An Optimisation Framework for Magnetic Resonance Fingerprinting

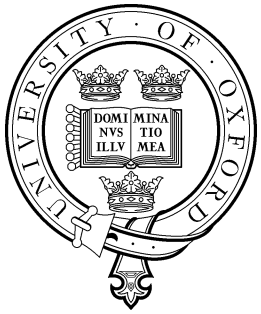

Jack Allen

Jesus College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary 2019

This thesis is dedicated to

Asher & Cyrus

Acknowledgements

Firstly I would like to thank my supervisors Prof. Peter Jezzard and Dr. James Kennedy for their wisdom and dedication. I also acknowledge the funding I received from the MRC and the EPSRC through the ONBI centre for doctoral training.

It's been a great help to have known the company and support of those at FMRIB. To name a few, I'm grateful to Dr. Tom Okell and Dr. Mark Chiew for conducting my progress vivas, to Dr. Johanna Vannesjo for guidance on implementing spiral gradient waveforms, to Dr. Lior Weizman for playing a crucial role in the spiral MRF reconstruction pipeline and to Dr. Jesper Andersson for his optimisation advice. I'm also grateful to Dr. Hongbae Jeong for helping me set up a computer for signal simulations and Dr. Robert Frost for his help regarding simultaneous multislice acquisitions.

The radiographers I have known at AVIC (Peter Manley, Alison Fletcher and Juliet Semple) have been crucial to my data collection and fun to work with.

Aside from those in Oxford, I thank Dr. Dan Ma (Case Western Reserve University, Ohio) and Dr. Jürgen Finsterbusch (University Medical Centre, Hamburg).

My final thanks go to my family for their love and support and to God for the opportunity to work on this thesis and for sustaining me.

Abstract

An Optimisation Framework for Magnetic Resonance Fingerprinting

Jack Allen

Jesus College, University of Oxford

A thesis submitted for the degree of *Doctor of Philosophy*

Hilary 2019

Magnetic Resonance Fingerprinting (MRF) is an efficient quantitative MRI technique for simultaneously mapping multiple tissue or system properties. This thesis presents a framework for optimising the acquisition parameter schedules of a Spin-Echo (SE) Echo Planar Imaging (EPI) MRF pulse sequence, with a view to an application to tracking oedema in acute stroke.

Reference quantitative MRI methods were implemented and used for validations of the SE EPI MRF acquisitions. For this purpose a custom phantom was built with relaxation parameters spanning those typical of *in vivo* grey and white matter. Inversion recovery and spin-echo sequences were used to provide gold standard maps of T_1 and T_2, respectively. A 3D DREAM protocol was used for B_1^+ mapping and gradient-echo sequences provided B_0 off-resonance measurements. Balanced SSFP and FISP MRF sequences were also implemented, with spiral readouts and acquisition parameters similar to those of literature examples.

The framework was designed to optimise the SE EPI MRF excitation flip angles, repetition times and the temporal positioning of inversion pulses. Many of the optimisation parameters were flexible and user-specified, such as the targeted T_1 and T_2 and the evaluation metric. Single-slice and 10-slice schedules were optimised using *fmincon* and genetic algorithms from MATLAB toolboxes and custom cost functions based on Monte Carlo error simulations.

Ten-slice SE EPI MRF schedules requiring 8 seconds and 20 seconds per slice were compared experimentally, with the longer schedules providing the better T_1 and T_2 accuracy and precision. However, both sets of schedules were unable to fit B_1^+ accurately, necessitating the use of a separately acquired B_1^+ map to correct the dictionary flip angles. Experimental phantom results with B_1^+ correction demonstrated agreement with the theoretical accuracy improvements predicted for the optimised schedules but undesired reductions in precision were observed. Those phantom results suggested that the *fmincon* schedule marginally outperformed the corresponding genetic algorithm schedule, showing maximum biases of 16.5% for T_1 and 4.6% for T_2 and maximum spreads of 16.3% for T_1 and 14.3% for T_2. The FISP MRF sequence provided the best all-round performance.

Contents

List of Figures

List of Tables

List of Equations

Preface

Contributions to this Thesis

The spin-echo EPI MRF sequence designed in this thesis was built on preliminary sequence development by Prof. Peter Jezzard. Dr. Mark Chiew provided EPI image reconstruction code, such as that used to apply the necessary phase corrections. Reconstruction code for producing spiral MRF images and parameter maps was provided by Dr. Lior Weizman. Dr. Robert Brand designed the phantom structure.

Dr. Dan Ma (Case Western Reserve University, Ohio) supplied a spiral sampling trajectory and acquisition parameter schedules for bSSFP MRF. Dr. Jürgen Finsterbusch (University Medical Center, Hamburg) developed the code for converting single-band pulses to multi-band and made it available on the Siemens IDEA online discussion board.

Publications Arising from this Thesis

These pieces of work were produced during the same timeframe as the work in this thesis, although they are not necessarily represented in the following chapters.

J. Allen, J. Kennedy, P. Jezzard. Estimating MRF Time Point Importance by Random Forest Classification. In: ISMRM Workshop on Magnetic Resonance Fingerprinting. ; 2017.

J. Allen, J. Kennedy, P. Jezzard. A Comparison of Optimised Single-Shot MR Fingerprinting Pulse Sequence Designs. In: Proceedings of the 25th Annual Meeting of ISMRM. ; 2017. p. 1348.

Chapter 1

Introduction

Magnetic Resonance Fingerprinting (MRF) [1] is a quantitative Magnetic Resonance Imaging (MRI) technique with potential to improve the feasibility of quantitative MRI measurements in clinical environments. This thesis contains an approach for optimising the acquisition parameters of an MRF sequence, motivated by a potential application to stroke imaging. Details of the motivation are provided in this chapter, as well as a summary of the contents of the following chapters.

1.1 Motivation

MRF has been used to efficiently provide quantitative parameter maps of the brain, with an emphasis on a potential application to clinical imaging, a field where the associated speed and repeatability of MRF would be beneficial. There is scope for MRF to be useful in areas where fast diagnosis and treatment is crucial for improving the prognosis of a patient. Information gained in these scenarios can

be vital for ensuring accurate diagnosis and informing decisions about subsequent treatment. Qualitative images are widely used in clinical scanning, but quantitative measurements can provide more information on physiological changes and have the potential to increase the robustness of tissue classification across different scans and imaging centres. A reliable tissue property measurement can be compared to an expected range for a particular pathology.

Stroke imaging, specifically the assessment and monitoring of acute ischeamic stroke, is one area that may benefit from MRF. Worldwide, stroke is one of the leading causes of death and disability [2, 3]. Ischaemic strokes are caused by a blood vessel occlusion, due to small-vessel disease or an embolism, which leads to a localised reduction in blood supply to a portion of the brain. This starves the brain tissue of oxygen, leading to cell damage or death as the normal cellular processes are disrupted. Fast restoration of the blood flow increases the chances of tissue recovery and improves the long-term prognosis [3]. Vessel blockages are treated by the administration of pharmaceutical agents to dissolve the clot (thrombolysis) or recanalisation via an endovascular procedure (thrombectomy). When the blood flow has been restored the affected brain tissue will reperfuse. However, reperfusion injury is possible, leading to further tissue damage and oedema as the integrity of the blood-brain barrier can be compromised, causing the extracellular space to swell [4, 5]. The swelling causes a pressure build-up within the skull, which can cause a fatal displacement and compression of the brain.

Imaging techniques have been used in ischaemic stroke research trials to map the affected areas of tissue using qualitative contrasts [6, 7]. Fast acquisition of quantitative parameter maps through MRF could enable a greater range of perspectives on the patient-specific pathological changes without significantly in-

creasing the total protocol duration. Reductions in protocol duration could be made by using quantitative measurements to generate images with the contrasts used in conventional stroke imaging, thereby reducing the number of acquisitions required. The MR parameters T_1, T_2 and proton density are sensitive to fluid concentration and so could be used to track the development of oedema. However, qualitative images weighted in terms of these parameters do not enable the quantitative characterisation of oedema. To achieve this through water content measurements requires the calibration of T_1, T_2 and corrections for the transmission and receive sensitivity profiles. Monitoring the progression of oedema would be useful for stratifying patients and assessing drug treatment efficacy. Additionally, in MRF the proton density map is derived through the MRF parameter fitting process, meaning the T_1 and T_2 MRF fitting errors have a direct influence on the accuracy of the proton density map. Therefore, the primary objective of this thesis is to develop an optimisation framework to reduce the MRF T_1 and T_2 measurement error, with a view to achieving accurate proton density maps of oedema in clinical research studies of acute stroke.

As a secondary research topic, longitudinal studies could facilitate the estimation of the time since the stroke onset. For example, ischemic tissue T_1 and T_2 values have been found to correlate with the elapsed time since symptom onset [8, 9].

1.2 Thesis Outline

This thesis contains the implementation of conventional quantitative methods, a proposed method for optimising MRF acquisition parameters, and a demonstra-

tion of acquired parameter maps from scan data.

Chapter 2 contains an introduction to the fundamental physical principles of MRI, including the magnetic behaviour of tissue and a description of how images can be produced from magnetic resonance signals. There is also a discussion of commonly used scanning and modelling methods for making quantitative measurements of tissue properties. The chapter concludes with an introduction to the original MRF approach, including subsequent developments from the literature and practical considerations regarding MRF.

Chapter 3 describes commonly used quantitative MRI methods in increased detail, particularly those that were implemented for validation comparisons later in the thesis. Practical details of the implementations are included, along with example parameter maps and phantom measurements.

Chapter 4 introduces a framework for optimising the acquisition schedule of a multislice spin-echo MRF pulse sequence with fully sampled echo-planar image readouts. Comparisons are made between optimisation algorithms and the number of acquired timepoints.

Chapter 5 contains a practical implementation of the theoretical framework described in Chapter 4, with a validation of some of the results from Chapter 4 using phantom and *in vivo* scans.

The thesis concludes with Chapter 6, which discusses the work from the previous chapters and makes recommendations for further research.

Chapter 2

Background and Theory

2.1 Principles of Magnetic Resonance Imaging

This section introduces the physics of MRI, leading to an introduction to the fields of quantitative MRI and MRF. An overview of MRI concepts can also be found in [10]. More detailed descriptions and explanations are contained in [11, 12, 13]. The work of Cercignani *et al.* [14] is a useful resource for quantitative MRI.

2.1.1 Magnetisation

Subatomic nuclei, such as protons, possess a discrete quantum mechanical property called spin, which in the case of an atomic nucleus can combine to result in an overall nuclear spin. The magnetic moment $\boldsymbol{\mu}$ (Eq. 2.1) for a given nucleus is the product of the gyromagnetic ratio γ and the spin vector $\mathbf{S}$, with a specific γ for each species of nucleus.

$$\boldsymbol{\mu} = \gamma \mathbf{S} \tag{2.1}$$

MRI often uses protons in the form of hydrogen nuclei to make an image, relying on the interaction between an ensemble of protons and a magnetic field **B**. In the presence of a static magnetic field along the z-axis many of the spins will also align along the z-axis. The corresponding potential energy E depends on the z-axis component of the magnetic moment, μ_z, and the strength of the static field B_0 applied by the imaging system, as described in Eq. 2.2, where h is Planck's constant and the spin state I_z is $\pm 1/2$ for protons.

$$E = -\mu_z B_0 = -\gamma S_z B_0 = -\gamma \frac{h}{2\pi} I_z B_0 \tag{2.2}$$

The relationship in Eq. 2.2 means that in the case of hydrogen the nuclear spin can occupy one of two energy states with respect to B_0, with spins parallel to B_0 occupying a lower energy state than those antiparallel. The energy needed to move a spin to the higher energy state, δE, is described in Eq. 2.3 and can be provided by incident electromagnetic radiation with a specific angular frequency ω_0 (Eq. 2.4), where the energy of the electromagnetic radiation is proportional to ω_0. This frequency is known as the Larmor frequency and incident electromagnetic energy at the Larmor frequency is said to be on resonance. Groups of spins with the same Larmor frequency are called isochromats.

$$\delta E = \gamma \frac{h}{2\pi} B_0 \tag{2.3}$$

$$\omega_0 = \gamma B_0 \tag{2.4}$$

The thermal energy at the typical temperature range of the human body is such that at thermal equilibrium both spin states are occupied but a small majority populate the energetically favourable lower energy state (i.e. aligned with B_0). Despite there being only a small difference in the size of the two state populations the abundance of protons in human tissue makes MRI signal generation possible. The net magnetisation $\mathbf{M}$ of an ensemble of spins is defined as the sum of the magnetic moments per unit volume, meaning a net magnetisation is induced in the same direction as B_0. The vector $\mathbf{M}$ has components in the transverse (x and y) and longitudinal (z) directions (Eq. 2.5). Although MRI is underpinned by quantum mechanical effects this classical description of magnetisation is usually sufficient for describing the macroscopic behaviour of spins and will be used in the remaining sections and chapters of this thesis.

$$\mathbf{M} = \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} \tag{2.5}$$

2.1.2 The Bloch Equation

The evolution of the magnetisation vector is described classically by the Bloch equation (Eq. 2.6), which can be considered as the addition of three terms. The

following sections describe these terms in more detail.

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} + \frac{1}{T_1}(M_0 - M_z) - \frac{M_{xy}}{T_2} \tag{2.6}$$

2.1.2.1 Precession

The first term (Eq. 2.7) of the Bloch equation (Eq. 2.6) describes the rotation of the magnetisation vector in the presence of an externally applied magnetic field $\mathbf{B}$.

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} \tag{2.7}$$

The magnetisation vector experiences a torque in the presence of $\mathbf{B}$, causing $\mathbf{M}$ to precess around the axis of $\mathbf{B}$ in a clockwise direction at the Larmor frequency. The Larmor frequency for a proton at typical clinical magnetic field strengths is in the Radio Frequency (RF) range, at 42.6MHz per unit of magnetic field strength (Tesla). The local magnetic field can be altered by applying external field gradients. It also changes in the presence of B_0 inhomogeneities caused by system imperfections and subject geometry.

If the transverse magnetisation is non-zero the precession about the longitudinal axis induces current in receiver coils which is recorded as a Free Induction Decay (FID) signal, with real and imaginary parts. Images are often constructed using the absolute value of this complex data but the signal phase can also be informative. The FID is a sinusoidal wave with an amplitude that decays exponentially

over time as the spins relax and the system returns to thermal equilibrium.

2.1.2.2 Relaxation

At thermal equilibrium, M_{xy} is zero and M_z is denoted as M_0. For simplicity M_0 is often set to 1 in signal simulations. This makes it more straightforward to assess the values of M_{xy} and M_z relative to M_0, which is their maximum possible value. The relaxation rate of M_z depends on its value relative to M_0.

The remaining two terms in the Bloch equations (Eq. 2.8 and 2.10) describe how the magnetisation vector relaxes back to its thermal equilibrium orientation after a perturbation. The longitudinal and transverse components evolve according to Eq. 2.8 and Eq. 2.10, respectively. Methods to measure the time constants T_1 and T_2 are included in Section 2.2.1 and the following chapter.

Longitudinal Relaxation (T_1) The solution to the longitudinal term (Eq. 2.9) shows that $M_z(t)$ depends on the equilibrium magnetisation M_0, the starting magnetisation $M_z(0)$, the time delay t and time constant T_1.

$$\frac{dM_z}{dt} = \frac{1}{T_1}(M_0 - M_z) \tag{2.8}$$

$$M_z(t) = M_0 - [M_0 - M_z(0)] \exp\left(\frac{-t}{T_1}\right) \tag{2.9}$$

The longitudinal relaxation time constant T_1 depends on spin-lattice interactions, where energy is transferred between the spins and the neighbouring lattice

of molecules. This energy exchange depends on the movement of the molecules in the lattice and so is temperature dependent. It is also affected by the types of molecules in the lattice, including the concentration of unbound water. Additionally, T_1 increases with the static magnetic field strength.

Transverse Relaxation (T_2) The solution to the transverse relaxation term in the Bloch equation is given in Eq. 2.11, where $M_{xy}(t)$ depends on the initial transverse magnetisation $M_{xy}(0)$, a time delay t and the time constant T_2 which represents the rate of transverse magnetisation decay. This decay is caused by dipole interactions between the isochromat population and neighbouring spins of other nuclei and electrons, which leads to fluctuations in the local magnetic field. This causes the individual magnetisation vectors of the isochromats to become out of phase, shortening the transverse component in a process known as dephasing.

$$\frac{dM_{xy}}{dt} = -\frac{M_{xy}}{T_2} \tag{2.10}$$

$$M_{xy(t)} = M_{xy}(0) \exp\left(\frac{-t}{T_2}\right) \tag{2.11}$$

Dipole interactions depend on the distance between the spins and their relative orientation in terms of the direction of the main magnetic field of the scanner. Free water is able to tumble freely and rapidly and so experiences weak dipole coupling, leading to a relatively long T_2. Spins in biological structures are more restricted and so experience faster dephasing, leading to a shortened T_2. MRI

often makes use of spins in the intracellular or extracellular space, which have a T_2 long enough for the signal to be recorded during an imaging experiment and provide information about the tissue structure and function.

In practice, additional dephasing occurs due to local field distortions from sources such as system imperfections, subject geometry and tissue boundaries. This means M_{xy} decays according to T_2^* (Eq. 2.12), where T_2' is the decay constant due to reversible dephasing (Section 2.1.5). The presence of T_2' in the expression means T_2^* is shorter than T_2.

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{2.12}$$

2.1.3 Radio Frequency Pulses

For simplicity the behaviour of $\mathbf{M}$ can be considered in a frame of reference rotating at ω_0 about the B_0 axis, in which the precession at ω_0 about B_0 appears stationary. This perspective highlights changes in precession frequency relative to ω_0 and is used to isolate the effect of applied RF pulses. The magnetic field $\mathbf{B_1^+}$ of a RF pulse applied along one of the transverse axes in this rotating frame will alter the effective field experienced by the spins, effectively causing $\mathbf{M}$ to rotate about the pulse axis if the pulse is on resonance. For example, if $\mathbf{M}$ is originally aligned with B_0 it will rotate into the transverse plane. In the case of non-adiabatic pulses, such as sinc pulses, the rotation angle θ can be calculated as the integral of the pulse shape over its duration t, as described by Eq. 2.13, where B_1^+ is the pulse amplitude.

$$\Theta = \gamma \int_0^t B_1^+(t)dt \tag{2.13}$$

2.1.4 Spatial Localisation and k-Space

To form an image the spatial location of the acquired signal needs to be determined and mapped. In MRI the signal location is encoded by modulating the spatial distribution of the field strength such that the precession frequencies vary as a function of spatial position. The readout modules of a sequence then sample the signal and an image is decoded from the signal frequencies. Spatial encoding is performed by applying linear magnetic field gradients along the x, y and z axes, which have the effect of adding or subtracting from the static magnetic field B_0, depending on the spatial location.

2.1.4.1 2D Slice Selection

In two-dimensional imaging a specific slice is excited by applying a magnetic field gradient with magnitude G_z, causing the Larmor frequency to vary along the slice direction. A specific slice can be excited by applying an RF pulse with the corresponding frequency band. The desired slice thickness δz determines the necessary frequency bandwidth $\delta\omega$ (Eq. 2.14). The central frequency of the RF pulse must be tuned to the resonance frequency of the spins in the middle of the target slice, as displayed in Fig. 2.1, where the central frequency is $\gamma(B_0 + G_z Z_0)$ and Z_0 denotes the centre of the slice.

$$\begin{aligned}\delta\omega &= \gamma(B_0 + G_z(Z_0 + \frac{\delta z}{2})) - \gamma(B_0 + G_z(Z_0 - \frac{\delta z}{2})) \\ \delta\omega &= \gamma G_z \delta z\end{aligned} \tag{2.14}$$

Ideally a slice-selective excitation pulse will have a rectangular spatial frequency profile with sharply defined edges. This is not the case in practice, as is discussed in Section 2.3.3.2.

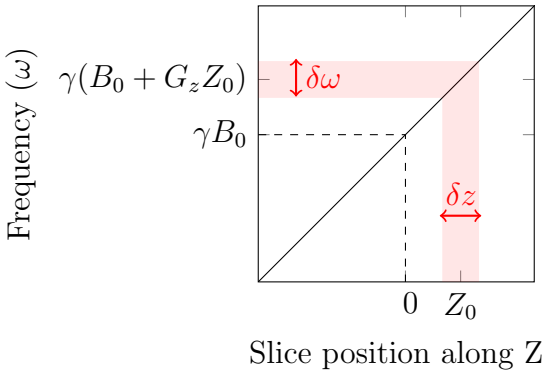

Figure 2.1: Slice-selective excitation. The relationship between the resonant frequency and the position and thickness of a target slice through the B_0 axis.

2.1.4.2 k-Space Sampling

The scanner system samples signal in the spatial frequency-space, otherwise known as k-space. Equation 2.15 describes the relationship between an applied gradient vector $\mathbf{G}$ and the signal S, where $\mathbf{k}(t)$ (Eq. 2.16) represents a 3D k-space with dimensions k_x, k_y and k_z. The majority of the MRI signal intensity comes from the low spatial frequency components of the data, which are near the origin of k-space. The high spatial frequency components contain information about regions of the sample with high resolution contrast, such as physical edges and boundaries. Equation 2.16 shows that the application of gradients for time t moves the system

through k-space, allowing the system to sample signal at multiple points in k-space. Once k-space has been sampled such that the Nyquist sampling criteria have been satisfied a discrete Fourier transform can be applied to the data to convert it to the spatial domain. This produces an image that corresponds to the spatial distribution of the magnetisation of the object in the scanner.

$$S(t) = \int \mathbf{M}(\mathbf{r}) \exp(i2\pi \mathbf{k}(t) \cdot \mathbf{r}) d\mathbf{r} = \int \mathbf{M}(\mathbf{r}) \exp\left(i\gamma \int_0^t \mathbf{G}(\tau) \cdot (\mathbf{r}) d\tau \right) d\mathbf{r} \quad (2.15)$$

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(\tau) d\tau \quad (2.16)$$

Figure 2.2 illustrates Eq. 2.16, showing an example of how the application of gradients enables movement through k-space during a single Repetition Time (TR) of a pulse sequence. The horizontal line in k-space in Fig. 2.2b can be sampled by acquiring signal at the same time as the positive gradient lobe of G_x in Fig. 2.2a. The gradient profiles in Fig. 2.2a are designed such that the FID from the RF excitation is initially dephased, before being *refocused* as a signal *echo* when the k-space trajectory reaches $k_x = 0$. In the sampling trajectory of Fig. 2.2 the G_y gradient moment causes an addition of phase to select a particular k_y line. In succeeding repetitions of the sequence in Fig. 2.2a the G_y moment could be decreased or increased to allow the sampling of different k_y lines.

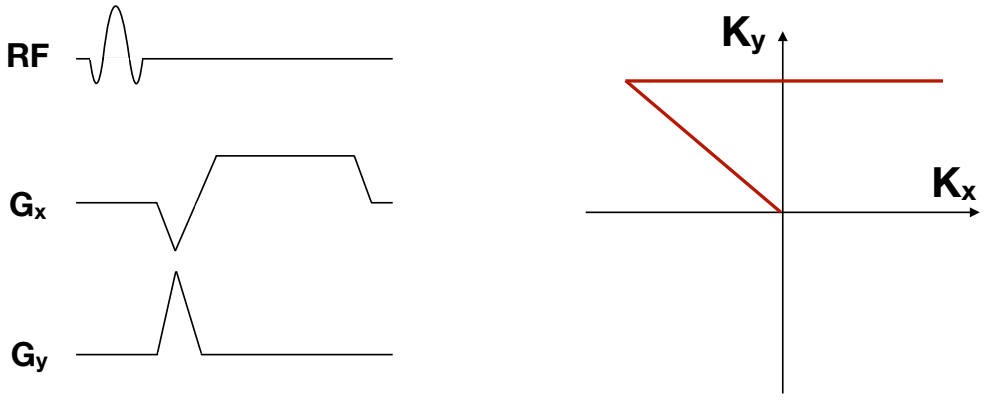

(a) (b)

Figure 2.2: An example of how G_x and G_y gradients move the system through k-space following an RF pulse. A simple pulse sequence for a single TR is shown in (a), with a sinc RF pulse following by gradients in the x and y directions. The resulting k-space trajectory is shown in (b).

In 2D imaging k-space is often sampled on a 2D Cartesian grid. Multiple repetitions of the sampling approach in Fig. 2.2 would be needed to sample different k_y lines to complete a grid. Alternative k-space trajectories sample the entire grid during a single TR and are known as snapshot acquisitions. One snapshot method is Echo Planar Imaging (EPI), which uses the principle illustrated in Fig. 2.2 to traverse through all the k_y lines during a single readout module.

The spacing between the sampled points Δk determines the Field Of View (FOV) of the reconstructed image, as $\Delta k = 1/\text{FOV}$. The highest spatial frequency sampled along a particular axis, k_{max}, determines the pixel width ΔW of the final 2D image in that direction, through the relationship $\Delta W = 1/2k_{max}$.

Non-Cartesian Sampling Although many methods sample a Cartesian grid of k-space, non-Cartesian methods have been used to acquire images quickly and may

have advantages in some applications, such as in cases where imaging acquisition time is limited. For example, spiral trajectories can be used for efficient sampling, often starting at the centre of k-space before moving outwards to sample the desired spatial frequency extent. However, the reconstruction of images from non-Cartesian trajectories, such as spirals, typically requires the data to be mapped to a grid before a Fourier transform is applied.

2.1.5 Spin Echoes

Section 2.1.2.2 introduced the concept of spin dephasing after a RF excitation. The dephasing caused by B_0 inhomogeneities can be reversed by playing a 180 degree refocusing pulse after the initial excitation, which flips the magnetisation about one of the transverse axes, inverting the phase distribution of the spins. Assuming the local field inhomogeneities are unchanged the spins will then begin to move back in phase and produce a signal peak (known as a spin echo) when they are exactly in phase. The elapsed time between the excitation pulse and the occurrence of an echo is known as the Echo Time (TE). For spin echo sequences TE is twice the time between the excitation pulse and the refocusing pulse. Spin echo refocusing reverses the decay governed by T_2^* but the refocused signal peak amplitude still depends on the T_2 of the imaged object. Spin-echo sequences are therefore used to produce signal weighted in terms of T_2, rather T_2^*.

2.1.6 Partial Volume Effects

The signal of each pixel in a two-dimensional MRI image comes from a 3D volume of the scanned object. Often the voxel dimensions are of the order of millimetres,

meaning that multiple tissue types contribute to the signal for each voxel. In this case the acquired signal from a voxel is a weighted average of the signal from the various different tissue types in the voxel. For example, if 90% of a voxel is tissue A and 10% is tissue B then the acquired signal from the voxel will mostly depend on the signal from tissue A. This is known as the partial volume effect and can provide errors and misleading parameter maps in quantitative MRI measurements, where tissue parameters are assigned to each voxel. This is a particularly significant issue in cases of acquisitions with relatively low resolution and when imaging tissue with fine detail and bordering regions with significantly different tissue properties [14]. To mitigate this the voxel dimensions can be reduced, but this reduces the acquired signal from each voxel.

2.2 Quantitative Magnetic Resonance Imaging

MRI has become an established and widely used tool for producing images for diagnosis and disease monitoring in clinical environments. Pathologies and differences between tissue types are often highlighted qualitatively via image contrast [15] but methods have been developed to produce robust quantitative measurements of tissue and system properties.

2.2.1 Single-Parameter Measurements

This section describes gold standard methods for mapping individual tissue and system parameters. Example data and results from implementations of these methods are provided in the next chapter. These techniques were used to obtain reference parameter maps for validation purposes later in the thesis.

2.2.1.1 Proton Density

The term Proton Density (PD) refers to the concentration of proton spins that contribute to the MRI signal. The signal magnitude S of a spin-echo experiment can be described by Eq. 2.17, which indicates that the signal intensity also depends on the relaxation time constants and a scaling factor C.

$$S = C \cdot PD(1 - \exp\left(\frac{-TR}{T_1}\right))(\exp\left(\frac{-TE}{T_2}\right)) \tag{2.17}$$

To isolate the effect of PD on S the effects of the other terms in Eq. 2.17 must be removed, or at least minimised. The T_1 and T_2 weighting of S can be reduced by setting TE $\ll T_2$ and TR $\gg T_1$. However, depending on the T_1 and T_2 of the sample it may be challenging to achieve negligible T_1 and T_2 contrast due to restrictions on the TE and TR, imposed by the sequence design. The scaling factor C in Eq. 2.17 represents the modulation of the proton density map by the profile of the receive coil gain [16, 14] B_1^-. At the field strength used in this thesis (3T) the removal of the B_1^- bias does not have a straightforward solution [17]. To measure water content PD measurements can be calibrated using a sample of pure water. In practice, brain PD measurements can be calibrated using voxels in the Cerebral Spinal Fluid (CSF) of the ventricles, an area with approximately the same proton concentration as pure water [16]. For much of this thesis PD is referred to as M_0, often normalised. Additionally, in this thesis the receiver bias has not been removed from M_0 maps, since the immediate focus of this work was on optimising for reduced T_1 and T_2 measurement error. A method for correcting for the receiver bias at 3T has been evaluated in the literature [18], which uses

the ratio of images acquired using a head coil and a body coil as receivers. This method would likely increase the total scan time needed. It would be valuable for future research to explore an appropriate method for correcting for the recieve bias of quantitative MRI maps without significiantly increasing the scan duration, either with the method from [18] or by an alternative technique.

Stroke and Proton Density As introduced in Chapter 1, stroke damage causes the local water content to increase, meaning that proton density mapping can provide indirect measurements of oedema progression and the effect of medication [16]. A proton density measurement technique with a measurement error smaller than the change in water content would allow detection of oedema changes over a longitudinal study. Qualitative images are unsuitable for this use as it is likely that the contrast would change between scans. Additionally, quantitative measurements of water content can be directly compared with normal values for a given tissue type (such as grey matter).

2.2.1.2 T_1 and T_2

Commonly used gold standard methods use spin-echo sequences with and without an initial inversion pulse to measure T_1 and T_2, respectively. To map T_1, fully relaxed initial magnetisation is assumed and an inversion pulse is applied to invert the sign of M_z. This is followed by an excitation pulse to rotate the magnetisation into the transverse plane before an image is acquired, using the shortest possible TE. The delay between the inversion and excitation pulse is known as the Inversion Time (TI). For a constant TE the magnetisation component at the time of the excitation pulse varies with TI, following the model in equation 2.18.

Equation 2.18 is a specific case of the longitudinal relaxation model described in Section 2.1.2.2. Figure 2.3 demonstrates examples of M_z recovery after an initial inversion. Using Eq. 2.18 it can be calculated that after TI $= 5T_1$ the longitudinal magnetisation M_z has returned to 98.7% of M_0 and can be regarded as fully relaxed. In experiments the time constant T_1 is determined by measuring points on the recovery curves in Fig. 2.3 and fitting a model based on Eq. 2.18.

$$M_z = M_0(1 - 2\exp\left(\frac{-TI}{T_1}\right)) \tag{2.18}$$

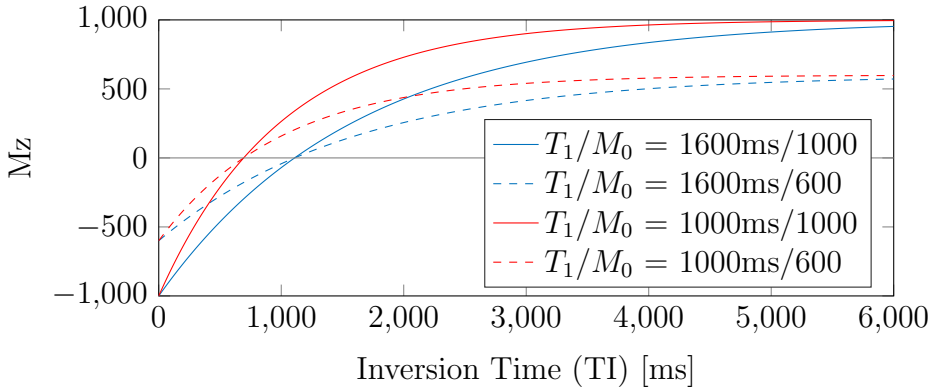

Figure 2.3: Examples of T_1 inversion recovery curves for different T_1 and proton density combinations.

The transverse magnetisation can be manipulated by varying the TE of a spin-echo experiment without an initial inversion pulse (Eq. 2.19). The parameter T_2 is measured by sampling points on the decay curve in Fig 2.4 and fitting a model based on Eq. 2.19.

$$M_{xy} = M_0\exp\left(\frac{-TE}{T_2}\right) \tag{2.19}$$

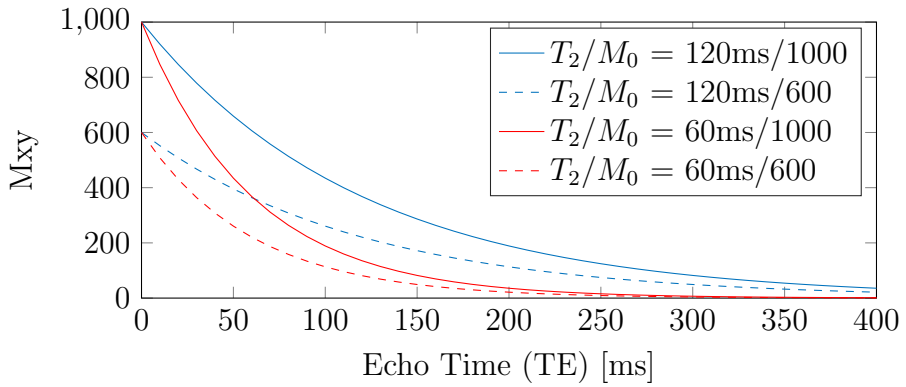

Figure 2.4: Examples of T_2 exponential decay curves for different T_2 and proton density combinations.

The gold standard T_1 and T_2 mapping approaches mentioned in this section often require scan times that are too long to be suitable for clinical applications. However, they provide good reference maps for validating new sequences using phantom acquisitions, where more scan time is available. There has been work to design parameter-mapping sequences with reduced scan durations, these include Look-Locker [19] and DESPOT1 [20] for T_1 and DESPOT2 [21, 22] for T_2. These examples do not require the magnetisation to fully recover between measurements and DESPOT1 and DESPOT2 can provide 3D T_1 and T_2 maps in a clinically feasible scan time [21, 22]. The DESPOT1 and DESPOT2 sequences use driven equilibrium signal from SPoiled Gradient Recalled (SPGR) and Steady State Free Precession (SSFP) sequences, respectively. The DESPOT2 T_2 measurements rely on a previously obtained T_1 measurement, which can be derived from the DESPOT1 data.

2.2.1.3 B_0 Off-Resonance and B_1^+

The resonant frequency relative to that of the externally applied B_0 is known as the *off-resonance* and can be calculated via the phase difference between gradient-echo images acquired at different echo times [23]. This is possible because over time the spins at each voxel location accrue phase according to the local field strength. Practical details of this method are provided in Chapter 3.

The homogeneity of the excitation RF field B_1^+ affects the measured signal intensity and so has an effect on quantitative imaging, particularly T_1 mapping. Several methods for mapping the actual flip angle delivered by the scanner have been designed and compared [24], where the actual flip angle is encoded in image ratios. The measured flip angle can be compared with the prescribed flip angle to quantify how much the system is over or under-delivering with respect to the requested B_1^+ strength. Some methods use signal intensities [25, 26, 27] whereas others require the image phase [28].

2.2.2 Multiple Parameters

Some sequences have been developed with the aim of making simultaneous measurements of multiple parameters in a clinically reasonable total scan time. The QRAPTEST [29] and QRAPMASTER [30] methods enable measurements of T_1, T_2, M_0 and B_1^+ with full brain coverage in approximately 5 minutes. An approach known as Multi-Parameter Mapping (MPM) [31] measures T_1, effective M_0, effective T_2^* and Magnetisation Transfer (MT). Another study demonstrated the acquisition of 3D T_1, T_2, T_2^* and magnetic susceptibility maps within 15 minutes of scan time, via a series of sequences [32]. Sequences that make use of the be-

haviour of spins in a steady-state regime have also been used for simultaneous multi-parameter mapping. These include the Double Echo Steady-State (DESS) [33] and Triple Echo Steady-State (TESS) [34] sequences. Simultaneous Multi-Slice (SMS) acquisitions have been used in conjunction with TESS [35] to allow the acquisition of whole brain T_1, T_2, M_0 and B_0 maps, with increased efficiency in terms of parameter error per unit scan time per slice.

2.2.3 Limitations of Conventional Methods

As mentioned in Section 2.2.1.2, a drawback of traditional quantitative imaging, particularly the gold standard methods, is that the required scan time is too long to be feasible for clinical applications. The long scan duration can be due to the need for signal-averaging and serial acquisitions to measure multiple tissue or system properties. The DESPOT1 and DESPOT2 methods have been introduced in 2.2.1.2 as more efficient single-parameter alternatives. However, they are sensitive to deviations from the prescribed flip angle, which are mainly caused by B_1^+ inhomogeneities. DESPOT1 is also sensitive to the actual prescribed flip angles [22]. A comparison of T_1 mapping methods [36] observed an overestimation of T_1 from DESPOT1.

Additionally, quantitative MRI methods in general must deal with sources of measurement error such as static field (B_0) inhomogeneity and deviations from the prescribed flip angle, as well as subject variability and movement. Imperfect excitation RF slice profiles also cause deviations in the achieved flip angle as a function of the position along the slice direction. This causes the acquired signal (integrated through the slice direction) to differ from modelled values if the

true slice profile is not modelled. The shape of the slice profile also depends on the prescribed flip angle. Additional acquisitions or post-processing to mitigate these inaccuracies can further increase the scan duration. Robust solutions to the problems associated with achieving accurate and precise quantitative parameter maps will help make quantitative MR more commonly used in clinics.

2.3 Magnetic Resonance Fingerprinting

The work in this thesis uses a technique called Magnetic Resonance Fingerprinting (MRF) [1]. This section describes the principles of MRF and introduces some of the published research that has been conducted to improve aspects of it.

2.3.1 Original Framework

MRF was developed to make simultaneous measurements of multiple tissue properties, such as T_1 and T_2, using a single acquisition. The method allows short scan durations and the original MRF publication showed an improvement in efficiency compared with two single-parameter sequences (DESPOT1 and DESPOT2) [1].

During a MRF experiment a sequence with a pseudo-random set of parameters is used to manipulate the spins in a sample and acquire images. In the ideal case each tissue type within a sample produces a unique signal timecourse through the images. This can be achieved by varying the time between subsequent RF pulses and the size of any RF-induced rotations in the net magnetisation. For a given pixel the signal timeseries over the whole experiment is then compared to a previously built dictionary of simulated signal patterns in order to find the closest match. The dictionary is compiled by simulating the signal expected from

the particular acquisition parameters and sequence structure for all combinations of the chosen tissue and system parameters (e.g. T_1 and T_2) over given ranges. Finally, the tissue parameters used to simulate the best matched dictionary time-course are assigned to the pixel. This process is repeated for every pixel. MRF is commonly implemented with interleaved undersampling of k-space to achieve fast data acquisition, where the k-space trajectory is rotated between each image acquisition to produce some incoherence between the underlying signal and the image artefacts.

2.3.2 Subsequent Developments

The original MRF proof-of-principle [1] was based on a balanced SSFP (bSSFP) sequence with an initial inversion pulse and produced maps of T_1, T_2, modulated M_0 and B_0 inhomogeneity, in an acquisition time of approximately 12 seconds per slice. TR and FA were varied and the TE was set as half of the corresponding TR, as is typical for bSSFP acquisitions. Subsequent developments have been made with regards to the sequence design, image reconstruction and parameter assignment.

FISP MRF An unbalanced version of the original bSSFP sequence (i.e. FISP) has become widely used for MRF due to its inherent insensitivity to B_0 inhomogeneity [37, 38]. Subsequent versions have extended FISP MRF to a 3D acquisition, giving whole brain T_1, T_2 and modulated M_0 maps in less than 5 minutes [39]. In the original implementation of FISP MRF the TE was fixed throughout the acquisition [37] but the FA schedule pattern was smoother than that used for the original bSSFP MRF work [1].

Additional Parameter Maps Alterations to the original MRF framework have enabled other tissue properties and system characteristics to be measured. Additional preparation pulses have been used to evoke sensitivity to diffusion [40, 41] and perfusion [42]. Christen *et al.* [40] used a cerebral vascular model to estimate maps of cerebral blood volume, mean vessel radius and blood oxygen saturation. In [40] a hybrid gradient-echo and spin-echo sequence was used to sample gradient echoes before and after the spin echo, as well as the spin echo signal itself. Signal evolution curves were acquired with and without contrast agent, with the ratio of the two curves used as a fingerprint. Anderson *et al.* [43] detected and quantified concentrations of multiple contrast agents after simultaneous administration. Hong *et al.* [44] simultaneously estimated T_1 , T_2 , $T_2{}^*$, M_0 and B_0 using a hybrid of FISP and spoiled gradient-echo squences. Methods have also been developed to account for, and map, the partial volume effect [45, 46, 47].

The B_1^+ field can be mapped by including an extra dimension in the dictionary. This approach requires that the effects of T_1 and B_1^+ variations can be distinguished. One way of ensuring this is to design the RF flip angle schedule [48, 49] such that the dictionary fingerprints vary differently as a function of B_1^+ and T_1. An alternative to accounting for B_1^+ variations in the dictionary is to use an externally acquired reference B_1^+ efficiency map to scale the flip angles during the production of the dictionary, generating a tailored dictionary for each pixel [39].

Dictionary Simulation When a dictionary is used to produce MRF parameter maps each fingerprint is assigned to one of the dictionary entries. This means that the sizes of the dictionary dimension increments affect the accuracy and precision

of the resulting parameter maps. For example, large steps in T_1 could worsen the parameter map accuracy if the true T_1 values are between the dictionary T_1 values.

The incorporation of slice profile simulations during the dictionary generation has been shown to improve parameter assignment accuracy [50, 51] for flip angle regimes often used in MRF sequences. Reductions in dictionary dimensionality [52] and clustering [53] have enabled more efficient dictionary searching, decreasing the computational cost of the matching process. One of the examples in [52] was the temporal compression of a bSSFP MRF dictionary with 363624 entries and 1000 timepoints. When 200 singular values were used to construct this compressed dictionary the matching time (using phantom data) reduced from 31 seconds to 7.3 seconds and there were only 5 pixels with T_1 or T_2 values that did not exactly match those obtained using the uncompressed dictionary. Each of these differences for those pixels was one step in the corresponding dictionary parameter resolution, highlighting the impact of the dictionary step size on the matching error.

Image Acquisition Simultaneous multislice acquisitions have been used to further reduce the acquisition time per slice [54, 55] and alternative sampling strategies have been explored, such as radial [56] and Cartesian trajectories.

Examples of Cartesian sampling in MRF have included the acquisition of single phase-encoding lines [57, 58] and a fully sampled EPI readout [59, 41] per TR. The sequence in [41] used pairs of refocusing pulses, allowing sensitivity to T_2. Acquiring a fully sampled image for each timepoint reduces the image artefacts and requires comparatively straightforward image reconstruction.

Parameter Map Reconstruction Reconstruction algorithms and data models have been used to produce parameter maps from the k-space data, with improved parameter accuracy [60, 61, 62]. Amthor *et al.* [63] applied a machine learning algorithm to the MRF matching process to provide tissue parameter maps. Various algorithms were compared for their performance in classifying fingerprints from a bSSFP MRF sequence, with an acquisition made substantially longer because of Cartesian sampling. The data were classified into clusters of T_1 and T_2 combinations with no exact assignment of B_0. Deep learning [64] has also been used to provide parameter estimates [65, 66].

Schedule Optimisation Cohen *et al.* [67] optimised the acquisition parameter schedule of a gradient-echo EPI sequence by minimising the dot products between the dictionary entries. In [67] a collection of algorithms was used to optimise MRF sequences with respect to the global dictionary dot product. An interior-point algorithm was shown to outperform others in terms of the final cost function minimum, the required function iterations and the scan duration of the optimised schedule. In simulations the represented algorithms provided relatively large T_2 percentage errors compared to T_1. However, the interior-point optimisation provided substantial improvements in T_2 compared to the other algorithms. MRF experiment simulations have also been used to estimate the parameter assignment error at each optimisation iteration and aim for a schedule that minimises this error, such as the work by Hamilton *et al.* [68] which used a genetic algorithm to optimise a FISP MRF schedule. Other work has used Monte Carlo simulations to assess the suitability of a given schedule [69].

2.3.3 Signal Modelling

This section includes points to consider when modelling the MRF signal, including comments on the signal modelling approach used for FISP MRF and a discussion on excitation slice profiles.

2.3.3.1 Extended Phase Graphs

Solutions to the Bloch equations are used to model the behaviour of isochromats and simulate the signals of bSSFP MRF experiments, using a series of matrix rotations and exponential relaxation terms. In FISP experiments the dephasing during each TR is not refocused and the acquired signal of each voxel is an average of all the spins inside it. The isochromat model used in bSSFP simulations is not practical for this application because the signal from a large number of isochromats would need to be simulated. The Extended Phase Graph (EPG) formalism is a practical way of modelling signal evolutions throughout unbalanced sequences such as FISP MRF [37], simulating the magnitude and temporal positioning of refocused signal echoes. The EPG formalism can model the effect of RF pulses and gradients on the magnetisation by considering intra-voxel phase distributions in terms of Fourier states, often in units of 2π. Figure 2.5 shows an example EPG diagram for a simple pulse sequence, with phase states on the y-axis and time on the x-axis. In Fig. 2.5 the second RF pulse refocuses a component of the magnetisation dephased by the first pulse, causing a signal echo. The rate of transition between states depends on the gradient amplitude. The effects of transverse and longitudinal relaxation on the magnitude of the signal echoes can also be modelled. More details about the principles behind extended phase graphs

can be found in [70].

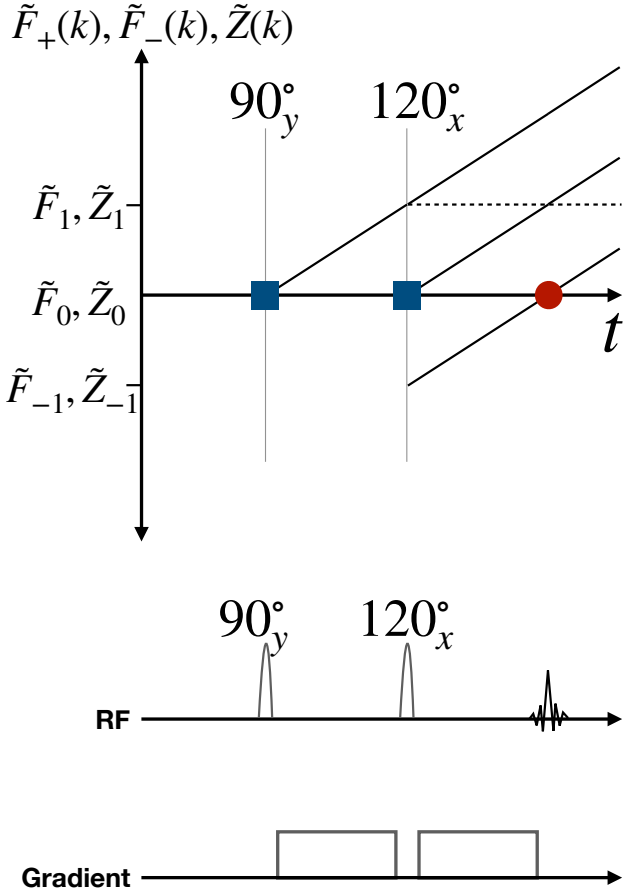

Figure 2.5: An example EPG diagram for a simple pulse sequence with two RF pulses and an applied gradient. In this case the first and second pulse deliver rotations about the y and x axes, respectively. The dephasing and rephasing states are represented by $\tilde{F}_+(k)$ and $\tilde{F}_-(k)$, respectively. The longitudinal states are represented by $\tilde{Z}(k)$. The parameter k represents k-space. Paths are shown to display the transition of the transverse (solid) and longitudinal (dashed) magnetisation between states. Note that in this case the longitudinal magnetisation is not dephased by the gradient. The crossing of the x-axis (red circle) represents a signal echo formation.

2.3.3.2 Excitation Slice Profiles

Pulse and Gradient Shapes An example slice-selective RF sinc pulse and the corresponding gradient from a balanced SSFP sequence is shown in Fig. 2.6. The RF pulse has the amplitude required to deliver a 90 degree rotation angle at the slice centre and the gradient has the amplitude required to excite a slice 5mm thick (full width half maximum). The gradient applied during the RF pulse has the effect of dephasing the transverse magnetisation across the slice, meaning that the phase of the transverse magnetisation depends on the position along the slice direction. The purpose of the gradient lobe after the RF pulse is to refocus this dephasing. The pre-dephaser lobe before the pulse is specific to balanced SSFP and is not present in FISP and spin-echo pulse sequences.

Non-Ideal Excitation Profile The effect of the RF pulse is to rotate the magnetisation into the transverse plane to produce signal. Ideally, and assuming zero transverse components before the pulse, the signal outside the slice after the slice-selective pulse will remain zero and the spins at all positions through the slice direction will be identically manipulated, producing a top-hat signal magnitude profile through the slice. In practice there is a gradual change in the signal magnitude across the slice edges (Fig. 2.7). Since the acquired signal is the signal integral across the slice the signal acquired in practice is lower than in the ideal case.

Simulation For small flip angles the relationship between the flip angle and the RF field is approximately linear, meaning the slice profile can be estimated by the Fourier transform of the RF pulse shape. For off-resonance pulses or high flip

angles the relationship becomes non-linear and numerical simulations are required to determine the slice profile [12]. The numerically simulated magnetisation profile resulting from the RF sinc pulse and gradient profiles from Fig. 2.6 is shown in Fig. 2.7. The pulse profile resembles a Gaussian shape because of the time-bandwidth product of the pulse. The profile of a sinc pulse with a higher bandwidth product would have visible lobes either side of the main peak. The magnetisation at a given position in the slice was a result of the net rotation at that point, which was calculated by simulating the effect of the RF pulse profile at discrete timepoints.

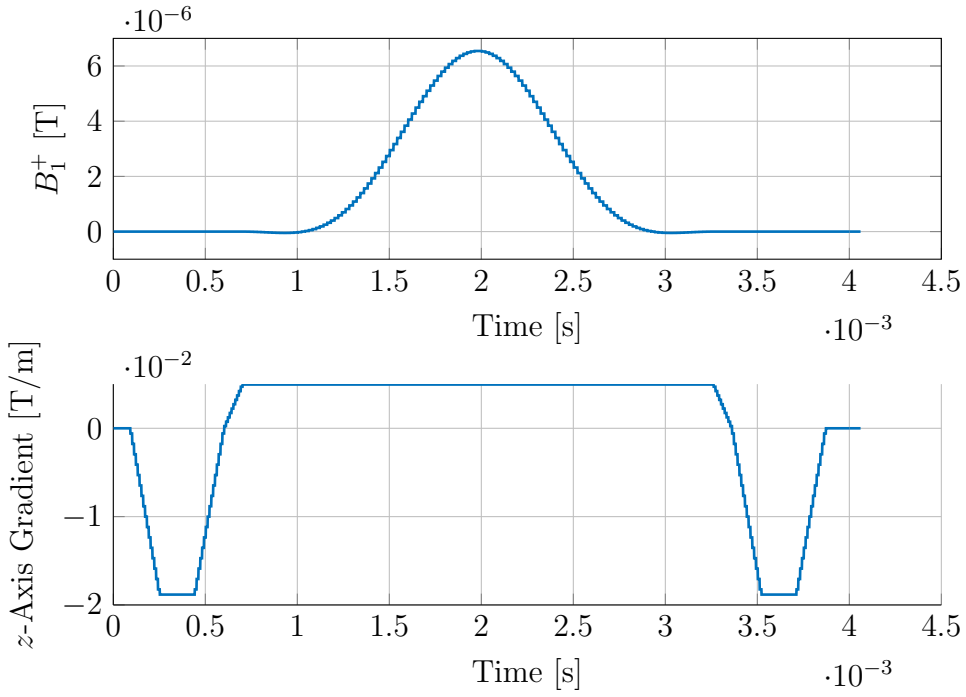

Figure 2.6: Example sinc pulse profile (top) and refocusing gradient profile (bottom) for a 90 degree flip angle and 5mm slice thickness, taken from a balanced SSFP sequence.

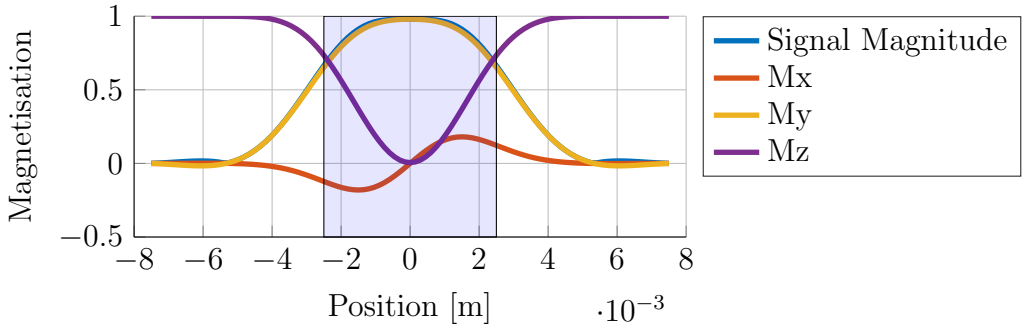

Figure 2.7: Magnetisation from the example RF pulse profile and z-axis gradients in Fig. 2.6.

2.3.4 Limitations

When choosing an image acquisition method for MRF it is important to consider the trade-off between the image quality and the number of images required. The majority of MRF methods use undersampled image acquisitions, allowing a large number of measurements in a reasonable timeframe. For example, the interleaved undersampling trajectory used in [1] required between approximately 11ms to 13.5ms per timepoint, although this did mean that the images contained severe undersampling artefacts. Despite the time-varying nature of the artefacts, due to a rotation of the trajectory between timepoints, the signal fluctuations caused by the artefacts meant that large numbers of timepoints (e.g. 1000) were required to obtain good fits. Fully sampled EPI images contain less severe artefacts and so can effectively provide a larger SNR. This means that although fully sampled EPI MRF sequences take longer to acquire each timepoint the overall scan duration will not necessarily be extended because fewer timepoints are needed.

Further advances may be made by achieving an increased signal-to-noise ratio at each acquired timepoint, sensitivity to additional tissue parameters or by

improving the accuracy and temporal efficiency of the acquisition. A key area of research is the choice of base sequence design and the optimisation of the sequence acquisition schedule, including the choice of algorithm and cost function.

Chapter 3

Implementation of Reference Measurement Methods

This chapter begins with details of a custom phantom (Section 3.1) that was constructed for characterising the scanner performance and testing sequence implementations in later sections of this thesis. This is followed by demonstrations of the various implementations of reference quantitative MRI methods to map T_1, T_2, excitation B_1 and B_0 (Sections 3.2 to 3.5). These methods were used later in the thesis to provide reference maps for validating the optimised MRF sequences, using data from phantom and *in vivo* scans. The final section describes how spiral MRF sequences from the literature were replicated and contains example results (Section 3.6). All the scan data in this chapter were acquired using a MAGNETOM Verio 3T scanner (Siemens Healthcare, Erlangen) and a 32-channel head coil. Unless stated otherwise, all image analysis and parameter map calculations were performed offline in MATLAB (Mathworks, Natick, MA, USA).

3.1 Custom Phantom

A phantom was constructed with T_1 and T_2 values spanning the range typically observed across grey and white matter in the *in vivo* human brain at 3T. There is considerable variation in the literature values [14, 71] but approximate ranges for grey matter have been reported as T_1 = 1300ms to 1700ms and T_2 = 80ms to 120ms. Measured white matter values are around T_1 = 800ms to 1000ms and T_2 = 60ms to 80ms. The phantom was used for all sequence development scans in this thesis, with the variation of relaxation properties across the phantom allowing exploration of the performance of the relevant mapping techniques.

Design The internal structure (designed originally as part of another study [72]) of the phantom was 3D-printed and held six 100ml cylindrical tubes (with conical ends) in a fixed position within a 4000ml cylindrical bottle. Figure 3.1 shows sagittal, coronal and axial views of the phantom.

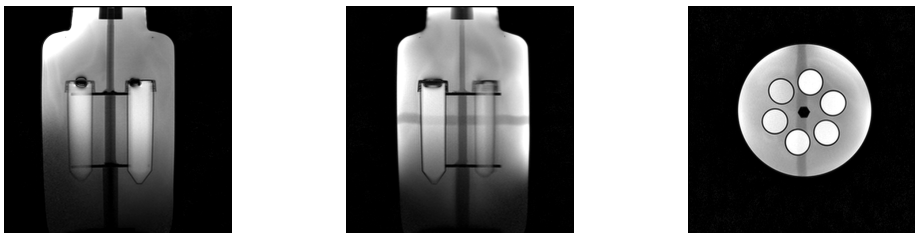

(a) Sagittal (b) Coronal (c) Axial

Figure 3.1: Siemens localiser scan views of the custom phantom.

Each of the six tubes contained solid agar gel doped with Nickel Chloride (NiCl), in varying concentrations such that each tube had different MR relaxation properties. The concentrations in [73] were used as a guide for the approximate concentrations needed to achieve the desired T_1 and T_2 range. Table 3.1 shows

the approximate concentrations of agar and Nickel Chloride used for each tube. The gels were formed by mixing powdered agar and Nickel Chloride with distilled water. The 4000ml bottle contained only distilled water. The T_1 and T_2 of the phantom tubes were characterised in Sections 3.2 and 3.3, showing that T_1 varied across the tubes from approximately 840ms to 1640ms (Table 3.2) and T_2 varied from approximately 60ms to 130ms (Table 3.3). These values were measured at scanner room temperature, which was approximately maintained for all subsequent acquisitions. Therefore, assuming the temperature was relatively consistent, these gold standard T_1 and T_2 measurements should be valid for all the subsequent fingerprinting experiments and comparisons.

Tube Index	**Agar [mg]**	**Nickel Chloride [mg]**	**Distilled Water [ml]**
1	1053	8	100
2	1209	10	100
3	1412	13	100
4	1583	19	100
5	1816	21	100
6	1997	24	100

Table 3.1: Concentrations in each tube of the phantom.

3.2 T_1: Inversion Recovery with a Variable Inversion Time

To map T_1 the gold standard inversion recovery method was used, as introduced in Chapter 2, which involved a variable post-inversion delay followed by a spin-echo EPI readout. Practical details of the implementation and example results are included in this section.

Signal Model In the previous chapter it was stated that after a complete inversion the longitudinal relaxation evolves according to the recovery curve in Eq. 2.18. However, this model assumes a perfect inversion of the magnetisation which in practice may not be achieved. Therefore it is necessary to modify the model to include a parameter for *inversion efficiency* β, as shown in Eq. 3.1 for the case of a magnitude (positive rectified) signal. The β parameter in Eq. 3.1 is the fraction of the fully relaxed M_z that has been inverted. In the case of a full inversion $\beta = 1$.

$$|S_{TI}| = \left| M_0(1 - 2\beta exp(\frac{-TI}{T_1})) \right| \tag{3.1}$$

Figure 3.2 demonstrates how the longitudinal magnetisation recovery curve is affected by changes in β or M_0. For example, in cases where $\beta < 1$ the signal passes through 0 at a shorter TI. This effect would lead to inaccurate T_1 measurements in these cases if β was not included in the model.

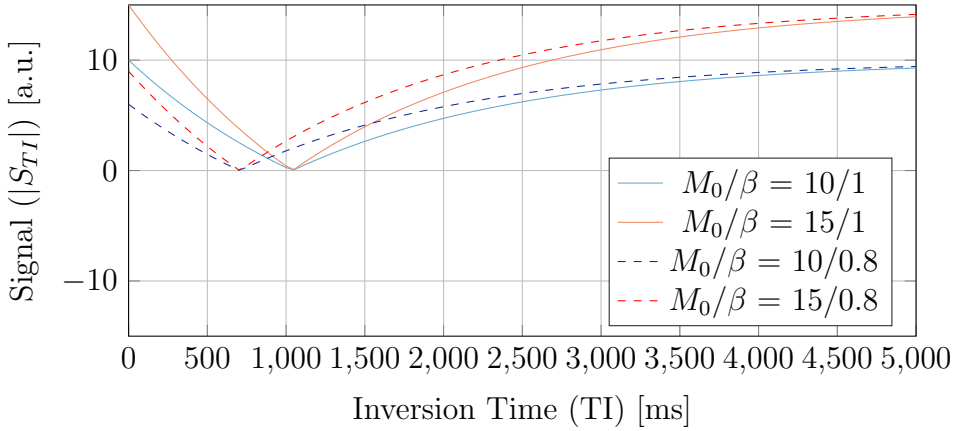

Figure 3.2: Simulated examples of how changes in proton density (reflected in M_0) and inversion efficiency β affect the longitudinal relaxation curve of the model in Eq. 3.1. Here $T_1 = 1500$ms for all the curves.

Data Acquisition Data were acquired for the T_1 fit using a Siemens inversion recovery SE EPI sequence, with a long TR (30000ms) to allow the phantom magnetisation sufficient time to fully relax between measurements. Using Eq. 2.9 it can be seen that 30 seconds after an inversion $M_z = 0.995M_0$ (for $T_1 = 5000$ms) and so is essentially fully relaxed. It was assumed that the maximum T_1 in the phantom would not exceed 5000ms, even in the regions of pure water. This assumption was made because CerebroSpinal Fluid (CSF) has high water content [14] and a T_1 of about 4200ms at 3T [74, 75]. The inversion time was varied for 12 measurements, using TI = [34, 70, 140, 280, 560, 750, 1120, 1500, 2240, 3000, 3500, 5000] ms. The use of 12 measurements resulted in a total measurement time of 12×30000ms = 360 seconds (6 minutes). To maximise the signal a short TE was used (21ms), kept constant to avoid T_2 contrast between the data points. The field of view was 250mm × 250mm and the voxel size was 3.9mm × 3.9mm × 5mm. A partial Fourier factor of 6/8 was used (in the phase-encoding direction) to accelerate the readout and enable a short TE. The readout bandwidth was 2298Hz per pixel.

Phantom T_1 Fitting The software package QMAP [76] was used to perform a non-linear least squares fit of Eq. 3.1 to the scan data, giving single-slice maps of T_1, M_0 and β. Examples of these maps for the custom phantom are provided in Fig. 3.3. Figure 3.4 shows the T_1 distribution across the six gel-filled tubes in the phantom and Fig. 3.5 shows two examples of the fitted recovery curves, with corresponding Bland-Altman plots in Fig. 3.6. Figure 3.6 shows that for both the examples only one of the fitted points was outside the 95% confidence interval. The T_1 mean and standard deviation for each tube is shown in Table

3.2. The phantom T_1 values in Table 3.2 approximately represent the range typically observed in human brain structures at 3T [14]. The distributions of values within each tube was relatively homogeneous: the measurement standard deviations were within ≈1-4%, with those of tube 1 having the worst precision. Table 3.2 shows that Tube 3 and Tube 4 had very similar T_1 values, despite having different doping concentrations (Table 3.1). This could have caused by the choice of TI, particularly around the point that the inversion recovery curve meets the x-axis (i.e. the signal null point). The timing of the null point depends on T_1. For magnitude data, if the recovery curve is not sufficiently sampled around the null point the T_1 may be incorrectly fitted. Additionally, noise rectification for low SNR magnitude data points around the null point could cause the temporal position of the null point be incorrectly determined. To overcome this the image phase information could be used and a model could be fitted to the real data rather than the magnitude data [77].

