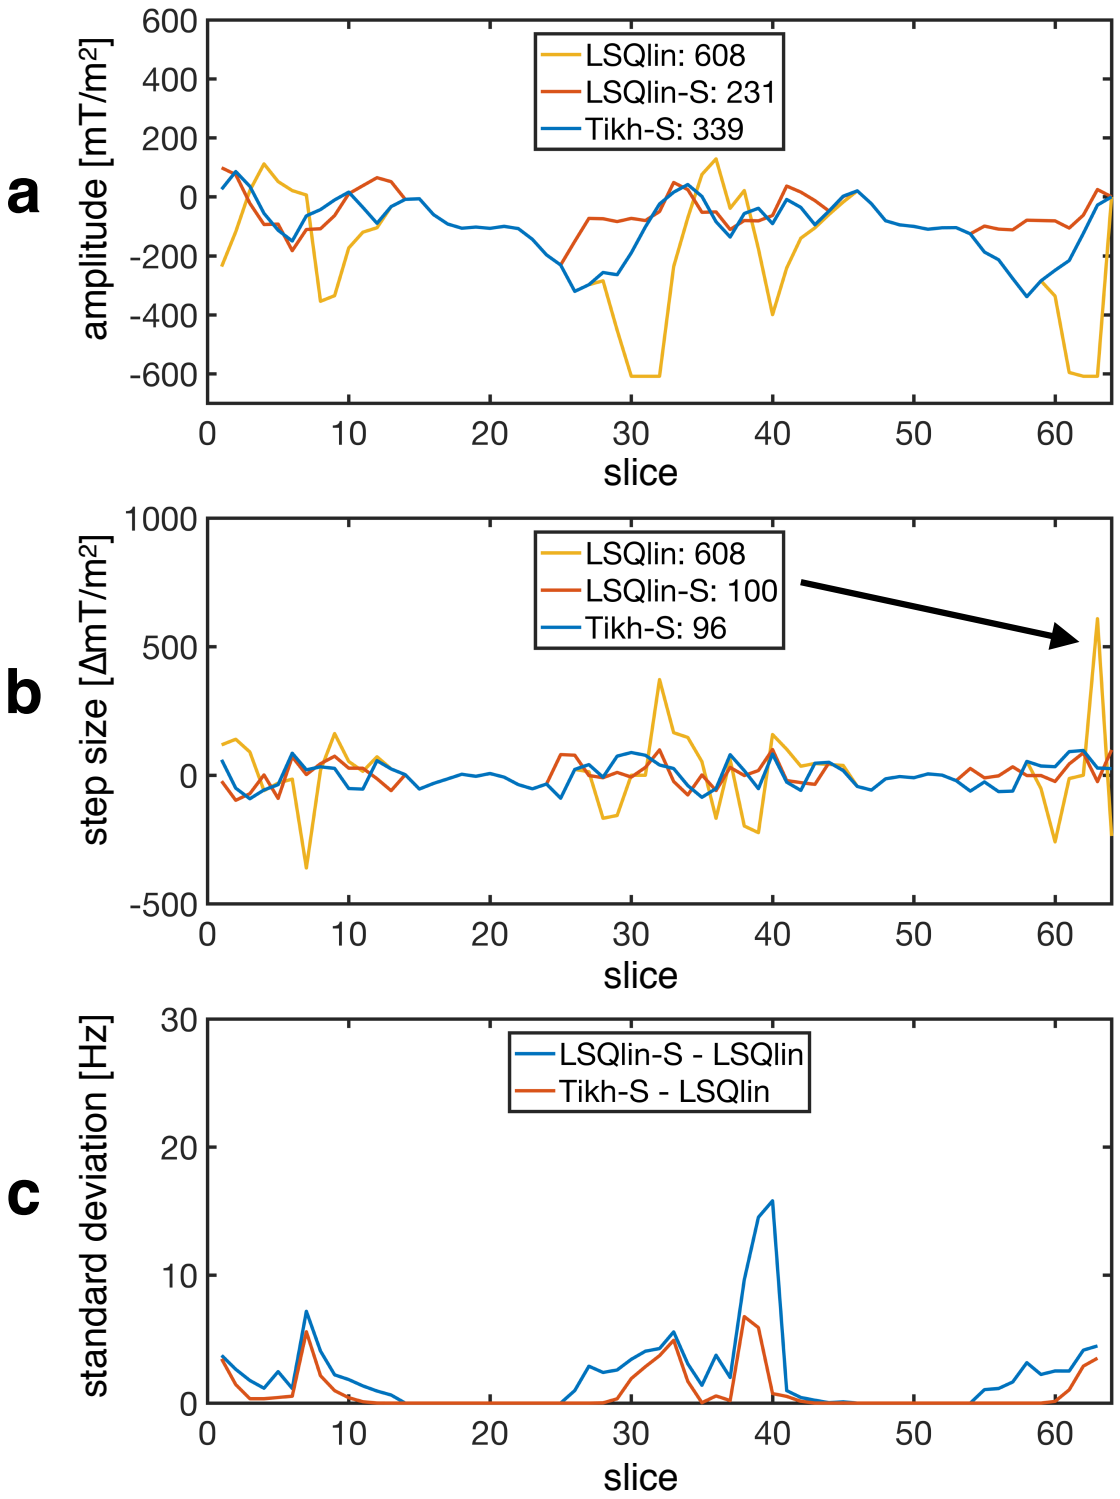


Figure 3.20: Comparison of the amplitude of shim settings, the size of shim steps and the residual B_0 standard deviation for the Z2 shim resulting from the LSQlin (yellow), LSQlin-S (blue), and Tikh-S (red) using an interleaved slice acquisition method.

0.4Hz on average on all of the slices for each subject in the 143 field map database. In contrast, the LSQlin degraded the standard deviation by 3 Hz. Moreover, the Tikh-S affected the shim settings on fewer slices than the LSQlin-S.

The outcomes of the Tikh-S and the LSQlin-S methods on optimising shim settings with an interleaved optimisation are displayed in Figure 3.22. Similar results were observed with the interleaved slice acquisition mode.

3.3.3 Discussion

We have demonstrated the feasibility of the Tikh-S and the LSQlin-S methods in mitigating the size of shim steps based on LSQlin method. The two shim step size optimisation methods increase the residual field standard deviation in return for reduced shim amplitudes and shim step size. The Tikh-S has superiority over the LSQlin-S. First, the Tikh-S reduced shim step size without degrading much standard deviation and affecting as many slices as the LSQlin-S. Second, the regularisation-based Tikh-S method could improve the conditioning of shim calculation and avoid the instability of shimming results.

In the implementation, we only constrained the size of shim steps. However, the single criterion on the size of shim steps is not feasible and some slices yielded much worse standard deviation of the residuals. To retain a good level of field homogeneity, a secondary criterion on the residual field standard deviation can be used.

The performance and the speed of the two methods depend on search steps. In order to converge to a local optimum, the amount of reduction in shim amplitudes in each step is preferred to be small. However, setting smaller search steps prolongs the computation of the optimisation. Reducing the search steps has different influences on the

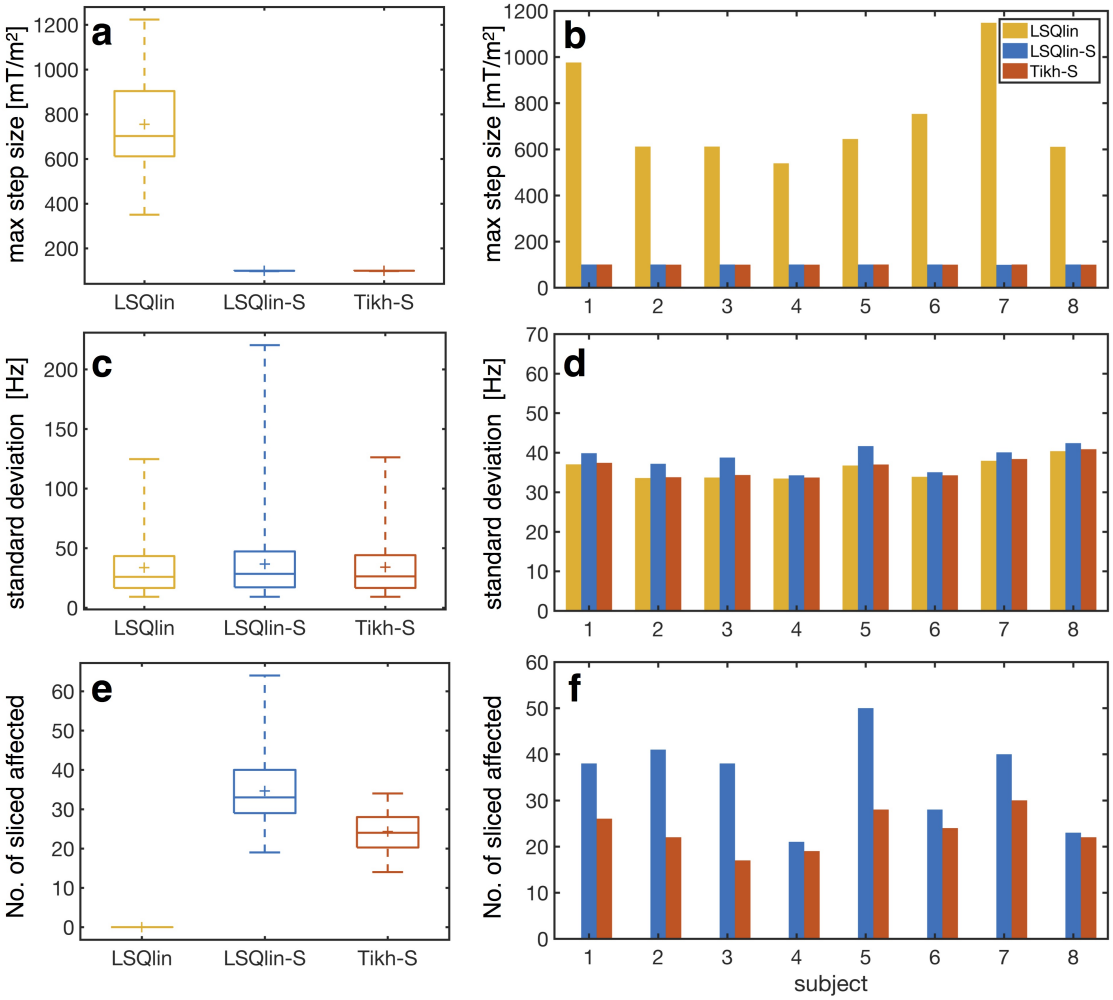


Figure 3.21: Comparison of the Tikh-S and LSQlin-S shim step size optimisation methods based on the LSQlin method with a sequential slice acquisition scheme. the maximum shim step size on second-order shim terms (**a, b**), the average residual B_0 field standard deviation (**c, d**), along with number of slices affected by the LSQlin-S and the Tikh-S (**e, f**). The boxplots (**a, c, e**) indicate the corresponding measures on all the 143 subjects in the field map database. The bar charts (**b, d, f**) indicate the measures for 8 example subjects in the database.

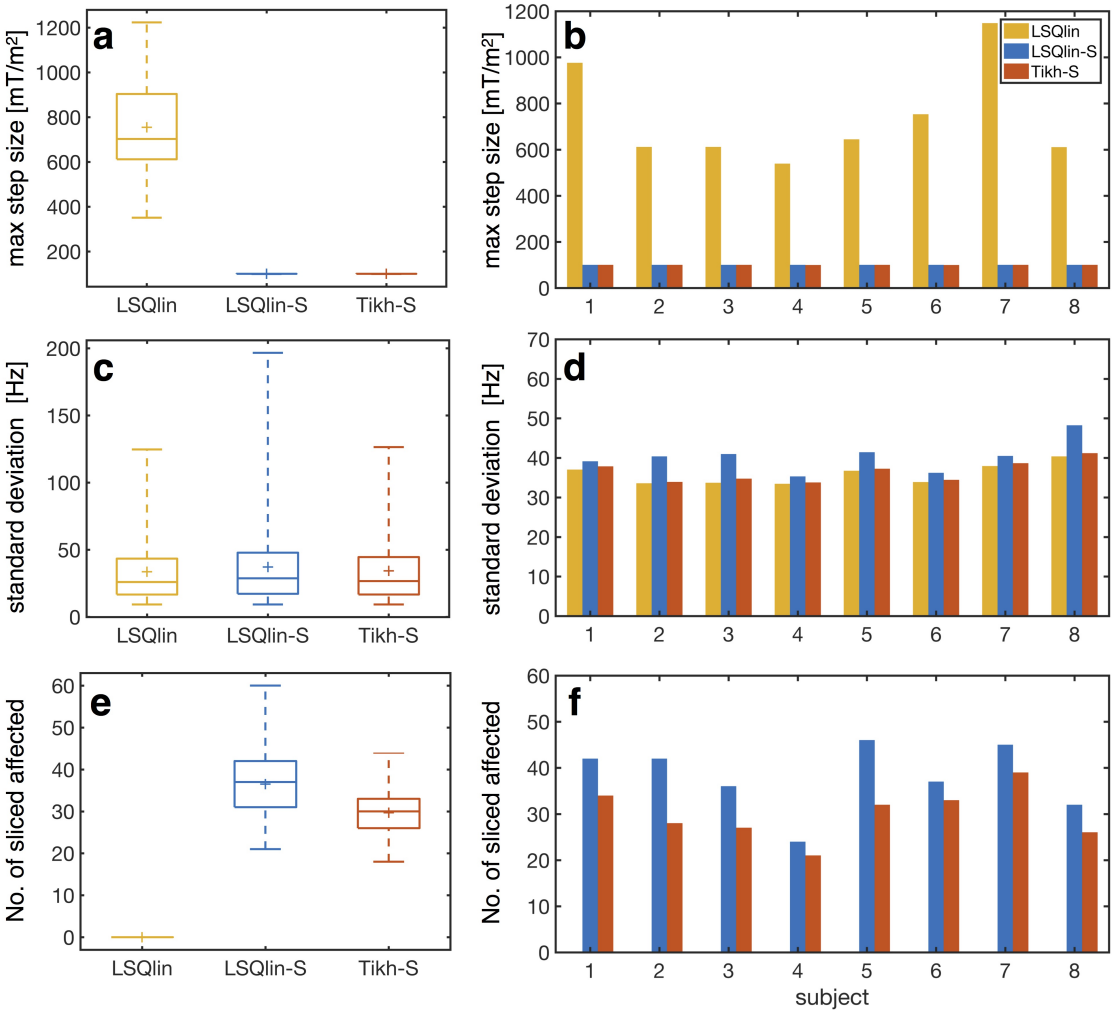


Figure 3.22: Comparison of the Tikh-S and LSQlin-S shim step size optimisation methods based on the LSQlin method with an interleaved slice acquisition scheme. the maximum shim step size on second-order shim terms (**a, b**), the average residual B_0 field standard deviation (**c, d**), along with number of slices affected by the LSQlin-S and the Tikh-S (**e, f**). The boxplots (**a, c, e**) indicate the corresponding measures on all the 143 subjects in the field map database. The bar charts (**b, d, f**) indicate the measures for 8 example subjects in the database.

optimisation results of the two methods. The LSQlin-S method involves iteratively searching for solutions that produce the shimming quality close to the unconstrained pesudoinverse within the currents limits. Increasing the number of search steps incurs the multiplication of the iterative shim calculation. In contrast, the Tikh-S has advantages of computing solutions that satisfy the amplitude constraints in an analytical way, which is inherently faster than the LSQlin. Increasing the amount of search steps improves the optimisation results of the Tikh-S and does not significantly increase the computational time.

A possible alternative method to optimise the size of shim steps is rearranging the order of slice acquisition. This method can be achieved by placing sets of shim settings between slice pairs associated with large shim step size. This optimisation could potentially mitigate the occurrence of abrupt shim changes across shim settings without hindering shimming quality. However this method has some drawbacks. First, reordering slices leads to inconsistent RF pulse profile for different slices. Second, in some cases, we can hardly find any combination of shim settings that yields lower shim step size for slices with the large shim values, for example, the shim step peak as shown in Figure 3.19b.

3.4 Effects of using various cost functions in dynamic shimming

Thus far we have only considered shimming targeting to minimise the standard deviation of the residual field. From the B_0 shimming perspective, the standard deviation metric is a measure of the variation of field offsets across pixels at different locations. As geometric distortion is linearly related to field offsets as indicated in **Chapter 1**, the minimisation of B_0 field standard deviation can be translated to a reduction

of geometric distortion artefacts. In practice, in low-order shim systems, the brain field distribution is relatively high in spatial order. The standard deviation optimisation often optimises regions of high field distortions, regions near sinuses, in order to minimise the residual field standard deviation. However, this optimisation may compromise B_0 field homogeneity in other regions of the brain. It will be a big problem for the whole brain shimming if the ROIs are located in those regions. A possible solution to this problem is using a weighting mask to give priority to regions in different parts of the brain for shimming [80].

Besides the standard deviation optimisation, signal intensity is an alternative measure that can be optimised in shimming. According to the signal intensity model as introduced in **Chapter 1 Section 3**, the image intensity can be modelled as a function of field gradients. In general, the standard deviation optimisation improves the field homogeneity and concurrently reduces field gradients. Additionally, the signal is also affected by some sequence related factors like echo time and the bandwidth in the phase encoding direction, for instance, severe signal loss occurs if the echoes are shifted out of the acquisition window. Therefore, in some cases, the reduction in standard deviation does not necessarily produce higher signal, if the intrinsic field off-resonance is still relatively high in spatial orders after shimming. Instead of minimising field inhomogeneity, maximising signal intensity can also be used as a target function for shimming.

For fMRI applications, the sensitivity of BOLD signal could also be used as a target for shimming [142]. The increase in popularity of whole brain resting state fMRI motivated the investigation into techniques to improve BOLD signal. In fMRI, BOLD signal is established from signal contrast induced by changes in T2* during blood oxygenation and deoxygenation due to neural activities. The sensitivity of BOLD signal in the EPI images to changes in T2* is referred to as BOLD sensitivity. A

realistic model of BOLD sensitivity is proposed by Deichmann R. [52, 53]. He models the BOLD sensitivity based on a measured field map and target EPI sequence parameters (e.g., echo time and bandwidth). Previous studies have demonstrated that using the BOLD sensitivity metric as the cost function of static shimming could yield improved BOLD signal for fMRI [129].

This section aims to evaluate the effectiveness and potential benefits of using different optimisation method in dynamic shimming through computational simulation using the field map database. This section is structured as follows: **Section 3.4.1** introduces the basic models to quantify signal intensity and BOLD sensitivity. **Section 3.4.2** presents the procedures to implement dynamic shimming that optimise for different cost functions. Simulation results of different shimming methods are comparably studied in **Section 3.4.3**.

3.4.1 Signal loss and BOLD sensitivity model

As described in **Chapter 1 Section 1.3**, the loss of signal intensity is related to the intra-voxel dephasing effects. The amount of dephasing effects depends on the local field gradients across voxels. Therefore, signal loss can be readily modelled by a function of local field gradients present in each voxel and image voxel size, which is given by

$$I_{loss} = \sqrt{\left(\frac{\partial B_0}{\partial x}\right)^2 D_x^2 + \left(\frac{\partial B_0}{\partial y}\right)^2 D_y^2 + \left(\frac{\partial B_0}{\partial z}\right)^2 D_z^2}, \tag{3.16}$$

where D_x, D_y, D_z indicate the sizes of image voxels in the frequency encoding, the phase encoding and the slice selective directions, respectively. $\frac{\partial B_0}{\partial x}$, $\frac{\partial B_0}{\partial y}$, and $\frac{\partial B_0}{\partial z}$ are the first order intra-voxel field derivatives in the corresponding directions. Thus, minimisation of the signal loss function can improve the image intensity in general.

The model of EPI image intensity is explicitly described in **Chapter 1 Section 1.3**, which depends on factors, including the local field gradients and the local echo times. Building on the image intensity model, the BOLD sensitivity is modelled by [51]

$$BS = I_{pred} TE_{eff}, \tag{3.17}$$

where the term TE_{eff} indicates the effective echo time, I_{pred} indicates the predicted image intensity. The equation indicates that BOLD imaging favours a long echo time which lead to more BOLD sensitivity, though signal attenuates with a longer echo time.

3.4.2 Implementation of different models in shimming

The performance of shimming algorithms targeting for optimising three costs, the standard deviation (LSQ-STD) (1), the gradient-related signal loss (LSQ-GRAD) (2), and the BOLD sensitivity (LSQ-BOLD) (3), were comparatively studied for slice-wise dynamic shimming. Simulation of these optimisation methods was performed in MATLAB.

The estimation of signal loss and BOLD sensitivity relies on accurate determination of local field gradients across each image voxel from field maps. To get good approximation of local field gradients, field gradients for each voxel in three imaging coordinate directions were calculated from a set of three voxels along each imaging direction centred on the target voxel through a linear fit to the field offsets. In practice, the acquired field maps are inherently noisy due to measurements errors. This problem can significantly affect calculated local field gradients. In order to minimise the influence of noise on measured field gradients, a 3D Gaussian filter with a kernel

size of 2 mm was used to smooth out noise in field maps. Field map pre-processing (brain extraction, phase unwrapping, noise filtering) was performed using FSL.

