Regional variation models of white matter microstructure

Gemma Louise Morgan

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Computer Science

University College London

2012

I, Gemma Louise Morgan, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Diffusion-weighted MRI (DW-MRI) is a powerful in vivo imaging technique that is particularly sensitive to the underlying microstructure of white matter tissue in the brain. Many models of the DW-MRI signal exist that allow us to relate the signals we measure to various aspects of the tissue structure, including measures of diffusivity, cellularity and even axon size. From histology, we know that many of these microstructure measures display distinct patterns of variation on length scales greater than the average voxel size. However very few methods exist that use this spatial coherence to inform and guide parameter estimation. Instead, most techniques treat each voxel of data independently. This is particularly problematic when estimating parameters such as axon radius which only weakly influence the signal, as the resulting estimates are noisy. Several methods have been proposed that spatially smooth parameter estimates after fitting the model in each voxel. However if the parameter estimates are very noisy, the underlying trend is likely to be obscured. These methods are also unable to account for spatial coupling that may exist between the various parameters.

This thesis introduces a novel framework, the Regional Variation Model (RVM), which exploits the underlying spatial coherence within white matter tracts to estimate trends of microstructure variation across large regions of interest. We fit curves describing parameter variation directly to the diffusion-weighted signals which should capture spatial changes in a more natural way as well as reducing the effects of noise. This allows for more precise estimates of a range of microstructure indices, including axon radius. The resulting curves, which show how microstructure parameters vary spatially through white matter regions, can also be used to detect groupwise differences with potentially greater power than traditional methods.

We begin by validating the framework using both simulated data and 'silver standard' ex vivo monkey data. We show that in both cases, the RVM is able to capture the underlying trends of parameter variation, even when the data acquisition is modest and the SNR is low. Next we show how the curves of parameter variation produced by the RVM can be used in groupwise studies to detect and localise parameter changes. Using a simulation study we demonstrate that the RVM is able to detect small effect sizes even when the number of 'subjects' is small. Our proposed method compares favourably to more established techniques such as region of interest (ROI) analysis and tract-based spatial statistics (TBSS): the RVM provides better spatial localisation than ROI analysis and has greater power to detect changes between small group sizes compared to TBSS. We then use the RVM to detect changes due to healthy ageing in the corpus callosum. Using only a modest number of subjects we are able to detect significant

increases in diffusivity with age in both the genu and the splenium. Finally, we demonstrate how the RVM can be used to improve estimation of low sensitivity parameters such as axon radius. Simulation studies show that the RVM can extract the correct trend of axon radius variation even at low gradient strengths, outperforming independent voxelwise estimation. We also estimate axon radius using the RVM and voxelwise estimation in the mid-sagittal corpus callosum in both ex vivo monkey and in vivo human data sets, acquired at a range of gradient strengths. As in the simulation study, the RVM is able to recover the expected trend of axon radius variation with less noise, even at lower gradient strengths. However, at the lowest clinical gradient strength of 40 mT/m, neither technique is able to extract the expected trends, indicating that there is no sensitivity to the range of axon radii found in the human corpus callosum at this gradient strength.

Acknowledgements

I want to thank my supervisor Professor Daniel Alexander, for hiring me for the project to begin with, but most importantly for providing very necessary guidance and advice over the past four years, without which this thesis would never have been written. I am also indebted to Dr Brandon Whitcher for all his help and encouragement and for guiding me through the relevant statistics, and Dr Gary Hui Zhang for his insightful comments and suggestions.

I am also very indebted to my examiners, Professor Derek Jones and Professor David Hawkes, for a tough, but ultimately enjoyable viva, and for their comments and suggestions which have helped to improve the final version of this thesis.

I would like to thank my secondary supervisor Professor Sebastien Ourselin for asking probing questions throughout my first year and transfer vivas that have helped to refine the project aims over the years. Without data for validation, the ideas presented wouldnt mean much. Therefore I'm very grateful to Dr Rex Newbould and Dr Tim Dyrby for supplying a lot of the data used here and all their help and suggestions.

Many thanks to Laura Panagiotaki, my best friend from the very start of the PhD, for keeping me sane in the office, having huge amounts of fun outside the office and graciously giving up her time, despite planning a wedding, to read through this thesis. I'd also like to thank the rest of the microstructure imaging group, particularly Ivana Drobnjak and Bernard Siow for their insights into MR physics, and the friends I've made here along the way: Matt Hall, Hubert Fonteijn, Torben Schneider, Kiran Seunarine, Ben Hall, Jenny Steeden, Ai Wern Chung.

Special thanks to Sara McCluskey, Julie Heyward and John Geddes - without their help and support I would not be here today.

Very special thanks to my husband, Shahrum. Thank you for understanding what it means to undertake a PhD, for reading through the drafts and most of all, for loving me.

Finally, I want to thank my parents, Margaret and Philip, my sister and brother-in-law, Claire and Steve, and my beautiful nephew, Jacob, for providing endless love, support and encouragement, and for making me the person I am today.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Medical imaging allows us to peer deep within the body; to examine organs and diagnose pathologies that were previously inaccessible, except using invasive surgeries and biopsies. Although invaluable for studying all regions and organs within the body, medical imaging provides a particularly necessary window into the brain, an organ whose deepest structures are inaccessible even to biopsies, and for which intervention carries a high risk of additional trauma. Using imaging we are able to examine both fine structural details and functional activity using different combinations of modalities and contrast enhancements, the exact choice of which depends on the pathology under investigation and the type of information we want to gather.

Magnetic resonance imaging (MRI) is one of the most commonly used imaging modalities, encompassing a wide range of contrast mechanisms. Using MRI, it is possible to obtain high resolution images of gross macroscopic anatomy (T_1 or T_2-weighted images), detect functional activity (functional MRI), measure blood flow (perfusion MRI) and calculate drug uptake rates using tracers (dynamic contrast-enhanced MRI). We can also gain greater insight into the microstructural composition of tissue using diffusion-weighted MRI.

Diffusion-weighted MRI is sensitive to the displacement of molecules, typically water molecules which are highly prevalent in the human body. Due to thermal fluctuations, these water molecules undertake random walks known as self-diffusion. In an open environment free of barriers, the particles can diffuse freely and are equally likely to move in all directions. However barriers or structures within the environment hinder or restrict the water molecules, influencing the overall dispersion pattern of the water. Diffusion MRI provides contrast by attenuating the overall MR signal based on the average amount of water dispersion along a specified direction. The greater the dispersion, the more the signal is attenuated. During a typical diffusion MRI scan, signals are measured along multiple directions in order to build up a three-dimensional dispersion profile, providing important information about the directionality as well as the amount of water diffusion.

Due to the way in which it provides contrast, diffusion MRI is particularly useful for studying white matter in the brain. White matter provides the network of cables which connect the grey matter processing centres within the brain, and consists of bundles, or tracts, of parallel axonal fibres. Each tract is highly directional and contains thousands of individual axons with diameters ranging from 0.1μm to

20μm [187]. On a standard T_1 or T_2 structural scan, white matter typically appears as a homogeneous region, and whilst we are able to differentiate white matter from grey matter and cerebrospinal fluid, we cannot distinguish the individual tracts or extract more fundamental microstructure features. However, due to its orientational dependence, diffusion MRI allows us to tease apart the white matter into separate fibre bundles with different directionalities. Furthermore, diffusion MRI is exquisitely sensitive to the underlying shape and geometry of white matter; the probability density functions describing water dispersion not only reflect the direction of the tract, but also the size of the axons, the amount of extracellular space and the presence of other cell bodies. This is potentially important for studying white matter diseases such as multiple sclerosis (MS) and schizophrenia, in which the structure of the axons and surrounding tissue is degraded. This alters the dispersion patterns of water molecules in affected regions of the brain. Therefore, if we can reliably extract specific, quantitative measures which reflect these changes from diffusion MRI, we can use them as biomarkers to improve understanding, diagnosis and treatments for white matter diseases.

Mathematical models which relate the measured signal attenuations to specific properties of interest allow us to obtain very detailed information about the structure of the tracts. One of the simplest and most common models used is the diffusion tensor (DT) model [20]. The DT model fits a covariance matrix to the measured signals, characterising the amount of diffusion within a given reference frame. From the DT we can derive useful measures such as the principal direction of diffusion, the mean diffusivity and the diffusion anisotropy [23]. These parameters can be useful indicators for white matter tissue development or pathology, but they are also reasonably non-specific. Therefore, there has been a move towards more specific markers using microstructure imaging, focusing on parameters such as axon diameters and densities [7, 15, 168]. However, the dependence of the MR signal upon these parameters is generally weak, particularly from data acquired using clinical hardware with low gradient strengths. As a result, these indices are particularly susceptible to noise and are difficult to estimate robustly, limiting their clinical utility.

The MR signals we measure are noisy for a number of reasons. Current fluctuations in the signal detection and processing equipment introduce electronic noise into the measurement. However, the main source of noise is the tissue we are imaging itself [118]. Ions such as sodium or potassium generate currents within the body that can cause additional electronic noise. Whereas careful calibration and fine tuning can reduce noise from the equipment, we cannot eliminate the noise from tissue. Because diffusion-weighted MRI provides contrast based on signal attenuation, the signal to noise ratio (SNR) of diffusion MR data is often lower than data acquired using other MR contrast mechanisms. This hinders our ability to estimate microstructure parameters, from simple markers such as anisotropy to more specific parameters such as axon radius indices.

