Magnetisation transfer in the context of multi-component MR fingerprinting

Sven Kwinkelenberg

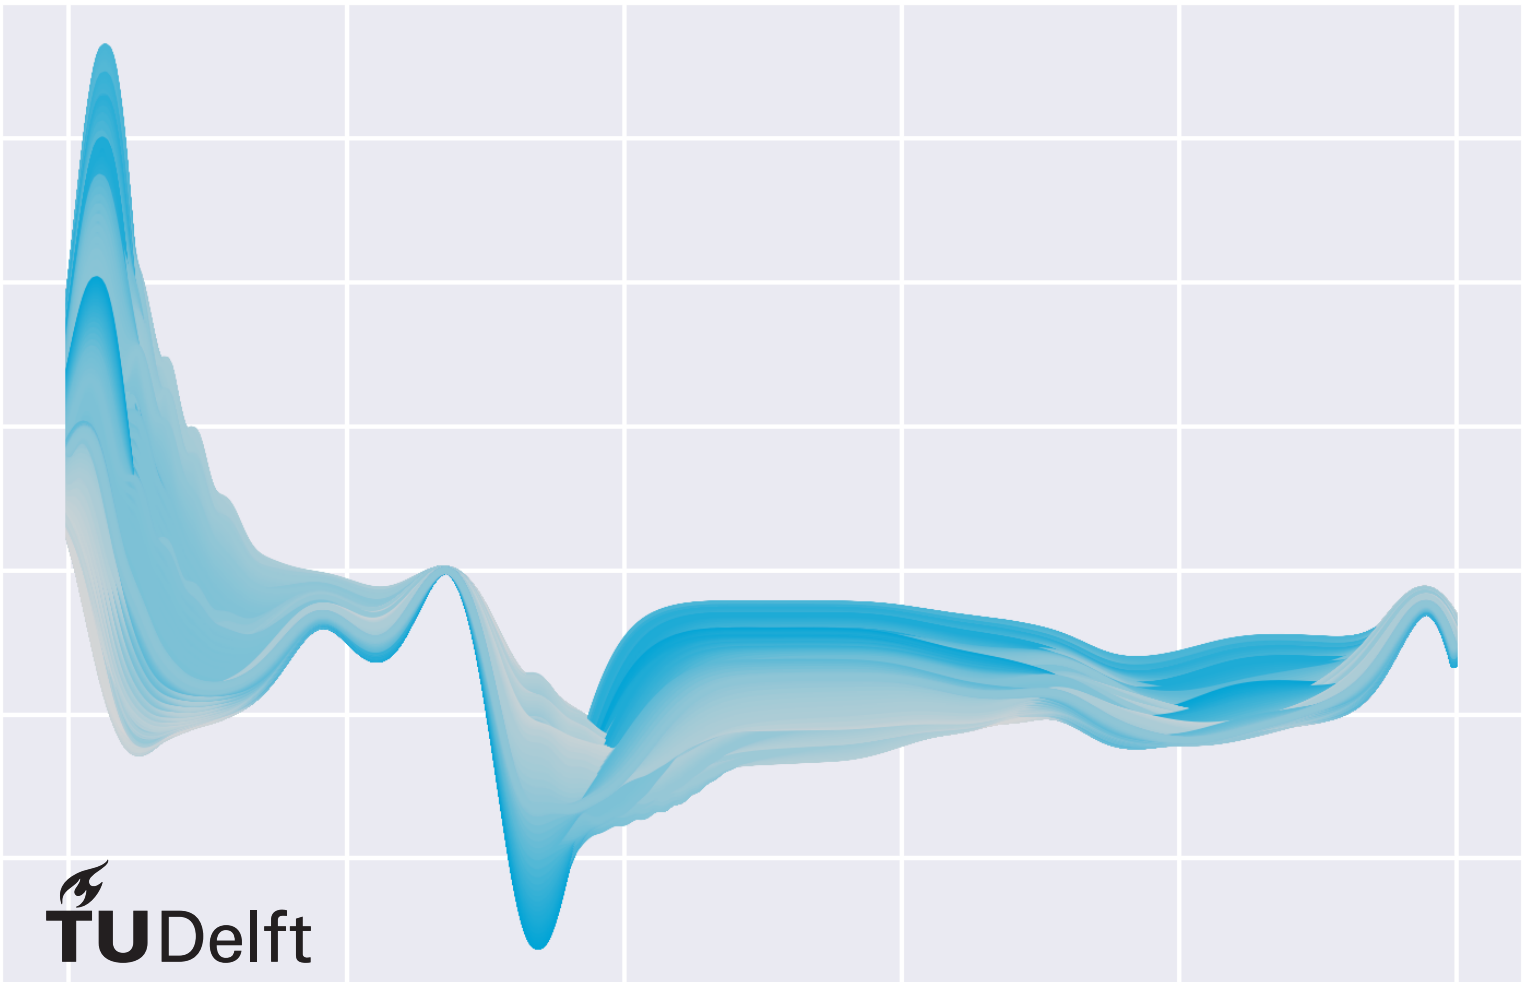

TU Delft

Magnetisation transfer in the context of multi-component MR fingerprinting

by

Sven Kwinkelenberg

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday May 23, 2023 at 3:00 PM.

Student number:	4719522	
Project duration:	Septempber 1, 2022 – May 23, 2023	
Thesis committee:	Dr. F. M. Vos,	TU Delft, supervisor
	Dr. Ir. M. A. Nagtegaal,	TU Delft, daily supervisor
	Dr. S. D. Weingärtner,	TU Delft
	Dr. Ir. M. Goorden,	TU Delft
External advisor:	Dr. Ir. D. H. J. Poot,	Erasmus MC

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

Abstract

Magnetic resonance fingerprinting is an MRI-based technique that allows for fast, simultaneous quantitative mapping of multiple tissue parameters. Multi-component MRF (MC-MRF) additionally allows for the mapping of multiple tissue components per voxel. Currently, most MC-MRF implementations ignore any nonlinear effects on signals resulting from multi-component systems, such as the effects introduced by magnetisation transfer (MT). Here, we investigate the effects of free pool to free pool magnetisation transfer on the accuracy of MC-MRF, with a focus on the application of this technique for myelin water fraction (MWF) imaging. Assessment of the different MC-MRF techniques is done through application of these algorithms to two-component numerical phantoms, where the signals from two interacting components were simulated using the EPG-X framework. Results show that MT has a negative effect on the accuracy of the acquired parameter estimates for all estimated parameters, resulting in biases and incorrect parameter estimates. Several adjusted methods for MC-MRF including magnetisation transfer were proposed and tested. Although theoretically improvements were expected, the used sequences showed to be inadequate for accurate MT estimates. More research into these techniques is still required to improve their performance and accuracy. A technique left mainly unexplored here is sequence optimisation to minimise the effects of magnetisation transfer on the resulting signals. A quick exploration, however, showed that this might be a viable approach for future research as well.

Contents

1 Introduction

Magnetic resonance imaging (MRI) is a widely used, non-invasive medical imaging technique, that sustains high resolution images (up to 1 mm resolution) of the body with high contrast. Contrary to e.g. X-ray imaging or PET scans, MRI does not expose the subject to ionising radiation. The specific imaging settings used during the acquisition of an MRI image allow the operator to tailor the image contrast. This combination of high resolution, low risk, and great flexibility makes MRI an invaluable tool in both clinical practice and medical research [1].

MRI is used in many fields where good soft-tissue contrast is important. Examples of this are cardiac imaging [2], abdominal imaging [3], and neuro imaging [4]. Especially in neuro imaging, the ability to choose the contrast of the images in advance and acquire different images with varying contrasts, allows for a good differentiation between tissue types such as white and grey matter. One branch of MRI neuro imaging is *myelin water imaging* [5]. In the brain, the neural pathways (axons) are wrapped in layers of myelin, a fatty tissue that insulates the axons from each other and increases the signal transmission speeds [5]. In between the myelin layers, small amounts of water are trapped, called the *myelin water*. Although the MRI signal obtained from lipids generally is very weak, the signal obtained from water is much stronger. This allows myelin water to be measured as an indirect substitute for myelin [6].

An active area of research for MRI is in the field of *multiple sclerosis* (MS) [7]. For MS, one of the main symptoms is *demyelination*, a decrease in the myelin content of the brain [8]. The reduction in myelin decreases or even stops signal transmission in the affected regions [7]. The myelin content of the brain is therefore an important parameter in MS research and diagnosis. Myelin water fraction (MWF) imaging is a technique that aims to determine the myelin water content of the brain as a fraction of the total tissue volume. In areas affected by demyelination, the myelin content will decrease, and with it the MWF. Therefore, determining the MWF can be used as a metric for research and early diagnosis of MS [9].

Different techniques to quantify the MWF through MRI already exist. Methods like multi-echo spin echo multi-component T_2 mapping aim to acquire quantitative T_2-relaxation time maps for multiple components in the brain. Determining the contribution of myelin water to the overall signal provides information about the MWF [10]. Similarly, multi-compartment Look-Locker utilises a quantitative multi-component T_1 mapping to estimate the myelin water fraction [11]. Furthermore, techniques like mcDESPOT aim to map both T_1 and T_2 simultaneously, and extract the MWF from this information [12].

This research focuses on *magnetic resonance fingerprinting* (MRF) [13] for the estimation of the myelin water fraction. MRF works by inducing a transient-state signal through application of a specialised pulse sequence. This transient-state signal is then compared to a dictionary of simulated signals, allowing for the simultaneous extraction of multiple tissue parameters [13]. The simultaneous mapping of multiple parameters prevents the need for multiple separate quantitative quantitative scans. This in turn decreases the total scan time, lowering the probability of motion artefacts. These factors make MRF an attractive technique for quantitative MRI.

Multi-component MRF (MC-MRF) aims to simultaneously extract tissue parameters for multiple components in a voxel, potentially allowing for the estimation of the MWF. Here, we explore the effects of *magnetisation transfer* [14] on the accuracy of the MWF estimates acquired through multi-component MRF. More specifically, we aim to answer two main questions: first, what is the extent of the effect of magnetisation transfer on the performance of the multi-component models? And secondly, assuming that magnetisation transfer introduces biases or inaccuracies, how can we include information about the exchange in the estimation models to get more accurate and reliable results?

The next chapter presents the relevant theory to explain the workings and challenges of this method for MWF mapping. Following that, the methods we used to answer the research questions are explained. In the remaining chapters, the results are shown, discussed, and a conclusion is drawn.

2

Theory

This chapter provides explanations of the basic theory used in this report.

2.1. Magnetic Resonance Imaging

In this report, we assume a general understanding of the mechanisms and techniques used in MRI. However, for the sake of completeness, some of the core concepts are briefly introduced here. A more complete and in-depth description of MRI can be found in [1].

In MRI, the subject is placed in a strong magnetic field called the B_0 field. Modern clinical MRI scanners typically use field strengths of 1.5 T or 3 T [1], while 7 T clinical scanners have been emerging more recently [15]. Conventionally, the B_0 field is taken to be in the z-direction. ($\boldsymbol{B_0} = B_0\hat{z}$). This strong field causes protons, which have a nonzero microscopic magnetisation, to either align with or against the B_0 field. A small majority aligns with the external field, inducing a macroscopic *magnetisation* of the subject. Due to their relative abundance in biological tissue and high sensitivity, this magnetisation is mainly the result of hydrogen protons. A second magnetic field perpendicular to the main field is used to change the direction of the magnetisation. This secondary field, or B_1-field, uses radio frequency (RF) pulses to *flip* the magnetisation around a certain axis. The strength and duration of this pulse result in a certain *flip angle* with the z-direction.

After an RF pulse, the magnetisation can consist of components both along and perpendicular to the B_0 field. The perpendicular components experience a torque from the main field. This results in a rotating motion around the z-axis, with a frequency given by $\gamma \times B_0$, in which γ is *gyromagnetic ratio*. The frequency with which the magnetisation rotates around the z-axis is the same RF frequency as the RF pulse that induces the flip. Generally, the magnetisation is described in the *rotating frame of reference*, a frame rotating with the same frequency as the transverse magnetisation. This allows us to write the magnetisation as two components:

$$\boldsymbol{M} = \boldsymbol{M_z} + \boldsymbol{M_{xy}}, \tag{2.1}$$

where $\boldsymbol{M_z}$ and $\boldsymbol{M_{xy}}$ are the *longitudinal* and *transverse* components respectively. Since the transverse magnetisation is rotating around the z-axis, this constitutes a time-varying magnetic field. This in turn induces an RF field that can be picked up by coils positioned around the subject. The magnitude of $\boldsymbol{M_{xy}}$ determines the strength of the induced field, with stronger fields for larger $\boldsymbol{M_{xy}}$. Since the frequency and phase of the measured signals depend on the strength of the external magnetic field, applying *gradients* to this field allows us to encode spatial information in the frequency and relative phase of the signals. MRI readouts sample from the so-called k-space, an analogue to the Fourier frequency domain. Sampling sufficient information from k-space allows us to reconstruct images through a transformation similar to the inverse Fourier transform. *Undersampling* the k-space will result in aliasing artefacts, and thus an inaccurate image reconstruction.

Relaxation effects

When in equilibrium, the magnetisation is aligned with the B_0 field. After excitation, the magnetisation will always return to its equilibrium state in the longitudinal direction. This process, called *longitudinal relaxation*, is characterised by an exponential with time constant T_1, called the longitudinal *relaxation time*. The

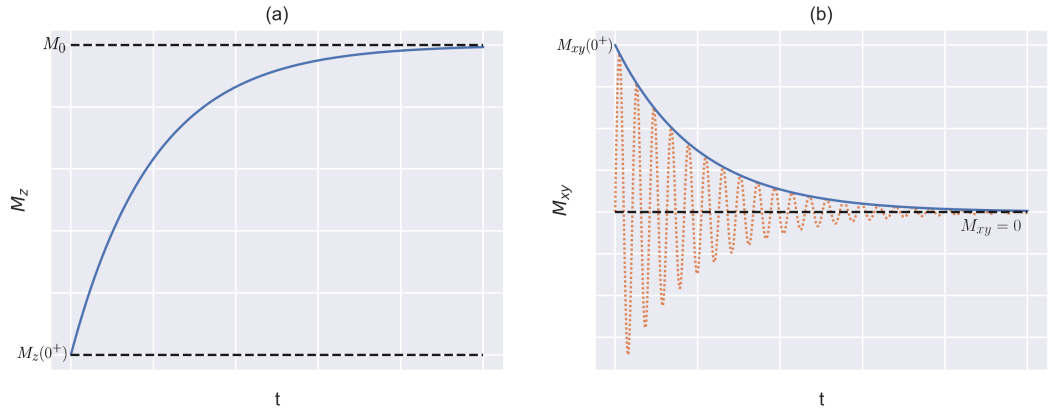

Figure 2.1: Illustrations of the effect of transverse and longitudinal relaxation. Figure (a) shows the results of longitudinal relaxation on the longitudinal component of the magnetisation. $M_z(0^+)$ is the longitudinal magnetisation directly after an RF pulse is applied, and M_0 is the equilibrium magnetisation. In (b) the effects of transverse relaxation are shown, where $M_{xy}(0^+)$ is the transverse magnetisation directly after an RF pulse. The blue line represent the evolution of $\boldsymbol{M_{xy}}$, while the orange dotted line represents the the time-evolution of one of the transverse components of the magnetisation.

relaxation time is a tissue parameter.

Simultaneously with longitudinal relaxation, the transverse magnetisation experiences *transverse relaxation*. As a result of interactions between individual atoms and inhomogeneities in the magnetic field, the spinning hydrogen protons *dephase* over time. This results in a decrease of the net transverse magnetisation over time. Similar to longitudinal relaxation, this effect can be modelled using an exponential with time constant T_2. This *transverse relaxation time* is also a tissue parameter.

The two relaxation parameters, together with the *proton density* (which determines the magnitude of the equilibrium magnetisation), mostly determine the contrast between tissues in MRI images.