The EPI-based signal intensity and the related BOLD sensitivity map estimation depends on the target EPI sequence settings. Simulations of the two models were performed on the field maps in the 143 subjects database, as described in **Chapter 2, Section 2.4.3**, targeting for optimising the EPI acquisition (voxel size=$2 \times 2 \times 2$ mm^3, TE=25 ms, bandwidth=1578 Hz, echo spacing=0.72 ms). In order to prevent severe geometric distortion, the residual B_0 field standard deviation was limited to not greater than 150% of the LSQ-STD method and the mean gradient in the phase-encoding direction was not greater than 1.5 Hz/pixel (5% compression of geometric distortion).

Signal loss and BOLD sensitivity maps are nonlinearly related to the field distribution. To find the shim settings optimised for different field quality measures, a Matlab embedded nonlinear optimisation algorithm was implemented. As nonlinear algorithms are sensitive to starting values, the global search interior-point method was implemented with multiple start points. In our implementation, for BOLD optimisation, the start values were set to the shim settings determined from the STD and the GRAD optimisation.

3.4.3 Comparison of different optimisation methods

The shimming quality was quantified using criteria of the residual field standard deviation (1), the predicted EPI image intensity (described in **Chapter 1 Section 1.3**), (2) and the predicted BOLD sensitivity (3) derived from the residual field maps after shimming using the LSQ-STD, LSQ-GRAD, and LSQ-BOLD methods on all of the

slices for each subject in the 143 field map database. Additionally, the nonparametric Wilcoxon signed-rank tests were carried out to compare the performance of shimming using the three methods.

Evaluation on a single subject

Figure 3.23, 3.24, and 3.25 show the simulated maps of residual B_0 field (a), image intensity (b), and BOLD sensitivity (c) with slice-wise shimming based on the LSQ-STD, the LSQ-GRAD, and the LSQ-BOLD method based on a field map acquired from an example subject in the database. As illustrated in Figure 3.23, the LSQ-STD method yielded lowest field inhomogeneity over all the slices - reflected in the reduced red pattern. However, for slice 25, the native field distribution in the inferior frontal lobe and temporal lobes, near the sinuses and ear cavities, exhibited higher spatial orders than other regions of the brain. The LSQ-STD method reduced the field inhomogeneity in the regions close to the sinuses while exacerbating field distortion in other regions (e.g. regions near ear cavities).

The corresponding image intensity maps with different shimming methods are shown in Figure 3.24. The white regions are associated with voxels shift outside acquisition window or complete signal loss where spins suffer from strong field gradients. The regions coloured with red correspond to high image intensity. The results show that the LSQ-GRAD optimisation consistently yielded good image intensity over all the slices. The LSQ-STD method was best in some regions, and worst in other regions. The LSQ-BOLD method provided relatively low image intensity compared the LSQ-STD and the LSQ-GRAD. This can be understood as that the reduced standard deviation in some regions (e.g. the frontal lobe) cannot be translated to gains in signal intensity.

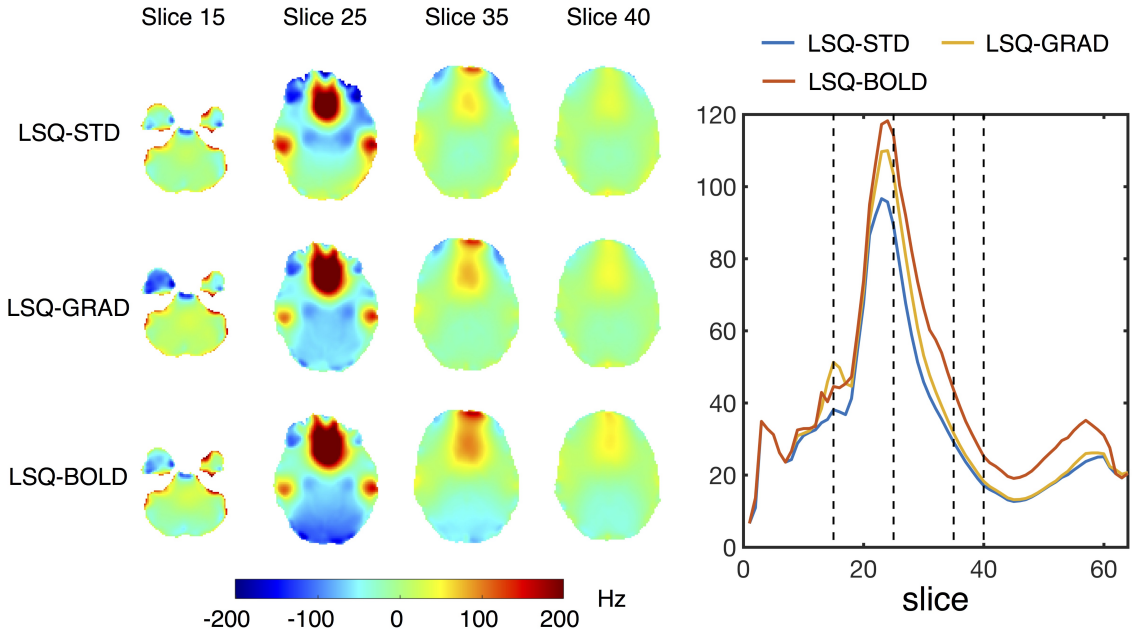


Figure 3.23: The residual B_0 field map resulting from slice-wise shimming using the LSQ-STD, LSQ-GRAD, and LSQ-BOLD methods on an experimentally acquired field map. TThe plot besides the images shows the corresponding the slice-wise residual field metric after shimming using the LSQ-STD (blue line), the LSQ-GRAD (yellow line), and the LSQ-BOLD (red line) methods.

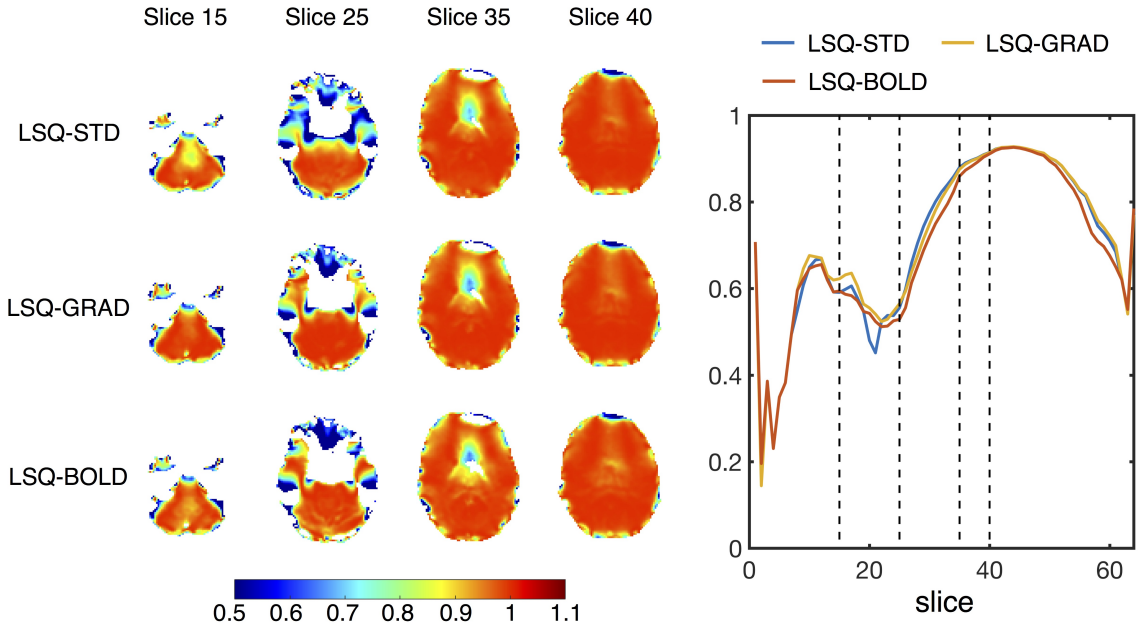


Figure 3.24: The image intensity map resulting from slice-wise shimming using the LSQ-STD, LSQ-GRAD, and LSQ-BOLD methods on an experimentally acquired field map. The plot besides the images shows the corresponding the standard deviation of slice-wise residual field after shimming using the LSQ-STD (blue line), the LSQ-GRAD (yellow line), and the LSQ-BOLD (red line) methods.

Figure 3.25 shows the simulated BOLD sensitivity maps using these optimisation methods. The green and blue regions of BOLD sensitivity maps reflect regions of low BOLD sensitivity. In the red and yellow regions, the BOLD signal is relatively more sensitive to changes of local frequencies. The LSQ-BOLD method produced improvements in the simulated BOLD sensitivity measures across all of the slices, though much degrading the residual field standard deviation. For slice 25, the LSQ-BOLD method optimised the regions in the posterior part of the brain while causing significant signal loss in the frontal lobe.

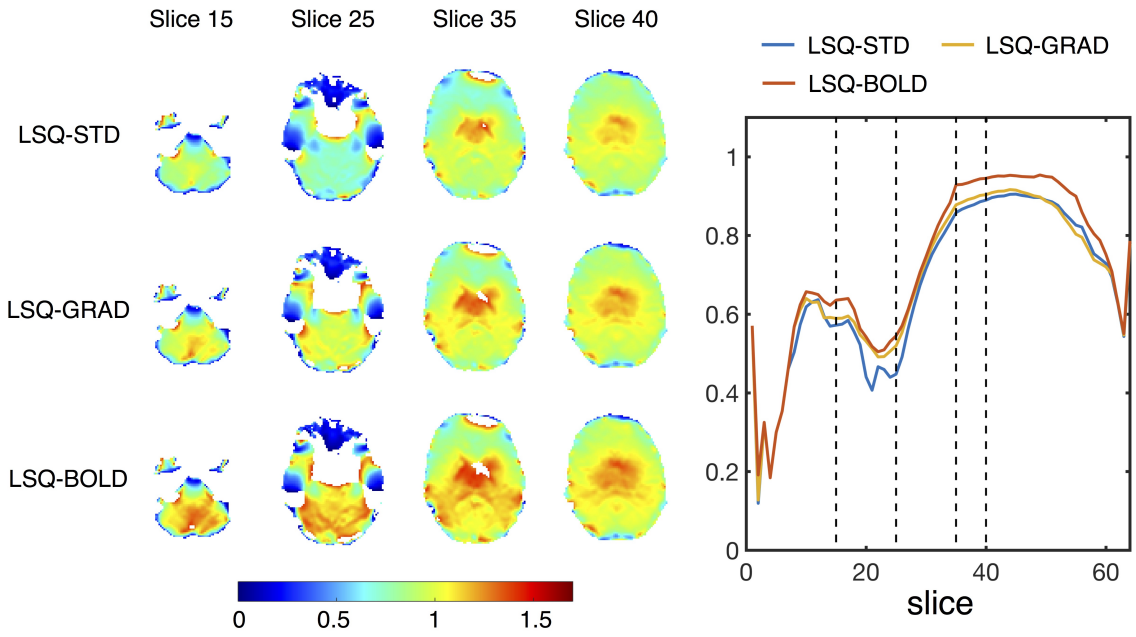


Figure 3.25: The BOLD sensitivity map resulting from slice-wise shimming using the LSQ-STD, LSQ-GRAD, and LSQ-BOLD methods on an experimentally acquired field map. The plot besides the images shows the corresponding the slice-wise BOLD sensitivity metric after shimming using the LSQ-STD (blue line), the LSQ-GRAD (yellow line), and the LSQ-BOLD (red line) methods.

Evaluation on a field map database

Figure 3.26 compares standard deviation (a), image intensity (b), and BOLD sensitivity (c) with slice-wise shimming using the LSQ-STD (blue boxplot), the LSQ-GRAD (yellow boxplot), and the LSQ-BOLD (red boxplot) methods for each subject in the 143 field map database. The nonparametric Wilcoxon signed-rank tests indicated there were significant differences between the performance of shimming using the three methods ($P<0.001$). The results confirm the findings above. The LSQ-STD method yielded the lowest residual field standard deviation, 30Hz, and comparable image intensity as the LSQ-GRAD method. However, it did not perform better than the LSQ-BOLD method in terms of BOLD sensitivity. The LSQ-BOLD method trades 10 Hz residual field standard deviation, compared to the LSQ-STD method, for the improved BOLD sensitivity. The LSQ-GRAD is a compromise of these two methods.

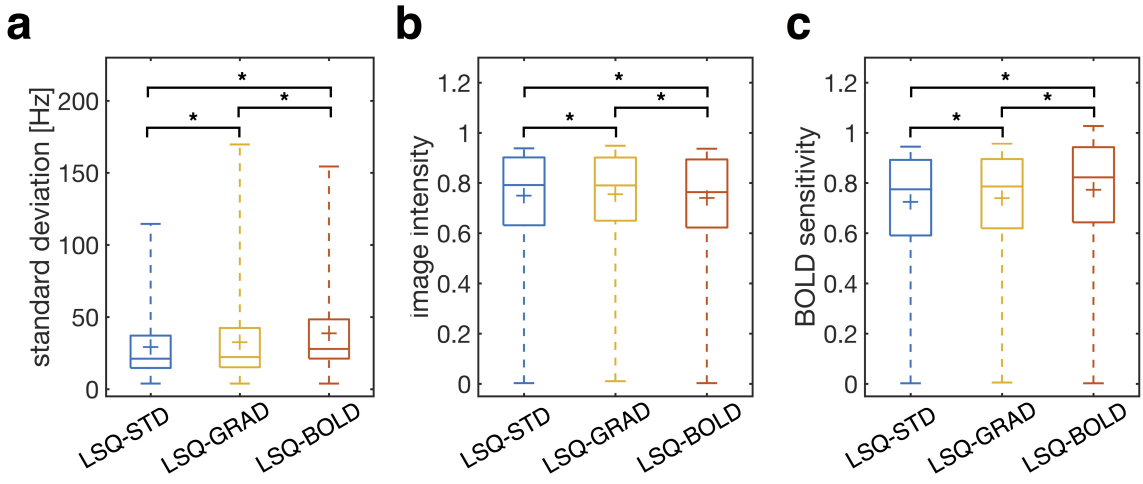


Figure 3.26: The slice-wise residual field standard deviation (a), signal intensity (b), and BOLD sensitivity (c) after shimming using the LSQ-STD, the LSQ-GRAD, and the LSQ-BOLD methods on all of the slices for each subject in the 143 field map database. The nonparametric Wilcoxon signed-rank tests resulted in $P<0.001$.

3.4.4 Discussion

We have evaluated the performance of the LSQ-STD, the LSQ-GRAD, and the LSQ-BOLD methods in slice-wise shimming. In general, the LSQ-STD method has tendency to homogenise regions of strong field distortions. For slices associated with the field inhomogeneity on a large spatial scale, minimisation of standard deviation also provides comparable results as the LSQ-GRAD method. However, for slice close to the frontal lobe, the LSQ-STD achieved minor improvements in BOLD sensitivity and image intensity. The LSQ-BOLD method optimises regions where the reduction of field inhomogeneity can be translated to gains of sensitivity of EPI to frequency shift, avoiding to over-optimise regions that hardly get improvement in signal. However, the gains in the BOLD sensitivity come at the cost of field homogeneity. The LSQ-GRAD method is a midway between the two methods.

There are some drawbacks of the LSQ-BOLD and the LSQ-GRAD methods.

First, nonlinear optimisation usually suffer from local minimum traps and is sensitive to starting values due to the lack of the global knowledge of the parameter space. To circumvent the local minimum challenges, careful selection of starting values and global minimum search methods are required. If the number of iterations is limited, the solutions will arrive at nonoptimal settings.

Second, the LSQ-BOLD method, though maximising the overall BOLD sensitivity, induced signal attenuation over a large region, as shown in Figure 3.26b. The intrinsic signal attenuation in MR image is unrecoverable which post challenges to post-processing. To counteract the effects, more weighting can be placed on regions where signal can be easily recovered without much degrading the overall BOLD sensitivity. Alternatively, constraints on signal intensity can be coupled in BOLD sensitivity optimisation method to prevent the occurrence of much signal loss.

Third, optimising signal intensity and BOLD sensitivity involve using sophisticated nonlinear optimisation methods for optimal solution search. The use of the nonlinear optimisation algorithms prolongs the shim current determination time, compared to the LSQ-STD method. The long shim calculation time is not ideal for shimming. To cut down the time spent in shim calculation, some method can be used, for instance, reducing the number of iterations in the optimisation, and selecting good start points for fast shim solution search. Alternatively, as shown in Figure 3.23, 3.24, and 3.25, different optimisation methods yielded comparable shimming results for some slices (e.g., slice 1$\sim$10). Thus, the LSQ-GRAD and the LSQ-BOLD methods in conjunction with the LSQ-STD method for bottom slices could yield fast shim calculation while not degrading much optimisation results.

Finally, the shim settings of the LSQ-BOLD method depend on the setup of EPI sequences (e.g., flip angles and image resolution). Usually, during MR scans, the sequence parameters are required to be optimised for different imaging conditions [143]. Thus, the LSQ-BOLD shim calculation is needed to be rerun each time which is time-consuming. Thus, using some routine EPI protocols for the BOLD shimming could reduce the needs of recalculating the shim currents.

Chapter 4

Template-based field map prediction for static and dynamic B_0 shimming

In typical MRI protocols, time is spent acquiring a field map to calculate the shim settings for best image quality. Although dynamic shimming provides a more homogeneous field than static shimming, it requires additional time to acquire a high-resolution field map ($\sim$ 4 minutes) for calculating shim settings. For many applications, particularly in the clinical setting, cutting exam time is of great importance. To shorten the time spent on the field mapping stage, instead of acquiring a field map, it is possible to predict the subject-specific field inhomogeneity by incorporating brain structural information with a model of field distortion in the brain. **Chapter 4** presents a fast template-based field map prediction method that yields near-optimal shims without measuring the field. The template-based prediction method uses prior knowledge of the B_0 distribution in the human brain, based on a large database of field maps acquired from different subjects, together with subject-specific structural information from a quick localiser scan. This chapter is comprised of two sections. **Section 4.1** presents the performance of static shimming using a fixed general field map template in comparison to a range of potential fast shimming methods. **Sec-**

tion 4.1 has been published as a paper in Magnetic Resonance in Medicine [144]. In **Section 4.2**, we propose adaptive field map prediction methods to assess whether extra information on the brain geometry and subject characteristics could yield a more accurate prediction of the field distribution for static and dynamic shimming.

4.1 Template-based field map prediction for static B_0 shimming

Determination of shimming settings for homogenising magnetic field relies on accurate mapping of the field distribution in the brain. As introduced in **Chapter 1**, the field mapping methods can be roughly classified into three categories: 3D volume based field mapping, projection-based field mapping and field map modelling. The 3D volume based methods rely on extracting the phase change information from 3D images. Different sequences have different effects on the resulting field maps. Field mapping using the conventional GRE based sequences is robust to distortion but time-consuming. EPI-based field map acquisition is efficient in time, however it suffers from signal loss and distortion. The bSSFP-based acquisition is not suitable for measuring the field inhomogeneity with high spatial orders due to its high sensitivity to dephasing artefacts. Instead of acquiring 3D field distribution for shimming, the projection-based methods determine optimal shim settings based on sparse sampling of the field distribution information along some k-space line projections. Although projection-based methods are fast, they provide limited spatial field distribution information. Thus, the projection-based methods are not suitable for dynamic shimming applications. The third types of methods are field map modelling. These methods model the subject-specific susceptibility induced fields based on the anatomical structure of an individual subject. This method requires an extra acquisition of the subjects' structural information which is not convenient for clinical applications.

A typical static B_0 shimming procedure for a human head can take approximately 30 s; however, the time for shimming will be scaled if the field map needs multiple acquisitions to accurately characterise the field. This is particularly the case if the subject is poorly shimmed to start with. In this section, we propose a rapid template-based field map prediction to estimate field inhomogeneity across the brain in a short timeframe, which can reduce the time spent on the field map acquisition at the shimming stage. The prediction method uses prior knowledge of the B_0 field distribution in the human brain based on a large database of field maps, together with the basic subject-specific structural information from a localiser scan. In clinical MR studies, a localiser is typically acquired for slice positioning, so the proposed method adds no time to the routine MR procedure. The performance of static shimming using the predicted field map method was evaluated on 143 subjects in simulation and 7 subjects in experiment.

4.1.1 Field map template generation and prediction

The field map prediction method relies on determining a field map template, and using that to estimate a particular subject's field map. The method therefore consists of two main stages: field map template generation and subject-specific field map prediction. The first stage is to determine a template of the characteristic susceptibility-induced field distribution in the brain, based on measurements from a wide range of subjects. In the second stage, the field map for a specific subject is predicted based on the template together with a quick structural scan (e.g. localiser). The method is illustrated step by step in Figure 4.1.