There are several options available for increasing the SNR of the data during acquisition, namely:

- increasing the magnetic field strength,
- increasing the voxel size,
- averaging data from multiple rescans,

- using multi-channel RF coils.

The magnetic field strength is an intrinsic property of the scanner, and therefore cannot be raised above its given value. Increasing the voxel size is undesirable as it reduces the spatial resolution of the acquired images. This makes it more difficult to identify specific structures within the images and increases the number of voxels with partial volume effects, i.e. voxels in which the measured signal is due to multiple tissue types. Finally, performing multiple rescans is often impractical. For example, the acquisition protocols used by Alexander et al [7] to estimate axon radius require a scan time of one hour; therefore repeating the scan would exceed human tolerance within the scanner. Rescanning can also be impractical for shorter DTI protocols, which only probe between 6 and 15 diffusion directions. In a clinical setting, one of the key aims during imaging is to minimise the time that patients spend within the scanner. Therefore, structural T_1 and T_2 scans are often acquired in conjunction with diffusion-weighted scans, reducing the time available to repeat any of the protocols.

Although options for improving the SNR of the data during acquisition are often limited, we can potentially improve the effective SNR of the data and thus the SNR of our estimated parameter maps by exploiting the dependencies between signals in neighbouring voxels. Many key white matter tracts exhibit regional patterns in axon size and density [2]. This means that neighbouring voxels have a high degree of spatial coherence in terms of underlying microstructure indices, and therefore should exhibit similar diffusion MR signals. These dependencies are ignored by naive but widely used voxel by voxel parameter estimation techniques; however if we harness this spatial information during parameter estimation, we can potentially improve the robustness of our fitted parameters to image noise.

1.1 Problem statement

Image noise affects all voxels within medical images, regardless of modality or contrast. Many neuroimaging studies performing model fitting on the individual voxel level under the assumption that each voxel of data is independent. This makes parameter estimation more susceptible to the effects of image noise, reducing the accuracy and precision of potential biomarkers derived from the data. In this thesis, we aim to address this problem by devising algorithms that exploit the spatial coherence of the underlying histological parameters, and thus the measured signals, in order to improve the SNR of estimated parameter maps. We then aim to investigate and demonstrate the experimental and practical benefit of these algorithms in a selection of exemplar applications.

1.2 Project aims

The overall goal of this work is to propose the regional variation model (RVM), a general framework that exploits the underlying spatial coherence of parameters during model fitting. The RVM estimates trends describing the spatial variation of tissue indices across a tract or region of interest (ROI) instead of on a voxel scale. It uses information about the shape and structure of the ROI to pool data from all voxels in a natural, intuitive way. The trends are fit directly to the diffusion MR signals from all voxels simultaneously. This should increase the effective SNR of our data, resulting in robust and reliable

estimates of the regional parameter trends. The specific aims of this PhD thesis are:

- to implement the framework,
- to demonstrate the successful use of the RVM framework in exemplar applications from diffusion MRI.

Although the focus of this work is the application of the RVM to diffusion MR data, we present the RVM as a general purpose framework so that it can be developed and used more broadly by other researchers in the future working with other modalities or contrast mechanisms.

1.3 Contributions

The contributions of this thesis are as follows:

- General formulation of the RVM framework.
- Implementation of the RVM framework. Throughout the thesis, we present three implementations of the RVM using different mathematical models of the diffusion MR signal. The diffusion models we present are:
  1. Ball and stick model [26]. A simple diffusion model to estimate intracellular volume fraction, fibre orientation and water diffusivity.
  2. ZeppelinCylinder model [6, 134]. This model is used to estimate the axon radius index, a summary statistic which corresponds to the mean of the axon radius distribution within the voxel weighted by axon volume, as well as the intracellular volume fraction, fibre orientation and water diffusivity.
  3. ZeppelinCylinderDot model [7, 134]. As above, but for ex vivo tissue in which a stationary water compartment is needed to account for the observed restriction.
- Validation of the RVM framework using the ball and stick implementation, with synthetic data. Using simulated data generated with a range of noise levels and gradient directions, we demonstrate that the RVM is able to accurately and precisely recover trends of multiple tissue indices simultaneously and investigate the conditions in which the RVM outperforms standard independent voxelwise estimation.
- Validation of the RVM framework using the ball and stick implementation, with high quality ex vivo monkey brain data. We subsample the full data set (which serves as a ground truth) in order to generate lower quality data sets containing measurements along fewer gradient directions. We demonstrate that as the number of gradient directions sampled decreases, parameter trends estimated using the RVM reflect the ground truth better than independent voxelwise parameter estimates.
- Demonstration of how the RVM can be used in population analyses in order to detect and localise groupwise differences. We compare the results obtained using the RVM with current state of the

art methods such tract-based spatial statistics and region of interest analysis. Using synthetic data with a range of noise levels and gradient directions, we investigate the acquisition conditions under which the RVM performs more favourably. We also demonstrate the use of the RVM in a group study of healthy ageing.

- Demonstration of how the RVM can be used in microstructure imaging to estimate an axon radius index. Using simulated data generated for a range of noise levels and gradient strengths, we investigate the minimum gradient strength required to recover axon radius trends using both the RVM and current state of the art estimation methods. We also present and compare parameter maps of axon radius index estimated using the RVM and independent voxelwise estimation, from both ex vivo and in vivo data acquired at a range of gradient strengths.

1.4 Thesis outline

Chapter 2 provides detailed background information on brain white matter, MRI and diffusion MRI as well as related work which aims to model trends of diffusion MRI indices within tracts. We also provide details of several mathematical techniques which underpin the methods developed within this thesis and are used throughout the work. Chapter 3 introduces the main methodological contribution of the thesis, outlining the regional variation model (RVM) framework before providing more detailed information about the key components of the technique. Chapter 4 is the first of three experimental chapters, which aims to validate the RVM using both synthetic and real MRI data. Chapters 5 and 6 present applications of the RVM using synthetic, ex vivo and in vivo brain data. Chapter 5 demonstrates how the trends estimated using the RVM can be used in population studies of healthy ageing, and compares the results with those obtained using well established, state of the art methods. Chapter 6 uses the RVM to estimate trends in axon radius index from synthetic, ex vivo monkey and in vivo human brain data acquired using a range of maximum gradient strengths. By comparing the parameter maps obtained using the RVM with those estimated using the current state of the art fitting techniques, we aim to determine under which conditions the RVM provides an advantage over voxelwise estimation methods. Chapter 7 summarises the main contributions of the work as a whole and draws together the key conclusions from all chapters. We assess the RVM relative to the contributions stated in section 1.3 and discuss several interesting avenues for future work which could extend and improve upon the existing implementation of the RVM.

Chapter 2

Background

This chapter begins with a description of the structure and function of healthy white matter tissue in the brain. We then discuss how this structure and function is impaired due to different types of pathology in order to motivate the need for specific and sensitive white matter biomarkers. We introduce standard techniques for estimating these types of white matter indices from histology as well as various imaging modalities, including MRI. We provide a brief overview of the physical principles behind the nuclear magnetic resonance (NMR) phenomenon, signal generation and image reconstruction, as well the different MRI constrast mechanisms available, before focusing on one technique, diffusion-weighted MRI, in particular. We discuss diffusion and diffusion-weighted imaging, before reviewing a wide range of analytical models commonly used to elucidate the measured diffusion-weighted MR signals and provide markers for white matter microstructure. As discussed in chapter 1 the key aim of this thesis is to develop parameter estimation algorithms which exploit the underlying microstructure coherence using both shape and spatial information. Therefore we give an overview of shape models used in white matter modelling and discuss recent studies that have attempted to combine macroscopic shape information with microstructure information from diffusion MRI to provide more specific markers of tissue change. We also briefly describe some spatial regularisation techniques that allow local spatial neighbourhood information to be incorporated into parameter estimation. The chapter closes with an overview of some mathematical techniques, B-splines and Bayesian model fitting, that are used extensively in the methods developed within this thesis.

2.1 White matter

Tissue in the brain can be broadly divided into two classes: grey matter and white matter. Within the skull, the brain is enclosed in a membrane called the meninges, which is filled with cerebrospinal fluid (CSF) [130]. The CSF provides a protective layer to cushion the brain, but also facilitates the transportation of chemicals within the brain [130]. Grey matter is mainly comprised of neuronal cell bodies [112], which are the information processing centres of the brain. Cell bodies which perform similar tasks are often grouped together to form larger functional areas called nuclei [130]. However the brain is a complex organ and performing even simple tasks often requires communication between several functional regions. Connections are needed that can carry and transmit signals between different

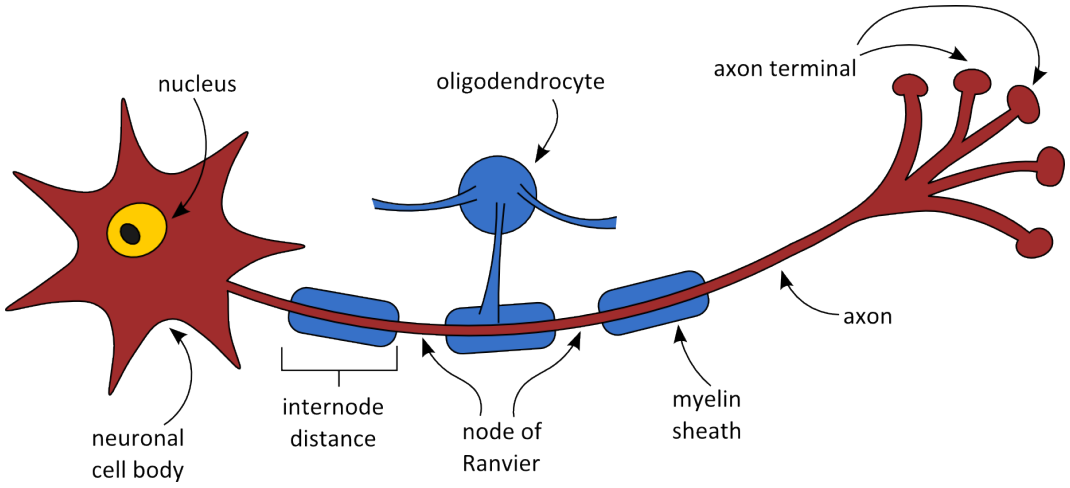

Figure 2.1: A schematic diagram showing an axonal projection from the neuronal cell body. Axons are often covered by myelin sheaths, produced by the oligodendrocyte glial cells, which increase the propagation speed of electrical signals along the axons.

grey matter areas quickly and efficiently. These connections are formed by axons, which are projections from the neuronal cell bodies. Regions of the brain which primarily contain axons are referred to as white matter. Glial cells, which are found within both grey and white matter regions, support and repair the surrounding tissue. Because this work is concerned with modelling microstructure variation within white matter, we now focus our discussion on the structure, function and pathology of the axons and glial cells that constitute white matter.