Quantitative imaging

The contrast of conventional clinical MRI images is based on differences in the measured signal strength. Different areas will produce different signal strengths based on their relaxation parameters and proton density. A radiologist is then able to distinguish between different tissues based on the contrast in the images. Conventional MRI is a *qualitative* imaging technique: the images show differences in signal strength, but this provides no information on the values of the underlying tissue parameters. Particularly, scans made on machines of different vendors, even using very similar imaging protocols, can result in different intensities for the same tissues [16]. The lack of a clear reference makes relating different scan results to each other only more difficult. For clinical practice, this is not a problem. The main focus there is on anatomical imaging, for which generally only a clear contrast between relevant tissues is needed. However, in medical research, *quantitative* results are invaluable for objectively comparing different scans to each other, and observing changes in values over time [16].

Currently, there already exist different MRI-based techniques for quantitative measurement. For example, the multi-echo spin-echo method , or a Look-Locker scheme for quantitative T_1 and T_2 measurement respectively [9]. Another technique is DESPOT, which is able to measure both relaxation parameters simultaneously [17]. An important measure in assessing the usability of these quantitative technique is *scan time*. Generally, quantitative measurements need long scan times. This increases the probability of movement artefacts in the final parameter maps, and drastically decreases patient comfort. It is therefore desirable to keep the scan time as low as possible without sacrificing accuracy.

2.2. Magnetic Resonance Fingerprinting

The main focus of this report will be on *magnetic resonance fingerprinting* (MRF) [13]. MRF is a quantitative MRI technique that aims to create multiple parameter maps *simultaneously*, while also being both fast and

accurate [18]. MRF hinges on two main imaging principles: efficient *parameter encoding* and *undersampling*.

2.2.1. Parameter encoding

Conventional MRI generally produces only one, or a few images. These images are acquired, aiming to have high contrast and resolution, allowing the physician to interpret them accurately. MRF, on the other hand, produces a *series* of images from a transient state signal. This transient state is obtained by continuously changing the imaging settings between readouts of a series. This assures that the signal never reaches a steady state. Generally for MRF, the acquisition parameter that is varied is the flip angle, resulting in a *flip angle series*. Advanced signal models allow the simulation of signals. As such, the tissue parameters are *encoded* in the simulated MRF signals. An example of a flip angle series, together with three example signals resulting from this series, is shown in Figure 2.2. [19].

Not shown in this figure is the use of an *inversion pulse*. This is an RF pulse with a flip angle of 180°, often at the start of a flip angle train. The inclusion of an inversion pulse allows for a good simultaneous encoding of the T_1 and T_2 times in the resulting signals [13].

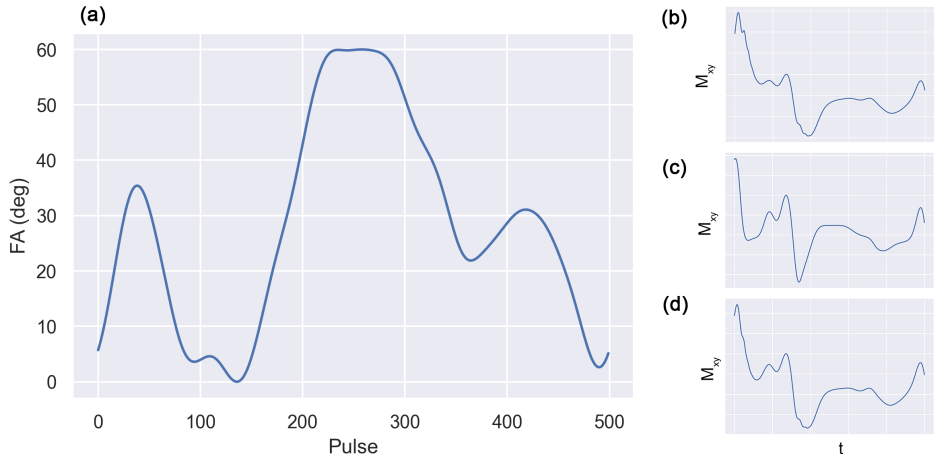

Figure 2.2: (a) The flip angle pattern of the gradient spoiled SSFP pulse sequence. The inversion pulse of 180° at the start of the pulse series is not shown. Along with the flip angle series, the simulated signals for (b) grey matter, (c) white matter, and (d) myelin water are shown. The relaxation parameters for these tissue types are given in Table 3.1.

2.2.2. Parameter matching

One way to retrieve the encoded parameters is through the use of a *signal dictionary*. This is a large collection of precomputed signals for a wide range of tissue parameter combinations. Comparing the measured signal to all dictionary signals allows for the retrieval of the tissue parameters.

First, we define the measured signal $\boldsymbol{s}$ as a series of M measurements s_j, with M the length of the pulse series:

$$\boldsymbol{s} = [s_1\ s_2\ \cdots\ s_M]\,, \tag{2.2}$$

and the signal dictionary D as

$$D = [\boldsymbol{d_1}\ \boldsymbol{d_2}\ \cdots\ \boldsymbol{d_N}]\,. \tag{2.3}$$

Here, each dictionary atom $\boldsymbol{d_i}$ is a vector of length M. The conventional measure of similarity between the signal and a dictionary atom is the *inner product*. Finding the best match in the dictionary for the measured signal now becomes:

$$\arg\max_i\ |\mathbf{d_i} \cdot \boldsymbol{s}|\,. \tag{2.4}$$

The relaxation parameters are retrieved from the best matching dictionary entry (or *atom*). This process allows MRF to make maps for multiple parameters simultaneously. This prevents errors from confounding factors between separate parameter measurements, as can be the case with separate T_1 and T_2 measurements.

2.2.3. Signal simulation

In order to create a signal dictionary, the expected signals for the different parameter combinations have to be simulated in advance. There are two main methods for simulating MRF signals. The first is through direct simulation of a large ensemble of spins [20][21]. The behaviour of the magnetisation during an MRI scan is described by the *Bloch equations*. *Bloch simulations* directly model these equations, simulating the system behaviour this way. The other simulation method is called *extended phase graphs* (EPG) [22]. It is a Fourier-based technique that allows for relatively simple and fast signal simulation.

In the EPG framework, magnetisation is described through so-called *configuration states*. For each small volume V, there is a set of configurations states that fully describe the magnetisation in that volume. The dephasing of the magnetisation as a result of the application of gradients is modelled implicitly through a *dephasing coordinate*. In addition to gradients, the configuration states are manipulated through *RF pulses* and relaxation effects.

For a given volume V, the set of configuration states comprises three separate states: $F_+(k)$, $F_-(k)$, and $F_z(k)$. $F_+(k)$ and $F_-(k)$ describe the behaviour of dephasing and re-phasing *transverse* magnetisation respectively, while $F_z(k)$ describes the *longitudinal* magnetisation. The parameter k is called the *dephasing coordinate*, and is defined as

$$\boldsymbol{k}(t) = \int_0^t \boldsymbol{G}(t')\,dt', \tag{2.5}$$

where $\boldsymbol{G}(t)$ is the gradient in the magnetic field applied to the volume V at time t. As gradients in the magnetic field can be in any direction, the dephasing coordinate in (2.5) is described as a vector. However, generally the phase- and frequency-encoding gradients in the x- and y-directions are negligible in strength for a small volume V when compared to the gradient in the z-direction. This means that the gradient can be treated as if in a single direction, reducing k to a scalar. Equation (2.5) shows that the application of a gradient results in a shift of the configuration states. For higher values of k, the magnetisation becomes more dephased. For $k = 0$, the magnetisation is *in phase*, resulting in a measurable RF signal.

When a pulse sequence is used with constant repetition times and gradients, such as a CPMG [23] or gradient spoiled SSFP [24] sequence, the configuration states and dephasing coordinate can be discretised. Since every shift as the result of a gradient has the same magnitude, the exact magnitude of the shift is no longer important. Only the number of shifts and their direction is sufficient to describe the system states. The discretised configurations states become F_k and Z_k, where k can have both positive and negative values. Still, only the states with $k = 0$ result in a non-zero signal contribution.

Discretising the configuration states allows us to define a *shift operator* $\boldsymbol{S}$ with properties

$$\boldsymbol{S}: \qquad F_k \rightarrow F_{k+\Delta k} \quad \text{and} \quad Z_k \rightarrow Z_k. \tag{2.6}$$

As (2.6) shows, the longitudinal magnetisation is left unaffected by gradients in this direction.

RF pulses are implemented in the EPG framework through a *transition operator*. For the following section, the flip of the magnetisation is assumed to be around the x-axis. Similar operators can be created for rotations around the other axes. The transition operator is created by first transforming from the configuration states back to the individual components of the magnetisation, applying a rotation matrix around the x-axis, and then transforming back to the configuration state description. Without further derivation presented here, the transition operator is defined as

$$\begin{bmatrix} F_k \\ F_{-k} \\ Z_k \end{bmatrix}^+ = \begin{bmatrix} \cos^2\frac{\alpha}{2} & e^{2i\phi}\sin^2\frac{\alpha}{2} & -ie^{i\phi}\sin\alpha \\ e^{-2i\phi}\sin^2\frac{\alpha}{2} & \cos^2\frac{\alpha}{2} & ie^{-i\phi}\sin\alpha \\ -\frac{i}{2}e^{-i\phi}\sin\alpha & \frac{i}{2}e^{i\phi}\sin\alpha & \cos\alpha \end{bmatrix} \begin{bmatrix} F_k \\ F_{-k} \\ Z_k \end{bmatrix}^- . \tag{2.7}$$

Here, the matrix on the right hand side of (2.7) is the *transition matrix* $\boldsymbol{T}(\alpha,\phi)$. Here, α is the flip angle of the pulse, and ϕ is its phase. The full derivation of this matrix, and a complete description of the EPG framework, can be found in [22]. A very important consequence of equation (2.7) is that an RF pulse can split one configuration state into multiple, and that configuration states can be mixed by the application of a pulse. This *partitioning effect* allows dephasing magnetisation to be converted into re-phasing magnetisation and vice

versa. This allows for the description of spin echos in the EPG framework.

When working with the discrete description of EPG, three configuration states exist for every value of k. When an RF pulse is applied, the transition matrix is applied to each set of three states separately. No mixing takes place between states of different k-values.

Finally, the effects of relaxation can also be described using a matrix operator. Representing the exponential decay effects of longitudinal and transverse relaxation as $E_1 = e^{-\tau/T_1}$ and $E_2 = e^{-\tau/T_2}$ respectively, the relaxation matrix is given by

$$\boldsymbol{R}(\tau, T_1, T_2) = \begin{bmatrix} E_2 & 0 & 0 \\ 0 & E_2 & 0 \\ 0 & 0 & E_1 \end{bmatrix}. \tag{2.8}$$

Applying this matrix to a set of configuration states models the effects of relaxation. As with the transition matrix, when dealing with the discretised representation of EPG, the relaxation operator is applied to each set of three states separately.

Now, the effects of a pulse sequence can be modelled simply by successively applying these operators. As an example, assume the following operations:

1. An RF pulse with flip angle α and phase ϕ
2. A spoiling gradient (in the z-direction)
3. A relaxation period with duration τ

For any set of three discrete configuration states with dephasing coordinate k, this can be represented as

$$\begin{bmatrix} F_k \\ F_{-k} \\ Z_k \end{bmatrix}^{t=\tau} = \boldsymbol{R}(\tau, T_1, T_2) \quad \boldsymbol{S} \quad \boldsymbol{T}(\alpha, \phi) \begin{bmatrix} F_k \\ F_{-k} \\ Z_k \end{bmatrix}^{t=0}. \tag{2.9}$$

This shows: in EPG, the operations done by the pulse sequence are implemented as operators applied successively to the configuration states.

2.2.4. Undersampling

As mentioned above, MRF relies on many readouts to acquire the signals. To keep the total scan time to a minimum, MRF utilises efficient k-space trajectories, where during each readout only a small section of k-space is sampled. For a single readout, this would result in highly undersampled images with severe artefacts. However, as each individual readout samples a *different* section of k-space, a full sampling is acquired over the course of the scan. This allows for very fast individual readouts, while still acquiring accurate parameter maps over the full volume of interest [13].

2.3. Multi-component MRF

Conventional MRF, using inner-product matching, neglects *partial volume effects*, where multiple tissue types coexist in a single voxel and all contribute to the final signal. *Multi-component MRF* (MC-MRF) aims to retrieve all tissue components in a voxel, as opposed to only the single best matching one [25].

The general assumption in MC-MRF is that the signal measured from a multi-component system is a linear combination of the signals coming from the individual components. Nonlinear influences on the total signal are not taken into account. To resolve the individual components from the dictionary, the linear combination of atoms that best resembles the measured signal needs to be determined. Defining this linear combination as

$$\tilde{s} = \sum_{i}^{N} c_i \boldsymbol{d_i} = D\boldsymbol{c}, \tag{2.10}$$

where the coefficients in $\boldsymbol{c}$ determine the approximation of the signal. Since the best approximation of the measured signal minimises the difference between this signal and the approximation, $\boldsymbol{c}$ is the solution of

$$\min_{\boldsymbol{c} \in \mathbb{R}^N_{\geq 0}} \|\boldsymbol{s} - D\boldsymbol{c}\|_2, \tag{2.11}$$

where an additional constraint is added that the coefficients should be non-negative [26]. This type of minimisation problem is known as a *non-negative least squares* problem. Generally, a problem like this is highly under-determined. Often, a regularisation term is used to promote a solution with a small number of nonzero coefficients [26]. Additional constraints can also be added, such as with SPIJN [25], which iteratively applies a *joint sparsity* constraint over all voxels. Here, "joint" refers to the fact that this method aims to create a sparse solution for both the individual voxels and the volume of interest as a whole.

2.4. Magnetisation transfer

Contrary to the assumption made in most implementations of MC-MRF, nonlinear effects are known to influence the total signal resulting from a multi-component system [27]. The magnetisations of the individual tissue components can interact with each other, transferring magnetisation from one component to another. This process, called *magnetisation transfer* (MT), modifies the overall shape of the compound signal when compared to the linear combination. In the context of MWF imaging, this particular type of magnetisation transfer is *free pool-free pool magnetisation transfer*. This is MT between two "pools" of water molecules (such as the water in white matter, and myelin water), as opposed to MT between a free and a *bound* pool. This latter case is explored in e.g. [28]. The inclusion of magnetisation in the description of a multi-component system gives rise the goals as formulated in the previous chapter: to what extent does the exchange influence MC-MRF results, and how can we include information about this exchange in the estimation methods for more accurate results?

3

Methods

This chapter explains the methods used during this project. Relevant parameters, values and techniques are mentioned here.

3.1. Simulation and estimation techniques

As explained in the previous chapter, magnetic resonance fingerprinting relies on a signal dictionary with a large collection of precomputed signals for many combinations of the relevant tissue parameters. During this project, signals were generated using the EPG framework, motivated by the pulse sequence type. More specifically, an extension of EPG, called EPG-X [29], allowed us to simulate not only the expected signals for a single component at a time, but also for multiple components simultaneously, including magnetisation transfer effects.