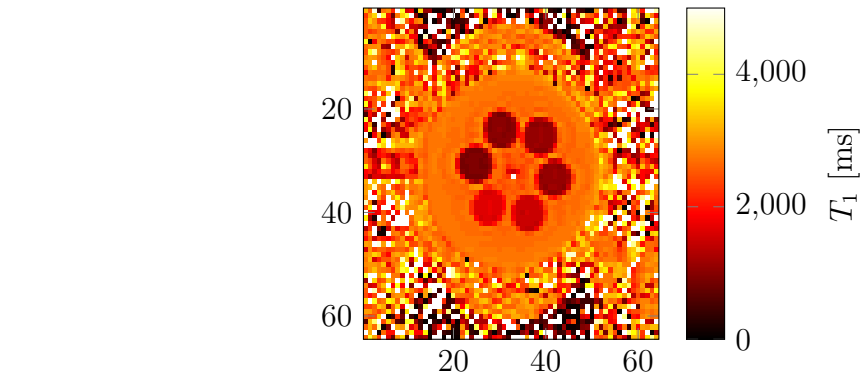

(a) Fitted T_1 map

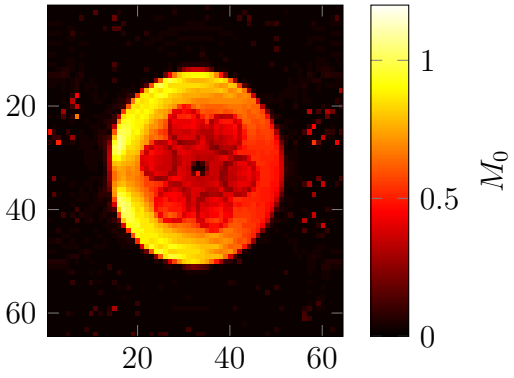

(b) Fitted T_1: M_0 map

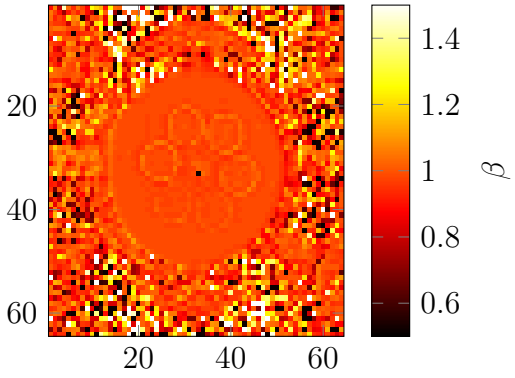

(c) Fitted T_1: Beta efficiency parameter

Figure 3.3: Fitted parameter maps of the phantom, obtained using the gold standard inversion recovery model.

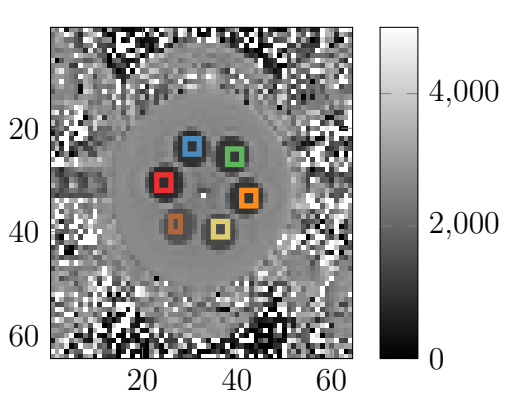

(a) ROI Positions

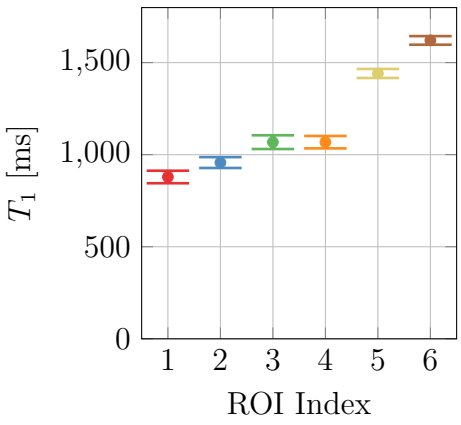

(b) ROI T_1 distributions

Figure 3.4: T_1 distributions in regions of the custom phantom.

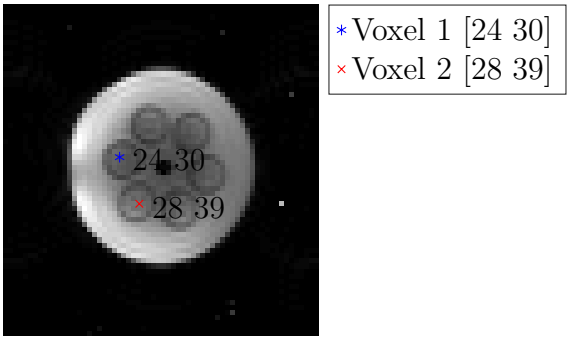

(a) Selected voxels

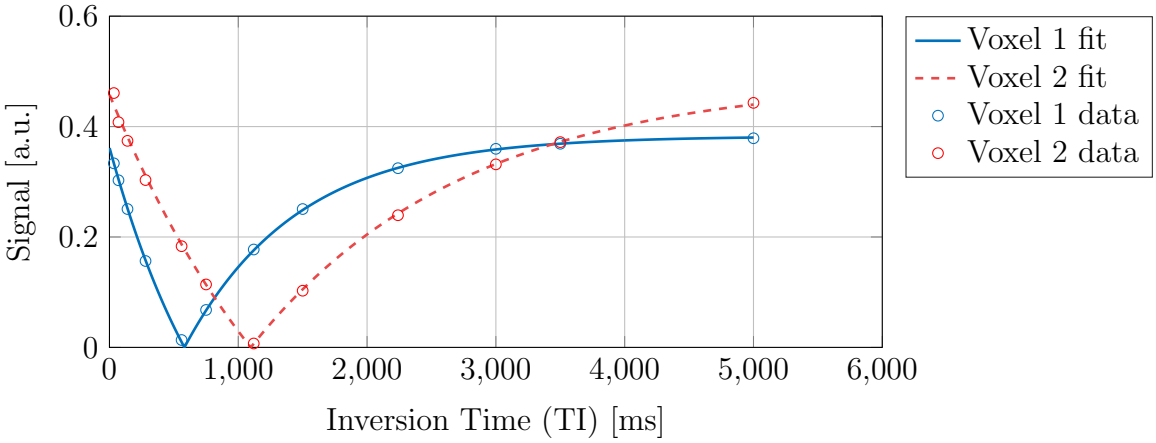

(b) T_1 fit for selected voxels.

Figure 3.5: T_1 fits for the voxels from the tubes with the longest and shortest T_1.

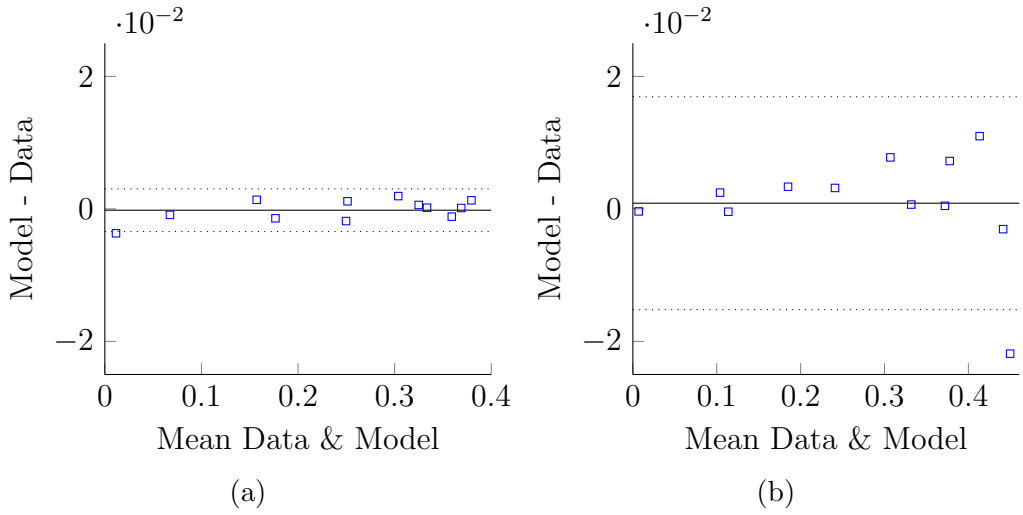

Figure 3.6: Bland-Altman plots for the T_1 fits for the voxels from the tubes with the shortest (a) and longest (b) T_1, showing the mean difference between the data and the fitted points (solid). The dashed lines mark the 95% confidence interval.

Tube Index	T_1 **[ms]**
1	879±34
2	957±30
3	1068±37
4	1068±34
5	1441±25
6	1621±23

Table 3.2: Measured T_1 for each of the phantom tubes (data from Fig. 3.4).

Minimum Scan Time The longest T_1 of the phantom tubes was measured as 1621ms ±23ms (Table 3.2). By inserting the upper value of this measurement into Eq. 2.9 one can calculate that 10000ms after a full inversion M_z has effectively recovered ($M_z = 0.995M_0$). This means that it is suitable for the gold standard protocol for mapping T_1 of the tubes to have a TR of 10000ms, reducing the scan time to 120 seconds (2 minutes) per slice. Since the slice thickness was 5mm, acquiring the equivalent of full-brain coverage (e.g. 20 slices with slice gap =

50%) with this protocol would require a scan time of 2400 seconds (40 minutes).

3.3 T_2: Variable Single-Echo Echo Time

For the reference T_2 measurements the gold standard single-echo spin-echo method was used, as introduced in Chapter 2, with various echo times.

Data Acquisition Multiple single-slice images were acquired, each with a different TE, using a Siemens spin-echo EPI sequence. The echo times and repetition times were (min:increment:max) TE = [21:4:69, 80, 90, 110, 200, 400, 800, 1200, 1600, 3200] ms and TR = 30000ms, making the total scan time 660 seconds (11 minutes). Other acquisition parameters were: field of view = 250mm × 250mm, in-plane pixel size = 3.9mm × 3.9mm, slice thickness = 5mm, readout bandwidth = 2298Hz per pixel and partial Fourier factor = 6/8.

Model and Fitting As each acquisition represented a measurement on the transverse magnetisation decay curve an exponential decay model was fitted (Eq. 3.2) to the magnitude images to obtain the time constant T_2. This model is identical to Eq. 2.19 except for the addition of an offset term c to account for positive signal offsets. For magnitude data with very low SNR (i.e. below approximately SNR = 3) a positive contribution to this offset comes from the rectification of noise during the imaging acquisition and reconstruction [78]. This means that the transverse relaxation curve "plateaus" when TE is long compared to T_2. Including a offset term allows better fitting of the data for these cases but introduces some bias for measurements with higher SNR (e.g. shorter TE), potentially introducing error in the fitted T_2. Other methods to account for noise rectification

include removing the plateaued points from the fit, as demonstrated for T_2^* in [79], or squaring the signal [80]. Although a simple offset correction was used in this thesis (Eq. 3.2) the methods [79] in [80] would be worth considering for future research.

The work of Milford *et al.* demonstrated that imperfect refocusing pulses can also contribute to the signal offset [81]. Although the work of [81] considered a T_2-mapping method with multiple echoes measured after each excitation pulse a similar effect could conceivably occur for the single-echo method performed in this chapter.

$$S_{TE} = M_0 exp(\frac{-TE}{T_2}) + c \tag{3.2}$$

To perform the fit the MATLAB *fit* function was used with the default algorithm. Fitted parameter maps of the custom phantom are shown in Fig. 3.7. Figure 3.8 shows the T_2 distribution across the six tubes in the phantom. Example fits for voxels from the tubes with longest and shortest T_2 are shown in Fig. 3.9, with corresponding Bland-Altman plots in Fig. 3.10. As for the T_1 results, Fig. 3.10 shows that the majority of the fitted points were within the 95% confidence interval. However, both plots in Fig. 3.10 show a positive bias.

Minimum Scan Time The scan time required by the T_2 protocol can be reduced by acquiring fewer measurements. Table 3.3 compares the tube T_2 values obtained from fitting with 22 and 9 measurements from the same data set. The 9 different TEs were [21, 33, 45, 65, 80, 110, 200, 400, 800] ms. Using 9 measurements reduces the scan duration for a single slice to 270 seconds (4 minutes 30

seconds), for TR = 30 seconds. The two sets of values in Table 3.3 are in very close agreement and approximately cover the literature values for grey and white matter [82, 71]. The measured standard deviations in the tubes were between ≈6-10% of the respective means. Furthermore, using the minimum TR from Section 3.2 (i.e. 10 seconds) the scan duration could be reduced to 90 seconds (1 minute 30 seconds) per slice. This would translate to a scan duration of 30 minutes to acquire 20 slices.

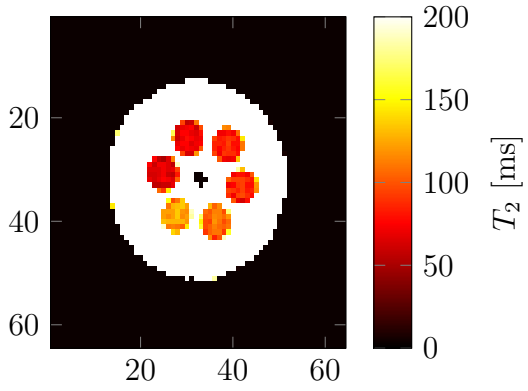

(a) T_2 map: colour scale to highlight the T_2 of the small tubes.

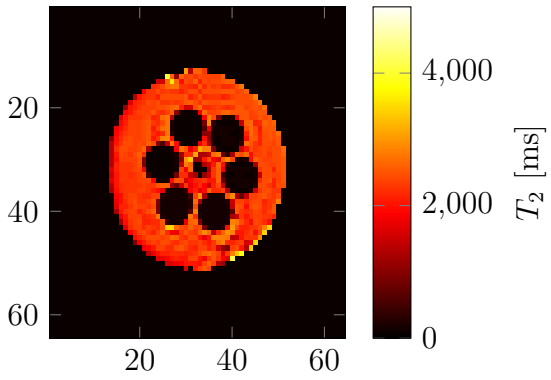

(b) T_2 map: colour scale to display long T_2.

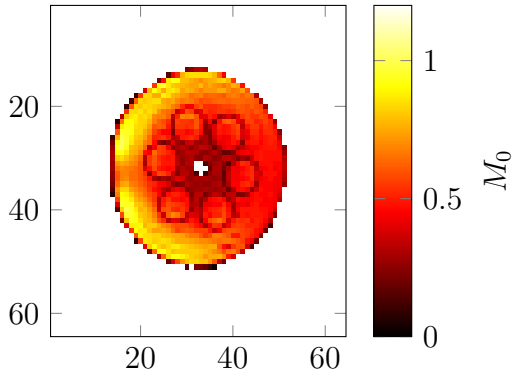

(c) Fitted proton density factor M_0.

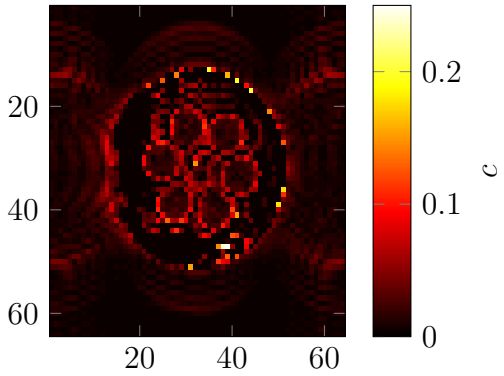

(d) Fitted noise factor c.

Figure 3.7: T_2 decay parameters for the custom phantom, fitted with the T_2 decay model.

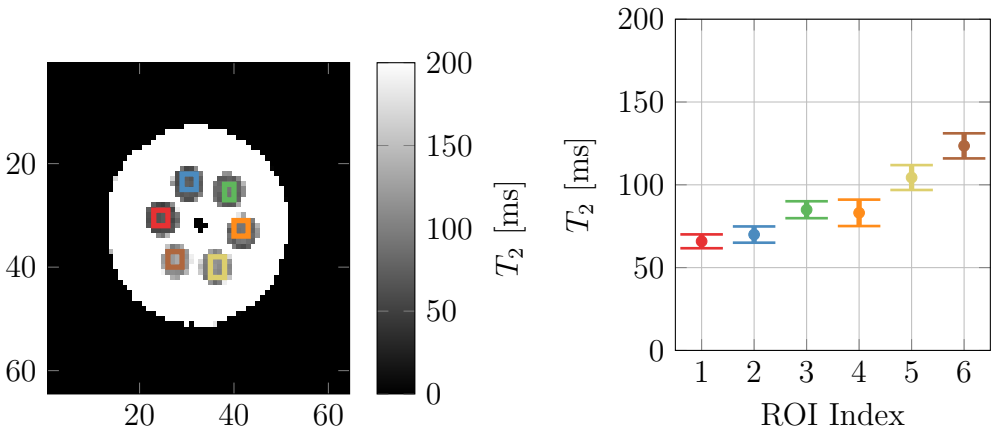

(a) ROI Positions

(b) ROI T_2 distributions

Figure 3.8: Distributions of T_2 within regions of the custom phantom, fitted to SE data acquired with various TEs.

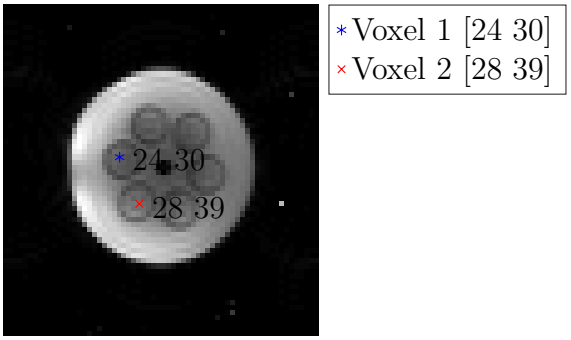

(a) Selected voxels.

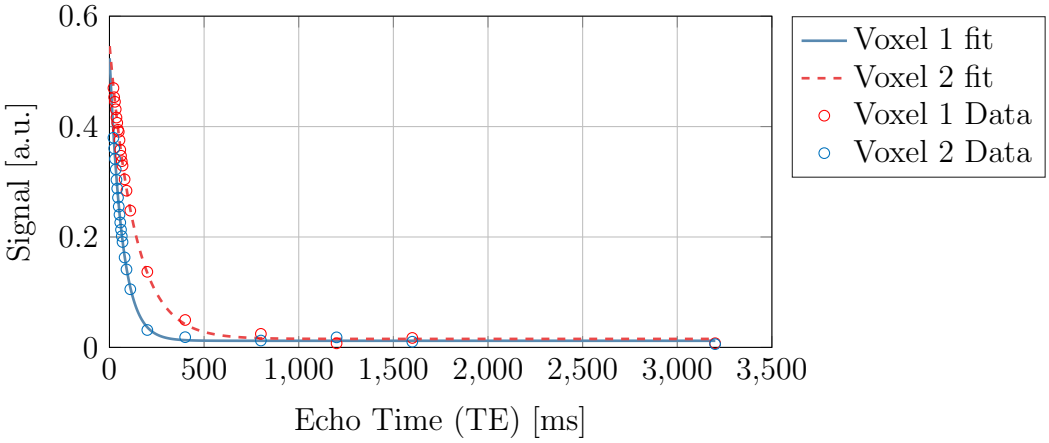

(b) T_2 fit for selected voxels (0 to 3500ms).

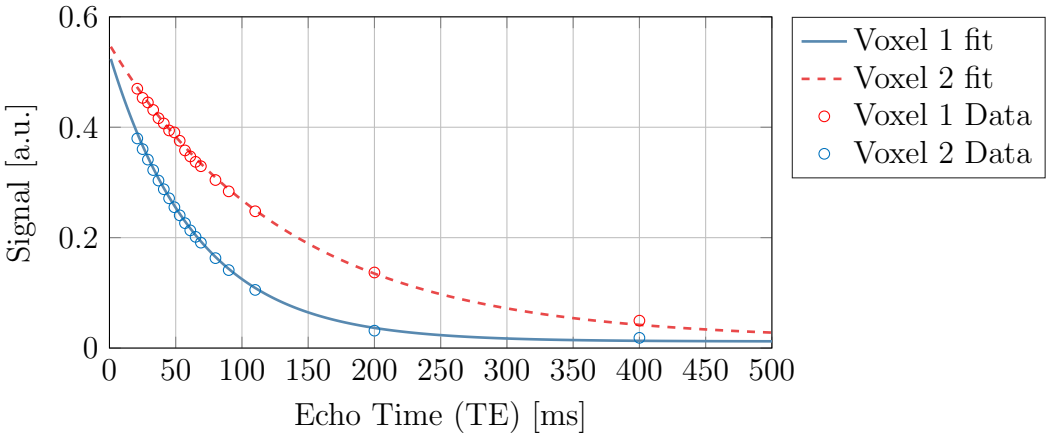

(c) T_2 fit for selected voxels (0 to 500ms).

Figure 3.9: Example T_2 fits (using 22 different TEs) for a voxel from the tubes with the longest and shortest T_2.

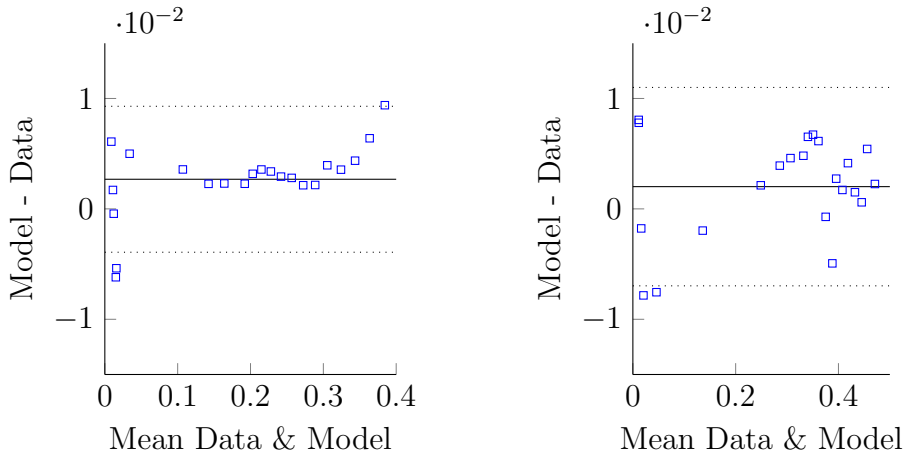

Figure 3.10: Bland-Altman plots for the T_2 fits for the voxels from the tubes with the shortest (a) and longest (b) T_2, showing the mean difference between the data and the fitted points (solid). The dashed lines mark the 95% confidence interval.

Tube Index	T_2 **from 22 measurements [ms]**	T_2 **from 9 measurements [ms]**
1	66±4	65±4
2	70±5	69±5
3	85±5	83±5
4	83±8	82±8
5	104±8	103±7
6	124±8	123±7

Table 3.3: Comparison of the phantom tube T_2 values.

3.4 B_1^+: DREAM

A 3D variant of the DREAM [26] sequence was chosen to map B_1^+. The original DREAM sequence was published by Philips Research Laboratories (Hamburg, Germany) but the 3D version used in this thesis (known as 3DREAM [27, 83]) was implemented by researchers at the German Center for Neurodegenerative Diseases (DZNE) in Bonn, Germany. It may be necessary to acquire MRF data and a B_1^+

map in series, to correct the flip angles used in the dictionary generation. The 3D DREAM sequence enables accurate mapping of the B_1^+ field from a few seconds of acquisition time, making it rapid enough to be potentially run immediately after one of the MRF acquisitions without significantly increasing the overall scan duration. The DREAM pulse sequence (Fig. 3.11) comprises a preparation pulse sequence for B_1^+ encoding, followed by a fast imaging segment to provide images. The preparation module uses two pulses with flip angle α_p to prepare a STimulated Echo (STE). There is a prephasing gradient between these two pulses and a spoiling gradient after the second of the two. The imaging module is a series of gradient-echo TRs with a low flip angle β. The gradient-echo refocuses the STE and a Free Induction Decay (FID) echo, the timings of which are determined by the gradients G_1 and G_2, respectively. The 3D DREAM implementation modified the readout to produce two 3D volumes, one from the STE and one from the FID.

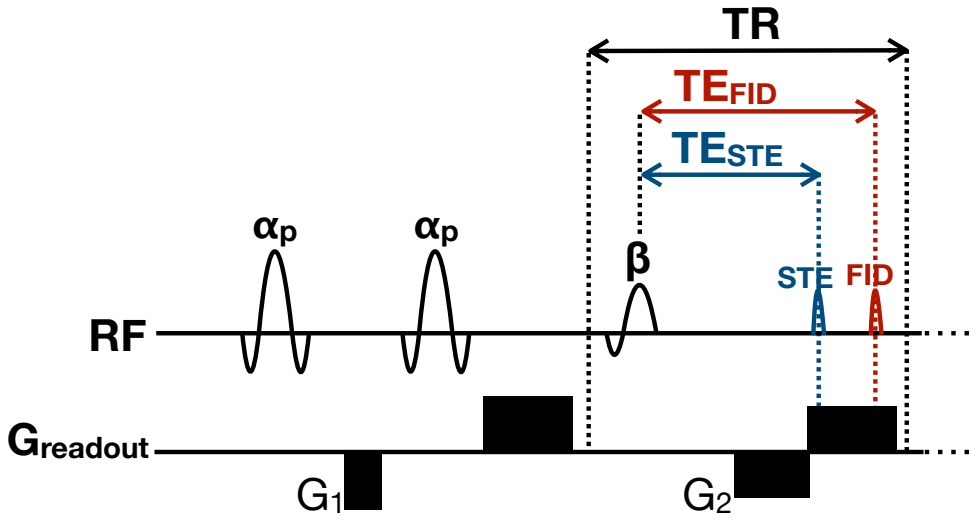

Figure 3.11: DREAM pulse sequence diagram, adapted from Nehrke *et al.* [26]. A B_1^+-encoding preparation sequence with flip angle α_p, followed by imaging TRs with flip angle β. STimulated Echo (STE) and Free Induction Decay (FID) echoes are formed during the gradient-echo readout.

Data Acquisition The STE and FID 3D slabs effectively contained 32 slices, each with a thickness of 7.5mm. The STE and FID echo times were 1.29ms and 2.69ms, respectively. The field of view was 256mm × 256mm and the in-plane pixel size was 3.9mm × 3.9mm. Other acquisition parameters were: TR = 5000ms, readout bandwidth = 750 Hz/pixel. The online reconstruction used a parallel imaging (GRAPPA) acceleration factor of 2 in the phase-encoding direction, acquiring 12 reference lines for this purpose. The sequence took approximately 5 seconds to run. The nominal preparation flip angle α_p was 60 degrees and the imaging flip angle β was 3 degrees.

Analysis To produce a map of the measured flip angle α_m for each slice of the volume the ratio calculation in [26] was applied to the slice magnitude images. This calculation is shown in Eq. 3.3, where I_{STE} and I_{FID} represent the STE and FID images, respectively. The derivation of Eq. 3.3 is provided in [26].

$$\alpha_m = \arctan\left(\sqrt{\frac{2I_{STE}}{I_{FID}}}\right) \tag{3.3}$$

Finally the measured α_m maps were converted to B_1^+ efficiency by dividing them by α_p. This meant voxels with $\alpha_m > \alpha_p$ had an efficiency greater than 1. Figure 3.12 shows a B_1^+ efficiency map for an example slice of the custom phantom, showing that the phantom structure causes the achieved flip angle to be higher than the prescribed flip angle (i.e. B_1^+ efficiency > 1). The B_1^+ efficiency appears to have a small variation across the different phantom tubes. Figure 3.13 shows an example *in vivo* B_1^+ efficiency map.

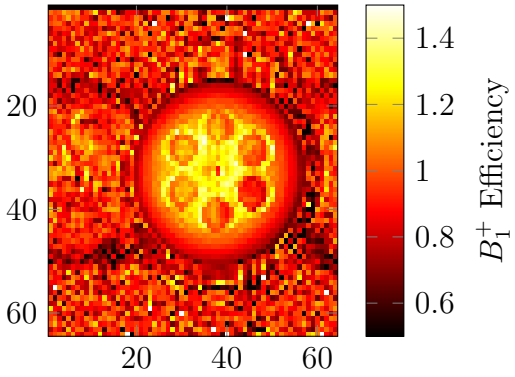

Figure 3.12: 3D DREAM B_1^+ efficiency map of the custom phantom, taken from a central slice of the 3D slab.

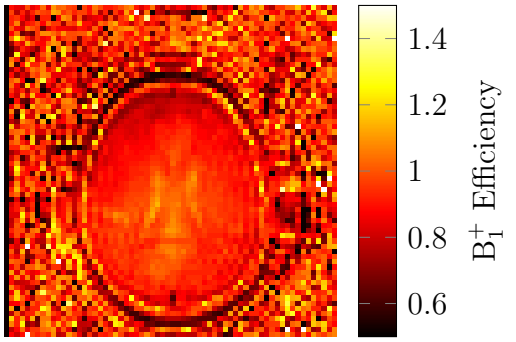

Figure 3.13: 3D DREAM B_1^+ efficiency map of an *in vivo* human brain, taken from a central slice of a 3D slab.

3.5 B_0: Double-Echo

To map the static magnetic field B_0 the double echo method [23] was used, with echo times TE_1 and TE_2. As shown in Eq. 2.4 the spin precession frequency $f_0 = \omega_0/2\pi$ depends on the local B_0 strength. Over time $\Delta TE = TE_2 - TE_1$, spins experiencing an off-resonance Δf accumulate phase $\Delta\phi$. This relationship can be rearranged to obtain Δf (Eq. 3.4) if the time difference and phase accumulation are known.

$$\Delta f = \frac{\Delta\phi}{2\pi\Delta TE} \tag{3.4}$$

Data Acquisition Two-dimensional slices were acquired using a Siemens field-mapping sequence with $TE_1 = 5.19$ms and $TE_2 = 7.65$ms. The voxel size was 2.3mm × 2.3mm in-plane, with 5mm slice thickness and 50% slice gap. The field of view was 300mm × 300mm, with a grid size of 128 × 128. Other parameters were: bandwidth = 260Hz per pixel, TR = 488ms and flip angle = 60 degrees. The total acquisition time was approximately 2 minutes 7 seconds for 17 slices (≈ 7.5 seconds per slice).

Analysis The Siemens online reconstruction associated with the sequence provided maps of $\Delta\phi$, which were then scaled between $-\pi$ and π, with phase wraps removed if needed, and used to calculate the off-resonance map (Eq. 3.4) in Hertz. Figure 3.14 shows an example off-resonance map of the phantom. Lobes of high off-resonance can be seen in the anterior portion of the slice in Fig. 3.14.

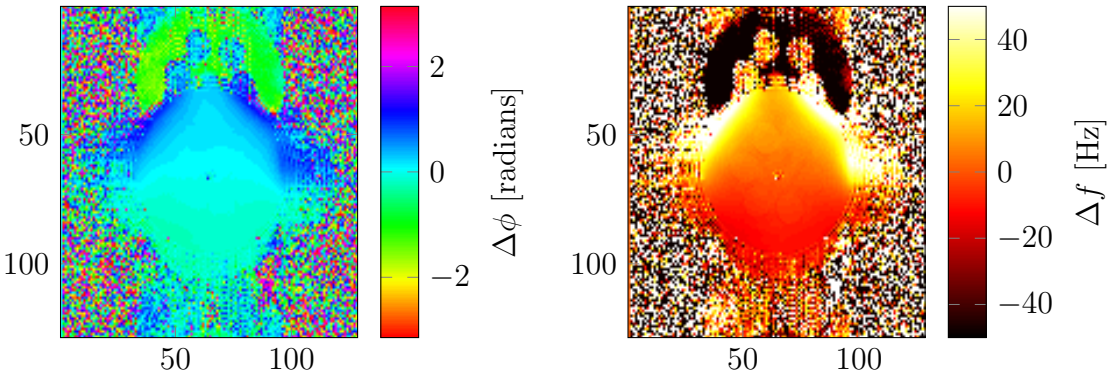

(a) Phase difference map.

(b) B_0 map of a central slice

Figure 3.14: Example B_0 off-resonance mapping of the custom phantom.

An example *in vivo* phase difference map is shown in Fig. 3.15a and the corresponding off-resonance map is shown in Fig. 3.15b. A region of relatively high B_0 can be seen in the anterior portion of the slice. Previous authors [1, 84, 85] also observed this spatial distribution of off-resonance in the brain.

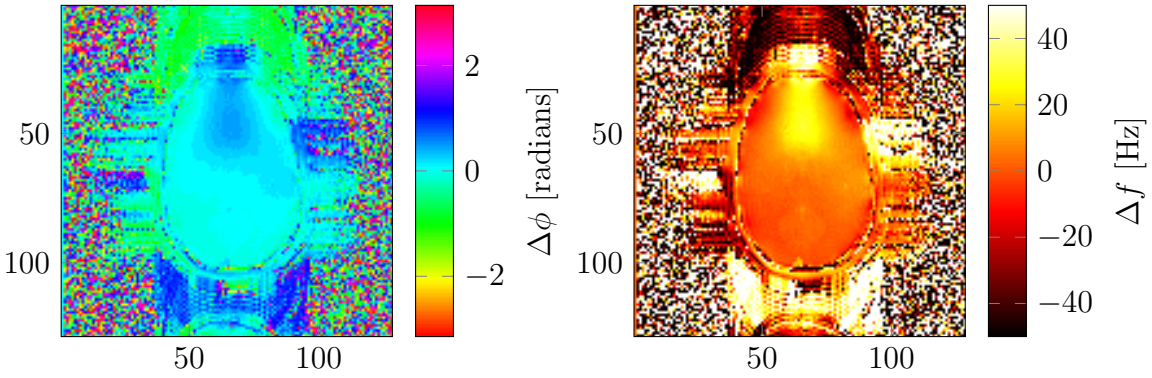

(a) *In vivo* phase difference map (b) *In vivo* double echo B_0 map

Figure 3.15: Example *in vivo* B_0 off-resonance mapping of the human brain.

3.6 MRF Sequences

This section includes details of the spiral MRF sequences created to replicate the original bSSFP [1] and FISP [37] MRF methods. Much of the MRF literature is based on these two sequence designs, making it appropriate to use these designs as validation tools. These sequences were made by editing a Fast Low Angle SHot (FLASH) sequence with a gradient-echo readout, within the Siemens IDEA pulse programming environment. This involved adding an inversion pulse at the beginning of the sequence and modifying the readout gradients to achieve a spiral sampling trajectory in k-space. The repetition times and excitation flip angles were designed to vary through the sequence according to a pre-defined pattern.

3.6.1 Inversion Pulse

To achieve the initial magnetisation inversion seen in [1] and [37] an adiabatic hyperbolic secant inversion pulse was positioned at the beginning of the sequence. In accordance with [1] and [37] the inversion was designed to perform a global inversion, rather than a slice-selective inversion. For simplicity a single dictionary was used for each slice acquired with the spiral MRF sequences, assuming fully relaxed magnetisation before the inversion. To satisfy this condition any multi-slice experiments with these sequences would need to allow sufficient time for full magnetisation recovery between slice acquisitions.

3.6.2 k-Space Trajectory

As in the original work [1] an undersampled spiral trajectory was used to sample k-space. The spiral was rotated by 7.5 degrees after each timepoint to cause the spatial distribution of the undersampled image artefacts to vary through time. The highest resolution possible with the scanner hardware and the trajectory was an in-plane pixel size of approximately 2.2mm $\times$ 2.2mm. This was limited by the maximum amplitude (24mT/m) and slew rate (180.18mT/m/ms) of the gradient system. To achieve this resolution a maximum gradient amplitude of 21mT/m was used, causing a maximum slew rate of 173mT/m/ms with the given gradient waveforms. This means that to achieve increased image resolution with this system the readout times would need to be increased to allow the imaging gradients time to move the system to higher frequencies in k-space. It is likely that this would necessitate longer TRs, which would alter the T_2 sensitivity. To account for possible discrepancies between the nonimal trajectory and that actually

achieved by the gradient coils the actual trajectory was measured (Fig. 3.16.) using the Duyn method [86]. The measured spirals were used in subsequent image reconstructions.

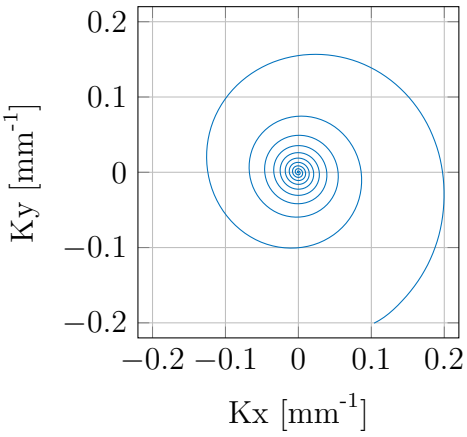

Figure 3.16: k-space sampling trajectory used to acquire images in the bSSFP and FISP MRF sequences. The trajectory was rotated by 7.5 degrees between timepoints. This trajectory produces an undersampled image but when it is used in conjunction with the other rotations k-space is effectively fully sampled. The example shown here would achieve 2.2mm × 2.2mm in-plane spatial resolution.

3.6.3 Image Reconstruction

To construct an undersampled image for each timepoint the Non-Uniform Fast Fourier Transform (NUFFT) MATLAB library in the Michigan Image Reconstruction Toolbox (MIRT) [87] was used. To do this the Density Compensation Function (DCF) [88] of the sampling trajectory was first calculated, a necessary step in order to appropriately weight the k-space samples along the trajectory. The measured trajectory and its DCF was used to perform an inverse NUFFT on each image from each channel, thus converting the channel data from spatial frequency-space to image-space.

A single complex image was reconstructed for each timepoint by combining

the images from each of the 32 channels via the adaptive combine method [89]. As part of this the mean images for each channel across the timepoints were used as inputs to generate the sensitivity maps.

Finally the parameter maps were produced by generating a dictionary of complex signal fingerprints and matching it to the data to find the best fits.

3.6.4 bSSFP MRF

Following the method used in the original MRF demonstration [1] the timing of each readout module was set such that it was centred within its TR. As in [1] the TRs and flip angles were varied through the sequence, meaning the readout timing had the same pattern of variation as the TRs, albeit downscaled. The pulse sequence diagram is shown in Fig. 3.17.

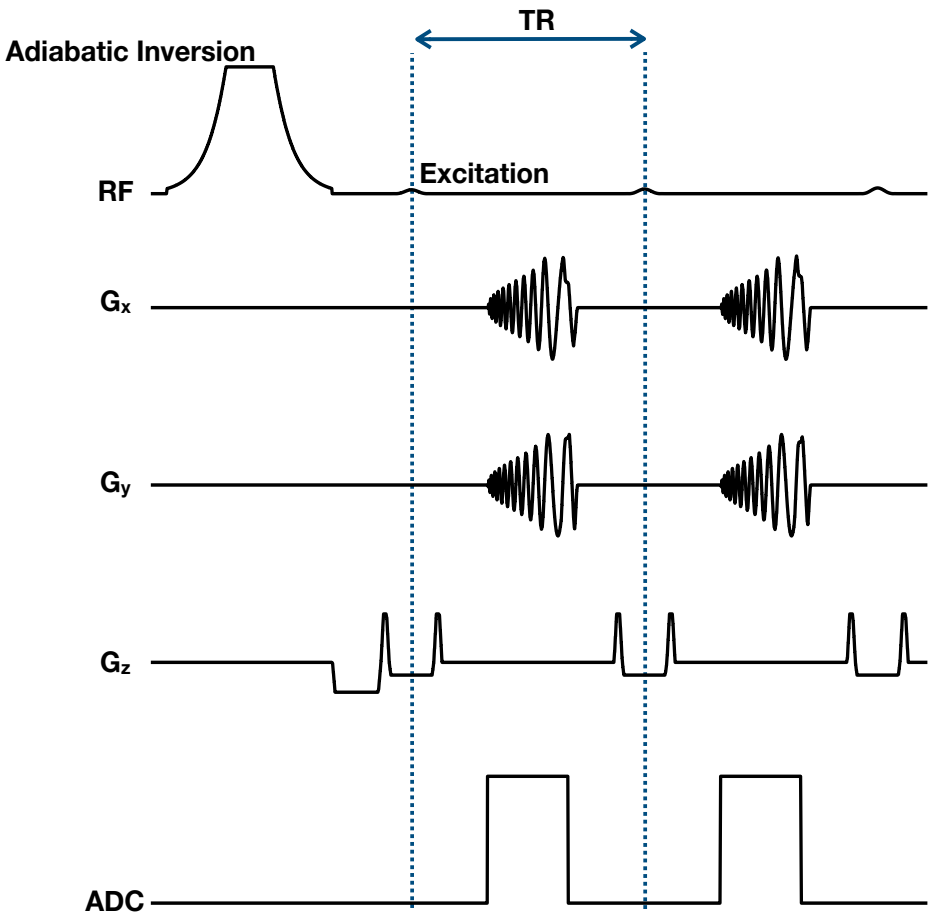

Figure 3.17: Pulse sequence diagram for the bSSFP MRF implementation. Vertical dotted lines mark the timings points for the start and end of the first TR and are positioned in the centre of the excitation pulses. Note the excitation pulses have considerably smaller amplitudes than the inversion pulse.

3.6.4.1 Acquisition Schedule

The bSSFP MRF acquisition schedule used in these studies (Fig. 3.18) was the same as that used in the original work [1], with varying TR and flip angle throughout the scan.

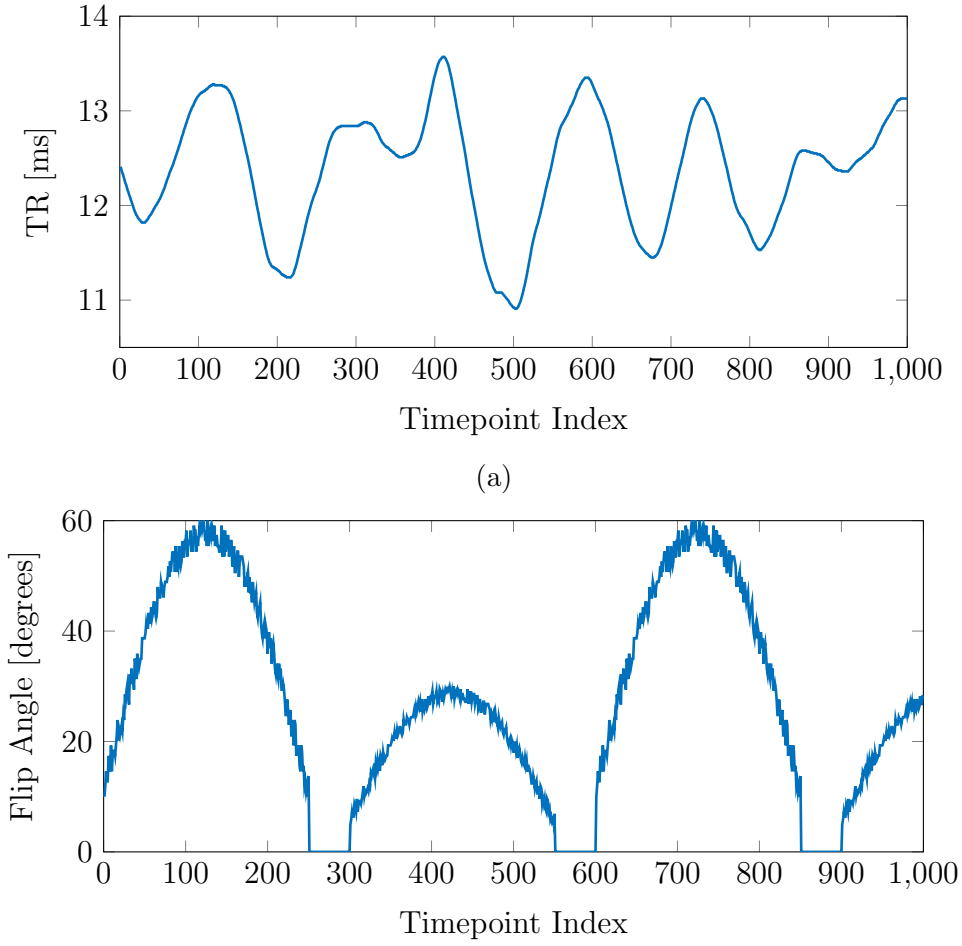

Figure 3.18: The repetition times (a) and flip angles (b) used for the bSSFP MRF sequence.

3.6.4.2 Slice Profile Simulation

Previous publications included the excitation slice profile in the dictionary simulation to improve the parameter map accuracy [39]. Here the same was done for the bSSFP MRF implementation, using a similar approach to [39]: the signal fingerprints were simulated for discrete points along the slice direction, using the sinc excitation pulse and slice selection gradient profiles from the scanner. Next the

signal was summed across the slice direction at each timepoint. This simulation was not performed for the adiabatic hyperbolic inversion pulse as it was assumed to supply a complete inversion across the slice. Instead, the initial magnetisation was set to -M_0, where $M_0 = 1$. Figure 3.19 demonstrates an example of how including the slice profile in the dictionary simulations affects the signal fingerprint shape. Since the excitation slice thickness was 5mm the simulation was run from -5mm to 5mm along the slice direction, with 0mm representing the middle of the slice.

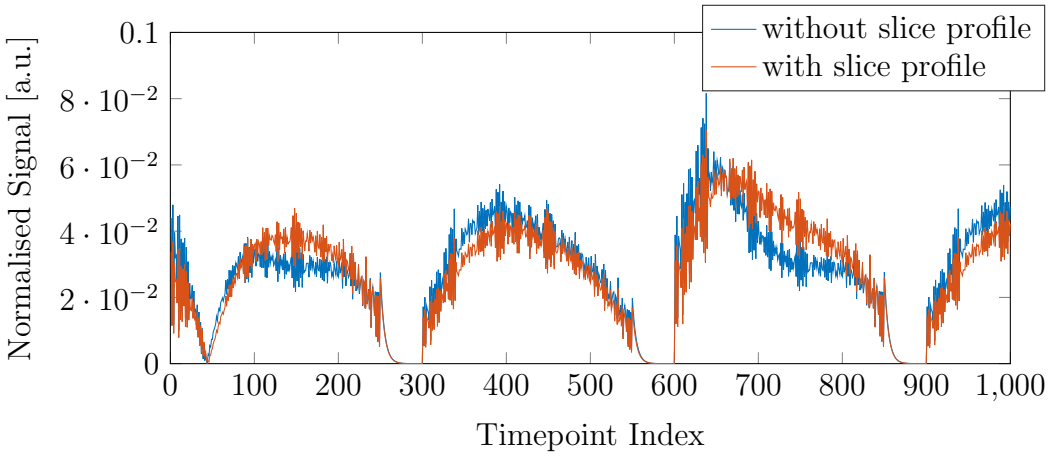

Figure 3.19: Example of the effect of the excitation slice profile on a simulated signal fingerprint ($T_1/T_2/\Delta f$ = 820ms/60ms/0Hz). Fingerprints are shown from simulations with and without the slice profile effect.

3.6.4.3 Fitted Parameter Maps

Figure 3.20 compares an acquired signal fingerprint from a voxel in one of the tubes with the best matched entry from a dictionary generated using slice profile simulations. Figure 3.21 displays parameter maps produced from matching the same dictionary to the entire image. Note that for these example parameter maps the k-space sampling trajectory was scaled such that the in-plane pixel size was

4mm × 4mm. Since bSSFP sequences are sensitive to off-resonance effects an off-resonance (Δf) dimension was including in the dictionary [1].

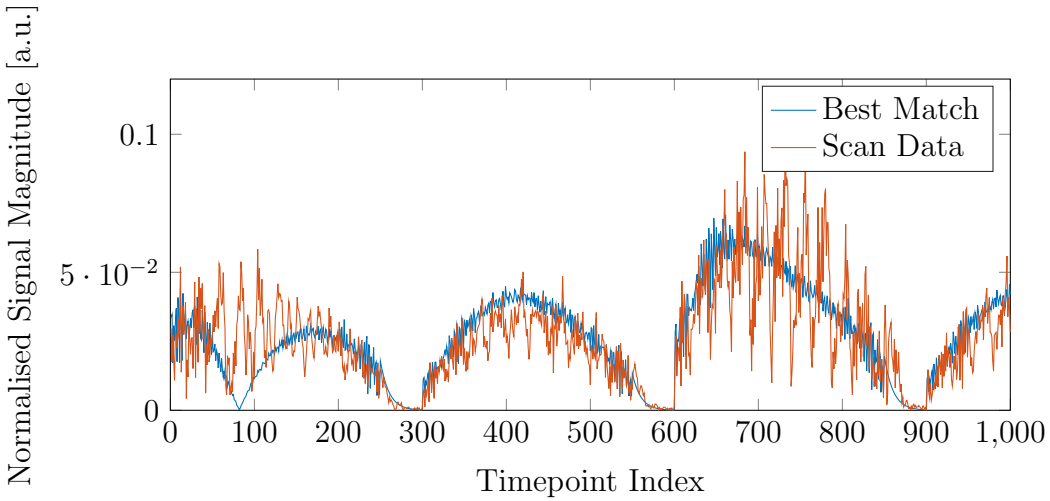

Figure 3.20: Dictionary fit with the bSSFP MRF fingerprint from an example voxel. Matched $T_1/T_2/\Delta f$ = 1600ms/128ms/13Hz.

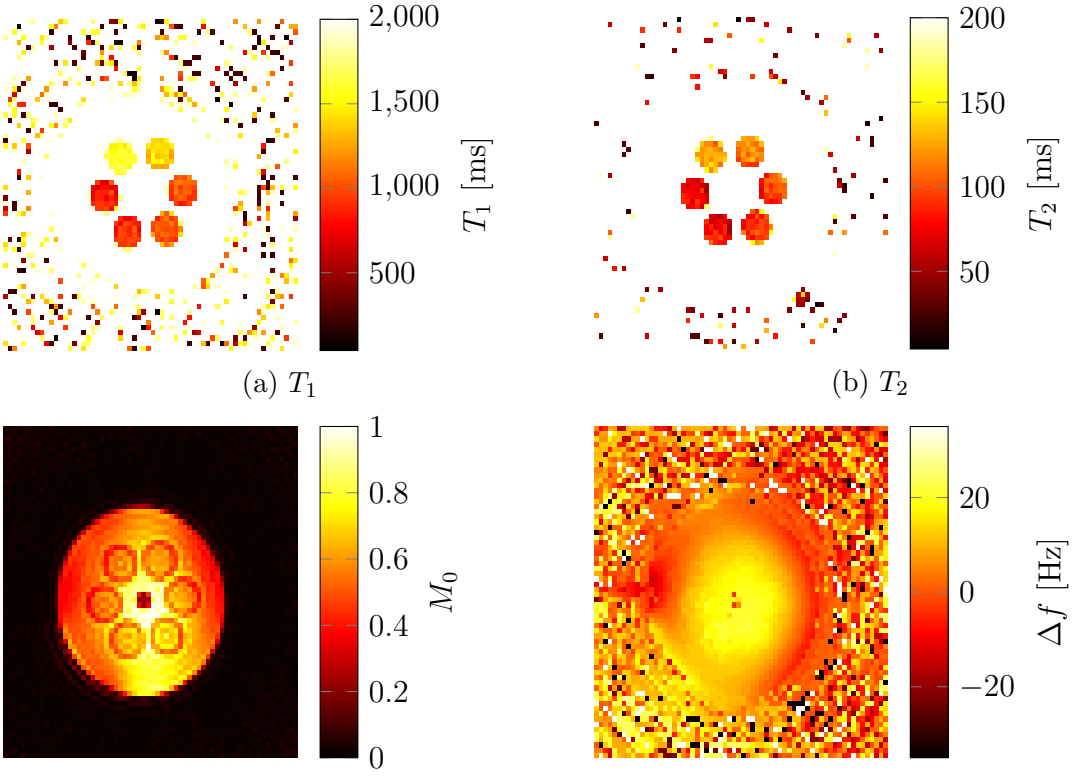

(a) T_1

(b) T_2

(c) Scaled proton density

(d) Off-resonance

Figure 3.21: Example bSSFP MRF parameter maps. The dictionary for this case was generated using slice profile simulations for each excitation pulse.

3.6.5 FISP MRF

By unbalancing the gradients in the slice direction the bSSFP sequence was converted to a SSFP sequence, otherwise known as FISP [37]. Another difference between bSSFP and FISP sequences is the position of the readout module within each TR. For FISP the TE is constant and the readout begins as soon as the rephasing gradient in the slice direction is turned off. The pulse sequence diagram for the FISP MRF implementation is shown in Fig. 3.22.

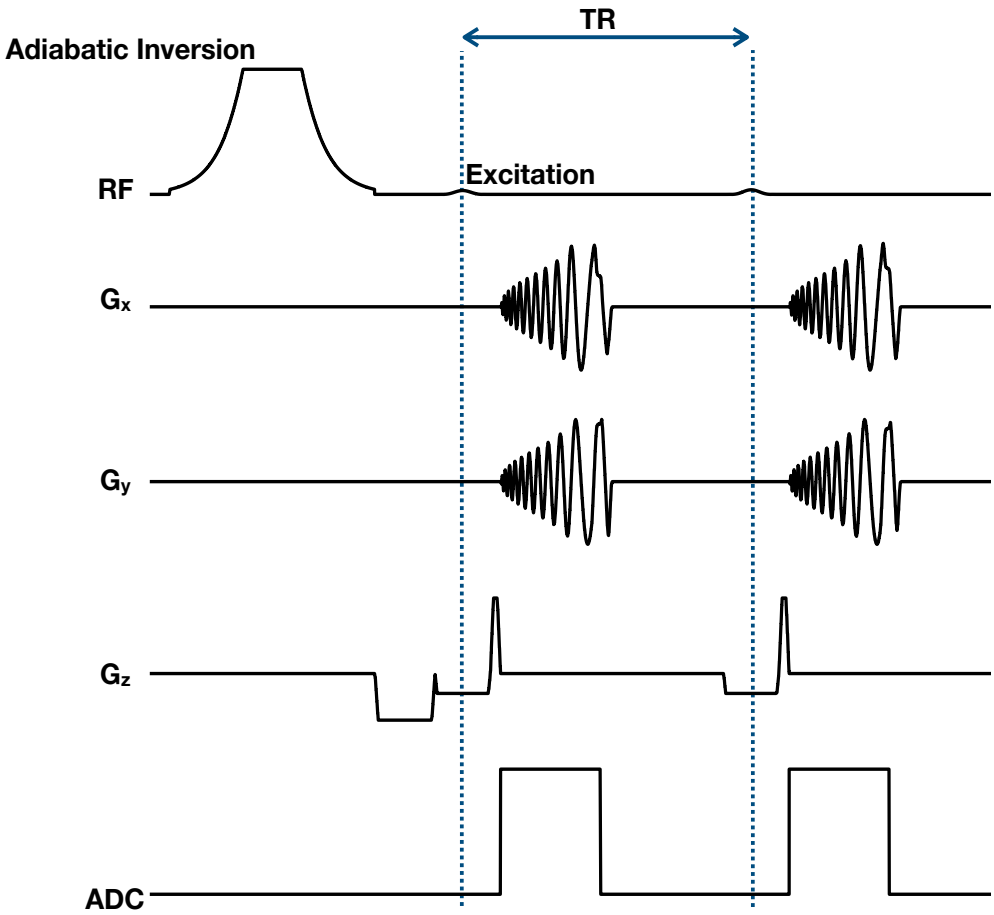

Figure 3.22: Pulse sequence diagram for the FISP MRF implementation. Vertical dotted lines mark the timings points for the start and end of the first TR and are positioned in the centre of the excitation pulses.

3.6.5.1 Flip Angle Schedule

The authors of the original FISP MRF demonstration [37] stated that to achieve smooth fingerprints with FISP sequences (and therefore improve the distinction between image artefacts and signal) it is particularly important that the acquisition parameters are smoothly varying. Hence in [37] the flip angle pattern was

designed to be smoother than that used in the bSSFP work [1]. The FISP sequence designed for this thesis utilised a flip angle pattern (Fig. 3.23) similar to that used in [37].

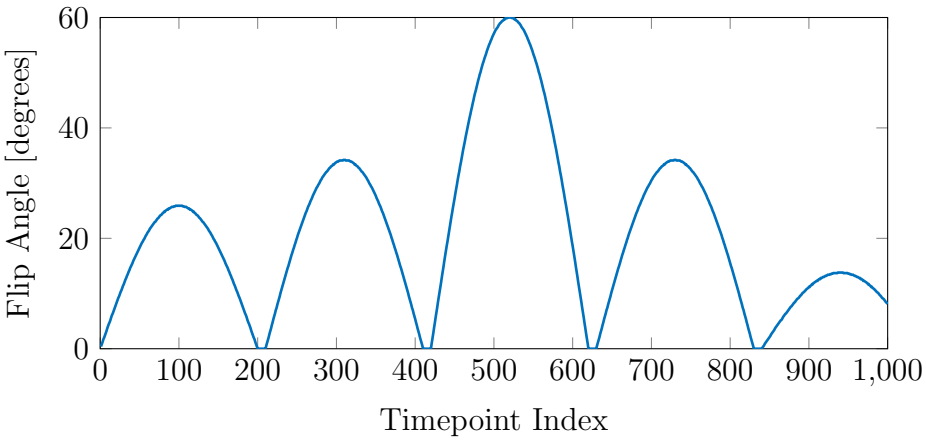

Figure 3.23: The flip angle pattern used for the FISP MRF implementation.

3.6.5.2 Fitted Parameter Maps

Here an example is given of an acquired FISP fingerprint from a voxel and its best matched dictionary entry (Fig. 3.24). In accordance with previous FISP MRF work an EPG implementation was used to simulate the FISP signal and make the dictionary. Figure 3.25 shows parameter maps produced from matching the dictionary to the image data, using in-plane pixel sizes of 4mm $\times$ 4mm. Since the FISP sequence is inherently insensitive to off-resonance [37] Δf was not included as a dictionary dimension.

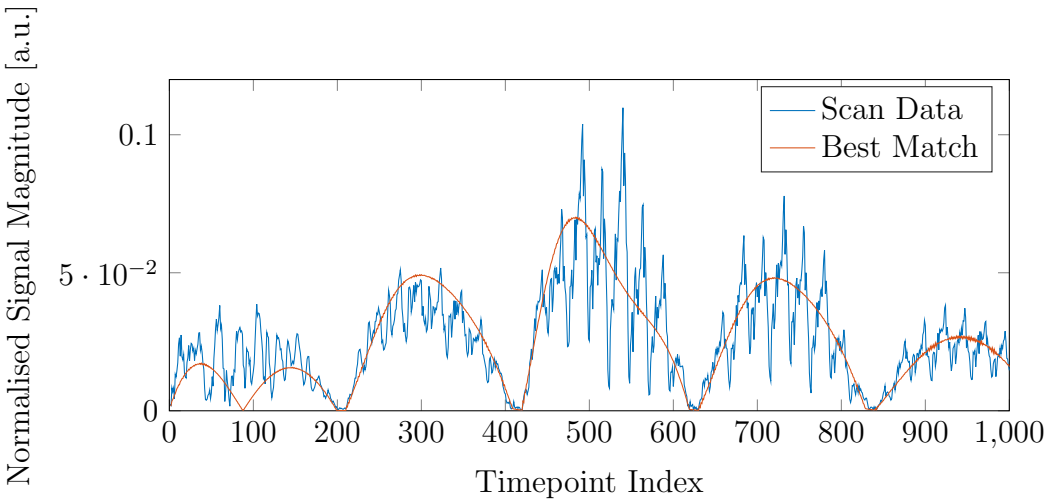

Figure 3.24: An example FISP MRF signal fingerprint and the best matched dictionary entry. The dictionary fingerprint was generated using T_1/T_2 = 1580ms/144ms.

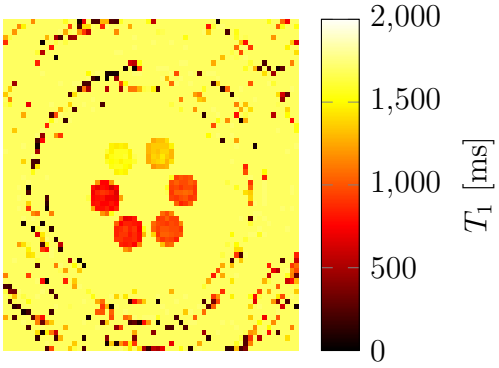

(a) T_1

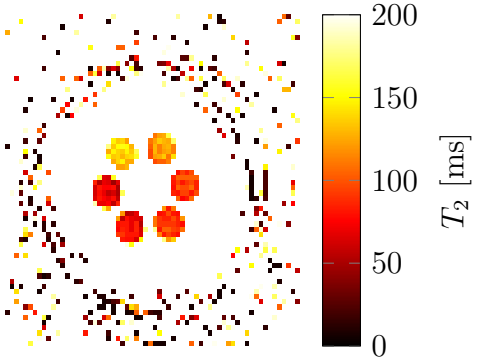

(b) T_2

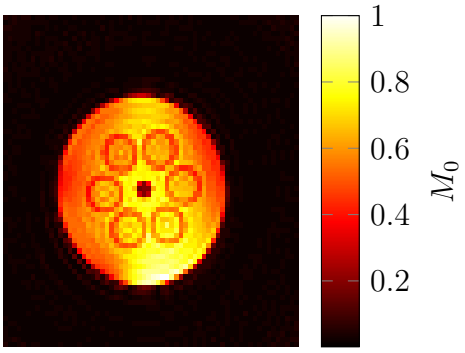

(c) Scaled proton density

Figure 3.25: Example FISP MRF parameter maps.

Chapter 4

Acquisition Schedule Optimisation for Snapshot MRF

4.1 Introduction

The choice of MRF acquisition parameters is an important problem to consider as it determines the shape of the fingerprint for a given set of tissue and system properties, which in turn affects the parameter estimation precision and accuracy. This chapter presents an approach for optimising the acquisition parameters for a multislice spin-echo EPI MRF acquisition to increase parameter-mapping accuracy and precision. The optimisation framework is able to target any chosen pairs of T_1 and T_2, such as those of typical *in vivo* grey and white matter in the human brain or the variation caused by injury or pathology. Section 4.2 describes the sequence components and details of the optimisation framework are provided in Section 4.3. Section 4.4 contains simulation results, which explore the theoretical performance of the optimised acquisition schedules for a wider range of T_1 and T_2.

The chapter concludes with a summary of the preceding sections and remarks on the theoretical performance of the framework (Section 4.5).

4.2 Sequence Design

Before choosing the acquisition parameters for a MRF experiment it is important to consider which fundamental variety of sequence to use as a base. In principle any variety of sequence can be used for MRF but each may have different practical advantages and limitations, such as sensitivity to different tissue or system properties. The framework presented here used a spin-echo EPI sequence as the foundation for a multislice MRF sequence, with fully sampled images at each timepoint. Figure 4.1 shows a diagram of the underlying pulse sequence and details about the pulse sequence are included in this section.