2.1.1 Structure and function

White matter primarily contains axons and glial cells. The axons are arranged in bundles of parallel fibres, known as tracts, with the glial cells (oligodendrocytes, astrocytes and microglia) nestled between the fibres [112]. As mentioned in the previous section, axons are elongated, approximately cylindrical projections from neuronal cell bodies that connect neurons within and between functional areas in grey matter. Figure 2.1 shows a schematic diagram of an axonal projection.

The key function of axons is to transmit chemical and electrical impulses between the neurons, which then process and interpret the signals [130]. Chemical signals are transported within the axons whereas electrical signals are propagated along a plasma membrane covering the axon. The length of axons in the central nervous system (CNS) ranges from micrometres to metres and their diameters vary from 0.1 μm to 20 μm [187]. The speed of the electrical impulses along the axon depends directly on axon diameter. Small diameter axons have high resistance resulting in slow signal speeds; large diameter axons have lower resistance and thus have higher signal velocities [181].

Nearly all axons in the CNS with diameters greater than 2 μm are covered in myelin sheaths. Myelin comprises layers of lipids and proteins [108] and is produced by the oligodendrocyte glial cells [130]. Specifically, myelin is formed when a projection from the oligodendrocyte cell flattens out and wraps around a segment of the axon, forming a spiralling sheath [187]. Figure 2.2, taken from [108],

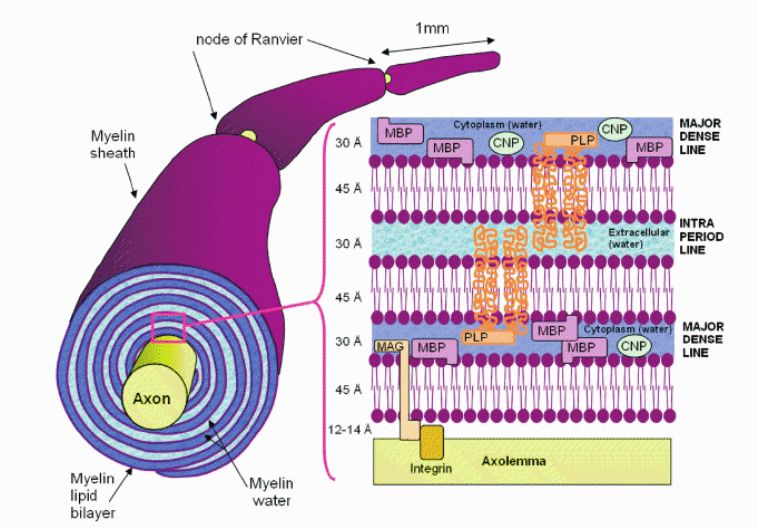

Figure 2.2: Most axons over 2μm in diameter are surrounded by myelin sheaths which comprise protein-lipid bilayers. This figure, taken from [108], illustrates the spiralling structure of the myelin which is produced as projections from the oligodendrocyte cells flatten out and wrap around the axons.

illustrates the structure of the myelin. One oligodendrocyte cell may have multiple projections which provide myelin sheaths for several neighbouring axons, as shown in figure 2.3. As each oligodendrocyte projection only provides myelin for a small segment of the axon, many oligodendrocytes are needed to myelinate the whole length of the axon. Between the myelin segments are uncovered sections called nodes of Ranvier [187] and the distance between adjacent nodes is termed the internode distance [130] (see figure 2.1). For myelinated fibres, the electrical impulses propagate by jumping between adjacent nodes [187], therefore the longer the internode distance the faster the signal propagates. The purpose of the myelin sheath is to prevent conduction leaks and to decrease the capacitance of the axon, thus increasing the signal speed further still [181]. The thicker the myelin sheath, the lower the capacitance. Both the internode distance and the thickness of the myelin sheath are proportional to the diameter of the axon [187], therefore increasing the speed of signal propagation along large diameter axons the most. As well as the oligodendrocyte glial cells, there are two other types of glial cells present in white matter: astrocytes and microglia [187]. The astrocytes maintain and support the structure of the myelin and the axons [187] whereas the microglia respond to damage and disease, removing dead cells and pathogens within the white matter [130].

The structure of the axons and myelin is related to the function of the grey matter regions they connect. As an example we consider the corpus callosum (CC), a large white matter tract that connects many different functional regions between the left and right hemispheres. The left and right prefrontal cortices are connected through the genu, at the front of the CC. These regions process complex decision making [85], which requires a large amount of information transfer although it does not need to be processed instantaneously. Consequently, the axon radii in the genu are small but densely packed [2, 85]. However, the left and right motor cortices, which are connected through the midbody of the CC, require

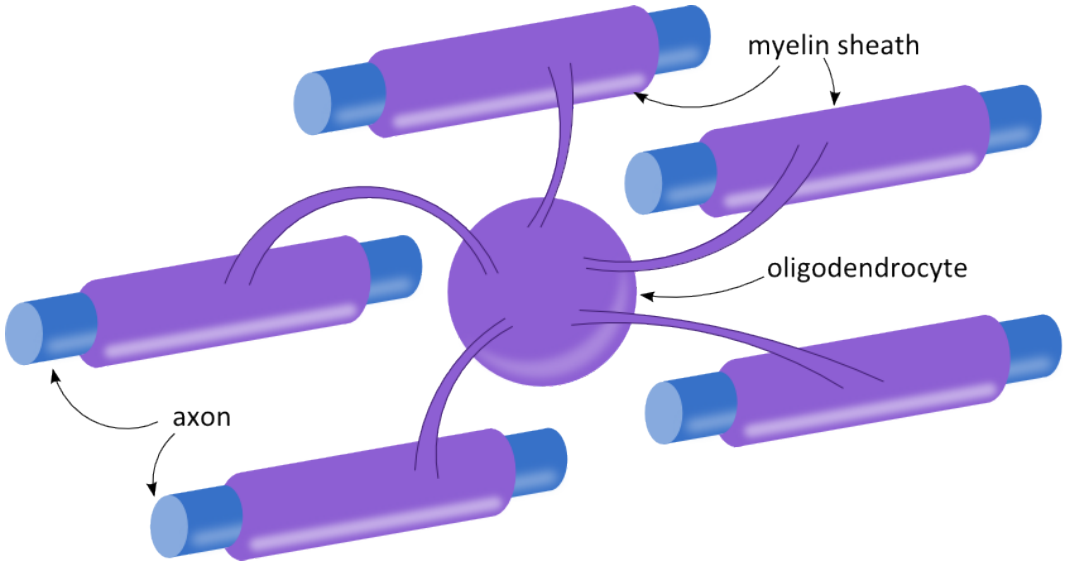

Figure 2.3: Oligodendrocytes can produce multiple projections, each of which flattens out to provide a segment of myelin sheath for an axon. A single oligodendrocyte can provide and support the myelin for between three and ten individual axons [187].

very fast signal processing which is reflected in the large size of axons in this region [2]. Axons at the back of the CC, in the splenium, which connect auditory and high level visual areas [85] are again smaller. This results in a distinct low-high-low pattern of axon radius (and corresponding high-low-high pattern in axon density) in the CC, which is consistent across subjects [2] and species [106, 131]. Patterns of axon radius and density variation can also be found in other white matter tracts such as the optic nerve [92] and spinal cord [53].

2.1.2 Pathology

Any damage to the myelin sheaths, the oligodendrocytes or the axons themselves can disrupt the electrical signals leading to abnormal brain function [187]. In this section, we explore some of the most common pathologies that occur in white matter and explain how they affect both structure and function. Axonal damage can be broadly divided into three categories: dysmyelination, demyelination and Wallerian degeneration. Both dysmyelination and demyelination are characterised by myelin damage followed by axonal death [181] whereas Wallerian degeneration is caused by axonal or neuronal death followed by secondary myelin decay [61].

Dysmyelination is the failure to form normal myelin during myelinogenesis (the process of myelination), resulting in lumpy, grainy or loose myelin sheaths [61]. These irregularities interrupt the transmission of electrical signals, resulting in impaired functional ability [181]. Dysmyelination is thought to be the cause of leukodystrophies, a group of illnesses characterised by gradual loss of high level functions such as motor ability, vision, hearing and behaviour [181]. The functional impairment is gradual as myelinogenesis is a slow process which continues until early adulthood [139]. For many of the functional networks affected, the axons which transport the electrical signals are not myelinated until late

childhood, so these symptoms can take years to manifest.