EPG-X extends the EPG framework, allowing multiple components at the same time to be simulated. Here, we only deal with two-component systems. The components will be referred to as component A and B respectively. Each component has its own set of configuration states: $\left[F_{+,A},\, F_{-,A},\, F_{z,A},\, F_{+,B},\, F_{-,B},\, F_{z,B}\right]$. To apply RF pulses to both components, a transition matrix $\boldsymbol{T}$ is defined by tiling the single-component transition matrix as defined in (2.7) twice diagonally. The remaining values in the resulting 6×6 matrix are set to zero:

$$\boldsymbol{T} = \begin{bmatrix} \boldsymbol{T}(\alpha,\phi) & 0 \\ 0 & \boldsymbol{T}_(\alpha,\phi) \end{bmatrix}. \tag{3.1}$$

The relaxation matrix is modified to include two new parameters: k_{ab} and k_{ba}, called the *exchange rates* or *transfer rates*. These parameters define the amount of magnetisation transferred from A to B, and the amount of magnetisation transferred from B to A respectively. They are included in the new relaxation operator as

$$\boldsymbol{R} = \begin{bmatrix} E_{2,A} - k_{ab} & 0 & 0 & k_{ba} & 0 & 0 \\ 0 & E_{2,A} - k_{ab} & 0 & 0 & k_{ba} & 0 \\ 0 & 0 & E_{1,A} - k_{ab} & 0 & 0 & k_{ba} \\ k_{ab} & 0 & 0 & E_{2,B} - k_{ba} & 0 & 0 \\ 0 & k_{ab} & 0 & 0 & E_{2,B} - k_{ba} & 0 \\ 0 & 0 & k_{ab} & 0 & 0 & E_{1,B} - k_{ba} \end{bmatrix}, \tag{3.2}$$

with $E_{1,A}$, $E_{2,A}$, $E_{1,B}$, and $E_{2,B}$ as defined in equation (2.8) for component A and B respectively. The influence of a gradient on the configuration states is unchanged, simply shifting the transverse magnetisation to the "next" configuration state.

A final additional parameter introduced in EPG-X is the *fraction* f. This parameter defines the fraction of component B in a two-component system. The fraction is defined on the interval [0 - 0.5], so that by definition component B is the smaller of the two components. With $\boldsymbol{M_0}$ representing the total equilibrium magnetisation of the system, f is used to define the the equilibrium magnetisations of the two tissue components as $\boldsymbol{M_0^A} = (1-f)\boldsymbol{M_0}$ and $\boldsymbol{M_0^B} = f\boldsymbol{M_0}$. To maintain thermal equilibrium, the exchange rates are related to each

other through the fraction:

$$k_{ba} = \frac{1-f}{f} k_{ab}. \tag{3.3}$$

Through this equation only one of the exchange rates has to be specified, along with f, to fully define the exchange in a two-component system.

The MRF pulse sequence used in this project is the sequence as described in [19]. This is a sequence consisting of 500 RF pulses with flip angles varying between 0° and 60°. The time between each pulse is constant at 15 ms. The phase of the RF pulses is zero for every pulse. Additionally, the sequence includes an initial *inversion pulse* to better encode both T_1 and T_2 effects in the signals [13]. This pulse also has a phase angle of zero, and an inversion time of 19 ms [19].

The used values for the relaxation parameters for *white matter*, *grey matter*, and *myelin water* are given in Table 3.1. The values are adapted from [29] and [30]. Additionally, the exchange rate between brain tissues is taken to be k_{ab} =2 s^{-1} [29]. The flip angle sequence, together with the simulated single-component signals for the three aforementioned tissues, is shown in Figure 2.2.

All EPG-X simulations were implemented in the Python programming language [31]. The implementation is an adaptation of the MATLAB [32] implementation as given in [29]. The code, as well as relevant simulation data, can be found at `https://gitlab.tudelft.nl/managtegaal/mep_sven`.

Table 3.1: The values of the relaxation parameters for the three main tissue types used in this report.

Tissue type	T_1 **(ms)**	T_2 **(ms)**
White matter	800	79
Grey matter	1459	93
Myelin water	400	20

3.2. Numerical phantoms

The performance of several different MC-MRF techniques was tested by applying these to two-component numerical phantoms. Signals for each voxel were computed using EPG-X, based on their parameters and a given pulse sequence. While EPG-X parameters are defined per voxel, the pulse sequence is applied phantom-wide.

Different phantoms were made for different purposes. The first phantom has a resolution of 10 by 10 voxels, with six parameters for each voxel: f, k_{ab}, $T_{1,A}$, $T_{2,A}$, $T_{1,B}$, and $T_{2,B}$. An increasing gradient from left to right is applied to the fraction f, with values in the range [0, 0.5]. An increasing gradient from the top to the bottom of the phantom is applied to the exchange rate k_{ab}, with values ranging from 0 to 2. The relaxation parameters of component A are chosen to represent white matter, while component B represents myelin water.

The second phantom just has a higher number of voxels: 40 by 40. Like the previous phantom, the fraction and exchange rate maps contain horizontal and vertical increasing gradients respectively, with identical ranges for the values. As such, this phantom is identical to the previously described phantom with respect to all the tissue parameters.

The last phantom has the same resolution as the second (40 by 40 voxels), but a different structure. An increasing gradient is still present in the fraction, ranging from 0 to 0.5. However, the exchange rate is now set at a constant value of 2 s^{-1} for the entire phantom. Component A now has the relaxation parameters of either white matter or grey matter, and component B has the parameters of either white matter or myelin water. This is done in such a way that all three combinations of these tissues are present in the phantom.

All relevant numerical phantoms used during this project are shown in Figure 3.1. The descriptions of the phantoms above correspond to phantom (a), (b) and (c) in the image respectively.

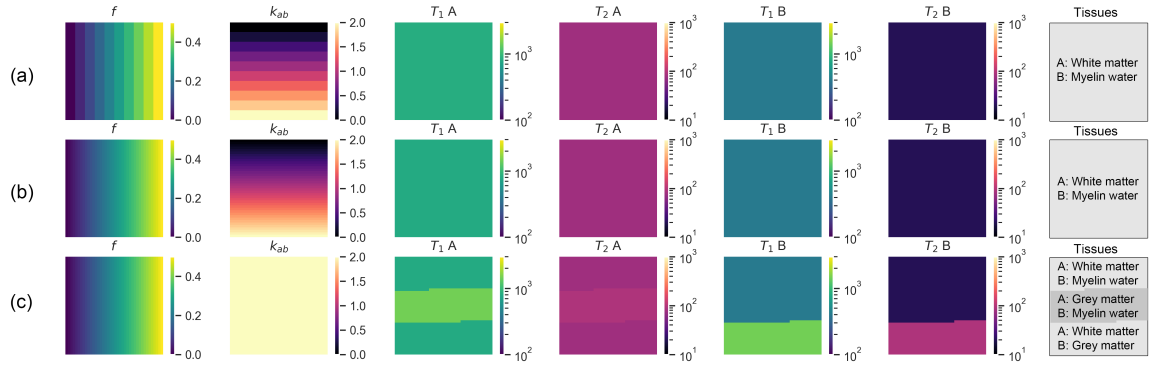

Figure 3.1: Overview of the phantom parameters for the three phantoms used in this project. Each row contains one phantom. The leftmost column gives an identifier for each phantom for reference in the text. The rightmost column gives an overview of the tissues represented in the phantom, for different sections if necessary.

Table 3.2: The parameters used for generating the single-component MRF dictionary. Both T_1 and T_2 follow an exponential scaling between their lowest and highest values.

Parameter	Min	Max	N	Units
T_1	100	3000	50	ms
T_2	10	1000	50	ms

3.3. Parameter mapping techniques

Several different parameter mapping techniques were applied to the numerical phantoms. This includes conventional MC-MRF techniques, and novel methods including more information about the magnetisation transfer. An overview of these techniques is given here.

Voxel-wise Non-Negative Least Squares

Non-negative least squares, as already described in equation (2.11), is the conventional method for MC-MRF. This method is applied voxel-wise, meaning that no information about neighbouring voxels is used during the fitting process.

NNLS was implemented using the `scipy.optimise.nnls` function [33]. The dictionary used was a single-component dictionary generated for a range of T_1 and T_2 combinations. The exact ranges of these values are given in Table 3.2. Important here is that by definition the transverse decay is faster than the longitudinal decay. Therefore, the final dictionary only retains atoms for which $T_1 > T_2$ holds. This resulted in a dictionary containing 2064 atoms in total.

We applied NNLS with this dictionary on phantom (a) (Figure 3.1), which resulted in a coefficient vector $\mathbf{c}$ for each voxel in the phantom. The sum of coefficients for each voxel was normalised to 1. Parameter maps for two components were obtained by thresholding the values for T_1 and T_2 into a slow (component A) and a fast (component B) water pool. The fast water pool consisted of all atoms with $T_1 < 550$ ms and $T_2 < 40$ ms. The slow pool contained all atoms with $T_1 > 550$ ms and $T_2 > 40$ ms. The fraction is then defined as the sum of coefficients for component B. Relaxation parameters for both components were computed as the weighted sum of the dictionary parameters, where the component-specific coefficients were normalised to a sum of 1.

SPIJN

As opposed to NNLS, the SPIJN algorithm takes the *entire* phantom into account. It solves the NNLS problem with ℓ_1-regularisation for a sparser solution for each voxel separately, and subsequently uses *all* acquired coefficients to iteratively reweigh the signal dictionary to promote sparsity over the entire region. The implementation of SPIJN as given in [25] was used during this research. Like voxel-wise NNLS, SPIJN was applied using the dictionary given in Table 3.2.

SPIJN was applied to the same phantom as NNLS was. The regularisation parameter λ was chosen in such

a way that applying SPIJN to this phantom resulted in two final tissue components, the same number as is present in the phantom. SPIJN results in a coefficient vector for each voxel. Parameters were retrieved from these coefficient maps using the same methods as with NNLS.

Two-component Dictionary

To improve mapping accuracy by taking into account the full effects of magnetisation transfer on a two-component system, *two-component dictionaries* were generated. Contrary to a conventional MRF dictionary, these dictionaries contain signals simulated for different combinations of two tissue components. This increases the number of included tissue parameters from two (T_1 and T_2) to six (f, k_{ab}, $T_{1,A}$, $T_{2,A}$, $T_{1,B}$, and $T_{2,B}$). Given a range for each of these parameters, dictionary signals were generated using EPG-X for every combination of the parameter values. Exceptions to this are *invalid* and *duplicate* parameter combinations, which are discarded from the dictionary. The three conditions under which an atom is discarded are:

1. $T_{1,A} < T_{2,A}$ or $T_{1,B} < T_{2,B}$ (invalid relaxation time combinations)
2. $T_{1,A} = T_{1,B}$, $T_{2,A} = T_{2,B}$, and $f \neq 0$ (duplicate with the atom where $f = 0$)
3. $T_{1,A} \neq T_{1,B}$ or $T_{2,A} \neq T_{2,B}$, *and* $f = 0$ (fraction of zero does not need two different components)

Following the above selection criteria, a full two-component dictionary could be generated without any duplicate signals.

Using this definition, we created two two-component dictionaries. The parameter ranges used to generate the atoms are given in Table 3.3. The first dictionary has variation in all six parameters. The second dictionary was made with the exchange rate at a fixed value of $2\,s^{-1}$. This allowed for a slightly higher sampling density of the remaining parameters. As a result of the increased number of parameters, both these dictionaries are significantly larger than the single-component dictionary, with respectively 259 120 and 414 305 atoms.

Table 3.3: The parameters used to generate the two two-component dictionaries. N is the number of different values the parameter has in the dictionary. Here, f and k_{ab} scale linearly between their lowest and highest values. The different relaxation times all scale exponentially between their lowest and highest values.

(a) The parameters and their minimum and maximum values used to generate the two-component dictionary with varying exchange rate. The dictionary contains a total of 259120 atoms.

Parameter	Min	Max	N	Units
f	0	0.44	5	-
k_{ab}	0	3	10	s^{-1}
$T_{1,A}$	100	3000	10	ms
$T_{2,A}$	10	1000	10	ms
$T_{1,B}$	100	3000	10	ms
$T_{2,B}$	10	1000	10	ms

(b) The parameters and their minimum and maximum values used to generate the two-component dictionary with fixed exchange rate. The dictionary contains a total of 414305 atoms.

Parameter	Min	Max	N	Units
f	0	0.47	10	-
k_{ab}	-	2	1	s^{-1}
$T_{1,A}$	100	2273	15	ms
$T_{2,A}$	10	227	15	ms
$T_{1,B}$	100	2273	15	ms
$T_{2,B}$	10	227	15	ms

The two-component phantoms were matched to the dictionaries using inner-product matching (equation (2.4)), here also called *dictionary matching*. Since every atom corresponds to two components, this results in a two-component parameter map without the need for NNLS. The inner-product matching was applied on phantom (b) using the first of the two dictionaries. The second dictionary was used for matching of the signals from phantom (c). The tissue parameters for both components, as well as the fraction and exchange rate, were directly extracted from the dictionary.

Additionally, the same dictionary and phantom combinations were used for *dictionary fitting*. To cope with an exponential increase of atoms due to the large number of pairs, the parameter sampling of the two-component dictionary is done relatively coarsely. Dictionary fitting, as opposed to dictionary matching, first interpolates between the dictionary atoms using B-spline interpolation [34]. It does this by treating the dictionary atoms as if distributed in a 6-dimensional parameter space. We used the implementation as given in [34] in MATLAB [32] to interpolate between the dictionary signals using a second-order B-spline interpolation. An initial match is found from the dictionary using inner-product matching, which serves as a starting

point for a more detailed matching on the interpolated dictionary. Additionally, a fitting with a manual initialisation close to the reference values was done. The initial relaxation parameters were chosen to be as close to the reference values as possible. The fraction was set at a constant $f = 0.21$ for the manual initialisation. Since the parameter space also includes all invalid parameter combinations, which would normally be discarded, an additional mask is added to the dictionary to mask out the invalid combinations for the initial matching step. The result of the dictionary fitting is a set of non-integer indices into the 6-dimensional parameter space. As the dictionary parameters were determined by sampling a parameter distribution (here linear or exponential) at integer intervals, the non-integer indices can be converted back to parameter values by sampling these distributions at the non-integer positions.

Clustering-based Regularisation

A parameter mapping method with a sparsity-like constraint was also implemented for the two-component dictionaries. It uses a *clustering* step to greatly reduce the amount of available dictionary atoms for matching. More specifically, we used *k-means clustering* [35]. The idea behind this is similar to SPIJN: in a biological sample, there exists generally a limited number of tissues. This also implies that there is a limited number of tissue *combinations*. Since every dictionary atom is a tissue combination, this translates to a limited number of dictionary atoms. The clustering step aims to reduce the number of combinations of k_{ab}, $T_{1,A}$, $T_{2,A}$, $T_{1,B}$, and $T_{2,B}$. The number of values for the fraction f is not limited by the clustering. This parameter determines *how* the tissues are combined, but does not influence the number of combinations. Note that in the above definition of a tissue combination, it is assumed that the exchange rate k_{ab} is particular to the tissue combination, and does not vary for the same combination of tissues.

The spatial regularisation is applied as follows: initially, an inner-product match is made following equation 2.4. Each of these matches represents a full set of estimated parameters. These top matches are then distributed as points in a 5-dimensional parameter space (one dimension for each parameter, with the exception of f). The k-means clustering algorithm groups these points in a predetermined number N_c of groups or *clusters*. In k-means clustering, the algorithm minimises the total mean distance of all the points in a cluster to the cluster *centre*. It iteratively changes the centre positions until a minimum is reached. After the clustering, the N_c cluster centres represent N_c tissue combinations, given by the atoms closest to the cluster centres. The reduced dictionary that is used for the final matching step consists of all dictionary atoms with these parameter combinations. Additionally, the atoms with the same parameters as component A and $f = 0$ are retained to allow the fraction to be zero for these combinations.

Parameter mapping with the spatial regularisation was done using both dictionaries given in Table 3.3. As with the dictionary matching, the first dictionary was applied to phantom (b), the second dictionary to phantom (c). The k-means clustering algorithm was implemented using the `sklearn.cluster.KMeans` function [36]. The parameter values were extracted from the two-component dictionary directly.

Correction Terms

The two-component dictionaries were made to fully include all information about magnetisation transfer effects in the dictionary. A downside of this is that this results in very large dictionaries even for a relatively coarse parameter sampling. Another approach to including information about the magnetisation transfer in the dictionary is as follows. The isolated effects of magnetisation transfer on the signals of a multi-component system can be written as

$$\boldsymbol{m} = \boldsymbol{s} - \tilde{\boldsymbol{s}}, \tag{3.4}$$

where $\boldsymbol{s}$ and $\tilde{\boldsymbol{s}}$ are the signal of a multi-component system with and without magnetisation transfer respectively. Using equation (3.4), we can rewrite the solution of the NNLS problem (2.10) with magnetisation transfer as

$$\boldsymbol{s} = \sum_{i}^{N} c_i \boldsymbol{d_i} + \boldsymbol{m}. \tag{3.5}$$

Now, we assume that the total magnetisation transfer effect $\boldsymbol{m}$ can be approximated by a weighted sum of atom-specific MT correction terms $\boldsymbol{m_i}$. Additionally, we assume that the influence of these correction terms is directly related to the coefficient of their corresponding atom. Filling in these assumptions in the previous

Table 3.4: The parameters used for the generation of the single-component dictionary with correction terms. Both T_1 and T_2 follow an exponential scaling between their lowest and highest values.