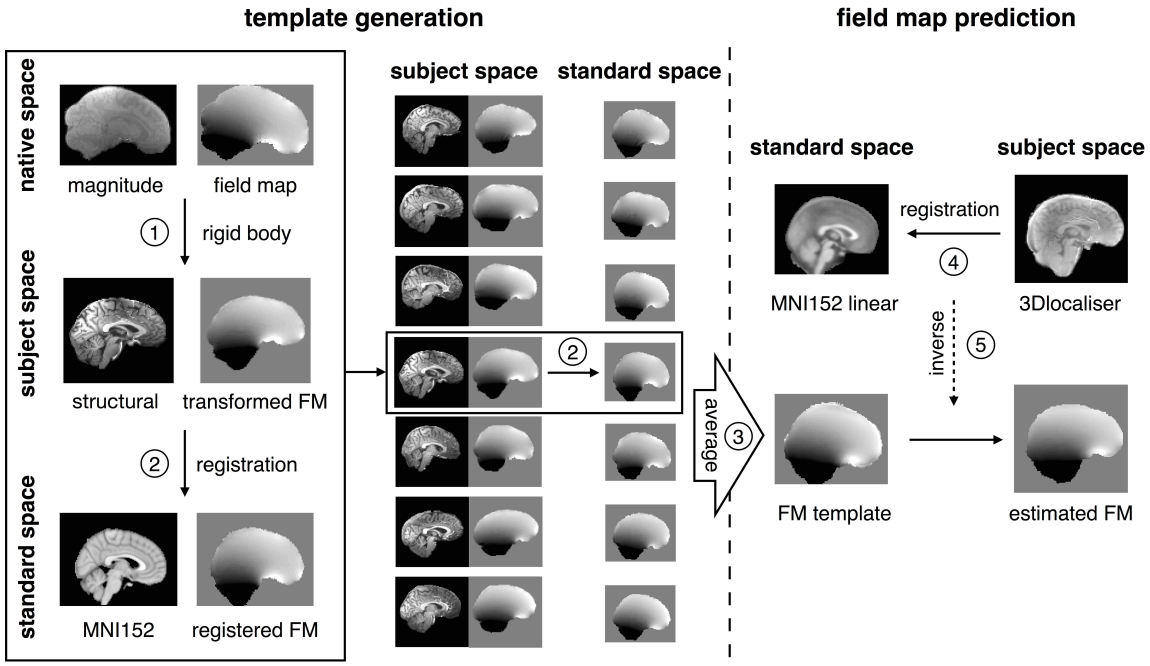


Figure 4.1: Procedures for template-based field map prediction, which consists of two main stages: template generation and field map prediction. The field map template generation is conducted before the experiment. For each subject in the database, (1) the magnitude image of the field map acquisition and the field map are aligned to the T1-weighted image and then (2) the aligned field map is warped to standard space by applying the registration transformation from the T1-weighted image to the MNI (Montreal Neurological Institute) template. (3) The field map template is obtained by averaging all the warped field maps in standard space. The field map prediction is conducted on the day of scan. (4) When the quick localiser is acquired, it is registered to standard space. Following this, (5) the generated inverse registration is applied to the general field map template to get the estimated field map corresponding to the scanned subject.

Field map template generation

Using the field map database described in **Chapter 2 Section 2.4.3**, a characteristic field-distribution template used for field map prediction was generated as follows. First, each T1-weighted structural image and the magnitude image from the field map acquisition were brain-extracted using BET [106]. Then the magnitude image from the field map data was aligned to the T1-weighted image using the rigid-body registration in FLIRT (FMRIB's Linear Image Registration Tool) [145], and the resulting transform was applied to the field map. Second, the T1-weighted image was moved from subject space to standard space, as defined by the MNI (Montreal Neurological Institute) template [146], using FLIRT with 12 degrees of freedom. The corresponding field maps were warped to standard space by applying the registration transformation obtained from the T1-weighted image. The reason for using the T1-weighted structural image as an intermediate for the registration to standard space, instead of just the field map magnitude image, is that T1 images acquired in the experiments cover the entire brain, thereby providing more structural information for the registration. In the final step, the field map template was generated by averaging all of the registered field maps in standard space.

Field map prediction

The subject-specific field distribution is estimated based on the field map template and an acquired structural image. For an individual subject, the structural image was brain-extracted and resampled down to 2 mm isotropic, equivalent to the MNI standard images, to speed up computations of the further processing steps. The resampled structural was registered to standard space with FLIRT. Finally, the predicted field

map was obtained by transforming the field map template to subject space using the inverse of the structural-to-standard space registration transformation.

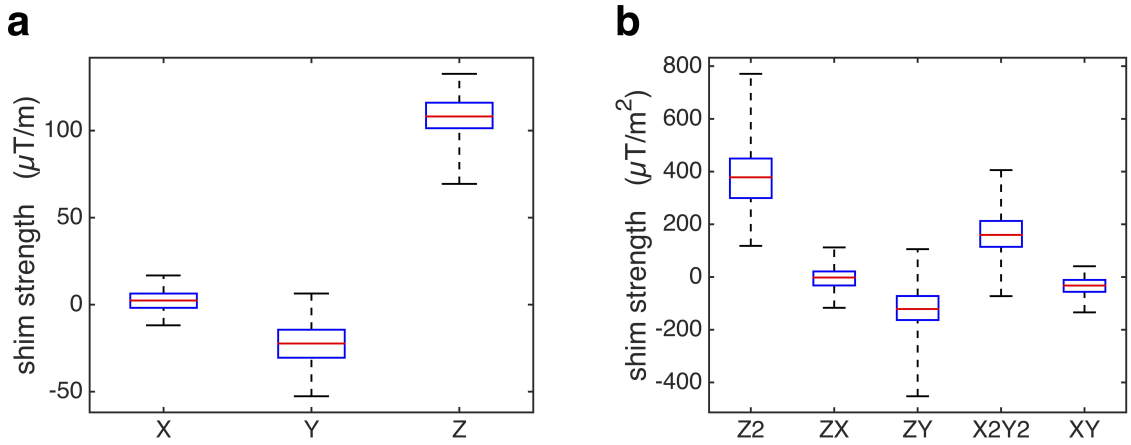


Figure 4.2: First-order (a) and second-order (b) shims calculated based on measured field maps for all 143 subjects in the database. For the box plot of each shim term, the five line markers from bottom to top represent the minimum, 25th percentiles, mean, 75th percentiles, and maximum of the shim value distribution. The x-, y-, and z-axis correspond to the left-right, anterior-posterior, and superior-inferior directions of the subject coordinates, respectively.

Figures 4.2a and 4.2b show the range and the average of individual shims settings of the first-order and second-order terms, respectively, calculated based on the measured field map, for all 143 subjects in the database. It is evident that the most prominent shim offset is on shims having strong spatial field dependence in the inferior-superior direction (z-axis), especially the Z and Z2 shims. The second dominant shim offset is in the anterior-posterior direction (y-axis), whereas the least is along the left-right direction (x-axis). The first-order shims show relatively high consistency among different subjects, whereas second-order shims vary more relative to the mean across subjects.

Figures 4.3b and 4.3c present the mean and standard deviation, respectively, of the field distribution from all 143 registered field maps in standard space, together with the field compensated for by the averaged fixed shim (Figure 4.3d). The mean

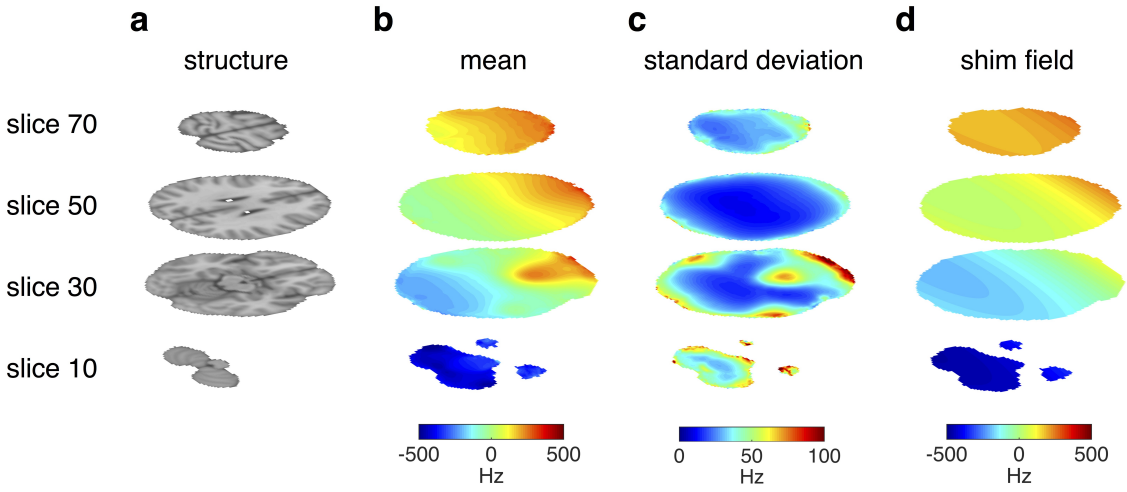


Figure 4.3: Example slices of the MNI brain structural template (a), mean (b) and standard deviation (c) of all 143 registered field maps, and the field compensated for by the shim field (d) generated by the averaged fixed shim.

field distribution indicates the common field features across different subjects in the database. This is dominated by a gradient in the inferior-superior and the anterior-posterior direction, on top of that there is a more localised field deviation in the regions of the orbitofrontal lobe. The standard deviation represents the variation in field distribution among different subjects. The largest intersubject field variations are in regions near the frontal sinuses and the ear cavities. The field distribution in these regions has higher-order spatial definition than can be generated by second-order shim coils for whole-brain shimming. Thus, the mean shim compensation field largely represents the long-ranging field gradients, as is evident from Figure 4.3d.

4.1.2 Simulations of the template-based method

This method produces an estimate of the B_0 field distribution in the individual subject, which can then be used for shimming. To calculate shim currents, measured field distributions from the first- and second-order shims coils were fit to the predicted field

map using a linear least-squares approach. The goal was to minimize B_0 inhomogeneity within the region defined by the brain extraction mask. The feasibility of using the template-based field map prediction for shimming was evaluated in comparison to a range of other potential shimming strategies in both simulation and experiment. To quantify the performance of the different approaches, the standard deviation of the residual field after shimming was calculated in each case.

The field map database contains a rich amount of information, so to evaluate what different features are contributing to shimming performance, a range of scenarios was simulated. For these simulations, we term this method the "average registered field map." In the simulations, the brain structural information available in the T1-weighted image was used to transform the template field map from standard space to the space of the subject of interest. The results were compared with a shim based on the actual field map for that subject (i.e., the "measured field map").

We might expect that some brains in the database would act as a very good predictor of field and others much less so. To assess this, each individual field map in the database was transformed to the space of the subject of interest, and this field map was used to calculate shim terms. We term this the "registered field map."

These registered field maps incorporate information about the brain geometry of the individual subjects. To evaluate whether this geometry information delivers better field map prediction for the shim settings calculation, an assessment without such information is useful. To do this we took the optimum shim for each individual subject in the database, without any transformation, and applied this to the subject of interest. We term this the "fixed shim." Furthermore, we calculated the average of all of the fixed shim values in the database to yield a template shim setting, here termed the "average fixed shim."

Then, to evaluate to what extent shimming quality will be improved using information from the database, a Monte Carlo simulation of random shims was implemented. The shims were sampled uniformly in the range between the minimum and the maximum of individual fixed shim values across all subjects in the database. We term this the "random shim."

Finally, the default vendor-calibrated shim for an oil phantom, the "tune-up shim," was used as a comparison to demonstrate the benefits of using field map information for shimming. Table 4.1 summarises all of the comparisons.

To quantify the performance of the different shimming methods, a leave-one-out cross-validation analysis was applied to each subject in the database. For each cross-validation step, 142 of 143 subjects were assigned into a training group used to generate a field map template. The remaining one subject was used to evaluate the shimming performance of each of the shimming methods in Table 4.1. The evaluation yielded one shim setting per subject for the averaged registered field map, the averaged fixed shim, the tune-up shim, and the measured field map, whereas for the individual registered shim, individual fixed shim and random shim, a whole distribution was generated per subject. The full set of simulations was repeated for each subject in the database. All simulations were carried out in the MATLAB R2015a environment (MathWorks Inc, Natick, MA, USA).

Figure 4.4 compares the residual B_0 standard deviation following simulated shimming using the shim calculation methods listed in Table 1, for an example subject in the database. The pink, red, blue, and black lines indicate the residual B_0 standard deviation using the measured field map, the averaged registered field map, the averaged fixed shim, and the tune-up shim, respectively. Using the averaged registered field map, the B_0 standard deviation within the brain ROIs is reduced from 152 Hz (tune-up shim) to 47 Hz, only 2 Hz worse than the measured field map method. The

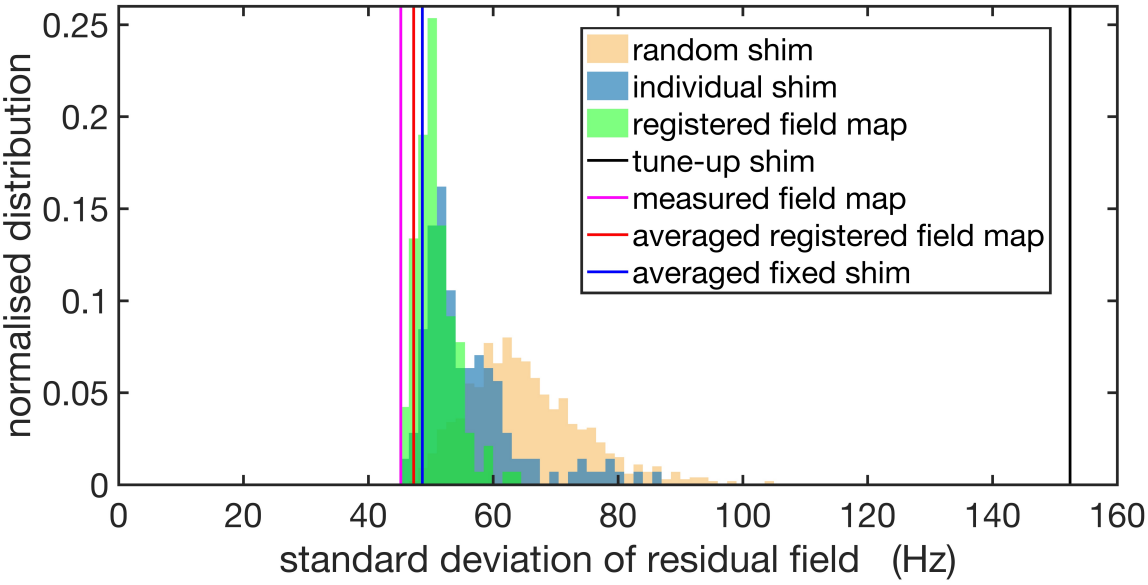


Figure 4.4: Comparison of simulated shimming performance using different shimming methods indicated in Table 4.1 for an example subject in the database. The x-axis shows the standard deviation of residuals after shimming. The y-axis indicates normalized distribution of the number of samples producing certain shimming results. The vertical reference lines from left to right indicate the standard deviation of residuals after shimming with the measured field map of the selected subject (pink line), the averaged registered field map (red line), the averaged fixed shim (blue line), and the tune-up shim (black line). The histograms indicate the shimming results of using individual registered field maps from the other subjects in the database (blue histogram), individual fixed shims (red histogram), and random shims (orange histogram).

averaged fixed shim performed almost as well in this subject, yielding 49 Hz residual B_0 standard deviation.

The green, blue, and orange histograms (normalized to unit area) show the residual B_0 standard deviation using individual registered field maps, individual fixed shims and random shims, respectively. Even the random shims yielded a distribution of residual standard deviation shifted toward much lower values than the tune-up shim, which is unintuitive at first glance, but is likely caused by the range of shim settings being limited to the maximum and minimum values observed in the database. Using the individual fixed shims shifted the mode of the distribution further toward lower residual standard deviation, likely because these represent a "valid" combination of shim terms corresponding to a real subject. Finally, the individual registered field maps yielded a narrower distribution and reduced the tail of poorly performing shims. Interestingly, the averaged fixed shim and the averaged registered field map both yielded lower residual standard deviation than the mode of the corresponding distributions.

Term	Definition
Measured field map	Measured spatial field distribution of an individual subject
Individual registered field map	Field map generated by registering a measured field map of a particular subject to another subject
Averaged registered field map	Average of registered field maps for all of the subjects in the database
Individual fixed shim	Shims calculated based on the measured field map of an individual subject
Averaged fixed shim	The average of the individual fixed shims across all of the subjects in the database
Random shim	Monte Carlo simulation of shims sampled uniformly at random in the range between the minimum and the maximum of individual fixed shim values across all of the subjects in the database
Tune-up shim	Calibration shim performed by manufacturer to optimize field homogeneity in an oil phantom

Table 4.1: Terminology Used for Describing Different Shimming Methods.

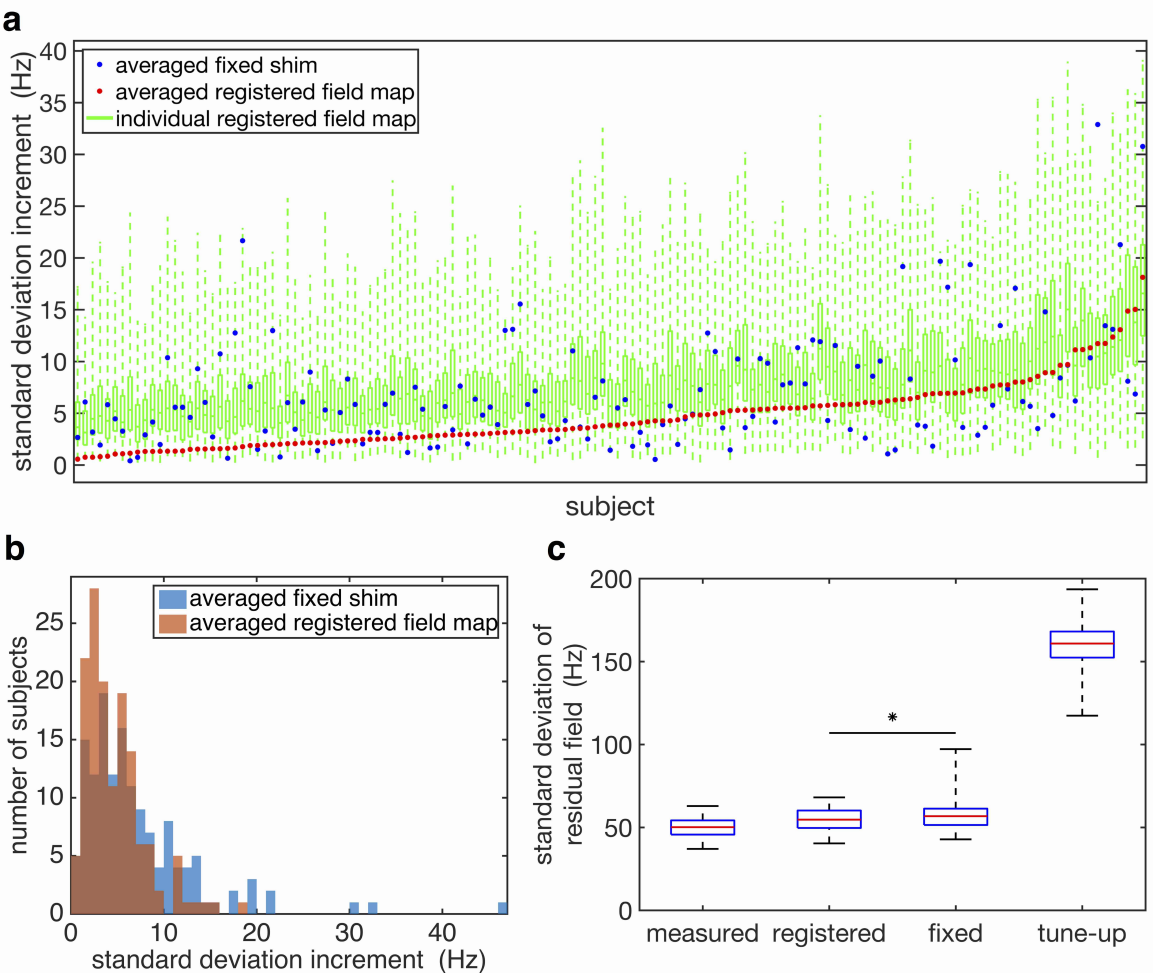


Figure 4.5: Comparison of the simulated shimming performance of the different shimming methods across all subjects in the database.a: Relative increase in standard deviation of the residual field with the averaged registered field map (red dots), the averaged fixed shim (blue dots), and the individual registered field maps (green box plots), compared to using the measured field map. The subjects are sorted according to the values of the relative standard deviation of the averaged registered field map.b: Distribution of relative increase in standard deviation of the residual field using the averaged registered field map (red histogram) and the averaged fixed shim (blue histogram), compared to using the measured field map.c: Distribution of standard deviation of the residual field using the measured field map, the averaged registered field map, the averaged fixed shim, and the tune-up shim for all subjects in the database. The nonparametric Wilcoxon signed-rank test resulted in $P < 0.001$.

The results from all 143 trials from the leave-one-out simulation are depicted in Figure 4.5. In 60% of the cases (86 subjects), the averaged registered field map provided standard deviation of residuals no more than 5 Hz higher than those based on a measured field map. Poor performance of the averaged registered field map generally correlated with a larger spread and higher average of the residual standard deviation of the individual registered field maps. This result has the intuitive interpretation that if the registered field maps used to generate the field map template are not representative of the scanned subject, the averaged registered field map cannot provide good shimming results. However, for all subjects in the database there existed a subset of individual registered field maps that provided a good prediction of the field. In 64% of the cases (91 subjects), the averaged registered field map provided lower standard deviation than the averaged fixed shim. Comparing the distribution of the residual standard deviation shows that the averaged registered field map reduced the tail of poorly performing shim settings compared to the averaged fixed shim (Figure 4.5).