Demyelination is the loss of previously healthy myelin around the axon [61]. This degeneration can be caused either by damage to the myelin itself or disruption to the oligodendrocytes supplying and supporting the myelin [181]. Demyelination begins as the myelin slowly unravels around the axon. As the myelin loosens, it is less effective at insulating the electrical impulses travelling along the plasma membrane, leading to a reduction in propagation speed [86]. Demyelination occurs in many white matter diseases, such as multiple sclerosis (MS) and schizophrenia, although the mechanisms causing it can be very different. In MS, the myelin loosens due to inflammation of the axons exerting pressure on the tightly bound myelin layers [86], whereas in schizophrenia, demyelination is thought to occur as mutated oligodendrocytes can no longer support and repair the myelin when the protein-lipid layers start to break apart [179].

Wallerian degeneration can either be axonal or neuronal. Axonal Wallerian degeneration is caused by direct damage to the axon, usually due to traumatic injury. The portion of the axon distal to the damage site decays, although the neuronal cell body from which the axon projects continues to function and may develop additional pathways to replace the damaged axon [61]. Neuronal Wallerian degeneration begins with the death of the neuronal cell body. After cell necrosis, all axonal projections from the cell decay. For both forms of Wallerian degeneration, the myelin sheaths remain intact for some time (up to months or years) after the axons' death [61] but eventually break down leaving debris amongst the remaining white matter pathways.

The types of damage we have discussed here all affect the propagation of electrical and chemical signals along axons, resulting in impaired function. However, the mechanisms behind the tissue changes are different and therefore result in different types of white matter disorders characterised by a range of symptoms. In order to fully understand, differentiate between and track the progression of these disorders, we need to be able to distinguish and localise these subtle changes.

2.1.3 Studying white matter

In this section we discuss some of the techniques that are currently used to study white matter diseases, potentially allowing us to distinguish between the different forms of pathology described in the previous section. Traditionally, histology has been used to determine the bulk properties and microstructure of both grey and white matter. Careful dissection has allowed mapping of the major white matter pathways [167], whilst the advent of microscopy and tissue staining has allowed quantities such as axon size and density [2, 84] and glial cell density [178] to be measured. Electron microscopy is able to differentiate between pathologies such as dysmyelination and demyelination and to localise where these tissue changes occur. Histology is often seen as the gold standard; however there are some problems associated with it. First, the time delay between sample acquisition and tissue fixation is crucial as tissue decay, which begins immediately after death, can confound results [161]. Second, samples are usually preserved in either formaldehyde or paraffin which is known to shrink and distort tissue [148]. Third, obtaining samples for clinical use via biopsy is often difficult, painful and sometimes, in the case of brain tissue, impossible to obtain. Finally, longitudinal studies are not possible using histology.

Over the past 30 years or so, in vivo imaging techniques such as positron emission tomography (PET), computed tomography (CT) and MRI have alleviated the need for painful biopsies to a certain extent. PET is particularly useful for imaging physiology and function whereas CT is a very good structural imaging modality [188]. Using MRI, we can image both structure and function. One of the key advantages of using MRI over PET or CT is that it does not require the use of any ionising radiation [188]. Within MRI, there are numerous methods available for providing image contrast that enable us to extract different structural information about the tissue. Some of the most commonly used image contrasts are T_1-, T_2-, proton density (*pd*-) and diffusion-weighting (there are also techniques available to measure flow [76], perfusion [35], susceptibility [78] and other phenomena that we do not discuss here). Although the underlying physical mechanisms of these different types of contrast are intrinsically linked to the microstructure of the tissue, we cannot extract this microstructure information from T_1-, T_2- or *pd*-weighted images directly. However, because diffusion-weighted MR is sensitive to the displacement of water molecules, by measuring the average amount of diffusion along different directions we can make inferences about the structure of the environment in which the water is diffusing [21]. Because of this diffusion-weighted MRI, along with magnetisation transfer (MT) and relaxometry [108], is one of the key techniques for studying white matter microstructure. In the next sections, we describe the general principles behind MRI and the different image contrasts available, before focusing on diffusion MRI and its application to the study of white matter in the brain.

2.2 Magnetic Resonance Imaging

MRI uses the NMR signal from individual nuclei in atoms and molecules to build images of the internal structure and function of macroscopic objects [33]. In this section, we present a simple semi-classical description of the NMR phenomenon (although strictly a quantum effect, classical physics may be used when dealing with large ensembles of particles [113]), and signal generation. We briefly discuss how signals are localised and images are reconstructed, before finishing the section with a description of some of the different image contrasts available.

2.2.1 Foundations of MRI

The nuclear magnetic moment

The NMR phenomenon occurs due to a property of the nucleus called spin $\mathbf{J}$ [113]. Classically, $\mathbf{J}$ is visualised as the rotation of the nucleus about its own axis, as shown in figure 2.4, although from a quantum mechanical perspective it is simply angular momentum. This angular momentum, in conjunction with the electrical charge of the nucleus, induces a magnetic field known as the nuclear magnetic moment $\boldsymbol{\mu}$ [113]. The nuclear magnetic moment is related to $\mathbf{J}$ via

$$\boldsymbol{\mu} = \gamma \mathbf{J}, \tag{2.1}$$

where γ is the gyromagnetic ratio, a nucleus dependent constant [190]. The magnitude of the nuclear magnetic moment μ is

$$\mu = |\boldsymbol{\mu}| = \gamma \hbar \sqrt{I(I+1)} \tag{2.2}$$

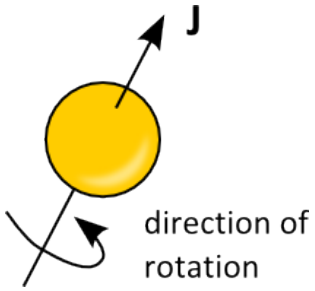

Figure 2.4: A nucleus with spin vector **J**. Classically spin is visualised as the rotation of the nucleus about the axis of **J**

where I is the nuclear spin quantum number and $\hbar$ is Planck's constant divided by 2π. I depends on the number of protons Z within the nucleus and its mass number A. The nuclear spin quantum number can be

- half integral if A is odd
- integral if A is even and Z is odd
- zero if both A and Z are even.

If I=0, the nucleus is NMR inactive [113]. Water, the main constituent of tissue in the human body, contains two hydrogen (^{1}H) atoms within the molecule. Water is NMR active as the nuclear spin quantum number for hydrogen is $\frac{1}{2}$ [190], and we shall focus on this simple but important example during the following sections.

Bulk magnetisation

With no external magnetic field present, the direction of $\boldsymbol{\mu}$ is random due to thermal movements, so within an ensemble of spins the magnetic moments cancel each other out resulting in zero net magnetisation [190]. However, during an NMR experiment, we want to measure the bulk nuclear magnetisation; therefore we need to find a way of aligning the magnetic moments and generating a net magnetisation vector. This can be done using an external magnetic field $\mathbf{B}_0$ with magnitude B_0 and direction $\hat{\mathbf{z}}$. When this external field is turned on, two things occur:

- the magnetic moments precess about the $\hat{\mathbf{z}}$ axis
- the magnetic moments align relative to the direction of the $\mathbf{B}_0$ field, generating a bulk magnetisation vector.

We now discuss these effects in turn.

First, the nuclear magnetic moments precess about $\mathbf{B}_0$ [33], i.e. they rotate about the axis of the $\mathbf{B}_0$ field. This occurs because the spins experience a torque $\mathbf{T}$ due to the $\mathbf{B}_0$ field. This torque is given by

$$\mathbf{T} = \frac{d\mathbf{J}}{dt} = \frac{1}{\gamma}\frac{d\boldsymbol{\mu}}{dt} = \boldsymbol{\mu} \times \mathbf{B}_0. \tag{2.3}$$

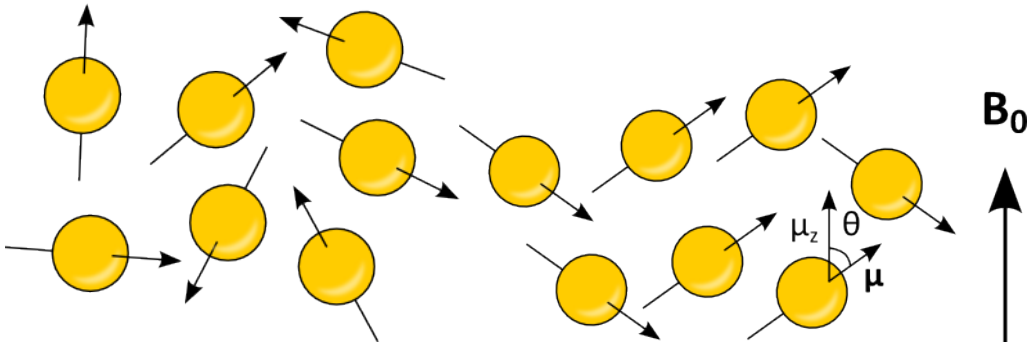

(a) (b)

Figure 2.5: (a) shows the random orientations of **J** in an ensemble of spins. After the external $\mathbf{B}_0$ field is switched on in (b), the orientations of **J** are discretised and they must align either parallel or antiparallel to the direction of the field.