Parameter	**Min**	**Max**	***N***
T_1	100	3000	100
T_2	10	1000	100

equation gives:

$$\boldsymbol{s} = \sum_i^N \left(c_i \boldsymbol{d_i} + f(c_i)\boldsymbol{m_i}\right), \tag{3.6}$$

where $f(c_i)$ is a function mapping the coefficient of an atom to the coefficient of the correction term. Assuming an ℓ_1-normalised coefficient vector, this function was chosen to be $f(c_i) = -2\left(c_i - \frac{1}{2}\right) + \frac{1}{2}$. This is a parabola with zeros at 0 and 1, and its peak at 0.5 with value 0.5. This function is chosen to ensure that at very small or very large coefficients, the influence of the correction factor is minimal. When it is very low, the component is present in such a small quantity that the influence of its transfer on the total signal should be very small. Likewise, if the coefficient is very high, there are only very small amounts of other tissue present to transfer with, also resulting in minimal MT influence. The closer the coefficient gets to 0.5, the higher the influence of the correction factor should be. The peak is placed at $\frac{1}{2}$ so that at a two-component system with both coefficients at 0.5, both components have equal contribution to the correction term, with a total of 1.

A dictionary with included correction terms was generated following Table 3.4. The dictionary signals were simulated in the same way as with the conventional dictionary (Table 3.2). To create the correction terms, for any atom with relaxation parameters $T_{1,i}$ and $T_{2,i}$, the two-component signals $\boldsymbol{s_{i,n}}$ and $\widetilde{\boldsymbol{s}}_{\boldsymbol{i,n}}$ were simulated with a *random* second component with parameters $T_{1,n}$, $T_{2,n}$, and a fraction of 50 %, after which the MT effect for this particular combination was determined. The final correction term was then defined as the mean value of these MT effects:

$$\boldsymbol{m_i} = \frac{1}{N}\sum_{n=1}^{N}\left(\boldsymbol{s_{i,n}} - \widetilde{\boldsymbol{s}}_{\boldsymbol{i,n}}\right), \tag{3.7}$$

with N the number of random tissue combinations. An example of three of these correction terms, including their standard deviations, is shown in Figure 3.2.

Parameter mapping was done using an iterative method. The initial mapping was done using NNLS, without any contribution of the correction terms. The coefficients from this initial mapping were then used to determine the coefficients of the correction terms. The correction terms were subtracted from the signal, resulting in a new signal for the next iteration:

$$\boldsymbol{s_{n+1}} = \boldsymbol{s_0} - \sum_i^N f(c_{i,n})\boldsymbol{m_i}. \tag{3.8}$$

Here, $\boldsymbol{s_{n+1}}$ is the signal for iteration $n+1$, $\boldsymbol{s_0}$ is the original signal vector, and $c_{i,n}$ are the NNLS coefficients from iteration n.

The method was tested by applying it to phantom (a) in Figure 3.1, which was modified to have a fixed exchange rate of $k_{ab} = 2\,\mathrm{s}^{-1}$. This was done for a single iteration of equation (3.8), as well as for three iterations.

3.4. Error definitions

Throughout this research, all results were acquired numerically. Therefore, they are free of measurement and noise errors. By estimating parameters, however, we can still define *estimation errors*. For any parameter θ and its estimated value $\widetilde{\theta}$, the estimation error is defined as

$$E_{est} = \widetilde{\theta} - \theta. \tag{3.9}$$

This definition is consistent with the notion that an overestimation of a parameter results in a positive error, and an underestimation results in a negative value.

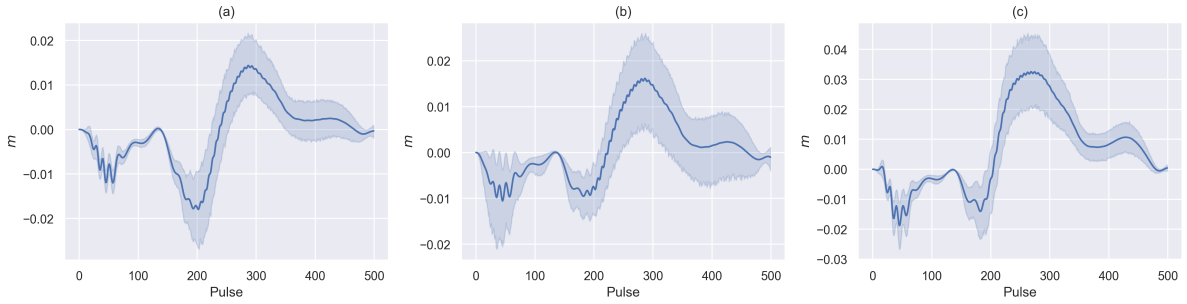

Figure 3.2: The mean correction terms of (a) white matter, (b) grey matter, and (c) myelin water. The shaded area represents the standard deviations. All results here are for 15 tissue combinations.

The *interpolation error* is the error made by the B-spline interpolation on a dictionary. This error is defined for each dictionary parameter separately. Assuming a dictionary containing n separate parameters, the "position" in parameter space can be written as $\boldsymbol{\nu} = [\theta_1\ \theta_2\ \cdots \theta_n]$, with θ_i any parameter. The interpolation error at a parameter position $\boldsymbol{\nu}$ is then defined as

$$E_{int}(\boldsymbol{\nu}) = \|\boldsymbol{s}(\boldsymbol{\nu}) - \widetilde{\boldsymbol{s}}(\boldsymbol{\nu})\|_2\,, \tag{3.10}$$

where $\boldsymbol{s}(\boldsymbol{\nu})$ and $\widetilde{\boldsymbol{s}}(\boldsymbol{\nu})$ are the simulated and interpolated signal at position $\boldsymbol{\nu}$ in parameter space respectively. The interpolation error is assumed to be largest at either the extreme values of the different parameters (i.e. their smallest and largest values), or in the centre of the dictionary. The total interpolation error for a parameter is then acquired by evaluating (3.10) for all points *in between* the dictionary samples at these extreme edges. Additionally, the error is evaluated at the centermost point of the dictionary in between the samples of the parameter of interest. For an n-parameter dictionary, this means evaluations along 2^{n-1} edges, plus a single evaluation in the middle. The interpolation error for a given parameter is then defined as the maximum value of the interpolation errors at the evaluated positions.

4

Results

This chapter presents the results of the experiments previously described. To improve the readability of this chapter, tables containing estimation errors of the parameter maps are presented in Appendix A. Also included there are additional mapping results using a different pulse sequence, including errors.

4.1. NNLS and SPIJN maps

Figure 4.1 shows the result of an NNLS fit on a single signal. This fit was made with the dictionary as generated following Table 3.2. This figure shows that NNLS is able to very closely represent the reference signal as a linear combination of dictionary atoms.

The results of the parameter mapping with both NNLS and SPIJN on phantom (a) are shown in Figure 4.2. The mean estimation errors, along with the standard deviations for the parameter estimations are given in Table A.1. Since both NNLS and SPIJN do not permit the mapping of the exchange rate, no maps for this are presented in Figure 4.2. The exchange rate map of the reference phantom is shown to more clearly illustrate the effects of the varying exchange rate on the resulting parameter estimations. The figure shows a clear bias in the retrieved fraction maps for both NNLS and SPIJN when compared to the fraction map of the phantom. For NNLS, this bias increases as the value for the exchange rate increases. SPIJN shows a more constant bias over the entire phantom. The relaxation parameters show a bias toward lower values as well. For NNLS, this underestimation increases with both exchange rate and fraction. As a result of the sparsity constraint, SPIJN shows no such trend. The mean estimation error of the fraction for each row of the phantom for both NNLS and SPIJN is given in Table 4.1.

Figure 4.3 shows three fraction maps made for three different values of the exchange rate. The maps are made using two-component signals consisting of white matter and myelin water components. As can be seen, the fraction maps show no bias when no exchange is present (k_{ab} =0 s^{-1}), as is expected. For larger values of

Table 4.1: The mean estimation errors for the fraction estimation made NNLS and SPIJN and shown in Figure 4.2. The values are given for each row in the phantom, with the corresponding exchange rate.

k_{ab} ($s^{-}1$)	**NNLS (%)**	**SPIJN** (%)
0.00	1.30	-13.7
0.22	-6.38	-13.0
0.44	-9.47	-12.4
0.67	-11.7	-11.9
0.89	-13.4	-11.4
1.11	-14.6	-11.0
1.33	-15.6	-10.7
1.56	-16.3	-10.4
1.78	-17.0	-10.1
2.00	-17.7	-9.81

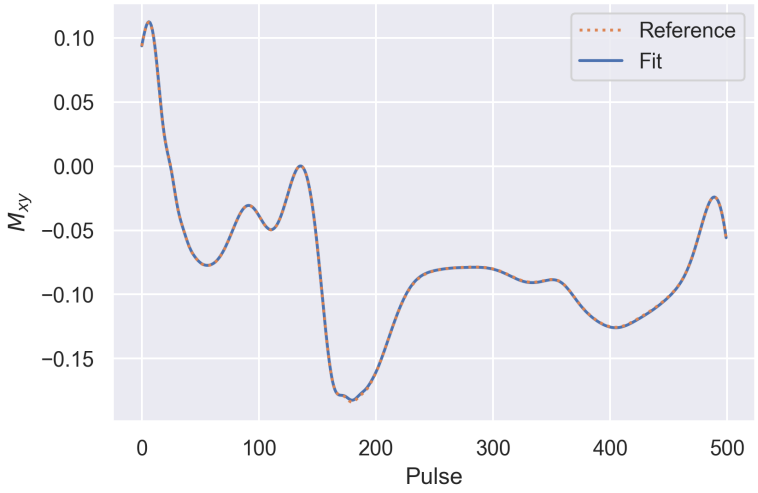

Figure 4.1: The result of an NNLS fit on a two-component signal (blue). The reference signal is shown as an orange dotted line. The original signal is simulated for 72 % white matter and 28 % myelin water, with an exchange rate of $k_{ab} = 1.11$. The fitted signal consists of the weighted sum of six dictionary atoms. All other atoms were given a coefficient of zero by the NNLS algorithm.

the exchange rate, NNLS shows an increasing bias towards lower fraction values. The bias of SPIJN remains relatively constant, but this method still shows a clear tendency to underestimate the fraction values as a result of magnetisation transfer.

4.2. Two-component dictionary

In Figure 4.4, the results for the dictionary described in Table 3.3a, applied to phantom (b), are shown. This figure shows the maps acquired using both the inner-product matching and the clustering-based method. Figure 4.5 shows the results of the same techniques for the second dictionary, with parameters given in Table 3.3b.

When comparing the middle row to the top row of Figure 4.4, it becomes clear that the parameter maps created by the inner-product matching do not accurately resemble the reference maps. The gradients present in the reference maps, for both the fraction and the exchange rate, cannot be recognised from the parameter estimation maps. Additionally, the relaxation parameter maps do not show the uniformity of the phantom, with the largest variation in the retrieved values in the relaxation parameters of component B. The bottom row of Figure 4.4 shows the results of the clustering-based method. As a result of the clustering, the relaxation parameter estimations are now uniform. The fraction map now does show a gradient. However, this gradient is in the wrong direction, now increasing from right to left instead of left-to-right. Additionally, the found relaxation values for component B seem to fit better to the reference values of component A, and vice versa.

Observing the middle row of Figure 4.5, the fraction estimate map does not resemble the reference map. There seems to be no clear bias; the fraction is both under- and overestimated equally. Additionally, the reference shows that in each column, all voxels have the same fraction, independent of the tissue combination. The fraction estimation, however, shows different fractions for different tissue combinations. The relaxation parameter maps show rapid spatial variation in the retrieved values, which is not present in the reference. Similar to the results of Figure 4.4, the variation seems to be the largest for the parameters of component B. The results of the clustering-based method again show more smooth results. However, the retrieved fraction still changes with different tissue combinations, and shows clear errors. Additionally, large errors still exist in the relaxation parameter maps.

The estimation errors and the standard deviations for the maps in figures 4.4 and 4.5 are given in tables A.2 and A.3 respectively.

To gain further insight in the behaviour of the clustering-based method, four *component plots* were made.

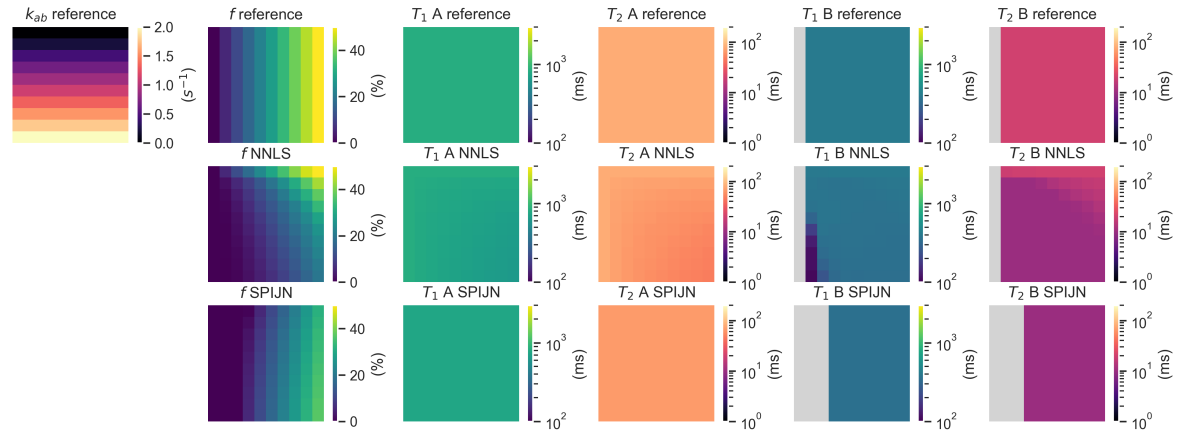

Figure 4.2: The results of the parameter mapping using both NNLS and SPIJN. The T_1- and T_2-boundary values used to separate the two components were 550 ms and 40 ms respectively. The first row of maps in this image contains the reference values as present in the numerical phantom. The middle row shows the maps as acquired using single-voxel NNLS. The bottom row shows the maps as acquired using SPIJN. Important to note are the light grey areas in the maps for the relaxation times of the second component. These are the result of the mapping returning a fraction of 0 %. When this is the case, there is no second component, and therefore no parameter mapping possible for component B of those voxels.

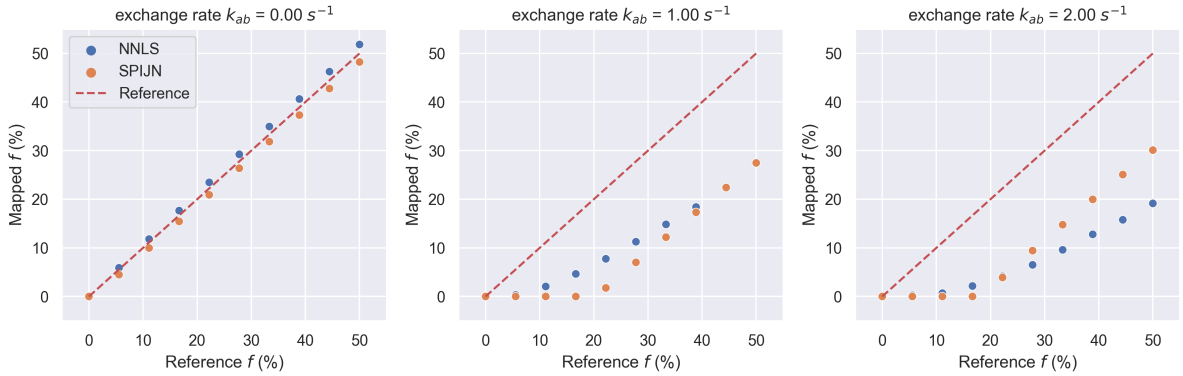

Figure 4.3: Fraction maps for different values of the exchange rate k_{ab}. Mapped values are acquired using both NNLS and SPIJN. All signals are simulated for a combination of white matter (component A) and myelin water (component B). The red dotted line shows the expected trend in the fraction map.