As illustrated in Figure 4.5c, the mean standard deviation of residuals across all subjects in the database was 50 Hz using the measured field map, compared with 160 Hz for the tune-up shim. Using the averaged registered field map for shimming brings the standard deviation of residuals down to, on average, 55 Hz. The average fixed shim was also effective in homogenising the magnetic field over the brain, yielding a mean standard deviation of 57 Hz. To evaluate whether the averaged registered field map statistically outperforms the fixed shim, a nonparametric Wilcoxon signed-rank test was carried out. This demonstrated that there is a significant ($P < 0.001$) improvement in shimming performance using the averaged registered field map, compared to the averaged fixed shim.

4.1.3 Experimental evaluation of the template-based field map prediction

To evaluate the performance of the proposed template-based shimming in a realistic setting, experiments were conducted on seven volunteers who were not included in the database of field maps. In the experiments, a routine 3D localiser (FOV=260 $\times$ 260 $\times$ 201.6 mm^3, voxel size=1.4 $\times$ 1.4 $\times$ 1.4 mm^3, TE=1.56 ms, TR=3.6 ms, GRAPPA factor=3, acquisition time=19 s) was acquired to get the brain geometric information of the scanned subject, to generate the predicted field map. The template-based shimming was compared to a number of different methods used in the simulations. For each method, shimming performance was quantitatively evaluated by the standard deviation in field maps (voxel size=2 $\times$ 2 $\times$ 2 mm^3, TE1=4.08 ms, ΔTE=1.02 ms, slices=80, flip angle=39°, TR=640 ms, FOV=220 $\times$ 220 mm^2, acquisition time=2:24 min) acquired with the respective shim settings, and qualitatively evaluated by signal loss and distortion artefacts in EPI images (voxel size=2 $\times$ 2 $\times$ 2 mm^3, slices=80, TR=4610 ms, TE=25 ms, FOV=220 $\times$ 220 mm^2, GRAPPA factor=2, echo spacing=0.72 ms). A B_0 field map and an EPI volume was acquired with each of the following shim settings:

- Tune-up shims;
- Shims calculated by the routine vendor-supplied shimming method. This measures the field distribution in the subjects' brains with a dual-echo steady-state sequence, and calculates the shim settings designed to minimize field gradients over the brain. The procedure can be performed iteratively, updating the baseline shim settings to the newly calculated ones in each step, incurring an increase in the time required for shimming. Here, we used shim settings resulting from a single iteration and from three iterations, for comparison;

- The averaged fixed shim from the database;
- The proposed template-based shimming (i.e., shims calculated based on the averaged registered field map, using the 3D localiser to transform the map to subject space); and
- Shims calculated based on a measured B_0 field map (parameters as defined previously).

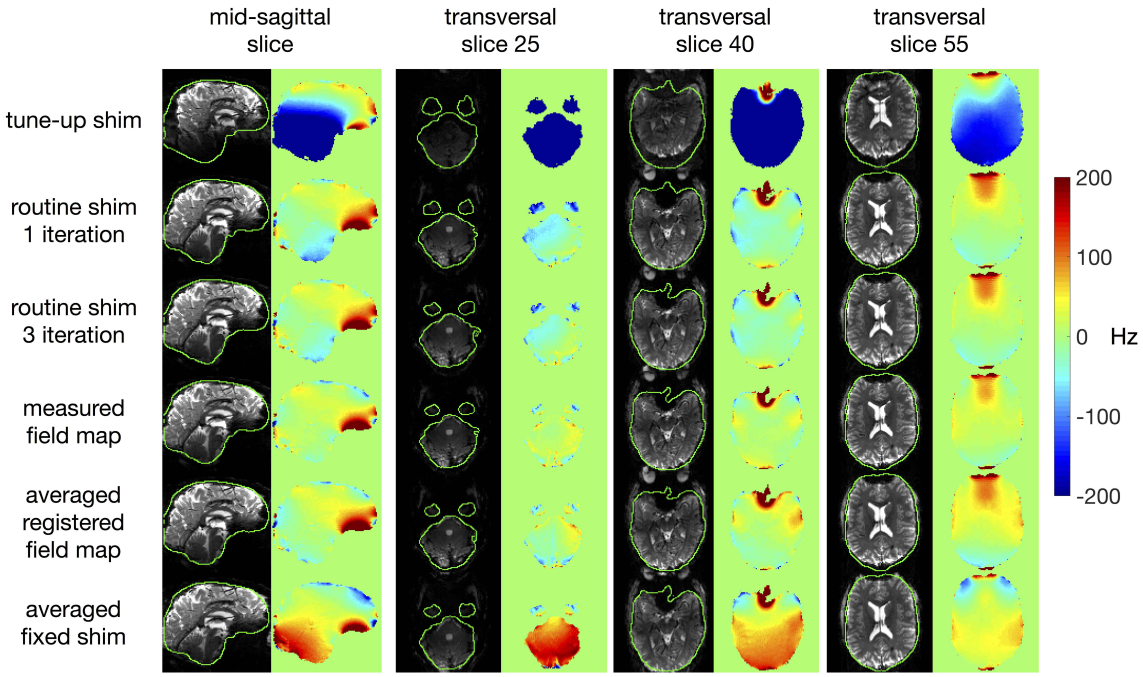


Figure 4.6: Field maps and echo-planar images acquired using different shimming methods in one subject (subject C) not included in the database. The columns from left to right indicate mid-sagittal images (a) and transversal images at slice number 25, 40, and 55 (b-d).The rows (from top to bottom) indicate the shimming results using (i) the tune-up shim, (ii) the vendor-supplied routine shim with oneiteration, (iii) the vendor-supplied routine shim with three iterations, (iv) the measured field map, (v) the averaged registered field map(i.e., the template-based field map prediction), and (vi) the averaged fixed shim.

Figure 4.6 shows sagittal and transversal views of B_0 field maps and EPI images acquired on one of the subjects (subject C) using different shimming strategies. To visualise the magnitude of the EPI distortion artefacts, brain outlines based on the

magnitude images of the field map acquisition are overlaid. In the EPI images, substantial gain in signal and reduction of distortion is visible for all shimming methods compared with the tune-up shim. In the B_0 field maps, the least improvement is observed in regions where air meets tissue (sinuses and ear cavities), as these cannot be fully compensated with static second-order shimming only. For most regions of the brain, the shimming performance of using the averaged registered field map is comparable to the vendor routine shim and the shimming based on the measured field map. However, somewhat lower field homogeneity is observed in the frontal and temporal lobes, using the averaged registered field map. As the largest intersubject field variation is located in these regions (Figure 4.3c), a general field template is always unlikely to perform equally well there. In this example case, the averaged fixed shim provided a more homogeneous field than the other methods in the frontal lobes. However, this came at the cost of overall lower field homogeneity in other regions and was not consistent across all subjects.

Figure 4.7 shows the experimentally measured B_0 field maps and EPI images with shimming based on the averaged registered field map and the measured field map for all seven subjects (A-G). In five of the subjects (A, C, D, E, and G), the averaged registered field map produced results that were very comparable to using a measured field map. For subjects B and F, poor field inhomogeneity was seen at the back of the brain when using the template method. This is most likely because the subjects were scanned with large head rotation angles, which makes the template-based prediction less reliable. Quantified shimming results of the B_0 field homogeneity using the different shimming methods are indicated in Figure 4.8. For all subjects (A-G), both the averaged registered field map and the averaged fixed shim delivered lower standard deviation of the residual field than the tune-up shim. Specifically, static shimming using the averaged registered field map reduced the standard deviation from 176 to 59 Hz on average, which deviated from results using the measured field maps by 9

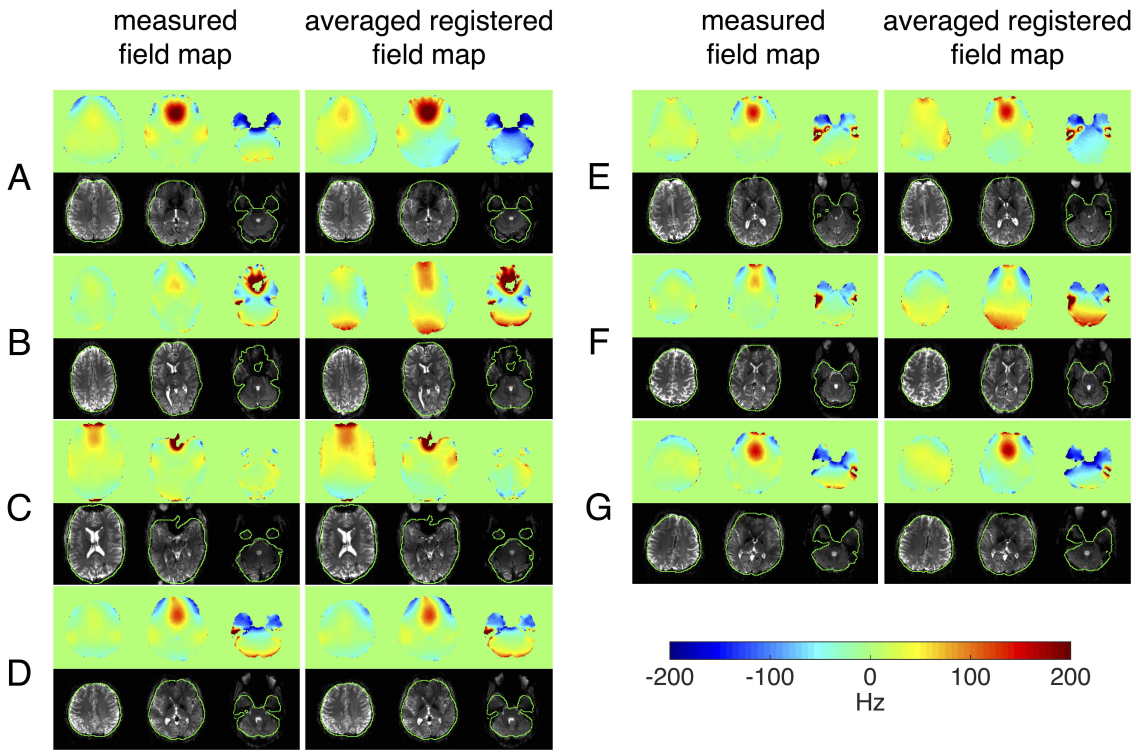


Figure 4.7: Experimentally acquired field maps and echo-planar images shimmed using the measured field map and the averaged registered field map for seven subjects (A-G) not included in the database.

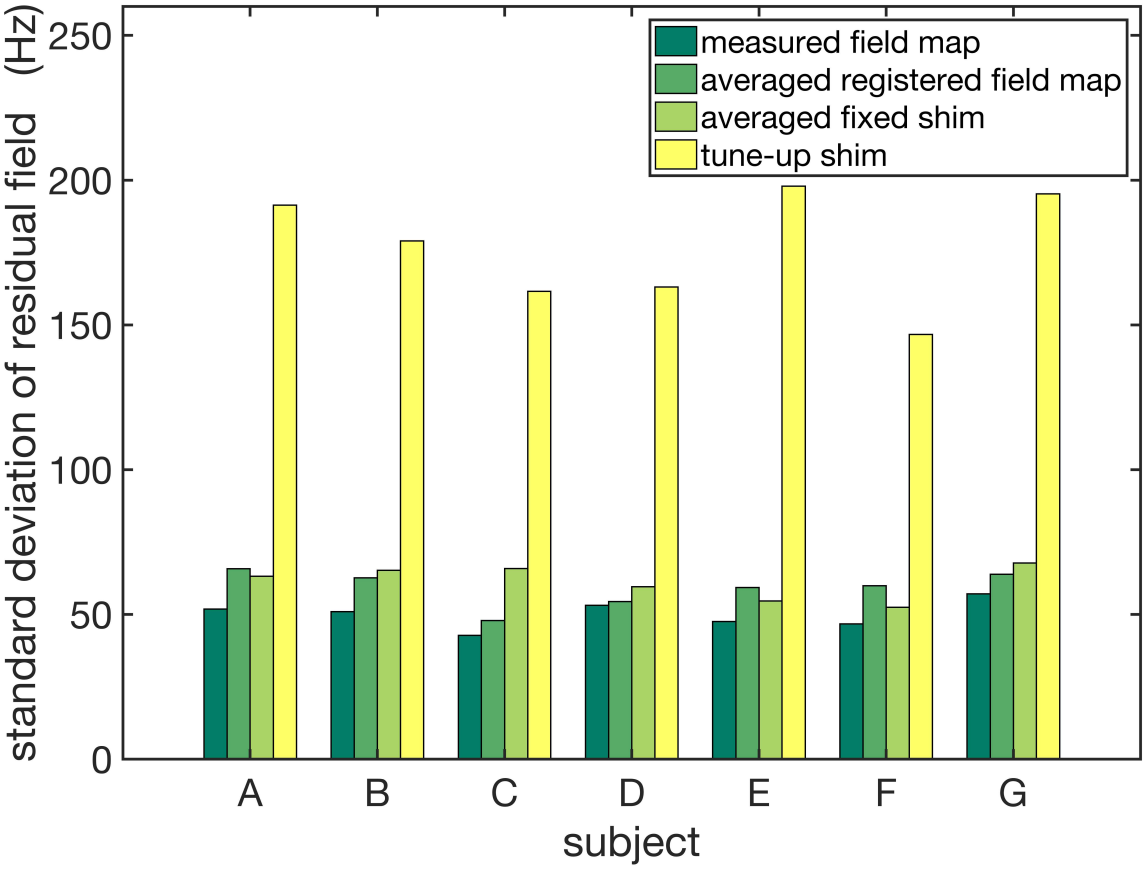


Figure 4.8: Standard deviation of the residual field within the brainafter shimming using the measured field map, the averaged regis-tered field map, the averaged fixed shim, and the tune-up shim for seven subjects (A-G) not included in the database.

Hz. The averaged fixed shim provided a residual standard deviation of 61 Hz on average, comparable to the averaged registered field map. These results suggest that an averaged registered field map offers useful information on the field distribution of the brain for second-order static shimming.

4.1.4 Third-order shimming

All evaluations so far were performed with up to second-order spherical harmonic shimming. However, some MR systems have up to third-order shim hardware, which is not available on our system. To assess whether the field map template approach provides meaningful spatial field information for third-order shimming, additional simulations were performed on the data from the seven subjects not included in the database. The simulations compared second- and third-order shimming based on the measured field map and the averaged registered field map.

The residual standard deviation after simulated second- and third-order shimming based on the averaged registered field map and the measured field map is shown in Figure 4.9. Third-order shimming leads to a further reduction of the field inhomogeneity by 3.4 Hz on average when based on the measured field map, and by 2.1 Hz when based on the averaged registered field map.

4.1.5 Considerations

We have presented a rapid template-based field map prediction approach capable of providing a quick estimation of the field distribution of scanned subjects using only a field map database and a localiser scan. The field distribution in an individual

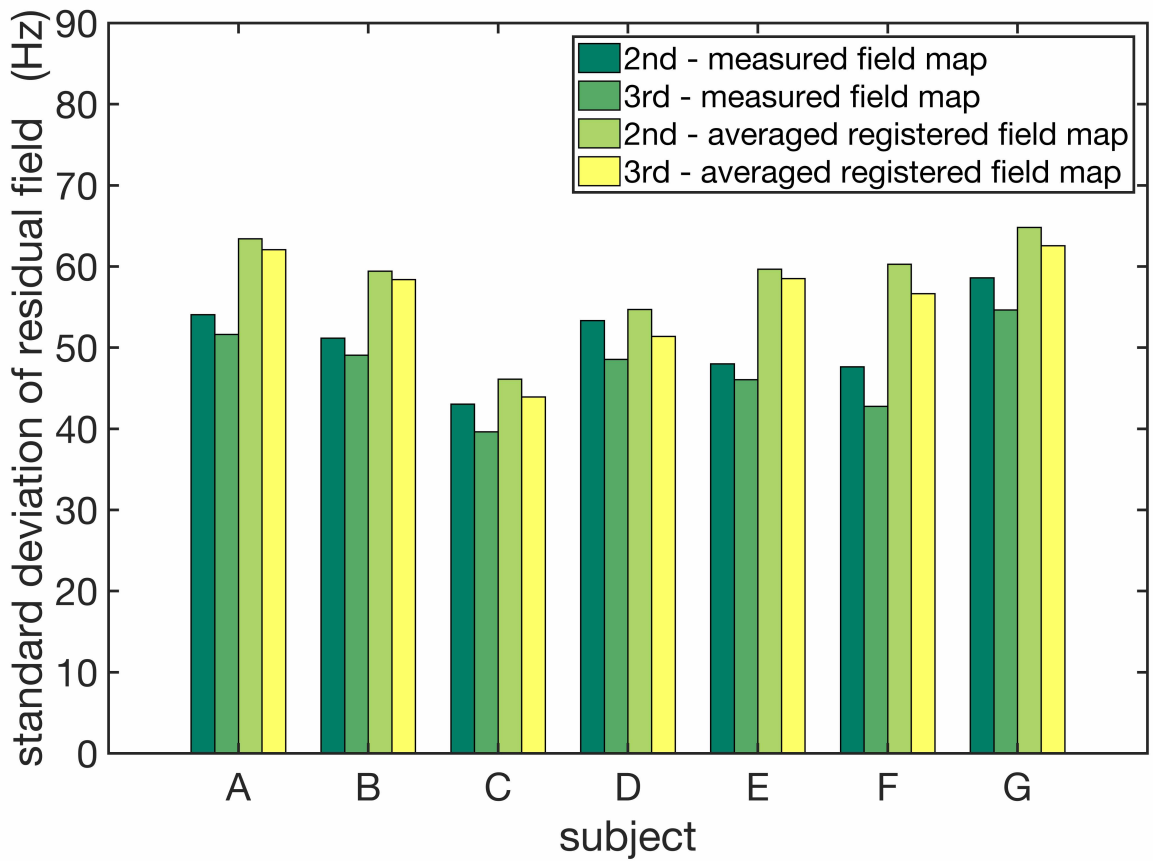


Figure 4.9: Standard deviation of the residual field within the brainafter second- and third-order simulated static shimming using themeasured field map and the averaged registered field map forseven subjects (A-G) not included in the database.

subject's brain is determined by the susceptibility distribution in the head and body, and the positioning angle relative to the main magnetic field. A simple rigid-body transformation of a template B_0 field map would therefore not be expected to provide a highly accurate predictor of the field. However, the results shown here suggest that such an approach may be feasible for the purpose of whole-brain shimming. Both simulations and experiments demonstrate that shimming based on the field map template can substantially reduce the large-scale field inhomogeneity over the brain. In practice, if the shim quality provided by the prediction method is insufficient in individual cases, it still provides a good initial estimate of the shim. The method would therefore reduce the need for multiple iterations in the shim determination.

The reasons behind the feasibility of using a general field map template for static shimming are two-fold. First, in human populations, there is a high degree of similarity of the field distribution in the brain among different subjects. Averaging the registered field maps smoothes out individual differences, while maintaining the dominant field distribution. In addition, second- or third-order spherical harmonic shim fields are relatively lower in spatial order compared with the field across the brain. As demonstrated here, the largest intersubject field variations are in areas of localised field that cannot be shimmed with low-order spherical harmonic shimming. Conversely, the low-order field component that can be addressed by spherical harmonic shimming is fairly consistent among subjects.

The dominant field inhomogeneity over the brain is a gradient along the z-direction, which is associated with the field induced by the body, especially structures above the shoulders (e.g., the jaw and neck) [72]. Because of the high similarity of the low-order fields across subjects, especially along the z-axis, even just the range of individual shim settings over the database carries useful information for shimming. This was demonstrated by the Monte Carlo simulation of random shims, which even

in the worst cases reduced field inhomogeneity within the brain by almost a factor of two compared with the tune-up shims, which are calculated from a phantom. Surprisingly good shim performance was also achieved with the fixed shim approach, using no subject-specific information. The mean residual field inhomogeneity of the fixed shim was almost as low as that of the averaged registered field map, and it even provided better shim quality in individual cases. Therefore, both methods appear to be capable of delivering good whole-brain shimming in most subjects, within a very short timeframe. The fixed shim has the advantage that it is as quick and straightforward to implement as standard shim setting on any system. The performance of the fixed shim, however, was somewhat less robust than for the prediction method. In simulation, two of three subjects showed better shim performance with the averaged registered field map, and the variance of the residual field inhomogeneity was lower. Despite the small difference in average residual field inhomogeneity between the two approaches, it was statistically highly significant.

The main difference between the two methods is that the averaged registered field map incorporates information about the subject's head size and position. Because the center of the shim field is fixed in the scanner, the fixed shim is sensitive to differences in head positioning. In our data, the fixed shim approach was probably helped by the fact that the very tight head coil on our 7T scanner leaves little room for differences in head positioning. The provision of a very clear land marking cross on the head coil also reduces variability in positioning. The information about the subject's geometry in the prediction method comes from a quick localiser scan, which for most MR applications would be acquired anyway. The prediction method therefore yields a small, but significant, gain, at virtually no extra cost in time. In our current "offline" implementation, the total time for field map prediction, along with shim determination, just takes a few seconds. This is about one order of magnitude faster than the vendor routine shimming method, which takes approximately 30 s per it-

eration. Furthermore, graphics processing unit-based parallel computing techniques, and more integrated workflows with the scanner software, would help to accelerate the field map prediction and shim calculation processes even further [147].