The frequency of precession ω_0, also known as the Larmor frequency, is

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

Second, the presence of the $\mathbf{B}_0$ field restricts the orientation of $\boldsymbol{\mu}$ relative to $\mathbf{B}_0$ to one of 2*I*+1 discrete directions, each of which has its own distinct energy level [33]. It is important to note that it is only the component of $\boldsymbol{\mu}$ parallel to $\mathbf{B}_0$, in this case μ_z, that is restricted; the transverse components μ_x and μ_y are still random [113]. Figure 2.5 shows the effect of the $\mathbf{B}_0$ field on an ensemble of $\boldsymbol{\mu}$ for a nucleus with $I=\frac{1}{2}$, resulting in two possible spin states. The allowed values of μ_z are given by

$$\mu_z = \gamma \hbar m_I \tag{2.5}$$

where m_I is the magnetic quantum number [113]. The magnetic quantum number can take one of $2I+1$ discrete integer values (corresponding to the $2I+1$ possible orientations) in the range $[-I, I]$. The angle θ between $\boldsymbol{\mu}$ and $\mathbf{B}_0$ is given by

$$\cos\theta = \frac{\mu_z}{\mu} = \frac{\gamma\hbar m_I}{\gamma\hbar\sqrt{I(I+1)}} = \frac{m_I}{\sqrt{I(I+1)}}. \tag{2.6}$$

We can also calculate the energy level *E* corresponding to each distinct orientation of $\boldsymbol{\mu}$ using the formula

$$E = -\boldsymbol{\mu}\cdot\mathbf{B}_0 = -\mu_z B_0 = -\gamma\hbar m_I B_0. \tag{2.7}$$

In the case of hydrogen, $I = \frac{1}{2}$ so there are two possible orientations and energy levels when a hydrogen nucleus is placed in an external magnetic field. These are:

- parallel to $\mathbf{B}_0$, also called spin up. In this case, $m_I = I = \frac{1}{2}$. The angle between the spin up state and $\mathbf{B}_0$ is $\theta_\uparrow = 54.74°$ and the energy level of this state is $E_\uparrow = -\frac{1}{2}\gamma\hbar B_0$.
- antiparallel to $\mathbf{B}_0$, also known as spin down. For this case, $m_I = -I = -\frac{1}{2}$. The angle between the spin up state and $\mathbf{B}_0$ is $\theta_\downarrow = 125.26°$ and the energy level of this state is $E_\downarrow = \frac{1}{2}\gamma\hbar B_0$.

Although the spin up state has a lower energy level than the spin down state, when a large ensemble of N spins are placed in the $\mathbf{B}_0$ field, the population difference between the two states is small [190]. The ratio between the number of parallel spins $N_\uparrow$ and antiparallel spins $N_\downarrow$ is described by a Boltzmann ratio [33]

$$\frac{N_\uparrow}{N_\downarrow} = \exp\left(\frac{\Delta E}{k_B T}\right), \tag{2.8}$$

where

$$\Delta E = E_\downarrow - E_\uparrow = \gamma \hbar B_0 \tag{2.9}$$

is the difference in energy between spin states, k_B is Boltzmann's constant and T is the temperature of the spin system. Therefore, for a sample of hydrogen ($\gamma = 2.675 \times 10^8 s^{-1}\mathrm{T}^{-1}$) at room temperature ($T = 300\mathrm{K}$) within a standard clinical magnetic field ($B_0 = 3\mathrm{T}$), with a total spin population $N = 2000000$, the population difference between the two states is only about 20. However, despite this very small difference, it does result in a bulk magnetisation vector $\mathbf{M}$ [190], which is calculated by summing all the nuclear magnetic moments as

$$\mathbf{M} = \sum_{n=1}^{N} \boldsymbol{\mu}_n = \sum_{n=1}^{N} \mu_{x,n}\hat{\mathbf{x}} + \sum_{n=1}^{N} \mu_{y,n}\hat{\mathbf{y}} + \sum_{n=1}^{N} \mu_{z,n}\hat{\mathbf{z}}. \tag{2.10}$$

However as we mentioned earlier, the transverse components μ_x and μ_y are random, so over the whole ensemble of spins they sum to zero [113] and thus the total magnetisation vector comprises the longitudinal μ_z components only. Therefore

$$\mathbf{M} = M_z\hat{\mathbf{z}} = \sum_{n=1}^{N} \mu_{z,n}\hat{\mathbf{z}} = \sum_{n=1}^{N_\uparrow} \frac{1}{2}\gamma\hbar\hat{\mathbf{z}} - \sum_{n=1}^{N_\downarrow} \frac{1}{2}\gamma\hbar\hat{\mathbf{z}} = \frac{1}{2}(N_\uparrow - N_\downarrow)\gamma\hbar\hat{\mathbf{z}} = M_0\hat{\mathbf{z}} \tag{2.11}$$

where M_z is the net longitudinal component of magnetisation, which has an initial magnitude of M_0. The net magnetisation vector $\mathbf{M}$ is aligned with $\mathbf{B}_0$, however unlike the individual magnetic moments from which it is generated, it is stationary and does not precess about $\mathbf{B}_0$.

2.2.2 Signal generation and detection

In the previous section we discussed some of the fundamental principles behind NMR including spins, magnetic moments and the generation of a bulk magnetisation vector. However, in order to image a sample we need to measure this magnetisation. We cannot measure $\mathbf{M}$ directly as the receiver coil within the MR system is only sensitive to oscillating magnetisations perpendicular to $\mathbf{B}_0$ [81], and $\mathbf{M}$ is a stationary magnetisation parallel to $\mathbf{B}_0$. However, by rotating $\mathbf{M}$ into the plane perpendicular to $\mathbf{B}_0$, where it precesses about the z axis, we can induce a measureable voltage in the receiver coil, which is the NMR signal [118]. In the next sections we discuss the key steps - excitation and relaxation - required to measure and detect this signal.

Excitation

The bulk magnetisation vector $\mathbf{M}$ is rotated into the plane transverse to $\mathbf{B}_0$ using a radiofrequency (RF) pulse perpendicular to the z axis [118]. The RF pulse generates an oscillating magnetic field

$$\mathbf{B}_1(t) = B_1(t)[\cos(\omega_{RF}t)\,\hat{\mathbf{x}} + \sin(\omega_{RF}t)\,\hat{\mathbf{y}}], \tag{2.12}$$

where $B_1(t)$ ($B_1(t) \ll B_0$) is an envelope function which determines the maximum magnitude of the field and ω_{RF} is the frequency of oscillation. Common choices for $B_1(t)$ include rectangular, sinc [113] and Gaussian [33] functions. For resonance to occur, ω_{RF} must be tuned to the precessional frequency of the spins being imaged [81]. The resonance condition states that

$$\hbar\omega_{RF} = \Delta E = \hbar\omega_0, \tag{2.13}$$

therefore the required frequency of the RF pulse is equal to the Larmor frequency [113].

During the RF pulse, some of the spins in the lower energy spin up state absorb energy ΔE and are excited into the higher energy spin down state. Similarly, some of the spins in the higher energy spin down state are stimulated by the RF pulse to emit energy ΔE and fall into the lower energy spin up state. However, because there are initially more spin up than spin down protons, during the pulse more spins undergo excitation than stimulated emission so there is a net increase in $N_\downarrow$ and a net decrease in $N_\uparrow$. As shown in equation 2.11, the magnitude of the longitudinal magnetisation M_z depends on the population difference between the spin up and spin down states. Therefore as the spins absorb energy and $N_\uparrow$ and $N_\downarrow$ reach equilibrium, M_z becomes zero.

Meanwhile, the RF pulse forces the magnetic moments to align with and precess about $\mathbf{B}_1$ as well as $\mathbf{B}_0$. This introduces a transverse phase coherence between the spins, which generates a net transverse magnetisation

$$M_{xy} = M_x + iM_y, \tag{2.14}$$

where M_x and M_y are the individual transverse components of $\mathbf{M}$ and $i = \sqrt{-1}$. The precessional frequency ω_1 of $\mathbf{M}$ about $\mathbf{B}_1$ is

$$\omega_1(t) = \gamma B_1(t). \tag{2.15}$$

Because $B_1(t)$ is much smaller than B_0, the spins precess about the applied $\mathbf{B}_1$ field much more slowly than about the external $\mathbf{B}_0$ field.

The overall effect of these two processes is that $\mathbf{M}$ spirals away from the z axis, towards the xy plane. This process is shown in figure 2.6. The angle α introduced between $\mathbf{M}$ and the z axis during the spiralling motion is known as the flip angle, and its size depends on both $\omega_1(t)$ and the duration of the RF pulse τ_{RF}. The flip angle is given by

$$\alpha = \int_0^{\tau_{RF}} \omega_1(t)dt. \tag{2.16}$$

For the key diffusion MR pulse sequence we consider in section 2.3, the typical value for the flip angle is $\alpha = \frac{\pi}{2}$, also known as a 90° or $\frac{\pi}{2}$ pulse. During a $\frac{\pi}{2}$ pulse $\mathbf{M}$ is tipped $\frac{\pi}{2}$ away from $\mathbf{B}_0$ into the xy plane, resulting in a purely transverse magnetisation. After the pulse, $\mathbf{M}$ is contained completely within the xy plane and precesses about the $\mathbf{B}_0$ field at frequency ω_0.