These plots show the positions of the top match per voxel after inner-product matching with the dictionary with fixed exchange rate. The matches are shown in parameter space for several combinations of relaxation parameters. The results of the clustering step are shown by colour coding the matches based on their cluster. Additionally, the cluster centres are shown as stars, and the positions of the reference values are shown as plus signs. The component plots are shown in Figure 4.7. These plots show that the top matches are not distributed in three obviously separate clusters. Instead, the formation more closely resembles a single "cloud" of points. As can be most clearly seen in the bottom left and bottom right plots in this figure, the cluster centres do not always properly coincide with the reference values.

Finally, the results of the parameter mapping made using dictionary fitting are shown in Figure 4.6. The estimation errors are given in Table A.4. The middle row in this figure shows the results of fitting using the results of inner-product matching as initialisation. The bottom row shows the results of dictionary fitting with the manual initialisation. The interpolation errors for the different dictionary parameters are given in Table 4.2.

Comparing the middle rows of Figure 4.5 and 4.6, similar trends in the retrieved relaxation parameter maps can be observed. Over- and underestimations occur in roughly the same locations in both figures for all relaxation parameters. The fraction map in Figure 4.6 does not show the gradient that is present in the reference map. Additionally, some fractions returned by the dictionary fitting method are *negative* (shown in grey). Following the definition of the fraction, negative values do not make physical sense.

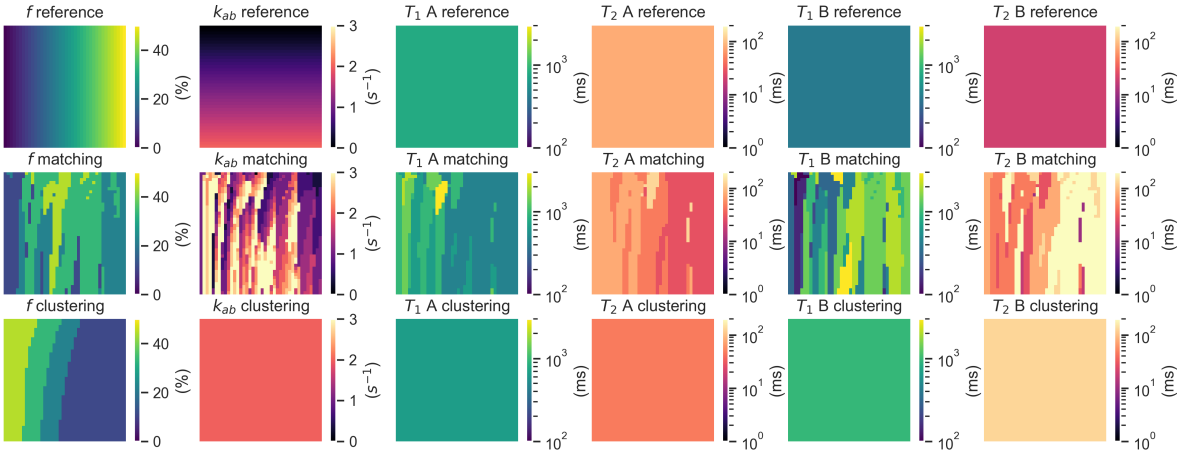

Figure 4.4: The parameters of the numerical phantom (top row) and the parameter maps made using inner-product matching (middle row) and clustering-based mapping (bottom row) on this phantom. All maps were made using a two-component dictionary with parameters as described in Table 3.3a.

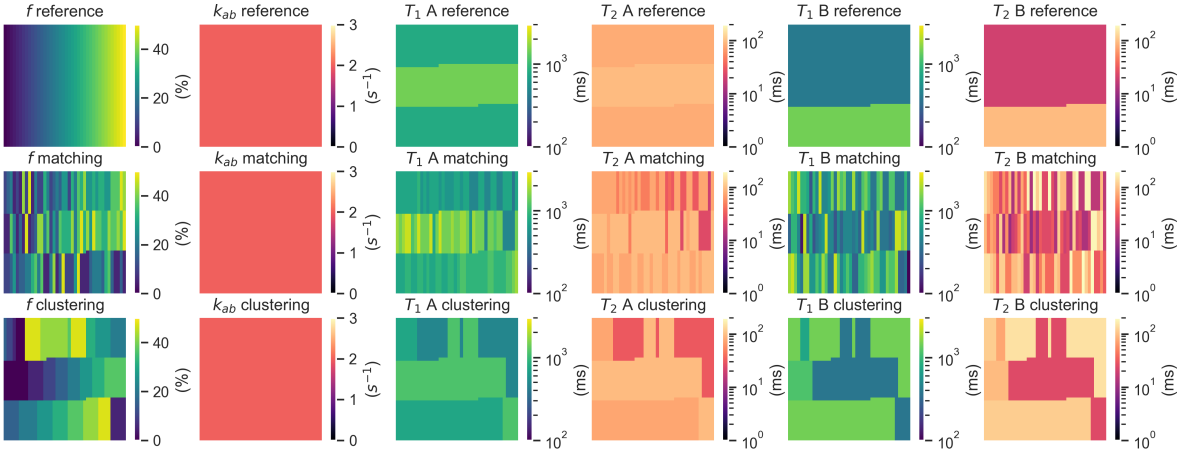

Figure 4.5: The parameters of the numerical phantom (top row) and the parameter maps made using the inner-product matching (middle row) and clustering-based mapping (bottom row) on this phantom. All maps were made using a two-component dictionary with fixed exchange rate. The parameter ranges are described in Table 3.3b.

Table 4.2: The interpolation errors for the B-spline interpolation of the two-component dictionary as given in Table 3.3b.

Parameter	N	N_{int}	E_{int}
f	10	144	0.0655
$T_{1,A}$	15	224	0.0097
$T_{2,A}$	15	224	0.0044
$T_{1,B}$	15	224	0.0035
$T_{2,B}$	15	224	0.0017

To provide further insight in the quality of the parameter encoding for the two-component dictionary, plots were created that show the correlation between different atoms. Figure 4.8 shows the correlation for a single dictionary atom with its surrounding atoms, for (a) an atom in the dictionary given by Table 3.3a and (b) an atom in the dictionary given by Table 3.3b. Particularly the top right and bottom left plots of each tissue combination in this figure show very high correlations between different atoms. For example, from the top right plot of (a) in Figure 4.7, any value of the fraction f would match almost equally well as any other for a given value of $T_{2,B}$. The correlation plots here are only shown for a small number of parameter combinations.

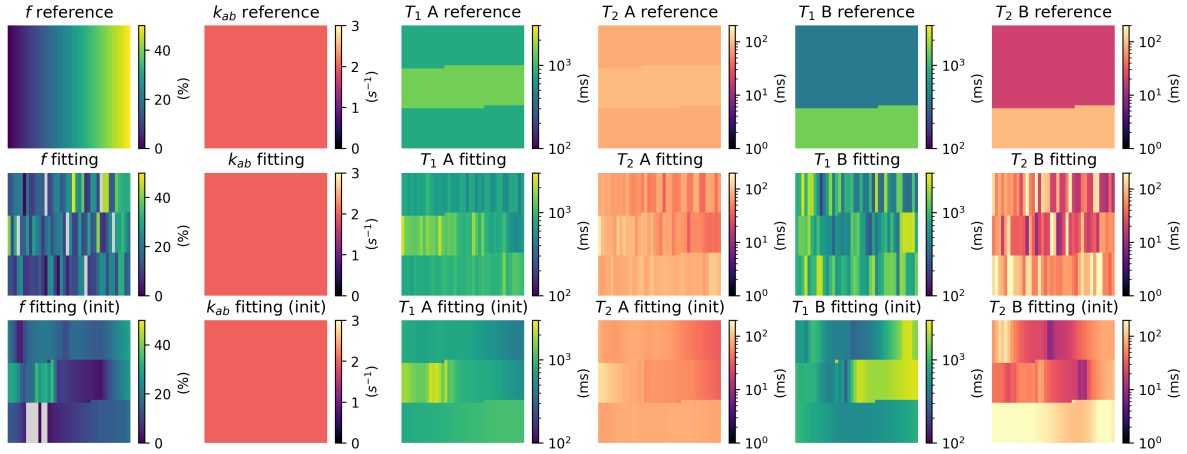

Figure 4.6: The parameters of the numerical phantoms (top row) and the parameter maps made using the dictionary fitting method. The middle row uses the result of inner-product matching as initialisation before the fitting. The bottom row uses a custom initialisation, with relaxation parameters as close to the reference values as possible, and a fraction of 0.21. The dictionary used is the two-component dictionary with fixed exchange rate. The values coloured grey are values outside of the expected range.

Full correlation maps for these atoms can be found in Appendix C.

4.3. Correction factor

Lastly, the results of parameter estimation using the dictionary with correction terms are presented here. Figure 4.9 shows the numerical phantom, together with the maps created using the correction-based method for 1 and 3 iterations. The estimation errors and standard deviations of the parameter maps are given in Table A.5. For both the middle and bottom row, the fraction map shows no visible gradient. Additionally, the relaxation parameter maps for $T_{1,A}$ show that this method either correctly estimates the relaxation time, or severely overestimates it. The map for $T_{2,B}$ shows a clear overestimation compared to the reference. $T_{1,B}$ shows a general underestimation of this parameter. Finally, $T_{2,B}$ shows both over- and underestimation of this parameter. Additionally, none of the relaxation parameter maps exhibit the uniformity expected when compared to the reference maps.

Figure 4.10 shows the result of a single iteration of the correction-based method on the signal from the voxel with position (4, 0) of the phantom. This signal was generated for a fraction of $f = 0$, where the only component is white matter. The figure shows the original signal, the modified signal after one iteration following equation (3.8), and the NNLS fit on this modified signal. The signal after one iteration shows clear modifications when compared to the original signal. Following the description of this technique however, for $f = 0$, there should not be any modification of the signal. The NNLS fit made to the modified signal is shown to closely approximate the signal.

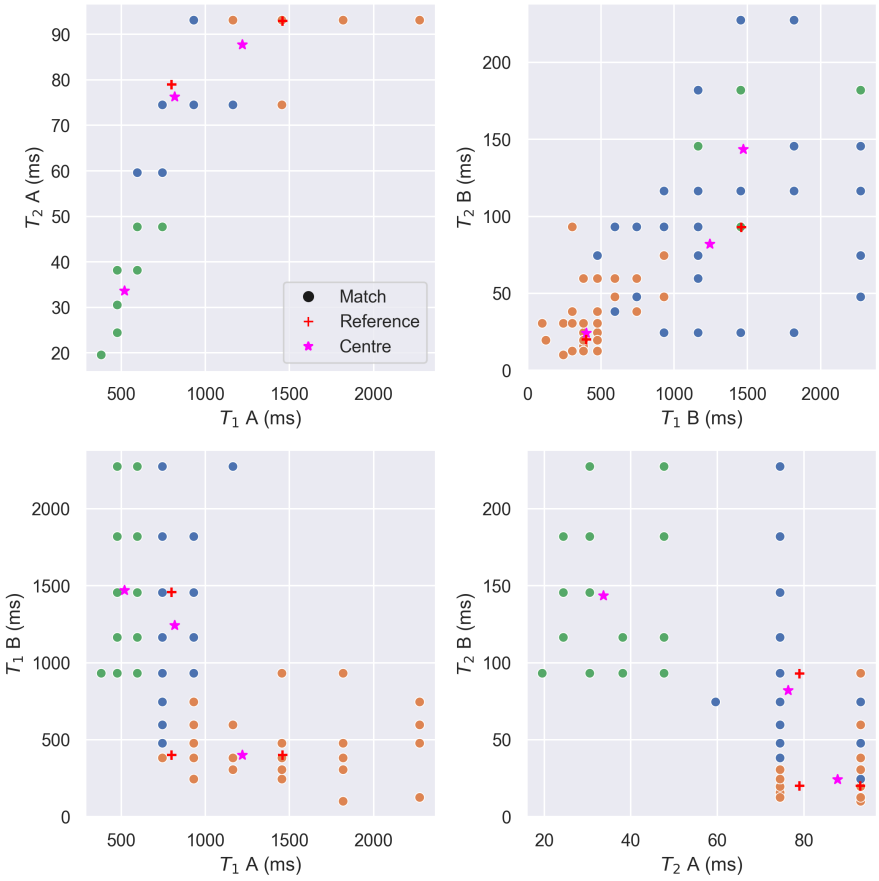

Figure 4.7: The components resulting from the initial inner-product matching on the numerical phantom as shown in Figure 4.5. All dots of the same colour have been assigned to the same cluster by the k-means clustering algorithm. The cluster centres are shown as stars, and the reference values are shown as plus signs.

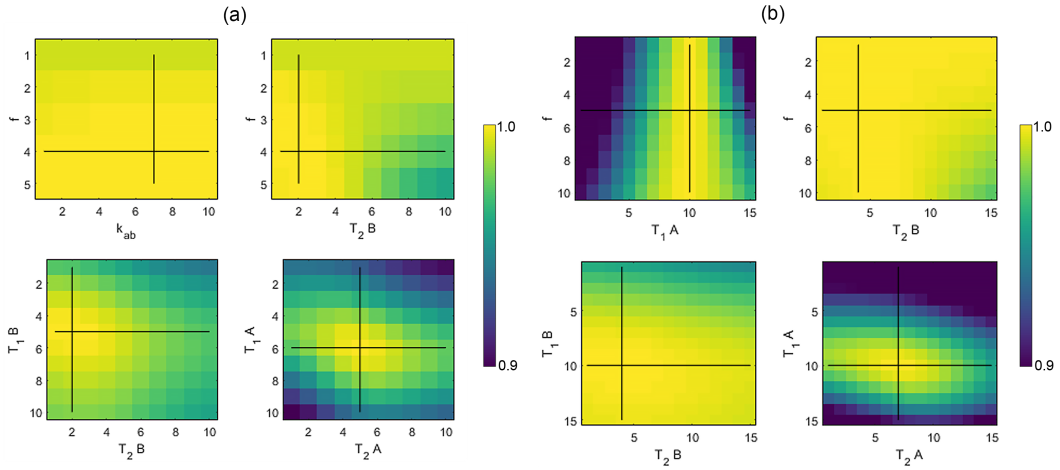

Figure 4.8: The correlation of one dictionary atom with its surrounding atoms for the atoms resembling the combination of white matter and myelin water for (a) the dictionary with all varying parameters. The fraction of this atom is chosen at 33 %, with an exchange rate of $2\,\mathrm{s}^{-1}$. The axes show the dictionary indices for the corresponding parameters. The range for f is [1, 5], the index range for all other parameters is [1, 10]. In (b), the correlation is given for the same tissue combination, but for the two-component dictionary with fixed exchange rate. The fraction here is chosen at 25 %. The axes show the dictionary indices for the corresponding parameters. The range for f is [1, 10], the index range for all other parameters is [1, 15]. The parameter values can be reconstructed from the indices using tables 3.3a and 3.3b respectively.