The proposed template-based field map prediction was implemented here for ultrahigh-field brain imaging, using second-order shimming. Simulations indicate that systems equipped with third-order shim hardware could also benefit from the averaged registered field map. At higher field strength, shimming becomes increasingly important. However, the prediction method in itself should be equally applicable to imaging at lower field strength, which is more clinically relevant. Future work can be carried out to evaluate the performance of the prediction method on standard clinical 3T systems.

In conclusion, the template-based field map prediction technique joins satisfactory shimming performance with speed in shim determination. A wide range of MR applications could therefore potentially benefit from using this method for shimming of the human brain.

4.2 Adaptive template-based field map prediction for static and dynamic B_0 shimming

Section 4.1 introduced a template-based field map estimation method based on an individual subject's quick localiser scan and a large brain field map database to accelerate the field map acquisition stage for static shimming.

The performance of the averaged registered template-based prediction is degraded if the scanned subject is not well represented by subjects in the database. As shown in Figure 4.5, as the overall shimming performance provided by the individual field

maps in the database becomes more variable, the quality of shimming provided by the average template becomes worse. However, there are always some registered field maps that deliver satisfactory shimming results. If these subjects can be identified based on structural information, such as using the parameters of the linear registration matrix, it may be possible to improve the field map template by averaging the field maps of a subset of selected individuals that represent a good match. Head rotation angles may provide particularly relevant information for subject selection. Head positioning [148], in particular tilting [149], changes the susceptibility-induced B_0 field, which can lead to shim changes of more than 5% [150].

In this section, we propose an adaptive template-based field map prediction method to improve the accuracy of field map template estimation by identifying the field maps in the database that are similar to a particular subject. The adaptive field map prediction method follows the basic framework of the averaged prediction method introduced in **Section 4.1**. It relies on determining an adaptive field map template to predict a subject-specific field map.

The adaptive field map prediction method consists of two steps: field-distribution similarity determination and adaptive field map prediction. The first stage is to create a model to determine the field-distribution similarity between a particular brain and the brains in the 143 field map database. The second stage is to generate an adaptive field-distribution template based on the field maps similar to the particular subject, and using the generated template together with a quick structural scan to predict the field map for the specific subject.

Section 4.2.1 explores the correlations of field map similarity between different subjects with changes in various subject-specific information, like brain geometry, weight, and age. In **Section 4.2.2**, we investigate two methods, a linear model and a Gaussian model, to model the field map similarity using the subject-specific information.

In **Section 4.2.3**, we evaluate the performance of static and dynamic shimming using the two methods as well as identifying the challenges in the current methods.

4.2.1 Correlation between field map similarity and different factors

To assess what factors are good predictors of field distribution, the correlations between different metrics and between-subject field map similarity are determined. The procedures to determine the correlations are illustrated in Figure 4.10.

1. **Field template generation:** As described in **Section 4.1**, the field maps acquired from different subjects were warped to standard space by applying their corresponding registration transformation from the T1-weighted image to the MNI template. Each subject in the database has a generated registration transformation matrix from subject space to standard space. Each registration matrix contains transformation parameters of rotation, scale, and skew. In registration, the rotation parameters function to tilt and rotate the brains, the scale parameters magnify or minify the size of the brains, and the skew parameters twist the brain structure. Each set of rotation, scale, and skew transformation contain three parameters which indicate the corresponding geometric operation about the x-axis (the left-right direction), y-axis (the anterior-posterior direction), and z-axis (the inferior-superior direction), respectively.

2. **Field map similarity quantification:** To quantify the field distribution similarity between a particular subject and subjects in the database, each individual field map in the database was transformed to the space of the specific subject of interest. The similarity between the field map of the specific subject and the

subjects in the database was quantified by the root mean square error of the field map difference. This step was repeated for all the brains in the database. We then obtained a matrix S to indicate the between-subject field map similarity, as shown in Figure 4.10.

3. **Correlation calculation:** These registration parameters generated in Step 1 contain information on the brain geometry. Besides the registration transformation parameters, some other factors may also carry information about the field distribution information in the brain. For example, a registration quality metric indicates how well the registered brain matches the MNI template. As the brain size changes as a function of age [151], subjects of different age ranges have different field distribution in the brain. Additionally, the subjects' weight could be an indication of the body structure which affects the field distribution in the brain. Thus, to find the correlation between the difference in the field distribution in the brain from different subjects and the different subject-specific parameters (e.g., the rotation (TX-R), scale (TX-C), and skew (TX-K) transformation parameters from subject space to standard space, the age, the weight, and the registration quality metric (CorM)), we calculated a distance matrix DM_{ns} for each subject, s, in the database, which is given by

$$DM_{sn} = \begin{bmatrix} TX\text{-}R_{sn} \\ TX\text{-}C_{sn} \\ TX\text{-}K_{sn} \\ age_{sn} \\ weight_{sn} \\ CorM_{sn} \end{bmatrix} = \begin{bmatrix} TX\text{-}R_{s} \\ TX\text{-}C_{s} \\ TX\text{-}K_{s} \\ age_{s} \\ weight_{s} \\ CorM_{s} \end{bmatrix} - \begin{bmatrix} TX\text{-}R_{n} \\ TX\text{-}C_{n} \\ TX\text{-}K_{n} \\ age_{n} \\ weight_{n} \\ CorM_{n} \end{bmatrix} \tag{4.1}$$

where n indicates the brain index in the field map database. Based on the distance matrix for all the brains in the database, the correlation between different subject-specific parameters and the between-subject field map similarity matrix, S, was calculated.

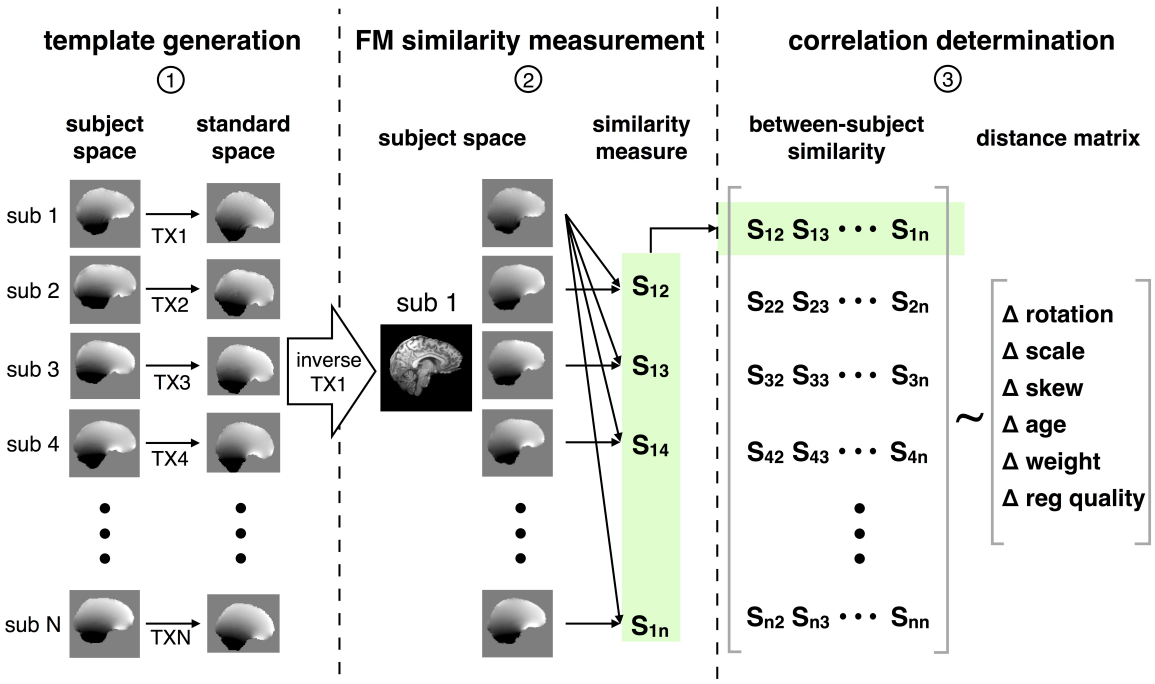


Figure 4.10: Procedures for field map similarity estimation for one example subject in the field map database, which consists of three main stages: template generation, field map similarity measurement, and field map similarity prediction. (1) Field template generation. (2) The field map similarity is determined by calculating the RMSE of the field map difference between the selected subject and the rest of subjects in the database. (3) The distance matrix is determined by calculating difference in subject-specific information, like age and weight. The correlation between field map similarity and distance matrix is calculated.

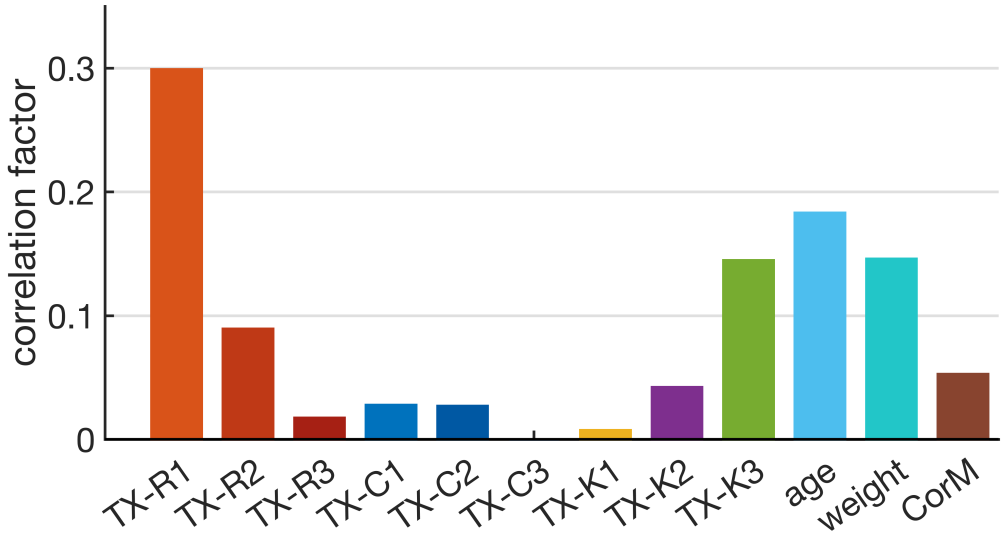


Figure 4.11: Comparison of correlation between the similarity of the field distribution across the different brains in the field map database and various metrics of registration parameters, the age, the weight, and the registration quality metric. TX-R, TX-C, and TX-K indicate the rotation, the scale, and the skew transformation matrix, where the index 1, 2, and 3 indicate the geometric operation in the sagittal plane, the transverse plane, and the coronal plane.

Figure 4.11 shows the correlation between field map similarity and the different subject-specific parameters. The most prominent factor that affects the field distribution in the brain is the rotation matrix, especially the brain rotation about the x-axis. The second dominant factor is the skew matrix, especially the skew about the z-axis, whereas the least is the scale matrix. Moreover, the age and the weight of subjects also show considerable correlations with the field map similarity. These parameters could potentially be used to approximate how much the field-distribution in the brains differ from each other.

In order to validate the findings on the correlations, experiments were conducted on three subjects to assess the field distribution in the brain with varied head rotation angles. In the experiments, subjects were instructed to rotate the head in different positions and a GRE field map was acquired in each position. Specifically, each subject was instructed to hold the head in the normal position first, and then tilt

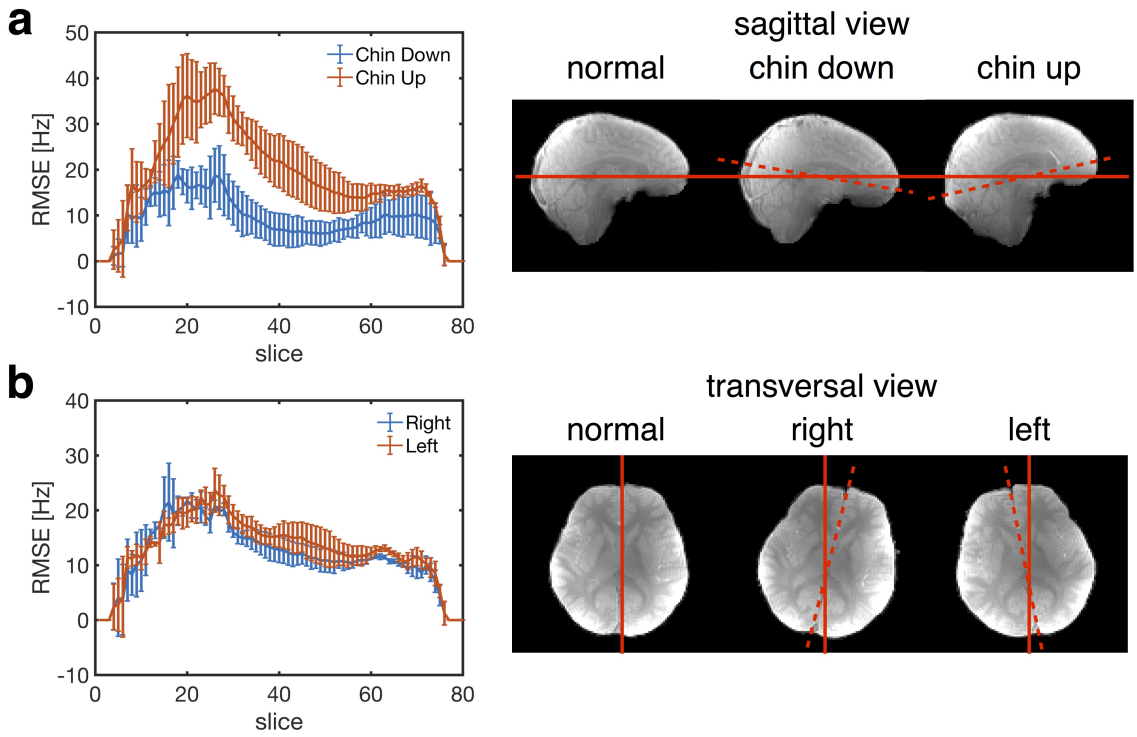


Figure 4.12: Influences of head rotation on the field distribution in the brain over three experimentally acquired subjects. Root-mean-square error of the difference between the field distributions in the brains measured in the norm position and the rotated brain. The plot (a) shows the mean (solid line) and the range (vertical bars) for head rotation with chin down (blue line) and chin up (red line) in the sagittal plane. The plot (b) shows field distributions result from right (blue line) and left (red line) head rotation in the transverse plane.

the head forward, backwards, left and right between different field map acquisitions. After that, the acquired field maps were brain-extracted and unwrapped. The four field maps acquired with head rotation were registered to the field map acquired in the normal position using FLIRT with rigid-body transformation. The difference between field maps acquired in the rotated head positions and the normal position was quantified by the root mean square error as described by Equation 4.1.

The field distribution deviation in the brain resulting from the head rotation is illustrated in Figure 4.12. Figure 4.12a and 4.12b show the root-mean-square error between the field deviations induced by head rotation about the x-axis and about the z-axis, and the normal head position over all of the slices across the brain, respectively. Head rotation around the x-axis caused significant changes of field inhomogeneity in the brain. It is worth noting that tilting the head backwards affected the field distributions almost double the amount of that caused by rotating the head forwards. The head rotation about the z-axis shows a significant but relatively small field variations. As the head coil has limited space, the head rotation about the y-axis was assumed to be negligible. Additionally, the largest field deviation occurred around the slice 25 where the sinuses are located.

4.2.2 Adaptive field map prediction

The adaptive field map template based prediction relies on determining a subset of the field maps in the database showing high similarity to a selected subject using the subject-specific parameters (e.g., registration parameters, age, and weight). There are various methods that can be used to find similar field maps. Here we propose to use two methods to achieve this goal:

1. Using a linear regression to model the relationship of the field similarity between a specific subject and the subjects in the database, with the corresponding distance in the subject specific information.

2. Using the agglomerative hierarchical clustering algorithm to cluster the field maps based on their similarity to each other. A Gaussian mixture model is then utilised to model each cluster based on the subject specific information.

Field map prediction using a linear model

The linear model method relies on determining a model that generates a metric indicating the similarity of the field map between a specific subject and the subjects in the database, based on subject-specific distance parameters. The metric is then used to select field maps in the database that are predicted to be most similar to the specific subject.

Using the field similarity matrix and distance metrics described in **Section 4.2.1**, an adaptive field map template was generated by the following steps step by step, in Figure 4.13.

1. **Parameterising a linear model**. The subject-specific metrics showing high correlation to the field map similarity (e.g., registration metrics TX-R1, TX-R2, TX-K3, age, and weight) were used to create a distance matrix (DM). The matrix was then fit to the field map similarity matrix using least-squares optimisation. The derived parameters correspond to contribution of subject-

specific parameters to the field map similarity measures.

$$min\{\begin{bmatrix} S_{12} \\ S_{13} \\ \vdots \\ S_{n-3n-2} \\ \vdots \\ S_{nn-1} \end{bmatrix} - \begin{bmatrix} DM_{12} \\ DM_{13} \\ \vdots \\ DM_{n-3n-2} \\ \vdots \\ DM_{nn-1} \end{bmatrix} \cdot \begin{bmatrix} L_{TX-R1} \\ L_{TX-R2} \\ L_{TX-K3} \\ L_{age} \\ L_{weight} \end{bmatrix}\}, \tag{4.2}$$

where S indicates similarity measures between different brains, and L denotes the linear model parameters for different factors.

2. For a specific subject, its brain structural image was registered to standard space with FLIRT.

3. The registration parameters together with the subject' age and weight information were extracted from data to form a subject-specific DM-S.

4. The DM-S derived in Step 2 was subtracted to those obtained from subjects in the database to generate a distance matrix, DM-NS, which contains the difference of those parameters. Following that, the distance matrix was multiplied by the LSQlin model parameters, L, to generate a set of field map similarity indexes for subjects in the database. The adaptive field map template was generated by averaging 25 registered field maps in standard space associated with high similarity indices to the specific subject.

5. The final predicted field map was obtained by transferring the created adaptive field map template to the space of the subject by applying the inverse registration transformation.

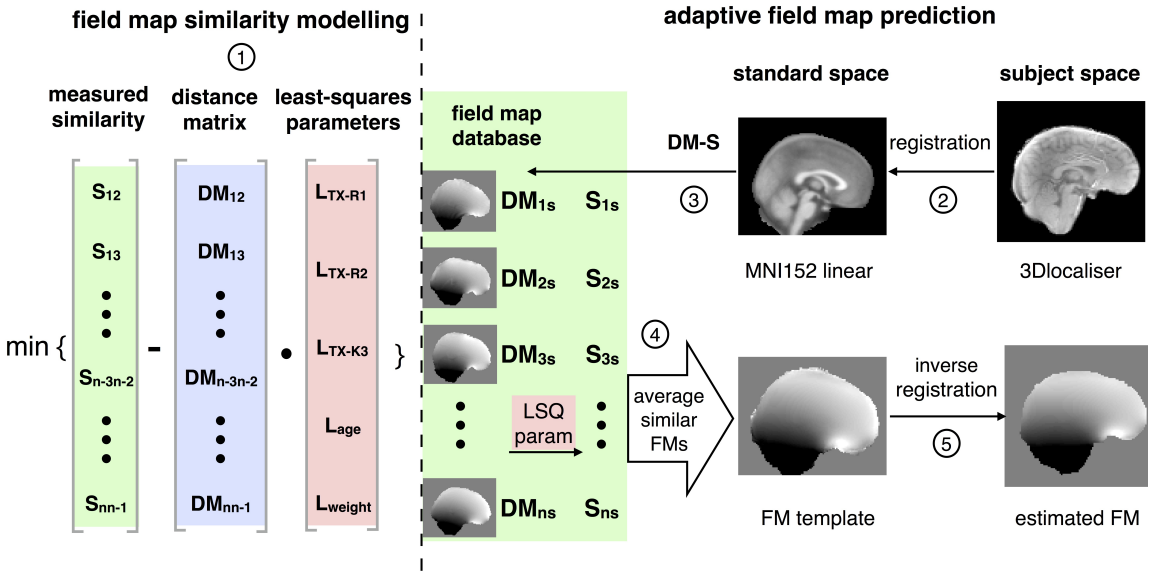


Figure 4.13: Procedures for adaptive template-based field map prediction using a linear regression model: (1) In the linear regression modelling phase, the linear least-squares model parameters was obtained by fitting the similarity matrix to the distance matrix, DM. (2) When the brain structure of a specific subject was acquired, it was registered to standard space. (3) The generated registration parameters together with subject-specific information were subtracted to that of subjects in the database used to generate DM-S. Following that, (4) the generated DM-S was used with the linear model to find most similar brains and averaging the field maps predicted to be similar to the scanned subject. (5) The inverse registration matrix was applied to the created field map template to get the estimated field map corresponding to the specific subject.

Field map prediction using Gaussian mixture models

The Gaussian mixture model finds the probability of a specific subject belonging to a group of field maps that are similar to each other. The steps to generate an adaptive field map template using the Gaussian mixture models method are as follows.

1. **Parameterising Gaussian models**. All the field maps in the database were classified to a few clusters using agglomerative hierarchical clustering, based on the between-subject field-distribution similarity. The number of clusters was set to 5 in the implementation, due to the limited size of the field map database and the relatively large number of subject specific factors used for modelling. A subject-specific correlation matrix (CM) was created based on the subject-specific parameters showing high correlation to the field map similarity, which is given by

$$CM_n = \begin{bmatrix} TX\text{-}R1_n \\ TX\text{-}R2_n \\ TX\text{-}K3_n \\ age_n \\ weight_n \end{bmatrix} \tag{4.3}$$

   In each cluster, the registered field maps in standard space were averaged to generate a template. The generated field template was representative to the field maps in their corresponding cluster. With the assumption that the metrics are Gaussian, the CM parameters of the subjects assigned in each cluster were fit to a mixture of Gaussian distributions. Thus, the data were then represented by the parameters of Gaussian distributions, mean and standard deviation.