Relaxation

When the RF pulse is switched off, the spin system relaxes back towards its original state. The bulk magnetisation vector $\mathbf{M}$ spirals back towards its prior orientation parallel to $\mathbf{B}_0$, resulting in the recovery of

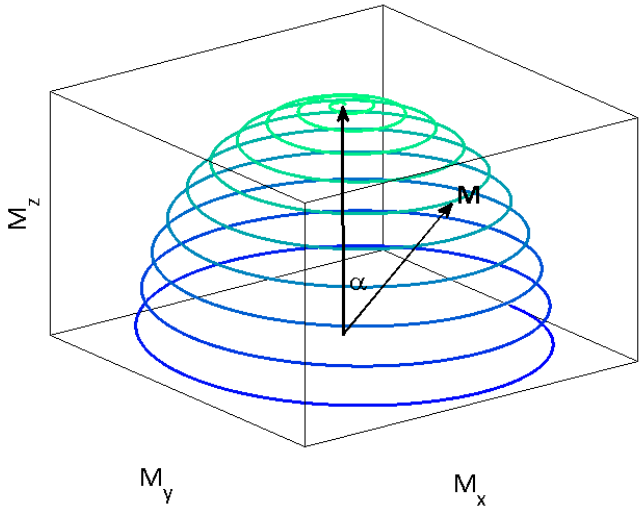

Figure 2.6: During the RF pulse, the bulk magnetisation vector $\mathbf{M}$ spirals away from the $\mathbf{B}_0$ axis and into the xy plane.

longitudinal magnetisation M_z and the destruction of transverse magnetisation M_{xy} [33]. This process, shown in figure 2.7, is known as free precession and it can be described by the Bloch equations [30]

$$\frac{d\mathbf{M}(t)}{dt} = \gamma \mathbf{M}(t) \times \mathbf{B}(t) - \frac{(M_z(t) - M_0)\,\hat{\mathbf{z}}}{T_1} - \frac{(M_x(t)\,\hat{\mathbf{x}} + M_y(t)\,\hat{\mathbf{y}})}{T_2}, \tag{2.17}$$

where T_1 and T_2 are characteristic relaxation times.

T_1 is known as the longitudinal or spin-lattice relaxation time, and it is the timescale for the recovery of longitudinal magnetisation. Spin-lattice relaxation occurs when spins that absorbed energy during the RF pulse interact with magnetic moments from neighbouring nuclei, causing them to emit energy and return to the lower energy spin up state [118]. This results in an overall loss of energy from the spin system. This process restores the population difference $N_\uparrow - N_\downarrow$, which results in the recovery of the longitudinal component of $\mathbf{M}$. By solving the Bloch equation for M_z [118], we find that its evolution as a function of time t is governed by the equation

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

therefore T_1 is the time required for M_z to recover 63% of its initial value. The recovery curve for M_z is shown in figure 2.8. The value of T_1 depends on both tissue type and B_0 [33]. T_1 is long for fluids such as CSF and is shortest for fatty tissues such as white matter [118]. Typical values of T_1 for different types of brain tissue, measured at B_0=1.5 T are shown in table 2.1.

T_2 is the transverse or spin-spin relaxation time, and it is the timescale for the destruction of the transverse component of magnetisaton M_{xy} [33]. Following the RF pulse, during which the spins gain transverse phase coherence, the spins begin to dephase again [118]. This occurs because the magnetic field that each spin experiences is altered by the magnetic moments of neighbouring spins [118]. As the spins move around relative to each other, the instantaneous magnetic fields they experience are constantly

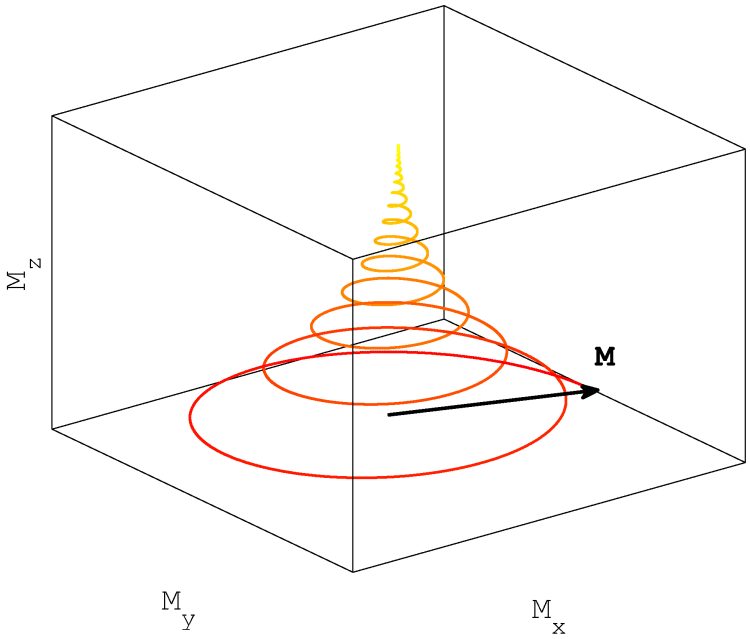

Figure 2.7: After the RF pulse is switched off, the bulk magnetisation vector $\mathbf{M}$ undergoes relaxation and returns to its original state parallel to $\mathbf{B}_0$.

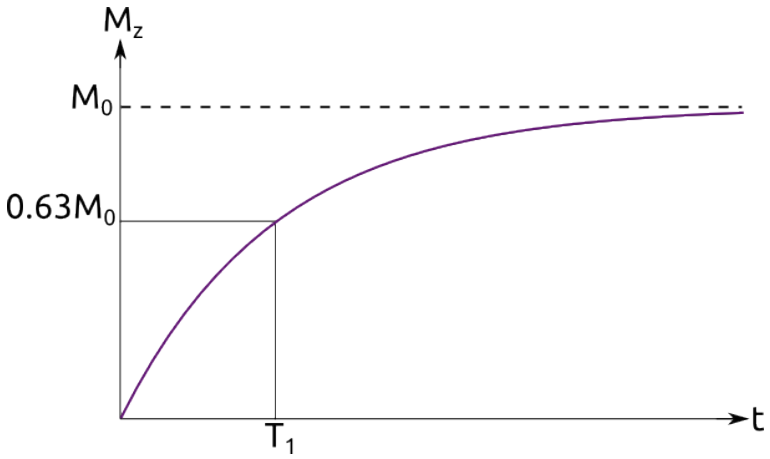

Figure 2.8: The recovery of M_z after the RF pulse is governed by the time constant T_1, also known as the spin-lattice relaxation time, which is the time taken for the longitudinal magnetisation to recover 67% of its initial value.

Tissue	T_1 (ms)	T_2 (ms)
CSF	3000	-
Grey matter	1100	90
White matter	560	80

Table 2.1: Example in vivo T_1 and T_2 relaxation times for different tissue types measured using $B_0 = 1.5\text{T}$.

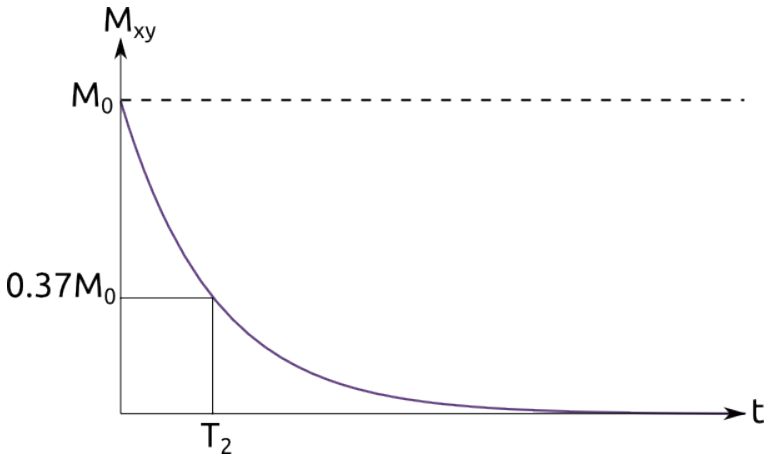

Figure 2.9: The destruction of M_{xy} is governed by the time constant T_2, also known as the spin-spin relaxation time, which is the time taken for the longitudinal component of magnetisation to drop to 37% of its value directly after the RF pulse.

changing. Because the frequency of precession is related to the magnetic field strength via equation 2.4, if the magnetic field strength increases, the spins precess faster; if the field strength decreases, they precess more slowly. Spins that precess at different frequencies start to dephase and it is this dephasing which causes the loss of transverse magnetisation. Unlike spin-lattice relaxation, there is no net loss of energy during spin-spin relaxation. By solving the Bloch equations for the transverse components of $\mathbf{M}$, we find that the evolution of M_{xy} as a function of t is

$$M_{xy}(t) = M_0 \exp\left(-i\omega_0 t\right) \exp\left(-\frac{t}{T_2}\right), \tag{2.19}$$

therefore T_2 is the time taken for the transverse magnetisation to be reduced to 37% of its initial value M_0 after the RF pulse. The corresponding curve for the destruction of M_{xy} is shown in figure 2.9. The value of T_2 also depends on tissue type and B_0. As for T_1, T_2 is long for fluids such as CSF and short for fatty tissues such as white matter. Typical values for T_2 are shown in table 2.1. From table 2.1 we note that for human tissue, T_1 is always longer than T_2 [118].