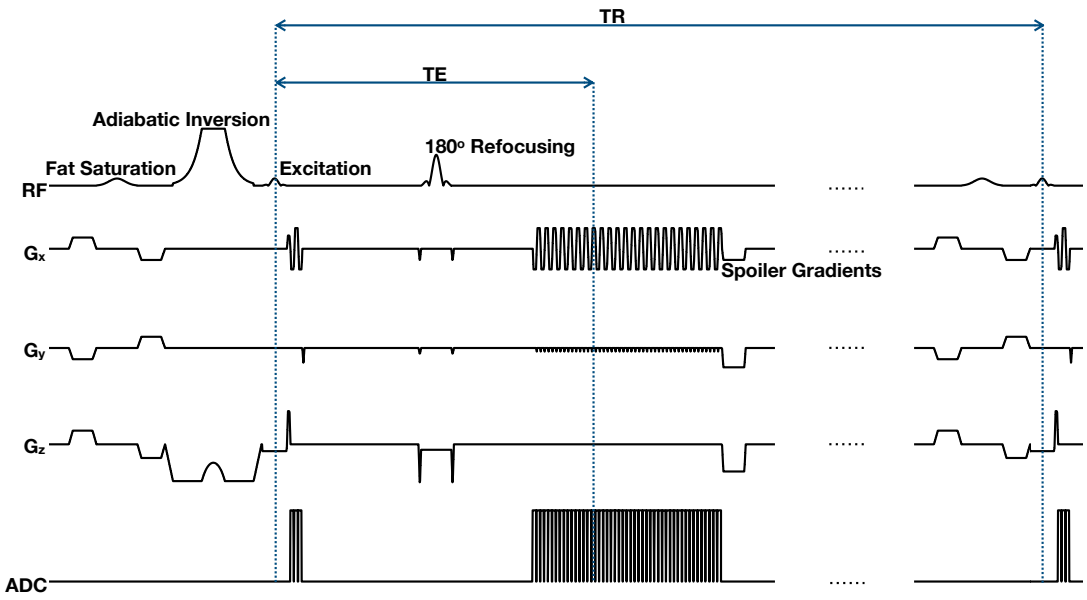

Figure 4.1: Diagram for the spin-echo EPI MRF pulse sequence. An optional adiabatic inversion pulse is positioned between the fat saturation pulse module and the excitation pulse module, shown here immediately before the labelled TR begins. The inversion pulse is a Frequency Offset Corrected Inversion (FOCI) pulse, with a FOCI factor of 2. The three echoes acquired immediately after the excitation module are for EPI phase correction of the snapshot image acquired later in the TR. At the end of the TR the next fat saturation pulse module is played, followed by the next optional inversion pulse. The optimisation framework decides which TRs will include an inversion pulse. In the example in this diagram the inversion before the second TR has been omitted to show how the fat saturation timing shifts in this case.

4.2.1 Fully Sampled EPI

At each timepoint k-space was fully sampled in order to avoid the artefacts associated with undersampled trajectories, providing the opportunity of increased SNR in the fingerprint measurements. This also allowed a more straightforward image reconstruction pipeline than would be necessary for a non-Cartesian sampling strategy such as undersampled interleaved spirals.

4.2.2 T_2 Sensitivity: Spin-Echo Refocusing

Previous EPI MRF work mainly used gradient-echo acquisitions [59], meaning the acquisitions were sensitive to B_0 inhomogeneities and T_2^*. Spin refocusing ensures sensitivity to T_2, rather than T_2^*, and is particularly important for cases of long TEs, where the signal would otherwise be more heavily T_2^*-weighted. Pairs of 180 degree refocusing pulses were used in [41], simultaneously mapping T_2 and diffusion by including diffusion-encoding gradients before and after each refocusing pulse. In contrast, a single refocusing pulse was deployed in this thesis, enabling T_2 sensitivity and access to shorter TE than was possible with the double refocusing scheme of [41] which in turn allowed the possibility of increased signal. This was achievable because diffusion effects were not modelled in the dictionary, with an emphasis placed on T_1 and T_2 measurement, rather than diffusion. TE was varied through the acquisition to generate sensitivity to T_2.

4.2.3 T_1 Sensitivity

To create T_1 sensitivity the sequence manipulated M_z by varying the excitation Flip Angle (FA) and using strategically positioned inversion pulses.

4.2.3.1 Flip Angle and Repetition Time

As discussed in Chapter 3, T_1 contrast is affected by the TR length, since unless TR $\gg T_1$ the longitudinal magnetisation component at the end of a given TR period contributes to the signal in the next TR. Since FA also affects M_z these two acquisition parameters can be used together to produce the required T_1 contrast. In the context of this thesis it is desirable to achieve a combination of FAs and TRs

that generate high T_1-sensitivity in a short scan time. In this work FA was varied through the sequence whereas TR was kept constant for the sake of simplicity and to allow the total scan duration to be fixed.

4.2.3.2 Inversion pulses

Use in SSFP MRF The early MRF implementations used a single global inversion pulse at the beginning of the sequence to generate T_1 sensitivity. After an initial inversion M_z passes through zero at a timepoint specific to the T_1 species. The same principle is used in Inversion Recovery (IR) T_1 mapping, the gold standard method introduced in Chapter 2. Figure 4.2 shows examples of this in the case of a bSSFP MRF simulation. The flip angles and repetition times used for the simulation are also provided in Fig. 4.2.

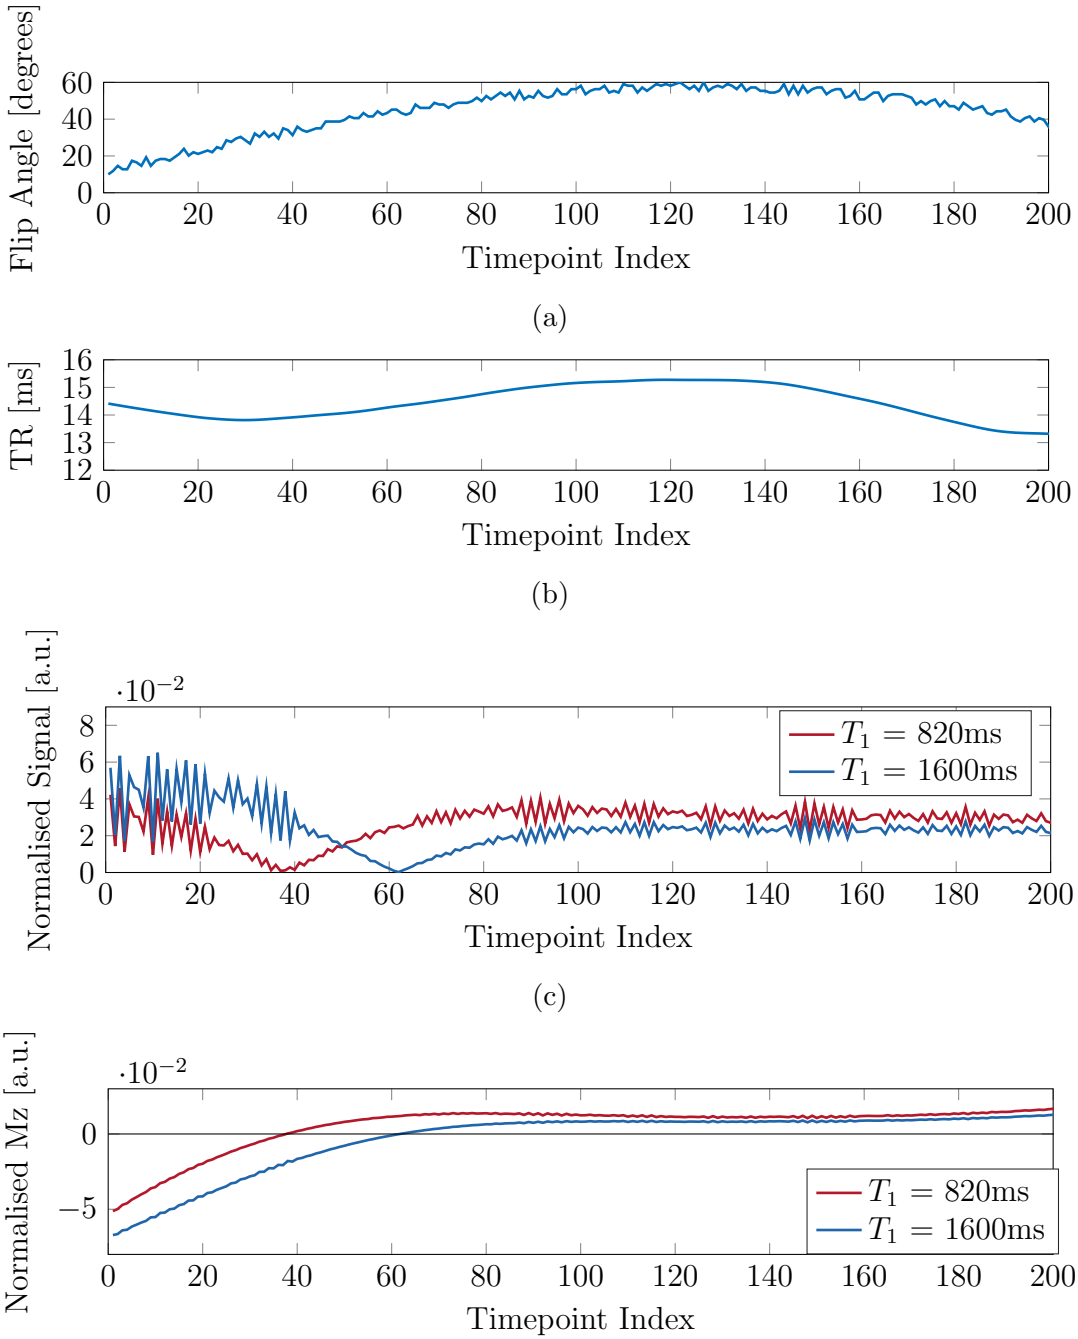

Figure 4.2: bSSFP MRF magnetisation simulations using similar flip angles (a) and repetition times (b) to that of [1], with the corresponding MRF signal (c) and longitudinal magnetisation M_z (d). M_z is shown to pass through zero part-way through the sequence, at a time dependent on T_1. Here $T_2 = 65$ms and $\Delta f = 0$Hz.

Use in Spin-Echo MRF The work of Rieger *et al.* [41] showed that inversion pulses can be used in conjunction with spin echoes to obtain MRF maps of T_1 and T_2, using global inversion pulses played at various timepoints during the sequence. In [41] the inversion pulses were explicitly credited with improving T_1 contrast but the use of two refocusing pulses per TR also altered M_z. In general, a single refocusing pulse inverts M_z and so can essentially negate a preceding inversion pulse (assuming $TE \ll T_1$). However, the presence of a second refocusing pulse in [41] effectively reversed this cancellation, meaning contrast from the inversion pulse was mostly restored. Note the inversions were not exactly negated or restored, due to the magnetisation recovery during the time between the pulses. Figure 4.3 demonstrates this phenomenon, plotting the simulated M_z evolution over a single TR with and without an initial inversion pulse.

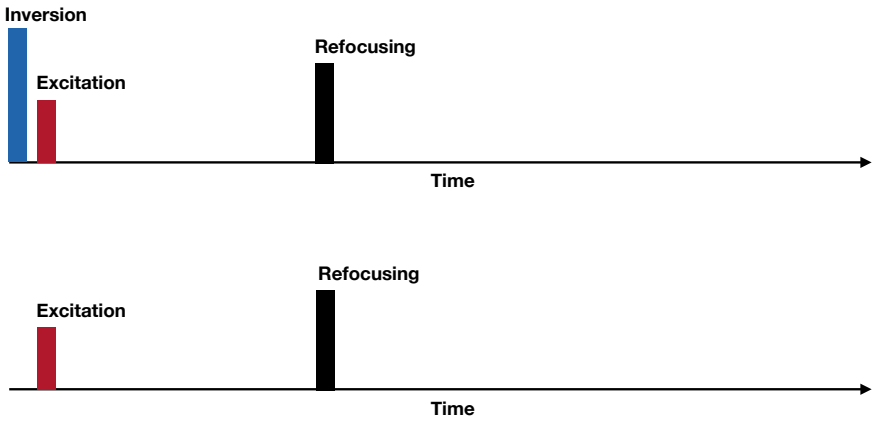

(a) Two TR cases: pre-inversion pulse flag switched on (top) or off (bottom)

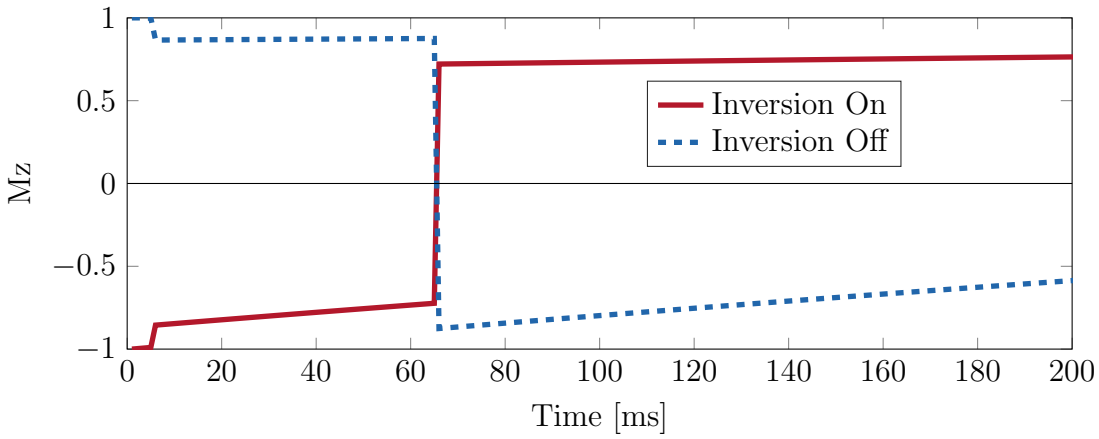

(b) Interplay between inversion and refocusing pulses: simulated M_z for a single TR

Figure 4.3: Demonstration of the interplay between inversion and refocusing pulses of an inversion recovery spin-echo sequence in terms of M_z. A symbolic pulse sequence diagram of a single TR with and without an inversion pulse (a), aligned with the simulated M_z evolution for the two cases (b). Here $T_1/T_2 = 800\text{ms}/60\text{ms}$.

Use in this Work Here slice-selective pre-inversion pulses were used to enhance T_1 sensitivity. The optimisation algorithm determined whether or not an inversion pulse should be played before the spin-echo excitation pulse, thus varying the temporal positioning of the inversion pulses throughout the sequence. Ideally the inversion pulses would be positioned such that their contributions to T_1 contrast

would not be immediately negated by the refocusing pulses. Whenever an inversion pulse was not played the preceding fat saturation module was shifted to be immediately before the excitation pulse module, as illustrated in Fig. 4.1 where an inversion pulse has been omitted from the end of the example TR. Each time an inversion pulse did not immediately precede a TR an effective inversion was achieved because of the refocusing pulse. Therefore if a TR was played out with no preceding inversion followed by a block of TRs with inversions the M_z evolved in an inversion recovery-style through that block of the sequence, as illustrated in Fig. 4.4.

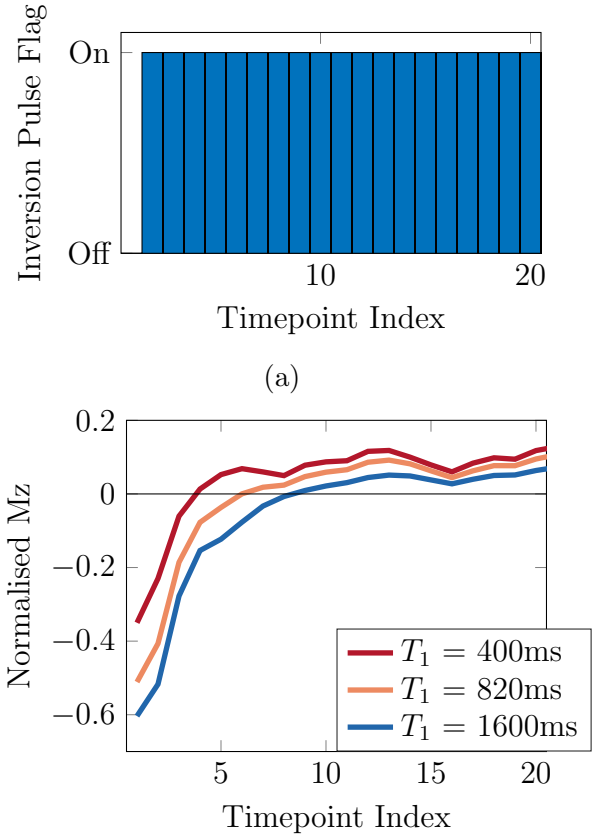

Figure 4.4: The inversion pulse flag pattern (a) used to invoke inversion recovery-style M_z behaviour and simulated examples of the resulting longitudinal magnetisation evolution (b). Here $T_2 = 65$ms.

Pulse Type Slice-selective inversion pulses were used in this work to ensure the magnetisation evolution of each slice was manipulated independently. Using slice-selection avoided conflict between the schedules of individual slices during the optimisation and allowed the same optimised schedule design to be used across multiple slices. Specifically, hyperbolic-secant adiabatic pulses with Frequency Offset Corrected Inversion (FOCI) factor 2 were used. This type of pulse is designed to be tolerant to B_1^+ field inhomogeneity and provide uniform inversion

across a given slice with well defined slice edges [90].

Safety Limits If a large number of inversion pulses are played too close together temporally the energy deposition safety limits imposed by the scanner may be exceeded, causing the scan to stop during the acquisition or fail to start. Specifically, the scanner would stop the scan to avoid exceeding a predefined maximum Specific Absorption Rate (SAR) for a given time window, where SAR is the radio frequency power absorbed per unit mass of tissue in units of W/Kg. Therefore, whether or not these safety limits are exceeded would have a dependence on the TR length, as well as whether an inversion pulse was played in the preceding TR. A long TR would allow a power-intensive series of TRs with inversion pulses to be played in succession but could reduce the T_1 sensitivity. Ideally the excitation flip angles, TR length and inversion pulse flags would interact to produce maximum T_1 sensitivity while staying below the allowed energy deposition limits. The (fixed) TR length in this work was chosen such that inversion pulses could be played before most of the TRs if the optimisation algorithm required it without the sequence exceeding these limits.

4.2.4 Multislice Acquisition

The sequence acquired blocks of several TRs for each slice before moving to the next slice, cycling through the slices in ascending order starting from Slice 1. An alternative slice ordering method is to acquire the slices in an interleaved manner, which can ensure minimal interaction between adjacent slices. Here, this was not found to be necessary and the slices were acquired sequentially. Once the same block of the acquisition schedule had been used to acquire data for all slices the

sequence returned to Slice 1 and acquired the next block of the schedule. The sequence continued until the entire schedule had been used for all slices. Figure 4.5 illustrates this acquisition strategy, using a 10-slice scheme as an example. The nth TR for a particular slice used the acquisition parameters from the nth timepoint of the schedule and each TR was the same length. This meant each slice experienced the same acquisition schedule, as illustrated by Fig. 4.5b, thus simplifying the dictionary generation and matching processes.

Acquiring blocks of TRs for each slice enabled extended effective TRs for the final timepoint of each block, which occurred at periodic positions during the sequence. The effective TR of the final timepoint of each block was longer than the other TRs in the block because other slices were being acquired, meaning it depended on the total number of slices and the block length. During these periods the magnetisation of each slice had time to at least partially recover while other slices were being acquired, creating opportunities for periods of increased signal and image SNR at the succeeding timepoints. Therefore the fingerprint of a 10-slice scheme was likely to be different to that of a single-slice acquisition, leading to a T_1 sensitivity change.

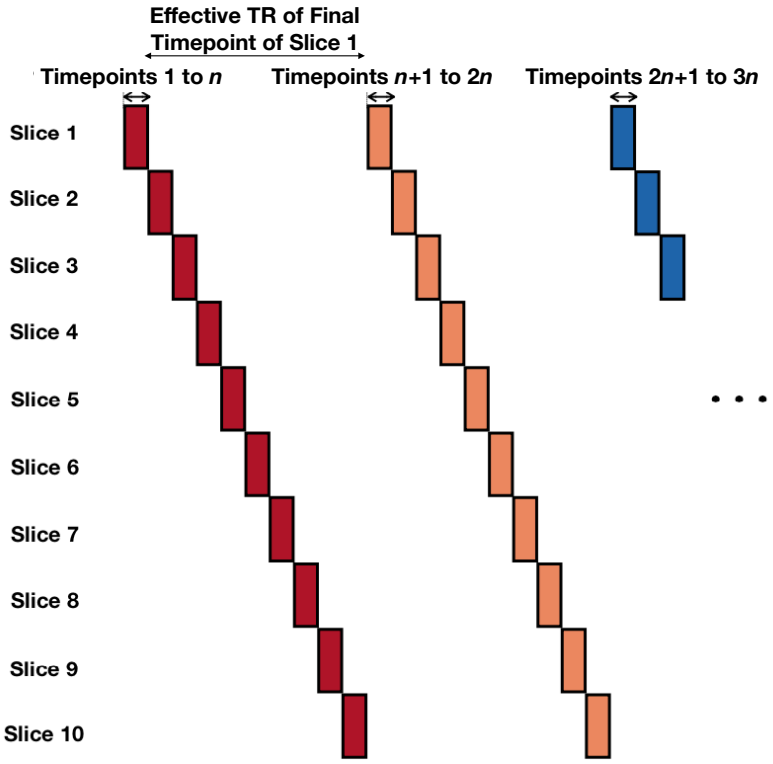

(a)

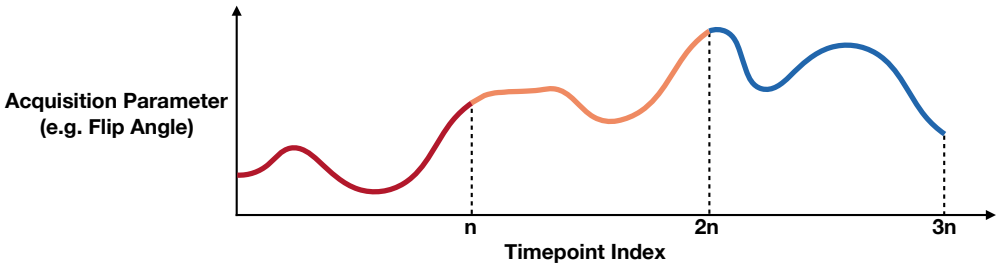

(b)

Figure 4.5: Schematic diagram showing the slice acquisition order through the sequence for a 10-slice scheme (a) and its relation to an example acquisition parameter schedule (b). A block of n timepoints was acquired for Slice 1, before acquiring the same for Slice 2 and the remaining slices. Each block comprised multiple TRs, all of the same length. Each timepoint block in (a) is colour-matched to the corresponding schedule section in (b). The effective TR of the final timepoint of the blocks is labelled to demonstrate the relationship between the effective TR and the number of slices.

Spoiling By using a fully sampled image acquisition for each timepoint it is possible to change the excitation slice between timepoints. However, effective spoiling must be performed to ensure the signal acquired from each slice is independent of the other slices. Previous multislice EPI MRF work used time-invariant gradient spoiling in conjunction with RF spoiling. The next chapter of this thesis contains the use of a hexagonal spoiling scheme, which was been shown to be effective in previous work [91]. RF spoiling was not used for the SE EPI MRF sequences in this thesis.

4.3 General Optimisation Framework

This section begins with an introduction to the necessary considerations for designing a sequence optimisation approach (Section 4.3.1). Next is a description of the method used to optimise the SE EPI MRF acquisition parameters, including how the parameters were varied (Section 4.3.2), the optimisation metrics (Section 4.3.3) and the chosen algorithms (Section 4.3.4).

4.3.1 Introduction

The choice of which acquisition parameters to optimise, and which to fix, should be guided by what kind of sensitivity is desired. For example, T_2 sensitivity can be achieved by varying the TE across measurements, as discussed earlier. One of the benefits of MRF is its potential sensitivity to multiple tissue or system properties in a single scan. In theory, flexibility in all the sequence acquisition parameters can be exploited to produce sensitivity to the targeted tissue or system properties. The goal of MRF schedule design is to produce contrast at each timepoint, between

the targeted tissue or system properties. For example, for T_1 and T_2 mapping it is important to achieve good contrast between the signal time courses of different T_1 and T_2 species.

Timepoint-Wise Optimisation Versus Analytical Functions In the initial demonstrations of MRF, abrupt changes in the schedule parameters were avoided because of the artefacts from the undersampled image acquisitions. The sequence was designed such that the fingerprints were spatially incoherent with respect to the image artefacts, therefore making them distinguishable from the artefacts during the dictionary matching. If, instead, fully sampled images are acquired for each image acquisition the image artefacts are less severe, meaning it is potentially less important for the fingerprints to vary smoothly. In this case the acquisition parameters for each timepoint could be treated independently and optimised on a timepoint-wise basis, altering the acquisition parameters of the image acquired at each timepoint to minimise the cost function for the schedule as a whole. However, timepoint-wise optimisation would require a large parameter space, potentially making it more challenging to find a global minimum. Additionally, multiple small and sharp signal changes between timepoints may resemble the fluctuations caused by thermal noise, which could reduce the measurement precision and make the acquisition less tolerant to increased levels of noise or periods of reduced signal.

An alternative to timepoint-wise optimisation is to set the parameters of the sequence using analytical functions [68]. By doing this it is possible to generate continuous FA and TE patterns determined by a small number of function parameters. Depending on the parameter limits and the analytical expression these patterns may also be smooth. Using analytical functions instead of timepoint-wise

optimisation would greatly reduce the dimensionality of the optimisation problem. Even with fully sampled images it may be beneficial for the fingerprints to vary according to smooth analytical functions, to provide incoherence between the signal fingerprints and thermal noise, as well as any subtle artefacts caused by the image reconstruction.

Slice Acquisition Order Another aspect of a multislice acquisition that could be optimised is the slice acquisition order. Altering the slice order in a pseudo-random manner could lead to increased signal and improved acquisition efficiency. Changes to the slice order would change the magnetisation evolution and would therefore alter the T_1 sensitivity. Figure 4.6a provides an example of a pseudo-random slice acquisition scheme. In the framework used in this study the slice acquisition order was fixed. This made it more straightforward to design a scheme where each slice experienced the same schedule and reduced the dimensionality of the optimisation. It also made the matching process less computationally demanding, as only a single dictionary was required. As introduced in Section 4.2.4, in this study the slices were acquired in ascending order (Fig. 4.5) rather than in the interleaved arrangement shown in Fig. 4.6b.

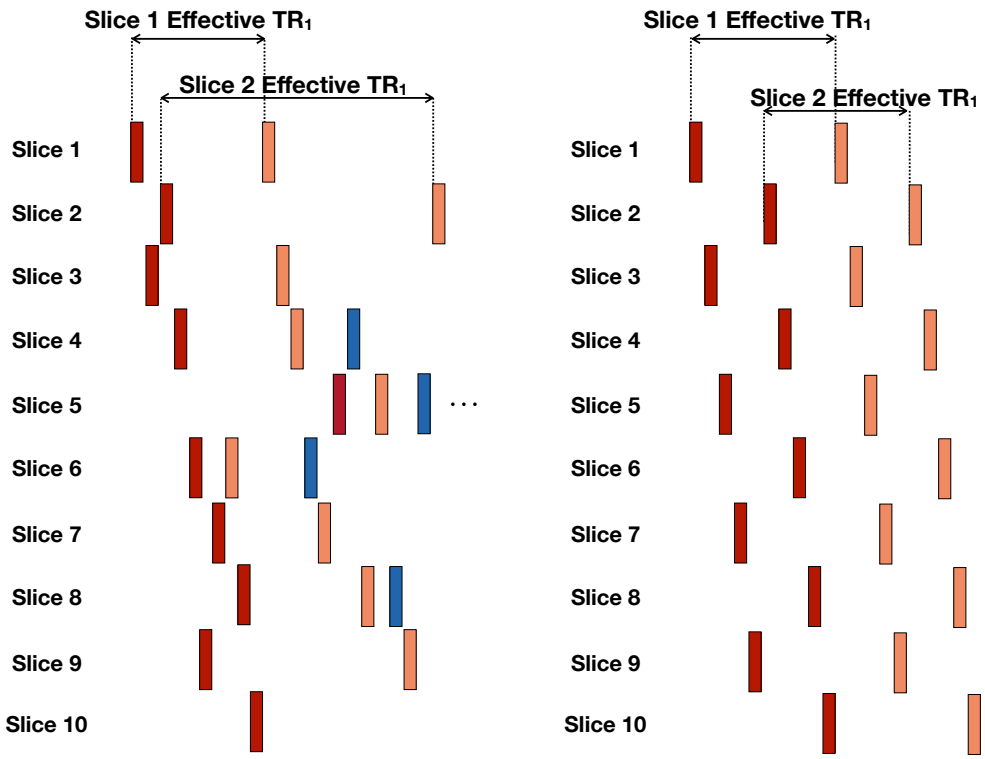

(a) (b)

Figure 4.6: Examples of pseudo-random (a) and interleaved (b) slice acquisition orders. In the pseudo-random scheme the effective TRs vary across the slices, producing differences in signal fingerprints for different slices. The interleaved scheme results in identical effective TRs for each slice. In this thesis, the slices were acquired sequentially, in ascending order. Blocks of TRs were acquired for each slice, as illustrated by Fig. 4.5.

Parameter Space Dimensionality Reducing the parameter space dimensionality is likely to reduce the computational cost of each iteration, leading to a reduced overall optimisation time. This is particularly useful for the development process of building an optimisation framework. However, it is acceptable for the fully developed optimisation framework to require an extended period of time to run as it only needs to be performed once for a given set of conditions. Optimis-

ing in a reduced parameter space may also improve the ability of the algorithm to reach a global solution, a benefit that would be advantageous during the final optimisation.

4.3.2 Optimisation Parameters

This section contains details of how the acquisition parameter schedules were varied during the optimisation. The framework simultaneously optimised three types of parameters for the sequence: excitation flip angle, echo time and inversion pulse positioning.

4.3.2.1 Analytical Flip Angle and Echo Time Functions

In this thesis, analytical functions were used to control the echo time and flip angle patterns. These patterns were designed by modulating a Gaussian distribution (Eq. 4.1) with a swept-frequency sinusoidal function known as a chirp (Eq. 4.2). This enabled the generation of smoothly varying acquisition parameter schedules with adjustable oscillation frequency and temporal shift. The minimum and maximum values of the modulated patterns were altered via the spread and positioning of the Gaussian function. The Gaussian function allowed more degrees of freedom for variation in the schedules by allowing oscillations with varying amplitudes through the sequence. Therefore, the FA and TE patterns for the entire sequence could be optimised by adjusting only five parameters for each of these two acquisition parameter types. Using analytical functions in this way provided a relatively low-dimensional way of controlling the sequence and appears to have not been used in previous MRF optimisation publications.

The parameter I in Eq. 4.1 and Eq. 4.2 represents the timepoint index. The parameters m and σ in Eq. 4.1 are the mean and standard deviation of the distribution, respectively. The Gaussian distribution was centred on the timepoint $I = m$.

$$y_{Gaussian}(I) = \exp\left(\frac{-(I-m)^2}{2\sigma^2}\right) \tag{4.1}$$

The chirp function frequency (Eq. 4.2) was designed to vary through the schedule, reaching the transition frequency f_{chirp} at timepoint I_{chirp}. The parameter f_0 was the initial frequency. For chirp functions in general the waveform frequency naturally increases with time for cases of $f_0 < f_{chirp}$. However, in this work the waveform was forced to increase in frequency regardless of the relative magnitudes of f_0 and f_{chirp} by inverting the time axis and swapping f_0 and f_{chirp} in the case of $f_0 > f_{chirp}$.

$$y_{chirp}(I) = cos(2\pi\frac{(f_{chirp}-f_0)}{3I_{chirp}^2}I^3 + 2\pi f_0 I) \tag{4.2}$$

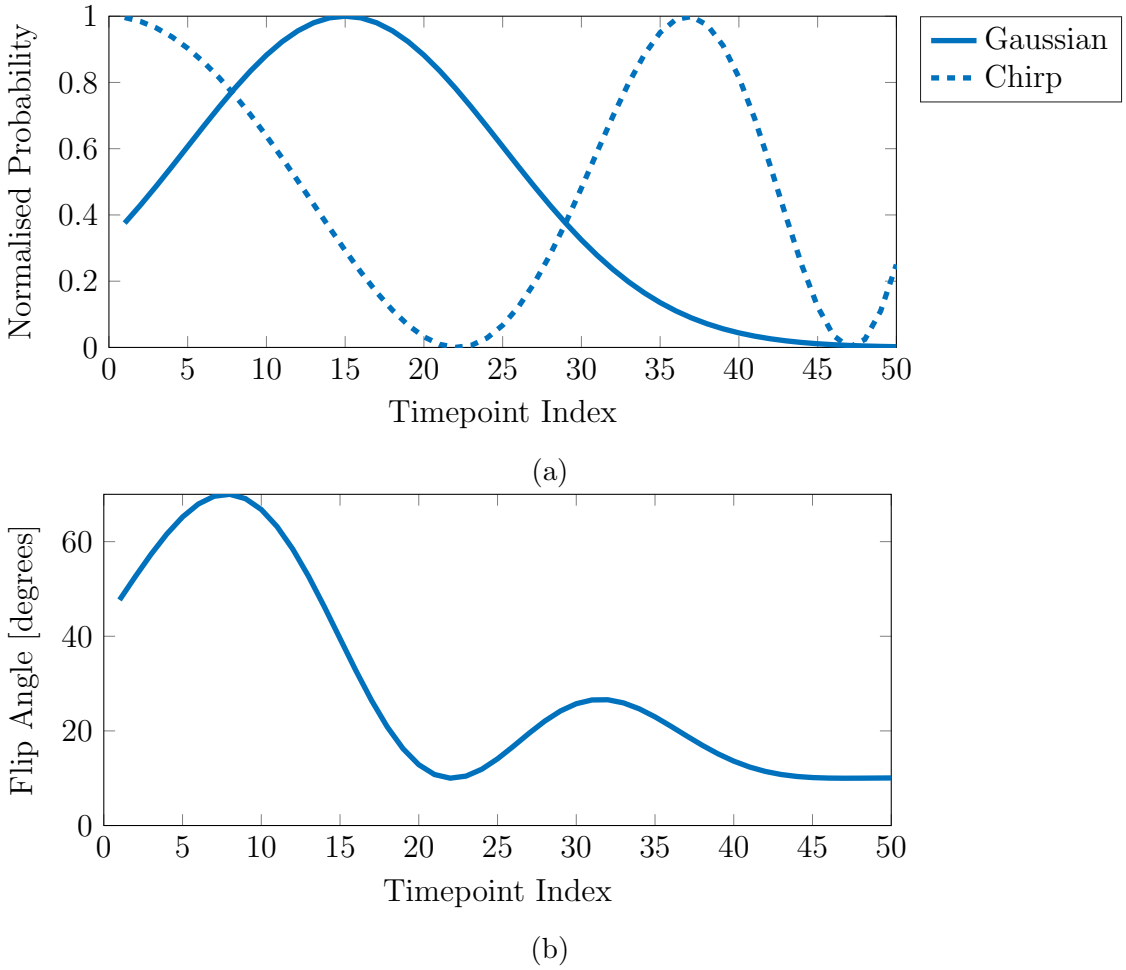

Figure 4.7: Example analytical functions (a) and a demonstration of how they were multiplied and scaled to produce a flip angle pattern (b). The same approach was used to design the echo time pattern.

Repetition Time The TR was fixed throughout the acquisition and was long enough to accommodate the maximum TE set before optimising. The framework was designed to have the option of adjusting the TR in accordance with changes in the maximum TE during the optimisation. This would allow the sequence duration to alter and so would provide some freedom for the algorithm to optimise for increased scan efficiency if desired.

TR Block Length The choice of block length could affect the T_1 contrast, as long blocks would result in a long effective TR for the final timepoint of each block. A suitable block length would make it possible to achieve low T_1 measurement error within the pre-set limits of the excitation flip angle, inversion pulse positioning and number of slices. The number of timepoints per block was fixed during each optimisation and so T_1 sensitivity was generated by varying the flip angles and inversion pulse positioning.

4.3.2.2 Inversion Pulse Flags

The algorithm optimised which TRs did not have a preceding inversion by switching the inversion pulses on or off using a binary flag, thereby manipulating M_z to minimise the T_1 error. An inversion pulse would be played before the relevant timepoint with a switched on flag. The maximum number of timepoints with the flag switched off was set before each optimisation. Figure 4.8 shows an example inversion flag schedule. Theoretically it was possible for the optimisation to produce an inversion flag pattern where only a single timepoint did not have an inversion. However, with a T_1 error metric this scenario was unlikely to occur, as it would probably lead to greater T_1 error compared to a situation with inversions distributed through the sequence, which would produce more occurrences of inversion recovery-style behaviour.

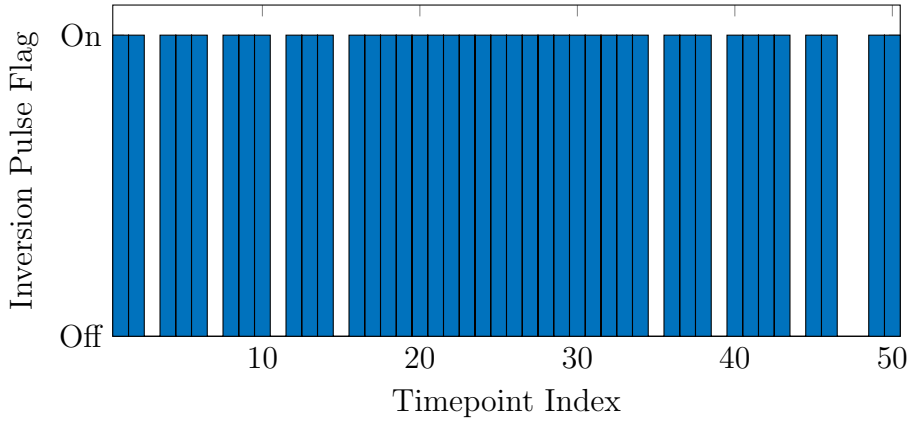

Figure 4.8: Example inversion flag position pattern for a single-slice acquisition.

4.3.2.3 Target T_1 and T_2

The optimisation framework was designed to minimise the measurement error of target T_1 and T_2 values. To accomplish this an optimisation dictionary was generated using T_1 and T_2 values above and below the target values, with uniform parameter increments in the two dimensions.

4.3.3 Choice of Cost Function

Orthogonality The accuracy of MRF relies on each set of target tissue or system properties producing a specific pattern of signal evolution (i.e. fingerprint) over the sequence. Since MRF experiments conventionally take the dot product of the acquired data and a dictionary of simulated fingerprints it is intuitive to approach the optimisation problem by minimising the dot product of the fingerprints [69]. Ideally all fingerprints would be orthogonal to each other, meaning the dot product between fingerprints from different combinations of tissue and system parameters would be zero. In such a situation the dot product of the dictionary

with itself would produce a matrix with all off-diagonal elements equal to 0 and all diagonal elements equal to 1 (i.e. the identity matrix). However, solely using the dot product as a metric does not account for image noise and may produce fingerprints with low SNR and hence poor accuracy of the measured parameters.

4.3.3.1 Monte Carlo Error Simulations

Monte Carlo simulations have been proposed as an effective way of determining the suitability of a MRF schedule, by estimating the parameter matching error distribution for a given noise level [69]. The optimisation framework in this thesis used a similar method at each iteration to simulate the MRF experiment.

Optimisation Iteration First an optimisation dictionary was built using T_1 and T_2 around the target values, thus aiming to ensure good fingerprint performance over a range of T_1 and T_2 values centred about a target pair of physiologically plausible values. Multiple synthetic data timecourses were then generated for each target T_1 and T_2 by adding Gaussian noise to the corresponding dictionary time course which was simulated using solutions to the Bloch equations. Finally, the synthetic data were fitted to the optimisation dictionary (using the dot product [1]) and the parameter errors were calculated with respect to the target T_1 and T_2. Repeating this process within each iteration with hundreds of different noise samples produced an error distribution for T_1 and T_2 individually. The distributions provided a numerical estimate of the systematic parameter assignment bias and spread for a given acquisition schedule. The cost function was based on these statistics. Schedules optimised using this kind of cost function will provide fingerprints with increased orthogonality and sufficient signal magnitude

to overcome the synthetic noise added during the process.

Root Mean Squared Error (RMSE) The cost function score for the initial examples was the RMSE of the matched T_1 distribution across all the target T_1 and T_2 pairs. Optimising with this metric has the effect of increasing the T_1 precision and accuracy. Using this metric an ideal acquisition schedule would lead to a distribution of T_1 error centred on zero (i.e. perfect accuracy) with zero standard deviation (i.e. perfect precision) for the target spin species. Alternatively, the cost function could be adapted to optimise for accuracy and precision independently by minimising the mean or standard deviation of the simulated error distributions. Examples with T_2 RMSE were also tested.

4.3.4 Choice of Algorithm

This work compared the optimisation performance of a Genetic Algorithm (GA) with that of an interior-point algorithm (*fmincon*, MATLAB, The MathWorks, Natick, MA). In work by Cohen *et al.* [67] the *fmincon* function was shown to outperform other algorithms in MRF schedule optimisation but was not compared with a GA.

Algorithm Optimisation Since the results are affected by the algorithm settings there is a degree of optimisation of the algorithm itself that can be performed, by adjusting the initial conditions, limits and stopping criteria. This is important to consider when comparing algorithms and means that it may be possible to further reduce the cost score. The optimisation of the algorithm is beyond the scope of this thesis but future research in this area would be valuable.

4.3.4.1 Genetic Algorithm

Genetic algorithms are modelled on the process of genetic evolution, with a population of parameters that evolves with each iteration towards a set that provides lower cost function values. They were developed to solve computationally costly problems [92, 93] and are theoretically able to escape local minima and provide a global search. This class of algorithm is also suitable for problems with large parameter spaces [94, 95]. In the field of MRF schedule optimisation GAs have received little use, with the timepoint-wise optimisation in [68] appearing to be the only example in the literature.

The framework in this thesis made use of the MATLAB (The MathWorks, Natick, MA) GA implementation to move towards an optimal set of acquisition parameters. This software is able to constrain particular parameters to integers, which is useful for optimising timepoint indices. In the context of this work the GA first randomly generated an initial population of acquisition schedules, where each schedule was known as a *genome*. Each genome was a vector of parameters to control the sequence acquisition parameters. The algorithm then calculated the *fitness* (i.e. the cost function value) of each genome and created a new generation of genomes. The new generation was built by mutating and swapping portions of the genomes with the best fitness (i.e. lowest cost function value), which was based on the relaxation parameter measurement error. The evolution continued until the algorithm reached pre-set limits, such as the maximum number of generations or the minimum fitness improvement.

Figure 4.9 contains the example fitness scores for a population of 100 schedules before and after optimisation of a single-slice acquisition. Here the cost function

score was the percentage T_1 RMSE, using target $T_1/T_2 = 1600\text{ms}/120\text{ms}$. After optimisation the population of schedules was more homogeneous, reflected in the narrowing of the cost function score distribution (Fig. 4.9b compared to Fig. 4.9a).

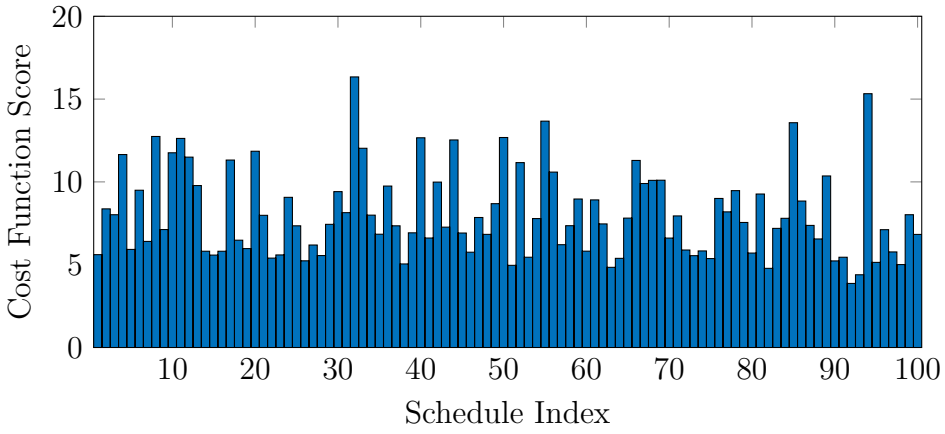

(a) After zero generations

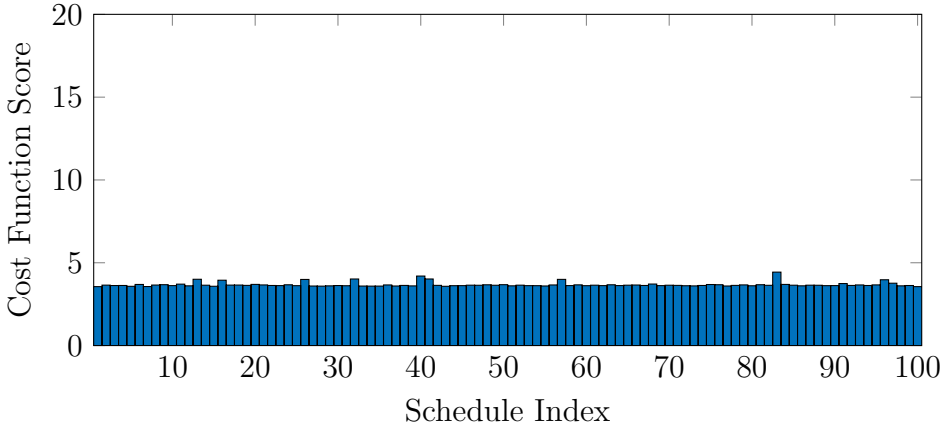

(b) After 20 generations

Figure 4.9: Example of the cost function scores for all schedules in the population, before (a) and after (b) a genetic algorithm optimisation. A single-slice optimisation with T_1 RMSE as the cost function metric, targeting $T_1/T_2 = 1600\text{ms}/120\text{ms}$.

Figure 4.10 shows an example evolution of the population cost function scores from the same optimisation as Fig. 4.9. The population appears to tend towards

the schedule with the lowest score, with a gradual change in the lowest score itself.

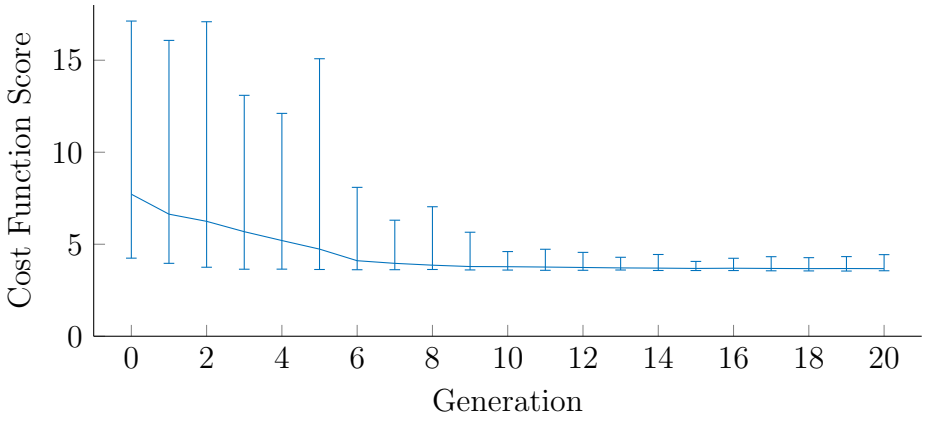

Figure 4.10: Genetic algorithm example: the evolution of the population cost function score distribution. The low and high points on the vertical bars represent the best and worst cost function score of the population, respectively. The middle point is the mean. The data shown here are from the same single-slice optimisation used for Figs. 4.9a and 4.9b.

4.3.4.2 Interior-Point (*fmincon*) Algorithm

At each iteration the interior-point algorithm assesses the gradient of a given cost function in different directions in the parameter space, before moving towards the parameters which give the most reduction in the cost function score. In this application the algorithm probed the change in cost function score for small changes in the acquisition parameters, before making adjustments to the schedule. Figure 4.11a shows the number of calculations performed per iteration of an example optimisation, with the corresponding cost function scores provided in Fig. 4.11b. Figure 4.11a illustrates that the extent of the parameter space exploration can vary across the iterations. Here an increased number of evaluations represents a more thorough testing of the local parameter landscape.

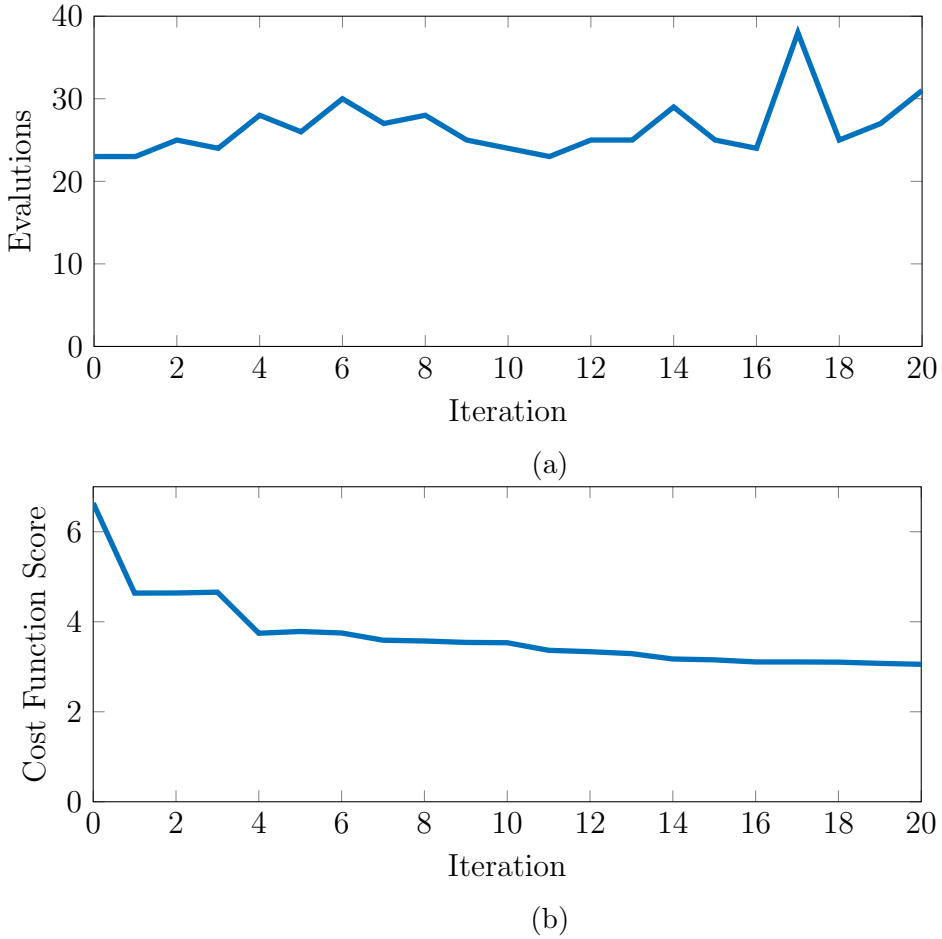

Figure 4.11: Example of an *fmincon* optimisation over 20 iterations. The algorithm performed multiple evaluations of the cost function for each iteration (a) before determining the optimal direction to move in the parameter space based on the lowest evaluation. The cost function score for the parameter set after each iteration is shown in (b).

4.4 Simulation Results

4.4.1 Single-Slice Scheme

As an initial example this work demonstrates the optimisation of a single-slice scheme with 50 timepoints, targeting T_1 and T_2 errors separately before consid-

ering them in combination. Figure 4.12 shows the acquisition schedules before and after T_1 optimisation with the GA and *fmincon* algorithms independently. For the GA results the first schedule of the initial population is compared with the optimised schedule. Figure 4.13 shows the same comparison but with schedules optimised in terms of T_2 error. Figure 4.14 shows the results of a combined T_1 and T_2 optimisation. The optimisation dictionary for these results used (min:increment:max) T_1 = [1200:10:2000]ms and T_2 = [90:1:150]ms. The maximum number of generations or iterations was set at 20 and the GA population comprised 100 schedules. The excitation flip angles were bound between 10 degrees and 70 degrees. The echo time (TE) limits were 21ms and 150ms. The maximum number of inversion flags switched off was fixed at 20% of the number of timepoints (i.e. 10 for a schedule of 50 timepoints). The synthetic data generated during the optimisation were created by adding complex Gaussian-distributed noise to the complex dictionary. The fingerprint magnitudes were used for the matching, with the complex noise set to possess a standard deviation of 0.015 after the magnitude conversion. Since the maximum simulated signal was 1 this meant the maximum SNR was $1/0.015 = 67$, corresponding to realistic SNR values that could be achieved in practical EPI readouts.

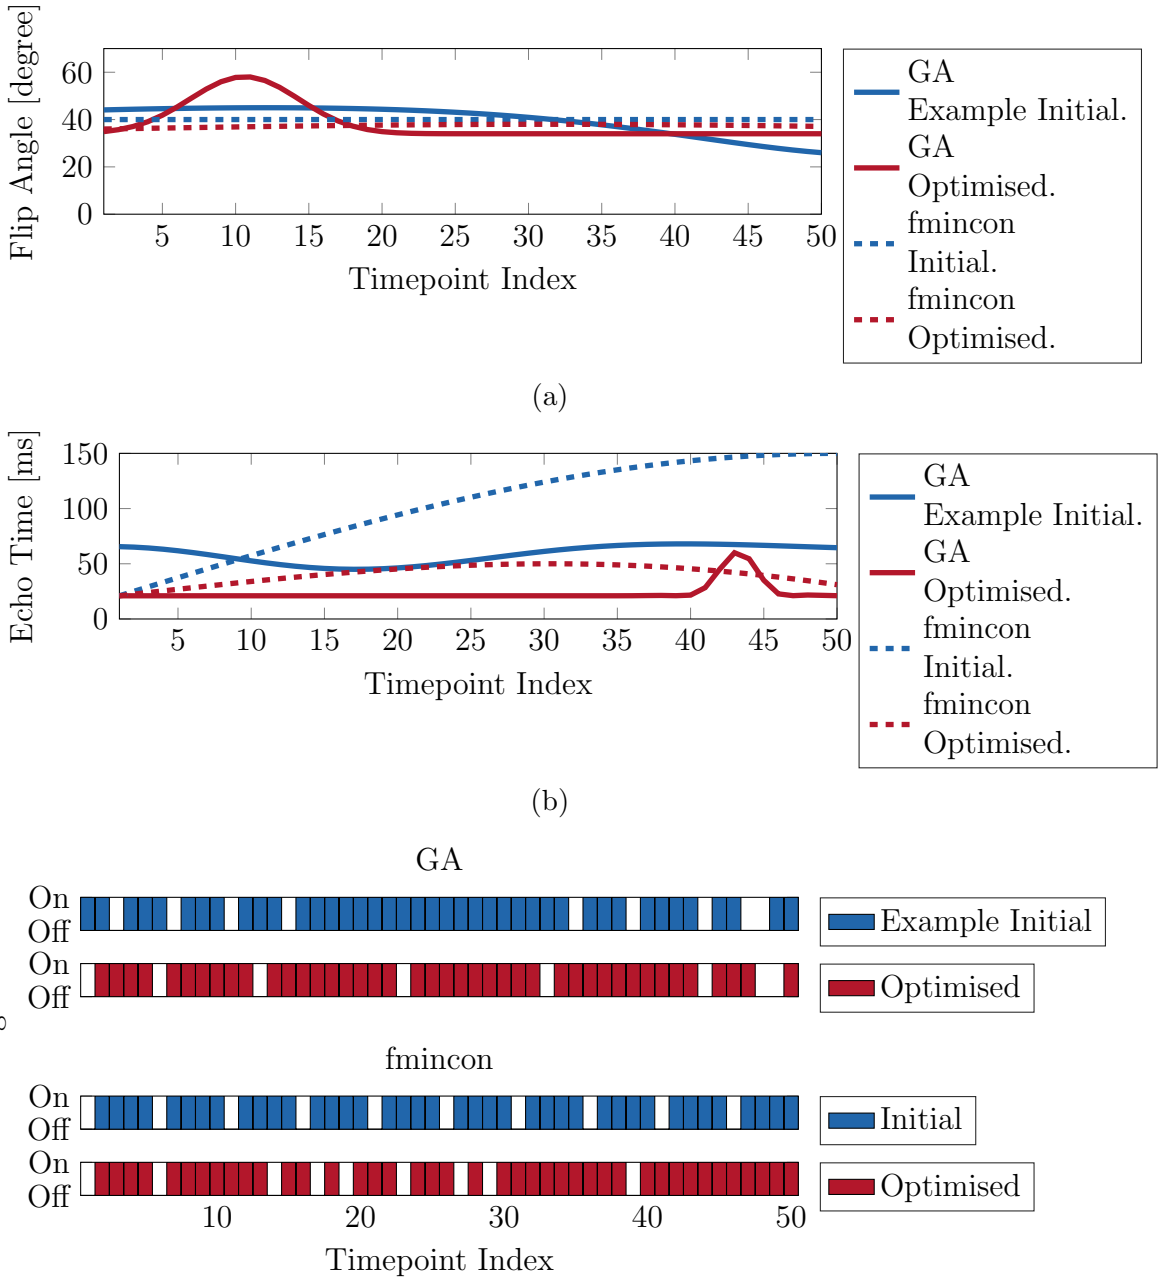

Figure 4.12: Single-slice T_1 RMSE optimisation with target T_1/T_2 = 1600ms/120ms. Initial (blue) and optimised (red) patterns for the flip angles (a), echo times (b) and inversion flags (c).

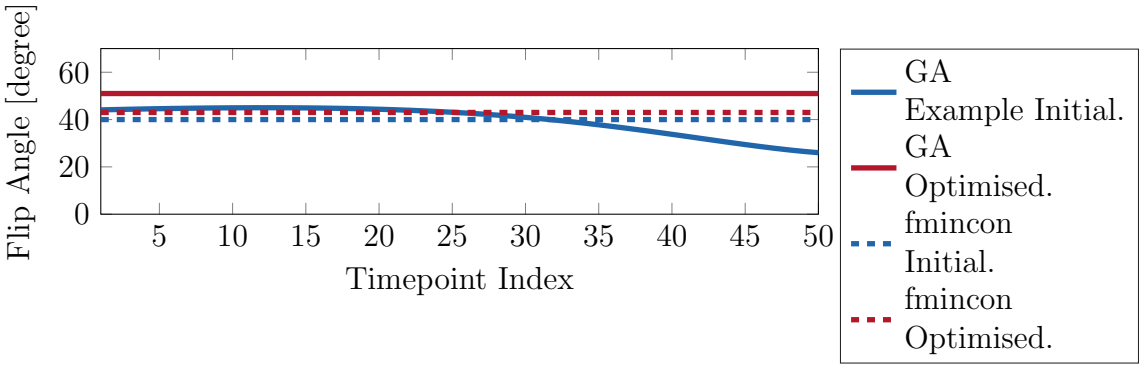

(a)

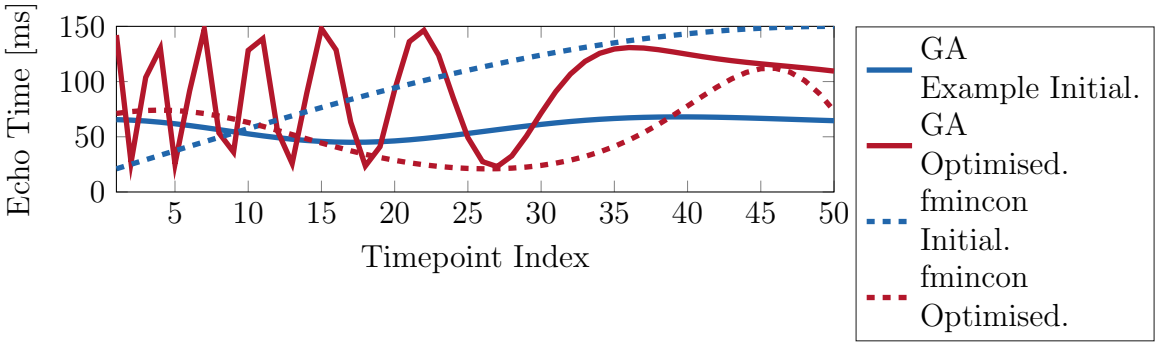

(b)

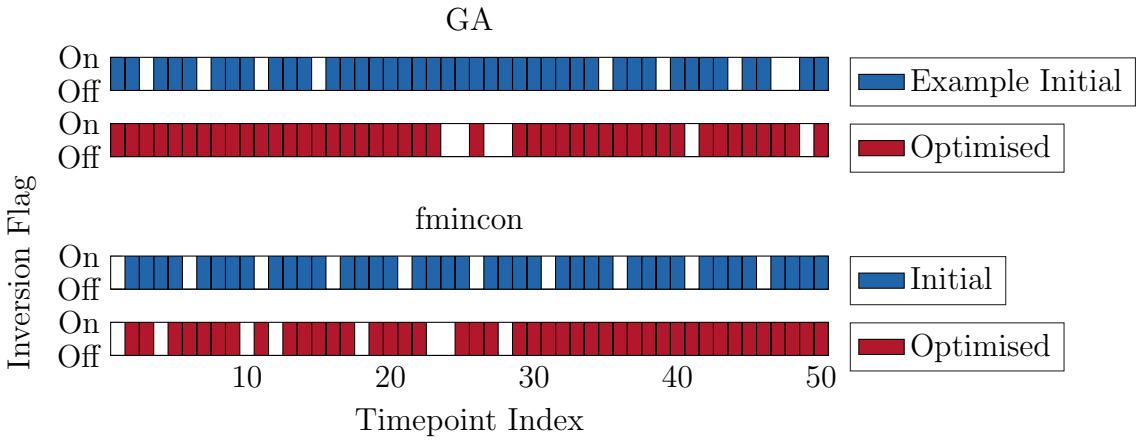

(c)

Figure 4.13: Single-slice T_2 RMSE optimisation, with target T_1/T_2 = 1600ms/120ms. Initial (blue) and optimised (red) patterns for the flip angles (a), echo times (b) and inversion flags (c).