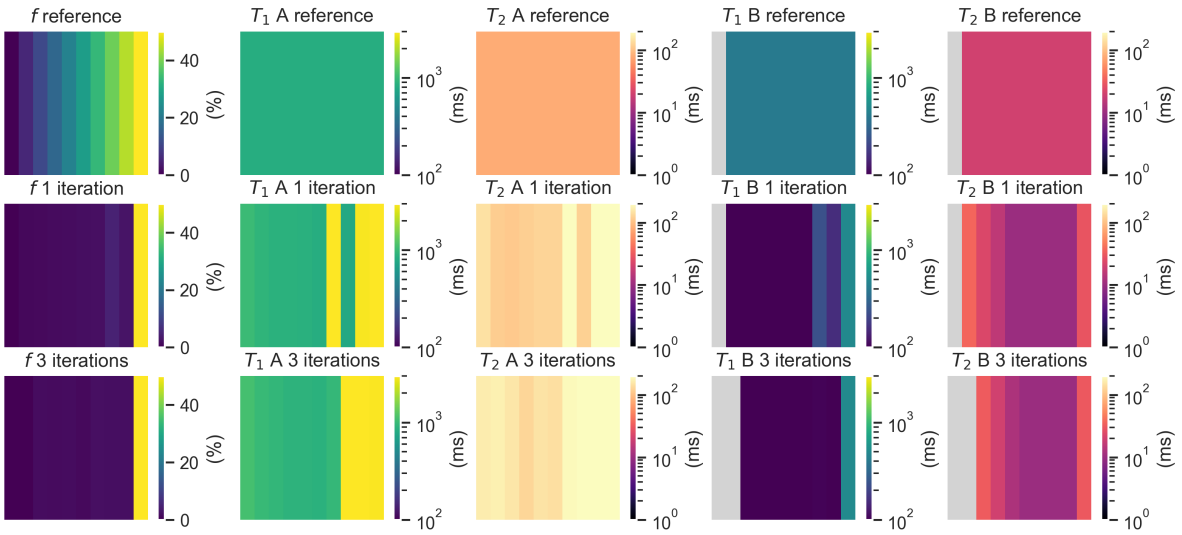

Figure 4.9: The parameters of the numerical phantom (top row) and the parameter maps made using the dictionary with correction terms for one iteration (middle row) and three iterations (bottom row). Grey areas in the B-component maps represent voxels where the mapped fraction is zero, so that there is no second component. The T_1- and T_2-boundary values used to separate the two components were 550 ms and 40 ms respectively.

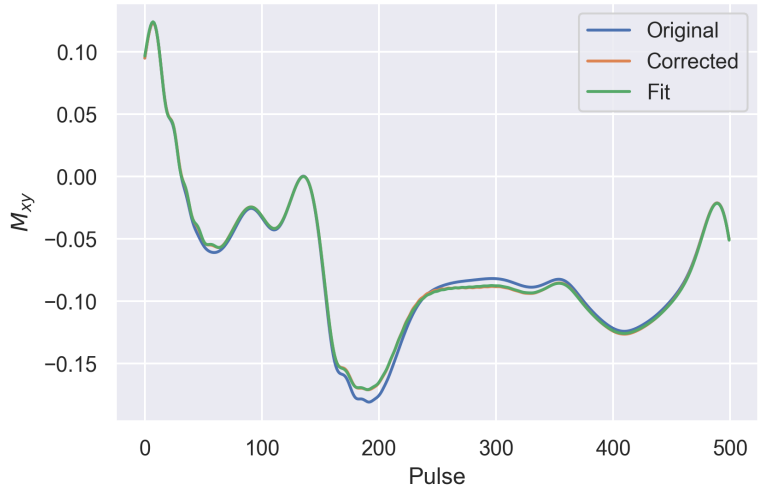

Figure 4.10: The result of one iteration of the correction term-based method. The original signal is the signal from the voxel with position (4, 0) in phantom (c), with a fraction of $f = 0$. The modified signal is shown with label "Corrected", the original signal before any correction is labelled "Original". The NNLS fit on the corrected signal is labelled "Fit".

5
Discussion

5.1. Magnetisation transfer in a linear multi-component model

Figure 4.1 shows that the fits made to the signals using NNLS are able to closely approximate the reference signal, even in the presence of magnetisation transfer. Figure 4.2 was made with the intention of studying the effects of varying strengths of magnetisation transfer on accuracy of the parameter mapping using NNLS and SPIJN. Focusing on the middle row of this image, it becomes clear that the magnetisation transfer has a measurable effect on the results. The fraction map of the voxel-wise NNLS shows a consistent underestimation of the fraction, with a larger bias for larger values of the exchange rate. At the same time, a biasing effect can be seen in the maps of the relaxation parameters of both component A and B. The bottom row, showing the parameter estimates of SPIJN, show a similar bias in the fraction map. This shows that the small changes to the signal as a result of MT result in NNLS estimating slightly different atoms and coefficients to still accurately approximate the target signal. This consequently leads to errors in fraction and relaxation parameters.

The more constant bias in the SPIJN fraction map (Figure 4.2), as opposed the NNLS map, suggests that most of the fraction bias is the result of an incorrect estimation of the relaxation parameters. When there is no variation in the selected components over the phantom, as is the case with SPIJN, the change in signal due to magnetisation transfer only leads to a small variation in retrieved fractions. When the components are allowed to vary between voxels, more variation becomes visible. The same trend can be seen in Figure 4.3, both for NNLS and SPIJN.

The result that magnetisation transfer in a linear multi-component model can influence the obtained results is in line with previous literature findings for similar techniques and problems. In [17], the effect of free pool-free pool magnetisation transfer on mcDESPOT is investigated. The research shows that in the presence of this two-pool exchange, it is not possible to obtain accurate parameter estimates using mcDESPOT.

The research in [37] shows increased estimation errors in both T_1 and T_2 using both an inversion recovery and saturation recovery MRF sequence. The focus of this research is on *semisolid* magnetisation transfer, instead of the free pool-free pool transfer here.

In [28], a sequence optimisation is used to optimise the precision in the estimates of the free pool parameters for an MRF acquisition. After that, a neural network was trained to do the parameter estimation step. However, this network is primarily trained to retrieve the free pool parameters, as opposed to the parameters of both pools. Additionally, like [37], this paper deals with semisolid magnetisation transfer, with only a single free pool. Here, we deal with two free pools.

5.2. Two-component dictionary

A two-component dictionary could have solved the two problems of the NNLS and SPIJN estimates. The first problem consists of the inaccuracies in the parameter mapping as a result of the magnetisation transfer. The second is the inability to map the exchange rate as a tissue parameter. The two-component dictionary that was created for this contained information about the effects of magnetisation transfer by explicitly modelling

the MT in the dictionary atoms. Assuming a sufficient parameter encoding, this would allow the dictionary to uniquely determine all six parameters simultaneously, while also taking the MT effects into account.

Figure 4.8 shows the correlation of a single dictionary atom with the surrounding atoms for several parameter combinations, with part (a) for the dictionary described above. Ideally, a correlation plot such as this would show a single high-correlation value for the atom of interest with itself, and low correlations for all other combinations. These plots, however, show high inter-atom correlation over the entire parameter range of the dictionary. This suggests that this dictionary does not allow distinguishing between different atoms, even when the parameter values of these atoms are far apart.

The results shown in Figure 4.4 further support this hypothesis. The large jumps in value between adjacent voxels show that a relatively small change in the signal can lead to a large change in the retrieved parameter values from the dictionary. This is a direct result of the high similarity between the atoms. In this figure, and the following parameter maps as well, the estimation error for the relaxation parameters of component B always seems to be larger than those of component A. This is not surprising, however, as component B has a smaller contribution to the final signal than A.

The clustering-based method, as applied in the bottom row of Figure 4.4, was an attempt at spatial regularisation. The underlying assumption here was that if a small number of tissue combinations is present in the phantom, the top matches of the individual voxels would "cluster" around the parameters of these tissue combinations. Even with large inter-atom correlation, which would introduce additional variation in the retrieved parameters, finding these groups and their corresponding "central parameters" might allow to find the correct tissue combinations. Assuming that the limited number of tissue combinations is the result of a limited number of tissue components present, this sets restrictions on all parameters but f.

The bottom row of Figure 4.4 shows that the clustering-based algorithm was also not able to create accurate maps.

The results presented thus far show that the inter-atom correlation of the two-component dictionary does not allow for an accurate parameter mapping. Figure 4.8 shows a particularly high correlation for f and k_{ab}. The second two-component dictionary, with a fixed value for the exchange rate, was created to eliminate errors as the result of this correlation, hopefully improving the parameter estimates for the remaining parameters (f in particular). The correlation maps in (b) in Figure 4.8, however, still show large inter-atom correlation for the remaining parameters. Combining this with the results in Figure 4.5, this shows that the remaining parameters are still not sufficiently and uniquely encoded in the dictionary signals.

The inaccurate parameter maps for the clustering-based method (Figure 4.5) can be explained using Figure 4.7. It becomes clear from this figure that the assumption that the top matching results would group around the reference values is incorrect. Where the reference values contain three distinct tissue combinations, the component plots show one wide "cloud" of points. This makes it difficult for the clustering algorithm to accurately retrieve the tissue combinations through the cluster centres. Especially in the top-left and bottom-right plot of Figure 4.7, the cluster centres and the reference points do not line up.

However, we do expect that with a better parameter encoding, the clustering-based method would be able to create accurate parameter maps. Even with the relatively continuous cloud of initial matches, the clustering is still regularly able to place the cluster centres close to the reference values. With better parameter encoding, we expect the groups to be more defined, which could then result in better clustering results. It might be worth looking into different clustering algorithms as well. K-means clustering requires a pre-set number of clusters. For general applications, it is better not to predefine the number of tissue components in advance.

The results of the dictionary fitting in Figure 4.6 show that the dictionary fitting does not provide a distinct advantage over inner-product matching. The same errors that are made by the inner-product matching are still present after the fitting on the interpolated dictionary. The second set of parameter maps, made using a manual initialisation before fitting, shows that even when starting at approximately the correct parameter values, the high inter-atom correlation dominates the results.

An additional problem shows when using dictionary fitting: the retrieved parameter values are no longer confined by the bounds of the parameters in the dictionary. This is most clearly visible in the retrieved fraction maps, where the grey values represent *negative* fractions returned by the fitting. Following the definition of the fraction, a negative value does not make physical sense. These negative values are the result of the interpolation extending beyond the given parameter ranges. This can be resolved by restricting the fitting algorithm to only fit on values within the initially given ranges.

The interpolation errors given in Table 4.2 are relatively high for the different parameters when compared to the target error value in [34]. This would suggest that, in order to accurately use dictionary fitting for a two component dictionary, either the parameter ranges have to be decreased, or the parameter sampling density has to be increased for the same ranges. Both options would lead to an increase in the size of the dictionary. For a two-component dictionary, a relatively small increase in the parameter sampling density will lead to a large increase in the number of dictionary atoms. When compared to the size of a conventional single-component dictionary (here 2064 atoms), even the smallest of the two two-component dictionaries is more than a hundred times as large, while having a coarser parameter sampling. A larger number of atoms leads to greatly increased simulation times for the dictionary. This is increased further by the fact that EPG-X is computationally more demanding than EPG. In order to create larger, more densely sampled two-component dictionaries, the implementation of the simulation has to be highly optimised. Additionally, larger dictionaries require more memory to store and process. This increases both the computational cost and processing time of matching with such a dictionary. Even with a highly optimised matching algorithm, the feasible size of the dictionary will be limited by the computational capabilities of the available computing system. One way to decrease dictionary size, which was not implemented during this project, is compression based on singular value decomposition (SVD) [38]. This would lower dictionary size and matching time.

5.3. Correction terms

The correction term-based method would allow us to reuse the MC-MRF algorithms for single-component dictionaries. However, implementation of the method with correction terms based on NNLS shows that such a general approach did not enable accurate parameter mapping, as shown in figure 4.9.

A first explanation of this is found in Figure 3.2. This figure suggests that the shape of the mean single-component correction term is relatively well-defined. However, the magnitude can exhibit a large variation, as can be seen from the standard deviation in the figure. Depending on the specific tissue combination, this might result in a significant overcompensation for the MT effects by the correction term. This effect increases with more iterations.

Another problem with this approach becomes clear in Figure 4.10. It shows that single-component signals, for which no MT is present at all, still have some correction applied to them. The explanation for this lies in the NNLS estimation. Often when a single-component signal does not perfectly coincide with a dictionary atom, NNLS will approximate the signal using *multiple* dictionary atoms. For each of these atoms, a correction term is now incorrectly applied to the signal. An approach to solving this problem would be to include a sparsity constraint *before* applying a correction on the signal. This would decrease the probability of multiple nonzero coefficients for single-component signals.

Finally, this method suffers from the same disadvantage as the other single-component dictionary-based methods in this report: it does not facilitate the retrieval of a parameter map for the exchange rate. More specifically, this method assumes a constant exchange rate throughout the sample, which is used to simulate the correction terms.

5.4. Further research

As shown in this report, magnetisation transfer can have negative effects on the quality of parameter estimates made with MC-MRF techniques. The methods tested in this report all show biases and errors in the estimated values in the presence of magnetisation transfer. Explorations of new or modified techniques, evaluated using numerical phantoms, have shown that the current implementations of these techniques are not able to create accurate parameter maps. More research is necessary to turn these techniques into useful solutions for more accurate MWF mapping.

An important technique left mainly unexplored in this report is *sequence optimisation*. It might be possible to specifically tailor the pulse sequence to minimise the effects of magnetisation transfer. Important to keep in mind here is that any changes to the pulse sequence should keep compromises to the parameter encoding to a minimum. Another option is to use sequence optimisation to improve the parameter encoding of the two-component signals to make the two-component dictionary a more viable technique. This would still require the aforementioned optimisations in software performance as well.

Appendix A includes parameter maps made with a two-component dictionary using a different pulse sequence, made to maximise T_1 and T_2 accuracy in MRF mapping with a specific undersampling pattern [39]. This was done to make sure that the inaccurate maps as shown in the results are not only the consequence of the used pulse sequence. The maps in the Appendix show similar errors to those made with the sequence shown in 2.2. This similar poor performance of two different pulse sequences indicates that a specialised pulse sequence is necessary to properly use a two-component dictionary.

6

Conclusion

Magnetic resonance imaging (MRI) is a versatile, non-invasive medical imaging technique capable of producing high-resolution, high-contrast in vivo images of biological tissue. MRI has been widely applied in neuroimaging, such as in the field of MS (multiple sclerosis) research. An important symptom of MS is demyelination, which is the breakdown of myelin sheaths in the brain. Myelin water fraction (MWF) imaging focuses on mapping the myelin water content of the brain. This myelin water fraction can be used as a proxy measure for myelin content, which generally has a very weak NMR signal.

This research focused on the use of multi-component magnetic resonance fingerprinting (MC-MRF) for estimating the MWF. This is conventionally done by constructing a linear combination of MRF dictionary atoms to approximate the measured signal. However, a nonlinear interaction, called magnetisation transfer, influences the total signal shape of the measured signal. This is not taken into account in current, state-of-the-art methods. Here, we focused on studying the effects of magnetisation transfer in the context of MC-MRF, with a focus on MWF imaging.

First, numerical phantoms were created, with each voxel containing two components, representing either white matter, grey matter, or myelin water. Two MC-MRF algorithms, voxel-wise non-negative least squares (NNLS) and SPIJN, were applied to this phantom. A single-component signal dictionary was used, which was simulated for a gradient-spoiled SSFP MRF pulse-sequence.

The parameter maps made using NNLS and SPIJN show that the magnetisation transfer has a clear biasing effect on the estimates. Both the retrieved fraction maps and the relaxation parameter maps for all components show a bias towards the lower values, with increased bias for higher values of the exchange rate. The more constant bias for SPIJN, as opposed to the increasing bias in NNLS, suggests that the larger estimation errors in the fraction are primarily the result of errors in the estimation of the relaxation parameters. When using the correct components, MT has a smaller influence on the fraction estimation.

Both the dictionary matching and dictionary fitting maps made using a two-component dictionary resulted in similar errors. The inaccurate results for both methods are caused by a large inter-atom correlation in the two-component dictionaries: a small change in measured signal will lead to a large change in retrieved parameter values. This means that the pulse sequence used is not able to properly encode every parameter individually. We suggest further research into the optimisation of the parameter encoding for a two-component dictionary. Additionally, more work has to be done to optimise the implementation of the software to make this a feasible technique.