Signal detection

The signal we measure during the NMR experiment is the transverse magnetisation generated by the RF pulse [118]. Following the RF pulse, $\mathbf{M}$ is in the plane perpendicular to $\mathbf{B}_0$, precessing about the z axis at the Larmor frequency. Due to Faraday's law of induction, this oscillating magnetisation induces an oscillating voltage in the receiver coil [113]. This voltage is known as the free inducation decay (FID) signal [118]. The maximum amplitude of the FID signal depends on the initial magnetisation M_0 and the flip angle α, and its frequency of oscillation is ω_0. As the FID signal is a measure of transverse magnetisation, it decays quickly due to the dephasing of spins and subsequent loss of M_{xy}. We discussed the destruction of transverse magnetisation due to spin-spin relaxation in section 2.2.2. However there are other factors that affect how fast the spins dephase. As well as intrinsic effects such as spin-spin relaxation, there are also extrinsic effects such as inhomogeneities in the $\mathbf{B}_0$ field

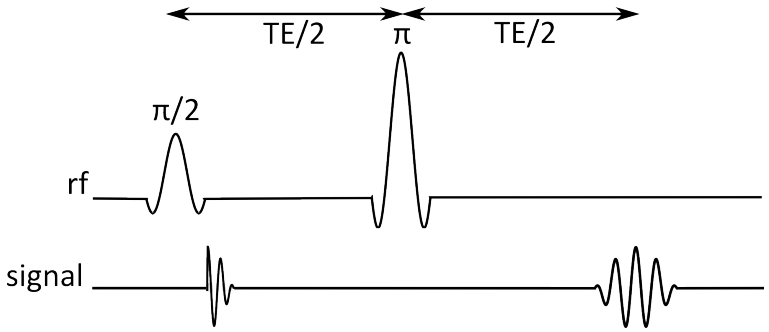

Figure 2.10: The spin echo sequence introduces a π refocusing pulse, which reverses the phases of the spins. This reintroduces the transverse phase coherence between the spins at time TE, forming the spin echo.

and magnetic susceptibility effects [81]. Typically, the dominant effect is due to the inhomogeneities in $\mathbf{B}_0$ [81]. Producing a completely uniform $\mathbf{B}_0$ field, particularly when there is a sample in the scanner, is not possible and there are always small fluctuations ΔB in the external field [190]. Because the frequency of precession is proportional to B_0 [113], these inhomogeneities cause the spins to precess at different frequencies. This increases the rate at which the spins dephase and thus the rate at which M_{xy} is destroyed. This altered rate for the destruction of transverse magnetisation can be described by a new time constant $T_2^* < T_2$ [113] where

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \gamma \Delta B. \tag{2.20}$$

The larger the field inhomogeneities are, the smaller T_2^* becomes, resulting in more rapid dephasing of the spins. Therefore the FID signal, which decays according to the T_2^* relaxation rate, decays very quickly. However, in practice we never measure the FID signal directly. Instead we generate echoes of the signal. In the next section we describe the formation of one type of echo, the spin echo, in greater depth. We focus on the spin echo sequence rather than gradient echo sequences as the spin echo sequence forms the basis of the key diffusion MR pulse sequence that we introduce in section 2.3.

Spin echoes

A schematic diagram of the spin echo (SE) sequence is shown in figure 2.10. After the inital $\frac{\pi}{2}$ pulse, the spins are allowed to dephase (due to both intrinsic and extrinsic effects) until time τ. At this time a π pulse is used to reverse the relative phases of the spins. However, the directions and frequencies of precession do not change [33]. After another time period τ, the phase differences that accumulated between the spins due to the $\mathbf{B}_0$ inhomogeneities are eliminated and the spins regain their transverse coherence, producing the spin echo. The π pulse does not eliminate the effects of spin-spin relaxation so the decay of the spin echo is governed by the T_2 relaxation rate [33]. The time 2τ at which the echo occurs is also known as the echo time TE [190].

2.2.3 Image generation

Using the sequences and methods described so far, when performing an NMR experiment on a sample, all spins within that sample would be excited during the RF pulse and we would measure one average

signal over the whole volume. In order to create an image we need to be able to spatially localise the signals. In the next sections we briefly describe how this is done and how the images are reconstructed from the signals we record.

Spatial localisation

During an MR acquisition the signals are spatially localised using magnetic field gradients [141]. These gradients $\mathbf{G}(t) = (G_x(t), G_y(t), G_z(t))$ are generated by coils inside the scanner and can be applied along the x, y and z directions [118]. Their purpose is to alter the magnetic field strength within the scanner as a function of position [141] so that at location $\mathbf{r} = (x, y, z)$ the magnetic field strength B is

$$B = B_0 + \mathbf{G} \cdot \mathbf{r}. \tag{2.21}$$

When the gradient is switched on, the precessional frequencies of the spins are altered and they accumulate phase differences dependent on their positions within the scanner. This allows us to encode phase and frequency information directly into the signals we measure [113]. Based on our knowledge of $\mathbf{G}$ we can use this information to determine the spatial location of the signals [118]. Three different gradients are used to encode the spatial information. These are the slice select gradient, the phase encoding gradient and the frequency encoding gradient. We briefly describe these three gradients in turn.

The purpose of the slice select gradient is to limit the spins excited during the RF pulse to a two dimensional slice within the sample [141]. The slice select gradient pulse, which is applied simultaneously with the RF pulse, alters the resonant frequency of the spins as a function of position along the gradient direction, also called the slice direction. Therefore only those spins whose resonant frequencies fall within the bandwidth of the RF pulse are excited and produce a signal. The position of the slice within the sample can be altered by changing the frequency of the RF pulse and the slice thickness can be controlled by altering either the bandwidth of the RF pulse or the strength of the gradient [118]. The slice direction is typically aligned with the z axis resulting in axial image slices; however it can also be aligned with the x or y axes or an arbitrary direction to produce sagittal, coronal or oblique slices [118].

Once a slice has been excited along the z axis, we then need to spatially localise the signals within the plane using phase encoding directions along the y direction and frequency encoding gradients along the x direction. The phase encoding gradient $G_y(t)$ is turned on after the RF pulse for a short time τ_{PE}. During this time, the precessional frequencies of the spins are dependent on position along the phase encoding direction y. After the gradient pulse, the frequency of the spins reverts to the Larmor frequency, however each spin has a phase angle

$$\Phi(y) = y \int_0^{\tau_{PE}} \gamma G_y(t) dt \tag{2.22}$$

associated with it, which is dependent on its y position. The frequency encoding direction $G_x(t)$ is turned on during the MR signal readout at time TE for time τ_{FE}. The precessional frequencies of the spins during the pulse is therefore dependent on position along x so that the resonant frequency ω during the readout is

$$\omega(x) = \omega_0 + x \int_{TE}^{TE+\tau_{FE}} \gamma G_x(t) dt. \tag{2.23}$$

Because the MR signal oscillates at the same frequency as the resonant frequency of the spins, the frequency dependence on x is encoded directly into the frequency of the signals we record.

The measured signal $S(t)$ after spatial encoding [113] is

$$S(t) = \iint I(x,y) \exp\left(-2\pi i y \int_0^{\tau_{PE}} \gamma G_y(t)\, dt\right) \exp\left(-2\pi i x \int_{TE}^{TE+\tau_{FE}} \gamma G_x(t)\, dt\right) dxdy. \tag{2.24}$$

where $I(x,y)$ is the image function we want to measure. We can define spatial frequencies k_x and k_y [113] such that

$$k_x = \int_{TE}^{TE+\tau_{FE}} \gamma G_x(t) dt, \tag{2.25}$$

$$k_y = \int_0^{\tau_{PE}} \gamma G_y(t) dt. \tag{2.26}$$

Substituting these spatial frequencies into equation 2.24 gives

$$S(t) = \iint I(x,y) \exp(-2\pi i k_y y) \exp(-2\pi i k_x x)\, dxdy = F(k_x, k_y). \tag{2.27}$$

This shows that there is a Fourier relationship between positions (x,y) and spatial frequencies (k_x, k_y). Therefore we can gain information about the sample by measuring the signal at many values of k_x and k_y, also known as k-space. This is convenient because, as equations 2.25 and 2.26 show, we can easily traverse k-space by manipulating the gradients $G_x(t)$ and $G_y(t)$.

There are numerous different strategies for sampling k-space, including spiral, radial, linear and rectilinear sampling [163]. Because diffusion-weighted MRI is particularly susceptible to phase effects [163], it is preferable to use a sampling scheme which covers the whole of k-space during a single excitation. Therefore, the most common method for collecting diffusion MRI data is Single Shot Echo Planar Imaging (SS-EPI) [141]. This technique uses phase and frequency encoding gradients to traverse k-space in a rectilinear zig zag pattern. We show the SS-EPI pulse sequence and resulting k-space coverage in figure 2.11. After the first slice has been acquired, data from subsequent slices can be obtained by repeating the pulse sequence using RF pulses with different frequencies until the whole sample has been imaged. The time gap between successive RF excitations is called the repetition time TR [33].

Image reconstruction

The final step in MR imaging is to reconstruct the image $I(x,y)$ from the raw k-space data. From equation 2.27, we see that the k-space data $F(k_x, k_y)$ is related to $I(x,y)$ via a two-dimensional inverse Fourier transform [113]. By sampling enough points in k-space, we can invert this Fourier relationship for each slice of data to reconstruct the image [113]. The results of the discrete Fourier transform have both a real and imaginary component, which can be combined to form both magnitude and phase images [118].