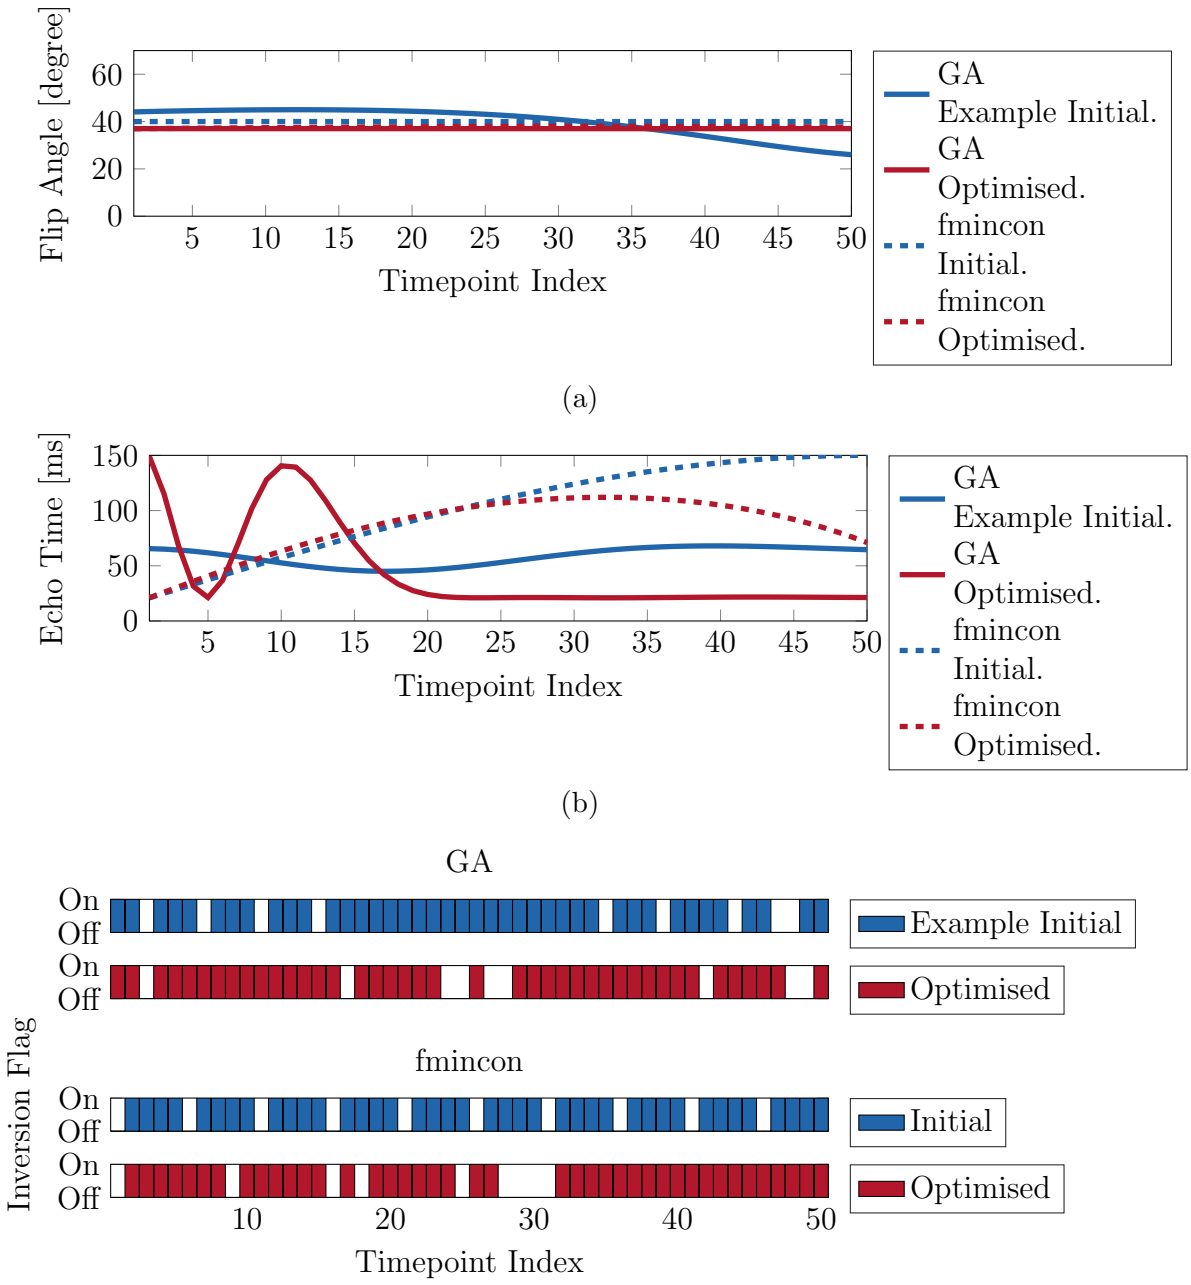

Figure 4.14: Single-slice combined T_1 and T_2 RMSE optimisation, with target T_1/T_2 = 1600ms/120ms. Patterns for the flip angles (a), echo times (b) and inversion flags (c).

4.4.1.1 Simulated Validation

Simulations were performed to predict the performance of the optimised schedules for a range of relaxation parameters other than those used during the optimisation. For each simulation the acquisition schedule was used to simulate synthetic data and a dictionary, in a similar way to each optimisation iteration. The dictionary for the validation was created with T_1 and T_2 similar to those used to match scan data: T_1 = [min:increment:max] [20:20:2000]ms and T_2 = [min:increment:max] [4:4:200]ms. This covers a broader range than the optimisation dictionary and contains larger increments. Figures 4.15 and 4.16 display the error from simulated testing using single-slice schedules optimised using the genetic and *fmincon* algorithms, respectively. The schedules used to generate Figs. 4.15 and 4.16 were both optimised for combined T_1 and T_2 RMSE. Figure 4.17 highlights the error for pairs of T_1 and T_2 similar to the measured values of the tubes in the custom phantom introduced in Chapter 3.

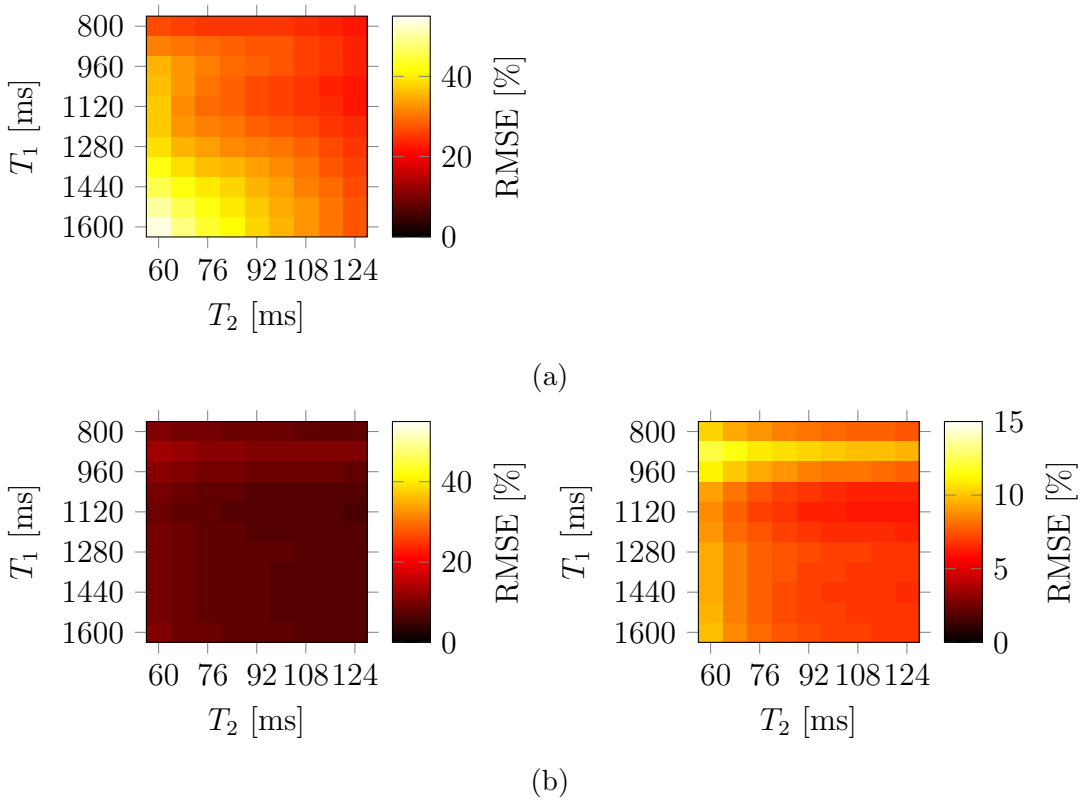

Figure 4.15: Genetic algorithm optimisation: simulated validation of a single-slice scheme, for the first member of the initial population (a) and the schedule optimised (b) in terms of combined T_1 and T_2 error. An additional plot with an expanded scale is shown for the optimised error maps, with an appropriate colour scale to highlight differences across the map.

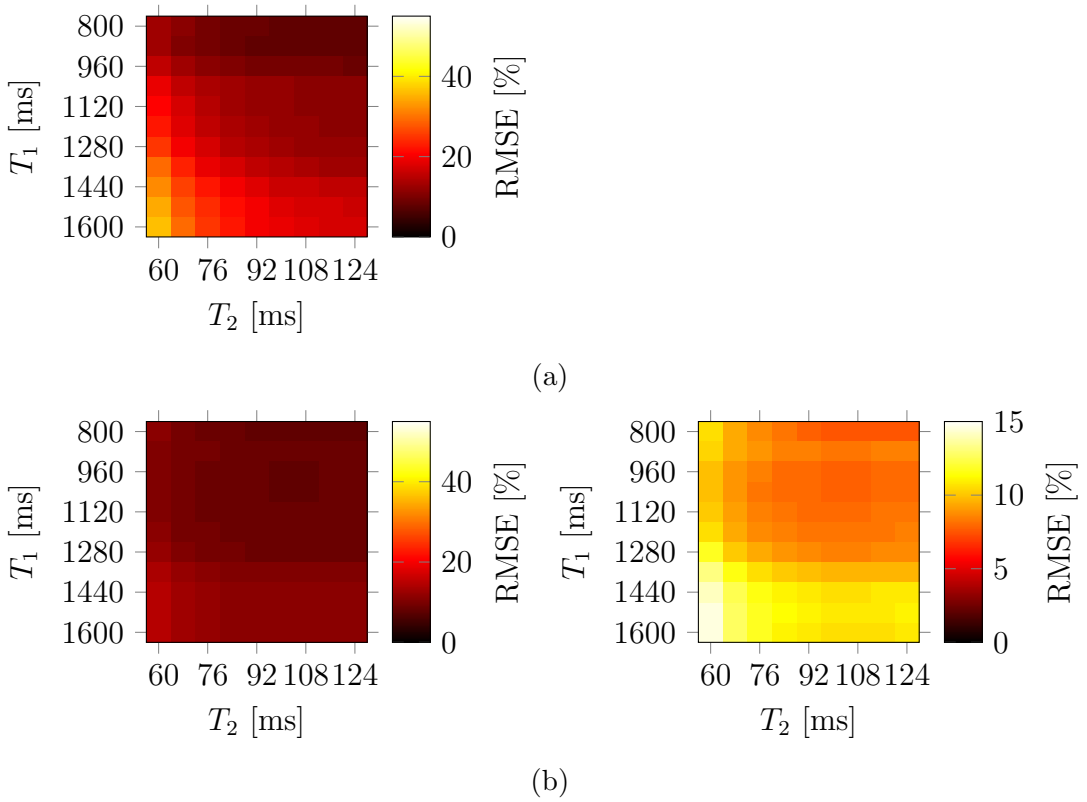

Figure 4.16: *fmincon* algorithm optimisation: simulated validation of a single-slice scheme, for the initial (a) and optimised schedules (b) in terms of combined T_1 and T_2 error.

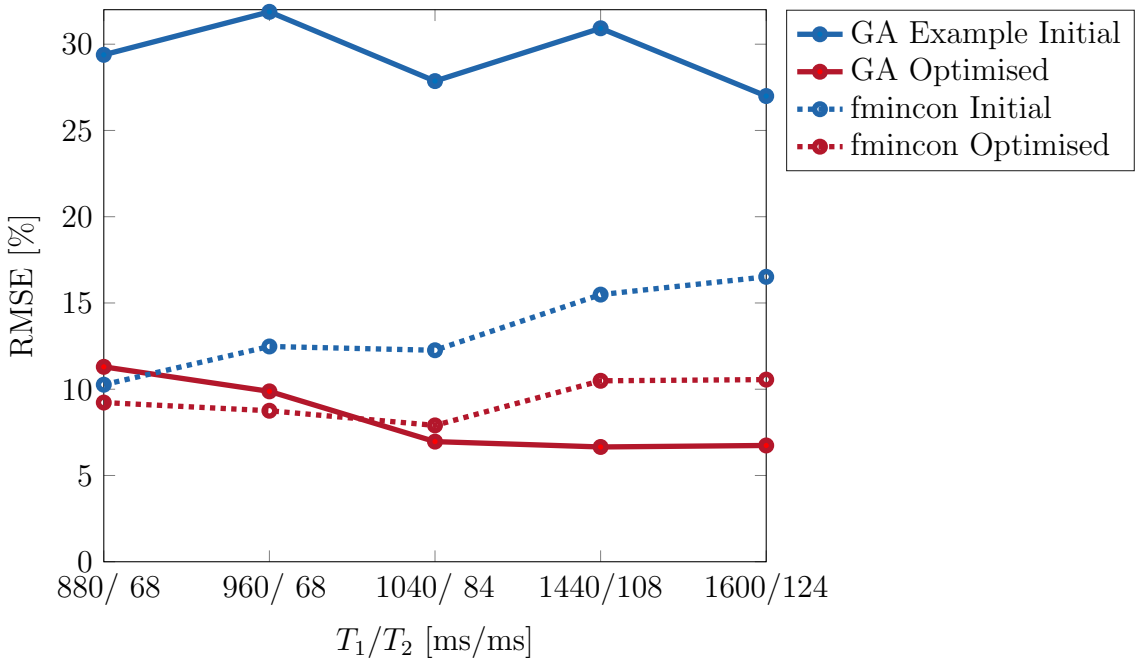

Figure 4.17: Simulated error for relaxation parameters similar to those of the custom phantom, using the single-slice schedules. Scan duration per slice: 50 × 400ms = 20 seconds.

4.4.2 10-Slice Scheme

This section includes results for a 10-slice scheme, using the acquisition approach described in Section 4.2.4 and portrayed in Fig. 4.5. The number of timepoints per block was set as 20% of the number of timepoints per slice. For example, for an acquisition of 20 timepoints per slice 4 timepoints were acquired in each block before moving to the next slice. The stopping criteria and acquisition parameter limits were the same as those used for the single-slice results. Figure 4.18 shows results from a 10-slice optimisation designed to minimise T_1 RMSE for T_1/T_2 = 1600ms/120ms. Figure 4.18a demonstrates the approximately linear relationship between the total optimisation time and the number of timepoints in the schedule,

particularly for the genetic algorithm. Figure 4.18b shows that the optimised cost decreases as the number of timepoints is increased, approximately following an exponential decay. Similar trends are apparent for results from an optimisation to minimise the combined T_1 and T_2 error (Fig. 4.19).

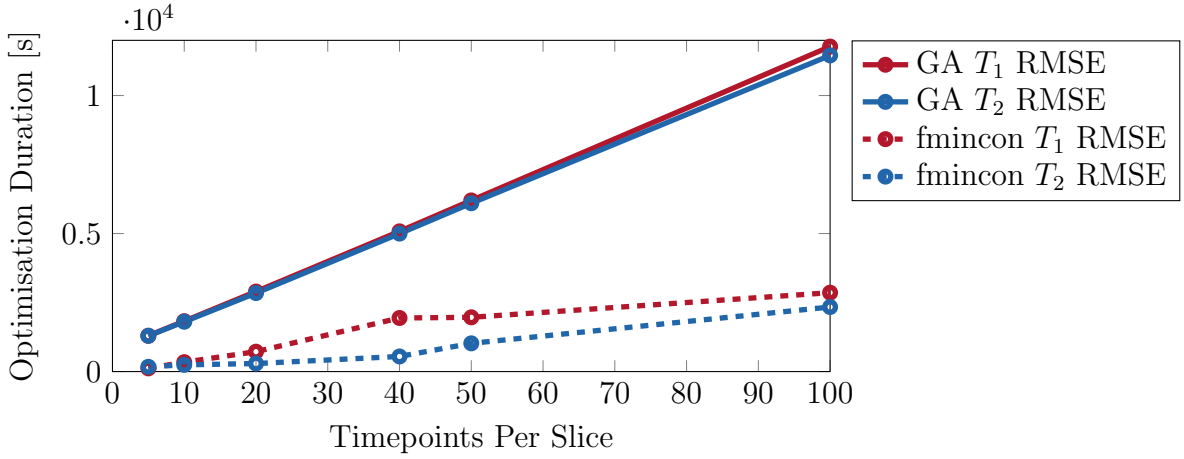

(a)

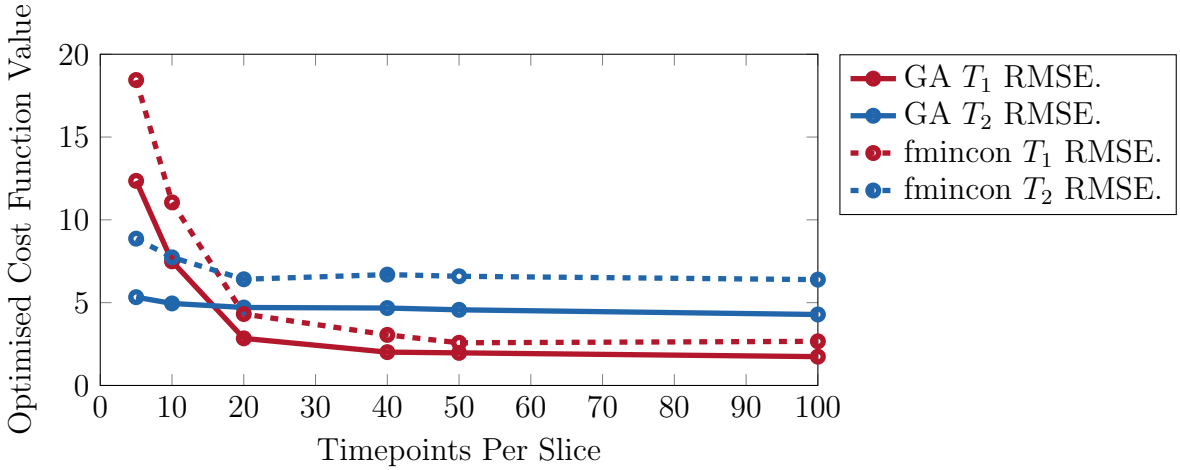

(b)

Figure 4.18: Optimisation duration (a) and final cost function (b) for a 10-slice sequence, minimising T_1 RMSE (red) and T_2 RMSE (blue) individually, for target T_1/T_2 = 1600ms/120ms. Each timepoint has a TR of 400ms, meaning that the total sequence duration for a 50 timepoint schedule would be 200 seconds (i.e. 3 minutes 20 seconds).

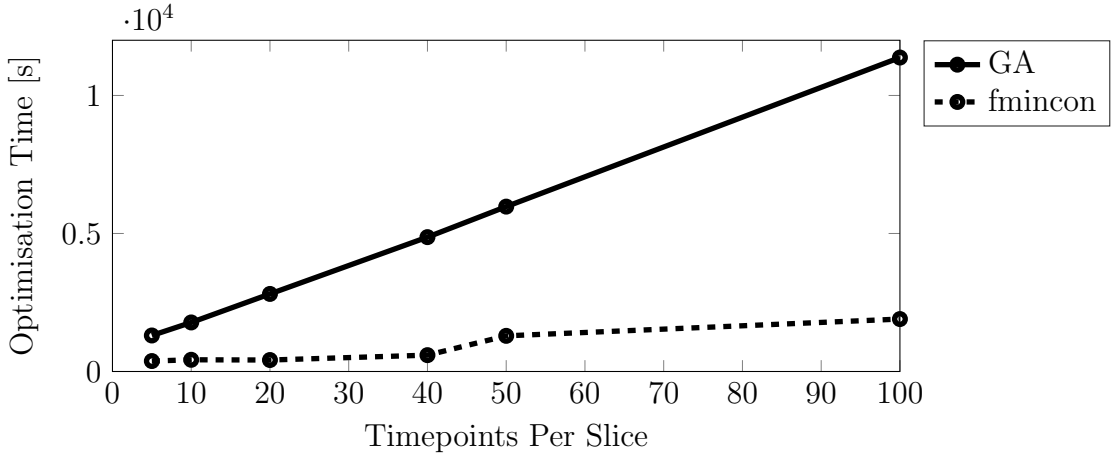

(a)

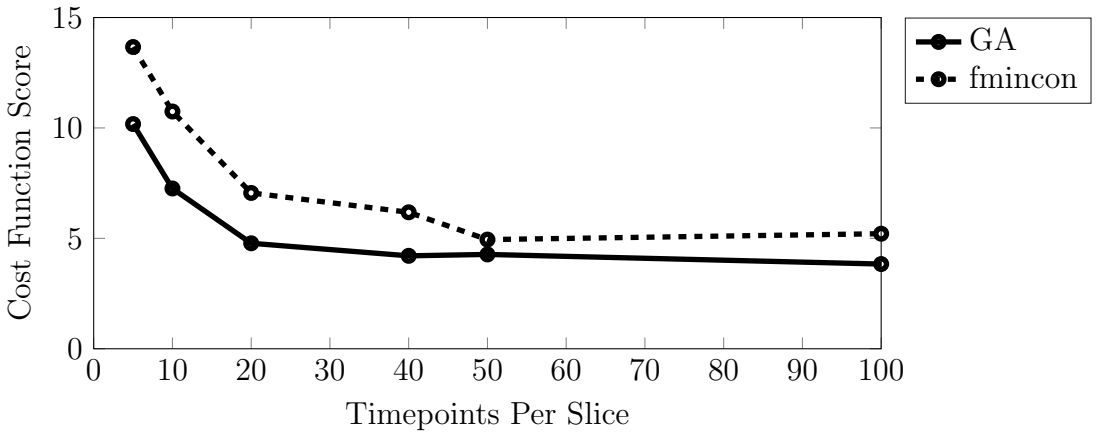

(b)

Figure 4.19: Optimisation duration (a) and final cost function (b) for the combined T_1 and T_2 RMSE of a 10-slice scheme, using various quantities of fingerprint timepoints.

Figure 4.20 compares initial and optimised schedules for the optimisations from Fig. 4.19 performed using 20 timepoints per slice. Figure 4.21 contains the same comparison for the optimisation with 50 timepoints per slice.

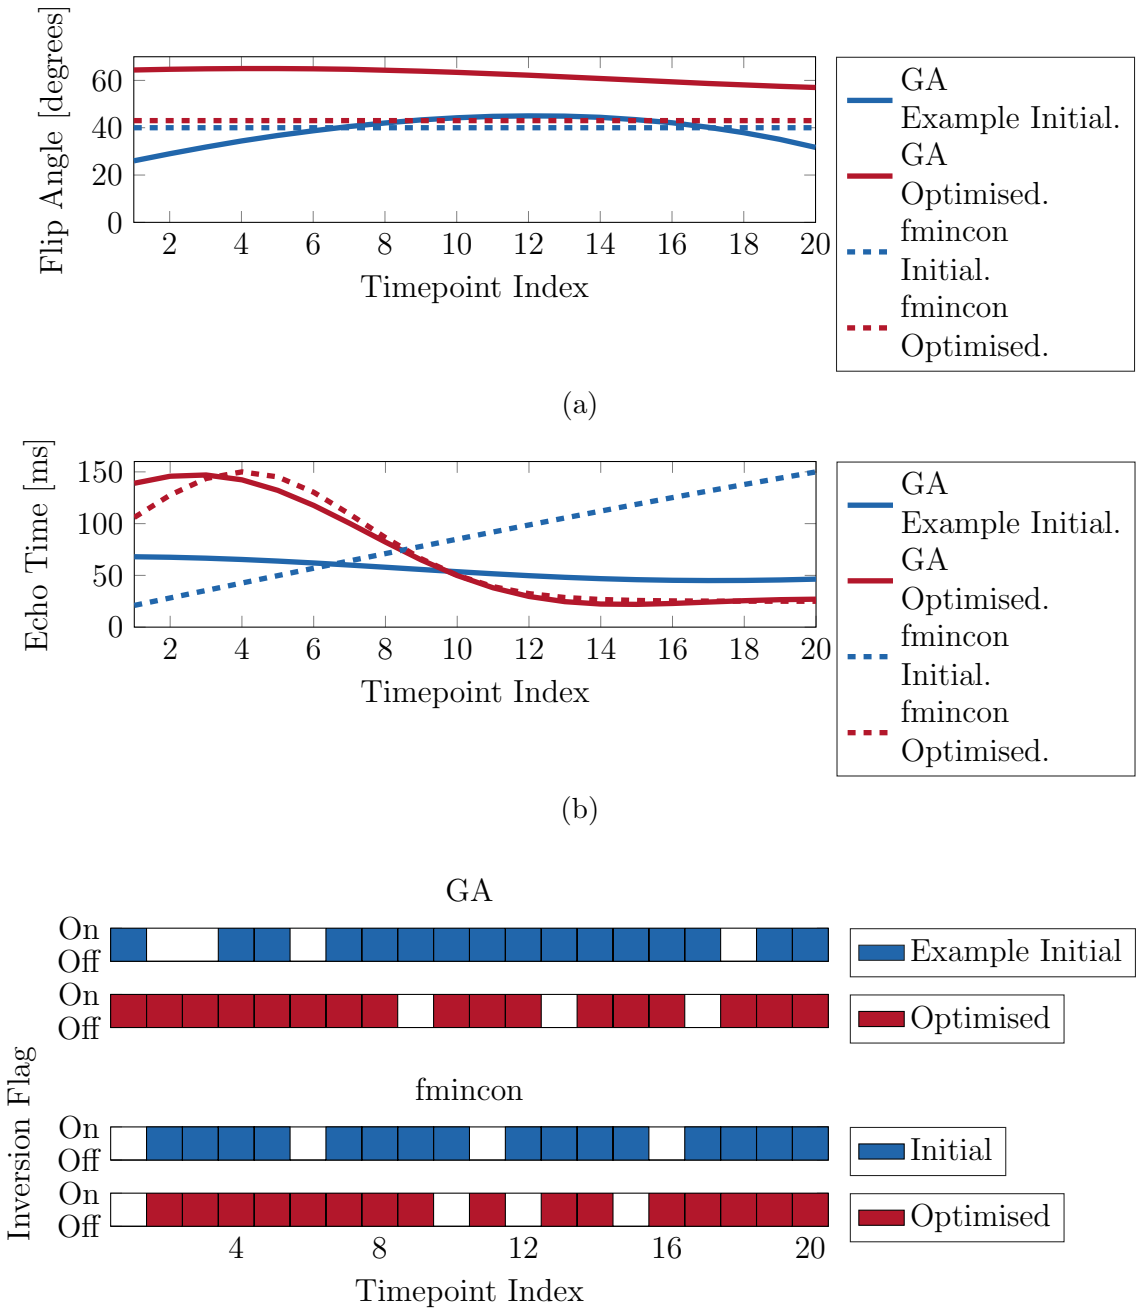

Figure 4.20: Flip angle (a), echo time (b) and inversion flag (c) schedules from optimisations for combined T_1 and T_2 RMSE, using a 10-slice scheme with 20 timepoints for each slice. For this simulation each TR block consisted of 4 TRs, with TR = 400ms. This meant 4 timepoints were simulated before moving to the next slice and the last TR of each block was effectively 10 × 400ms = 4000ms.

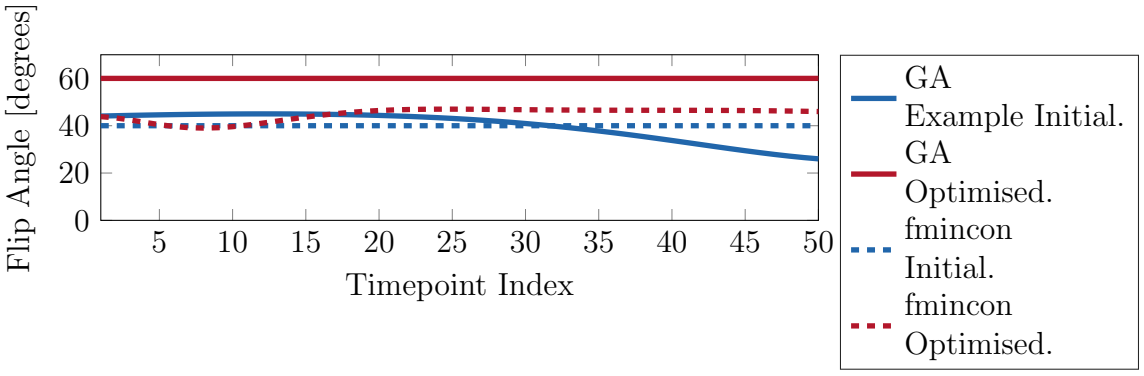

(a)

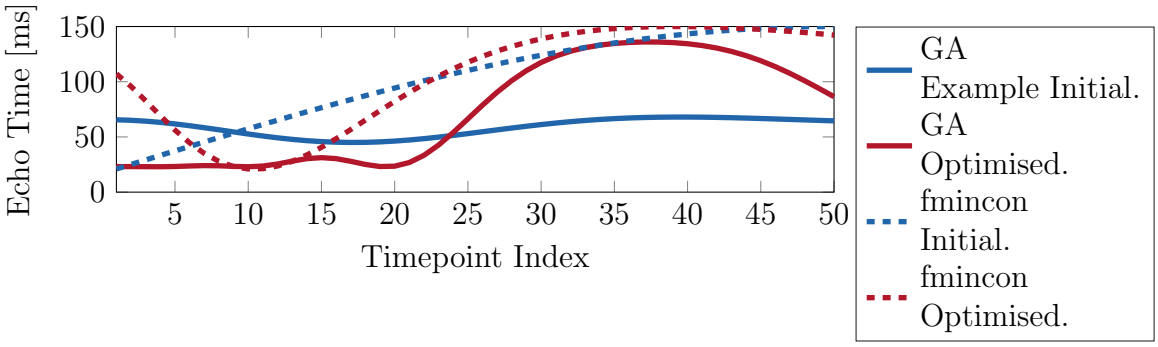

(b)

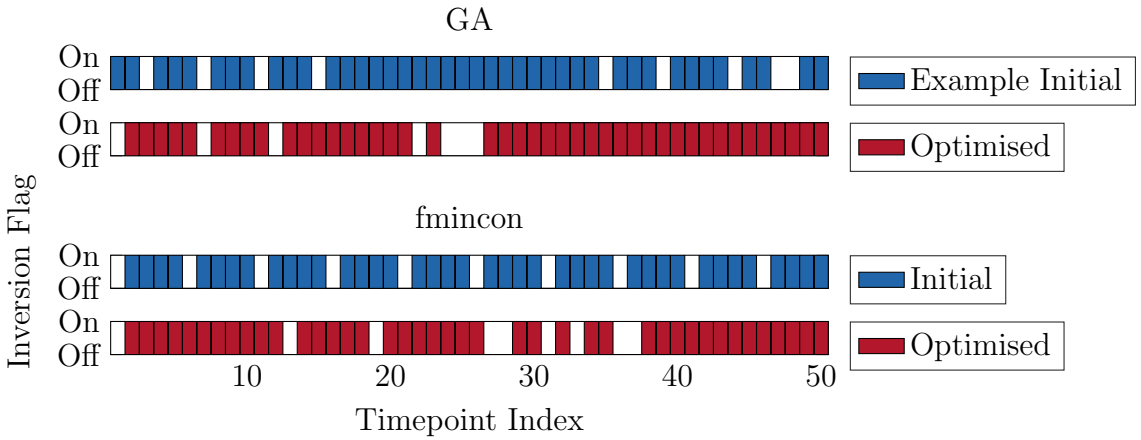

(c)

Figure 4.21: Flip angle (a), echo time (b) and inversion flag (c) patterns from optimisations for combined T_1 and T_2 RMSE, using a 10-slice scheme with 50 timepoints for each slice.

4.4.2.1 Simulated Validation

Simulated Validation: 20 Timepoints per Slice Figures 4.22 and 4.23 contain simulated error maps for pairs of T_1 and T_2 below those used in the 10-slice optimisation with 20 timepoints per slice (scan duration per slice: $20 \times 400\text{ms} = 8$ seconds). The results in Figs. 4.22 and 4.23 are from the genetic and *fmincon* optimisations, respectively. Figure 4.24 shows the simulated error for relaxation parameters similar to those of the custom phantom, using 20 timepoints per slice.

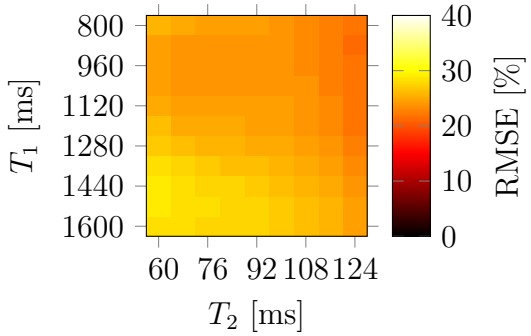

(a)

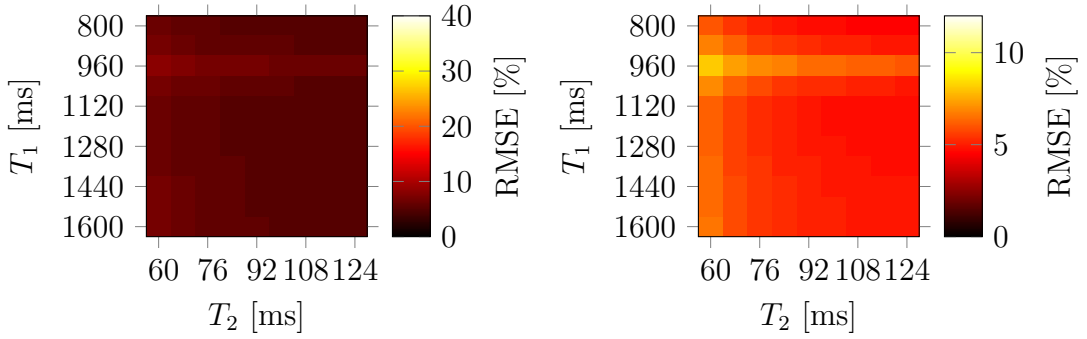

(b)

Figure 4.22: Genetic algorithm optimisation: simulated validation for a member of the initial schedule population (a) and the optimised schedule (b) for the 10-slice scheme with 20 timepoints per slice, optimised for combined T_1 and T_2 error.

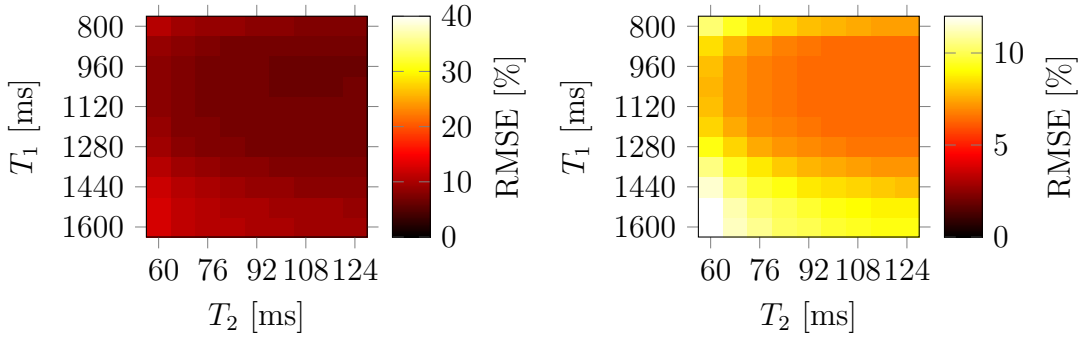

(a)

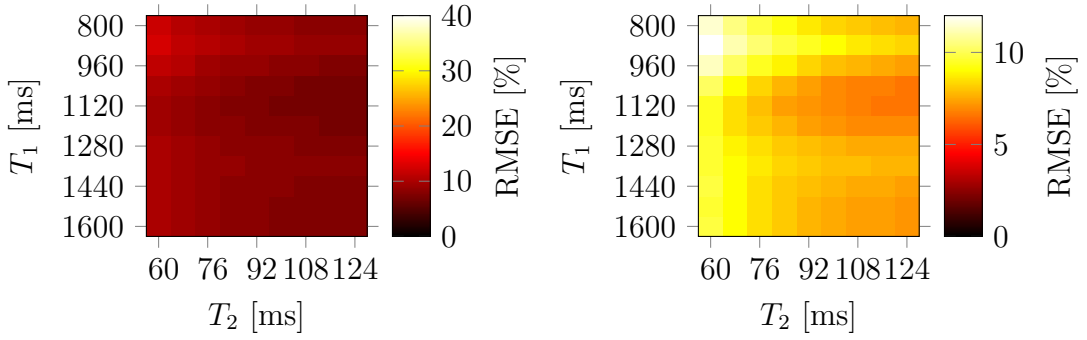

(b)

Figure 4.23: *fmincon* optimisation: simulated validation for the initial (a) and optimised (b) schedules of a 10-slice scheme with 20 timepoints per slice, optimised for combined T_1 and T_2 error.

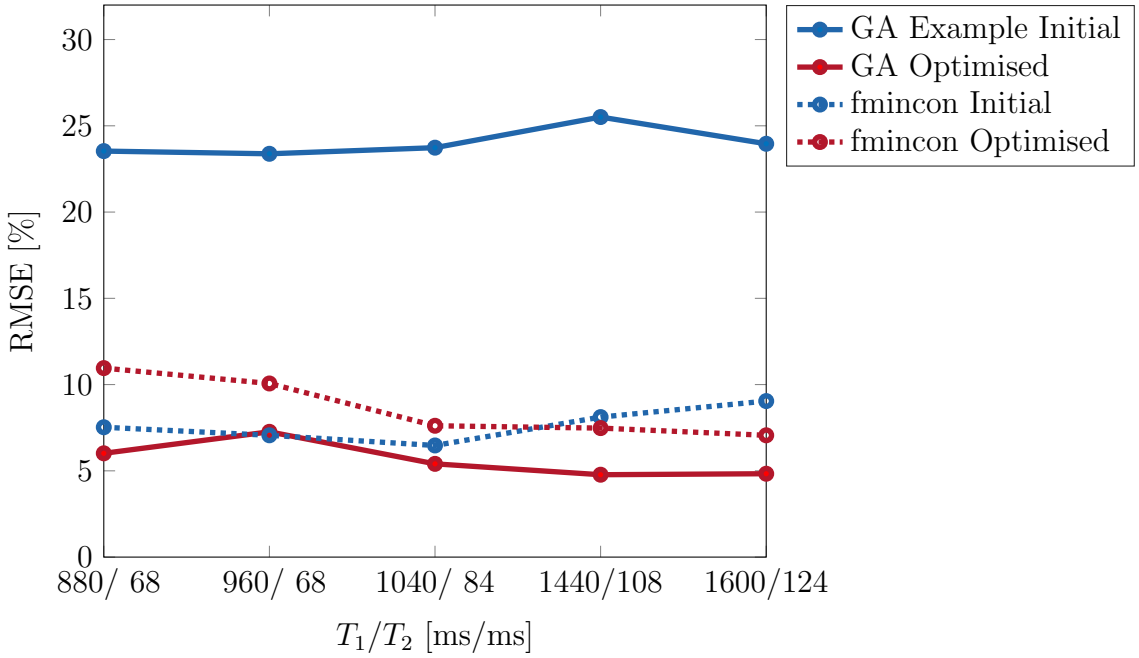

Figure 4.24: Simulated error for relaxation parameters similar to those of the custom phantom, using 20 timepoints per slice.

Simulated Validation: 50 Timepoints per Slice Figures 4.25 and 4.23 show simulated results for 50 timepoints per slice (scan duration per slice: $50 \times 400\text{ms} = 20$ seconds). Figures 4.25 and 4.23 were generated using schedules from the genetic and *fmincon* optimisations, respectively. Results for a selection of relaxation values close to those of the custom phantom tubes are provided in Fig. 4.27.

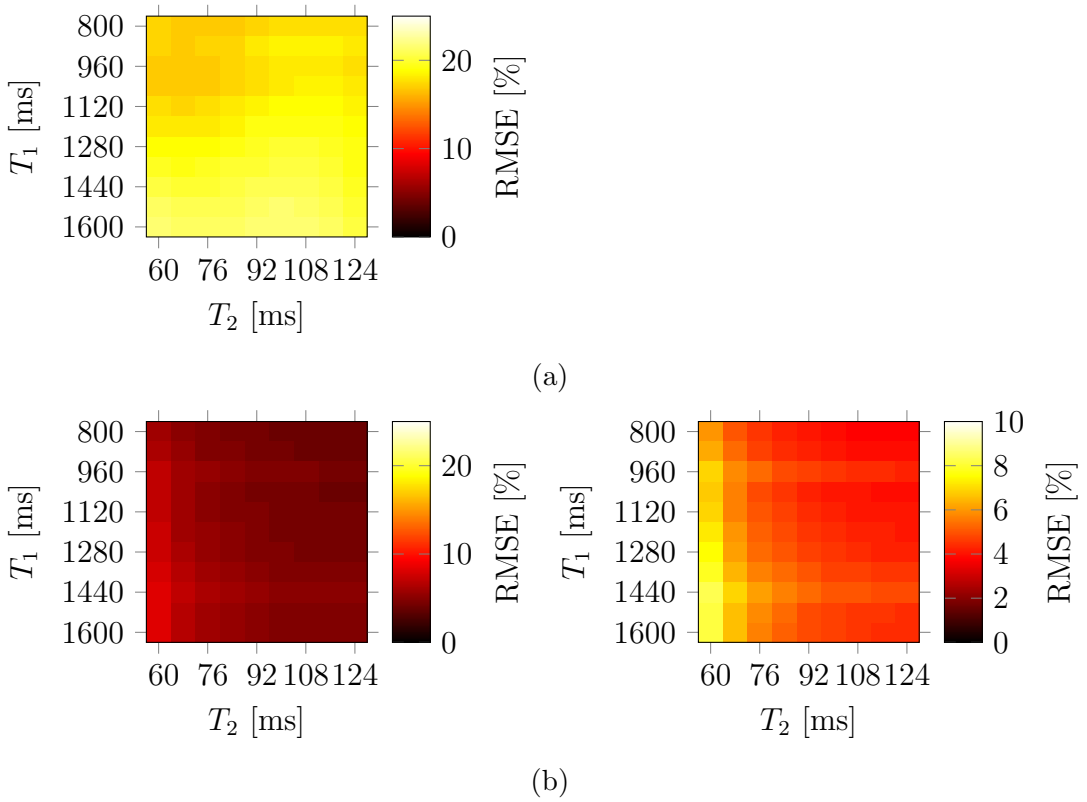

Figure 4.25: Genetic algorithm optimisation: simulated validation for a member of the initial schedule population (a) and the optimised schedule (b) for a single-slice scheme with 50 timepoints per slice, optimised for combined T_1 and T_2 error.

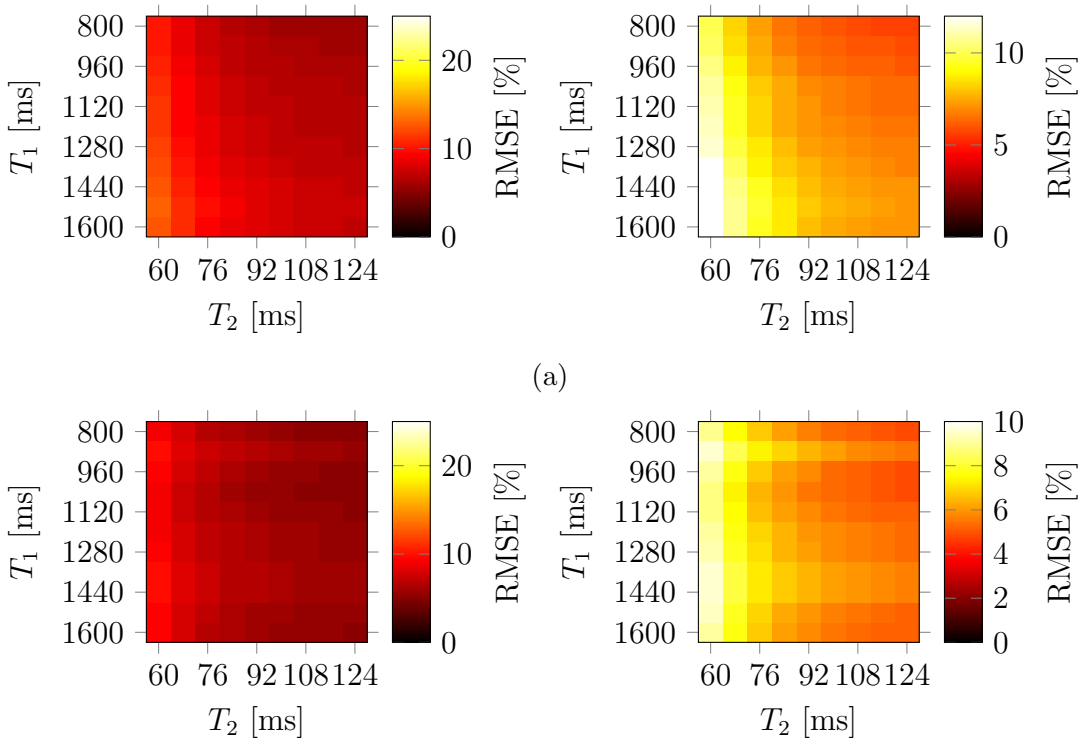

Figure 4.26: *fmincon* optimisation: simulated validation for the initial (a) and optimised (b) schedules of a 10-slice scheme with 50 timepoints per slice, optimised for combined T_1 and T_2 error. Two colour ranges are used to highlight the error differences within each map.

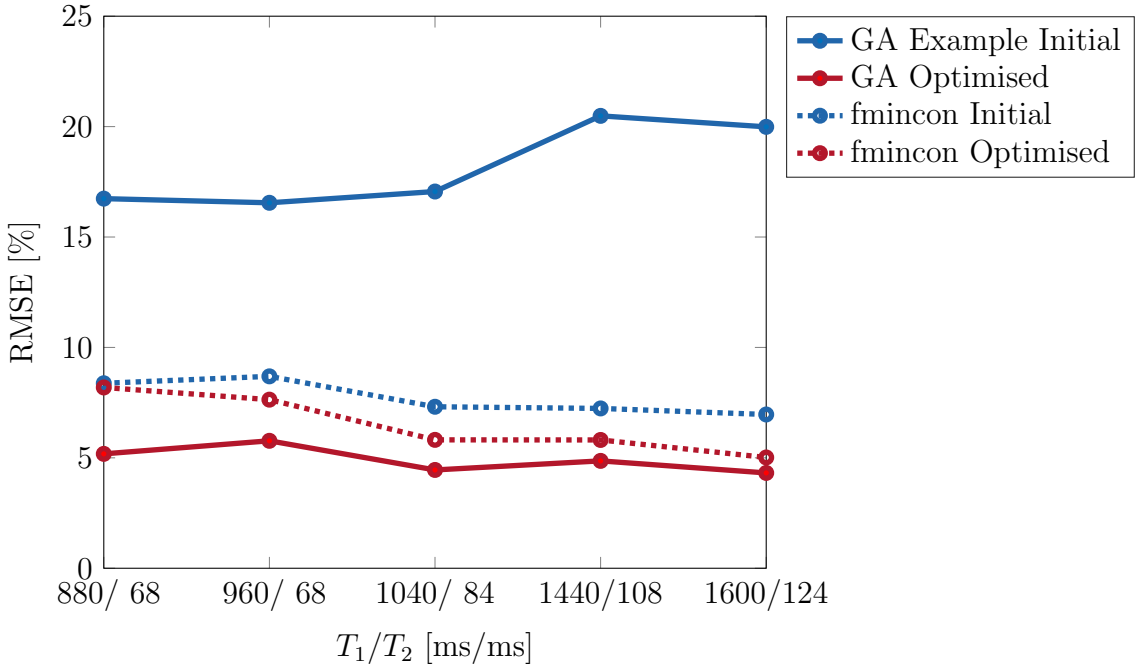

Figure 4.27: Simulated error for relaxation parameters similar to those of the custom phantom, using 50 timepoints per slice.

4.5 Discussion and Conclusions

This chapter demonstrated a framework for optimising the acquisition parameters of a multislice single-shot spin-echo MRF sequence. Schedules were optimised for single-slice and 10-slice schemes. A comparison was made between the optimisation results of interior-point (i.e. *fmincon*) and genetic algorithms.

The single-slice schedules optimised for T_1 and T_2 individually provided insights regarding the important schedule characteristics for mapping the two different parameters. The schedule optimised for T_1 rather than T_2 possessed a flat echo time over the majority of the timepoints, with low absolute values. This is intuitive as echo time variation alone does not contribute to T_1 contrast and is not

varied in gold standard T_1 mapping. The trend for optimal inversion flag positioning was less clear: the inversion flags for the GA-optimised schedule were more evenly distributed through the sequence than the corresponding initial schedule but the *fmincon* schedule possessed the opposite characteristic. As expected, the echo time patterns exhibited greater variation in the single-slice schedules optimised for T_2 rather than T_1. The optimised flip angle schedules for the two algorithms were both fairly flat and had similar magnitudes although the T_1-optimised schedule displayed a significant rise near the beginning of the sequence, which was not present in the T_2-optimised schedule. The GA inversion flag schedule optimised for T_2 was significantly less evenly distributed than the GA schedule optimised for T_1, intuitively suggesting the periods of inversion recovery were less important for T_2 measurements. The GA echo times and inversion flag positioning for combined T_1 and T_2 optimisation appeared to display attributes of the schedules optimised for T_1 and T_2 individually, although this was less clear for the *fmincon* results.

In the simulated validations there were some combinations of T_1 and T_2 where the initial schedules outperformed the optimised, particularly for the *fmincon* 10-slice schedules. It would be valuable for future work to investigate the performance of schedules optimised for these pairs of relaxation values as well as for multiple pairs in a single optimisation. For the target T_1 and T_2, the simulated validations showed that both algorithms successfully reduced the combined T_1 and T_2 error for schedules of 20 or 50 timepoints. This was evident in both the single and 10-slice cases, with the genetic algorithm providing the greatest minimisation for both schemes. Although the framework targeted a single specific pair of relaxation parameters the simulations predicted reasonable errors for T_1 and T_2 similar to

those of the other tubes in the phantom. The GA-optimised errors, going from the longest (the target) to the shortest T_1/T_2 values in the phantom, were 4.3% to 5.2% RMSE for the 10-slice scheme with 50 timepoints and 4.8% to 6.0% for 20 timepoints. The corresponding *fmincon* ranges were 5.0% to 8.2% and 7.1% to 11.0% for 50 and 20 timepoints, respectively. This suggests the framework would produce acceptable RMSE across the six compartments of the phantom if used to acquire scan data and that the GA-optimised schedules would achieve the best performance across the phantom, including the lowest error for the target T_1/T_2. To validate the trends presented in this chapter between unoptimised and optimised schedules, as well as between algorithms, experimental comparisons are provided in the following chapter.

Chapter 5

Multislice Snapshot MRF Implementation and Validation

5.1 Introduction

This chapter describes the implementation of the SE EPI MRF sequence optimised in the previous chapter. It includes practical details of the acquisition protocols (Section 5.2) and image reconstruction method (Section 5.3) required to produce the parameter maps. It also includes quantitative comparisons of the SE EPI MRF sequences with gold standard parameter-mapping methods and literature MRF sequences. Each section focuses on a specific topic and contains the methods, results and conclusions relating to the work on that topic. Section 5.4 contains the results of using the unoptimised and optimised 20-timepoint schedules from Chapter 4 with intrinsic B_1^+ fitting. Section 5.5 follows the same method as the work in Section 5.4 but uses the 50-timepoint schedules from Chapter 4, rather than the 20-timepoint schedules. Section 5.6 explores the effect of using a separate

B_1^+ map to correct the dictionary flip angles, using the 50-timepoint schedules.

Supplementary work for this chapter is provided in Appendices A and B. Appendix A contains work on the importance of considering the inversion pulse efficiency. Appendix B looks at the relationship between phantom parameter maps and the slice position relative to the phantom.

5.2 Data Acquisition

This section describes the protocols used to acquire data to validate optimised 10-slice schedules from Chapter 4. Schedules with 20 and 50 timepoints were used to explore the performance trade-off between parameter error and scan time. Later sections show results for phantom data and examples of *in vivo* maps from asymptomatic human volunteers.

In the phantom acquisitions Slice 1 (of 10) was positioned such that it sampled a cross-section of the phantom tubes, avoiding the supporting structure. The tubes were assumed to be homogeneous, given the production process, meaning any perpendicular cross-section of the tubes was assumed to possess the same relaxation properties.

The *in vivo* parameter maps presented were obtained using the optimised schedules tested on the phantom. Data from using the 20-timepoint and 50-timepoint schedules were acquired during separate sessions, with a different asymptomatic subject for each of the sessions. All scans were performed in accordance with local research ethics agreements.

5.2.1 Protocols

The phantom introduced in Chapter 3 was scanned with SE EPI MRF schedules optimised for 10 slices and a target $T_1/T_2 = 1600\text{ms}/120\text{ms}$, using the genetic and *fmincon* algorithms described in Chapter 4. To reduce the interaction between slices in the multislice acquisitions a 50% slice gap was used. For all sequence comparisons the following image acquisition parameters were used: FOV = 256mm, grid = 64x64, slice thickness = 5mm. A partial-Fourier factor of 6/8 was used to reduce the image acquisition time and enable shorter echo times. However, this leads to a trade-off between spatial resolution and echo time. Increasing the spatial resolution increases the duration of the EPI readout, which in turns increases the shortest achievable echo time. This means the minimum echo time used in this thesis would not be possible with a smaller pixel size, assuming no additional acceleration methods were applied. Therefore the SE EPI MRF sequence in this thesis compromises spatial resolution in favour of a relatively short minimum echo time, with zero-padding used to fill the missing region of k-space. Alternative filling methods are available, such as Projection Onto Complex Sets (POCS) [96] and homodyne detection [97], which provide images with reduced blurring compared to zero-padding.

5.2.2 Gradient Spoiling

To allow the signal to be accurately simulated using the Bloch equations and to avoid signal interference between slices it was important to spoil the transverse magnetisation at the end of each TR. The work of Hess *et al.* [91] demonstrated that varying the spoiling gradient moment with each timepoint, with a gradient

amplitude pattern that repeats every six TRs, is an effective way of spoiling transverse magnetisation. This thesis used a hexagonally-varying spoiling gradient scheme similar to that used in [91]. The scheme was implemented by varying the spoiling gradient amplitude in the x and y directions through the sequence, while keeping the z-direction gradient unchanged. This is described in Eq. 5.1, which shows how the x and y gradient amplitudes G_x and G_y varied as a function of the timepoint index I, whereas the z-direction amplitude G_z was kept at the pre-set maximum gradient amplitude G_{max}. For example, for the first timepoint $G_x = 0.5 \cdot G_{max}$ and $G_y = 0.866 \cdot G_{max}$.

$$\begin{aligned} G_x(I) &= cos(60I) \cdot G_{max} \\ G_y(I) &= sin(60I) \cdot G_{max} \\ G_z(I) &= G_{max} \end{aligned} \tag{5.1}$$

5.3 Image Reconstruction

For all the MRF sequences timepoint images were reconstructed offline using the raw scanner data. Images were reconstructed from each coil channel which were then combined to produce a single complex image for each timepoint. Since the SE EPI MRF sequence in this thesis sampled k-space on a Cartesian grid a discrete Fourier transform was directly applied to the acquired data to produce the coil images. As in the image reconstructions in Chapter 3 an adaptive combine method [89] was used to estimate a sensitivity map for each coil channel and combine the channel images. Parameter maps were produced by matching the magnitude

signal fingerprints to magnitude signal dictionaries. For the dictionaries in this chapter the T_1 dimension range was [min:increment:max] [20:20:2000]ms and the T_2 range was [min:increment:max] [4:4:200]ms.

5.4 20 Timepoints (8 Seconds) per Slice

This section includes results from using 20-timepoint SE EPI MRF schedules with a B_1^+ efficiency dictionary dimension to fit B_1^+ efficiency during the matching process.

5.4.1 20 Timepoints (8 Seconds) per Slice: Introduction

It is desirable for the scan duration of the SE EPI MRF sequence to be short but still provide a low parameter matching error. The bSSFP and FISP implementations in this thesis require approximately 12 seconds of scan time per slice, a duration that reflects the typical literature examples of these methods. To be comparable to these methods it is desirable for the SE EPI MRF scan duration to be similar. The 20-timepoint schedules tested in the simulations in Chapter 4 satisfy this, requiring only 8 seconds of scan time per slice.

5.4.2 20 Timepoints (8 Seconds) per Slice: Methods

This section explores whether the optimised 20-timepoint schedules from Chapter 4 outperform the initial 20-timepoint schedules and whether using 20 timepoints is sufficient to enable the framework to overcome noise in the signal fingerprints.

Variations in B_1^+ can be incorporated into the signal model using measure-

ments from a separately acquired B_1^+ map in series with the MRF acquisition. However, in the interest of avoiding an increase in scan duration, B_1^+ efficiency was included as a dictionary dimension. The B_1^+ dictionary dimension was made by making a global dictionary with a flip angle scaling factor dimension, as well as dimensions for T_1, T_2 and inversion efficiency. The B_1^+ scaling factors ranged from 0.6 to 1.4, in increments of 0.1. This represented a B_1^+ efficiency range of 60% to 140%, in steps of 10%. An efficiency of 100% represents a one-to-one relationship between the prescribed flip angle and the delivered flip angle. For example, a prescribed flip angle of 50 degrees and a delivered flip angle of 55 degrees would mean that the efficiency was 110%. The increments and ranges of the T_1, T_2 and inversion efficiency dimensions of the dictionary were the same as those noted earlier in this Chapter. Quantitative assessments were made of the effect that fitting for B_1^+ variations had on the resulting T_1, T_2 and M_0 maps. Fitting with a B_1^+ dimension provided a B_1^+ map which was compared with an externally acquired map obtained using the 3D DREAM sequence. However, the accuracy and precision of the T_1 and T_2 measurements was the focus of this work.

The optimised and unoptimised 20-timepoint schedules from Chapter 4 were used to produce T_1, T_2, M_0, inversion efficiency and B_1^+ maps, using schedules from the genetic and *fmincon* optimisations targeting T_1/T_2 = 1600ms/120ms. Inversion efficiency was included as a dictionary dimension because of the conclusions of Appendix A. The unoptimised schedules are labelled as *initial* in this Section. Note that the genetic algorithm starts with a population of initial schedules and the initial schedule used in this section is the example initial schedule from Chapter 4. The T_1 and T_2 maps were compared with single-slice gold standard, bSSFP and FISP acquisitions with the same phantom. Comparisons

of the optimised schedules were also performed *in vivo*, comparing the T_1 and T_2 values of the white and grey matter by manually drawing masks on the T_1 maps. Measurement accuracy, precision and efficiency were used as comparison metrics for the phantom data, with measurement efficiency defined according to Eq. 5.2 [1], where μ is the measurement mean, σ is the measurement standard deviation and T_{slice} is the sequence duration per slice. From Eq. 5.2 it is clear that the measurement efficiency depends on the measurement precision but not the accuracy.

$$Efficiency = \frac{\mu}{\sigma\sqrt{T_{slice}}} \tag{5.2}$$

The scan durations of the phantom acquisitions in this section are provided in Table 5.1. The gold standard (reference) T_1 protocol had a longer scan duration than the T_2 protocol because it recorded more data points.

Reference T_1 **[s]**	**Reference** T_2 **[s]**	**SE EPI MRF [s]**	**bSSFP [s]**	**FISP [s]**
360	270	8	12.4	12.4

Table 5.1: The 20-timepoint SE EPI MRF scan duration per slice compared with the single-slice validation scans: T_1 and T_2 gold standard (reference) and bSSFP and FISP MRF.

5.4.3 20 Timepoints (8 Seconds) per Slice: Results and Discussion

Phantom T_1 Maps Figures 5.1 and 5.2 show the T_1 maps from the genetic and *fmincon* optimisations, respectively. The T_1 maps vary across the slices, with notable differences between the central slices and Slices 1 to 4, for the unoptimised

and optimised sets of T_1 maps. Ideally all slices would experience the same perturbation as Slice 1 and would sample identical cross-sections of the phantom, but the work in Appendices A and B shows that this may not be the case in practice. Figure 5.3 shows the T_1 distributions from Slice 1 of the data in Figs. 5.1 and 5.2, using manually drawn regions in each of the 6 tubes of the phantom.

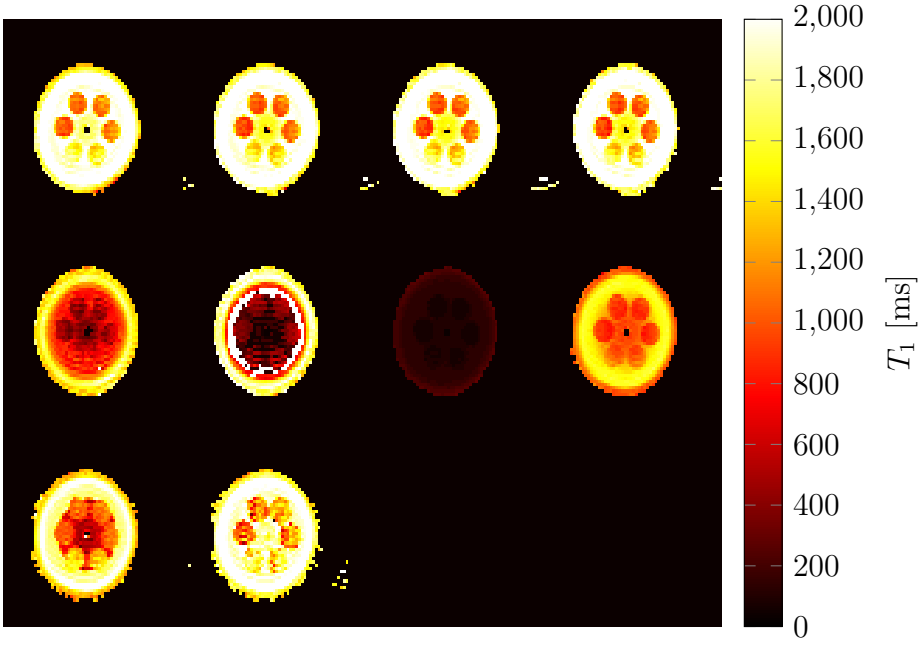

(a)

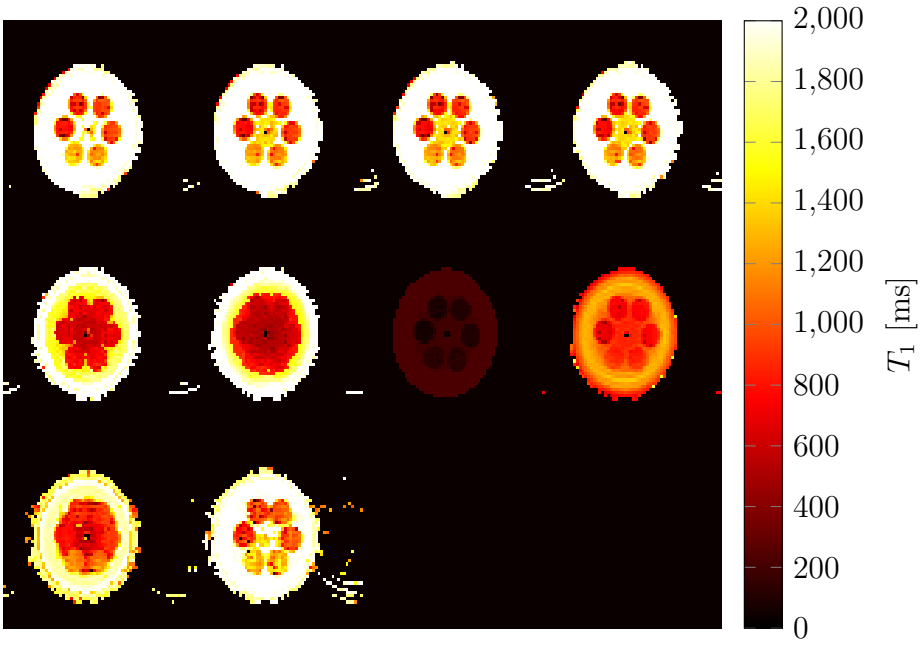

(b)

Figure 5.1: T_1 maps from genetic algorithm example initial (a) and optimised (b) 20-timepoint schedules with B_1^+ fitting.