For the correction term-based method, the results of an NNLS step are used to compute the correction on a signal. However, this does allow correction on single-component signals when no correction is expected. For compensation of potential magnetisation transfer effects, we suggest more research in the implementation of a sparsity constraint before applying the signal correction. Additionally, a more structured approach to the simulation of the correction terms might result in less overcompensation.

A final suggestion for further research is to optimise the pulse sequence to minimise the effects of magnetisation on the signals, without sacrificing parameter encoding. An important challenge here is to do this optimisation for all atoms simultaneously, which is a computationally heavy task.

Bibliography

[1] Jerry L. Prince and Jonathan M. Links. *Medical imaging signals and systems*. Pearson, Boston, 2 edition, 2015.

[2] Jan-Ray Liao, John M. Pauly, Thomas J. Brosnan, and Norbert J. Pelc. Reduction of motion artifacts in cine MRI using variable-density spiral trajectories. *Magnetic Resonance in Medicine*, 37(4):569–575, 1997. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910370416.

[3] Yong Chen, Yun Jiang, Shivani Pahwa, Dan Ma, Lan Lu, Michael D. Twieg, Katherine L. Wright, Nicole Seiberlich, Mark A. Griswold, and Vikas Gulani. MR Fingerprinting for Rapid Quantitative Abdominal Imaging. *Radiology*, 279(1):278–286, April 2016.

[4] Edward F. Jackson, Lawrence E. Ginsberg, Don F. Schomer, and Norman E. Leeds. A review of MRI pulse sequences and techniques in neuroimaging. *Surgical Neurology*, 47(2):185–199, February 1997.

[5] Gian Franco Piredda, Tom Hilbert, Jean-Philippe Thiran, and Tobias Kober. Probing myelin content of the human brain with MRI: A review. *Magnetic Resonance in Medicine*, 85(2):627–652, 2021. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.28509.

[6] Alex L. MacKay and Cornelia Laule. Magnetic Resonance of Myelin Water: An in vivo Marker for Myelin. *Brain Plasticity*, 2(1):71–91, January 2016. Publisher: IOS Press.

[7] R. Dobson and G. Giovannoni. Multiple sclerosis – a review. *European Journal of Neurology*, 26(1):27–40, 2019. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ene.13819.

[8] Wolfgang Brück, Andreas Bitsch, Herbert Kolenda, Yvonne Brück, Michael Stiefel, and Hans Lassmann. Inflammatory central nervous system demyelination: Correlation of magnetic resonance imaging findings with lesion pathology. *Annals of Neurology*, 42(5):783–793, 1997. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/ana.410420515.

[9] Eva Alonso-Ortiz, Ives R. Levesque, and G. Bruce Pike. MRI-based myelin water imaging: A technical review. *Magnetic Resonance in Medicine*, 73(1):70–81, 2015. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.25198.

[10] Alex Mackay, Kenneth Whittall, Julian Adler, David Li, Donald Paty, and Douglas Graeb. In vivo visualization of myelin water in brain by magnetic resonance. *Magnetic Resonance in Medicine*, 31(6):673–677, 1994. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910310614.

[11] Christian Labadie, Jing-Huei Lee, William D. Rooney, Silvia Jarchow, Monique Aubert-Frécon, Charles S. Springer Jr., and Harald E. Möller. Myelin water mapping by spatially regularized longitudinal relaxographic imaging at high magnetic fields. *Magnetic Resonance in Medicine*, 71(1):375–387, 2014. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.24670.

[12] Sean C.L. Deoni, Brian K. Rutt, Tarunya Arun, Carlo Pierpaoli, and Derek K. Jones. Gleaning multicomponent T1 and T2 information from steady-state imaging data. *Magnetic Resonance in Medicine*, 60(6):1372–1387, 2008. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.21704.

[13] Dan Ma, Vikas Gulani, Nicole Seiberlich, Kecheng Liu, Jeffrey L. Sunshine, Jeffrey L. Duerk, and Mark A. Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, March 2013. Number: 7440 Publisher: Nature Publishing Group.

[14] Tom Hilbert, Ding Xia, Kai Tobias Block, Zidan Yu, Riccardo Lattanzi, Daniel K. Sodickson, Tobias Kober, and Martijn A. Cloos. Magnetization transfer in magnetic resonance fingerprinting. *Magnetic Resonance in Medicine*, 84(1):128–141, 2020. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.28096.

[15] Tanja Platt, Mark E. Ladd, and Daniel Paech. 7 Tesla and Beyond. *Investigative Radiology*, 56(11):705–725, November 2021.

[16] Vikas Gulani and Nicole Seiberlich. Quantitative MRI: Rationale and Challenges. In Nicole Seiberlich, Vikas Gulani, Fernando Calamante, Adrienne Campbell-Washburn, Mariya Doneva, Houchun Harry Hu, and Steven Sourbron, editors, *Advances in Magnetic Resonance Technology and Applications*, volume 1 of *Quantitative Magnetic Resonance Imaging*, pages xxxvii–li. Academic Press, January 2020.

[17] Christopher L. Lankford and Mark D. Does. On the inherent precision of mcDESPOT. *Magnetic Resonance in Medicine*, 69(1):127–136, 2013. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.24241.

[18] Megan E. Poorman, Michele N. Martin, Dan Ma, Debra F. McGivney, Vikas Gulani, Mark A. Griswold, and Kathryn E. Keenan. Magnetic resonance fingerprinting Part 1: Potential uses, current challenges, and recommendations. *Journal of Magnetic Resonance Imaging*, 51(3):675–692, 2020. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.26836.

[19] K Sommer, T Amthor, Peter Koken, J Meineke, and Mariya Ivanova Doneva. Determination of the optimum pattern length of MRF sequences. *Proc. Intl. Mag. Reson. Med.*, 25, 2017.

[20] ML Lauzon. A Beginner's Guide to Bloch Equation Simulations of Magnetic Resonance Imaging Sequences. page 18.

[21] Christian Guenthner, Thomas Amthor, Mariya Doneva, and Sebastian Kozerke. A unifying view on extended phase graphs and Bloch simulations for quantitative MRI. *Scientific Reports*, 11(1):21289, October 2021. Number: 1 Publisher: Nature Publishing Group.

[22] Matthias Weigel. Extended phase graphs: Dephasing, RF pulses, and echoes - pure and simple. *Journal of Magnetic Resonance Imaging*, 41(2):266–295, 2015. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.24619.

[23] H. Y. Carr and E. M. Purcell. Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Physical Review*, 94(3):630–638, May 1954. Publisher: American Physical Society.

[24] Yun Jiang, Dan Ma, Nicole Seiberlich, Vikas Gulani, and Mark A. Griswold. MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout. *Magnetic Resonance in Medicine*, 74(6):1621–1631, 2015. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.25559.

[25] Martijn Nagtegaal, Peter Koken, Thomas Amthor, and Mariya Doneva. Fast multi-component analysis using a joint sparsity constraint for MR fingerprinting. *Magnetic Resonance in Medicine*, 83(2):521–534, February 2020.

[26] Sunli Tang, Carlos Fernandez-Granda, Sylvain Lannuzel, Brett Bernstein, Riccardo Lattanzi, Martijn Cloos, Florian Knoll, and Jakob Assländer. Multicompartment magnetic resonance fingerprinting. *Inverse Problems*, 34(9):094005, July 2018. Publisher: IOP Publishing.

[27] Nicole Seiberlich, Vikas Gulani, Adrienne Campbell, Steven Sourbron, Mariya Ivanova Doneva, Fernando Calamante, and Houchun Harry Hu. *Quantitative Magnetic Resonance Imaging*. Academic Press, November 2020. Google-Books-ID: 65ngDwAAQBAJ.

[28] Jakob Assländer, Cem Gultekin, Andrew Mao, Xiaoxia Zhang, Quentin Duchemin, Kangning Liu, Robert W. Charlson, Timothy Shepherd, Carlos Fernandez-Granda, and Sebastian Flassbeck. Rapid quantitative magnetization transfer imaging: utilizing the hybrid state and the generalized Bloch model, July 2022. arXiv:2207.08259 [physics].

[29] Shaihan J. Malik, Rui Pedro A.G. Teixeira, and Joseph V. Hajnal. Extended phase graph formalism for systems with magnetization transfer and exchange. *Magnetic Resonance in Medicine*, 80(2):767–779, 2018. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.27040.

[30] Janaka P. Wansapura, Scott K. Holland, R. Scott Dunn, and William S. Ball Jr. NMR relaxation times in the human brain at 3.0 tesla. *Journal of Magnetic Resonance Imaging*, 9(4):531–538, 1999. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291522-2586%28199904%299%3A4%3C531%3A%3AAID-JMRI4%3E3.0.CO%3B2-L.

[31] Python Software Foundation. Python 3.10.9, 2022.

[32] The MathWorks Inc. MATLAB version R2022B, 2022.

[33] The SciPy project. SciPy version 1.9.3, 2022.

[34] Willem van Valenberg, Stefan Klein, Frans M. Vos, Kirsten Koolstra, Lucas J. van Vliet, and Dirk H. J. Poot. An Efficient Method for Multi-Parameter Mapping in Quantitative MRI Using B-Spline Interpolation. *IEEE Transactions on Medical Imaging*, 39(5):1681–1689, May 2020. Conference Name: IEEE Transactions on Medical Imaging.

[35] John A. Hartigan. *Clustering Algorithms*. Wiley, 1975. Google-Books-ID: cDnvAAAAMAAJ.

[36] Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, and David Cournapeau. Scikit-learn: Machine Learning in Python. *MACHINE LEARNING IN PYTHON*.

[37] Yuguang Meng, Jesse Cheung, and Phillip Zhe Sun. Improved MR fingerprinting for relaxation measurement in the presence of semisolid magnetization transfer. *Magnetic Resonance in Medicine*, 84(2):727–737, 2020. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.28159.

[38] Debra F. McGivney, Eric Pierre, Dan Ma, Yun Jiang, Haris Saybasili, Vikas Gulani, and Mark A. Griswold. SVD Compression for Magnetic Resonance Fingerprinting in the Time Domain. *IEEE transactions on medical imaging*, 33(12):2311–2322, December 2014.

[39] D Heesterbeek. Mitigating undersampling errors in magnetic resonance fingerprinting by sequence optimization.

A

Supplementary results

In this appendix, additional results are presented. This information is presented separately from the main Results chapter to improve readability of this report.

A.1. NNLS and SPIJN

Here, the estimation errors for NNLS and SPIJN are presented. They are given in Table A.1

Table A.1: The mean estimation error and standard deviation of the parameter maps made using NNLS and SPIJN as shown in Figure 4.2. The units of each parameter are given in the rightmost column.

(a) The mean errors and standard deviations (σ) for the parameter maps made using voxel-wise NNLS.

Parameter	Mean	σ	Units
f	-12	8	%
$T_{1,A}$	-86	59	ms
$T_{2,A}$	-14	9	ms
$T_{1,B}$	-51	58	ms
$T_{2,B}$	-8	3	ms

(b) The mean errors and standard deviations (σ) for the parameter maps made using SPIJN.

Parameter	Mean	σ	Units
f	-11	4	%
$T_{1,A}$	-83	0	ms
$T_{2,A}$	-13	0	ms
$T_{1,B}$	-51	0	ms
$T_{2,B}$	-11	0	ms

A.2. Two-component dictionary

Below, the estimation errors for the various maps made using the two-component dictionaries are presented. The first table (Table A.2) shows the errors for the maps made with the dictionary with varying exchange rate. Table A.3 shows the errors for the maps made with the dictionary with fixed exchange rate. Table 4.2 shows the errors for the maps made using the dictionary fitting technique.

Table A.2: The mean error and its standard deviation for the parameter maps made using inner-product matching and the clustering-based method using the two-component dictionary with varying exchange rate. The units of each parameter are given in the rightmost column.

(a) The mean estimation errors and standard deviations (σ) for the parameter maps made using inner-product matching.

Parameter	Mean	σ	Units
f	2	15	%
$T_{1,A}$	-78	440	ms
$T_{2,A}$	-26	25	ms
$T_{1,B}$	781	746	ms
$T_{2,B}$	213	283	ms

(b) The mean estimation errors and standard deviations (σ) for the parameter maps made using the clustering-based technique.

Parameter	Mean	σ	Units
f	-19	27	%
$T_{1,A}$	-138	0	ms
$T_{2,A}$	-32	0	ms
$T_{1,B}$	565	0	ms
$T_{2,B}$	109	0	ms

Table A.3: The mean error and its standard deviation for the parameter maps made using inner-product matching and the clustering-based method with the two-component dictionary with fixed exchange rate. The units of each parameter are given in the rightmost column.

(a) The mean estimation errors and standard deviations (σ) for the parameter maps made using inner-product matching.

Parameter	Mean	σ	Units
f	-8	15	%
$T_{1,A}$	1	344	ms
$T_{2,A}$	-8	20	ms
$T_{1,B}$	143	739	ms
$T_{2,B}$	25	64	ms

(b) The mean estimation errors and standard deviations (σ) for the parameter maps made using the clustering-based technique.

Parameter	Mean	σ	Units
f	-0.3	19	%
$T_{1,A}$	-86	365	ms
$T_{2,A}$	-11	25	ms
$T_{1,B}$	98	808	ms
$T_{2,B}$	27	71	ms

Table A.4: The mean estimation error and standard deviation (σ) for the parameter maps made using dictionary fitting on the interpolated dictionary. The units of each parameter are given in the rightmost column.

(a) The mean estimation errors and standard deviations (σ) for the parameter maps made by fitting on the interpolated dictionary. The initialisation used here are the results of the dictionary matching.

Parameter	Mean	σ	Units
f	-7	15	%
$T_{1,A}$	-92	322	ms
$T_{2,A}$	-10	21	ms
$T_{1,B}$	217	690	ms
$T_{2,B}$	30	59	ms

(b) The mean estimation errors and standard deviations (σ) for the parameter maps made by fitting on the interpolated dictionary. The initialisation used here is a custom initialisation close the the target values.

Parameter	Mean	σ	Units
f	-11	16	%
$T_{1,A}$	-90	368	ms
$T_{2,A}$	-13	20	ms
$T_{1,B}$	215	929	ms
$T_{2,B}$	46	49	ms

A.3. Correction term

The estimation errors for the maps made using the dictionary with correction term are presented in Table A.5.

Table A.5: The errors and standard deviations of the parameter maps acquired using the dictionary with correction terms. The units of each parameter are given in the rightmost column.

(a) The mean estimation errors and standard deviations (σ) for the parameter maps acquired after one iteration.

Parameter	Mean	σ	Units
f	-15	20	%
$T_{1,A}$	675	958	ms
$T_{2,A}$	307	401	ms
$T_{1,B}$	-234	127	ms
$T_{2,B}$	-3	9	ms

(b) The mean estimation errors and standard deviations (σ) for the parameter maps acquired after three iterations.

Parameter	Mean	σ	Units
f	-15	21	%
$T_{1,A}$	729	932	ms
$T_{2,A}$	331	386	ms
$T_{1,B}$	-248	135	ms
$T_{2,B}$	-3	9	ms

A.4. Different pulse sequence

Figure A.2 shows the results of inner-product matching and the clustering-based method using a two-component dictionary with phantom (c) (Figure 3.1). The pulse sequence used is shown in Figure A.1. This is a gradient spoiled SSFP sequence that is optimised to minimise MRF mapping errors for T_1 and T_2, with a pulse train of 400 pulses [39]. The sequence is preceded by a 180° inversion pulse with 15 ms inversion time. The two-component dictionary was made with a fixed value for the exchange rate at $k_{ab} = 2\,\mathrm{s}^{-1}$, with parameter ranges as given in Table 3.3b. The estimation errors and standard deviations are given in Table A.6.