2. For a specific subject, its structural map was registered to standard space and yielded a registration matrix.

3. The registration parameters together with subjects' age and weight information were extracted from data to generate a subject-specific CM-S.

4. The probabilities of the individual subject belonging to the Gaussian models from the different clusters were computed based on the CM-S derived in Step 3. The sum of these probabilities is 1. An adaptive field map template in standard space was generated by averaging the templates (GMM-A, B, C, D and E) of each cluster with weighting based on the probabilities derived in Step 4.

5. The predicted field map was then obtained by transferring the generated adaptive field map template to subject space.

4.2.3 Evaluation of adaptive methods in simulation

Figure 4.15 shows the RMSE map of the field difference between the averaged registered template described in **Section 4.1** and the measured field maps (b), the RMSE improvement using the adaptive method with a linear model (c), and a Gaussian mixture model (d), compared to the averaged registered template. Figure 4.15b indicates the averaged registered field map template was less able to predict field distribution in the regions near the sinuses and the ear cavities. In the RMSE improvement map, as shown in Figure 4.15c and 4.15d, the blue patterns represent the regions where the field maps predicted by the adaptive methods were more similar to the measured field maps. The results indicate that the two adaptive methods showed comparable performance and were able to better identify the field in the back of the brain. This suggests that the adaptive prediction methods can provide additional information on the field inhomogeneity of a specific subject beyond that contained in an averaged registered field map. It is not surprising that the adaptive methods did not yield better prediction in the regions close to the sinuses and the ear cavities, as the whole brain geometry information used in the prediction does not account for local subject-specific variables, like the size of the sinuses.

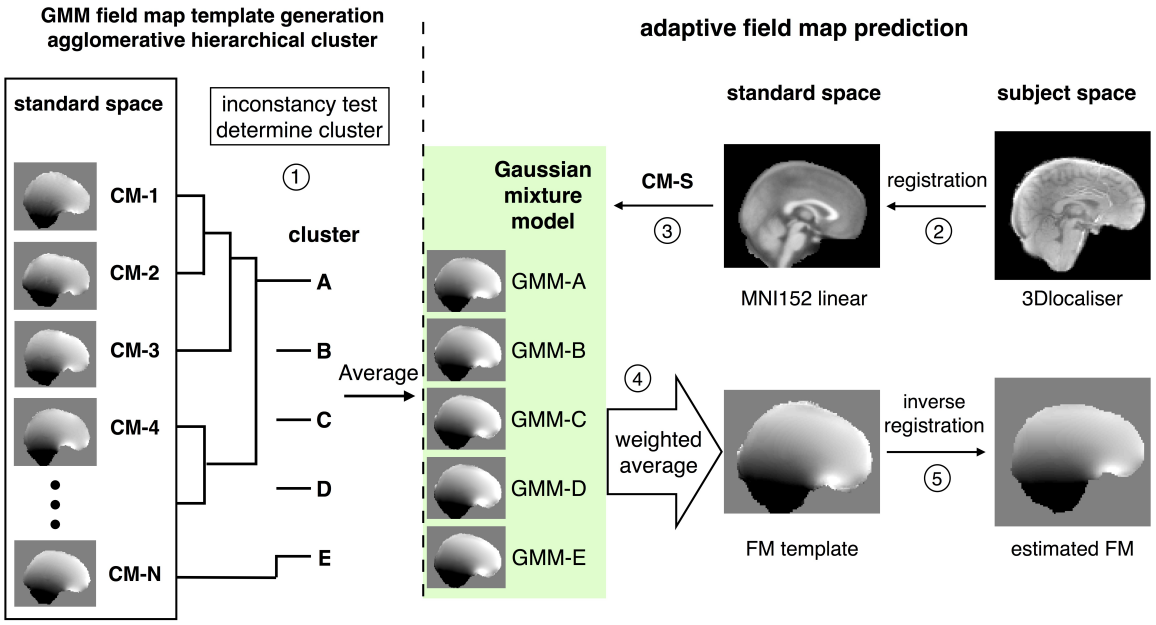


Figure 4.14: Procedures for adaptive template-based field map prediction using Gaussian models: (1) In the field map clustering stage, field maps in the database were classified to a few subgroups based on between-subject similarity. Field maps in each cluster were averaged to generate a field map template. The corresponding CM were fit to Gaussian distributions to generate a Gaussian mixture model (2) When the brain structure of a specific subject was acquired, it was registered to standard space. (3) The generated registration parameters together with subject-specific information were used to generate a CM-S. Following that, (4) the probabilities of the specific subject belonging to Gaussian models derived from all the clusters were computed based on the generated CM-S parameters. An adaptive field map template was generated by averaging the templates (GMM-A, B, C, D and E) with weightings based on the derived probabilities. (5) The inverse registration matrix was applied to the created field map template to get the estimated field map corresponding to the specific subject.

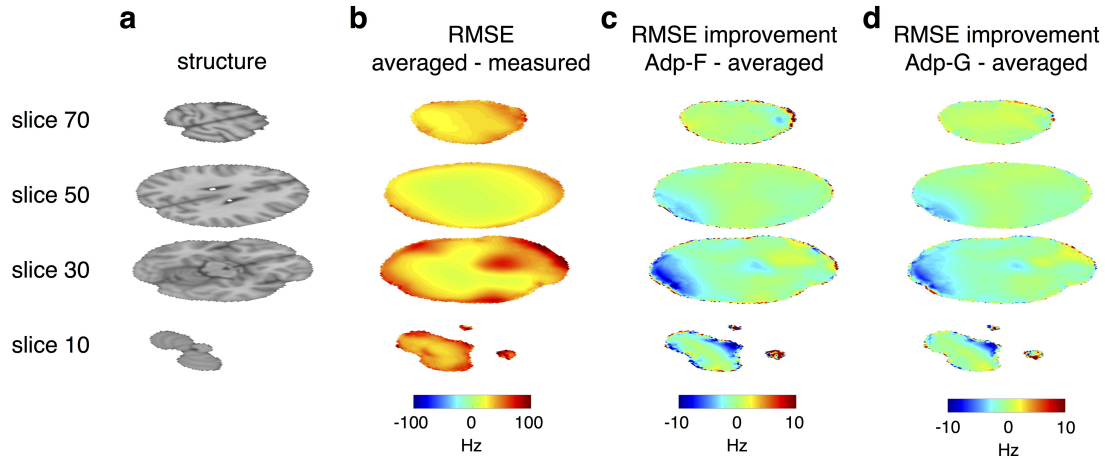


Figure 4.15: Example slices of the MNI brain structural template (**a**), the RMSE between the averaged registered field map template and the measured field maps in standard space, (**b**) the RMSE improvement of using the adaptive methods based on a linear model (**c**), and Gaussian models (**d**) compared to the averaged registered field map template.

A leave-one-out cross validation technique was implemented to assess the performance of static and dynamic shimming using the adaptive field map prediction methods in simulation. The performance of these methods was quantified by the standard deviation of the residual fields after shimming.

As illustrated in Figure 4.16, both the adaptive methods delivered slightly lower standard deviation of the residual fields in static and slice-wise B_0 shimming, compared to the averaged registered template. However, there was still considerable residual field inhomogeneity. For static shimming, the mean standard deviation of the residuals across all the subjects in the database was 50 Hz using the measured field map. Using the adaptive approaches yielded on average 54 Hz standard deviation of the residuals, compared to 55 Hz for the averaged registered field map. The nonparametric Wilcoxon signed-rank tests indicate there is a significant difference between the averaged and the adaptive methods (**$P<0.001$). However, there are no significant difference between the Adp-F and the Adp-G methods. For slice-wise dynamic shimming, shimming based on the measured field maps yielded 40 Hz standard deviation

of the residual fields. The adaptive methods provided 1 Hz improvement on the mean residual field standard deviation, compared to the averaged registered field map, 48 Hz. However, the nonparametric Wilcoxon signed-rank tests show that there is no significant difference between the averaged and the adaptive methods (*P>0.05).

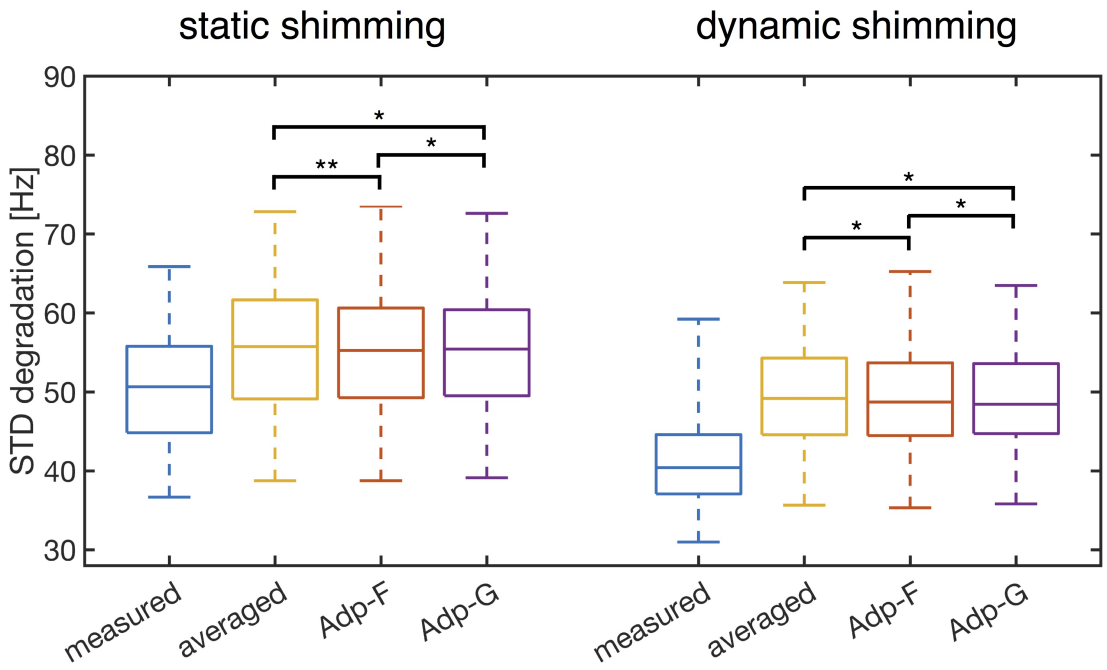


Figure 4.16: Distribution of standard deviation of the residual field using the measured field map, the averaged registered field map, the adaptive method with a linear model, and the adaptive method with Gaussian models for all subjects in the database. The nonparametric Wilcoxon signed-rank test results *P>0.05, **P<0.001

4.2.4 Discussion

This section investigated the performance of a linear model and a Gaussian mixture model for identifying a subset of field maps in the database that better resemble the field distribution of a specific subject not included in the database. The results demonstrated that the adaptive methods slightly outperformed the averaged registered template. The ability of the adaptive methods to better predict the field

variations was found in the back of the brain, compared to the averaged registered field map. It suggests that the adaptive methods have some utility. Nevertheless, the simple strategy of averaging registered field maps is relatively robust and easy to implement.

The adaptive methods can be improved in many aspects.

First, using more representative field distribution predictors has potential to further improve the accuracy of field map prediction. The correlation factors used in the model (e.g., head rotation and skew) only capture the large-scale brain geometry information. The most significant in-brain susceptibility-induced field inhomogeneity is located in the regions near the sinuses and the ear cavities, and at the periphery of the brain close to the skull [33, 152]. Thus, it is expected that inclusion of the size of sinuses and the shape of the brain might yield better prediction on the field distribution in the brain. Additionally, the chest and neck shape have dominant effects on the field distribution in the brain. Thus, the subjects' chest cavities and the size of neck could also be good predictors of the field distribution.

The linear model and the Gaussian mixture model may be inadequate to interpret the field map similarity between different subjects. In practice, the field map similarity might be nonlinearly related to the correlation parameters. Similarly, for the Gaussian model method, the actual distribution of the correlation parameters can be non-Gaussian. Further investigations into how these factors are correlated with the field map distribution could help to improve the accuracy of the prediction.

The prediction relies on a standard brain atlas as reference, to which different subjects' structural images and field maps are registered for template generation. The registration to the standard brain atlas, however, is not representative for all population groups. For subjects who have obvious anatomical changes as a result of

neuropathology or implants, registration to a standard template is unlikely to be sufficient. Standard templates could be stratified for specific target population groups to improve registration accuracy.

Due to the limited size of the database, we parameterised models only using the subject-specific parameters showing relatively high correlations to the field map similarity metrics. However, these parameters only reveal some parts of the field distribution information in the brain. Increasing the database size allows us to include more parameters in the model which may facilitate the selection of matching subjects. Furthermore, with a larger database of field maps, incorporation of advanced machine learning approaches, like the random forest and the neural network method, could potentially yield better prediction.

Chapter 5

Conclusion

Dynamic shimming is an advanced B_0 field homogenisation technique for high-quality MR imaging at ultra-high field. This technique is implemented by optimising field inhomogeneity in small image volumes through real-time updating of shim settings during image acquisition. A promise of dynamic shimming technique is that it facilitates magnetic field homogenisation to a higher degree by making the best use of existing shim coils. In practice, hardware issues, such as eddy currents, and software issues, like ill-conditioned shim calculation problem and long field map acquisition time, compromise the performance of dynamic shimming. In order to improve the efficiency of hardware implementation and the performance of dynamic shimming, this thesis investigated methods to optimise the shim system and shim calculation algorithms, and to accelerate the the shimming process. Investigation into these technical aspects could make an impact and enable a wider acceptance of dynamic shimming techniques.

Chapter 2 focused on methods to address hardware issues in order to produce high-quality shim fields for dynamic shimming. Rapidly altering shim currents induces strong eddy currents that perturb desired shim fields spatially and temporally. We used the pre-emphasis method to compensate for eddy current effects. Determination

of pre-emphasis parameters requires the knowledge of eddy-current-induced fields. We evaluated the performance of two popular strategies, the image-based method and a state-of-art field camera, to measure and characterise eddy-current-induced fields. Our findings suggest that pre-emphasis parameter calibration using the image-based measurement method, though time-consuming, still accounts for the majority of eddy current effects and delivered comparable ECC results as the SIRF measurements.

Chapter 3 investigated shim optimisation algorithms methods to improve the robustness of shimming solutions to errors in field measurements and the conditioning of shim determination problems. We proposed to use a Tikhonov-regularised least-squares algorithm in conjunction with the L-curve method to achieve this goal. In comparison with a range of regularised methods, the results suggested that the proposed method simultaneously improves the current efficiency and the robustness of shim calculation to noise, while preserving good shimming quality. Additionally, the method could potentially reduce hardware demands on high performance shim amplifiers for dynamic shimming applications. Eddy current effects can be significant even with proper ECC. To mitigate the eddy current effects with specifically designing the shim waveform or RF pulse, we proposed to use the Tikh method to reduce the size of shim steps by constraining the shim current amplitude, through manipulating regularisation levels. The results demonstrated the effectiveness of the regularised shim step size optimisation in reducing the size of shim steps, while not significantly degrading theoretical shimming quality. In addition, in this chapter, we also evaluated the effectiveness and potential benefits of dynamic shimming using the signal loss model and the BOLD sensitivity model. The results suggested that the BOLD sensitivity optimisation improved the BOLD sensitivity measure at the cost of the residual field standard deviation and the signal intensity. The signal loss model is a compromise to the BOLD sensitivity and the standard deviation optimisation.

Chapter 4 explored a new direction of field map estimation, predicting a field map of a subject without acquiring any field map, to accelerate the field mapping stages of shimming. The field map prediction method used prior knowledge of the B_0 distribution in the human brain, based on a large database of field maps acquired from different subjects, together with subject-specific structural information from a quick localiser scan. Our results demonstrated that static B_0 shimming based on predicted field maps performed almost as well as shimming based on measured field maps. To investigate whether using subject-specific information in prediction could improve the field map estimation accuracy, we extended the averaged template method by incorporating factors, such as brain geometry, age and weight. Specifically, we proposed two adaptive field map prediction methods, based on a linear model and a Gaussian model, to generate templates from a subset of subjects in the field map database. The results suggested that subject-specific information contains some information about the field distribution in the brain. Incorporating the information in the field map prediction yielded a field homogeneity slightly better than the averaged registered method. Nevertheless, the averaged registered field map still largely accounted for field distribution in the brain.

5.1 Future work

With the growing MR applications at higher fields, advances in shimming techniques have caught great attention and are highly valued by physicists and clinicians. Future work of dynamic shimming techniques could be carried out in the following aspects.

5.1.1 Eddy current compensation

Minimising eddy current effects induced by rapid shim switching is one of most important topics for dynamic shimming. To deliver high-quality dynamic shimming results, endeavour to improve the fidelity of shim fields are worthy. In our implementation of ECC, we adopted exponential models to characterise eddy-current-induced fields. However, in practice, the ECC quality is limited by the number of pre-emphasis terms used which are not capable of delineating the whole field dynamics induced by eddy currents. Additionally, the SIRF measurement outweighs the interpretation ability of the exponential models. Thus, the image-based eddy current characterisation yielded similar ECC performance as using a field camera. To make the full potential of SIRF measurements, non-parametric models can be used for more accurate eddy-current-induced field characterisation, especially for higher-order and cross-term eddy current term compensation. Alternatively, incorporation of the eddy current measurement in image reconstruction could also potentially improve the image quality [153].

5.1.2 Shim calculation optimisation using regularisation methods

Regularised shim optimisation methods investigated in **Chapter 3** aimed to improve the robustness of dynamic shimming and reduce its hardware demands on shim amplifier performance and high ECC quality. In this work, we used the L-curve method for optimal regularisation parameter selection. A possible future work is to explore the performance of using other regularisation parameter searching method (e.g. Generalised Cross Validation), and compare with the L-curve method. Besides this, the performance of using proposed Tikh-L method in B_0 shimming for MR spectroscopy could be further investigated.

Our dynamic shimming implementation only involved the 1st and 2nd shim coils. Inclusion of higher shims, though delivering better shimming quality, increases the difficulty in delivering robust shimming results [108]. Thus, further evaluation of the performance of the proposed regularisation methods could be conducted for high-order dynamic shimming. Besides spherical harmonic shimming applications, further examination of the proposed regularised methods could be carried out for the non-spherical harmonic shimming techniques, like reducing the eddy current effects in multi-coil techniques and minimising shim materials in passive shim design [154].

5.1.3 The template-based field-map prediction

The template-based field map prediction method we developed in **Chapter 4** is a novel field map acquisition strategy and extends the current field map acquisition framework. The current prediction method, though efficient in time, is not capable of delivering field maps as accurate as the conventional field mapping methods. The prediction accuracy can be improved in many ways. From the experience of research, incorporating field map modelling method in template-based field map prediction could yield better field map estimation accuracy.

For the template-based method, the field map prediction RMSE map, as shown in **Chapter 4**, Figure 4.5, illustrates that the template-based method can well approximate the large scale field inhomogeneity, however the field inhomogeneity in the local brain structure cannot be well estimated. The dominant errors in the predicted field map is mostly attributed to the susceptibility induced field inhomogeneity in the regions of sinuses. It can be expected that capturing the field distribution features near the sinuses could potentially yield further improvements on the field map estimation accuracy. In contrast, for the field map modelling method, as the large scale field inhomogeneity over the brain contributes from the susceptibility distribution above

the torso (e.g. neck) [155]. Thus, the field modelling method requires knowledge of tissue structure above the torso to accurately model the field distribution in the brain. In practice, most brain scans only acquire the brain structure, due to the limited anatomical coverage of MR coils, or the limited scan time. Thus, with the acquisition of the brain structure only, the field map modelling methods, though estimating the in-brain local susceptibility-induced field inhomogeneity information, while not capable of estimating the large-scale field inhomogeneity. A combination of the template-based prediction and field map modelling could make full use of the advantages of the two methods, while decoupling their disadvantages. Specifically, instead of acquiring the whole brain structural information, it might be sufficient to model the local field distribution using the structural information located near sinuses only, based on a reference template CT image of sinuses.

Additionally, numerous studies have highlighted the role of machine learning methods in MRI and researchers showed keen interests in predicting data from previous measurements, like magnetic resonance fingerprinting. Exploring various machine learning algorithms in field map prediction could be a new topic in MR research.

5.1.4 Potential applications of the template-based prediction method

The concept of the template-based prediction could be extended and used in other applications. It is feasible to use the rapid template-based method for solving nonlinear optimisation problem in shim current calculation. The initial results for static shimming have been published as an abstract in International Society for Magnetic Resonance in Medicine 2017.

For the application of shimming for maximising BOLD sensitivity, the use of the nonlinear optimisation method can be relatively time-consuming and suffer from lo-

cal minimum problems. The nonlinear optimisation process can be accelerated by using prior knowledge of nonlinear optimisation results on a large database of field maps. This information can then be used in a linear optimisation to deliver shim quality that is closer to the result of a non-linear optimisation. The proposed optimisation method relies on determining a weighting mask to be used in a linear least-squares optimisation. The choice of weighting is determined by identifying those regions where targeting field uniformity will also increase BOLD sensitivity. we have demonstrated that shimming based on this identification yielded similar BOLD sensitivity improvement as direct BOLD optimisation in a short time frame, while still outperforming the LSQ-STD method. However, there is a major factor hindering the performance of proposed method. Optimisation in the whole database was performed based on the specific EPI parameters combined with the specific shim volume. Any changes in EPI parameters will require optimisation to be rerun for the whole database.