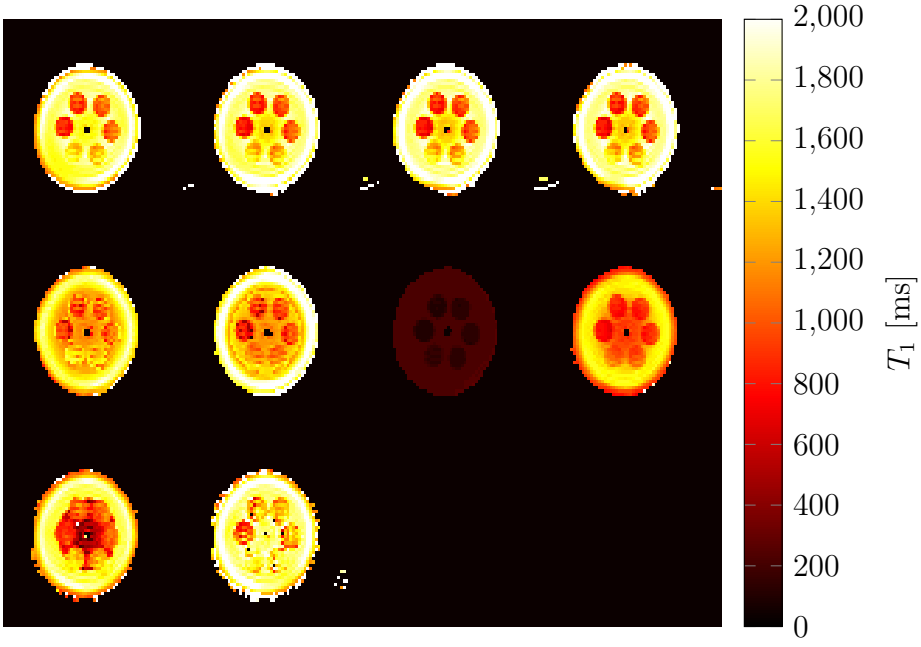

(a)

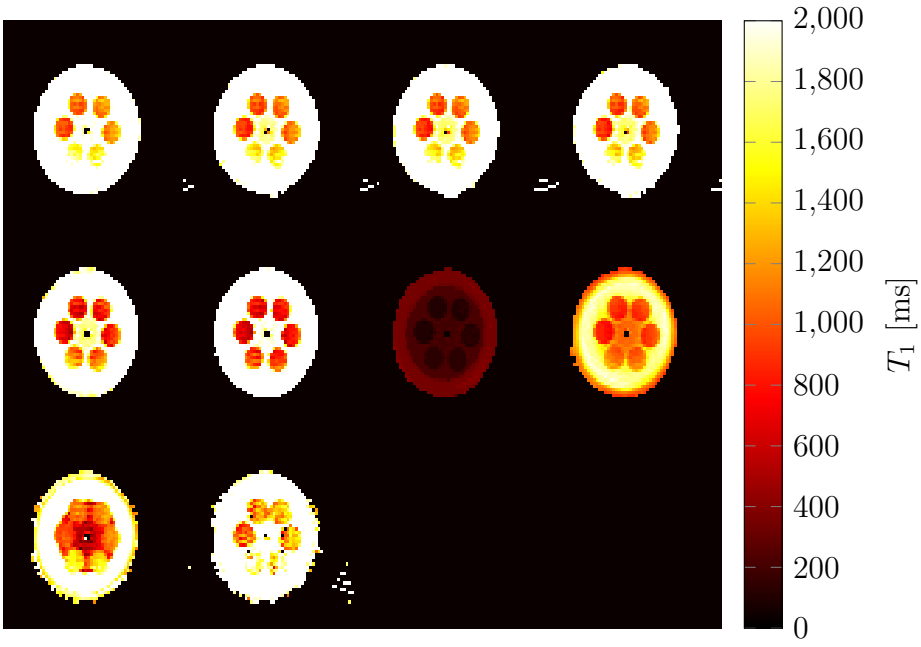

(b)

Figure 5.2: T_1 maps from *fmincon* algorithm initial (a) and optimised (b) 20-timepoint schedules with B_1^+ fitting.

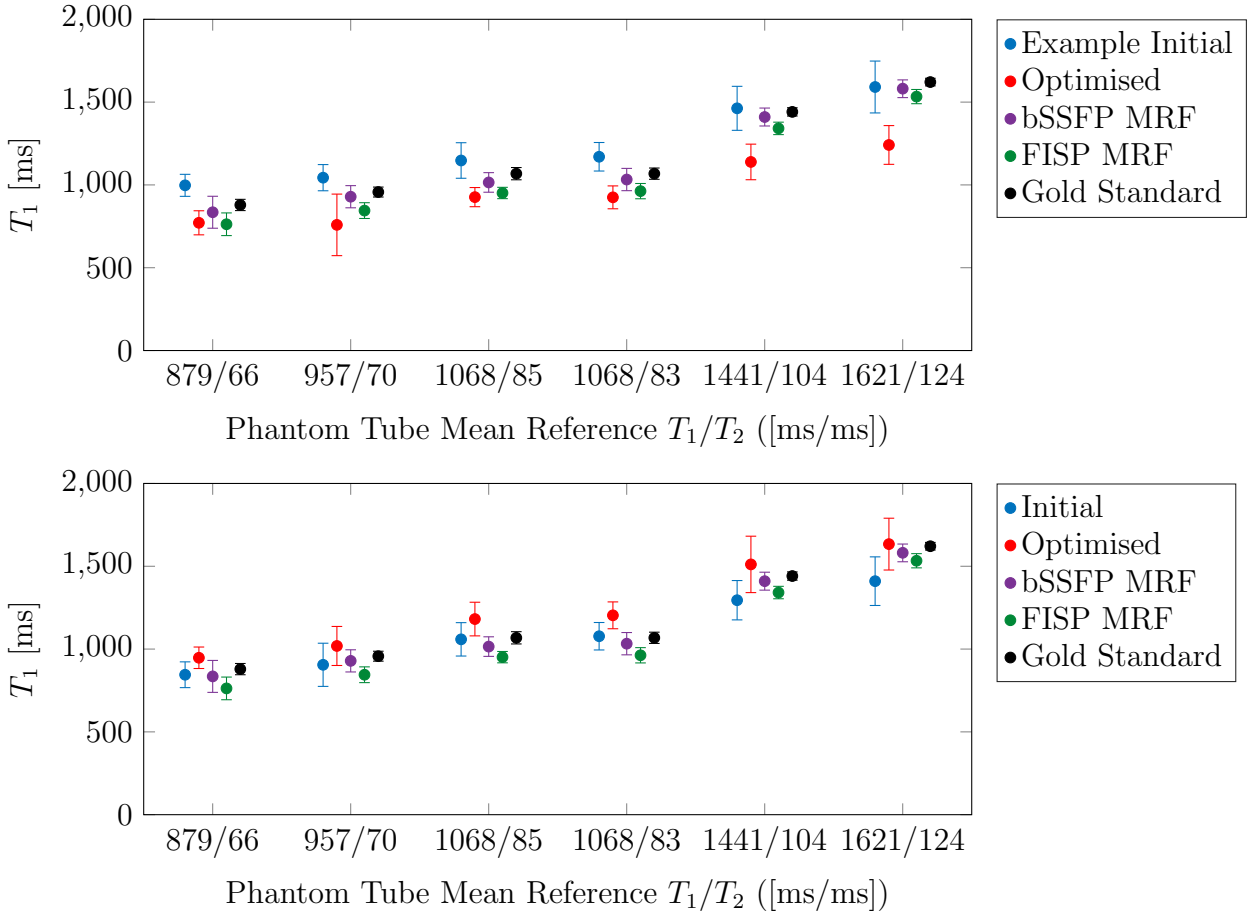

Figure 5.3: Comparison of the SE EPI MRF T_1 values for Slice 1 and reference maps, for the 20-timepoint genetic (top) and *fmincon* (bottom) schedules. Results are shown for each of the 6 phantom tubes.

Phantom T_1 Analysis Quantitative analysis of the T_1 distributions in Fig. 5.3 is provided by Tables 5.2, 5.3, 5.4, 5.5, 5.6 and 5.7. Each row of the Tables contains data from one of the tubes in the phantom. For context the mean reference T_1 or T_2 measurement is provided in the left column of each Table.

Tables 5.2 and 5.3 display the SE EPI MRF T_1 measurement bias for the genetic and *fmincon* results, respectively. The bias is a measure of the accuracy of the techniques and was calculated by taking the difference between the SE EPI

MRF and gold standard (i.e. reference) mean in the form of a percentage of the reference. Note that for the results from each algorithm it would be desirable to have a reduction in the bias for the optimised schedule compared to the corresponding initial schedule, because that would signify an increase in accuracy as a result of the optimisation. Each row of the tables represents results from one of the 6 tubes in the phantom, with the lowest bias (i.e. highest accuracy) highlighted in red. Table 5.3 shows that the bSSFP measurements were the most accurate for the majority of the tubes, but that the highest accuracy for the two tubes with the longest reference T_1 values was actually provided by the unoptimised schedule. In each of the 6 tubes the optimised case degraded the accuracy, contradicting the optimisation simulations.

Reference [ms]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
879	+13.5	-12.3	-5.0	-13.2
957	+9.1	-20.7	-2.9	-11.7
1068	+7.4	-13.3	-5.0	-11.0
1068	+9.6	-13.4	-3.3	-9.8
1441	+1.5	-21.0	-2.2	-6.9
1621	-1.8	-23.4	-2.5	-5.4

Table 5.2: Genetic algorithm T_1 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean, using 9 measurements for the gold standard method. The smallest absolute bias for each row is highlighted in red.

Reference [ms]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
879	-3.9	+7.8	-5.0	-13.2
957	-5.4	+6.4	-2.9	-11.7
1068	-0.9	+10.6	-5.0	-11.0
1068	+0.9	+12.7	-3.3	-9.8
1441	-10.1	+4.9	-2.2	-6.9
1621	-13.0	+0.8	-2.5	-5.4

Table 5.3: *fmincon* T_1 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean, using 9 measurements for the gold standard method. The smallest absolute bias for each row is highlighted in red.

Tables 5.4 and 5.5 display the SE EPI MRF T_1 measurement spreads for the genetic and *fmincon* results, respectively. A small spread is desirable. The spread was defined as the standard deviation as a percentage of the corresponding mean. This metric reflects the measurement precision: a small spread means high precision. The FISP MRF measurements are the most precise for all but the shortest reference T_1 (i.e. Row 1). In the genetic algorithm results (Table. 5.4) the unoptimised schedule leads to the highest precision for this tube, whereas for the *fmincon* results the optimised schedule outperforms the other acquisitions for this tube.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
879±3.9	6.6	9.5	11.5	8.9
957±3.1	7.6	24.5	7.2	5.7
1068±3.5	9.3	6.3	5.8	3.6
1068±3.2	7.4	7.5	6.5	4.8
1441±1.7	9.1	9.4	3.8	2.8
1621±1.4	9.9	9.4	3.4	2.8

Table 5.4: Comparison of the genetic algorithm T_1 spreads with the gold standard (reference), bSSFP and FISP measurements. The spread was calculated as the percentage standard deviation. The smallest spreads of the MRF sequences are highlighted in red. The reference column shows the means and percentage spreads of the gold standard measurements.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
879±3.9	9.2	6.9	11.5	8.9
957±3.1	14.4	11.6	7.2	5.7
1068±3.5	9.5	8.6	5.8	3.6
1068±3.2	7.7	6.7	6.5	4.8
1441±1.7	9.2	11.3	3.8	2.8
1621±1.4	10.4	9.5	3.4	2.8

Table 5.5: Comparison of the *fmincon* T_1 spreads with the gold standard (reference), bSSFP and FISP measurements. The spread was calculated as the percentage standard deviation. The smallest spreads of the MRF sequences are highlighted in red. The reference column shows the means and percentage spreads of the gold standard measurements.

Tables 5.4 and 5.5 display the SE EPI MRF T_1 measurement efficiency for the genetic and *fmincon* results, respectively. A large efficiency score is desirable. The efficiency of each scan was calculated using the duration per slice, provided in Table 5.1. In both sets of results the FISP acquisition was the most efficient for each of the tubes apart from that with the shortest T_1. For the genetic algorithm results the initial schedule was the most efficient, whereas for the *fmincon* results it was the optimised schedule.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
879	1.4	5.3	3.7	2.5	3.2
957	1.7	4.7	1.4	3.9	5.0
1068	1.5	3.8	5.6	4.9	7.9
1068	1.7	4.8	4.7	4.4	5.9
1441	3.0	3.9	3.8	7.4	10.0
1621	3.7	3.6	3.8	8.5	10.1

Table 5.6: Genetic algorithm T_1 efficiency per slice, with 20 timepoints per slice and fitted B_1^+. The largest efficiency for each row is highlighted in red. The reference T_1 means are shown in the first column.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
879	1.4	3.8	5.2	2.5	3.2
957	1.7	2.5	3.1	3.9	5.0
1068	1.5	3.7	4.1	4.9	7.9
1068	1.7	4.6	5.3	4.4	5.9
1441	3.0	3.8	3.1	7.4	10.0
1621	3.7	3.4	3.7	8.5	10.1

Table 5.7: *fmincon* T_1 efficiency per slice, with 20 timepoints per slice and fitted B_1^+. The largest efficiency for each row is highlighted in red. The reference T_1 means are shown in the first column.

Phantom T_2 Maps Figures 5.4 and 5.5 show the T_2 maps from the genetic and *fmincon* optimisations, respectively. Figure 5.6 compares the T_2 distributions from the tubes of Slice 1 with those of single-slice gold standard, bSSFP and FISP measurements.

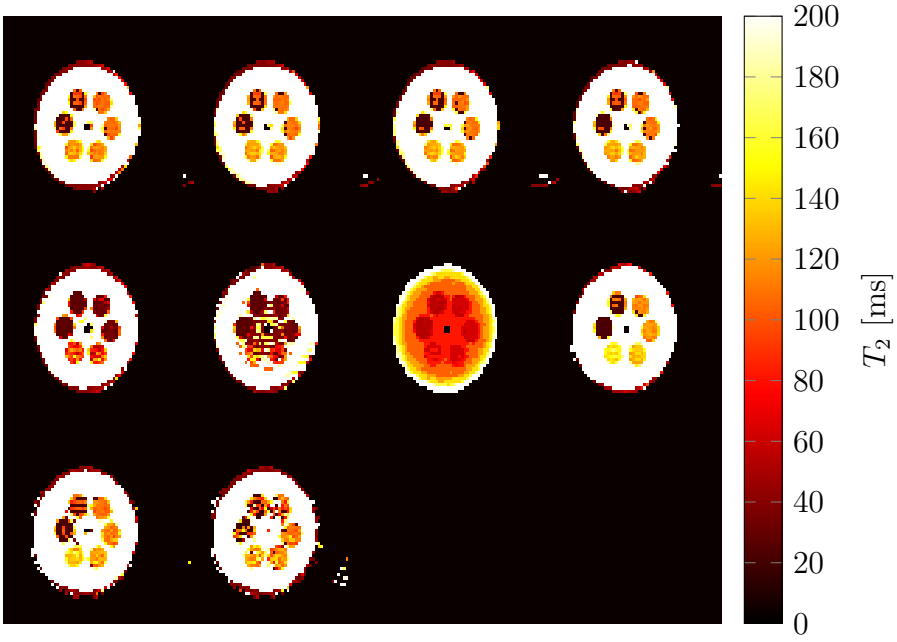

(a)

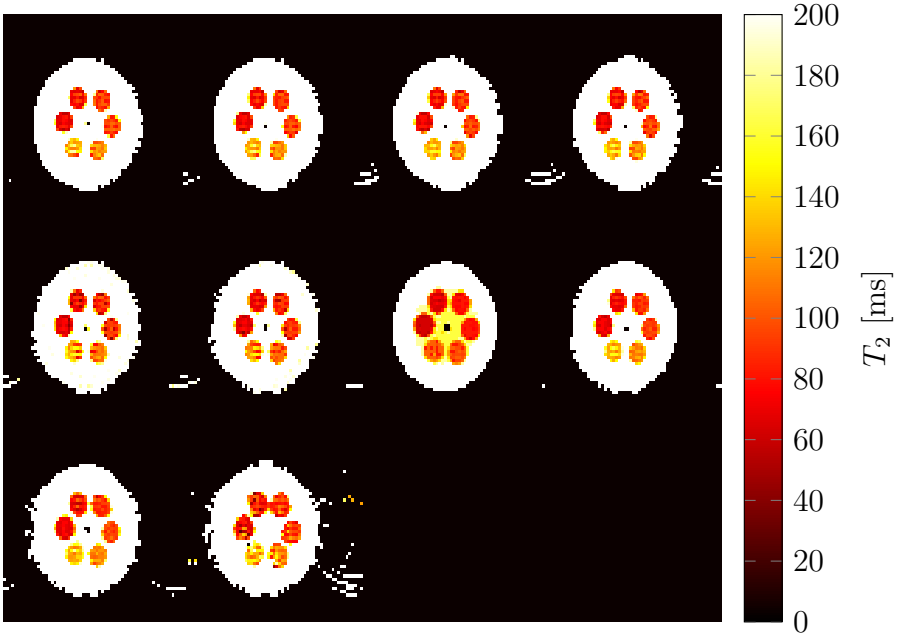

(b)

Figure 5.4: T_2 maps from genetic algorithm example initial (a) and optimised (b) 20-timepoint schedules with B_1^+ fitting.

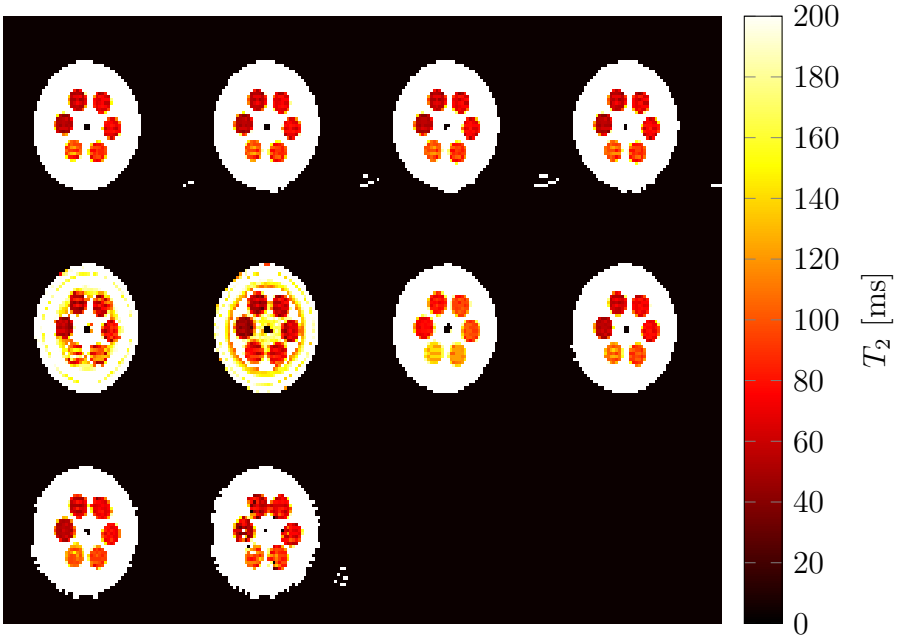

(a)

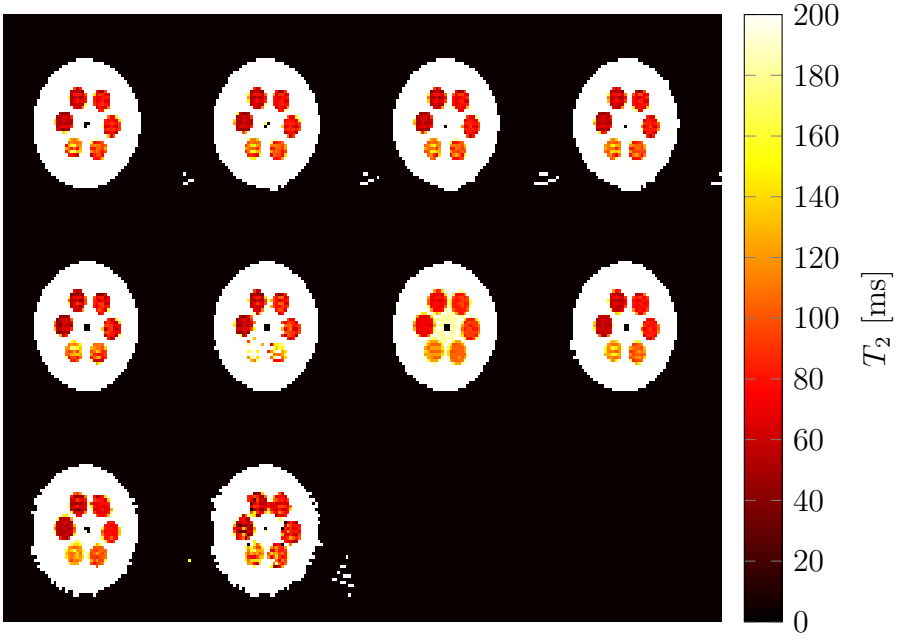

(b)

Figure 5.5: T_2 maps from *fmincon* algorithm initial (a) and optimised (b) 20-timepoint schedules with B_1^+ fitting.

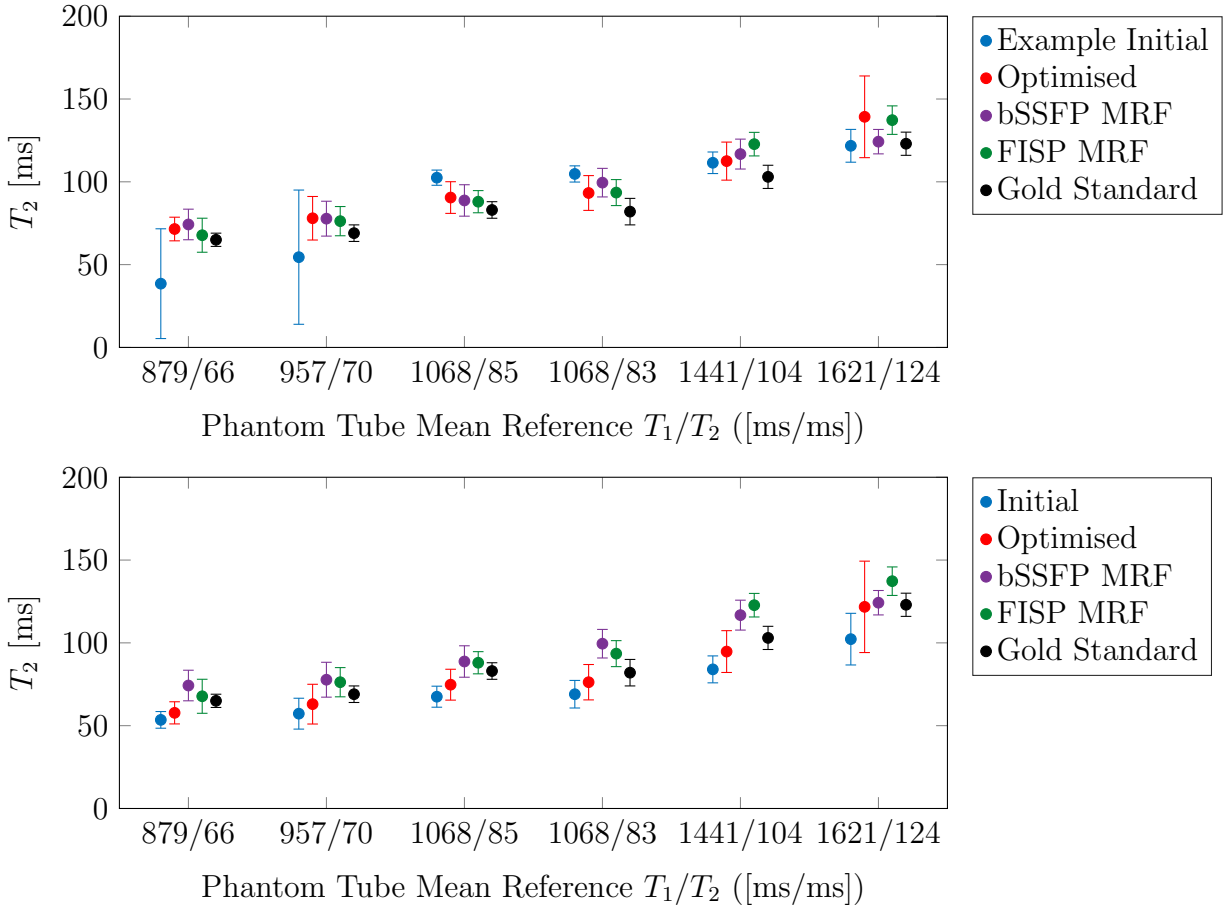

Figure 5.6: Comparison of the SE EPI MRF T_2 values for Slice 1 and reference maps, for the 20-timepoint genetic (top) and *fmincon* (bottom) schedules. Results are shown for each of the 6 phantom tubes.

Phantom T_2 Analysis Quantitative analysis of the T_2 distributions in Fig. 5.6 is provided by Tables 5.8, 5.9, 5.10, 5.11, 5.12 and 5.13.

Tables 5.8 and 5.9 show the bias of the T_2 measurements for the genetic and *fmincon* algorithms, respectively. The biases were calculated in the same way as those of the T_1 measurements. For both sets of data the optimised schedules successfully reduced the T_2 bias compared to the unoptimised schedules.

Reference [ms]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
65	-40.8	+10.0	+13.8	+4.6
69	-21.0	+13.0	+13.0	+10.1
83	-23.5	+9.0	+7.2	+6.0
82	+27.8	+13.8	+22.0	+14.6
103	+8.3	+9.2	+13.6	+19.4
123	-1.0	+13.3	+0.8	+11.4

Table 5.8: Genetic algorithm T_2 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean, using 9 measurements for the gold standard method. The smallest absolute bias for each row is highlighted in red.

Reference [ms]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
65	-17.7	-11.1	+13.8	+4.6
69	-17.0	-8.7	+13.0	+10.1
83	-18.7	-10.0	+7.2	+6.0
82	-15.9	-7.0	+22.0	+14.6
103	-18.4	-8.0	+13.6	+19.4
123	-16.8	-1.0	+0.8	+11.4

Table 5.9: *fmincon* T_2 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean, using 9 measurements for the gold standard method. The smallest absolute bias for each row is highlighted in red.

Tables 5.10 and 5.11 show the precision of the T_2 measurements for the genetic and *fmincon* algorithms, respectively, calculated in the same way as for the T_1 measurements. In the genetic algorithm results each of the sequences outperforms the others for at least one of the tubes. In the *fmincon* results FISP performs the best for most of the tubes. Overall, the optimised SE EPI schedules are outperformed by the initial schedules for most of the tubes.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
65±6.2	86.8	9.7	12.2	14.7
69±7.2	74.5	16.7	14.1	11.8
83±6.0	4.9	11.1	10.1	8.0
82±9.8	4.8	11.8	9.0	8.5
103±6.8	6.3	9.7	7.7	5.7
123±5.7	8.2	18.0	5.6	6.6

Table 5.10: Comparison of the genetic algorithm T_2 spreads with the gold standard (reference), bSSFP and FISP measurements. The spread was calculated as the percentage standard deviation. The smallest spread of the MRF sequences is highlighted in red. The reference column shows the means and percentage spreads of the gold standard measurements.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
65±6.2	9.3	12.1	12.2	14.7
69±7.2	15.8	19.0	14.1	11.8
83±6.0	9.0	12.0	10.1	8.0
82±9.8	11.6	14.5	9.0	8.5
103±6.8	9.5	13.7	7.7	5.7
123±5.7	15.7	23.0	5.6	6.6

Table 5.11: Comparison of the *fmincon* T_2 spreads with the gold standard (reference), bSSFP and FISP measurements. The spread was calculated as the percentage standard deviation. The smallest spread of the MRF sequences is highlighted in red. The reference column shows the means and percentage spreads of the gold standard measurements.

Tables 5.12 and 5.13 show the efficiency of T_2 measurements from the genetic and *fmincon* schedules, respectively. Overall the optimised schedules were less efficient than the initial schedules.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
65	1.0	0.4	3.6	2.3	1.9
69	0.8	0.5	2.1	2.0	2.4
83	1.0	7.3	3.2	2.8	3.6
82	0.6	7.4	3.0	3.2	3.3
103	0.9	5.7	3.6	3.7	5.0
123	1.1	4.3	2.0	5.0	4.3

Table 5.12: Genetic algorithm T_2 efficiency. The largest efficiency for each row is highlighted in red.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
65	1.0	3.8	2.9	2.3	1.9
69	0.8	2.2	1.9	2.0	2.4
83	1.0	3.9	2.9	2.8	3.6
82	0.6	3.0	2.4	3.2	3.3
102	0.9	3.7	2.6	3.7	5.0
123	1.1	2.3	1.5	5.0	4.3

Table 5.13: *fmincon* T_2 efficiency. The largest efficiency for each row is highlighted in red.

Overall, the preceding results show that the optimised schedules were more efficient than the corresponding initial schedules and the *fmincon* optimisation improved T_2 accuracy. However, the results regarding the other metrics were less successful, with the optimised schedules providing reduced accuracy and precision for many of the tubes in terms of T_1 and T_2.

The following sets of results contain maps from the *in vivo* scans, including the other parameters derived from dictionary matching: B_1^+, inversion efficiency and M_0.

In Vivo T_1 **Maps** Figures 5.7 and 5.8 show the *in vivo* T_1 maps from the schedules optimised via the genetic and *fmincon* algorithms, respectively. Unlike

the phantom maps the central *in vivo* slices have similar parameter maps to the inferior slices, with the exception of Slice 1. Slice 1 shows lower values compared to the other slices, for the genetic and *fmincon* schedules. Figure 5.9 shows the T_1 maps from the bSSFP and FISP MRF acquisitions. Figure 5.10 shows the gold standard reference map acquired during the same scanning session. Table 5.14 shows the T_1 values of the gold standard (reference), genetic, *fmincon*, bSSFP and FISP maps, obtained from manually drawn white and grey matter masks. The SE EPI MRF analysis was performed on Slice 5 of the 10 slices.

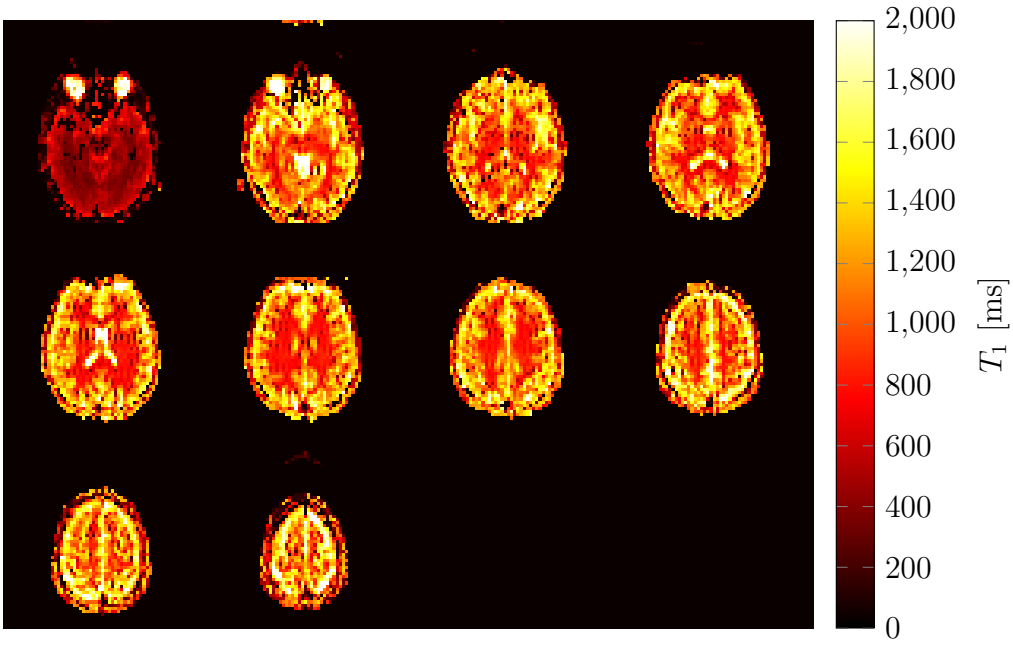

Figure 5.7: *In vivo* SE EPI MRF T_1 with 20 timepoints per slice, using the optimised genetic algorithm schedule.

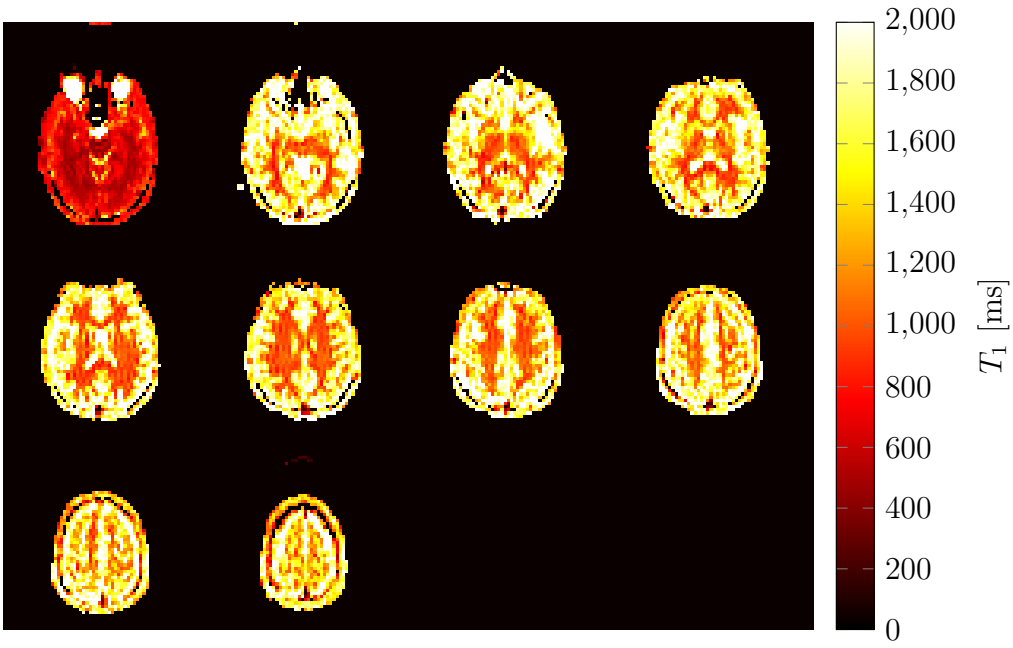

Figure 5.8: *In vivo* SE EPI MRF T_1 with 20 timepoints per slice, using the optimised *fmincon* schedule.

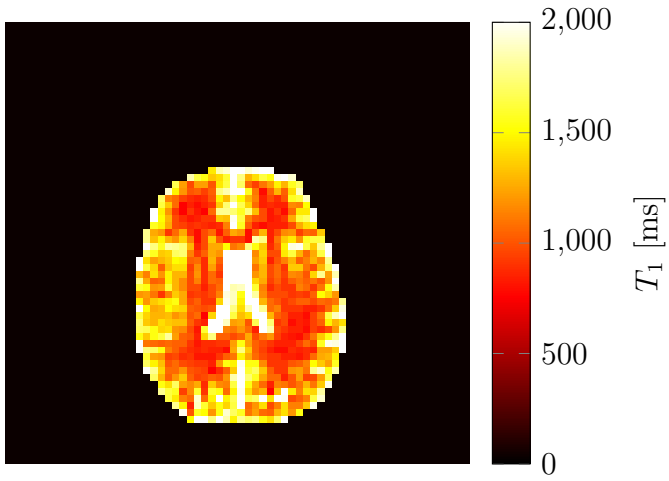

(a)

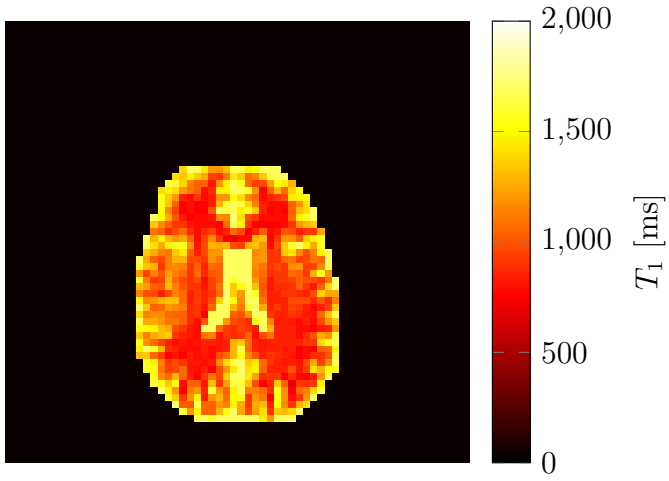

(b)

Figure 5.9: *In vivo* bSSFP (a) and FISP (b) MRF T_1 maps (masked) obtained via the acquisitions during the same scanning sesssion as the 20-timepoint SE EPI MRF acquisitions.

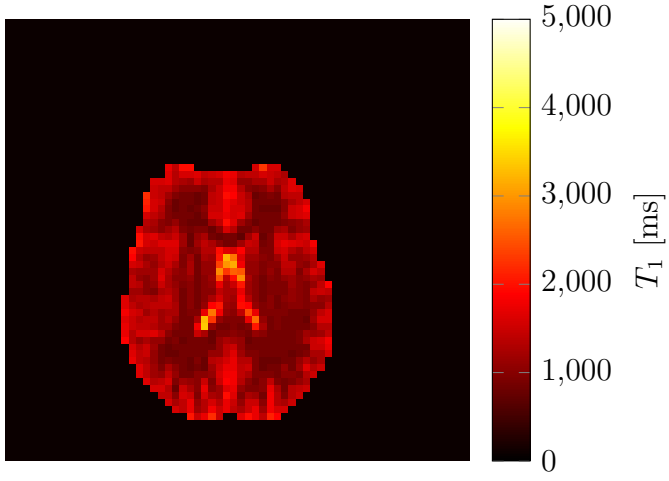

(a)

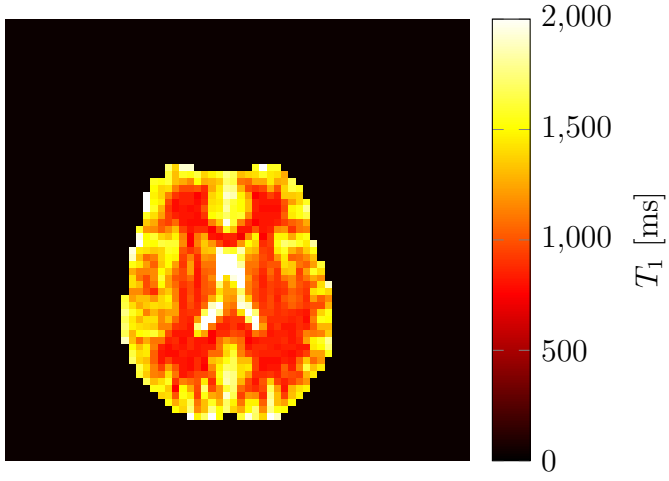

(b)

Figure 5.10: *In vivo* gold standard T_1 maps (masked) from the 20-timepoint spin-echo EPI MRF session. Two colour scales are shown: the range of T_1 used during the fitting (a) and the range of the bSSFP, FISP and spin-echo EPI MRF maps (b).

Reference [ms]	**Genetic [ms]**	***fmincon* [ms]**	**bSSFP [ms]**	**FISP [ms]**
926±180	793±215	1018±208	990±199	909±156
1359±249	1249±297	1660±320	1433±315	1312±239

(a)

Reference [ms]	**Genetic [%]**	***fmincon* [%]**	**bSSFP [%]**	**FISP [%]**
926	-14.4	+9.9	+6.9	**-1.8**
1359	-8.0	+22.1	+5.4	**+3.4**

(b)

Table 5.14: *In vivo* T_1 values from manually-drawn masks (a). The biases of each MRF measurement are also shown (b) as percentages of the gold standard reference mean (Column 1). The lowest bias for each row is shown in red.

Table 5.14 shows that all the measurements possess similar spreads to the gold standard but the spiral MRF measurement means are closest to the reference mean, as reflected by the bias calculations shown in the form of the difference from the reference mean as a percentage of the reference mean. Here the FISP results show the highest accuracy as they have the smallest bias. The optimised *fmincon* schedule is more accurate for the shorter T_1 (white matter) but the genetic algorithm schedule is more accurate for the grey matter.

***In Vivo* T_2 Maps** Figures 5.11 and 5.12 show the T_2 maps obtained from the 20-timepoint *in vivo* SE EPI MRF acquisitions optimised via the genetic and *fmincon* algorithms, respectively.

Figure 5.13 shows the bSSFP and FISP MRF T_2 maps acquired in the same session. A distinct band of exceptionally low T_2 is visible in the anterior region of the bSSFP T_2 map (Fig. 5.13a), reflecting the *banding* artefact commonly observed in bSSFP imaging [37].

Figure 5.14 shows the gold standard reference T_2 map acquired during the same scanning session, as well as the fitted map of the corresponding addition factor

c. The parameter c was introduced in the T_2 relaxation equation in Chapter 3 and can be affected by the measurement noise and the signal magnitude. In Fig. 5.14b relatively large values of c are present in the ventricles and the periphery of the brain. This could be due to exceptionally long T_2 caused by high water content, which would lead to high signal even at the longest TE.

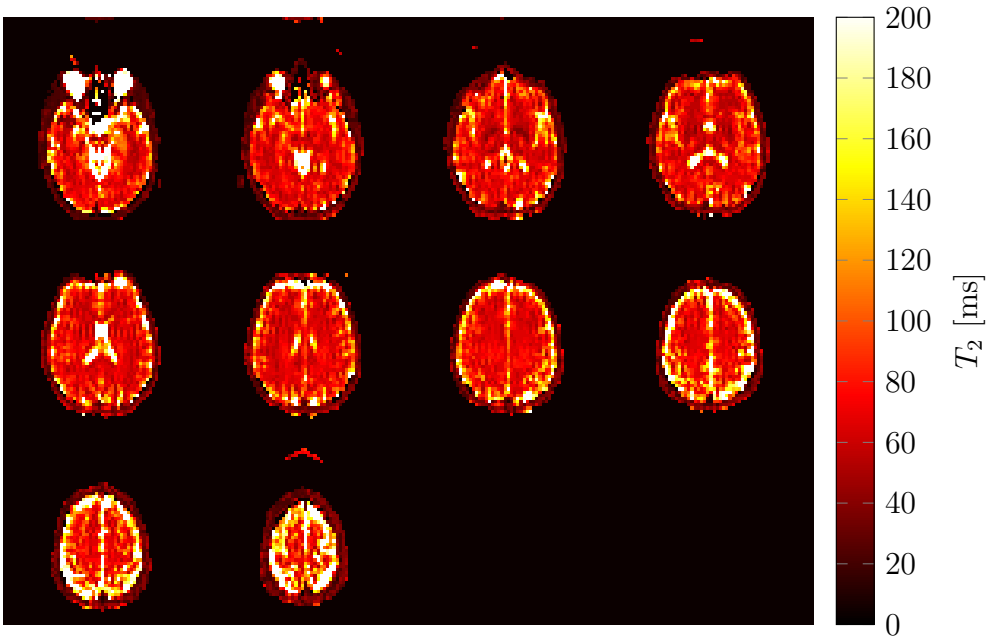

Figure 5.11: *In vivo* SE EPI MRF T_2 map from the 20-timepoint schedule optimised via the genetic algorithm.

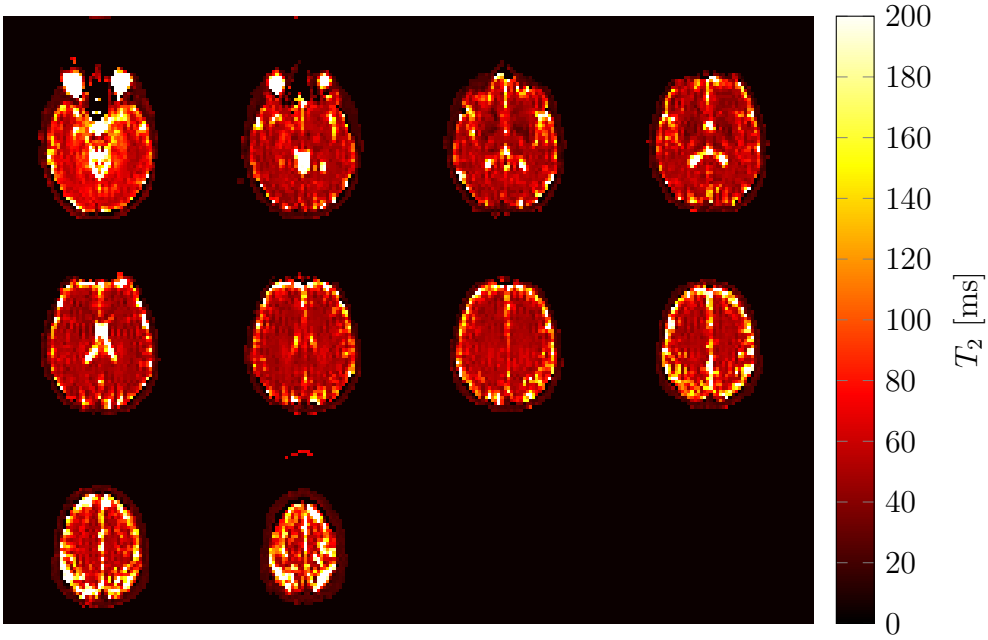

Figure 5.12: *In vivo* SE EPI MRF T_2 map from the 20-timepoint schedule optimised via the *fmincon* algorithm.

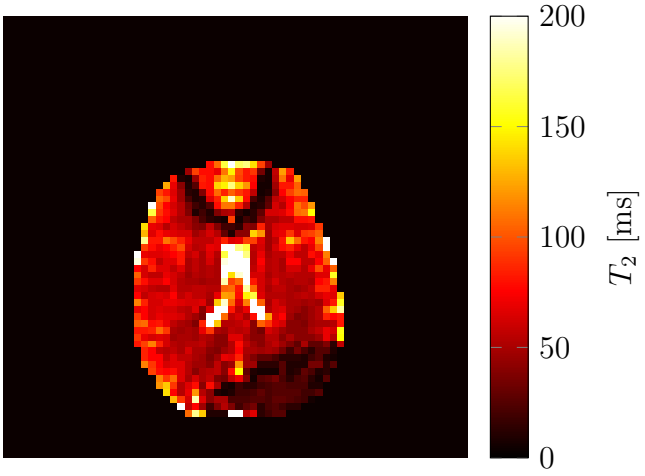

(a)

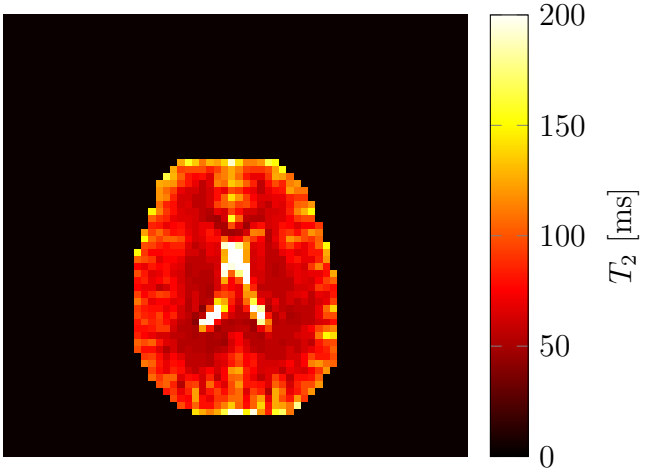

(b)

Figure 5.13: *In vivo* bSSFP (a) and FISP (b) T_2 maps (masked) from the 20-timepoint spin-echo EPI MRF session.

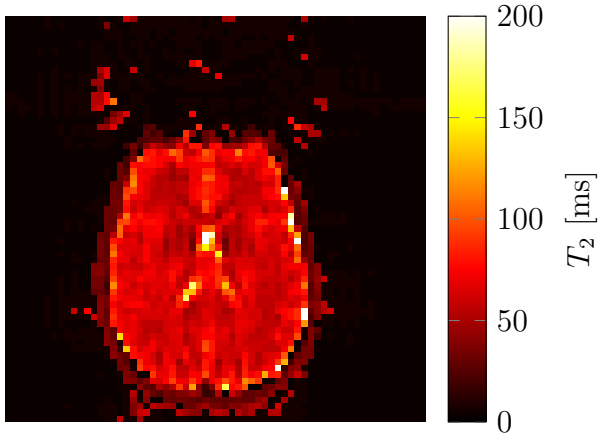

(a)

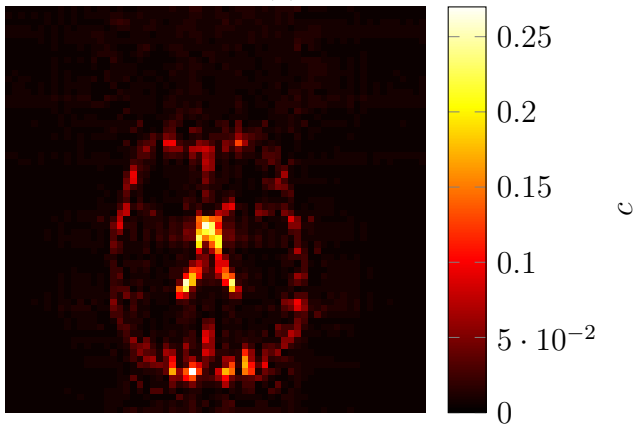

(b)

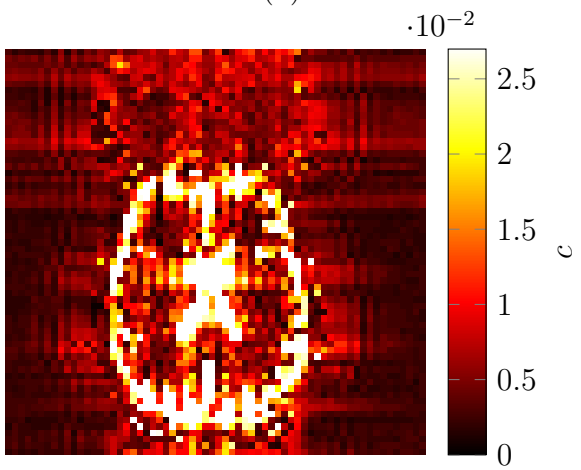

(c)

Figure 5.14: *In vivo* gold standard T_2 (a) and addition parameter c (b) from the 20-timepoint spin-echo EPI MRF session. The c map is also shown with a lower colour range (c).

Table 5.15 shows each of the *in vivo* MRF T_2 measurements compared to the gold standard (reference) measurements in the left-hand column. The spreads of each MRF measurement are larger than those of the gold standard reference. As shown by the T_2 accuracies in the same Table the bSSFP and FISP scans were the most accurate, although the FISP acquisition showed low accuracy for grey matter (reference T_2 mean = 68ms).

Reference [ms]	**Genetic [ms]**	***fmincon* [ms]**	**bSSFP [ms]**	**FISP [ms]**
58±9	68±24	51±20	48±20	63±16
68±11	88±35	63±32	68±41	94±25

(a)

Reference [ms]	**Genetic [%]**	***fmincon* [%]**	**bSSFP [%]**	**FISP [%]**
58	+17.2	-12.1	-17.2	**+8.6**
68	+29.4	-7.4	**0.0**	+38.2

(b)

Table 5.15: *In vivo* T_2 values (a) from manually-drawn masks of white matter (top row) and grey matter (bottom row). The biases are also shown (b), in the form of the difference between the means of the MRF and corresponding gold standard (reference) measurement, as a percentage of the gold standard mean (Column 1). The smallest absolute bias for each row is in red.

In Vivo M_0 **Maps** Figure 5.15 shows maps of scaling factors obtained from the gold standard T_1 and T_2 fits, which reflect the M_0 spatial distribution. There is contrast between the M_0 of the white matter and the ventricles due to differing water concentrations. Figures 5.16 and 5.17 show the normalised M_0 maps obtained from the optimised genetic and *fmincon* acquisitions, respectively. As seen in the gold standard map, a boundary between the ventricles and the white matter is visible. Figure 5.18 shows the bSSFP and FISP MRF M_0 maps from the same scanning session. Evidence of the banding artefact observed in the bSSFP T_2 map is present in the M_0 map (Fig. 5.18a).

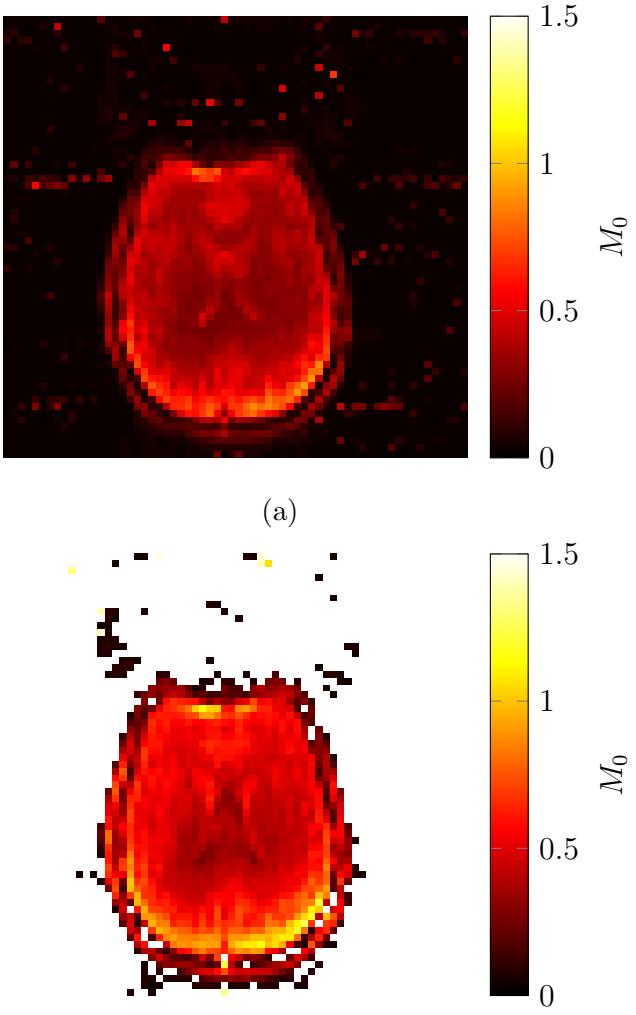

Figure 5.15: *In vivo* M_0 maps obtained from the gold standard T_1 (a) and T_2 (b) fits.

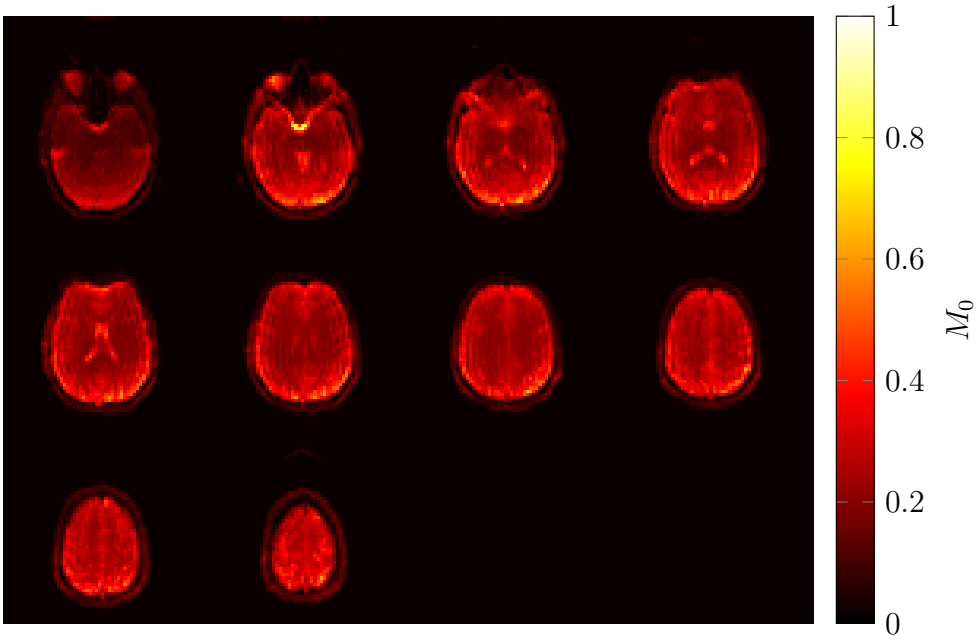

Figure 5.16: *In vivo* SE EPI MRF: 20 timepoints per slice, showing the derived M_0 from the optimised genetic algorithm schedule.

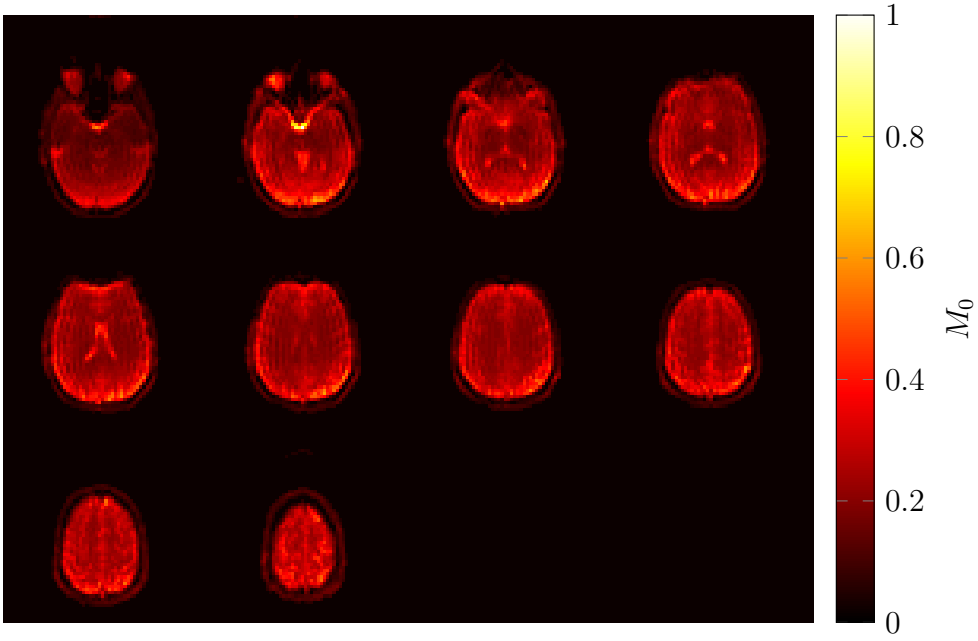

Figure 5.17: *In vivo* SE EPI MRF: 20 timepoints per slice, showing the derived M_0 from the optimised *fmincon* schedule.

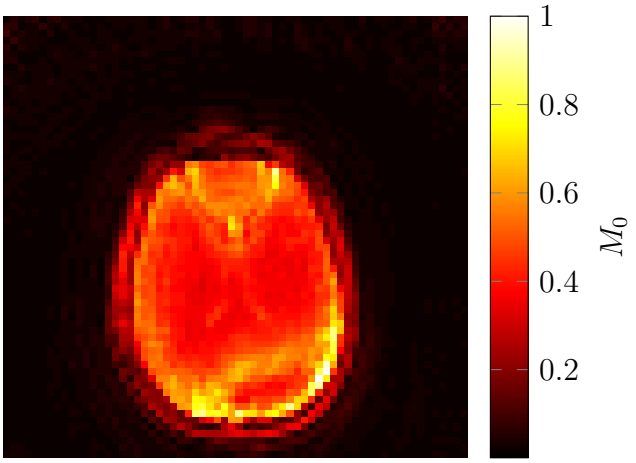

(a)

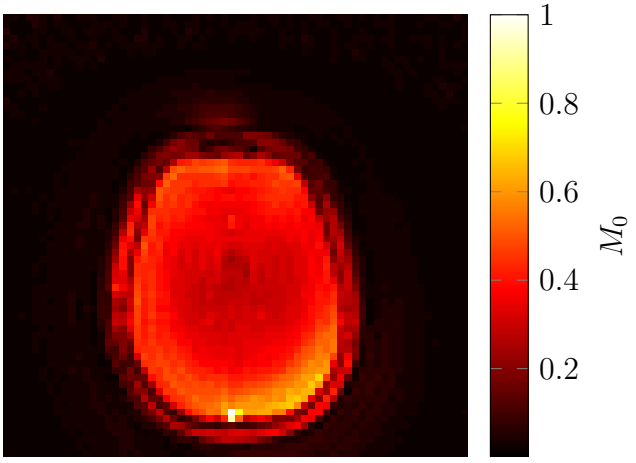

(b)

Figure 5.18: *In vivo* M_0 maps obtained from the bSSFP (a) and FISP (b) spiral MRF acquisitions from the same session as the 20-timepoint SE EPI MRF scans.

***In Vivo* Inversion Efficiency Maps** Figures 5.19 and 5.20 show the fitted inversion efficiency maps from the optimised genetic and *fmincon* 20-timepoint schedules, respectively. The maps show that complete inversions were fitted for most of each slice of the genetic algorithm data. Marginally lower inversion efficiency scores were fitted to the *fmincon* data. Although the phantom data showed large inter-slice differences in the fitted inversion efficiency, the *in vivo* maps shown

here are relatively consistent across the slices. This is further evidence that the phantom itself has an impact on the inversion efficiency, as explored in Appendix B.

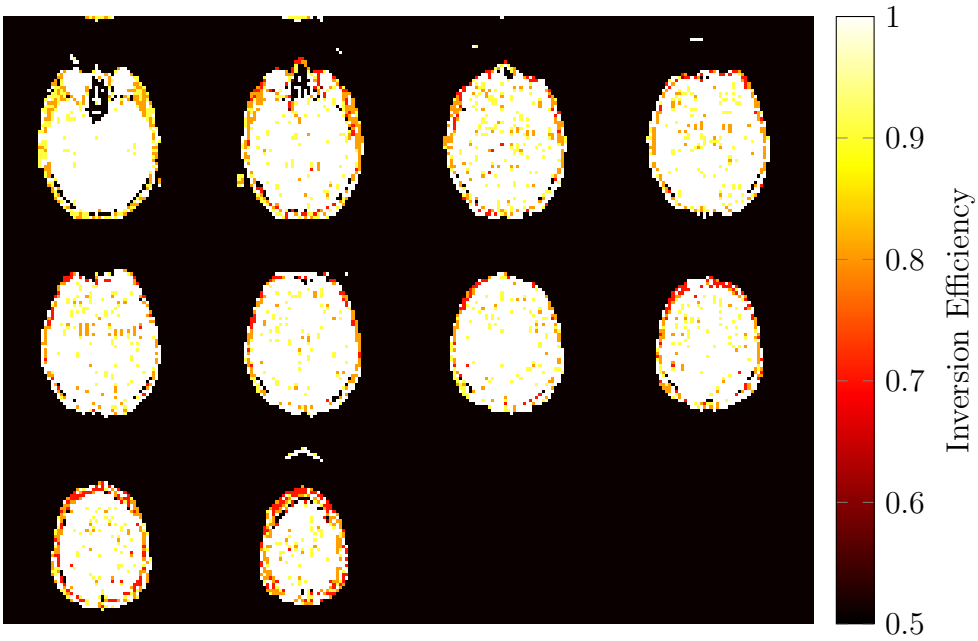

Figure 5.19: *In vivo* SE EPI MRF: 20 timepoints per slice, showing the inversion efficiency from the optimised genetic algorithm schedule.