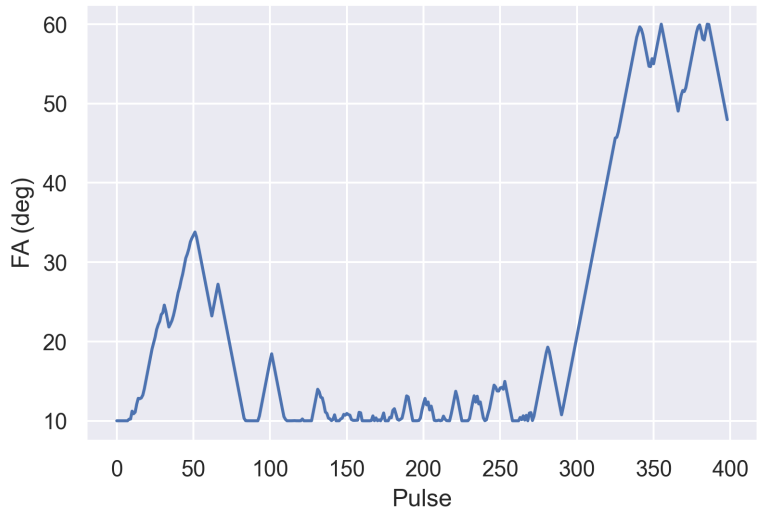

Figure A.1: The flip angle pattern of the SSFP pulse sequence optimised for minimal MRF T_1 and T_2 estimation errors. The 180° inversion pulse at the beginning of the sequence is not shown.

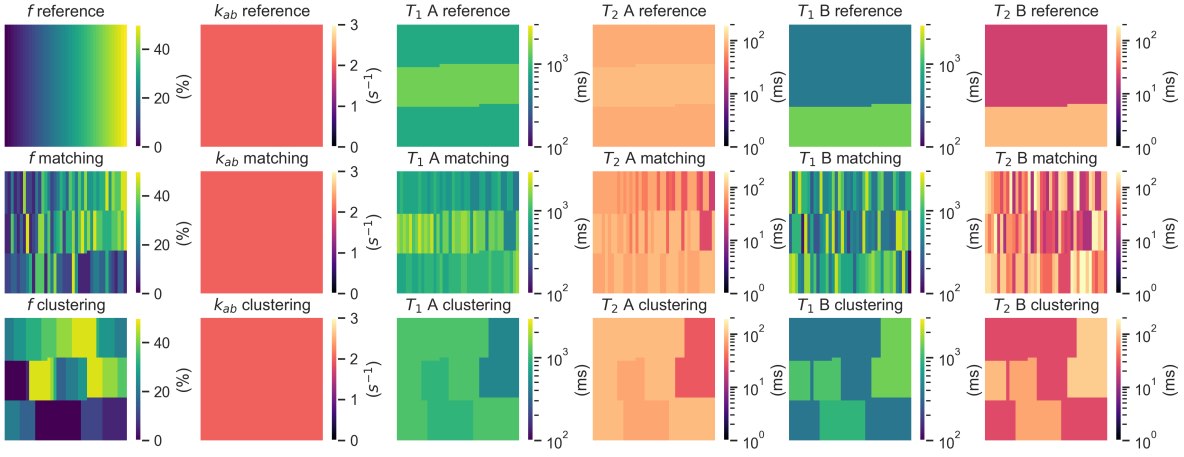

Figure A.2: The parameter maps of the numerical phantom (top row) and the parameter maps made using inner-product matching (middle row) and the clustering-based method (bottom row) on this phantom.

Table A.6: The mean error and the standard deviation of the parameter maps shown in Figure A.2. The units of the parameters are given in the rightmost column.

(a) The mean estimation errors and standard deviations for the parameter maps made using inner-product matching.

Parameter	**Mean**	σ	**Units**
f	1	25	%
$T_{1,A}$	13	345	ms
$T_{2,A}$	10	21	ms
$T_{1,B}$	-165	764	ms
$T_{2,B}$	-30	62	ms

(b) The mean estimation errors and standard deviations for the parameter maps made using the clustering-based method.

Parameter	**Mean**	σ	**Units**
f	1	20	%
$T_{1,A}$	30	455	ms
$T_{2,A}$	6	26	ms
$T_{1,B}$	0	751	ms
$T_{2,B}$	-11	56	ms

B

Sequence optimisation

A quick exploration of the possibility to use sequence optimisation to decrease the influence of magnetisation transfer was done near the end of this project. Here, the approach and findings are shown to illustrate the claims in the discussion.

Sequence parameterisation

The problem of sequence optimisation was approached as a *minimisation problem*. Specifically, the flip angle series was to be optimised such that the effects of the magnetisation transfer were minimal. Starting point was the flip angle series used in this the rest of this report. Since this pulse series has a length of 500 pulses, optimising the flip angle of *every* pulse would be computationally very expensive. Instead, the flip angle train was parameterised using 25 *knots*, between which a cubic interpolation was applied.

The initial positions of the knots were determined by applying an optimisation that minimised the squared distance between the original flip angle sequence and the interpolated sequence. An additional constraint was added to prevent the interpolated sequence from adopting negative values. The result of this approximation is shown in Figure B.1.

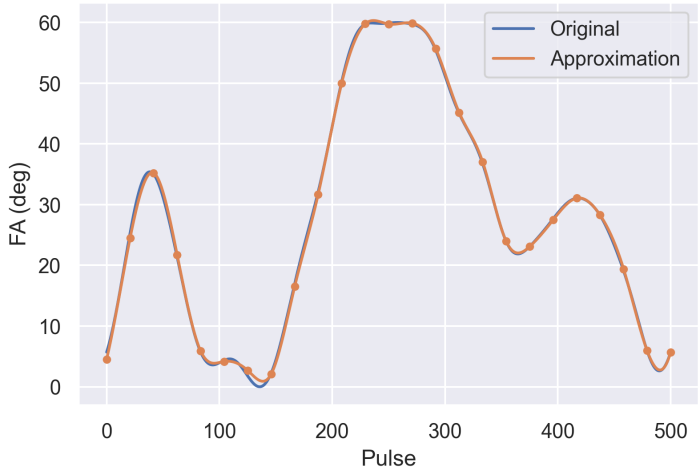

Figure B.1: The flip angle series, approximated using a cubic interpolation on 25 knots (shows as dots here). The original flip angle train is shown for reference.

Cost function and constraints

Since the goal of this optimisation problem was to minimise the effects of MT, the cost function was defined as the maximum absolute MT effect:

$$\text{cost} = \max|\boldsymbol{s} - \widetilde{\boldsymbol{s}}|, \tag{B.1}$$

where $\boldsymbol{s}$ and $\widetilde{\boldsymbol{s}}$ are the signals simulated with and without magnetisation transfer respectively.

Additionally, constraints were used to ensure the correct behaviour of the optimisation. First, similar to the creation of the pulse sequence optimisation, the interpolated flip angle sequence was not allowed to be negative. Secondly, the mean deviation of the knots from their original position was not allowed to be more than 15°. This prevents the solver from finding the "trivial solution". This is a flip angle sequence of only 0° RF pulses. Since there would be no signal then, the MT effect would be zero as well. This is obviously not the desired solution.

Results

Optimisation was implemented using the `scipy.optimise.minimise` method [33]. Since this was done mainly as a *proof of concept*, the cost function was evaluated only for signals resulting from the combination of white matter and myelin water.

Starting from the cubic approximation of the original pulse sequence, the minimisation was applied for 25 iterations, with a total of 628 cost function evaluations. The resulting optimised pulse series is shown in Figure B.2.

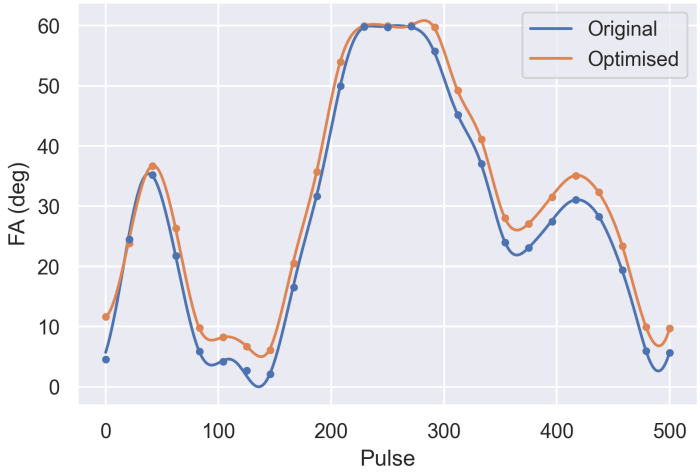

Figure B.2: The results of 25 iterations of minimisation on the parameterised flip angle series. The original (approximated) flip angle series is shown for reference. The dots represent the knots used for the cubic interpolation.

The optimised flip angle train reduced the value of the cost function from 0.00822 to 0.00804. This shows that it is possible to reduce the effects of magnetisation transfer through sequence optimisation. However, since this is only an example of optimisation with respect to a single tissue combination, the extent to which the optimisation could work for a larger number of tissue combinations simultaneously is still uncertain. Additionally, this small reduction in cost for only two tissues took several hours of computation time. An improvement in computational efficiency is also needed to further test this approach.

C

Full correlation maps

In this appendix, the full correlation maps for the two two-component dictionaries are given. They show the correlation of a single dictionary atom with all surrounding atoms. The atom of interest for the two figures presented here is the atom representing the combination of white matter and myelin water. Figure C.1 shows the correlation maps for the two-component dictionary given in Table 3.3a. Figure C.2 shows the correlation maps for the two-component dictionary given in Table 3.3b.

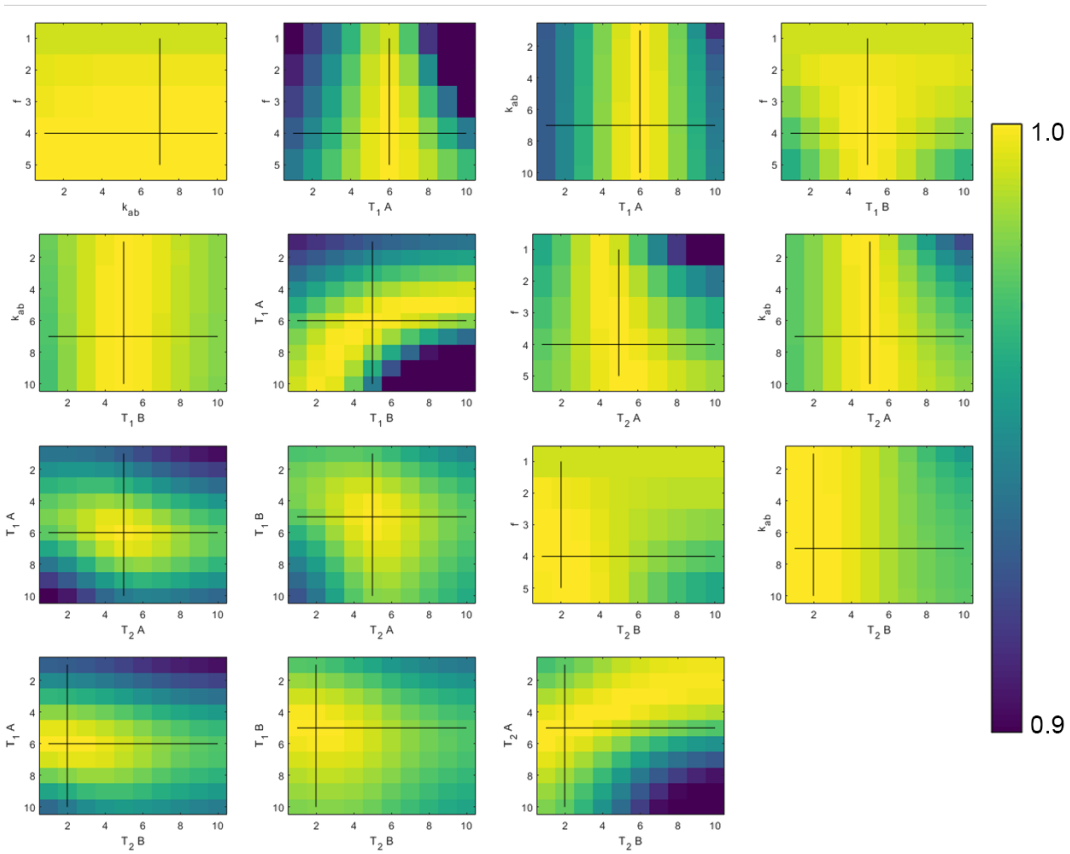

Figure C.1: Correlation plots showing the correlation of one dictionary atom with its surrounding atoms. The crosshair shows the position of the atom of interest. This atom represents the combination of white matter and myelin water, with an exchange rate of 2 s^{-1} and a fraction of 33 %. The axes represent the indices of the corresponding parameter in the dictionary. For the fraction, the range is [1 - 5]. For all other parameters, the range is [1 - 10]. The values of the parameters can be retrieved from the indices using Table 3.3a.

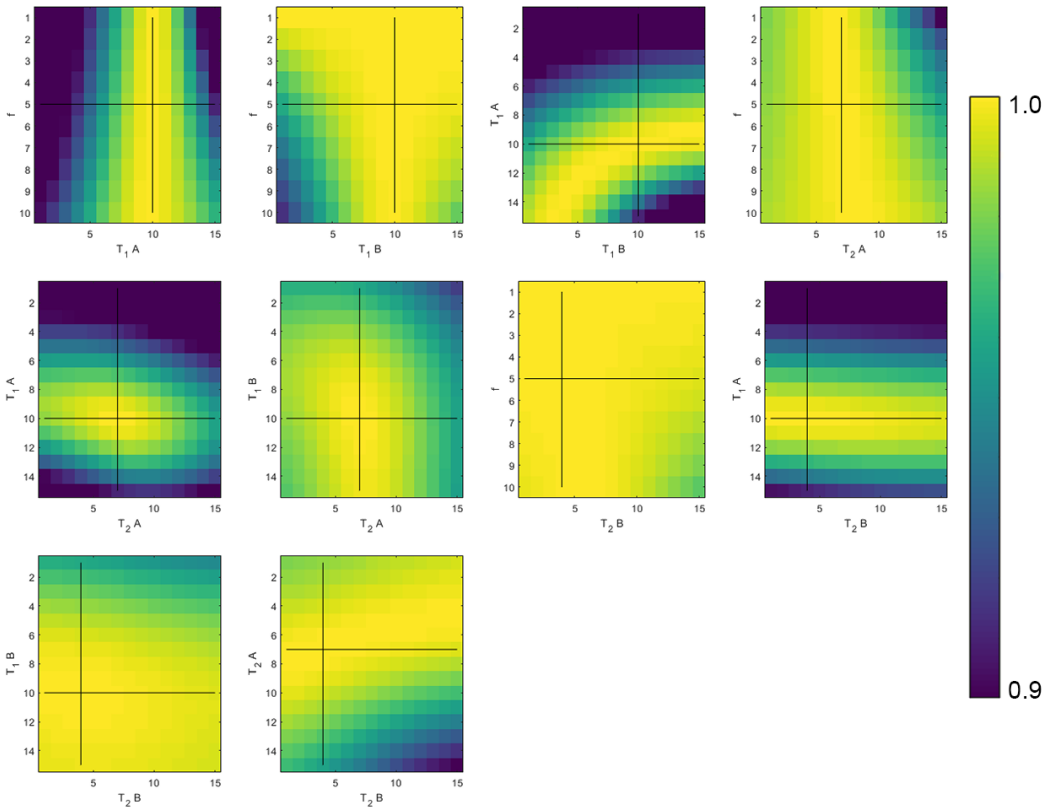

Figure C.2: Correlation plots showing the correlation of one dictionary atom with its surrounding atoms. The cross shows the position of the atom of interest. This atom represents the combination of white matter and myelin water, with an exchange rate of $2\,s^{-1}$ and a fraction of 25 %. The axes represent the indices of the corresponding parameter in the dictionary. For the fraction, the range is [1 - 10]. For all other parameters, the range is [1 - 15]. The values of the parameters can be retrieved from the indices using Table 3.3b.