The proposed method dramatically reduced the shim calculation time for BOLD sensitivity optimisation and provided satisfactory shimming results. In addition, dynamic shimming would benefit from this technique tremendously, as the shim calculation time scales with the number of slices used.

5.1.5 Potential applications of dynamic shimming

This thesis evaluated the performance of dynamic shimming on the human brain. The increase in popularity of ultra-high field fMRI motivated the investigation into the human central nervous system located in the brainstem [156] [157], and the spinal cord [97]. However, shimming for these regions is challenging. The field distribution in these regions is intrinsic highly inhomogeneous due to large susceptibility artefacts. Physiological noise and subject motion lead to significant temporal-varying bulk magnetic susceptibility shift in those regions [158]. Previous research has attempted to

characterise the field variations induced by the physiological noise [159]. Incorporating those information allows real-time compensation of the respiration-induced time-varying field drifts in the brain [104, 160]. Thus, dynamic shimming that incorporates the model of the spatiotemporal field fluctuations induced by physiological noise and movement in spinal cord [161] [162] could potentially produce good compensation of field inhomogeneity.

For high-field MRI, the benefits of dynamic shimming outweigh the costs on system complexity. Recently, a 7T MRI scanner has been cleared by U.S. Food and Drug Administration for clinical imaging. Integrating a pipeline for automatic implementation of dynamic shimming into routine shimming procedure would bring the enormous convenience for clinicians to enjoy the benefits of the techniques. Besides 7T, much research at even higher field strength, like 11.7T and 14T has highlighted the need for dynamic shimming techniques to further improve the capability of higher-order shims. We believe that dynamic shimming could make great impacts in MRI in the near future.

Appendices

A A list of spherical harmonics up to fourth orders

A list of real-valued spherical harmonic basis functions up to 4th order, and their maximum value on a sphere with 10 cm radius is shown in Table A.1.

B Schematic diagrams of EPI sequences for dynamic shimming

The modified EPI sequence enables real-time shim updating. Figure B.1a shows the schematic diagrams of image acquisition protocol running on the Master MPCU (MPCU0) which conducts the standard image acquisition and the first-order shim corrections. Figure B.1b shows the example schematic diagrams of second-order shim control waveforms running on the Slave MPCU1. The shim waveforms running on MPCU 1 is synchronised to the protocol running on the MPCU0. The highlighted region indicate a event block of acquiring a slice of an EPI image.

index	m	n	basis function	maximum value on unit sphere
0	0	0	1	1
1	1	-1	x	1
2	1	0	y	1
3	1	1	z	1
4	2	-2	xy	0.5
5	2	-1	zy	0.5
6	2	0	$3z^2 - (x^2 + y^2 + z^2)$	2
7	2	1	zx	0.5
8	2	2	$x^2 - y^2$	1
9	3	-3	$3x^2y - y^3$	1
10	3	-2	xyz	0.193
11	3	-1	$y(5z^2 - (x^2 + y^2 + z^2))$	1.377
12	3	0	$5z^3 - 3z(x^2 + y^2 + z^2)$	2
13	3	1	$x(5z^2 - (x^2 + y^2 + z^2))$	1.377
14	3	2	$x^2z - y^2z$	0.385
15	3	3	$x^3 - 3xy^2$	1
16	4	-4	$xy(x^2 - y^2)$	0.25
17	4	-3	$zy(3x^2 - y^2)$	0.325
18	4	-2	$xy(6z^2 - y^2 - x^2)$	0.643
19	4	-1	$zy(4z^2 - 3y^2 - 3x^2)$	1.056
20	4	0	$3(z^2 + y^2 + x^2)^2 - 30(z^2 + y^2 + x^2)z^2 + 35z^4$	8
21	4	1	$xz(4z^2 - 3y^2 - 3x^2)$	1.056
22	4	2	$(x^2 - y^2)(6z^2 - y^2 - x^2)$	1.286
23	4	3	$xz(x^2 - 3y^2)$	0.325
24	4	4	$y^4 - 6x^2y^2 + x^4$	1

Table 1: Maximum value of the spherical harmonics on the surface of a 10-cm-radius sphere

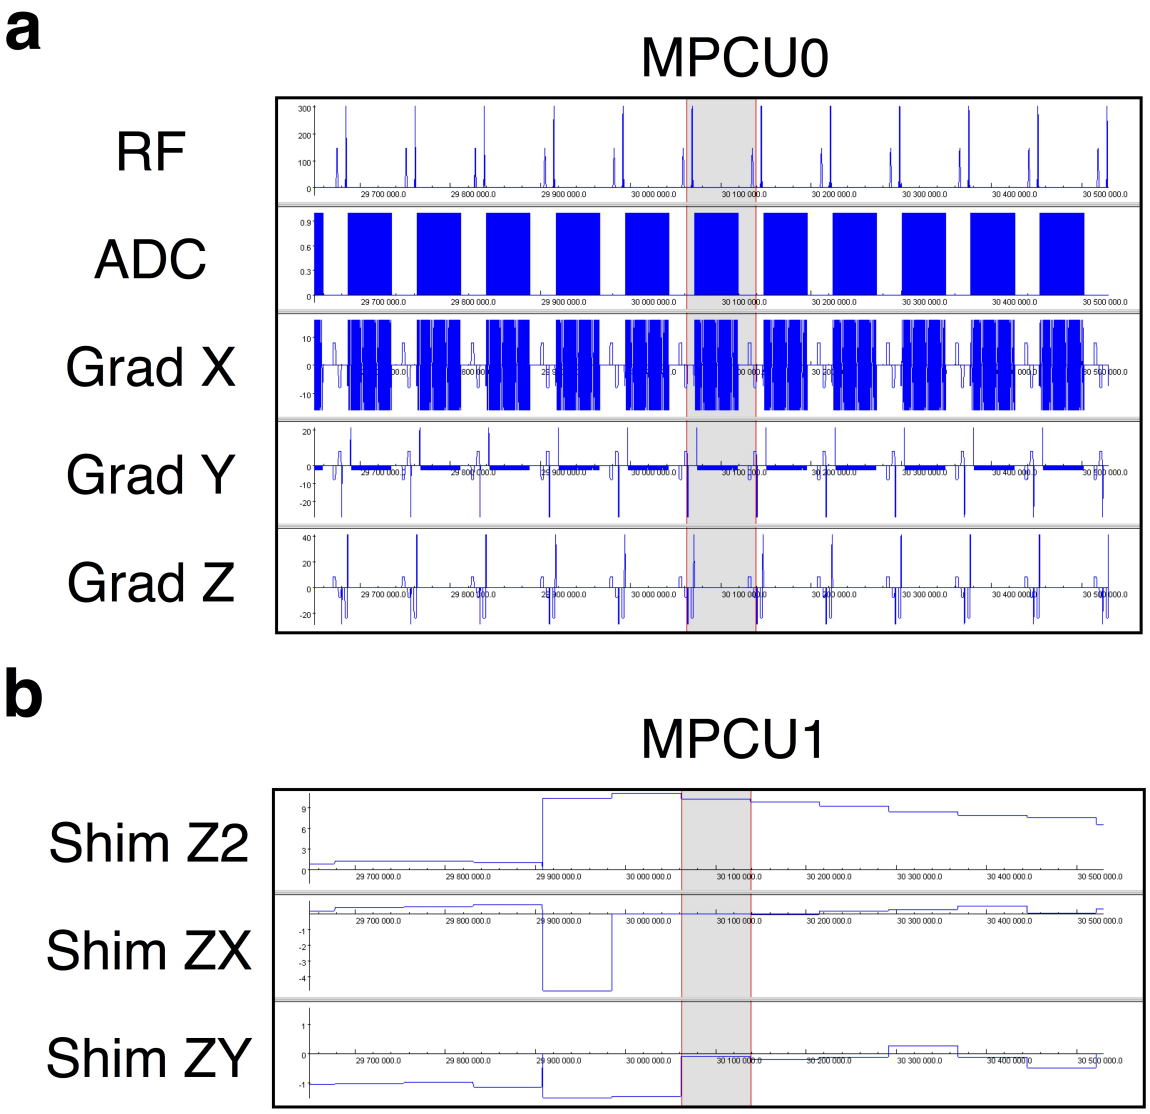


Figure 1: Schematic diagrams of the modified EPI sequences running on MPCU0 (a) and MPCU1 (b). The grey boxes indicate an EPI acquisition event block. It is visible that the shim waveforms running on MPCU1 is synchronised with the imaging protocol running on MPCU0

Bibliography

[1] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, and W. U-Minn HCP Consortium. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.

[2] A. M. Peters, M. J. Brookes, F. G. Hoogenraad, P. A. Gowland, S. T. Francis, P. G. Morris, and R. Bowtell. T2* measurements in human brain at 1.5, 3 and 7 t. *Magn Reson Imaging*, 25(6):748–53, 2007.

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

[4] R. Pohmann, O. Speck, and K. Scheffler. Signal-to-noise ratio and mr tissue parameters in human brain imaging at 3, 7, and 9.4 tesla using current receive coil arrays. *Magn Reson Med*, 75(2):801–9, 2016.

[5] S. Trattnig, W. Bogner, and S. Gruber. Clinical applications at ultrahigh field (7 t). where does it make the difference? *NMR Biomed*, 2015.

[6] A. G. van der Kolk, J. Hendrikse, J. J. Zwanenburg, F. Visser, and P. R. Luijten. Clinical applications of 7 t mri in the brain. *Eur J Radiol*, 82(5):708–18, 2013.

[7] A. G. Webb and P. F. Van de Moortele. The technological future of 7 t mri hardware. *NMR Biomed*, 29(9):1305–1315, 2016.

[8] L. Umutlu, M. E. Ladd, and M. Forsting. 7 tesla mr imaging: Opportunities and challenges. *Rofo*, 2014.

[9] T. Nakada, H. Matsuzawa, H. Igarashi, Y. Fujii, and I. L. Kwee. In vivo visualization of senile-plaque-like pathology in alzheimer's disease patients by mr microscopy on a 7t system. *J Neuroimaging*, 18(2):125–9, 2008.

[10] M. Metcalf, D. Xu, D. T. Okuda, L. Carvajal, R. Srinivasan, D. A. Kelley, P. Mukherjee, S. J. Nelson, D. B. Vigneron, and D. Pelletier. High-resolution phased-array mri of the human brain at 7 tesla: initial experience in multiple sclerosis patients. *J Neuroimaging*, 20(2):141–7, 2010.

[11] R. Beisteiner, S. Robinson, M. Wurnig, and M. Hilbert. Clinical fmri: evidence for a 7t benefit over 3t. *Neuroimage*, 2011.

[12] K. M. Koch, D. L. Rothman, and R. A. de Graaf. Optimization of static magnetic field homogeneity in the human and animal brain in vivo. *Prog Nucl Magn Reson Spectrosc*, 54(2):69–96, 2009.

[13] G. N. Chmurny and D. I. Hoult. The ancient and honorable art of shimming. *Concepts Magn Reson Part A Bridg Educ Res*, 2(3):131–149, 1990.

[14] E. Schneider and G. Glover. Rapid in vivo proton shimming. *Magn Reson Med*, 18(2):335–347, 1991.

[15] P. Webb and A. Macovski. Rapid, fully automatic, arbitrary-volume in vivo shimming. *Magn Reson Med*, 20(1):113–22, 1991.

[16] A. M. Blamire, D. L. Rothman, and T. Nixon. Dynamic shim updating: A new approach towards optimized whole brain shimming. *Magn Reson Med*, 36(1):159–165, 1996.

[17] G. Morrell and D. Spielman. Dynamic shimming for multislice magnetic resonance imaging. *Magn Reson Med*, 38(3):477–483, 1997.

[18] K. M. Koch, S. McIntyre, T. W. Nixon, D. L. Rothman, and R. A. de Graaf. Dynamic shim updating on the human brain. *J Magn Reson*, 180(2):286–96, 2006.

[19] F. Bloch, W. W. Hansen, and Martin Packard. Nuclear induction. *Physical Review*, 69(3-4):127–127, 1946.

[20] E. M. Purcell, H. C. Torrey, and R. V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69(1-2):37–38, 1946.

[21] K. Golman, L. E. Olsson, O. Axelsson, S. Mansson, M. Karlsson, and J. S. Petersson. Molecular imaging using hyperpolarized 13c. *Br J Radiol*, 76 Spec No 2:S118–27, 2003.

[22] W. D. Rooney, G. Johnson, X. Li, E. R. Cohen, S. G. Kim, K. Ugurbil, and Jr. Springer, C. S. Magnetic field and tissue dependencies of human brain longitudinal 1h2o relaxation in vivo. *Magn Reson Med*, 57(2):308–18, 2007.

[23] W. M. Spees, D. A. Yablonskiy, M. C. Oswood, and J. J. H. Ackerman. Water proton mr properties of human blood at 1.5 tesla: Magnetic susceptibility, t1, t2, t*2, and non-lorentzian signal behavior. *Magn Reson Med*, 45(4):533–542, 2001.

[24] M. A. Fernndez-Seara and F. W. Wehrli. Postprocessing technique to correct for background gradients in image-based r2* measurements. *Magn Reson Med*, 44(3):358–366, 2000.

[25] P. C. Lauterbur. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, 242(5394):190–191, 1973.

[26] P. Mansfield. Multi-planar image formation using nmr spin echoes. *Journal of Physics C: Solid State Physics*, 10(3):L55–L58, 1977.

[27] B. P. Sutton, D. C. Noll, and J. A. Fessler. Dynamic field map estimation using a spiral-in/spiral-out acquisition. *Magn Reson Med*, 51(6):1194–204, 2004.

[28] C. De Panfilis and C. Schwarzbauer. Positive or negative blips? the effect of phase encoding scheme on susceptibility-induced signal losses in epi. *Neuroimage*, 25(1):112–21, 2005.

[29] R. Turner. Gradient coil design: A review of methods. *Magn Reson Imaging*, 11(7):903–920, 1993.

[30] D. I. Hoult and D. Lee. Shimming a superconducting nuclear magnetic resonance imaging magnet with steel. *Review of scientific instruments*, 56(1):131–135, 1985.

[31] A. Belov, V. Bushuev, M. Emelianov, V. Eregin, Y. Severgin, S. Sytchevski, and V. Vasiliev. Passive shimming of the superconducting magnet for mri. *IEEE Transactions on Applied Superconductivity*, 5(2):679–681, 1995.

[32] F. Roméo and D. I. Hoult. Magnet field profiling: Analysis and correcting coil design. *Magn Reson Med*, 1(1):44–65, 1984.

[33] C. M. Collins, B. Yang, Q. X. Yang, and M. B. Smith. Numerical calculations of the static magnetic field in three-dimensional multi-tissue models of the human head. *Magn Reson Imaging*, 20(5):413–424, 2002.

[34] Roger J. Ordidge, Jay M. Gorell, Jean C. Deniau, Robert A. Knight, and Joseph A. Helpern. Assessment of relative brain iron concentrations using t2-weighted and t2*-weighted mri at 3 tesla. *Magn Reson Med*, 32(3):335–341, 1994.

[35] J. F. Schenck. The role of magnetic susceptibility in magnetic resonance imaging: Mri magnetic compatibility of the first and second kinds. *Med Phys*, 23(6):815–50, 1996.

[36] A. L. Sukstanskii and D. A. Yablonskiy. Theory of fid nmr signal dephasing induced by mesoscopic magnetic field inhomogeneities in biological systems. *J Magn Reson*, 151(1):107–17, 2001.

[37] R. A. Dashner, A. Kangarlu, D. L. Clark, A. RayChaudhury, and D. W. Chakeres. Limits of 8-tesla magnetic resonance imaging spatial resolution of the deoxygenated cerebral microvasculature. *Journal of Magnetic Resonance Imaging*, 19(3):303–307, 2004.

[38] W. van der Zwaag, S. Francis, K. Head, A. Peters, P. Gowland, P. Morris, and R. Bowtell. fmri at 1.5, 3 and 7 t: characterising bold signal changes. *Neuroimage*, 47(4):1425–34, 2009.

[39] J. R. Reichenbach, R. Venkatesan, D. A. Yablonskiy, M. R. Thompson, S. Lai, and E. M. Haacke. Theory and application of static field inhomogeneity effects in gradient-echo imaging. *Magn Reson Imaging*, 7(2):266–79, 1997.

[40] P. Jezzard, A. S. Barnett, and C. Pierpaoli. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magn Reson Med*, 39(5):801–812, 1998.

[41] R. A. De Graaf and C. Juchem. B0 shimming technology. In A. G. Webb and W. S. Price, editors, *New Developments in NMR*, volume 2016-January, pages 166–207. Royal Society of Chemistry, 2016.

[42] A. Deshmane, V. Gulani, M. A. Griswold, and N. Seiberlich. Parallel mr imaging. *Magn Reson Imaging*, 36(1):55–72, 2012.

[43] N. Chen and A. M. Wyrwicz. Optimized distortion correction technique for echo planar imaging. *Magn Reson Med*, 45(3):525–528, 2001.

[44] N. Chen and A. M. Wyrwicz. Correction for epi distortions using multi-echo gradient-echo imaging. *Magn Reson Med*, 41(6):1206–1213, 1999.

[45] R. Cusack, M. Brett, and K. Osswald. An evaluation of the use of magnetic field maps to undistort echo-planar images. *NeuroImage*, 18(1):127–142, 2003.

[46] P. Jezzard and R. S. Balaban. Correction for geometric distortion in echo planar images from b0 field variations. *Magn Reson Med*, 34(1):65–73, 1995.

[47] P. J. Reber, E. C. Wong, R. B. Buxton, and L. R. Frank. Correction of off resonance-related distortion in echo-planar imaging using epi-based field maps. *Magn Reson Med*, 39(2):328–330, 1998.

[48] Y. M. Kadah and X. Hu. Algebraic reconstruction for magnetic resonance imaging under b0 inhomogeneity. *IEEE Trans Med Imaging*, 17(3):362–70, 1998.

[49] P. Munger, G. R. Crelier, T. M. Peters, and G. B. Pike. An inverse problem approach to the correction of distortion in epi images. *IEEE Transactions on Medical Imaging*, 19(7):681–689, 2000.

[50] D. Cordes, P. A. Turski, and J. A. Sorenson. Compensation of susceptibility-induced signal loss in echo-planar imaging for functional applications. *Magn Reson Imaging*, 18(9):1055–68, 2000.

[51] R. Deichmann, J. A. Gottfried, C. Hutton, and R. Turner. Optimized epi for fmri studies of the orbitofrontal cortex. *NeuroImage*, 19(2):430–441, 2003.

[52] R. Deichmann, O. Josephs, C. Hutton, D. R. Corfield, and R. Turner. Compensation of susceptibility-induced bold sensitivity losses in echo-planar fmri imaging. *Neuroimage*, 15(1):120–35, 2002.

[53] N. Weiskopf, C. Hutton, O. Josephs, and R. Deichmann. Optimal epi parameters for reduction of susceptibility-induced bold sensitivity losses: a whole-brain analysis at 3 t and 1.5 t. *Neuroimage*, 33(2):493–504, 2006.

[54] N. Weiskopf, C. Hutton, O. Josephs, R. Turner, and R. Deichmann. Optimized epi for fmri studies of the orbitofrontal cortex: compensation of susceptibility-induced gradients in the readout direction. *MAGMA*, 20(1):39–49, 2007.

[55] S. Volz, E. Hattingen, C. Preibisch, T. Gasser, and R. Deichmann. Reduction of susceptibility-induced signal losses in multi-gradient-echo images: application to improved visualization of the subthalamic nucleus. *Neuroimage*, 45(4):1135–43, 2009.

[56] K. D. Merboldt, J. Finsterbusch, and J. Frahm. Reducing inhomogeneity artifacts in functional mri of human brain activation-thin sections vs gradient compensation. *J Magn Reson*, 145(2):184–91, 2000.

[57] I. R. Young, I. J. Cox, D. J. Bryant, and G. M. Bydder. The benefits of increasing spatial resolution as a means of reducing artifacts due to field inhomogeneities. *Magn Reson Imaging*, 6(5):585–590, 1988.

[58] G. Liu and S. Ogawa. Epi image reconstruction with correction of distortion and signal losses. *Magn Reson Imaging*, 24(3):683–9, 2006.

[59] F. Jens, M. Klaus-Dietmar, and H. Wolfgang. Direct flash mr imaging of magnetic field inhomogeneities by gradient compensation. *Magn Reson Med*, 6(4):474–480, 1988.

[60] Z. H. Cho and Y. M. Ro. Reduction of susceptibility artifact in gradientecho imaging. *Magn Reson Med*, 23(1):193–200, 1992.

[61] J. Tropp, K. A. Derby, C. Hawryszko, S. Sugiura, and H. Yamagata. Automated shimming of b0 for spectroscopic imaging. *Journal of Magnetic Resonance (1969)*, 85(2):244–254, 1989.

[62] J. Dagher, T. Reese, and A. Bilgin. Highresolution, large dynamic range field map estimation. *Magn Reson Med*, 2014.

[63] L. M. Klassen and R. S. Menon. Robust automated shimming technique using arbitrary mapping acquisition parameters (rastamap). *Magn Reson Med*, 51(5):881–7, 2004.