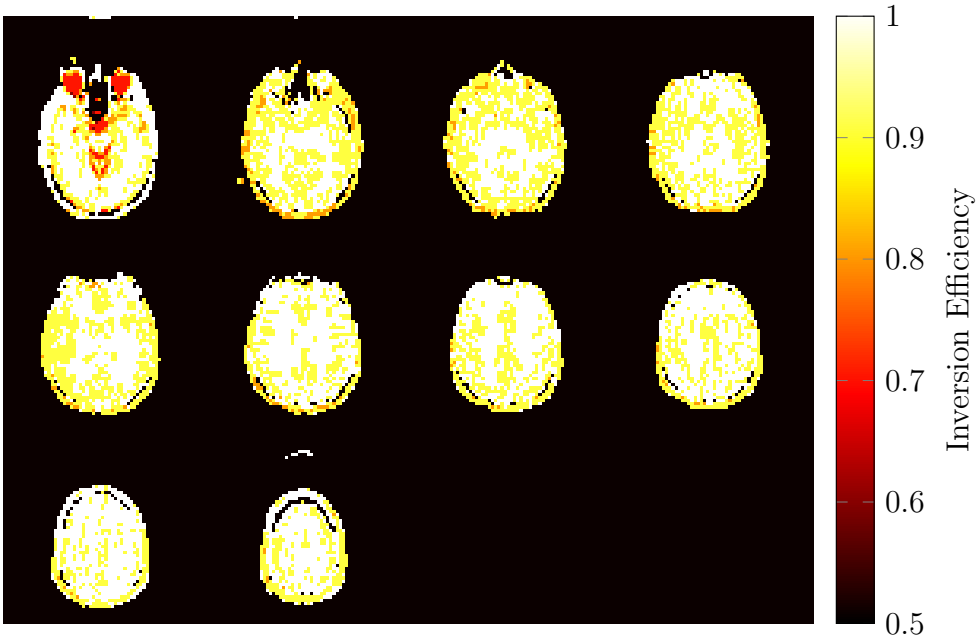

Figure 5.20: *In vivo* SE EPI MRF: 20 timepoints per slice, showing the inversion efficiency from the optimised *fmincon* schedule.

In vivo B_1^+ **Maps** Figures 5.21 and 5.22 show the B_1^+ efficiency obtained from the *in vivo* 20-timepoint SE EPI MRF schedules optimised via the genetic and *fmincon* algorithms, respectively. Figure 5.23 shows the 3D DREAM B_1^+ efficiency maps acquired during the same scanning session. The SE EPI MRF B_1^+ maps do not resemble the 3D DREAM maps, showing that the framework was not able to accurately derive B_1^+ when this was included in the dictionary.

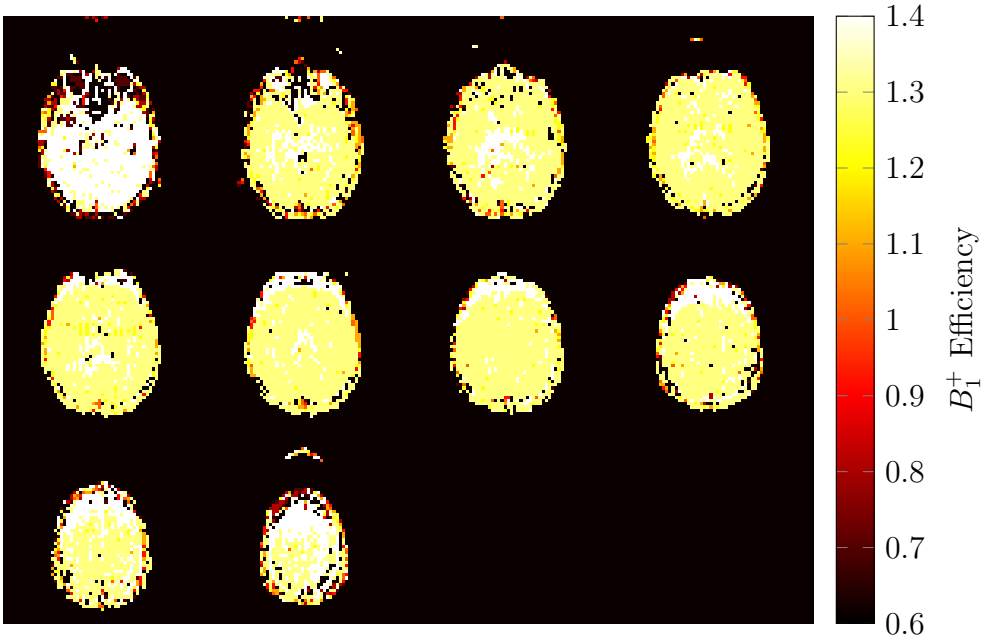

Figure 5.21: *In vivo* SE EPI MRF: 20 timepoints per slice, showing the B_1^+ efficiency from the optimised genetic algorithm schedule.

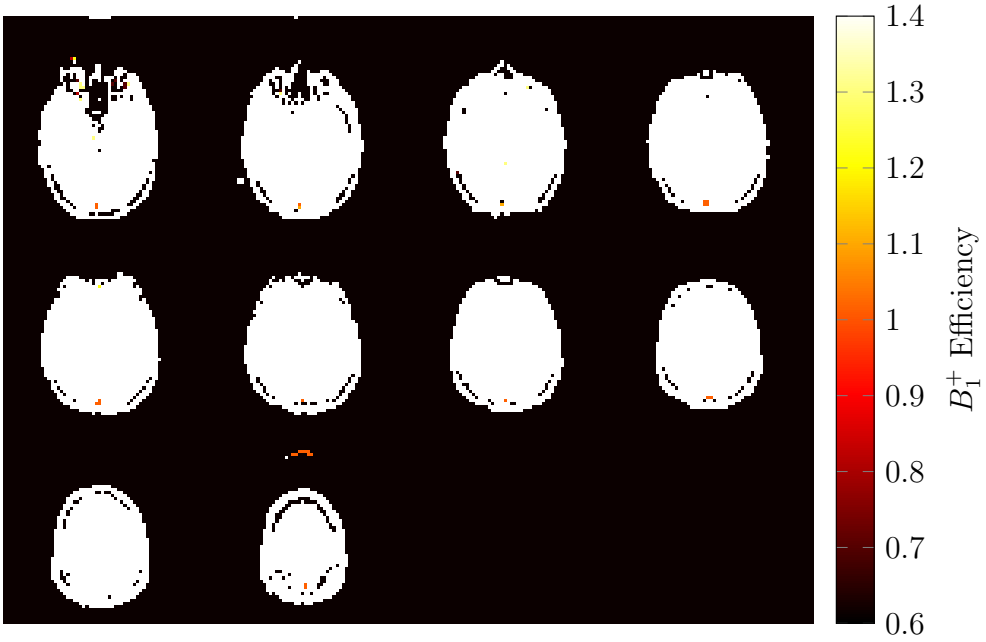

Figure 5.22: *In vivo* SE EPI MRF: 20 timepoints per slice, showing the B_1^+ efficiency: *fmincon* schedule.

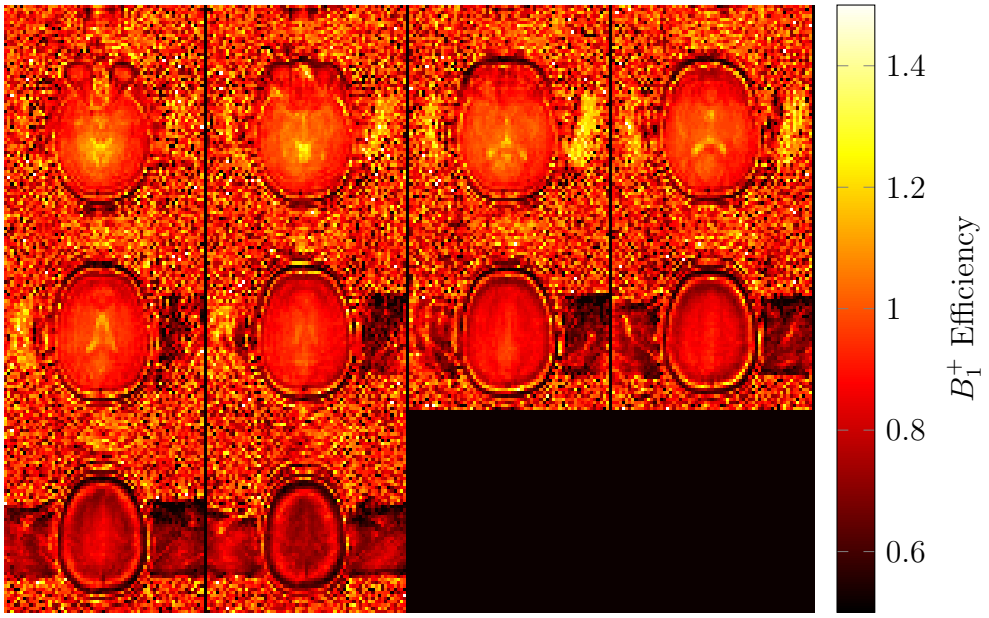

Figure 5.23: *In vivo* 3D DREAM B_1^+ efficiency maps, showing the central 10 slices, corresponding to the 10 SE EPI MRF slices.

In Vivo B_0 **Maps** Figure 5.24 shows the off-resonance map from the bSSFP MRF acquisition from the same session as the 20-timepoint SE EPI MRF *in vivo* acquisitions. As a validation the gold standard B_0 maps are shown in Fig. 5.25. Slice 5 was the closest gold standard slice to the bSSFP slice position. Overall the map at Slice 5 appears to show smaller off-resonance values than the bSSFP map. The gold standard and bSSFP maps both show a relatively large off-resonance change in the same anterior region but the bSSFP off-resonance in that region appears to be 'wrapped-around' from positive to negative frequency offsets.

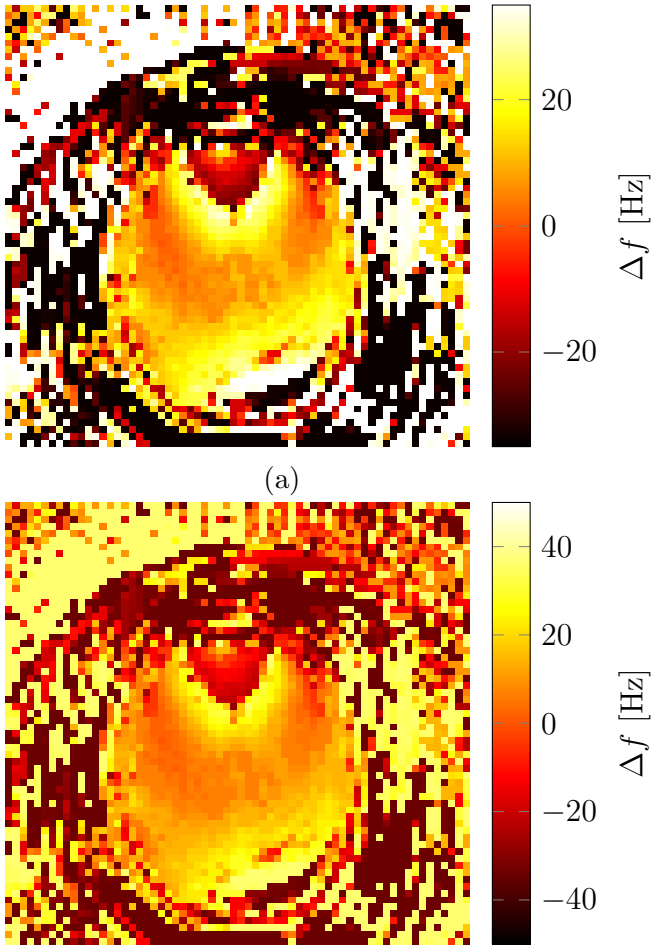

Figure 5.24: *In vivo* off-resonance maps obtained from the bSSFP spiral MRF acquisition, with colour scales matching the range of the off-resonance dictionary dimension (a) and the range used in the gold standard maps (b) shown in Fig. 5.25.

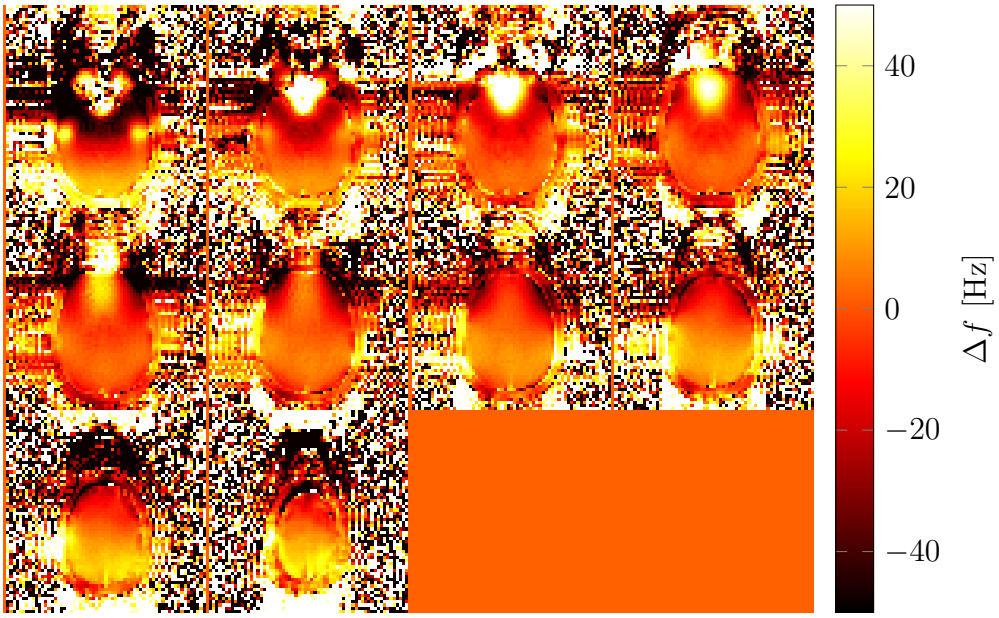

Figure 5.25: *In vivo* B_0 maps, showing the central 10 slices, corresponding to the 10 SE EPI MRF slices.

5.4.4 20 Timepoints (8 Seconds) per Slice: Conclusions

In this section the T_1 and T_2 maps from 20-timepoint SE EPI MRF schedules were quantitatively assessed and compared with measurements from the bSSFP and FISP spiral MRF implementations and the gold standard methods. In the interest of completeness, *in vivo* maps of the fitted MRF M_0, inversion efficiency and B_1^+ were also provided.

Phantom results showed that the optimised schedules were generally more efficient than the unoptimised schedules but in general there was poor agreement with the simulated validations in Chapter 4.

The phantom T_1 accuracy of the optimised genetic algorithm schedule was worse than the unoptimised schedule for most of the phantom tubes, including for the T_1 range targeted during the optimisations. Similar trends were observed

for the *fmincon* schedules, although for the *fmincon* schedules the bias improved from -13% to +0.8% for the targeted T_1. The optimisations provided marginal improvements in the spread for the target T_1, reducing it by 0.5% and 0.9% for the genetic and *fmincon* schedules, respectively. However, the precision of many of the other tubes worsened. The FISP sequence was the most precise for most of the T_1 species, with spreads ranging from 2.8% to 8.9% across all 6 tubes. These trends were reflected in the scan efficiency metric.

In terms of the phantom T_2, the optimised genetic schedule worsened the accuracy (-1.0% to +13.3%) for the target T_2 but increased the accuracy for the 4 tubes with the shortest T_2 (65ms, 69ms, 83ms and 82ms). The *fmincon* schedule optimisation successfully reduced the bias for the target T_2 from -16.8% to -1.0% and also improved the accuracy for the other 5 tubes. The genetic optimisation improved the spread for the T_2 = 65ms (86.8% to 9.7%) and 69ms (74.5% to 16.7%) but increased the spread for the 4 tubes with longer T_2. The *fmincon* optimisation increased the T_2 spread for all 6 tubes. As with T_1, the T_2 efficiency results reflected the T_2 precisions.

In the *in vivo* T_1 results the largest optimised genetic schedule bias was -14.4% and the largest *fmincon* bias was +22.1%. The spiral MRF sequences were more accurate, as their largest *in vivo* T_1 bias was +6.9% (bSSFP). The largest T_2 bias from the genetic schedule was +29.4% but for the *fmincon* schedule it was -12.1%. Interestingly, the FISP sequence had +8.6% error for white matter but a much larger bias (+38.2%) for grey matter. The largest bSSFP bias was -17.2%. The SE EPI MRF framework was not able to derive accurate B_1^+ maps from the 20-timepoint data.

Overall the results show that the optimised 20-timepoint SE EPI MRF sched-

ules with B_1^+ did not provide sufficient improvement over the unoptimised schedules. Future work could explore whether this could be achieved by using a separately acquired 3D DREAM B_1^+ map to produce a tailored dictionary for each voxel that incorporates the correct B_1^+ values. Additionally, using schedules with a larger number of timepoints could reduce the measurement error.

5.5 50 Timepoints (20 Seconds) per Slice

This section examines the performance of the 50-timepoint schedules optimised in Chapter 4.

5.5.1 50 Timepoints (20 Seconds) per Slice: Introduction and Methods

The experimental results for the optimised 20-timepoint SE EPI MRF schedules showed some improvements in efficiency (Section 5.4) but overall there was not sufficient agreement between the predictions of Chapter 4 and the experimental data. This section explores whether the optimised 50-timepoint schedules from Chapter 4 perform better. Increasing the number of timepoint measurements from 20 to 50 may enable the framework to overcome noise in the signal fingerprints. The methods used in this section were the same as those of Section 5.4 but with the 50-timepoint schedules from Chapter 4 instead of the 20-timepoint schedules.

5.5.2 50 Timepoints (20 Seconds) per Slice: Results and Discussion

Phantom T_1 Maps Figures 5.26 and 5.27 show the T_1 maps from the genetic and *fmincon* optimisations, respectively, fitted with a B_1^+ dimension included in the dictionary. Slice 7 is erroneous and Slices 9 and 10 show differences that can be attributed to the phantom structure, as explored in Appendix B. Analysis of the distributions in Slice 1 of maps is provided in Tables 5.16, 5.17, 5.18, 5.19, 5.20 and 5.21.

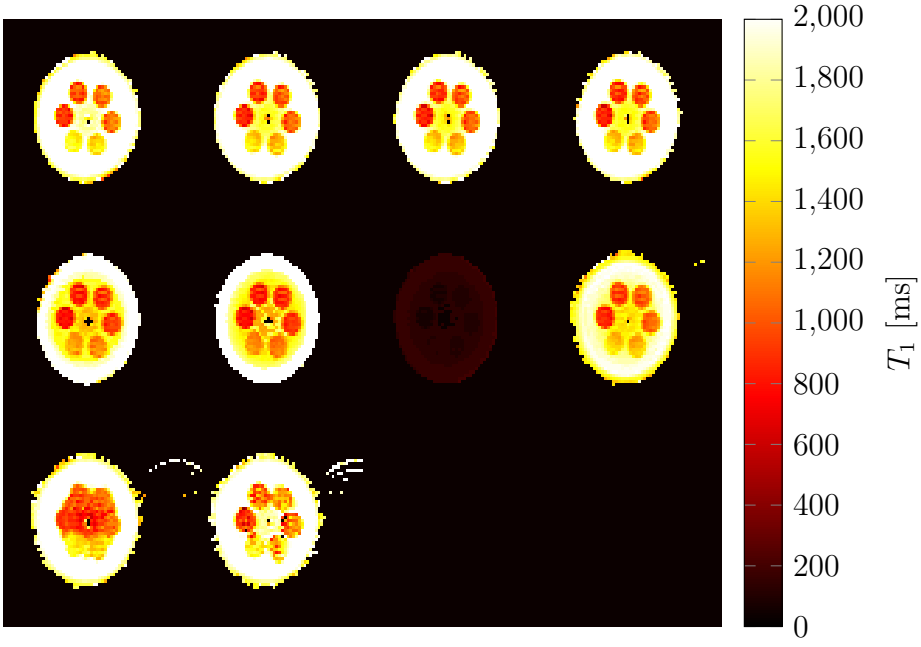

(a)

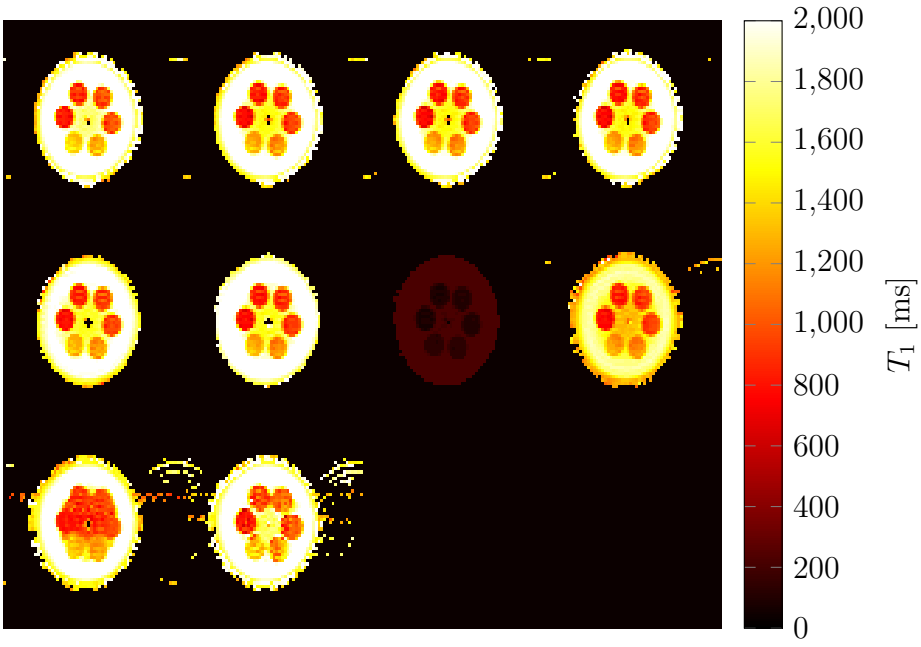

(b)

Figure 5.26: T_1 maps from the genetic algorithm example initial (a) and optimised (b) 50-timepoint schedules with B_1^+ fitting.

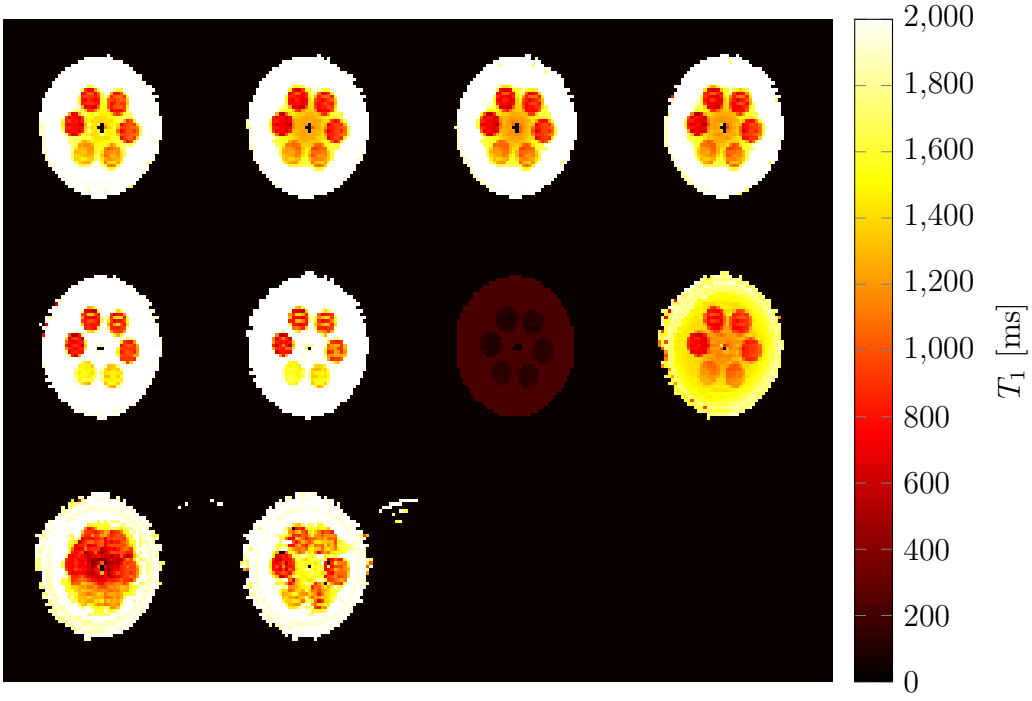

(a)

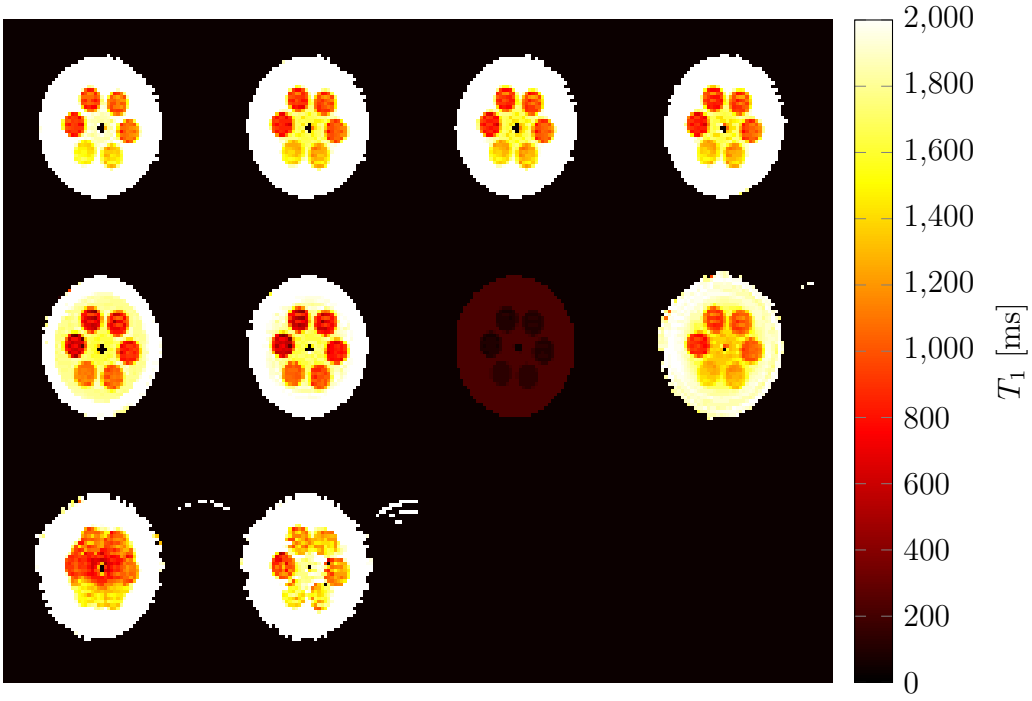

(b)

Figure 5.27: T_1 maps from the *fmincon* initial (a) and optimised (b) 50-timepoint schedules with B_1^+ fitting.

Tables 5.16 and 5.17 show the accuracy of the genetic and *fmincon* T_1 distributions, respectively. For all 6 tubes the optimised genetic schedule was less accurate than the initial schedule, whereas the *fmincon* optimisation successfully improved the accuracy. The *fmincon* schedule outperformed the bSSFP and FISP schedules in 4 of the 6 tubes.

Reference [ms]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
879	-1.0	-10.2	-5.0	-13.2
957	-1.0	-10.7	-2.9	-11.7
1068	+2.8	-9.2	-5.0	-11.0
1068	+3.7	-8.6	-3.3	-9.8
1441	-3.3	-12.9	-2.2	-6.9
1621	-7.1	-15.6	-2.5	-5.4

Table 5.16: Genetic algorithm T_1 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean. The smallest absolute bias for each row is in red.

Reference [ms]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
879	-15.7	-1.4	-5.0	-13.2
957	-16.1	-2.3	-2.9	-11.7
1068	-12.6	0.0	-5.0	-11.0
1068	-11.2	+0.8	-3.3	-9.8
1441	-18.5	-6.0	-2.2	-6.9
1621	-21.5	-9.2	-2.5	-5.4

Table 5.17: *fmincon* T_1 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean.

Tables 5.18 and 5.19 show the measurement precisions, in the form of the spread of the distributions. Both optimised schedules improved the T_1 precision for the tube with T_1 closest to the T_1 targeted in the optimisation (1600ms). Precision was also increased for most of the other tubes. Tables 5.20 and 5.21

show the calculated efficiency of the genetic and *fmincon* acquisitions, respectively. The optimisations improved efficiency for all but one of the tubes (mean T_1 = 1441ms). However, the efficiency decrease for this exception was small (3.2 to 3.1). Overall the FISP sequence provided the best performance in terms of precision and efficiency.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
879±3.9	9.3	6.3	11.5	8.9
957±3.1	12.9	10.1	7.2	5.7
1068±3.5	8.1	7.2	5.8	3.6
1068±3.2	5.8	5.1	6.5	4.8
1441±1.7	7.0	7.3	3.8	2.8
1621±1.4	6.1	5.6	3.4	2.8

Table 5.18: Genetic algorithm T_1 spread. The smallest spread of each row (other than that of the reference value) is in red.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
879±3.9	10.4	8.8	11.5	8.9
957±3.1	14.7	13.5	7.2	5.7
1068±3.5	9.7	8.6	5.8	3.6
1068±3.2	6.7	5.9	6.5	4.8
1441±1.7	8.8	8.1	3.8	2.8
1621±1.4	7.0	6.6	3.4	2.8

Table 5.19: *fmincon* T_1 spread. The smallest spread of each row (other than that of the reference value) is in red.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
879	1.4	2.4	3.5	2.5	3.2
957	1.7	1.7	2.2	3.9	5.0
1068	1.5	2.8	3.1	4.9	7.9
1068	1.7	3.9	4.4	4.4	5.9
1441	3.0	3.2	3.1	7.4	10.0
1621	3.7	3.7	4.0	8.5	10.1

Table 5.20: Genetic algorithm T_1 efficiency. The largest efficiency for each row is in red. The gold standard (reference) acquisition TR was 30s.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
879	1.4	2.2	2.6	2.5	3.2
957	1.7	1.5	1.7	3.9	5.0
1068	1.5	2.3	2.6	4.9	7.9
1068	1.7	3.3	3.8	4.4	5.9
1441	3.0	2.6	2.8	7.4	10.0
1621	3.7	3.2	3.4	8.5	10.1

Table 5.21: *fmincon* T_1 efficiency. The largest efficiency for each row is in red. The gold standard (reference) acquisition TR was 30s.

Phantom T_2 Maps Figures 5.28 and 5.29 show the T_2 maps from the genetic and *fmincon* optimisations, respectively. Quantitative comparisons of the measurements for Slice 1 are shown in Tables 5.22, 5.23, 5.24, 5.25, 5.26 and 5.27.

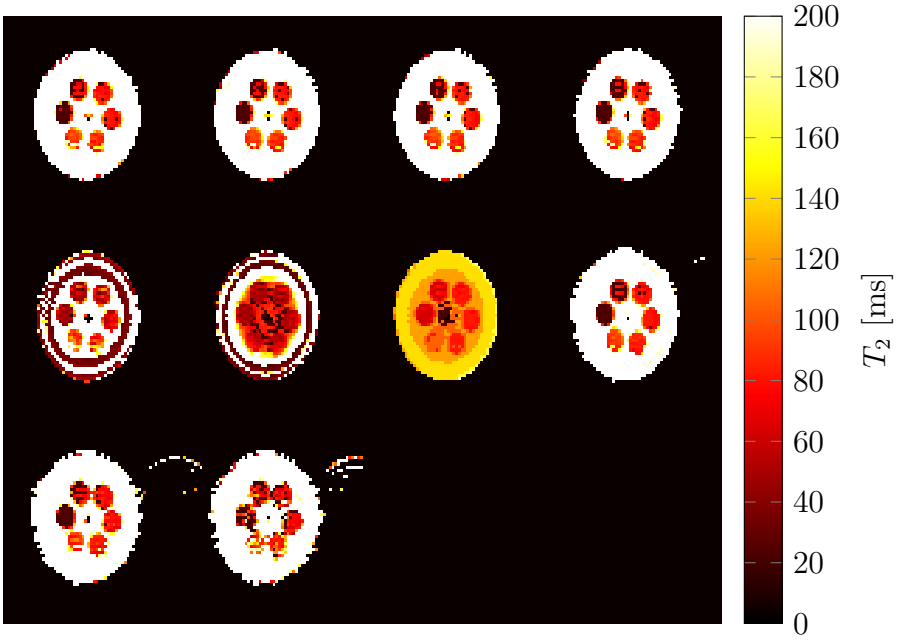

(a)

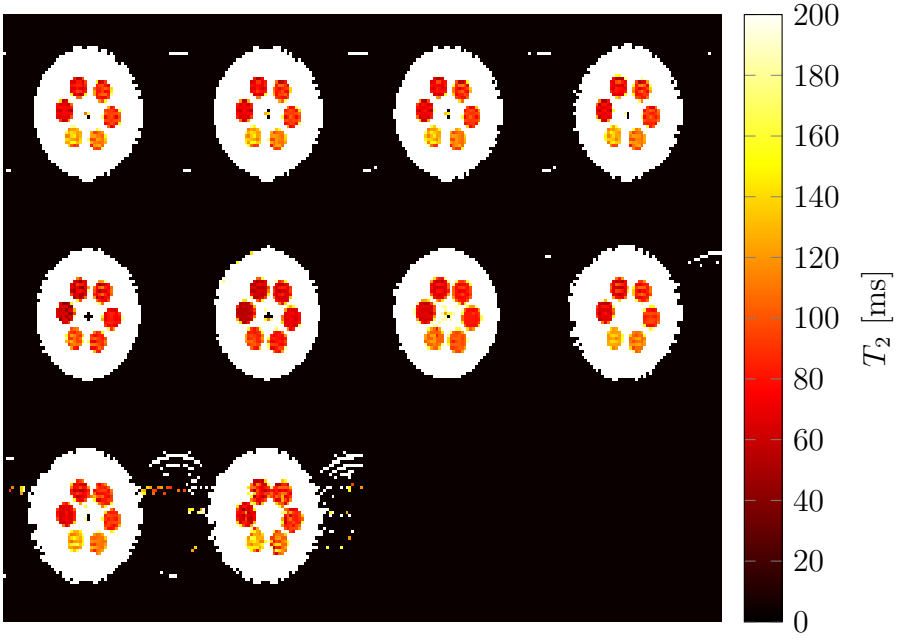

(b)

Figure 5.28: T_2 maps from the genetic algorithm example initial (a) and optimised (b) 50-timepoint schedules with B_1^+ fitting.

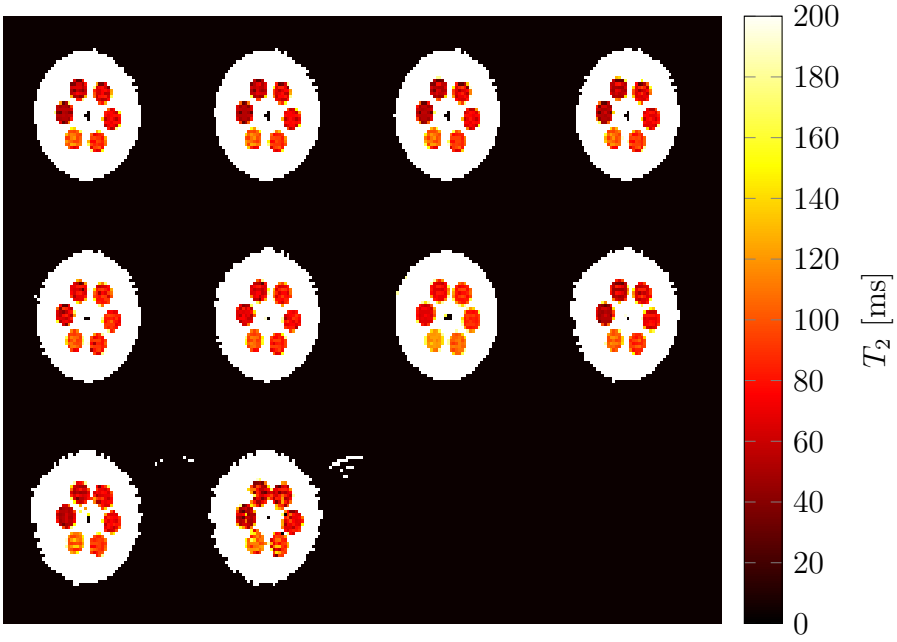

(a)

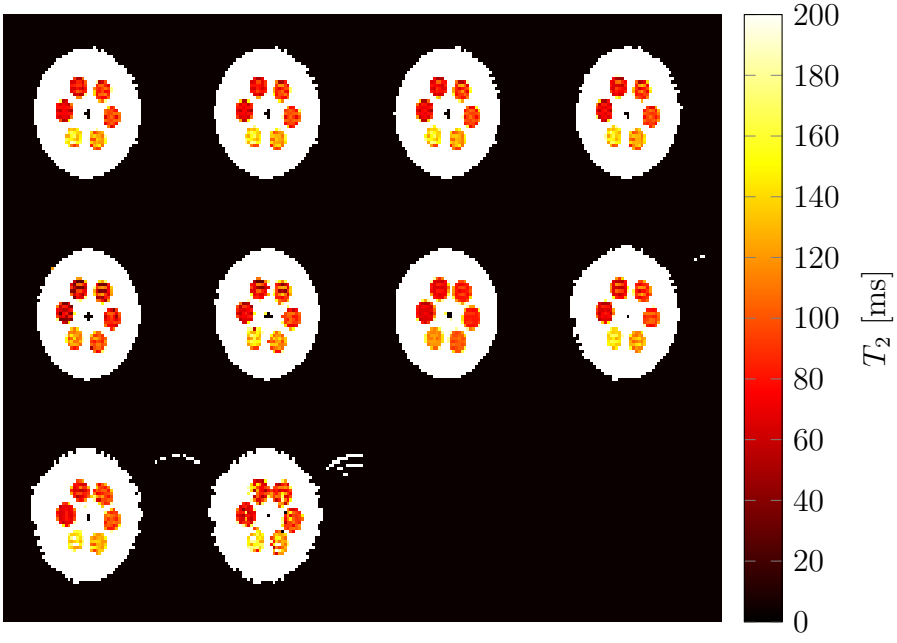

(b)

Figure 5.29: T_2 maps from the *fmincon* initial (a) and optimised (b) 50-timepoint schedules with B_1^+ fitting.

Phantom T_2 Analysis Tables 5.22 and 5.23 show the biases of the genetic and *fmincon* schedules, respectively. The genetic optimisations improved the accuracy for all 6 tubes of the phantom, outperforming the bSSFP and FISP schedules for all but the longest T_2. The *fmincon*-optimised schedule was less successful but improved the accuracy for most of the tubes. However, this schedule did not improve the accuracy for the T_2 closest to the T_2 targeted during the optimisation (i.e. 120ms).

Reference [ms]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
65	-64.0	+2.8	+13.8	+4.6
69	-31.9	+4.1	+13.0	+10.1
83	-8.7	+6.0	+7.2	+6.0
82	-9.8	+8.5	+22.0	+14.6
103	-8.5	+6.8	+13.6	+19.4
123	-9.4	+3.3	+0.8	+11.4

Table 5.22: Genetic Algorithm T_2 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean. The smallest absolute bias for each row is highlighted in red.

Reference [ms]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
65	-24.0	+3.4	+13.8	+4.6
69	-23.5	+5.8	+13.0	+10.1
83	-18.7	+10.8	+7.2	+6.0
82	-15.6	+14.0	+22.0	+14.6
103	-16.1	+16.3	+13.6	+19.4
123	-17.1	+18.0	+0.8	+11.4

Table 5.23: *fmincon* T_2 bias: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean. The smallest absolute bias for each row is highlighted in red.

The T_2 spreads in Tables and 5.24 and 5.25 show that the genetic optimisation

increased the precision for the target T_2, but the *fmincon* schedule did not. The efficiency results in Tables 5.26 and 5.27 show a similar trend.

Reference [ms±%]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
65±6.2	47.8	9.0	12.2	14.7
69±7.2	57.4	15.3	14.1	11.8
83±6.0	6.6	11.4	10.1	8.0
82±9.8	20.3	9.0	9.0	8.5
103±6.8	28.7	11.8	7.7	5.7
123±5.7	37.8	10.2	5.6	6.6

Table 5.24: Genetic algorithm T_2 spread. The smallest spread for each row, other than that of the reference, is highlighted in red.

Reference [ms±%]	**Initial [%]**	**Optimised [%]**	**bSSFP [%]**	**FISP [%]**
65±6.2	14.3	11.9	12.2	14.7
69±7.2	22.6	19.2	14.1	11.8
83±6.0	14.9	13.0	10.1	8.0
82±9.8	10.1	11.7	9.0	8.5
103±6.8	12.8	15.0	7.7	5.7
123±5.7	9.8	13.1	5.6	6.6

Table 5.25: *fmincon* T_2 spread. The smallest spread for each row, other than that of the reference, is highlighted in red.

Reference [ms]	**Reference Efficiency**	**Initial**	**Optimised**	**bSSFP**	**FISP**
65	1.0	0.5	2.5	2.3	1.9
69	0.8	0.4	1.5	2.0	2.4
83	1.0	3.4	2.0	2.8	3.6
82	0.6	1.1	2.5	3.2	3.3
103	0.9	0.8	1.9	3.7	5.0
123	1.1	0.6	2.2	5.0	4.3

Table 5.26: Genetic algorithm T_2 efficiency. The largest efficiency for each row is highlighted in red.

Reference [ms]	Reference Efficiency	Initial	Optimised	bSSFP	FISP
65	1.0	1.6	1.9	2.3	1.9
69	0.8	1.0	1.2	2.0	2.4
83	1.0	1.5	1.7	2.8	3.6
82	0.6	2.2	1.9	3.2	3.3
103	0.9	1.7	1.5	3.7	5.0
123	1.1	2.3	1.7	5.0	4.3

Table 5.27: *fmincon* T_2 efficiency. The largest efficiency for each row is highlighted in red.

In Vivo T_1 **Maps** Figure 5.30 shows the gold standard map for a single slice during the same scanning session as the 50-timepoint SE EPI MRF acquisitions. The map was acquired along a plane with a small angle difference relative to the following spiral MRF and SE EPI MRF maps but remains a useful reference because white and grey matter masks were drawn for each of the maps. Figure 5.31 shows single-slice T_1 maps from the bSSFP and FISP MRF acquisitions. Figures 5.32 and 5.33 show the T_1 maps from the optimised genetic and *fmincon* acquisitions, respectively. Quantitative comparisons are shown Table 5.28.

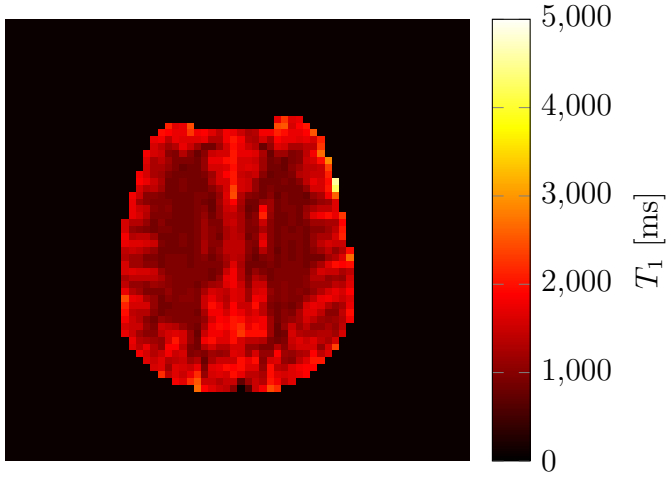

(a)

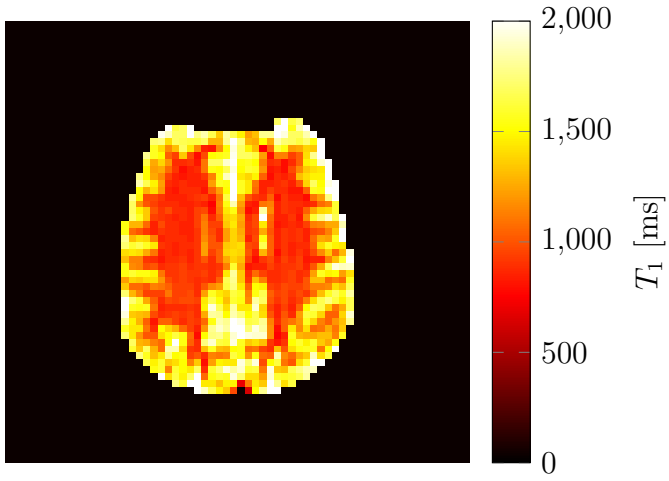

(b)

Figure 5.30: *In vivo* T_1 map (masked) of an asymptomatic volunteer, obtained via the gold standard T_1-mapping method introduced in Chapter 3, from the same sessions as the 50-timepoint SE EPI MRF acquisitions. Two T_1 maps are shown, each with a different colour range to highlight the tissue differences.

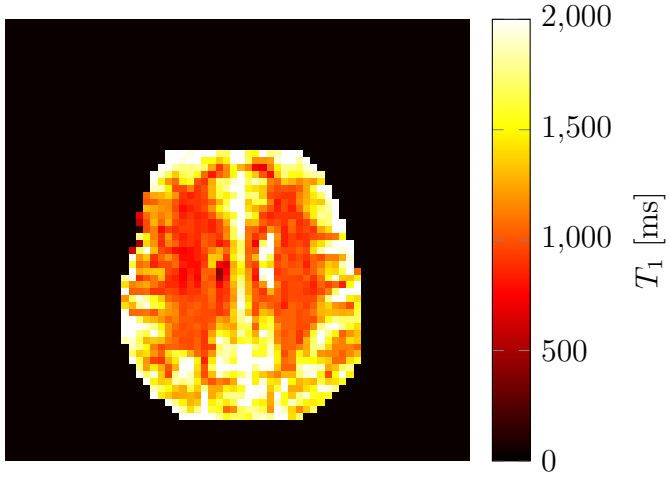

(a)

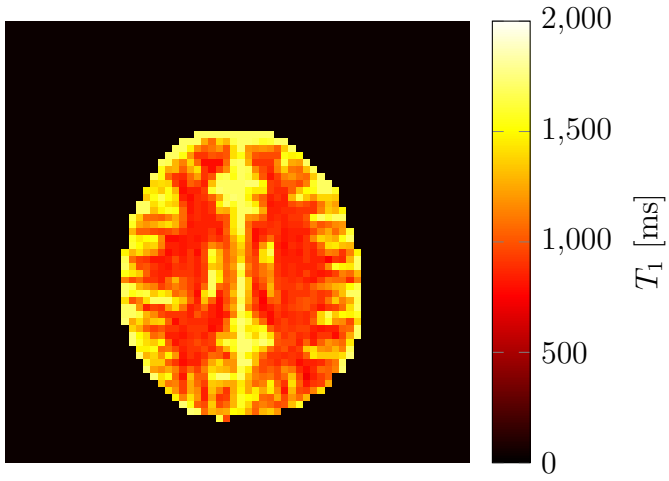

(b)

Figure 5.31: *In vivo* bSSFP (a) and FISP (b) MRF T_1 maps (masked) obtained from the same scanning session as the 50-timepoint SE EPI MRF acquisitions.

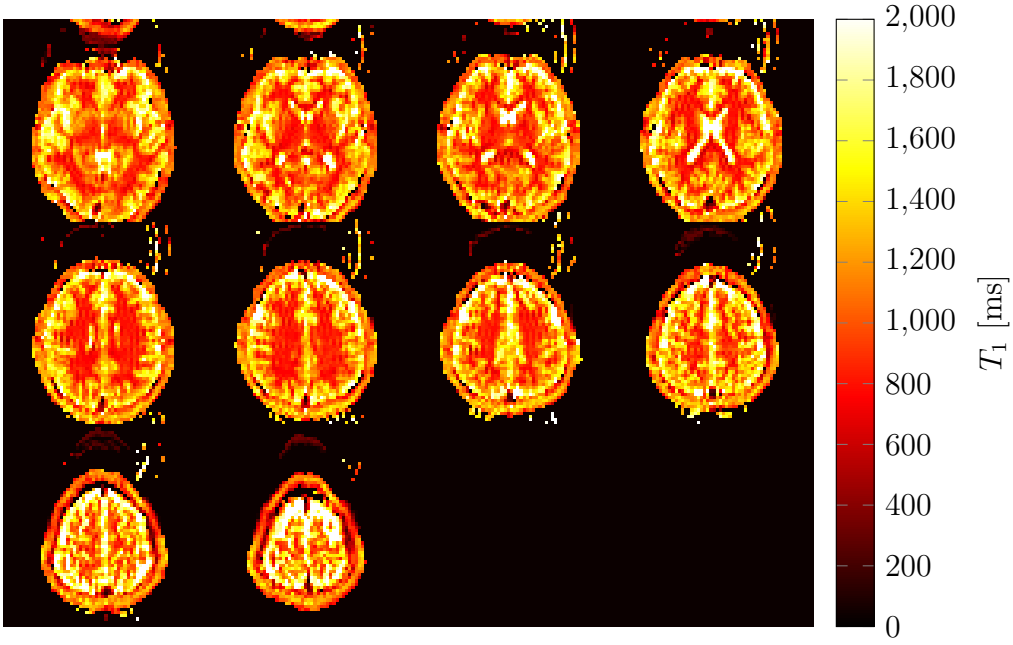

Figure 5.32: *In vivo* SE EPI MRF T_1 from the optimised 50-timepoint genetic algorithm schedule.

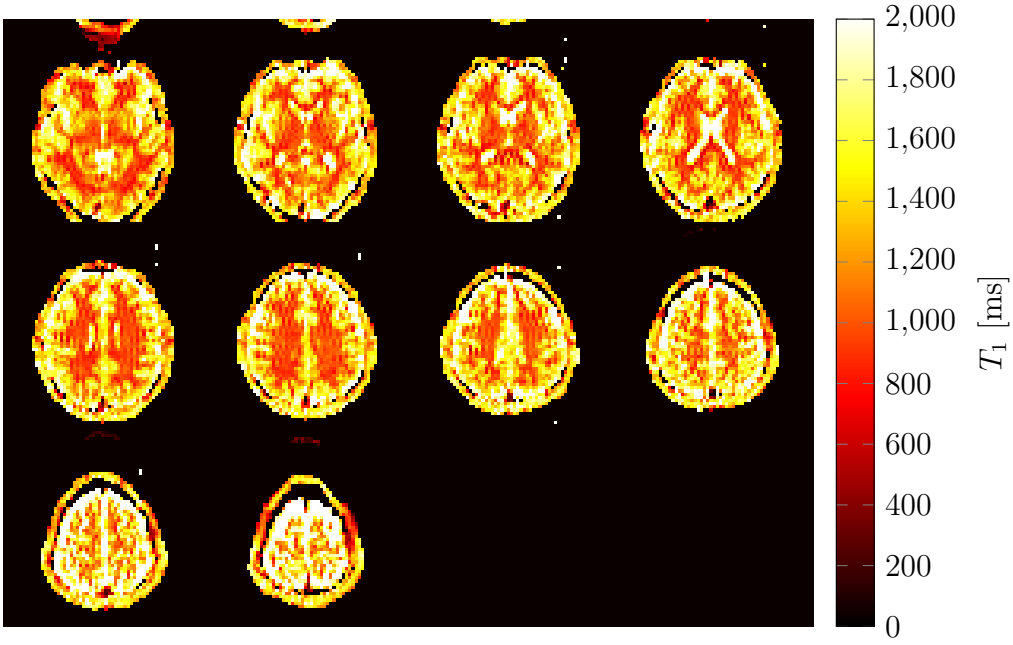

Figure 5.33: *In vivo* SE EPI MRF T_1 from the optimised 50-timepoint *fmincon* algorithm schedule.

In Vivo T_1 **Analysis** Table 5.28 shows the T_1 values from the grey and white matter masks, showing that SE EPI MRF measurements were similar to the gold standard. The bias results in the same Table show there was less than 10% error for the SE EPI MRF measurements.

Reference [ms]	**Genetic [ms]**	***fmincon* [ms]**	**bSSFP [ms]**	**FISP [ms]**
887±72	827±105	964±135	934±110	898±92
1472±365	1328±286	1501±290	1499±437	1328±250

(a)

Reference [ms]	**Genetic [%]**	***fmincon* [%]**	**bSSFP [%]**	**FISP [%]**
887	-6.8	+8.7	+5.3	**+1.2**
1472	-9.8	+2.0	**+1.8**	-9.8

(b)

Table 5.28: *In vivo* T_1 values (a) from manually-drawn grey and white matter masks of maps from the 50-timepoint SE EPI MRF schedules and the spiral MRF implementations. The biases are shown (b) as the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (Column 1), as a percentage of the gold standard mean. The smallest absolute bias for each row is in red.

In Vivo T_2 **Maps** Figure 5.34 shows the T_2 maps acquired using the gold standard (reference) method. Figure 5.35 shows the fitted T_2 maps from the bSSFP and FISP acquisitions, with a large banding artefact present in the bSSFP T_2 map (also seen in Section 5.4). The measurements from the optimised 50-timepoint SE EPI MRF genetic algorithm and *fmincon* acquisitions are shown in Fig. 5.36 and Fig. 5.37, respectively. Table 5.29 shows the white and grey matter T_2 distributions and biases.

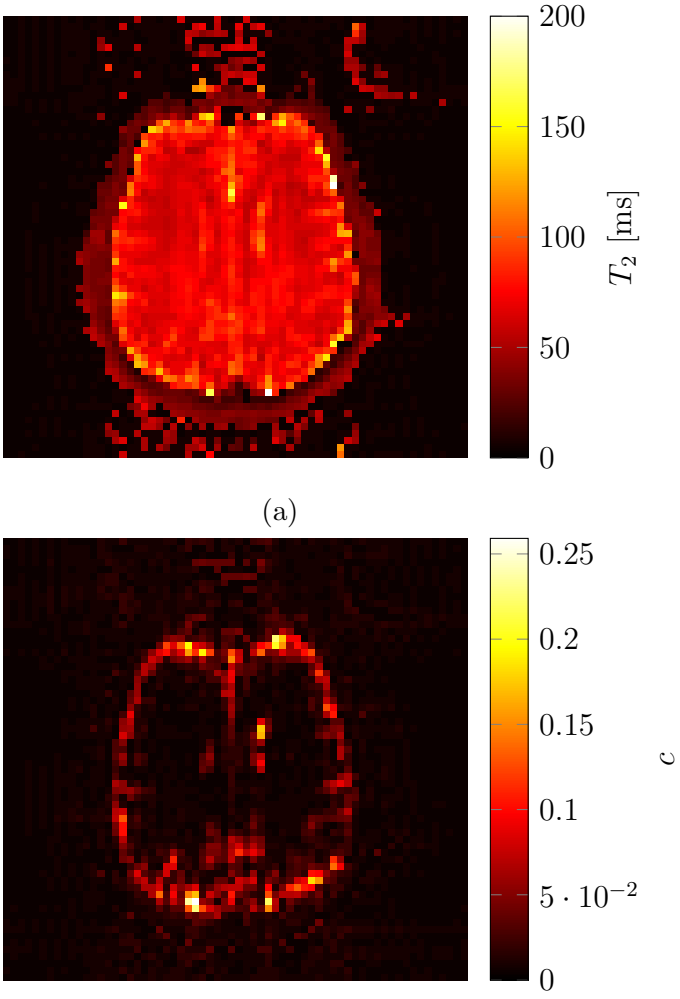

Figure 5.34: *In vivo* T_2 map (a) obtained from the gold standard T_2-mapping method, with a map of the corresponding noise term c (b).

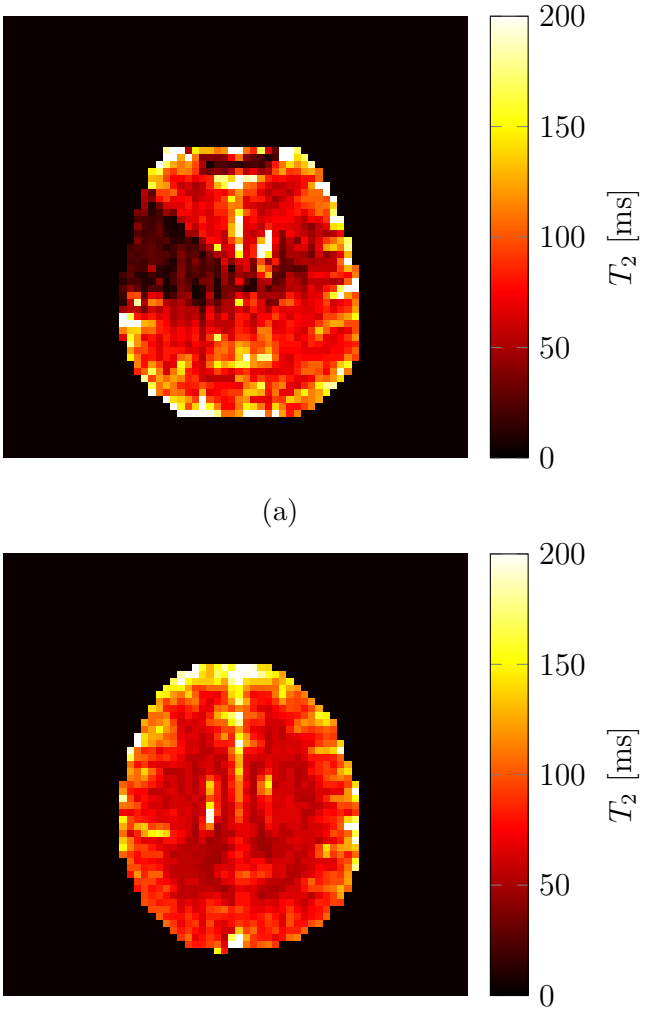

Figure 5.35: *In vivo* bSSFP (a) and FISP (b) MRF T_2 maps (masked), obtained via the acquisitions described in Chapter 3 with the excitation slice profiles included in the bSSFP signal simulations.

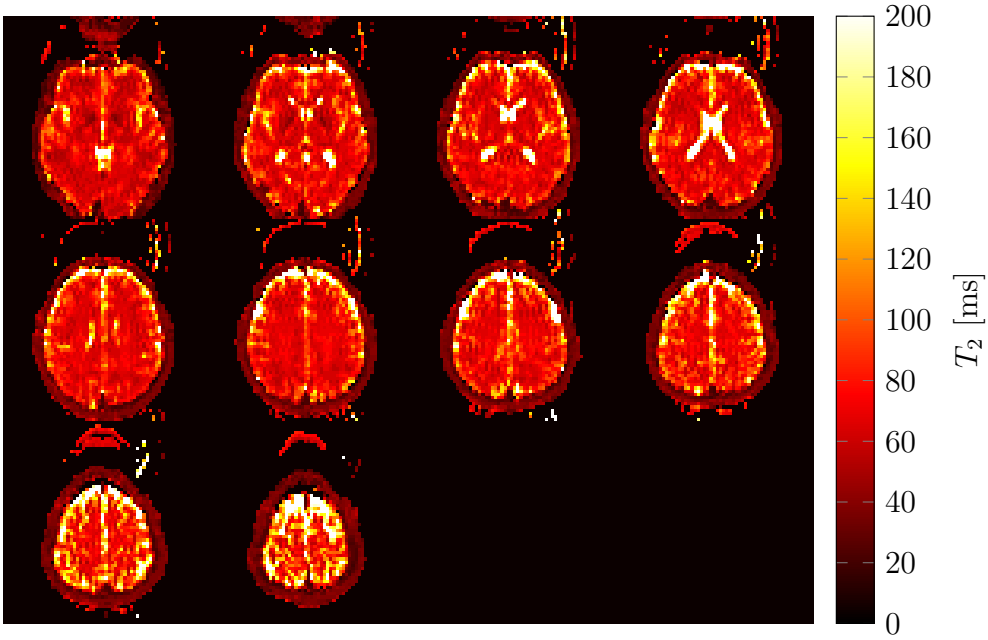

Figure 5.36: *In vivo* SE EPI MRF: T_2 maps for 10 slices, using the optimised genetic algorithm schedule with 50 timepoints.

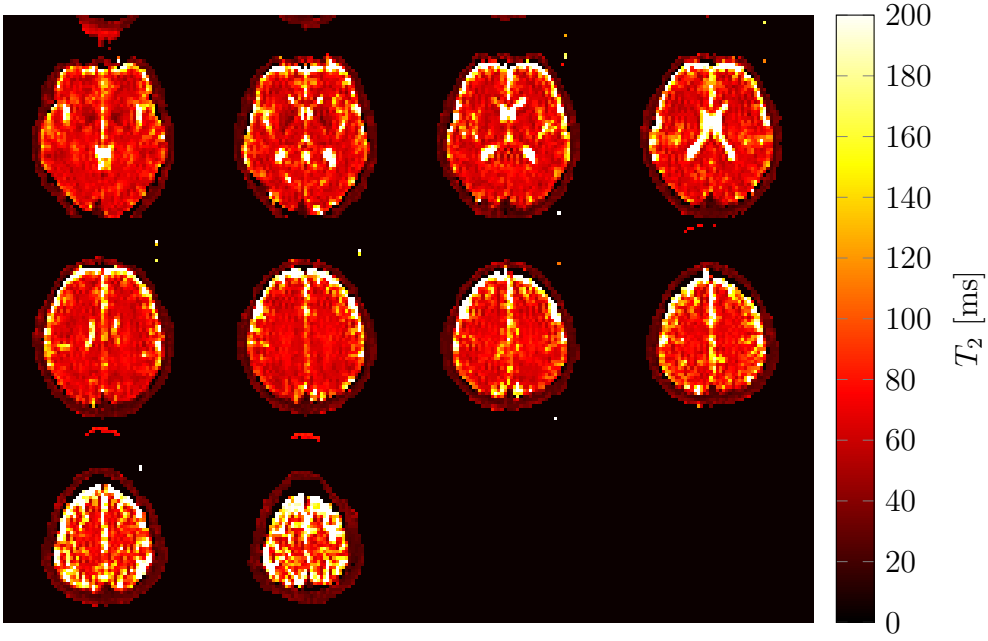

Figure 5.37: *In vivo* SE EPI MRF: T_2 maps for 10 slices, using the optimised *fmincon* schedule with 50 timepoints.

In Vivo T_2 **Analysis** Of all the acquisitions, the white and grey matter T_2 spreads from the genetic algorithm schedule were the most similar to the reference (Table 5.29). Furthermore, the genetic schedule showed the highest accuracy.

Reference [ms]	**Genetic [ms]**	***fmincon* [ms]**	**bSSFP [ms]**	**FISP [ms]**
64±7	64±10	65±14	48±22	62±11
76±16	87±10	91±39	91±46	99±33

(a)

Reference [ms]	**Genetic [%]**	***fmincon* [%]**	**bSSFP [%]**	**FISP [%]**
64	0.0	+1.6	-25.0	-3.1
76	+14.5	+19.7	+19.7	+30.3

(b)

Table 5.29: *In vivo* T_2 values (a) from manually drawn grey and white matter masks of maps from the 50-timepoint SE EPI MRF schedules and the spiral MRF implementations. The biases are shown (b) as the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (Column 1), as a percentage of the gold standard mean. The smallest absolute bias for each row is in red.

In Vivo B_1^+ **Efficiency Maps** Figure 5.38 shows the gold standard 3D DREAM B_1^+ maps. Figures 5.39 and 5.40 show the fitted B_1^+ maps obtained via the genetic and *fmincon* algorithms, respectively. As seen for the 20-timepoint schedules in Section 5.4, the fitted SE EPI MRF B_1^+ values are at the upper end of the dictionary range and do not resemble the gold standard maps from the same session (Fig. 5.38).

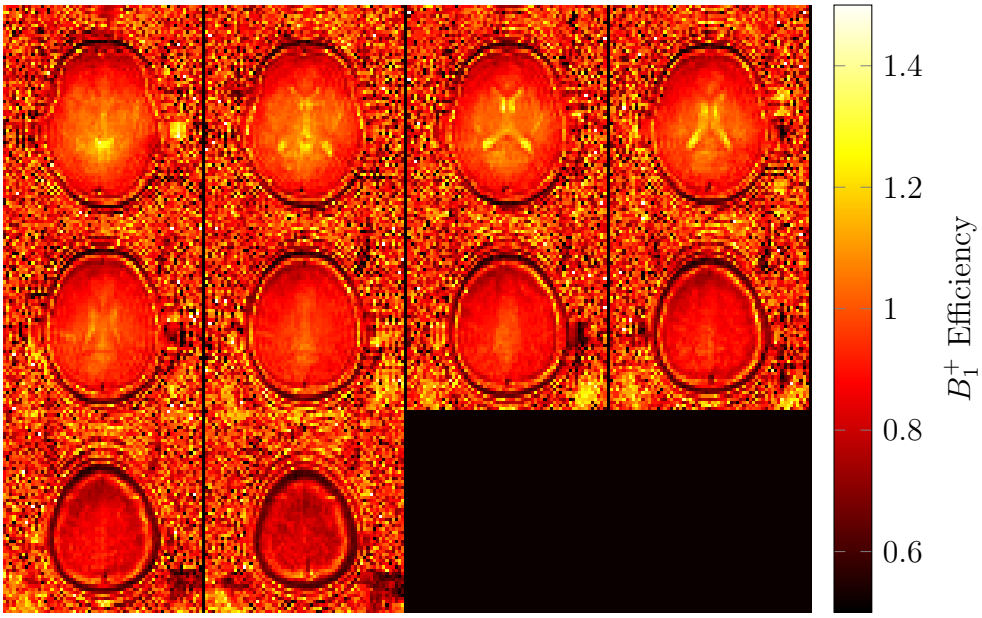

Figure 5.38: *In vivo* gold standard 3D DREAM B_1^+ efficiency maps. The central 10 slices are shown, which correspond to the 10 SE EPI MRF slices.

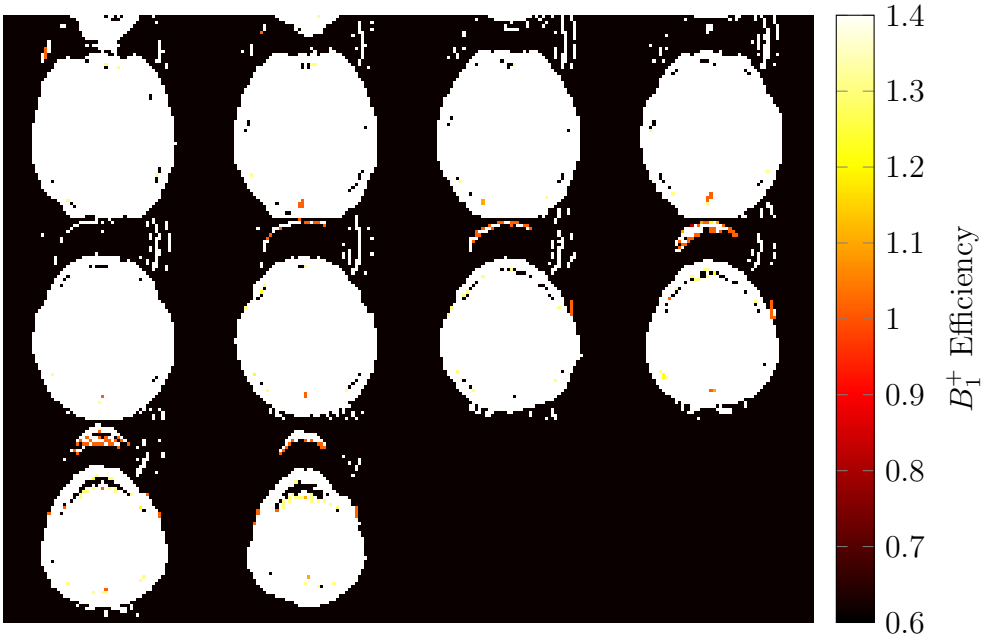

Figure 5.39: *In vivo* SE EPI MRF: the B_1^+ efficiency obtained using the optimised 50-timepoint genetic algorithm schedule.

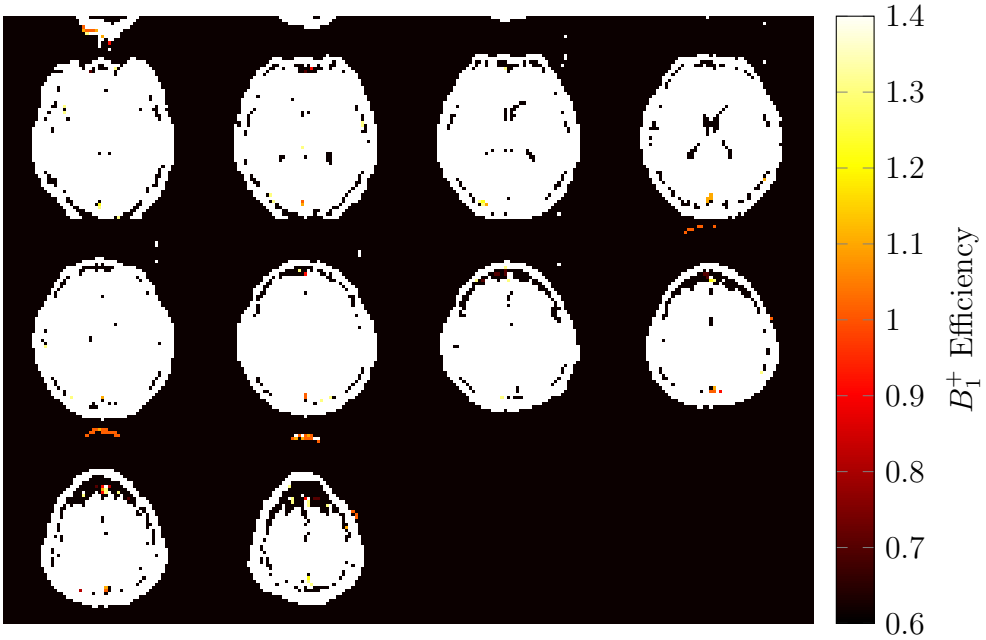

Figure 5.40: *In vivo* SE EPI MRF: the B_1^+ efficiency obtained using the optimised 50-timepoint *fmincon* schedule.

In Vivo M_0 **Maps** Figure 5.41 shows the scaling parameters from the gold standard T_1 and T_2 fits, which reflect M_0. Figure 5.42 shows the scaled M_0 maps from the bSSFP and FISP spiral MRF acquisitions. There is evidence of a banding artefact in the bSSFP M_0 map, as shown in the same location in the bSSFP T_2 map (Fig. 5.35a). Figures 5.43 and 5.44 show the fitted SE EPI MRF M_0 from schedules optimised using the genetic and *fmincon* algorithms, respectively.

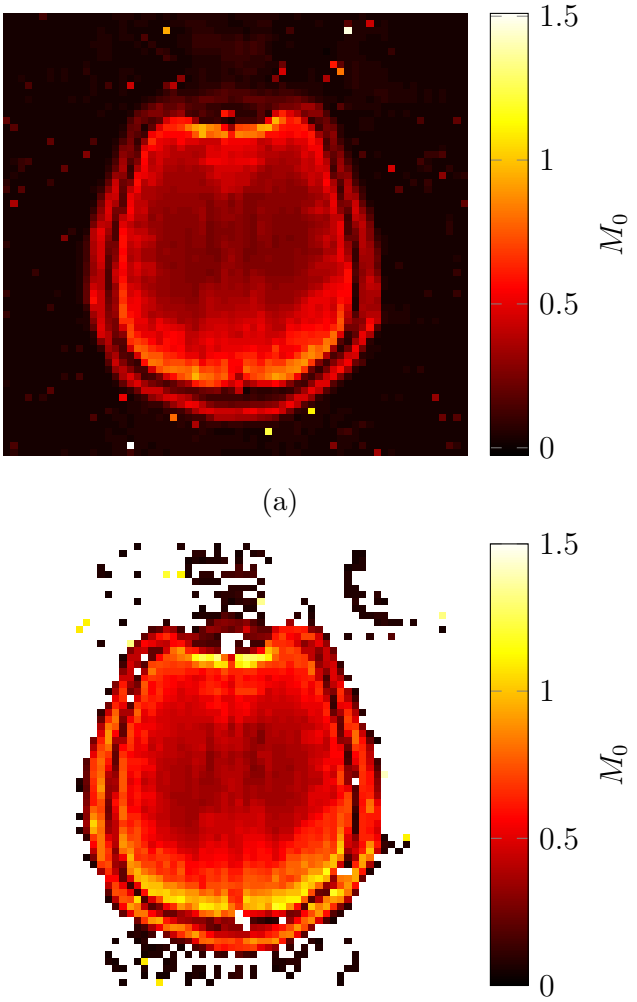

Figure 5.41: *In vivo* M_0 maps obtained from the gold standard T_1 (a) and T_2 (b) fits, using data from the same session as the 50-timepoint SE EPI MRF scans.

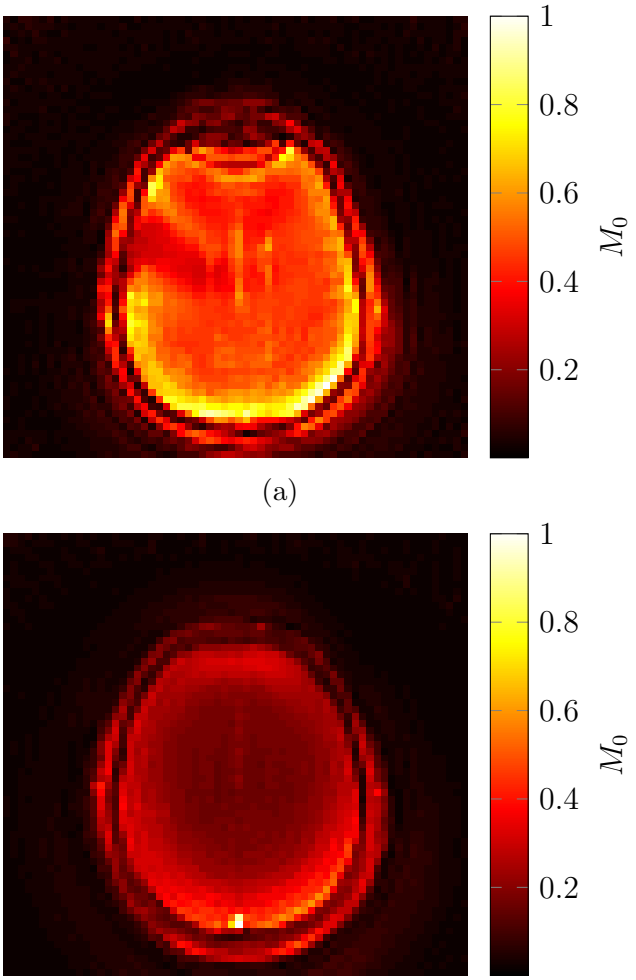

Figure 5.42: *In vivo* bSSFP (a) and FISP (b) MRF M_0 maps obtained via the acquisitions described in Chapter 3 with the excitation slice profiles included in the bSSFP signal simulations.

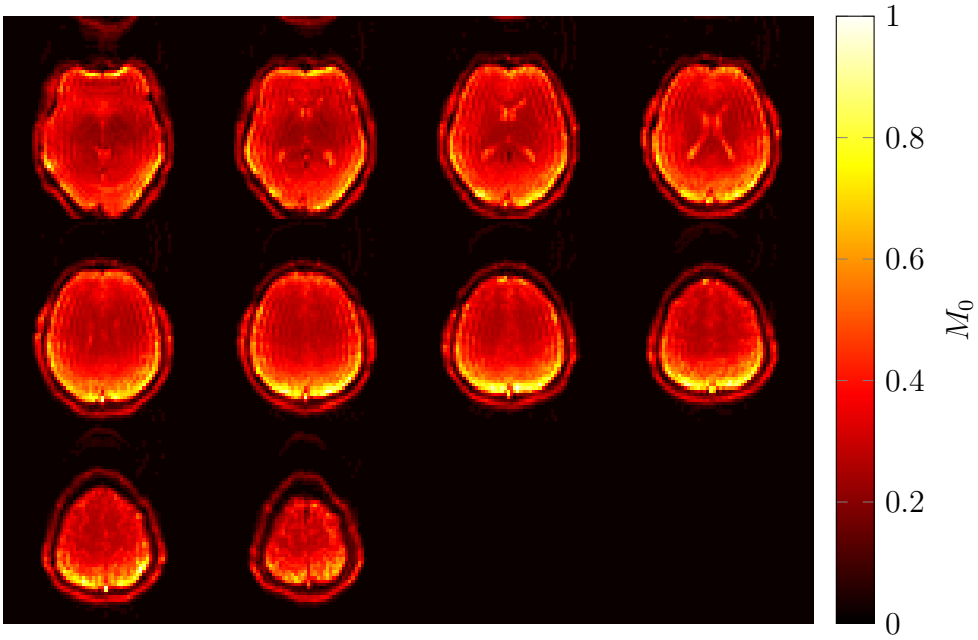

Figure 5.43: *In Vivo* SE EPI MRF normalised M_0 from the optimised 50-timepoint genetic algorithm schedule.

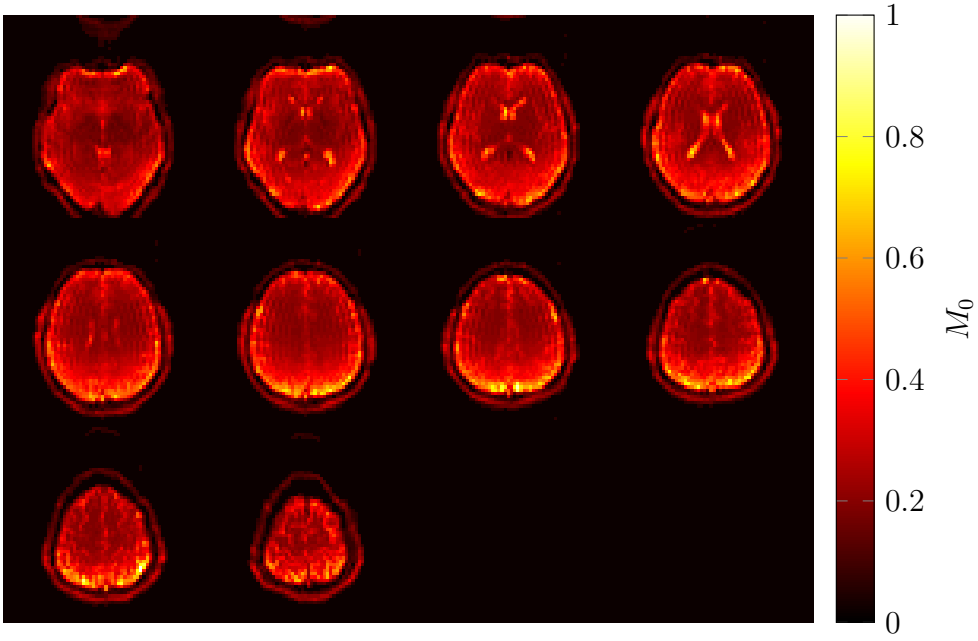

Figure 5.44: *In vivo* SE EPI MRF normalised M_0 from the optimised 50-timepoint *fmincon* schedule.

***In Vivo* Inversion Efficiency Maps** Figures 5.45 and 5.46 show the fitted inversion efficiency from the SE EPI MRF schedules optimised using the genetic and *fmincon* algorithms, respectively. The maps are similar to those derived for the 20-timepoint schedules (Section 5.4). Additionally, across the slices the *fmincon* inversion efficiency fits are consistently lower than those of the genetic algorithm.

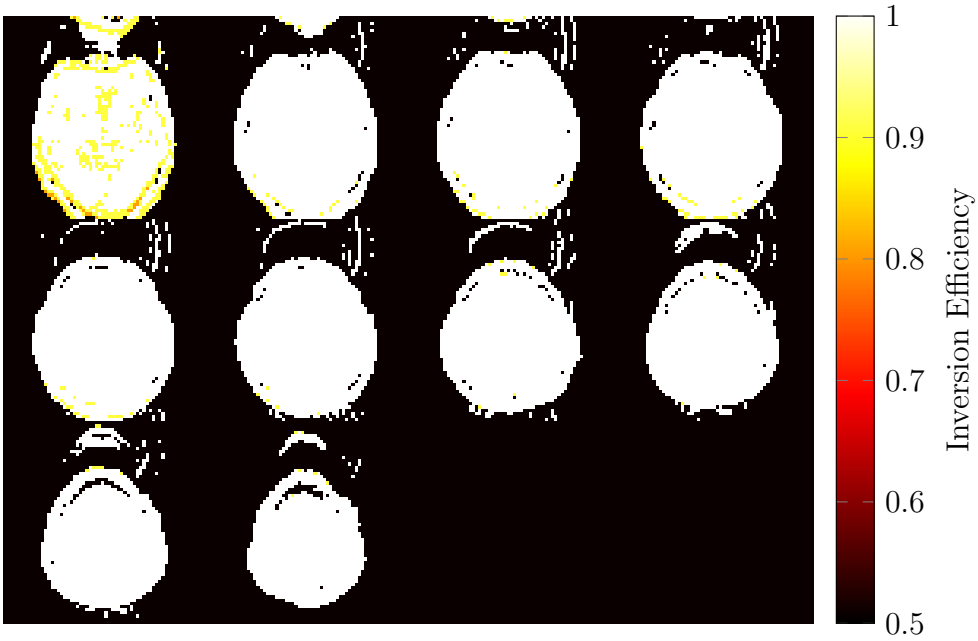

Figure 5.45: *In vivo* inversion efficiency fits from the optimised 50-timepoint genetic algorithm acquisition.

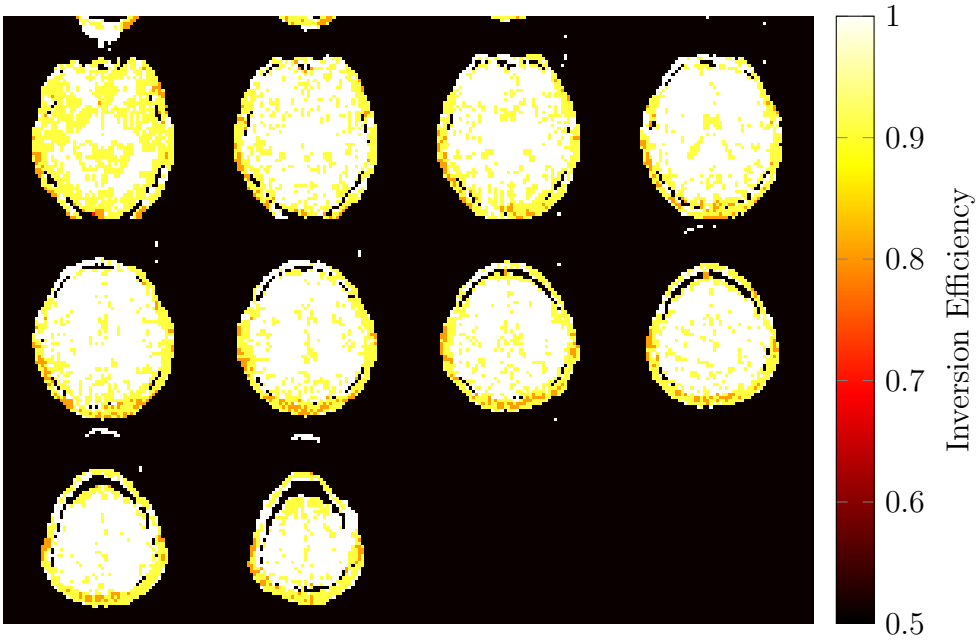

Figure 5.46: *In vivo* inversion efficiency fits from the optimised 50-timepoint *fmincon* acquisition.

5.5.3 50 Timepoints (20 Seconds) per Slice: Conclusions

Results from the optimised 20-timepoint SE EPI MRF schedules (Section 5.4) showed some improvements in efficiency for the optimised schedules but overall there was not sufficient agreement between the predictions of Section 4 and the experimental data. This section explored the performance of SE EPI MRF schedules with 50 timepoints.

The *fmincon*-optimised schedule showed improved T_1 accuracy, precision and efficiency for all of the phantom tubes, with a maximum bias magnitude of 9.2% and a maximum spread of 13.5%. The genetic algorithm optimisation was less successful, only improving the precision for some of the tubes and failing to improve

the accuracy for all the tubes.

The genetic algorithm optimisation improved the T_2 accuracy, with a maximum bias of +8.5%. The corresponding *fmincon* schedule worsened the accuracy for most of the tubes, with a maximum bias of +18.0%. A similar trend was present in the precision results. The largest spread from the optimised genetic schedule was 15.3% and the largest for the *fmincon* schedule was 19.2%.

Concerning the *in vivo* measurements, the optimised genetic algorithm schedule was more accurate (-6.8%) than the *fmincon* schedule (+8.7%) for grey matter T_1. The *fmincon* schedule was better for white matter T_1: +2.0% compared to -9.8% for the genetic schedule. For T_2 the genetic algorithm schedule was more accurate than all the other MRF acquisitions (0.0% for grey matter, +14.5% for white matter).

As in the 20-timepoint results (Section 5.4) the 50-timepoint schedule was not able to accurately fit B_1^+, potentially causing error in the fitted T_1 and T_2. This result highlights the importance of exploring the use of a separately acquired B_1^+ data in future research, to correct the dictionary flip angles on a voxel-wise basis.

5.6 50 Timepoints per Slice, with B_1^+ Correction.

5.6.1 50 Timepoints per Slice, with B_1^+ Correction: Introduction

Section 5.5 displayed discrepancies between the simulated validations (Chapter 4) and results from scan data when B_1^+ was fitted. It suggested that the acquisition

schedules were not able to produce accurate B_1^+ maps by including a B_1^+ dimension in the dictionary. This section explores the accuracy and precision achieved when using separately acquired B_1^+ values to correct the excitation and refocusing angles used in the dictionary creation, thus producing a tailored dictionary for each voxel.

5.6.2 50 Timepoints per Slice, with B_1^+ Correction: Methods

The unoptimised and optimised acquisition schedules from Section 5.5 were used in this section. As in Section 5.5 the schedules had 50 timepoints per slice, resulting in scan durations of 20 seconds per slice.

Reference B_1^+ maps were acquired using the 3D DREAM pulse sequence. As in Chapter 3 the B_1^+ efficiency value for each voxel was calculated by dividing the measured flip angle by the prescribed flip angle. A dictionary was generated for each voxel individually, scaling all the excitation and refocusing flip angles by the B_1^+ efficiency for the given voxel. The inversion pulses were assumed to be insensitive to B_1^+ deviations because they were adiabatic pulses. For this reason the flip angle scaling factor was not applied to the inversion pulses. After fitting the data to the dictionaries, accuracy and precision assessments were performed to compare the unoptimised and optimised schedules. This section focused on whether using B_1^+ corrections improved the agreement between the simulated validations from Chapter 4 and the scan data, rather than how efficient the scans were. For this reason measurement efficiency was not assessed in this section.

5.6.3 50 Timepoints per Slice, with B_1^+ Correction: Results and Discussion

Figure 5.47 shows the T_1 maps from the example initial and optimised genetic algorithm schedules. The slice indices start at the top left corner of the montage of maps and increase from left to right. For example, Slice 10 is represented by the map at Row 3 and Column 2. As seen in Section 5.5, Slice 7 is erroneous and the phantom structure is evident in Slices 9 and 10. There is consistency across the top row of maps (Slices 1 to 4) in each set of maps. As stated earlier in this chapter, these slice positions sample cross-sections of the phantom with similar contents. Tables 5.30 and 5.31 show quantitative analysis of the T_1 values of the phantom tubes from Slice 1, comparing the unoptimised and optimised schedules. The rows of results in each table correspond to the phantom tubes.

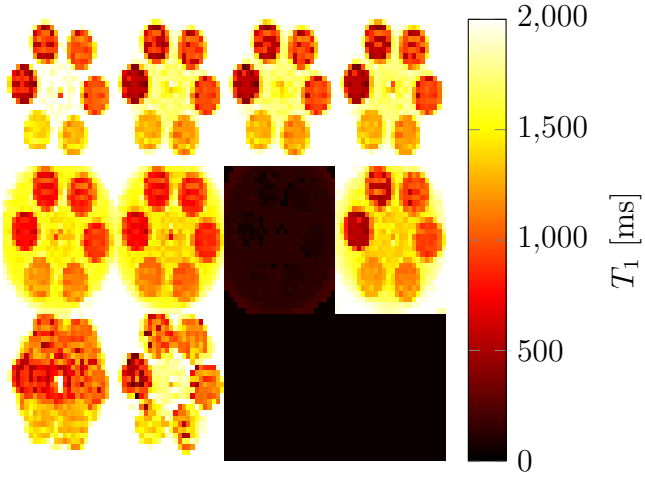

(a)

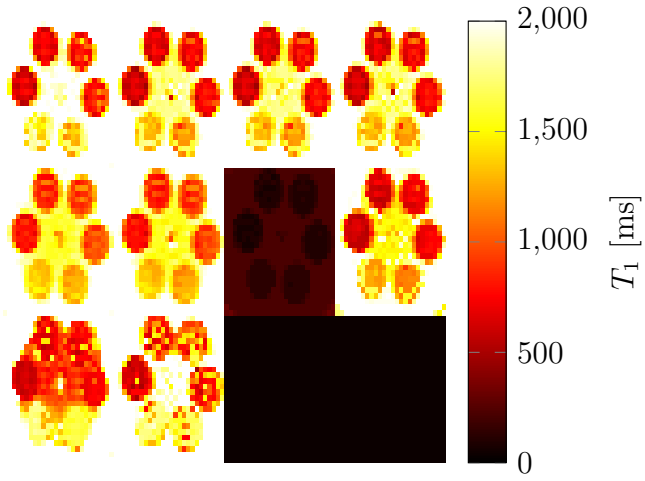

(b)

Figure 5.47: Phantom T_1 maps from the example initial (a) and optimised (b) genetic algorithm schedules, using B_1^+ correction.

Table 5.30 details the accuracy of the Slice 1 T_1 maps, in the form of biases, as percentages of the corresponding gold standard reference means. Ideally the bias would reduce as a result of the optimisation. For the tube with T_1 closest to that targeted during the optimisation (Row 6, T_1 = 1621ms) the optimised SE EPI MRF schedule provided the best accuracy. The bSSFP spiral MRF acquisition

was the most accurate for the other 5 tubes.

Reference [ms]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
879	-23.8	-25.4	-5.0	-13.2
957	-11.5	-27.1	-2.9	-11.7
1068	-6.2	-16.3	-5.0	-11.0
1068	-6.5	-18.0	-3.3	-9.8
1441	-12.5	-2.6	-2.2	-6.9
1621	-11.4	+1.5	-2.5	-5.4

Table 5.30: Genetic algorithm phantom T_1 biases with B_1^+ correction: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean. The smallest absolute bias for each row is in red.

Table 5.31 shows the T_1 precision of the genetic algorithm schedules in the form of the measurement spread, where a small spread is desirable. The optimised schedule failed to improve the T_1 precision for $T_1 = 1621$ms but was successful for the two lowest T_1 values ($T_1 = 879$ms and 957ms). The FISP spiral acquisition was the most precise for all the phantom tubes.

Reference [ms±%]	**Initial** [%]	**Optimised** [%]	**bSSFP** [%]	**FISP** [%]
879±3.9	27.9	10.3	11.5	8.9
957±3.1	25.7	9.2	7.2	5.7
1068±3.5	8.8	23.1	5.8	3.6
1068±3.2	5.5	17.0	6.5	4.8
1441±1.7	5.5	14.4	3.8	2.8
1621±1.4	6.1	12.3	3.4	2.8

Table 5.31: Genetic algorithm phantom T_1 spreads with B_1^+ correction. The smallest spread of each row (other than that of the reference value) is in red.

Figure 5.48 shows the T_1 maps from the unoptimised and optimised *fmincon* schedules, displaying erroneous slices at the same slice indices as seen in the genetic algorithm T_1 maps in Fig. 5.47. Tables 5.32 and 5.33 show quantitative analysis of the Slice 1 maps from Fig. 5.48.

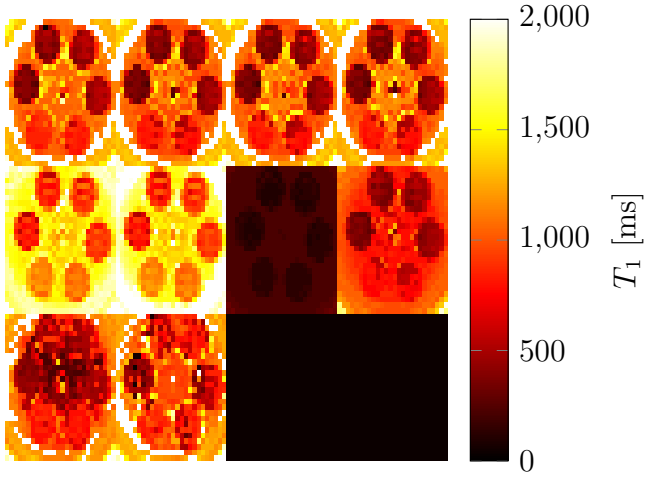

(a)

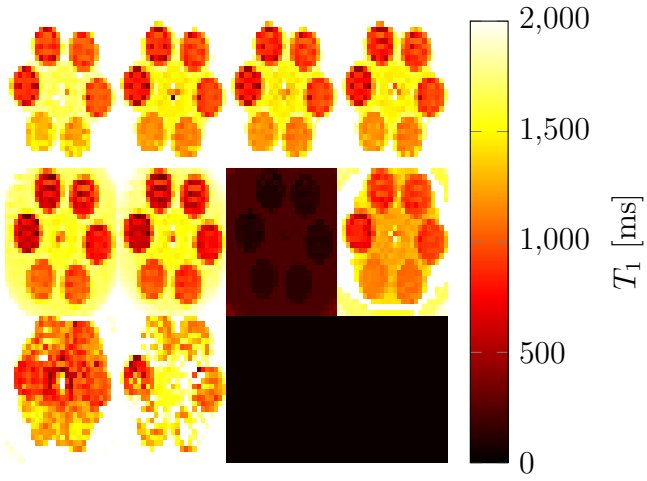

(b)

Figure 5.48: Phantom T_1 maps from the example initial (a) and optimised (b) *fmincon* schedules, using B_1^+ correction.

Table 5.32 shows the T_1 accuracy of the unoptimised and optimised *fmincon* schedules. The optimised schedule shows an improved accuracy compared to the initial schedule for all 6 reference T_1 values. The bSSFP sequence was the most accurate for all the tubes.

Reference [ms]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
879	-57.9	-7.5	-5.0	-13.2
957	-56.3	-4.4	-2.9	-11.7
1068	-55.2	-7.8	-5.0	-11.0
1068	-54.0	-7.3	-3.3	-9.8
1441	-45.5	-14.2	-2.2	-6.9
1621	-47.3	-16.5	-2.5	-5.4

Table 5.32: *fmincon* phantom T_1 biases with B_1^+ correction: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean.

Table 5.33 shows the T_1 spreads from the *fmincon* schedules. The optimised schedule failed to reduce the spread for $T_1 = 1621$ms but was successful for 4 of the other 5 tubes.

Reference [ms±%]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
879±3.9	12.3	16.3	11.5	8.9
957±3.1	12.8	10.6	7.2	5.7
1068±3.5	11.0	6.3	5.8	3.6
1068±3.2	7.9	3.9	6.5	4.8
1441±1.7	13.2	9.5	3.8	2.8
1621±1.4	4.9	7.4	3.4	2.8

Table 5.33: *fmincon* phantom T_1 spreads with B_1^+ correction. The smallest spread of each row (other than that of the reference value) is in red.

Figure 5.49 shows the T_2 maps from the unoptimised and optimised genetic algorithm schedules. Contrary to the T_1 maps, the T_2 maps from the optimised schedule show consistency across all the slices. Tables 5.34 and 5.35 show quantitative analysis of Slice 1 of each set of maps.

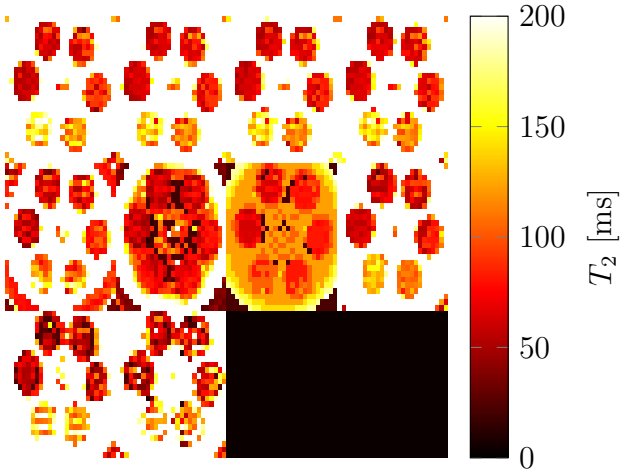

(a)

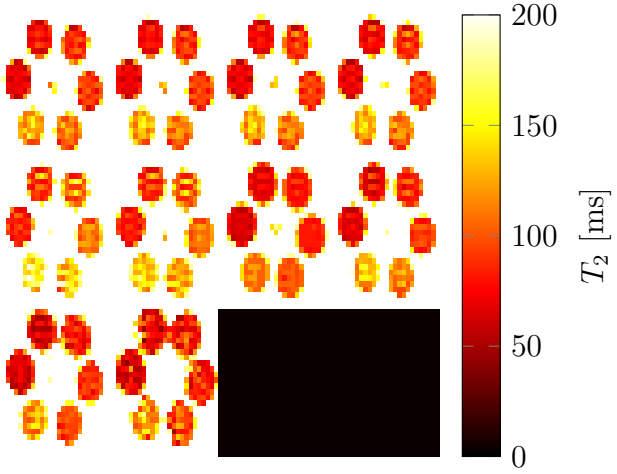

(b)

Figure 5.49: Phantom T_2 maps from the example initial (a) and optimised (b) genetic algorithm schedules, using B_1^+ correction.

Table 5.34 shows the T_2 biases from the genetic algorithm schedules. The optimised schedule successfully reduced the bias for the T_2 closest to that targeted during the optimisation ($T_2 = 123$ms). All the biases from the optimised schedule were less than 10%.

Reference [ms]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
65	-9.2	+4.6	+13.8	+4.6
69	-1.4	+8.7	+13.0	+10.1
83	-1.2	+4.8	+7.2	+6.0
82	+3.7	+9.8	+22.0	+14.6
103	+16.5	+4.9	+13.6	+19.4
123	+38.2	+3.3	+0.8	+11.4

Table 5.34: Genetic algorithm phantom T_2 biases: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean. The smallest absolute bias for each row is highlighted in red.

Table 5.35 shows the T_2 spreads from the genetic algorithm schedules. The optimised schedule successfully reduced the spread for all the tubes.

Reference [ms±%]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
65±6.2	12.5	7.9	12.2	14.7
69±7.2	24.5	13.9	14.1	11.8
83±6.0	27.6	11.4	10.1	8.0
82±9.8	15.8	9.6	9.0	8.5
103±6.8	12.8	12.5	7.7	5.7
123±5.7	14.6	13.7	5.6	6.6

Table 5.35: Genetic algorithm phantom T_2 spreads. The smallest spread for each row, other than that of the reference, is highlighted in red.

Figure 5.50 shows the T_2 maps from the unoptimised and optimised *fmincon* schedules. As seen in the T_2 results from the optimised genetic algorithm schedule, there is consistency across all the slices. This suggests that the reason for erroneous maps at Slice 7 in Figs. 5.47 and 5.48 was related to T_1 and the longitudinal magnetisation, rather than T_2.

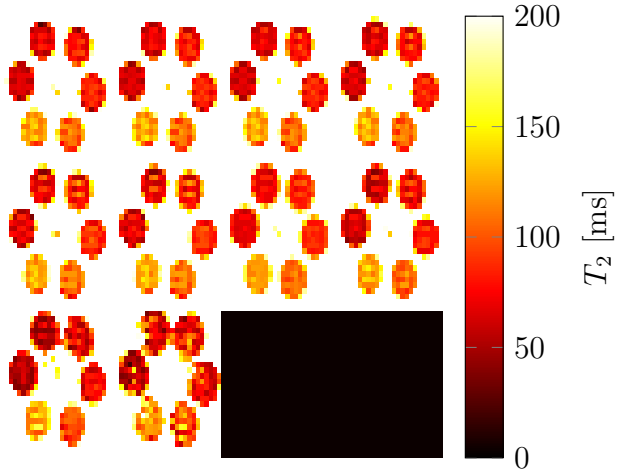

(a)

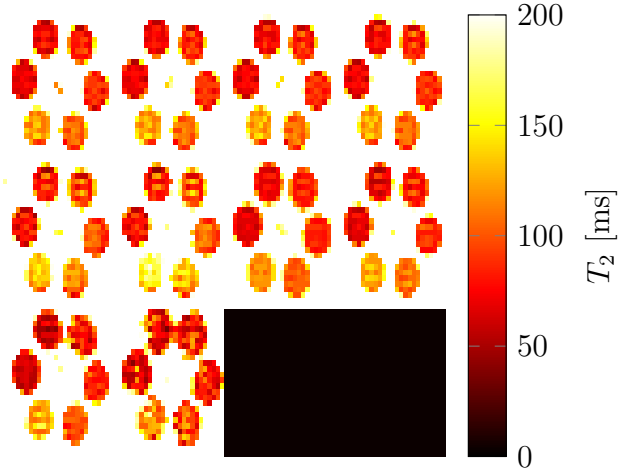

(b)

Figure 5.50: Phantom T_2 maps from the example initial (a) and optimised (b) *fmincon* schedules, using B_1^+ correction.

Table 5.36 shows the T_2 bias from the *fmincon* schedules. The biases from the optimised schedule are all below 5% but the optimised schedule was outperformed by the initial schedule for most of the reference T_2 values.

Reference [ms]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
65	-3.1	+4.6	+13.8	+4.6
69	-4.3	+4.3	+13.0	+10.1
83	-6.0	+0.0	+7.2	+6.0
82	-1.2	+3.7	+22.0	+14.6
103	+1.0	-1.0	+13.6	+19.4
123	-0.8	-1.6	+0.8	+11.4

Table 5.36: *fmincon* phantom T_2 biases: the difference between the mean of each distribution and the mean of the corresponding gold standard measurement (reference), as a percentage of the gold standard mean. The smallest absolute bias for each row is highlighted in red.

Table 5.37 shows the T_2 spreads of the *fmincon* schedules. The optimised schedule reduced the spread for all the reference T_2 means apart from $T_2 = 123$ms.

Reference [ms±%]	Initial [%]	Optimised [%]	bSSFP [%]	FISP [%]
65±6.2	9.9	7.5	12.2	14.7
69±7.2	15.4	14.3	14.1	11.8
83±6.0	13.3	12.4	10.1	8.0
82±9.8	11.6	11.1	9.0	8.5
103±6.8	10.6	9.0	7.7	5.7
123±5.7	8.2	9.1	5.6	6.6

Table 5.37: *fmincon* phantom T_2 spreads. The smallest spread for each row, other than that of the reference, is highlighted in red.

5.6.4 50 Timepoints per Slice, with B_1^+ Correction: Conclusions

This section presented phantom T_1 and T_2 measurements from the 50-timepoint schedules from Section 5.5 with B_1^+ correction, using B_1^+ maps obtained using the 3D DREAM acquisition.

The genetic and *fmincon* optimised schedules improved the T_1 accuracy for the tube with a similar T_1 to that targeted during the optimisation: for $T_1 = 1621$ms

the bias was reduced from -11.4% to +1.5% with the genetic algorithm and -47.3% to -16.5% with *fmincon*, compared to the corresponding initial schedules. However, the precision results for both algorithms showed a reduction in precision for this T_1 compared to the unoptimised schedules, reflected in the efficiency results. For the majority of the tubes the bSSFP sequence was the most accurate and the FISP sequence was the most precise, in terms of T_1.

Regarding the T_2 measurements, the genetic algorithm improved the accuracy for the T_2 similar to that targeted: for $T_2 = 123$ms the spread was reduced from +38.2% to +3.3%. The optimised *fmincon* schedule increased the T_2 bias to -1.6%, although the bias from the corresponding unoptimised schedule was already very small (-0.8%). Both optimised schedules produced biases under 10%: the largest bias for the optimised *fmincon* schedule was +4.6%, compared to +9.8% for the optimised genetic algorithm schedule. Small improvements in the T_2 spread were observed as a result of the genetic algorithm optimisation (14.6% to 13.7% for $T_2 = 123$ms) but the optimised *fmincon* schedule increased the spread (8.2% to 9.1% for $T_2 = 123$ms). In general the FISP acquisition outperformed the other schedules in terms of T_2 measurement error, although for some of the tubes the unoptimised schedules actually showed the highest accuracy.

Compared to using a dictionary dimension to fit B_1^+, using a separately acquired B_1^+ map to correct the flip angles improved the agreement between the experimental data and the T_1 accuracy increase predicted for the genetic algorithm optimisation from the simulations. However, with the B_1^+ correction the T_1 accuracy from the optimised *fmincon* schedule was worse than when B_1^+ was fitted (Section 5.5). Using B_1^+ corrections improved the accuracy for the targeted T_2 for the optimised *fmincon* schedules and maintained the absolute T_2 accuracy

of the optimised genetic algorithm schedule (3.3%).

Measurement efficiency was not assessed in this section but it is worth noting that acquiring the 3D DREAM data in addition to the SE EPI MRF data increases the total scan duration by at least 5 seconds, potentially affecting the measurement efficiency. For 10-slice acquisitions (as used in this section) this equates to 0.5 seconds of additional scan time per slice, which is equal to 2.5% of the 50-timepoint SE EPI MRF scan duration per slice. Furthermore, if a 3D DREAM acquisition was performed immediately before the MRF acquisition additional time may be needed to allow magnetisation recovery, depending on the starting magnetisation assumed by the signal simulations.

Overall, using a separately acquired B_1^+ map to correct the flip angles provided some accuracy improvements compared to fitting B_1^+ but the optimised schedules produced some undesirable reductions in precision, particularly for T_1.

5.7 Summary and Conclusions

This chapter presented the implementation and experimental validation of the spin-echo EPI fingerprinting pulse sequence and optimisation framework proposed in Chapter 4. Quantitative comparisons were made between the performance of genetic and *fmincon* optimisation algorithms with regard to the optimised acquisition parameter schedules, including an assessment of the performance of schedules with 20 and 50 timepoints per slice with B_1^+ fitted during the matching process. The optimised schedules with 20 and 50 timepoints were used in phantom and *in vivo* volunteer scans. The 50-timepoint schedules were also tested using a phantom and a separately acquired B_1^+ map.

The experimental results in this chapter confirmed the ability of the optimisation framework to reduce T_1 and T_2 measurement errors in some cases but there were examples of inconsistency between the observed experimental results and the simulated predictions in Chapter 4.

For the phantom data the 50-timepoint optimised schedules provided reductions in the combined measurement error for T_1/T_2 combinations similar to those targeted during the optimisation, both with B_1^+ fitting and with a separately acquired B_1^+ map. Error reductions were also observed for shorter T_1/T_2 combinations as a result of the optimisations. The 20-timepoint optimisations were less successful, with the genetic algorithm optimisation providing an increase in error for the phantom T_1/T_2 pairs closest to the target T_1/T_2 (1600ms/120ms), compared to the initial schedule. However, the 20-timepoint *fmincon* schedule was able to reduce the error for these phantom T_1/T_2 pairs. These observations indicate that the *fmincon* algorithm should be used for the cases described, contradicting the simulations in Chapter 4 that predicted that the genetic algorithm would provide the lowest error for both acquisition lengths. This inconsistency may have been caused in part by the presence of B_1^+ variations in the experimental acquisitions, as B_1^+ was not included in the optimisation steps. Using separately acquired B_1^+ measurements to correct the flip angles provided accuracy improvements compared to fitting B_1^+ but there was still an unexpected worsening of the precision for some cases.

The optimised schedules were successfully implemented *in vivo* and provided reasonable T_1 and T_2 maps. The quantitative comparisons of the *in vivo* data confirmed that the maps from the 50-timepoint schedules were more accurate than those of the 20-timepoint schedules.

The phantom and *in vivo* results suggest that the SE EPI MRF sequences required external B_1^+ corrections. Taking into account the T_1 and T_2 accuracy and precision with B_1^+ corrections it seems that the optimised *fmincon* schedule marginally outperformed the genetic algorithm schedule overall. The largest bias magnitudes across the phantom tubes from the optimised genetic algorithm schedule were -27.1% for T_1 and +9.8% for T_2. The corresponding spreads were 23.1% (T_1) and 13.9% (T_2). For the *fmincon* schedule the largest biases were 16.5% for T_1 and 4.6% for T_2, with maximum spreads of 16.3% for T_1 and 14.3% for T_2.

Overall, considering the phantom and *in vivo* results, the bSSFP sequence was the most accurate for T_1 and the optimised *fmincon* schedule with B_1^+ correction was the most accurate for T_2. The FISP sequence was the most precise and efficient for T_1 and T_2. Therefore the FISP sequence should be used in general for MRF acquisitions, especially since it did not suffer from the banding artefacts that were present in the bSSFP maps.

Future work could explore whether using EPG signal simulations would improve the SE EPI MRF signal model accuracy but this chapter, and work in the literature, suggests that Bloch simulations are sufficient for obtaining parameter maps comparable to those of reference methods. Additionally, it would be worth investigating the quantitative benefits providing by the inclusion of the inversion efficiency dictionary dimension.

All the MRF acquisitions had 4mm in-plane voxel sizes. Any increase in image resolution would likely require slight extensions to the TE and TR to accommodate the longer readout durations needed to sample higher k-space frequencies. This could be accounted for in the SE EPI MRF optimisations by increasing the lower limit of the TE and TR patterns.

The scan durations of the SE EPI MRF schedules used in this chapter were 8 seconds (20-timepoints) and 20 seconds (50-timepoints) per slice, excluding the 3D DREAM acquisition time. It may be beneficial to test SE EPI MRF schedules with scan times more similar to those of the bSSFP and FISP sequences (12.4 seconds). Simultaneous multislice acquisitions could be used to reduce the SE EPI MRF scan duration without reducing the number of timepoints, potentially making the SE EPI MRF measurement efficiency more comparable to that of the bSSFP and FISP sequences.

It is evident from the results in this chapter that further research is required to improve the SE EPI MRF optimisation framework. The choice of optimisation cost function metric may need to be reassessed in order to ensure improvements in both accuracy and precision, for T_1 and T_2.

Chapter 6

Summary and Future Work

This chapter summarises the preceding chapters and suggests future areas of research relating to the work presented in this thesis.

6.1 Thesis Summary

As discussed in Chapter 2, one of the key challenges in MRF research is how best to design the acquisition parameter schedule. The signal perturbation throughout the MRF acquisition is a crucial aspect of the MRF method and depends on these parameters, making it important to choose these parameters wisely. This thesis presented an approach for optimising the acquisition parameters of a spin-echo EPI MRF pulse sequence to reduce T_1 and T_2 measurement error. Motivation was drawn from a potential application to acute stroke imaging studies.

Chapter 3 detailed experimental methods for making reference quantitative measurements of the relaxation parameters T_1 and T_2, using principles introduced in Chapter 2. Double-echo off-resonance maps and 3D DREAM B_1^+ measurements

were also presented. A custom phantom was produced and used to obtain example parameter maps for the quantitative acquisitions. Using signal magnitudes for the gold standard T_1 measurements will have influenced the measured values but this effect was assumed to be negligible. In the 3D DREAM B_1^+ maps, contrast was observed between the phantom tubes, as well as between the tubes and the internal structure of the phantom. This could have been caused by differences in the T_1 values or dielectric properties of those structures, which exceeded the ability of the 3D DREAM sequence to avoid misrepresenting the B_1^+ field somewhat. Balanced SSFP MRF and FISP MRF sequences were also implemented and example relaxation parameter maps from the phantom were shown.

Chapter 4 provided details of the optimisation framework and the theoretical performance of unoptimised and optimised spin-echo EPI MRF schedules. Multiple optimisations were performed for single-slice and 10-slice regimes, using *fmincon* and genetic algorithms with Monte Carlo error simulations to reduce individual and combined T_1 and T_2 errors. Although the optimisations focused on T_1/T_2 = 1600ms/120ms the simulated validations also showed improvement for other relaxation parameter combinations. The results suggested that the framework would produce acceptable measurement errors across the six tubes in the phantom, with lower errors predicted for the optimised genetic algorithm schedule compared to that of *fmincon*.

Chapter 5 demonstrated the performance of the unoptimised and optimised schedules with 20 and 50 timepoints per slice, using phantom and asymptomatic *in vivo* human volunteer scan data to make quantitative comparisons between the schedules. As predicted by the simulations in Chapter 4, the 50-timepoint schedules generally performed better than the 20-timepoint schedules. However,

neither set of schedules were able to accurately fit B_1^+, prompting further phantom tests with the 50-timepoint schedules to explore the effect of using a separate B_1^+ map to adjust the dictionary flip angles for each voxel. Using B_1^+ correction from additionally acquired B_1^+ maps produced improved accuracy in general, although in some cases the optimised schedules unexpectedly increased the measurement bias. The results with B_1^+ correction suggested that in an experimental setting the optimised *fmincon* schedule marginally outperformed that of the genetic algorithm, with maximum bias magnitudes of 16.5% for T_1 and 4.6% for T_2, across the phantom tubes. The largest corresponding spreads were 16.3% (T_1) and 14.3% (T_2). The FISP MRF sequence performed the best overall but the 50-timepoint SE EPI MRF schedules did provide reasonable *in vivo* measurements.

6.2 Future Research

This section comments on possible areas for further study concerning the work in this thesis.

6.2.1 Patient Scans

To demonstrate the utility of the framework in a clinical environment optimised SE EPI MRF schedules should be used in a local acute stroke imaging study such as the AMICI study at the Acute Vascular Imaging Centre (AVIC) in Oxford.

6.2.2 Reproducibility

In this thesis it was assumed that the scan room temperature was constant across all the phantom scans and that this ensured the gold standard T_1 and T_2 measurements were valid for all the scan sessions. However, it is likely that some temperature fluctuations did occur. Therefore it would be worth assessing the temperature differences over a series of days and making gold standard T_1 and T_2 measurements on each of those days. It would also be worth investigating the stability of the SE EPI MRF acquisitions themselves by running the SE EPI MRF scans multiple times during a single scan session, as well as over multiple sessions.

6.2.3 Optimisation

Some cases in this thesis showed agreement between the predictions and the experimental data for the targeted T_1/T_2 but not for other T_1 and T_2 pairs. It would be valuable to explore the benefits of simultaneously optimising with multiple target T_1/T_2 pairs, such as with the shortest and the longest expected T_1/T_2 combinations. The TR durations and TR block lengths could also be modified or indeed optimised.

6.2.4 Simultaneous Multislice

Achieving greater brain coverage, whilst maintaining the imaging slice thickness used in this thesis (5mm), would require an increase in the number of acquired slices. However, optimising for an increased number of slices would be more computationally costly and would increase the total sequence duration, potentially reducing the measurement efficiency. Alternatively a 10-slice acquisition could

be repeated in series but this would also increase the total scan time. To be clinically practical the total scan duration should be kept short (e.g. less than 5 minutes). The 50-timepoint schedules were the longest of the SE EPI MRF schedules tested experimentally in this thesis, each with total scan durations of 3 minutes 20 seconds (excluding any additional 3D DREAM B_1^+ scans). Performing a simultaneous multislice acquisition would allow the slice coverage to be extended whilst maintaining a clinically acceptable total scan duration. As mentioned in Chapter 5, a SMS acquisition would also make the 50-timepoint scan duration per slice (20 seconds) more similar to that of the bSSFP and FISP sequences (12.4 seconds), thereby enabling a more direct comparison of the schedule performances. Simultaneously exciting 2 slices (i.e. MB factor = 2) would allow 20 slices to be acquired in the same acquisition time as the 10-slice schedules, plus a small amount of time for calibration images to be acquired, meaning a scan duration of 12 seconds per slice could be possible.

The following paragraphs present preliminary work towards achieving SMS acquisitions with the SE EPI MRF sequence, including the production of MultiBand (MB) RF pulses (required by SMS) and modifying the phase-encoding gradients.

6.2.4.1 Multiband Excitation

Using code provided by Dr. Jürgen Finsterbusch on the Siemens IDEA discussion board, the excitation and refocusing pulses of the SE EPI MRF pulse sequence in this thesis were given the capability of simultaneously perturbing 2 slices (i.e. MB factor = 2). Example single-band and multiband pulse amplitude profiles are shown in Fig. 6.1. The adiabatic inversion pulse would also need to be redesigned, which would require a different approach due to the mechanism adiabatic pulses

use to rotate the magnetisation vector.

(a) (b)

Figure 6.1: Example single-band (a) and multiband (b) excitation pulse amplitude profiles. The multiband pulse was designed to excite two slices (i.e. multiband factor = 2).

6.2.4.2 G_z Blips

A technique known as blipped-CAIPI [98] can be used to allow data from simultaneously excited slices to be more cleanly reconstructed into separate images, reducing the noise amplification that would occur otherwise. Gradient blips along the phase-encoding axis are used to create a controlled phase difference between adjacent lines of k-space, which in turn causes controlled shifting between the images of each slice, making it easier to distinguish between them. For example, for the 2-slice MB example in [98] the phase difference was chosen such that there was a FOV/2 shift between the uncorrected images of each slice. Figure 6.2 shows multiband RF pulses (MB factor = 2) in the context of a SMS SE EPI MRF pulse sequence. The G_z component in Fig. 6.2 possesses gradient blips before each line of k-space is acquired, designed to create the aforementioned blipped-CAIPI shift.

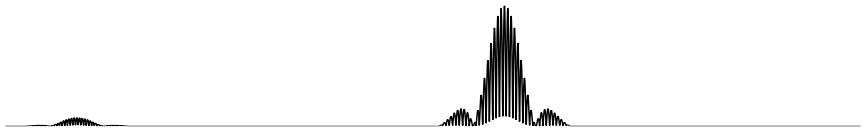

(a) RF

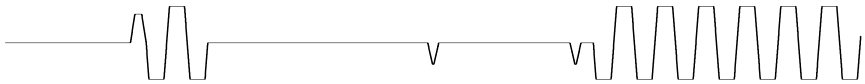

(b) G_x

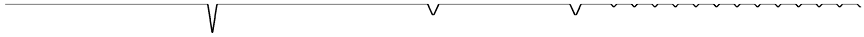

(c) G_y

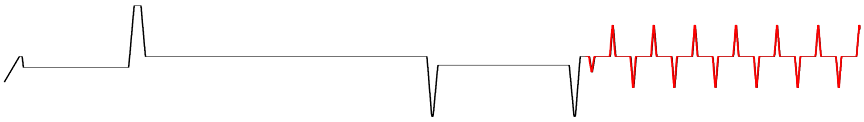

(d) G_z

(e) ADC

Figure 6.2: Example SE EPI MRF pulse sequence with multiband excitation and refocusing RF pulses, showing only a subsection of the TR. The RF (a), G_x (b), G_y (c), G_z (d) and ADC (e) components are shown. The RF (a) and G_z (d) components are different to those of a single-band acquisition but the G_x (b), G_y (c) and ADC (e) components are unchanged. The RF pulses are composite pulses, designed to simultaneously excite two slices. The G_z component has gradient *blips* (red) before each k-space line is sampled, with the aim of achieving the phase shifts required by blipped-CAIPI.

The sequence developments shown in Fig. 6.2 represent the initial steps towards a SMS implementation to allow increased slice coverage and improved scan efficiency for the SE EPI MRF sequence. To complete this work, multiband capability should be added to the inversion pulses, the blipped-CAIPI shift should be tested, and a specialised SMS image reconstruction pipeline should be developed.

References

[1] D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–92, 2013.

[2] G. J. Hankey. The global and regional burden of stroke. *The Lancet Global Health*, 1(5):e239–e240, 2013.

[3] G. J. Hankey. Stroke. *The Lancet*, 389(10069):641–654, 2017.

[4] C. Ayata and A. H. Ropper. Ischaemic brain oedema. *Journal of Clinical Neuroscience*, 9(2):113–124, 2002.

[5] M. Simard, T. A. Kent, M. Chen, K. V. Tarasov, and V. Gerzanich. Brain oedema in focal ischaemia: molecular pathophysiology and theoretical implications. *Lancet Neurology*, 6(3):258–268, 2007.

[6] M. Wintermark, G. W. Albers, J. P. Broderick, A. M. Demchuk, J. B. Fiebach, J. Fiehler, J. C. Grotta, G. Houser, T. G. Jovin, K. R. Lees, M. H. Lev, D. S. Liebeskind, M. Luby, K. W. Muir, M. W. Parsons, R. Von Kummer, J. M. Wardlaw, O. Wu, A. J. Yoo, A. V. Alexandrov, J. R. Alger, R. I. Aviv, R. Bammer, J. C. Baron, F. Calamante, B. C. V. Campbell, T. C. Car-

penter, S. Christensen, W. A. Copen, C. P. Derdeyn, E. C. Haley, P. Khatri, K. Kudo, M. G. Lansberg, L. L. Latour, T. Y. Lee, R. Leigh, W. Lin, P. Lyden, G. Mair, B. K. Menon, P. Michel, R. Mikulik, R. G. Nogueira, L. Ostergaard, S. Pedraza, C. H. Riedel, H. A. Rowley, P. C. Sanelli, M. Sasaki, J. L. Saver, P. W. Schaefer, P. D. Schellinger, G. Tsivgoulis, L. R. Wechsler, P. M. White, G. Zaharchuk, O. O. Zaidat, S. M. Davis, G. A. Donnan, A. J. Furlan, W. Hacke, D. W. Kang, C. Kidwell, V. N. Thijs, G. Thomalla, and S. J. Warach. Acute stroke imaging research roadmap II. *Stroke*, 44(9):2628–2639, 2013.

[7] G. W. J. Harston, N. Rane, G. Shaya, S. Thandeswaran, M. Cellerini, F. Sheerin, and J. Kennedy. Imaging biomarkers in acute ischemic stroke trials: a systematic review. *American Journal of Neuroradiology*, 36(5):839–843, 2015.

[8] S. Siemonsen, K. Mouridsen, B. Holst, T. Ries, J. Finsterbusch, G. Thomalla, L. Ostergaard, and J. Fiehler. Quantitative T2 values predict time from symptom onset in acute stroke patients. *Stroke*, 40(5):1612–1616, 2009.

[9] B. L. McGarry, H. J. Rogers, M. J. Knight, K. T. Jokivarsi, A. Sierra, O. H. Gröhn, and R. A. Kauppinen. Stroke onset time estimation from multispectral quantitative magnetic resonance imaging in a rat model of focal permanent cerebral ischemia. *International Journal of Stroke*, 11(6):677–682, mar 2016.

[10] D. B. Plewes and W. Kucharczyk. Physics of MRI: A primer. *Journal of Magnetic Resonance Imaging*, 35(5):1038–1054, 2012.

[11] R. W. Brown, E. M. Haacke, Y-C. N. Cheng, M. R. Thompson, and R. Venkatesan. Magnetic resonance imaging: physical principles and sequence design. 1999.

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

[13] D. G. Nishimura. *Principles of Magnetic Resonance Imaging*. Lulu, 1996.

[14] M. Cercignani, N. G. Dowell, and P. S. Tofts. *Quantitative MRI of the Brain*. Wiley Online Library, 2nd edition, 2018.

[15] R. G. González. Current state of acute stroke imaging. *Stroke*, 44(11):3260–3264, 2013.

[16] P. S. Tofts. *Quantitative MRI of the Brain: Measuring Changes Caused by Disease*. John Wiley & Sons, 2003.

[17] S. Volz, U. Nöth, and R. Deichmann. Correction of systematic errors in quantitative proton density mapping. *Magnetic Resonance in Medicine*, 68(1):74–85, 2012.

[18] Z. Abbas, V. Gras, K. Möllenhoff, A-M. Oros-Peusquens, and N. J. Shah. Quantitative water content mapping at clinically relevant field strengths: A comparative study at 1.5T and 3T. *NeuroImage*, 106:404–13, 2015.

[19] A. P. Crawley and R. M. Henkelman. A comparison of oneshot and recovery methods in T1 imaging. *Magnetic Resonance in Medicine*, 7(1):23–34, 1988.

[20] J. Homer and M. S. Beevers. Driven-equilibrium single-pulse observation of T1 relaxation. A re-evaluation of a rapid "new" method for determining NMR

spin-lattice relaxation times. *Journal of Magnetic Resonance*, 63(2):287–297, 1985.

[21] S. C. L. Deoni, B. K. Rutt, and T. M. Peters. Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Magnetic Resonance in Medicine*, 49(3):515–526, 2003.

[22] S. C. L. Deoni, T. M. Peters, and B. K. Rutt. High-resolution T1 and T2 mapping of the brain in a clinically acceptable time with DESPOT1 and DESPOT2. *Magnetic Resonance in Medicine*, 53(1):237–241, 2005.

[23] K. S. Nayak and D. G. Nishimura. Automatic Field Map Generation and Off-Resonance Correction for Projection Reconstruction Imaging. *Magnetic Resonance in Medicine*, 43(1):151–154, 2000.

[24] G. R. Morrell and M. C. Schabel. An analysis of the accuracy of magnetic resonance flip angle measurement methods. *Physics in Medicine and Biology*, 55(20):6157–74, 2010.

[25] V. L. Yarnykh. Actual flip-angle imaging in the pulsed steady state: A method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magnetic Resonance in Medicine*, 57(1):192–200, 2007.

[26] K. Nehrke and P. Börnet. DREAM - a novel approach for robust, ultrafast, multislice B1 mapping. *Magnetic Resonance in Medicine*, 68(5):1517–1526, 2012.

[27] D. Brenner, D. Tse, E. Pracht, T. Feiweier, R. Stirnberg, and T. Stöcker. 3DREAM - A Three-Dimensional Variant of the DREAM Sequence. In *Proceedings of the 22nd Annual Meeting of ISMRM*, volume 22, page 1455, 2014.

[28] G. R. Morrell. A phase-sensitive method of flip angle mapping. *Magnetic Resonance in Medicine*, 60(4):889–894, 2008.