[64] R. Gruetter and I. Tkac. Field mapping without reference scan using asymmetric echo-planar techniques. *Magn Reson Med*, 43(2):319–323, 2000.

[65] K. S. Nayak, C. Tsai, C. H. Meyer, and D. G. Nishimura. Efficient off-resonance correction for spiral imaging. *Magn Reson Med*, 45(3):521–524, 2001.

[66] D. H. Kim, E. Adalsteinsson, G. H. Glover, and D. M. Spielman. Regularized higher-order in vivo shimming. *Magn Reson Med*, 48(4):715–22, 2002.

[67] K. Scheffler. Fast frequency mapping with balanced ssfp: theory and application to proton-resonance frequency shift thermometry. *Magn Reson Med*, 51(6):1205–11, 2004.

[68] R. Gruetter. Automatic, localized in vivo adjustment of all first-and second-order shim coils. *Magn Reson Med*, 29(6):804–811, 1993.

[69] J. Shen, R. E. Rycyna, and D. L. Rothman. Improvements on an in vivo automatic shimming method [fastermap]. *Magn Reson Med*, 38(5):834–9, 1997.

[70] O. Reynaud, D. Gallichan, B. Schaller, and R. Gruetter. Fast low-specific absorption rate b0 -mapping along projections at high field using two-dimensional radiofrequency pulses. *Magn Reson Med*, 73(3):901–8, 2015.

[71] R. Gruetter and C. Boesch. Fast, noniterative shimming of spatially localized signals. in vivo analysis of the magnetic field along axes. *Journal of Magnetic Resonance (1969)*, 96(2):323–334, 1992.

[72] K. M. Koch, X. Papademetris, D. L. Rothman, and R. A. De Graaf. Rapid calculations of susceptibility-induced magnetostatic field perturbations for in vivo magnetic resonance. *Physics in Medicine and Biology*, 51(24):6381–6402, 2006.

[73] S. Li, G. D. Williams, T. A. Frisk, B. W. Arnold, and M. B. Smith. A computer simulation of the static magnetic field distribution in the human head. *Magn Reson Med*, 34(2):268–275, 1995.

[74] R. P. Dewal and Q. X. Yang. Volume of interest-based fourier transform method for calculation of static magnetic field maps from susceptibility distributions. *Magn Reson Med*, 75(6):2473–2480, 2016.

[75] C. Hutton, O. Josephs, J. Stadler, E. Featherstone, A. Reid, O. Speck, J. Bernarding, and N. Weiskopf. The impact of physiological noise correction on fmri at 7 t. *Neuroimage*, 57(1):101–12, 2011.

[76] M. J. Versluis, B. P. Sutton, P. W. de Bruin, P. Bornert, A. G. Webb, and M. J. van Osch. Retrospective image correction in the presence of nonlinear temporal magnetic field changes using multichannel navigator echoes. *Magn Reson Med*, 68(6):1836–45, 2012.

[77] V. Roopchansingh, R. W. Cox, A. Jesmanowicz, B. D. Ward, and J. S. Hyde. Single-shot magnetic field mapping embedded in echo-planar time-course imaging. *Magn Reson Med*, 50(4):839–43, 2003.

[78] M. Jenkinson. Fast, automated, n-dimensional phase-unwrapping algorithm. *Magn Reson Med*, 49(1):193–7, 2003.

[79] A. K. Funai, J. A. Fessler, D. T. Yeo, V. T. Olafsson, and D. C. Noll. Regularized field map estimation in mri. *IEEE Trans Med Imaging*, 27(10):1484–94, 2008.

[80] J. L. Wilson, M. Jenkinson, I. de Araujo, M. L. Kringelbach, E. T. Rolls, and P. Jezzard. Fast, fully automated global and local magnetic field optimization for fmri of the human brain. *NeuroImage*, 17(2):967–976, 2002.

[81] J. W. Pan, K. M. Lo, and H. P. Hetherington. Role of very high order and degree b0 shimming for spectroscopic imaging of the human brain at 7 tesla. *Magn Reson Med*, 68(4):1007–1017, 2012.

[82] R. Cusack, B. Russell, S. M. Cox, C. De Panfilis, C. Schwarzbauer, and R. Ansorge. An evaluation of the use of passive shimming to improve frontal sensitivity in fmri. *Neuroimage*, 24(1):82–91, 2005.

[83] J. L. Wilson, M. Jenkinson, and P. Jezzard. Optimization of static field homogeneity in human brain using diamagnetic passive shims. *Magn Reson Med*, 48(5):906–14, 2002.

[84] J. L. Wilson, M. Jenkinson, and P. Jezzard. Protocol to determine the optimal intraoral passive shim for minimisation of susceptibility artifact in human inferior frontal cortex. *NeuroImage*, 19(4):1802–1811, 2003.

[85] J. L. Wilson and P. Jezzard. Utilization of an intra-oral diamagnetic passive shim in functional mri of the inferior frontal cortex. *Magn Reson Med*, 50(5):1089–1094, 2003.

[86] J. L. Wilson, M. Jenkinson, and P. Jezzard. A bit of a mouth full: Susceptibility artifact reduction using diamagnetic passive shims. *Book of Abstracts: Tenth Scientific Meeting of the International Society for Magn Reson Med*, page 205, 2002.

[87] C. Juchem, B. Muller-Bierl, F. Schick, N. K. Logothetis, and J. Pfeuffer. Combined passive and active shimming for in vivo mr spectroscopy at high magnetic fields. *J Magn Reson*, 183(2):278–89, 2006.

[88] R. A. Österbauer, J. L. Wilson, G. A. Calvert, and P. Jezzard. Physical and physiological consequences of passive intra-oral shimming. *NeuroImage*, 29(1):245–253, 2006.

[89] J. L. Wilson. *Development of Functional Magnetic Resonance Imaging at High Field for Neuro-psychiatric Disorders*. PhD thesis, University of Oxford, 2003.

[90] J. J. Hsu and G. H. Glover. Mitigation of susceptibility-induced signal loss in neuroimaging using localized shim coils. *Magn Reson Med*, 53(2):243–248, 2005.

[91] C. Juchem, T. W. Nixon, S. McIntyre, D. L. Rothman, and R. A. De Graaf. Magnetic field homogenization of the human prefrontal cortex with a set of localized electrical coils. *Magn Reson Med*, 63(1):171–180, 2010.

[92] C. Juchem, P. Herman, B. G. Sanganahalli, P. B. Brown, S. McIntyre, T. W. Nixon, D. Green, F. Hyder, and R. A. de Graaf. Dynamic multi-coil technique (dynamite) shimming of the rat brain at 11.7t. *NMR Biomed*, 27(8):897–906, 2014.

[93] C. Juchem, T. W. Nixon, S. McIntyre, V. O. Boer, D. L. Rothman, and R. A. De Graaf. Dynamic multi-coil shimming of the human brain at 7 t. *Journal of Magnetic Resonance*, 212(2):280–288, 2011.

[94] C. Juchem, S. Umesh Rudrapatna, T. W. Nixon, and R. A. de Graaf. Dynamic multi-coil technique (dynamite) shimming for echo-planar imaging of the human brain at 7 tesla. *Neuroimage*, 105:462–72, 2015.

[95] C. Juchem, P. B. Brown, T. W. Nixon, S. McIntyre, D. L. Rothman, and R. A. de Graaf. Multicoil shimming of the mouse brain. *Magn Reson Med*, 66(3):893–900, 2011.

[96] R. A. De Graaf, P. B. Brown, S. McIntyre, D. L. Rothman, and T. W. Nixon. Dynamic shim updating (dsu) for multislice signal acquisition. *Magn Reson Med*, 49(3):409–416, 2003.

[97] J. Finsterbusch, C. Sprenger, and C. Büchel. Combined t2*-weighted measurements of the human brain and cervical spinal cord with a dynamic shim update. *NeuroImage*, 79:153–161, 2013.

[98] V. O. Boer, B. L. van de Bank, G. van Vliet, P. R. Luijten, and D. W. Klomp. Direct b0 field monitoring and real-time b0 field updating in the human breast at 7 tesla. *Magn Reson Med*, 67(2):586–91, 2012.

[99] C. Lemke, A. Hess, S. Clare, V. Bachtiar, C. Stagg, P. Jezzard, and U. Emir. Two-voxel spectroscopy with dynamic b0 shimming and flip angle adjustment at 7 t in the human motor cortex. *NMR Biomed*, 28(7):852–860, 2015.

[100] K. M. Koch, L. I. Sacolick, T. W. Nixon, S. McIntyre, D. L. Rothman, and R. A. De Graaf. Dynamically shimmed multivoxel 1h magnetic resonance spectroscopy and multislice magnetic resonance spectroscopic imaging of the human brain. *Magn Reson Med*, 57(3):587–591, 2007.

[101] Y. Zhao, A. W. Anderson, and J. C. Gore. Computer simulation studies of the effects of dynamic shimming on susceptibility artifacts in epi at high field. *J Magn Reson*, 173(1):10–22, 2005.

[102] C. Juchem, T. W. Nixon, P. Diduch, D. L. Rothman, P. Starewicz, and R. A. de Graaf. Dynamic shimming of the human brain at 7 tesla. *Concepts Magn Reson Part B Magn Reson Eng*, 37B(3):116–128, 2010.

[103] S. Bollmann, L. Kasper, S. J. Vannesjo, A. O. Diaconescu, B. E. Dietrich, S. Gross, K. E. Stephan, and K. P. Pruessmann. Analysis and correction of field fluctuations in fmri data using field monitoring. *Neuroimage*, 2017.

[104] P. van Gelderen, J. A. de Zwart, P. Starewicz, R. S. Hinks, and J. H. Duyn. Real-time shimming to compensate for respiration-induced b0 fluctuations. *Magn Reson Med*, 57(2):362–8, 2007.

[105] M. Poole and R. Bowtell. Volume parcellation for improved dynamic shimming. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 21(1-2):31–40, 2008.

[106] S. M. Smith. Fast robust automated brain extraction. *Hum Brain Mapp*, 17(3):143–55, 2002.

[107] S. Sengupta, E. B. Welch, Y. Zhao, D. Foxall, P. Starewicz, A. W. Anderson, J. C. Gore, and M. J. Avison. Dynamic b0 shimming at 7 t. *Magn Reson Imaging*, 29(4):483–96, 2011.

[108] A. Fillmer, S. J. Vannesjo, M. Pavan, M. Scheidegger, K. P. Pruessmann, and A. Henning. Fast iterative pre-emphasis calibration method enabling third-order dynamic shim updated fmri. *Magn Reson Med*, 75(3):1119–31, 2016.

[109] P. Mansfield and B. Chapman. Active magnetic screening of gradient coils in nmr imaging. *J Magn Reson (1969)*, 66(3):573–576, 1986.

[110] P. Mansfield and B. Chapman. Multishield active magnetic screening of coil structures in nmr. *J Magn Reson (1969)*, 72(2):211–223, 1987.

[111] W. M. Spees, N. Buhl, P. Sun, J. J. Ackerman, J. J. Neil, and J. R. Garbow. Quantification and compensation of eddy-current-induced magnetic-field gradients. *J Magn Reson*, 212(1):116–23, 2011.

[112] M. A. Morich, D. A. Lampman, W. R. Dannels, and F. D. Goldie. Exact temporal eddy current compensation in magnetic resonance imaging systems. *IEEE Trans Med Imaging*, 7(3):247–54, 1988.

[113] P. Jehenson, M. Westphal, and N. Schuff. Analytical method for the compensation of eddy-current effects induced by pulsed magnetic field gradients in nmr systems. *J Magn Reson (1969)*, 90(2):264–278, 1990.

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

[115] A. Bhogal, M. Versluis, J. Koonen, J. C. W. Siero, V. O. Boer, D. Klomp, P. R. Luijten, and H. Hoogduin. Image-based method to measure and characterize

shim-induced eddy current fields. *Concepts Magn Reson Part A Bridg Educ Res*, 42(6):245–260, 2013.

[116] R. E. Wysong and I. J. Lowe. A simple method of measuring gradient induced eddy currents to set compensation networks. *Magn Reson Med*, 29(1):119–121, 1993.

[117] N. De Zanche, C. Barmet, J. A. Nordmeyer-Massner, and K. P. Pruessmann. Nmr probes for measuring magnetic fields and field dynamics in mr systems. *Magn Reson Med*, 60(1):176–86, 2008.

[118] B. E. Hammer. Magnetic field mapping with an array of nuclear magnetic resonance probes. *Review of Scientific Instruments*, 67(6):2378–2380, 1996.

[119] H. Han, R. P. MacGregor, and B. J. Balcom. Pure phase encode magnetic field gradient monitor. *J Magn Reson*, 201(2):212–7, 2009.

[120] S. J. Vannesjo, B. E. Dietrich, M. Pavan, D. O. Brunner, B. J. Wilm, C. Barmet, and K. P. Pruessmann. Field camera measurements of gradient and shim impulse responses using frequency sweeps. *Magn Reson Med*, 72(2):570–83, 2014.

[121] S. Gross, C. Barmet, B. E. Dietrich, D. O. Brunner, T. Schmid, and K. P. Pruessmann. Dynamic nuclear magnetic resonance field sensing with part-per-trillion resolution. *Nat Commun*, 7:13702, 2016.

[122] S. J. Vannesjo, Y. Duerst, L. Vionnet, B. E. Dietrich, M. Pavan, S. Gross, C. Barmet, and K. P. Pruessmann. Gradient and shim pre-emphasis by inversion of a linear time-invariant system model. *Magn Reson Med*, 2016.

[123] D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM J Appl Math*, 11(2):431–441, 1963.

[124] N. Kickler, W. van der Zwaag, R. Mekle, T. Kober, J. P. Marques, G. Krueger, and R. Gruetter. Eddy current effects on a clinical 7t-68 cm bore scanner. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 23(1):39–43, 2010.

[125] Diego Hernando, Karl K. Vigen, Ann Shimakawa, and Scott B. Reeder. R2* mapping in the presence of macroscopic b0 field variations. *Magn Reson Med*, 68(3):830–840, 2012.

[126] M. J. Graves, P. Wong, A. N. Priest, R. T. Black, and D. J. Lomas. Slice offset frequency and shim adjustment for interactive steady-state free-precession (ssfp) imaging. *Journal of Magnetic Resonance Imaging*, 29(5):1230–1233, 2009.

[127] H. Wen and F. A. Jaffer. An in vivo automated shimming method taking into account shim current constraints. *Magn Reson Med*, 34(6):898–904, 1995.

[128] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews. Advances in functional and structural mr image analysis and implementation as fsl. *Neuroimage*, 23 Suppl 1:S208–19, 2004.

[129] E. Balteau, C. Hutton, and N. Weiskopf. Improved shimming for fmri specifically optimizing the local bold sensitivity. *Neuroimage*, 49(1):327–36, 2010.

[130] J. Shen. Effect of degenerate spherical harmonics and a method for automatic shimming of oblique slices. *NMR Biomed*, 14(3):177–183, 2001.

[131] P. C. Hansen. Regularization tools: A matlab package for analysis and solution of discrete ill-posed problems. *Numerical Algorithms*, 6(1):1–35, 1994.

[132] K. Song, Q. Bao, F. Chen, C. Huang, J. Feng, and C. Liu. Gradient shimming based on regularized estimation for b0-field and shim functions. *J Magn Reson*, 268:1–9, 2016.

[133] S. Nassirpour, P. Chang, A. Fillmer, and A. Henning. A comparison of optimization algorithms for localized in vivo b0 shimming. *Magn Reson Med*, 2017.

[134] G. Shou, L. Xia, F. Liu, M. Zhu, Y. Li, and S. Crozier. Mri coil design using boundary-element method with regularization technique: A numerical calculation study. *IEEE Transactions on Magnetics*, 46(4):1052–1059, April 2010.

[135] P. C. Hansen and D. P. OLeary. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 14(6):1487–1503, 1993.

[136] P. Craven and G. Wahba. Smoothing noisy data with spline functions. *Numerische Mathematik*, 31(4):377–403, Dec 1978.

[137] P. Jezzard, A. S. Barnett, and C. Pierpaoli. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magn Reson Med*, 39(5):801–812, 1998.

[138] J. Finsterbusch. Eddy-current compensated diffusion weighting with a single refocusing rf pulse. *Magn Reson Med*, 61(3):748–54, 2009.

[139] T. G. Reese, O. Heid, R. M. Weisskoff, and V. J. Wedeen. Reduction of eddy-current-induced distortion in diffusion mri using a twice-refocused spin echo. *Magn Reson Med*, 49(1):177–82, 2003.

[140] A. L. Alexander, J. S. Tsuruda, and D. L. Parker. Elimination of eddy current artifacts in diffusion-weighted echo-planar images: The use of bipolar gradients. *Magn Reson Med*, 38(6):1016–1021, 1997.

[141] S. Sengupta, M. J. Avison, J. C. Gore, and E. Brian Welch. Software compensation of eddy current fields in multislice high order dynamic shimming. *Journal of Magnetic Resonance*, 210(2):218–227, 2011.

[142] G. H. Glover. 3d z-shim method for reduction of susceptibility effects in bold fmri. *Magn Reson Med*, 42(2):290–9, 1999.

[143] C. Triantafyllou, R. D. Hoge, G. Krueger, C. J. Wiggins, A. Potthast, G. C. Wiggins, and L. L. Wald. Comparison of physiological noise at 1.5 t, 3 t and 7 t and optimization of fmri acquisition parameters. *Neuroimage*, 26(1):243–50, 2005.

[144] Y. Shi, S. J. Vannesjo, K. L. Miller, and S. Clare. Template- based field map prediction for rapid whole brain b0 shimming. *Magn Reson Med*, 80(1):171–180.

[145] M. Jenkinson, P. Bannister, M. Brady, and S. Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–41, 2002.

[146] G. Grabner, A. L. Janke, M. M. Budge, D. Smith, J. Pruessner, and D. L. Collins. Symmetric atlasing and model based segmentation: an application to the hippocampus in older adults. *Med Image Comput Comput Assist Interv*, 9(Pt 2):58–66, 2006.

[147] T. Y. Huang, Y. W. Tang, and S. Y. Ju. Accelerating image registration of mri by gpu-based parallel computation. *Magn Reson Imaging*, 29(5):712–6, 2011.

[148] P. Jezzard and S. Clare. Sources of distortion in functional mri data. *Hum Brain Mapp*, 8(2-3):80–5, 1999.

[149] J. M. Tyszka and A. N. Mamelak. Quantification of b0 homogeneity variation with head pitch by registered three-dimensional field mapping. *J Magn Reson*, 159(2):213–8, 2002.

[150] S. Clare, J. Evans, and P. Jezzard. Requirements for room temperature shimming of the human brain. *Magn Reson Med*, 55(1):210–214, 2006.

[151] M. Mennes, M. Jenkinson, R. Valabregue, J. K. Buitelaar, C. Beckmann, and S. Smith. Optimizing full-brain coverage in human brain mri through population distributions of brain size. *Neuroimage*, 98:513–20, 2014.

[152] O. Gonen and R. I. Grossman. The accuracy of whole brain n-acetylaspartate quantification. *Magn Reson Imaging*, 18(10):1255–1258, 2000.

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

[154] X. Kong, M. Zhu, L. Xia, Q. Wang, Y. Li, X. Zhu, F. Liu, and S. Crozier. Passive shimming of a superconducting magnet using the l1-norm regularized least square algorithm. *J Magn Reson*, 263:122–5, 2016.

[155] M. Jenkinson, J. L. Wilson, and P. Jezzard. Perturbation method for magnetic field calculations of nonconductive objects. *Magn Reson Med*, 52(3):471–477, 2004.

[156] A. K. Harvey, K. T. Pattinson, J. C. Brooks, S. D. Mayhew, M. Jenkinson, and R. G. Wise. Brainstem functional magnetic resonance imaging: disentangling signal from physiological noise. *Magn Reson Imaging*, 28(6):1337–44, 2008.

[157] R. Sclocco, F. Beissner, M. Bianciardi, J. R. Polimeni, and V. Napadow. Challenges and opportunities for brainstem neuroimaging with ultrahigh field mri. *Neuroimage*, 2017.

[158] J. T. Vaughan, C. J. Snyder, L. J. DelaBarre, P. J. Bolan, J. Tian, L. Bolinger, G. Adriany, P. Andersen, J. Strupp, and K. Ugurbil. Whole-body imaging at 7t: preliminary results. *Magn Reson Med*, 61(1):244–8, 2009.

[159] D. Raj, D. P. Paey, A. W. Anderson, R. P. Kennan, and J. C. Gore. A model for susceptibility artefacts from respiration in functional echo-planar magnetic resonance imaging. *Physics in Medicine and Biology*, 45(12):3809–3820, 2000.

[160] T. Verma and J. Cohen-Adad. Effect of respiration on the b0 field in the human spinal cord at 3t. *Magn Reson Med*, 72(6):1629–36, 2014.

[161] S. J. Vannesjo, K. L. Miller, S. Clare, and I. Tracey. Spatiotemporal characterization of breathing-induced b0 field fluctuations in the cervical spinal cord at 7t. *Neuroimage*, 167:191–202, 2017.

[162] P. F. Van de Moortele, J. Pfeuffer, G. H. Glover, K. Ugurbil, and X. Hu. Respiration-induced bo fluctuations and their spatial distribution in the human brain at 7 tesla. *Magn Reson Med*, 47(5):888–895, 2002.