[29] J. B. M. Warntjes, O. Dahlqvist, and P. Lundberg. Novel method for rapid, simultaneous T1,T2*, and proton density quantification. *Magnetic Resonance in Medicine*, 57(3):528–537, 2007.

[30] M. Warntjes, O. D. Leinhard, J. West, and P. Lundberg. Rapid magnetic resonance quantification on the brain: Optimization for clinical usage. *Magnetic Resonance in Medicine*, 60(2):320–329, 2008.

[31] N. Weiskopf, J. Suckling, G. Williams, M. M. Correia, B. Inkster, R. Tait, C. Ooi, E. T. Bullmore, and A. Lutti. Quantitative multi-parameter mapping of R1, PD*, MT, and R2* at 3T: a multi-center validation. *Frontiers in Neuroscience*, 7(95):1–11, 2013.

[32] G. Palma, E. Tedeschi, P. Borrelli, S. Cocozza, C. Russo, S. Liu, Y. Ye, M. Comerci, B. Alfano, M. Salvatore, E. M. Haacke, and M. Mancini. A novel multiparametric approach to 3D quantitative MRI of the brain. *PLoS ONE*, 10(8):e0134963, 2015.

[33] G. H. Welsch, K. Scheffler, T. C. Mamisch, T. Hughes, S. Millington, M. Deimling, and S. Trattnig. Rapid estimation of cartilage T2 based on double echo at steady state (DESS) with 3 Tesla. *Magnetic Resonance in Medicine*, 62(2):544–549, 2009.

[34] R. Heule, C. Ganter, and O. Bieri. Triple echo steady-state (TESS) relaxometry. *Magnetic Resonance in Medicine*, 71(1):230–237, 2014.

[35] R. Heule, Z. Celicanin, S. Kozerke, and O. Bieri. Simultaneous multislice triple-echo steady-state (SMS-TESS) T1, T2, PD, and off-resonance mapping in the human brain. *Magnetic Resonance in Medicine*, 80(3):1088–1100, 2018.

[36] N. Stikov, M. Boudreau, I. R. Levesque, C. L. Tardif, J. K. Barral, and G. B. Pike. On the accuracy of T1 mapping: Searching for common ground. *Magnetic Resonance in Medicine*, 73(2):514–522, 2015.

[37] Y. Jiang, D. Ma, N. Seiberlich, V. Gulani, and M. A. Griswold. MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout. *Magnetic Resonance in Medicine*, 74(6):1621–1631, 2014.

[38] J. Asslander, S. J. Glaser, and J. Hennig. Pseudo Steady-State Free Precession for MR-Fingerprinting. *Magnetic Resonance in Medicine*, 77(3):1151–1161, 2017.

[39] D. Ma, S. Coppo, Y. Chen, D. F. McGivney, Y. Jiang, S. Pahwa, V. Gulani, and M. A. Griswold. Slice profile and B1 corrections in 2D magnetic resonance fingerprinting. *Magnetic Resonance in Medicine*, 78(5):1781–1789, 2017.

[40] T. Christen, N. A. Pannetier, W. W. Ni, D. Qiu, M. E. Moseley, N. Schuff, and G. Zaharchuk. MR vascular fingerprinting: A new approach to compute cerebral blood volume, mean vessel radius, and oxygenation maps in the human brain. *NeuroImage*, 89:262–270, 2014.

[41] B. Rieger, M. Akçakaya, L. R. Schad, and S. Weingärtner. Simultaneous quantification of T1, T2 and Apparent Diffusion Coefficient using Magnetic Resonance Fingerprinting base on Echo Planar Imaging. In *Proceedings of the 26th Annual Meeting of ISMRM*, page 1023, 2018.

[42] P. Su, D. Mao, P. Liu, Y. Li, M. C. Pinho, B. G. Welch, and H. Lu. Multiparametric estimation of brain hemodynamics with MR fingerprinting ASL. *Magnetic Resonance in Medicine*, 78(5):1812–1823, 2017.

[43] C. E. Anderson, S. B. Donnola, Y. Jiang, J. Batesole, R. Darrah, M. L. Drumm, S. M. Brady-Kalnay, N. F. Steinmetz, X. Yu, M. A. Griswold, and C. A. Flask. Dual Contrast - Magnetic Resonance Fingerprinting (DC-MRF): A Platform for Simultaneous Quantification of Multiple MRI Contrast Agents. *Scientific Reports*, 7(1):8431, 2017.

[44] T. Hong, D. Han, and D-H. Kim. Simultaneous estimation of PD, T1, T2, T2*, and B0 using magnetic resonance fingerprinting with background gradient compensation. *Magnetic Resonance in Medicine*, 81(4):2614–2623, 2019.

[45] A. Deshmane, D. Ma, Y. Jiang, E. Fisher, N. Seiberlich, V. Gulani, and M. Griswold. Validation of Tissue Characterization in Mixed Voxels Using MR Fingerprinting. In *Proceedings of the 22nd Annual Meeting of ISMRM*, page 94, 2014.

[46] S. Tang, C. Fernandez-Granda, S. Lannuzel, B. Bernstein, R. Lattanzi, M. Cloos, F. Knoll, and J. Assländer. Multicompartment Magnetic Resonance Fingerprinting. *Inverse Problems*, 34(9):094005, 2018.

[47] D. F. Mcgivney, D. Ma, Y. Jiang, A. Deshmane, V. Gulani, N. Seiberlich, and M. A. Griswold. Partial volume mapping using magnetic resonance fingerprinting. *NMR in Biomedicine*, 32(5):e4082, 2019.

[48] G. Buonincontri and S. J. Sawiak. MR fingerprinting with simultaneous B1 estimation. *Magnetic Resonance in Medicine*, 76(4):1127–35, 2016.

[49] G. Buonincontri, R. F. Schulte, M. Cosottini, and M. Tosetti. Spiral MR fingerprinting at 7 T with simultaneous B1 estimation. *Magnetic Resonance Imaging*, 41:1–6, 2017.

[50] T. Hong, D. Han, M-O. Kim, and D-H. Kim. RF slice profile effects in magnetic resonance fingerprinting. *Magnetic Resonance Imaging*, 41:73–79, 2017.

[51] S-C. Chiu, T-M. Lin, J-M. Lin, H-W. Chung, C-W. Ko, M. Büchert, and M. Bock. Effects of RF pulse profile and intra-voxel phase dispersion on MR fingerprinting with balanced SSFP readout. *Magnetic Resonance Imaging*, 41:80–86, 2017.

[52] D. F. McGivney, E. Pierre, D. Ma, Y. Jiang, H. Saybasili, V. Gulani, and M. A. Griswold. SVD compression for magnetic resonance fingerprinting in the time domain. *IEEE Transactions on Medical Imaging*, 33(12):2311–2322, 2014.

[53] S. F. Cauley, K. Setsompop, D. Ma, Y. Jiang, H. Ye, E. Adalsteinsson, M. A. Griswold, and L. L. Wald. Fast group matching for MR fingerprinting reconstruction. *Magnetic Resonance in Medicine*, 74(2):523–528, 2015.

[54] H. Ye, D. Ma, Y. Jiang, S. F. Cauley, Y. Du, L. L. Wald, M. A. Griswold, and K. Setsompop. Accelerating magnetic resonance fingerprinting (MRF) using t-blipped simultaneous multislice (SMS) acquisition. *Magnetic Resonance in Medicine*, 75(5):2078–85, 2016.

[55] H. Ye, S. F. Cauley, B. Gagoski, B. Bilgic, D. Ma, Y. Jiang, Y. P. Du, M. A. Griswold, L. L. Wald, and K. Setsompop. Simultaneous multislice magnetic resonance fingerprinting (SMS-MRF) with direct-spiral slice-GRAPPA (ds-SG) reconstruction. *Magnetic Resonance in Medicine*, 77(5):1966–1974, 2017.

[56] M. A. Cloos, F. Knoll, T. Zhao, K. Block, M. Bruno, C. G. Wiggins, and D. K. Sodickson. Multiparametric imaging with heterogenous radiofrequency fields. *Nature Communication*, 7:12445, 2016.

[57] Y. Gao, Y. Chen, D. Ma, Y. Jiang, K. A. Herrmann, J. A. Vincent, K. M. Dell, M. L. Drumm, S. M. Brady-Kalnay, M. A. Griswold, C. A. Flask, and L. Lu. Preclinical MR fingerprinting (MRF) at 7 T: Effective quantitative imaging for rodent disease models. *NMR in Biomedicine*, 28(3):384–394, 2015.

[58] C. E. Anderson, C. Y. Wang, Y. Gu, R. Darrah, M. A. Griswold, X. Yu, and C. A. Flask. Regularly incremented phase encoding - MR fingerprinting (RIPE-MRF) for enhanced motion artifact suppression in preclinical cartesian MR fingerprinting. *Magnetic Resonance in Medicine*, 79(4):2176–2182, 2018.

[59] B. Rieger, F. Zimmer, J. Zapp, S. Weingärtner, and L. R. Schad. Magnetic resonance fingerprinting using echo-planar imaging: Joint quantification of

T1 and T2* relaxation times. *Magnetic Resonance in Medicine*, 78(5):1724–1733, 2017.

[60] B. Zhao, K. Setsompop, H. Ye, S. Cauley, and L. L Wald. Maximum Likelihood Reconstruction for Magnetic Resonance Fingerprinting. *IEEE Trans Med Imaging*, 35(8):1812–23, 2016.

[61] Z. Wang, H. Li, Q. Zhang, J. Yuan, and X. Wang. Magnetic Resonance Fingerprinting with compressed sensing and distance metric learning. *Neurocomputing*, 174:560–570, 2016.

[62] E. Y. Pierre, D. Ma, Y. Chen, C. Badve, and M. A. Griswold. Multiscale reconstruction for MR fingerprinting. *Magnetic Resonance in Medicine*, 75(6):2481–2492, 2016.

[63] T. Amthor, M. Doneva, P. Koken, J. Keupp, and P. Börnert. Comparison of Different Approaches of Pattern Matching for MR Fingerprinting. In *Proceedings of the 23rd Annual Meeting of ISMRM*, page 3394, 2015.

[64] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.

[65] O. Cohen, B. Zhu, and M. S. Rosen. Deep learning for fast MR Fingerprinting Reconstruction. In *International Society for Magnetic Resonance in Medicine (ISMRM)*, page 688, 2017.

[66] E. Hoppe, G. Körzdörfer, T. Würfl, J. Wetzl, F. Lugauer, J. Pfeuffer, and A. Maier. Deep learning for magnetic resonance fingerprinting: A new ap-

proach for predicting quantitative parameter values from time series. *GMDS*, pages 202–206, 2017.

[67] O. Cohen and M. S. Rosen. Algorithm comparison for schedule optimization in MR fingerprinting. *Magnetic Resonance Imaging*, 41:15–21, 2017.

[68] J. I. Hamilton, K. L. Wright, Y. Jiang, L. Hernandez-Garcia, D. Ma, M. Griswold, and N. Sieberlich. Pulse sequence optimization for improved MRF scan efficiency. In *Proceedings of the 23rd Annual Meeting of ISMRM*, volume 23, page 3386, 2015.

[69] K. Sommer, T. Amthor, M. Doneva, P. Koken, J. Meineke, and P. Börnert. Towards predicting the encoding capability of MR fingerprinting sequences. *Magnetic Resonance Imaging*, 41:7–14, 2017.

[70] M. Weigel. Extended phase graphs: Dephasing, RF pulses, and echoes - Pure and simple. *Journal of Magnetic Resonance Imaging*, 41(2):266–295, 2015.

[71] J. Z. Bojorquez, S. Bricq, C. Acquitter, F. Brunotte, P. M. Walker, and A. Lalande. What are normal relaxation times of tissues at 3 Tesla? *Magnetic Resonance Imaging*, 35:69–80, 2017.

[72] R. C. Brand, N. P. Blockley, M. Chappell, and P. Jezzard. Clinically relevant rapid 3D CEST imaging with hexagonal spoiling gradients, optimised B1, and symmetric z-spectrum sampling. In *Proceedings of the 24th Annual Meeting of ISMRM*, page 299, 2016.

[73] D. C. Zhu and R. D. Penn. Full-brain T1 mapping through inversion recovery fast spin echo imaging with time-efficient slice ordering. *Magnetic Resonance in Medicine*, 54(3):725–731, 2005.

[74] L. Chen and M. Bernstein. Measurements of T1 relaxation times at 3.0 T: implications for clinical MRA. In *Proceedings of the 9th Annual Meeting of ISMRM*, volume 9, page 1391, 2001.

[75] M. A. Bernstein, J. Huston, and H. A. Ward. Imaging artifacts at 3.0T. *Journal of Magnetic Resonance Imaging*, 24(4):735–746, 2006.

[76] University of Wisconsin - Madison. QMAP: Quantitative MRI Analysis Package. http://www.medphysics.wisc.edu/˜samsonov/qmap/, 2011.

[77] S. Clare and P. Jezzard. Rapid T1 mapping using multislice echo planar imaging. *Magnetic Resonance in Medicine*, 45(4):630–634, 2001.

[78] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995.

[79] T. He, P. D. Gatehouse, G. C. Smith, R. H. Mohiaddin, D. J. Pennell, and D. N. Firmin. Myocardial T2* Measurements in Iron Overloaded Thalassemia: An In Vivo Study to Investigate the Optimal Methods of Quantification. *Magnetic Resonance in Medicine*, 60(5):1082–1089, 2008.

[80] A. J Miller and P. M. Joseph. The Use of Power Images to Perform Quantitative Analysis on Low SNR MR Images. *Magnetic Resonance Imaging*, 11(7):1051–1056, 1993.

[81] D. Milford, N. Rosbach, M. Bendszus, and S. Heiland. Mono-exponential fitting in T2-relaxometry: Relevance of offset and first echo. *PLoS ONE*, 10(12):e0145255, 2015.

[82] G. J. Stanisz, E. E. Odrobina, J. Pun, M. Escaravage, S. J. Graham, M. J. Bronskill, and R. M. Henkelman. T1, T2 relaxation and magnetization transfer in tissue at 3T. *Magnetic Resonance in Medicine*, 54(3):507–512, 2005.

[83] T. Stöcker and E. D. Pracht. Whole brain B1 mapping using three dimensional DREAM. *Magnetic Resonance in Medicine*, 82(3):924–934, 2019.

[84] D. Ma. *Magnetic Resonance Fingerprinting*. PhD thesis, Case Western Reserve University, 2015.

[85] L. Leroi. *Quantitative MRI : towards fast and reliable T1 , T2 and proton density mapping at ultra-high field*. PhD thesis, Université Paris-Saclay, 2018.

[86] J. H. Duyn, Y. Yang, J. A. Frank, and J. W. van Der Veen. Simple correction method for k-space trajectory deviations in MRI. *Journal of Magnetic Resonance*, 132(1):150–153, 1998.

[87] J. A. Fessler. Michigan Image Reconstruction Toolbox, 2019.

[88] J. G. Pipe and P. Menon. Sampling Density Compensation in MRI : Rationale and an Iterative Numerical Solution. *Magnetic Resonance in Medicine*, 41(1):179–186, 1999.

[89] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin. Adaptive reconstruction of phased array MR imagery. *Magnetic Resonance in Medicine*, 43(5):682–690, 2000.

[90] G. S. Payne and M. O. Leach. Implementation and evaluation of frequency offset corrected inversion (FOCI) pulses on a clinical MR system. *Magnetic Resonance in Medicine*, 38(5):828–833, 1997.

[91] A. T. Hess and M. D. Robson. Hexagonal gradient scheme with RF spoiling improves spoiling performance for high-flip-angle fast gradient echo imaging. *Magnetic Resonance in Medicine*, 77(3):1231–1237, 2017.

[92] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT press, 1998.

[93] J. McCall. Genetic algorithms for modelling and optimisation. *Journal of Computational and Applied Mathematics*, 184(1):205–222, 2005.

[94] R. L. Haupt. An introduction to genetic algorithms for electromagnetics. *IEEE Antennas and Propagation Magazine*, 37(2):7–15, 1995.

[95] E. S. Yoshimaru, E. A. Randtke, M. D. Pagel, and J. Cárdenas-Rodríguez. Design and optimization of pulsed Chemical Exchange Saturation Transfer MRI using a multiobjective genetic algorithm. *Journal of Magnetic Resonance*, 263:184–192, 2016.

[96] E. M. Haacke, E. D. Lindskogj, and W. Lin. A Fast , Iterative , Partial-Fourier Technique Capable of Local Phase Recovery. *Journal of Magnetic Resonance*, 92(1):126–145, 1991.

[97] D. C. Noll, D. G. Nishimura, A. Macovski, and A. Detection. Homodyne Detection in Magnetic Resonance Imaging. *IEEE Trans Med Imaging*, 10(2):154–163, 1991.

[98] K. Setsompop, B. A. Gagoski, J. R. Polimeni, T. Witzel, V. J. Wedeen, and L. L. Wald. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic Resonance in Medicine*, 67(5):1210–1224, 2012.

Appendix A

Fitting Inversion Pulse Efficiency

This section investigates the inversion efficiency of the inversion pulse used in the SE EPI MRF experiments, with results from modelling this efficiency in the dictionary.

A.1 Fitting Inversion Pulse Efficiency: Introduction

Ideally the optional slice-selective inversion pulses before any given excitation pulse will invert the longitudinal magnetisation (M_z) along the slice direction. In practice the inversion efficiency depends on the design of the pulse.

By adapting the acquisition of the SE EPI MRF sequence a traditional inversion recovery experiment can be performed, where the acquired signal varies as a function of Inversion Time (TI). By performing this experiment with various inversion slice thicknesses and fitting a T_1 relaxation curve to the data it is possible to determine the inversion slice thickness needed to achieve complete inversion of

the longitudinal magnetisation. In the case of a complete inversion the fitted relaxation curve following the inversion will start at $-M_z(t = 0)$. If time is allowed for full relaxation before each measurement $M_z(t = 0)$ will equal $M_z(t = \infty)$.

A.2 Fitting Inversion Pulse Efficiency: Methods

Inversion Efficiency In this thesis inversion efficiency is defined as a fraction between 0 and 1, where an efficiency of 1 corresponds to a perfect inversion and an efficiency of 0 has no effect on M_z. An efficiency of 0.5 would represent a saturation, where M_z becomes zero. The effect of the inversion pulse is as follows: $M_z' = M_z - 2\beta M_z$, where β is the efficiency and M_z' is M_z after the inversion. By including this expression in the signal simulation it is possible to model the signal evolution in the case of $\beta < 1$. The parameter β can be assumed to be constant or varied through a dimension of the matching dictionary. The latter approach allows flexibility and potentially better fitting but increases the dictionary size, potentially making it impractically large in cases when the dictionary contains dimensions for many other types of dictionary parameters.

Inversion Recovery Using a schedule with variable TRs, an inversion recovery experiment was performed, with the custom phantom from Chapter 3 as the scanned object. The T_1 relaxation model $|M_z| = |M_0(1 - 2\beta \exp(-TI/T_1))|$ was then fitted to the data to obtain the inversion efficiency, in the form of β, for voxels in each of the six phantom tubes. Acquisitions were performed with various inversion slice thicknesses, as percentages of the imaging slice thickness (5mm).

SE EPI MRF Parameter Maps MRF parameters maps were obtained by fitting a dictionary to the data acquired using a 50-timepoint acquisition schedule optimised in Chapter 4 to reduce combined T_1/T_2 RMSE. Note the objective of producing these maps was not to compare the parameter matching error with maps from unoptimised schedules, but to compare the effects of including or omitting inversion efficiency from the fit. The inversion pulse slice thickness was set to 6mm (i.e. 120% of the imaging slice thickness).

First, the case of a perfect inversion was assumed (i.e. $\beta = 1$). As a comparison, matching was also performed with the inclusion of an inversion efficiency dictionary dimension. The B_1^+ efficiency was set to 1 in the matching process, meaning there was not a B_1^+ dimension and B_1^+ variations were not fitted. Examples of fitting B_1^+ during the SE EPI MRF matching process are explored in subsequent sections of this chapter. The β values in the inversion efficiency dimension of the dictionary were 0.5 to 1.0, in increments of 0.05.

A.3 Fitting Inversion Pulse Efficiency: Results and Discussion

Inversion Recovery Figure A.1 shows fitted curves for the phantom data acquired using an acquisition schedule design to mimic an inversion recovery experiment. The acquired signal is plotted for a range of inversion times and inversion pulse slice thicknesses.

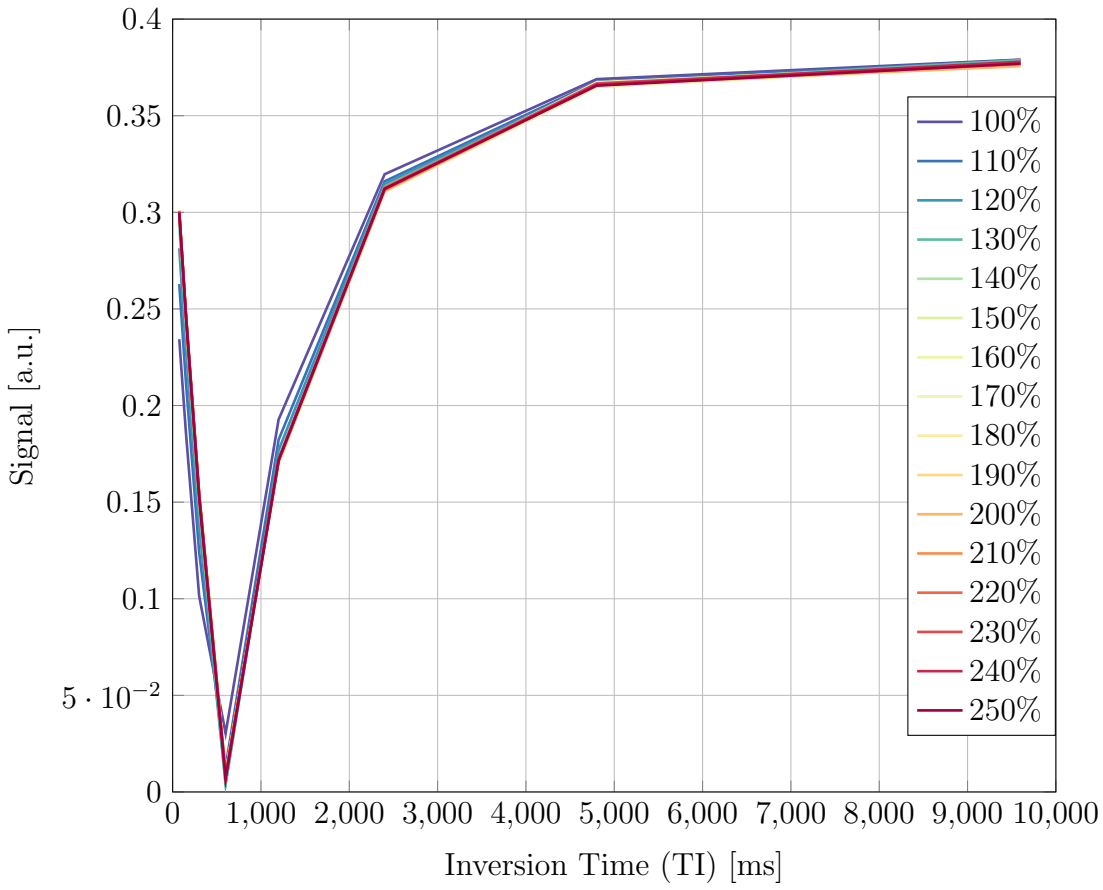

(a)

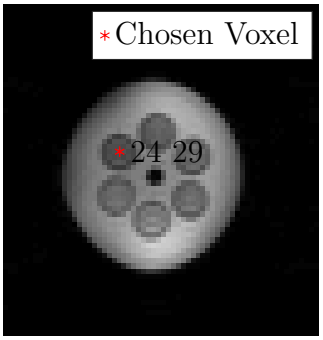

(b)

Figure A.1: The signal (a) from one voxel (b) in an inversion recovery experiment for a range of inversion slice thicknesses. The inversion slice thickness is represented as a percentage of the 5mm imaging slice thickness.

Figure A.2a shows the same data as Fig. A.1 but only the beginning of the signal curve, thus highlighting the differences in signal across the different inversion slice thicknesses. Figure A.2b plots the first data point for each of the curves in Fig. A.2a, suggesting that the minimum inversion thickness needed in order to achieve full inversion was approximately 140% of the imaging slice thickness (i.e. 1.4×5mm = 7mm). However, it must be noted that the inversion slice thickness threshold is independent of the imaging slice thickness.

The β values obtained from fitting the longitudinal relaxation model are shown in Fig. A.3. The numbered lines in Fig. A.3a correspond to the voxel labels in Fig. A.3b. The threshold inversion efficiency in Fig. A.3a is consistent with that of Fig. A.2b (i.e. ≈140%).

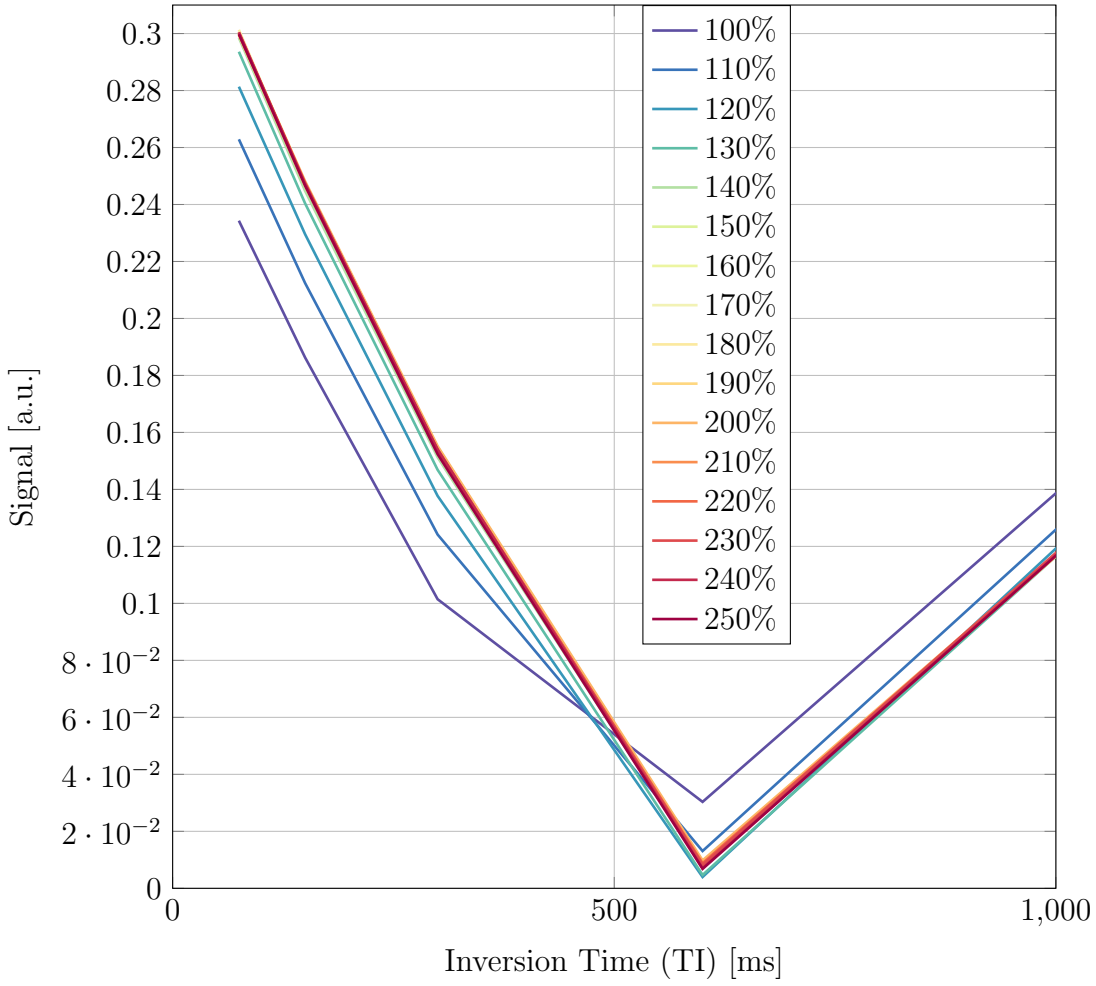

(a)

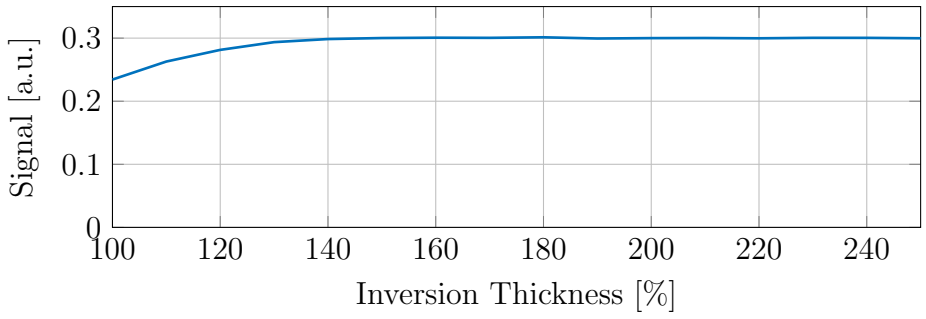

(b)

Figure A.2: The beginning of the inversion recovery curves (a) in Fig. A.1 and a plot of the first data point (b) of each data curve in Fig. A.1.

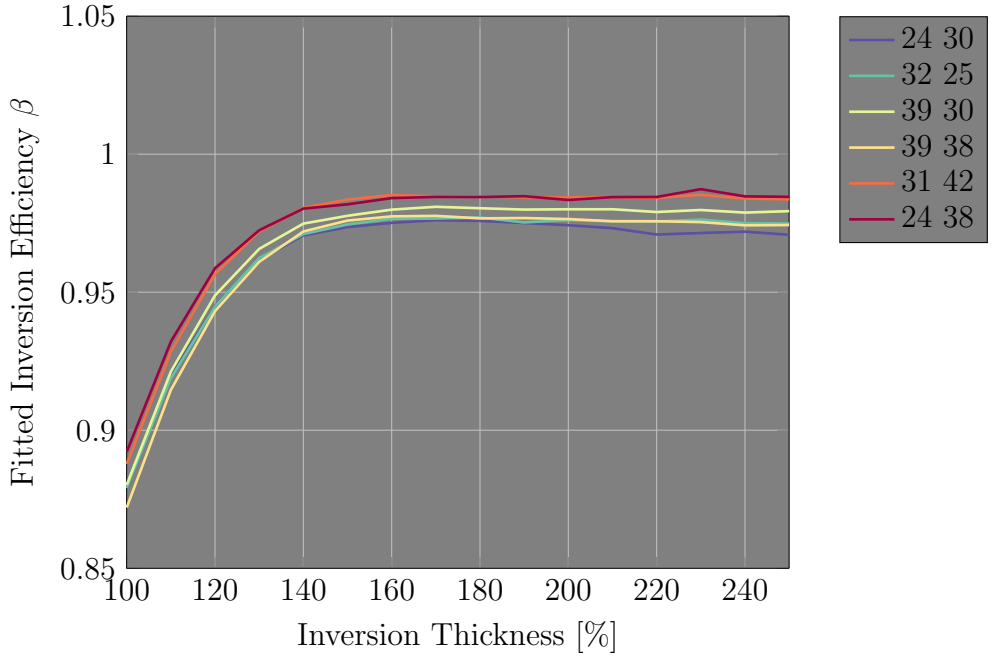

(a)

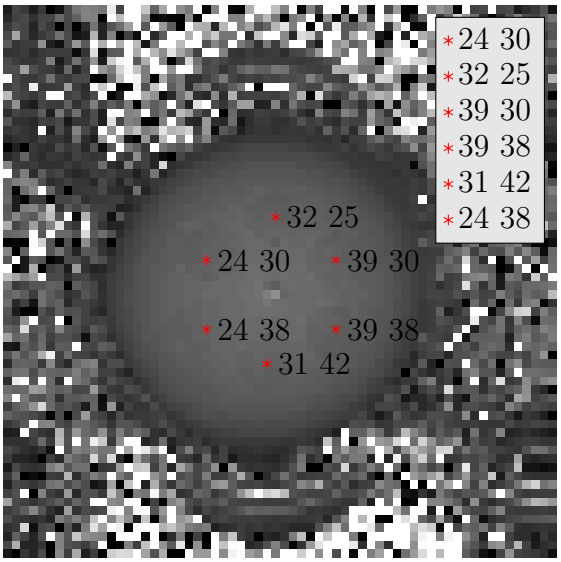

(b)

Figure A.3: Fitted inversion efficiency β values (a) from gold standard T_1 relaxation model fits, using the data used in Fig. A.2b, for a sample voxel (b) in each of the phantom tubes.

Figures A.4 and A.5 show T_1 maps obtained with MRF dictionaries without and with an inversion efficiency dimension, respectively. The slice indices start at the top left image and increase from left to right. For example, Slice 5 is represented by the map at Column 1 and Row 2 in each montage. Since the acquisition was designed such that each slice experienced the same acquisition schedule one would expect identical parameter maps for Slices 1 to 7, for identical phantom composition at each slice position. Given the slice group positioning the tube contents should indeed be the same across the slices. There is consistency across the T_1 maps for Slices 1, 2, 3 and 4 but the T_1 maps for Slices 5, 6 and 7 can be deemed incorrect as they have no resemblance to the gold standard from Chapter 3. This suggests that Slices 5, 6 and 7 did not actually experience the same schedule as the other slices. Appendix B discusses a possible reason for this in more detail. Including a β dictionary dimension allowed reasonable T_1 fits for Slices 5 and 6 and reasonable T_2 fits for Slices 5, 6 and 7, suggesting that the β parameter compensated for this effect during the matching process.

As seen in the T_1 results, the T_2 maps without inversion efficiency fitting were poor for Slices 5 and 6 in the absence of inversion efficiency fitting (Fig. A.6), with observable improvement as a result of the introduction of the β dimension (Fig. A.7).

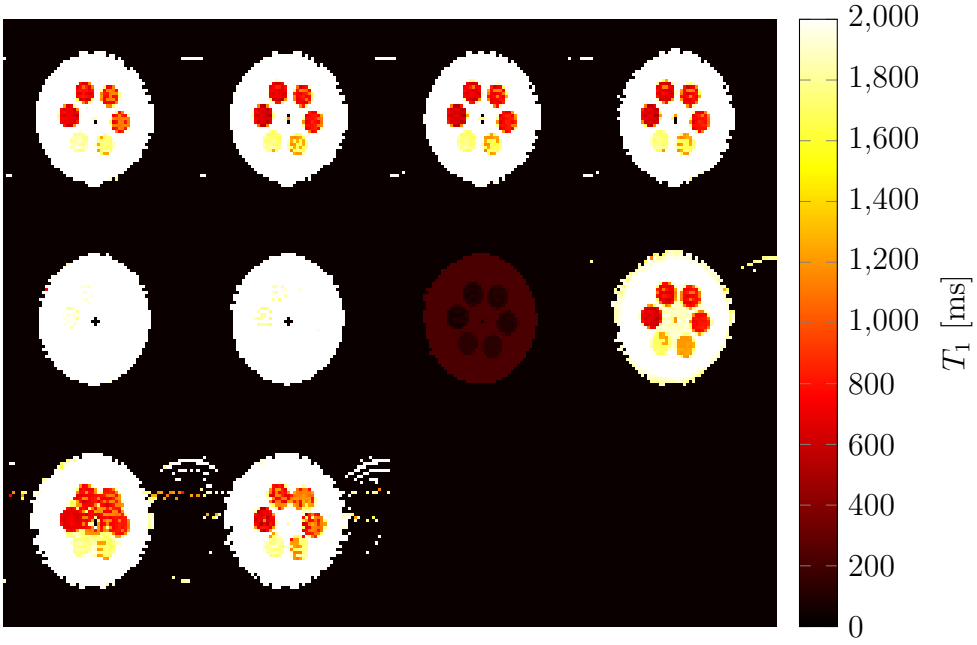

Figure A.4: SE EPI MRF without fitting for inversion efficiency: T_1.

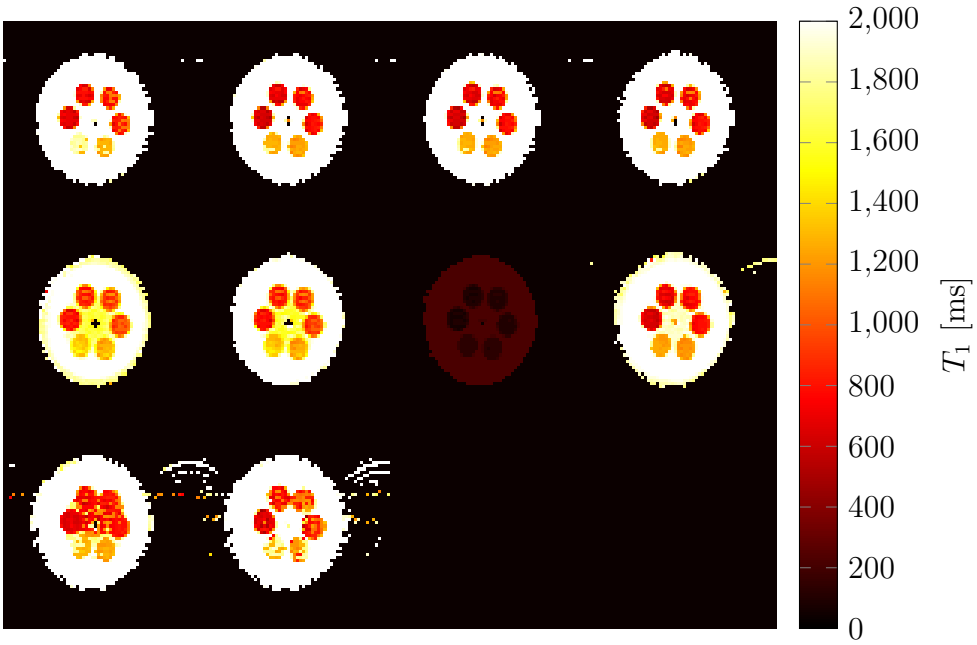

Figure A.5: SE EPI MRF fitting for inversion efficiency: T_1.

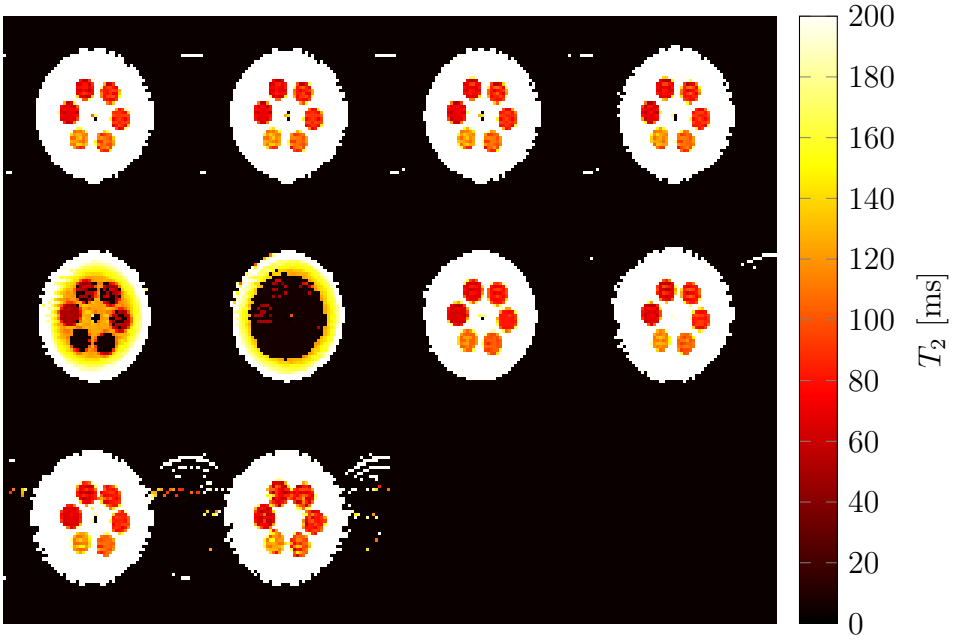

Figure A.6: SE EPI MRF without fitting for inversion efficiency: T_2.

Figure A.7: SE EPI MRF with fitting for inversion efficiency: T_2

It is likely that the differences in the T_1 and T_2 maps of Slices 9 and 10, compared to Slices 1 to 4, were caused by differences in the phantom structure. These structural differences are illustrated in the M_0 maps fitted without (Fig. A.8) and with (Fig. A.9) an inversion efficiency dimension. Slices 9 and 10 were in close proximity to the internal supporting structure of the phantom at the superior end of the 6 tubes, whereas Slices 1 to 4 sampled the tubes in a more inferior position where the structure is different. In particular, in Slice 9 the region of relatively low M_0 between the tubes was caused by the disc-shaped support. Slice 10 displays inhomogeneity within the tubes, presumably caused by bubbles formed during the production process since this slice position is near the tube lids.

Good consistency is visible between the M_0 maps produced without (Fig. A.8) and with (Fig. A.9) inversion efficiency fitting. There are visible differences between the inferior slices (Slices 1, 2, 3 and 4) and the central slices (Slices 5, 6 and 7), particularly in the case without inversion efficiency fitting.

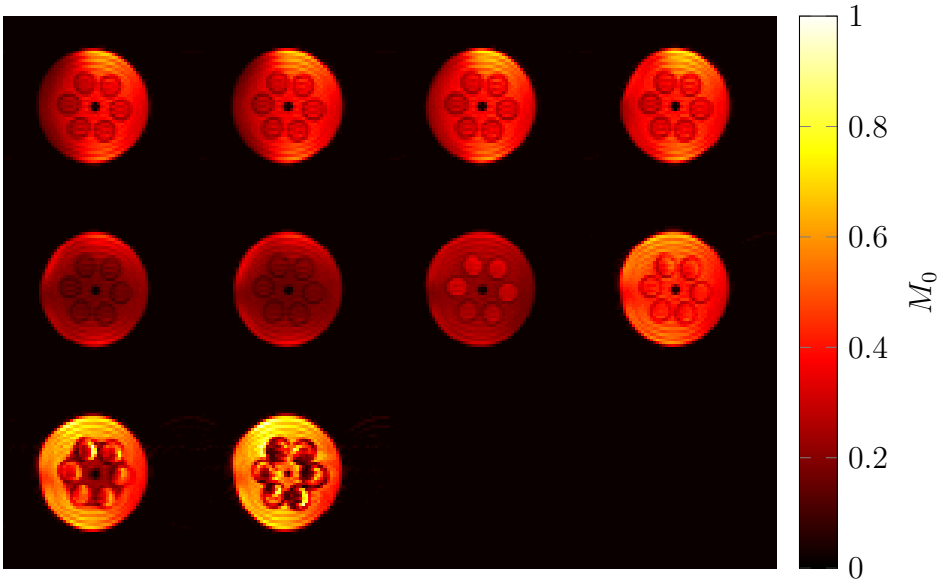

Figure A.8: SE EPI MRF without fitting for inversion efficiency: M_0

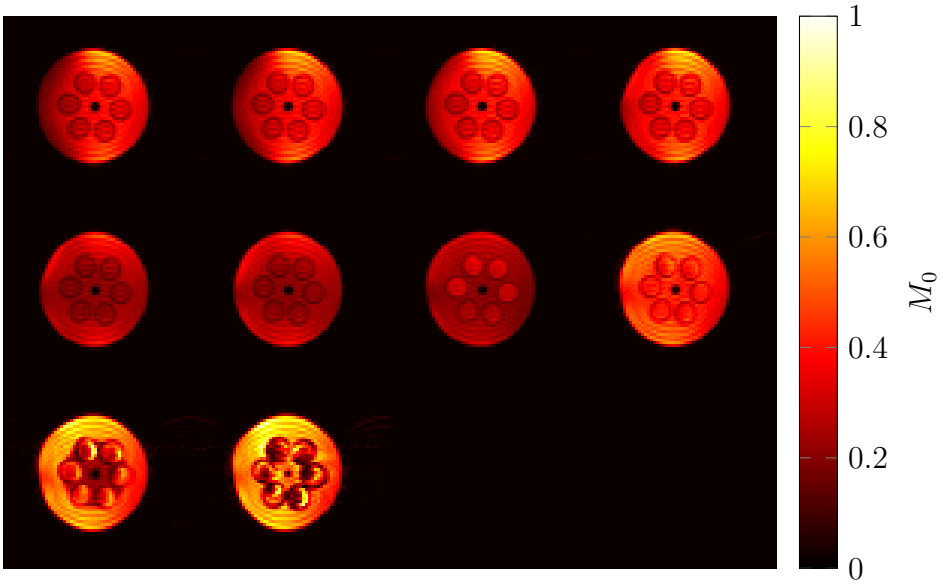

Figure A.9: SE EPI MRF with fitting for inversion efficiency: M_0

The SE EPI MRF inversion efficiency maps are shown in Fig. A.10. In accordance with the trend observed for the T_1, T_2 and M_0 results the fitted inversion efficiency maps for slices 5 and 6 are notably different to the other slices. In particular, their values are considerably lower than the other slices, at 0.6 to 0.7. In the tubes of the remaining slices the efficiency fits are approximately 0.95. Importantly, this is consistent with the efficiency obtained from the T_1 model fits in Fig. A.2b for an inversion slice thickness of 120% of the imaging slice thickness.

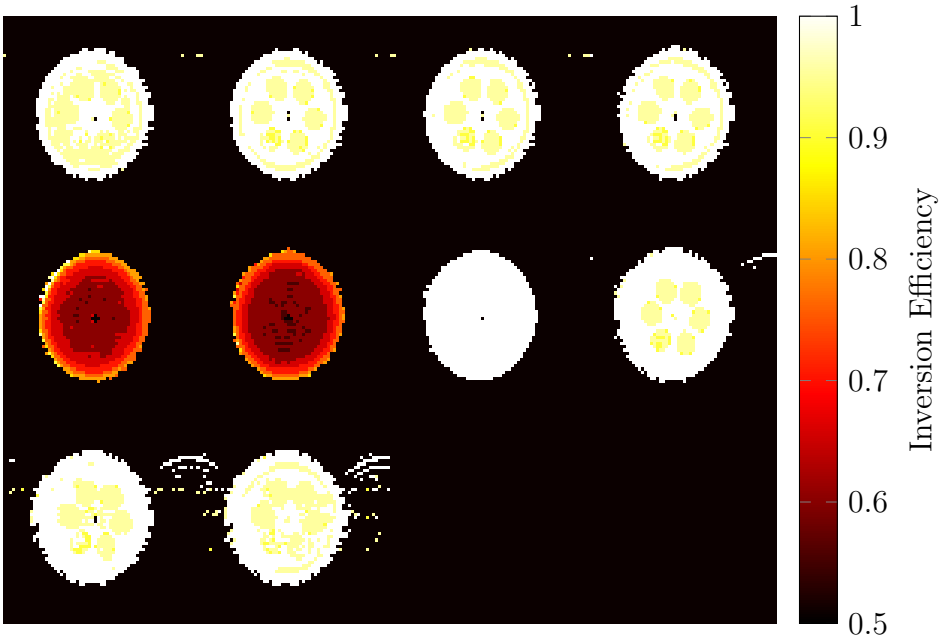

Figure A.10: SE EPI MRF fitted inversion efficiency β

A.4 Fitting Inversion Pulse Efficiency: Conclusions

The efficiency of the inversion pulse of the SE EPI MRF pulse sequence was examined and found to require a slice thickness of at least ≈140% of the imaging slice thickness to achieve maximum inversion. The maximum inversion measured across the phantom tubes was ≈98%. An inversion thickness of 120% (i.e. 6mm for a 5mm imaging slice) provided ≈95% inversion, making it a reasonable compromise for cases where the slice gap must be thin compared to the imaging slice thickness.

This section also investigated the effects of inversion efficiency on the SE EPI MRF parameter maps. Without an inversion efficiency dictionary dimension cen-

tral slices of SE EPI MRF acquisition produced erroneous T_1 and T_2 maps. This was partially corrected when an efficiency dimension was included, suggesting that the erroneous maps could have been caused by imperfect magnetisation inversions. In the inferior and superior slices the MRF-fitted phantom tube inversion efficiencies were similar to the fitted efficiency of the preceding experiments ($\approx$95% inversion), further evidence that the matching process was sensitive to inversion efficiency during the matching process. This means an inversion efficiency could potentially be used in cases with reduced slice gaps, where reduced relative inversion thicknesses would be required in order to avoid cross-talk between slices.

The results in this section show that incomplete inversions can occur during the SE EPI MRF acquisition but that the consequential effect on the signal time courses can be mitigated by fitting for inversion efficiency during the dictionary matching process. Accounting for this deviation improves the T_1 and T_2 parameter maps.

Appendix B

Effect of Slice Position

This section addresses the anomalous maps observed in the central slices for the SE EPI MRF implementation and explores a possible reason for this phenomenon by shifting the slice positions with respect to the phantom.

B.1 Effect of Slice Position: Introduction

Section A explored the effects of fitting for inversion efficiency within the MRF framework. However, notable differences were observed in the parameter maps of slices near the centre of the phantom compared to inferior slices, suggesting that the magnetisation of the central slices was not perturbed in the same way as the inferior slices. Although the inclusion of an inversion efficiency dimension in the dictionary mitigated this effect somewhat it was unclear whether the cause was a pulse sequence error or a location-specific interaction between the pulses, the hardware and the phantom. This section explores whether the erroneous maps were dependent on the slice positions relative to the phantom.

B.2 Effect of Slice Position: Methods

Two sets of 10-slice data were acquired, each with a different shift along the slice direction: 1) with the slice group centred at isocentre and 2) with the slice group shifted -15mm away from isocentre in the slice direction. The slice direction corresponded to the long axis of the phantom, with a negative shift representing a shift towards the inferior end of the phantom. The imaging slice thickness was 5mm, with slice gaps of 50% (i.e. 2.5mm), meaning the shift of -15mm was equivalent to a shift of 2 slice positions. The resulting parameter maps from the two acquisitions were compared.

B.3 Effect of Slice Position: Results and Discussion

Fitted M_0 Maps Figures B.1a and B.1b shows the fitted M_0 maps without and with a -15mm shift, respectively. Evidence of the slice shift is provided by Slices 9 and 10 in both sets of maps. At isocentre (Fig. B.1a) Slices 9 and 10 exhibit signs of the phantom structure described in Appendix A, whereas the shifted slices do not show this.

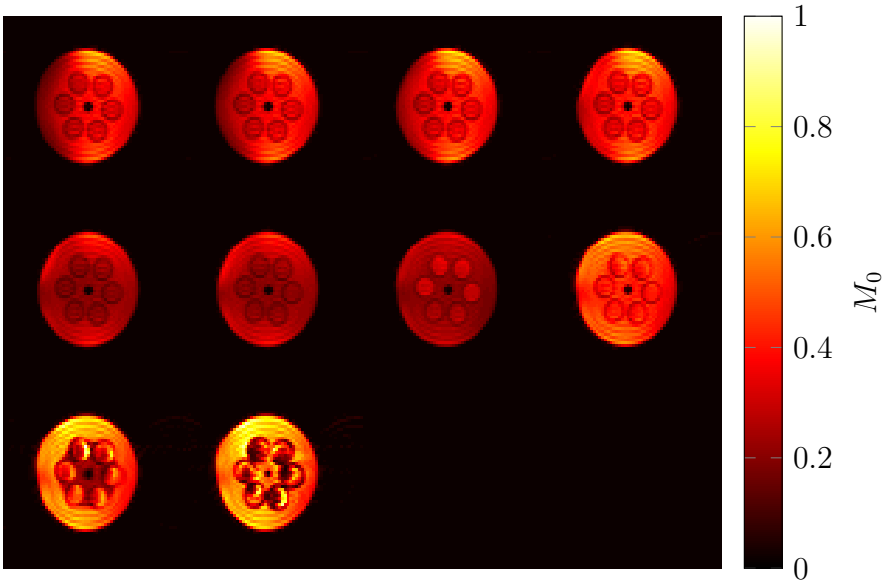

(a)

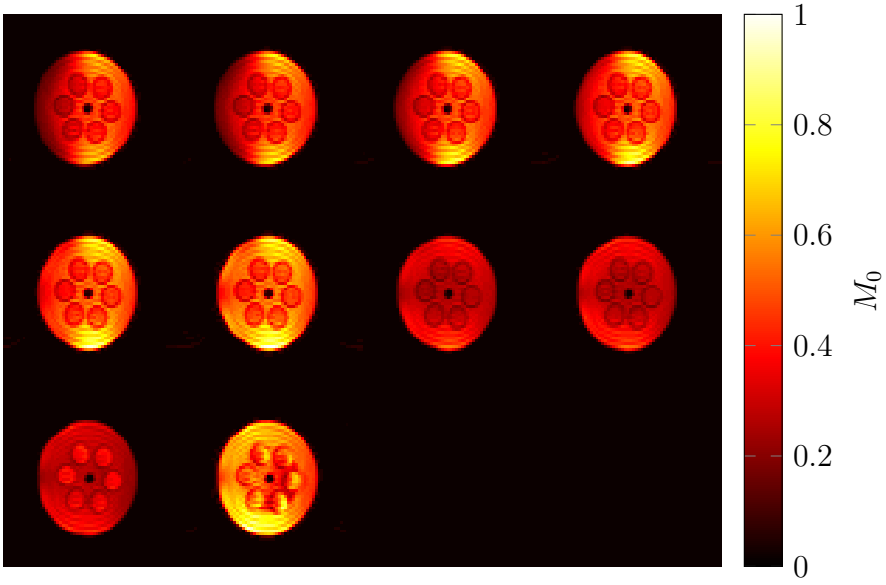

(b)

Figure B.1: SE EPI MRF fitted M_0 for slice groups at isocentre (a) and shifted -15mm in the slice direction (b).

Fitted T_1 and Inversion Efficiency Maps Figure B.2 shows fitted SE EPI MRF T_1 maps at isocentre and with the slice shift of -15mm, both with an inversion

efficiency dimension. The maps without the shift (i.e. those acquired at isocentre) are from Appendix A. For the maps acquired at isocentre (Fig. B.2a) Slice 7 shows notably lower T_1 values than the other slices, whereas in the case of the -15mm shift (Fig. B.2b) Slice 9 is the anomaly, meaning the index of the erroneous slice increased by 2 between the two sets of data. This result is in agreement with the imposed slice shift of -15mm. As mentioned in Appendix B.2, the shift of -15mm corresponded to a shift of 2 slice positions. Similar to the T_1 maps, the slices with erroneous inversion efficiency maps (Fig. B.3) are shifted by 2 slice indices as a result of the acquisition shift. The fitted T_2 maps are not shown here because the T_2 maps of the central slices had less pronounced differences than the T_1 an M_0 maps, meaning the effect of shifting the slices was less obvious.

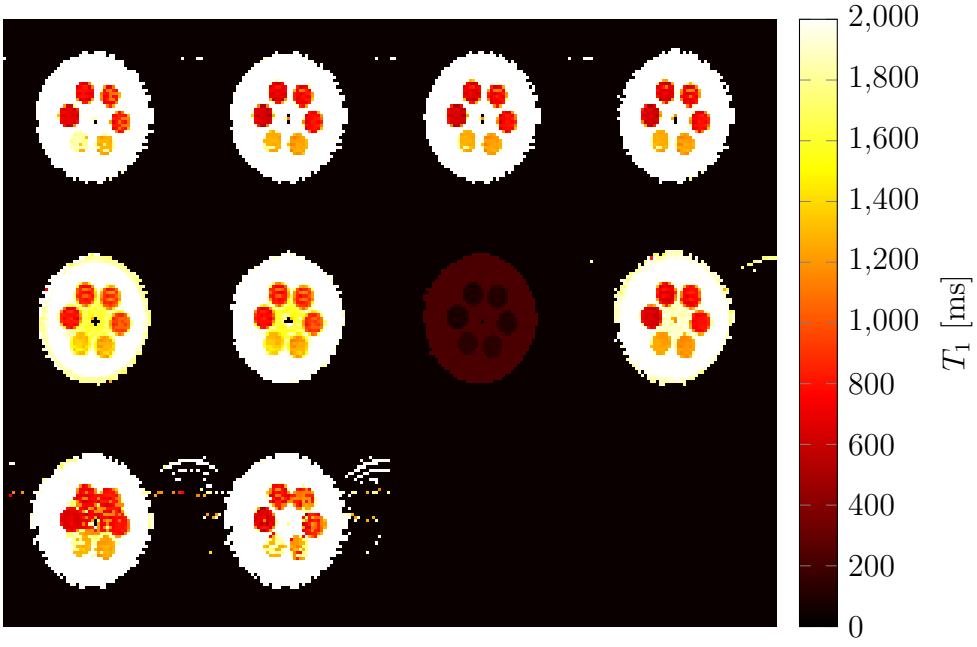

(a)

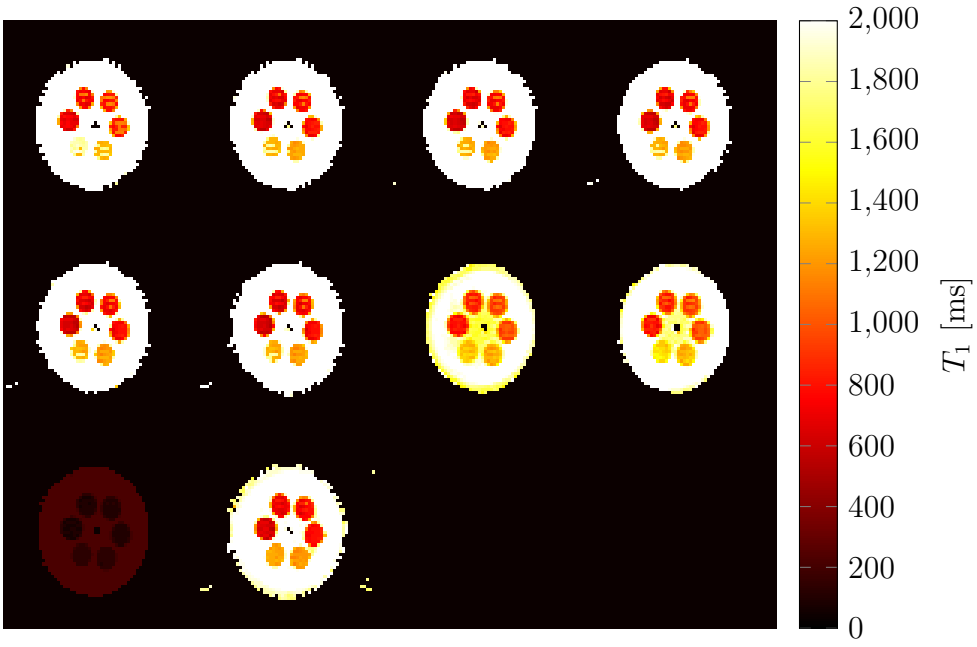

(b)

Figure B.2: SE EPI MRF fitted T_1 for a slice group at isocentre (a) and shifted -15mm in the slice direction (b).

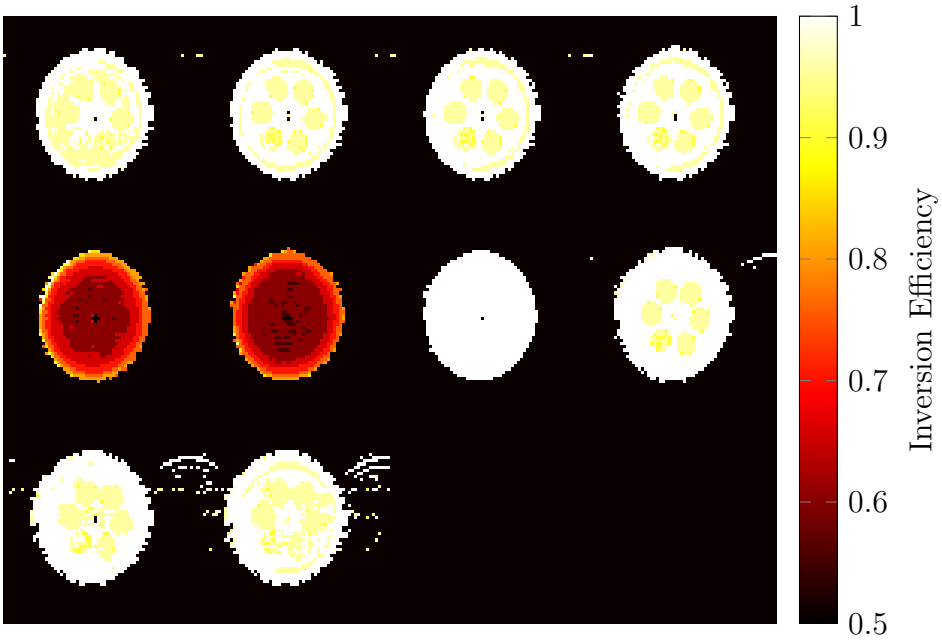

(a)

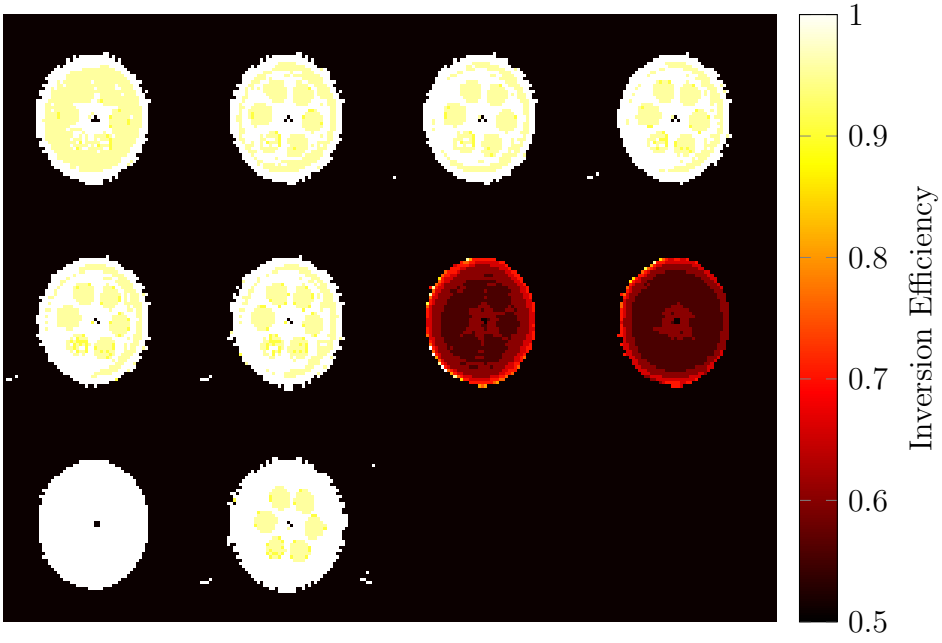

(b)

Figure B.3: SE EPI MRF fitted inversion efficiency for slice groups at isocentre (a) and shifted -15mm in the slice direction (b).

B.4 Effect of Slice Position: Conclusions

In this appendix, SE EPI MRF parameter maps were fitted to data from two slice groups at different positions along the slice direction to explore the reason for the erroneous maps observed in Sections 5.4 and 5.5. The results in this section demonstrated that the erroneous maps occurred at a fixed slice shift relative to the phantom, providing evidence that the inconsistency in the magnetisation perturbation across the slices was due to an interaction between the RF fields and the structure of the phantom, not a pulse sequence error. However, the exact reason for the inconsistency remains unclear.