2.2.4 Image contrast

The three parameters that we control during a spin echo pulse sequence are TE, TR and α. TE controls the amount of T_2 relaxation that occurs before the echo is generated [33] whereas TR controls the

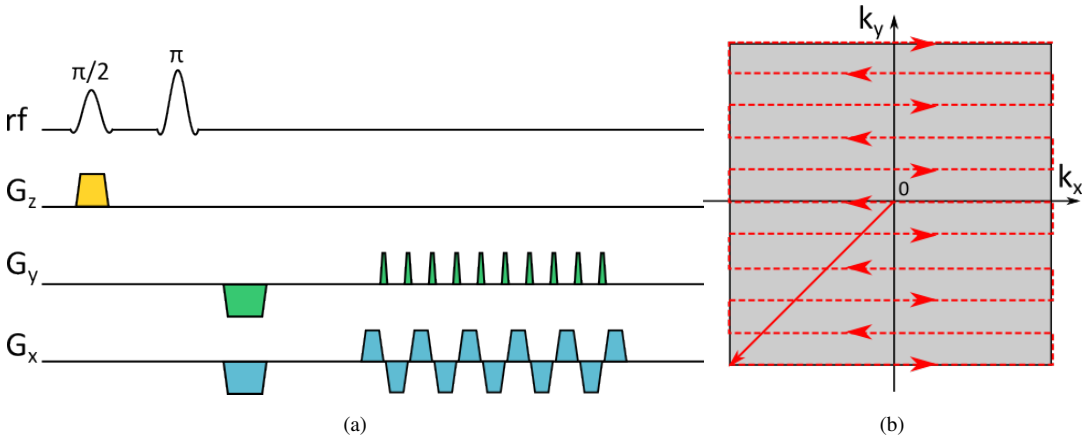

Figure 2.11: (a) shows the SS-EPI pulse sequence. The slice select gradient pulse G_z is switched on concurrently with the $\frac{\pi}{2}$ excitation pulse to select a slice in the z direction. The initial negative gradient lobes along G_x and G_y move us from the centre to the bottom left corner of k-space. The subsequent G_x and G_y pulses allow us to sample k-space in a rectilinear zig zag pattern as shown in (b). Data is acquired during the frequency encoding G_x gradients, but not during the phase encoding G_y gradients.

amount of T_1 relaxation that occurs before the next RF pulse [33]. The flip angle α determines the amount of transverse magnetisation generated by the RF pulse. However as it is set to $\alpha = \frac{\pi}{2}$ in the SE sequence, we do not discuss its role generating image contrast here. By using different combinations of TE and TR we can obtain three different image contrasts using the basic SE sequence. These are T_1-, T_2- and pd-weighting, which we now describe in turn.

In T_1-weighted images, the contrast is provided by maximising T_1 differences between tissues whilst minimising the effects of T_2 relaxation [190]. This is done using both short TR and short TE. As shown in table 2.1, T_1 is long for fluids such as CSF, shorter for watery tissue such as grey matter and shortest for fatty tissue such as white matter. Therefore in a T_1-weighted image, regions of CSF appear dark whilst white matter appears brightest.

To generate T_2-weighted images we must maximise the difference in T_2 relaxation between tissue types using long TE and minimise T_1 differences, which can be achieved using long TR [190]. As mentioned in section 2.2.2, CSF has a much longer T_2 than grey and white matter [118]. However, because T_2 is the timescale for the destruction of M_{xy}, a longer T_2 results in a larger signal. Therefore, unlike T_1-weighted images, CSF appears bright in T_2-weighted images whereas white matter appears dark.

Finally, to create pd-weighted images, which use the relative proton density of tissue as a contrast mechanism, we must minimise both T_1 and T_2 effects using a long TR and a short TE [190]. We show examples of T_1, T_2 and pd-weighted images in figure 2.12.

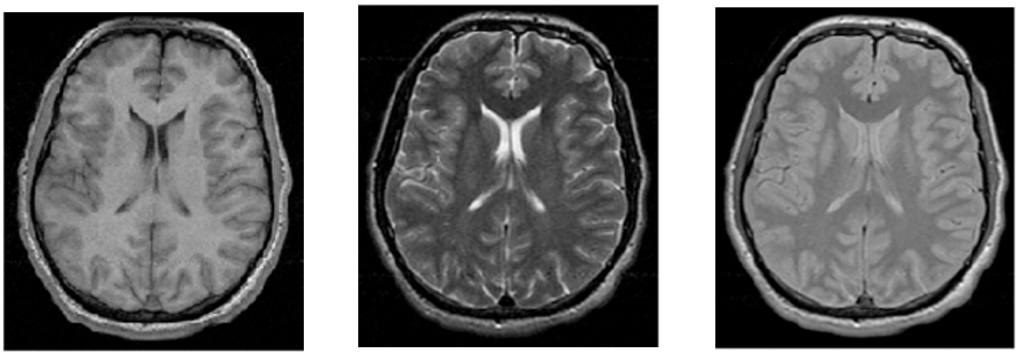

(a) (b) (c)

Figure 2.12: (a) shows a T_1-weighted image (TR/TE=14/5 ms). The CSF in the ventricles appears dark whilst white matter appears brightest. The opposite trend is observed in the T_2-weighted image (TR/TE= 4000/100 ms) shown in (b), in which the CSF appears brightest and regions of white matter are almost black. (c) shows a pd-weighted image (TR/TE=4000/15 ms). Due to the similar proton densities of different types of brain tissue, proton density images do not provide as much constrast between grey matter, white matter and CSF. Images courtesy of Ivana Drobnjak.

2.3 Diffusion MRI

We cannot achieve diffusion contrast by simply altering TE and TR as we did for the image contrasts described here; we need to modify the SE sequence further to allow for measurements of diffusion. In this section, we discuss the physical basis of diffusion-weighted imaging and pulse sequences for diffusion sensitisation in more depth.

Diffusion is the process by which particles move through a system due to random molecular motions [52]. The diffusion process was first described mathematically by Fick [37], who related the flux $\mathbf{F}$ of particles through the medium to the concentration of particles $n(\mathbf{r}, t)$ by

$$\mathbf{F} = -D\nabla n(\mathbf{r}, t), \tag{2.28}$$

where D is the diffusion coefficient. By considering the conservation of particles [52], we can relate the rate of change of concentration to $\nabla\mathbf{F}$ so that

$$\frac{\partial n(\mathbf{r}, t)}{\partial t} = D\nabla^2 n(\mathbf{r}, t). \tag{2.29}$$

Given that the motion of the particles is random, Einstein considered the particle concentration n as a probability distribution function (pdf) $\Pr(\mathbf{r}, t|\mathbf{r}_0, t_0)$, which describes the probability of a particle moving from position $\mathbf{r}_0$ at time t_0 to position $\mathbf{r}$ at time t. Using this probabilisitic description, and assuming that at $t_0 = 0$ the particle is at $\mathbf{r}_0 = 0$, Einstein solved equation 2.29 [37] to give

$$\Pr(\mathbf{r}, t) = (4\pi Dt)^{-\frac{3}{2}} \exp\left(-\frac{\mathbf{r}^2}{4Dt}\right). \tag{2.30}$$

This shows that when particles diffuse freely, their displacements can be described by a Gaussian function. Einstein also calculated the mean squared displacement $\langle\mathbf{r}^2\rangle$ of the particles as

$$\langle\mathbf{r}^2\rangle = 6Dt. \tag{2.31}$$

So how can we use diffusion to study white matter? White matter tissue accounts for approximately 50% of all brain tissue [63], and within white matter the fractional water content has been estimated at around 70% [9]. These water molecules naturally self-diffuse due to random thermal motions. However, unlike the free diffusion process described above, the architecture of the brain tissue hinders and restricts the dispersion of water molecules, so that instead of depending solely on D and t, the mean squared displacement and particle pdf depends on the directionality and lengthscales of the tissue. For example, grey matter tissue in the mature cortex comprises unordered tangles of neuronal cell bodies with limited spatial coherence or directionality on the length scale of a voxel. This hinders the diffusion of water molecules equally in all directions resulting in an approximately isotropic pdf, but with a small $\langle \mathbf{r}^2 \rangle$. On the other hand, white matter tissue is much more ordered. As described in section 2.1.1, it comprises parallel bundles of long, cylindrical axons. The diffusion of water molecules trapped inside the axons is restricted on the lengthscale of their diameters, whilst the dispersion of water molecules around the bundles is hindered. In both cases, due to the ordered, macroscopic structure of the tracts, it is easier for molecules to diffuse parallel to the axons than perpendicular to them, resulting in an anisotropic particle displacement pdf. Therefore, if we can estimate the pdf, we can infer properties about the axons such as directionality and size. To do this, we need to sensitise the MR signal we measure to the diffusion process and derive expressions which allows us to relate these signals and pdfs to the tissue properties we want to estimate.

2.3.1 Diffusion sensitisation

Diffusion MRI aims to provide image contrast based on the average amount of water diffusion along a given direction, as described above. In order to do this, the MR signal must be sensitised to the distance and direction that water molecules move between the initial $\frac{\pi}{2}$ excitation pulse and echo formation at time TE. Unlike the T_1-, T_2- and pd-weighted images discussed in section 2.2.4, this cannot be achieved by simply tuning TE and TR. It requires the addition of two equal area magnetic field gradient pulses (they are typically rectangular in shape) to the standard spin echo sequence, one before and one after the π refocusing pulse. This modified pulse sequence, first introduced by Stejskal and Tanner in 1965 [169], is known as the Pulsed Gradient Spin Echo (PGSE) sequence and it is illustrated in figure 2.13. The sequence is characterised by four key parameters: the duration of the pulses δ, the time separation between the leading edges of the two pulses Δ, the magnitude of the gradient G and its direction $\hat{\mathbf{G}}$. The PGSE sequence is the standard pulse sequence used in diffusion-weighted MRI experiments, although other sequences in common use include twice refocused spin echo (TRSE) [151], stimulated echo (STEAM) [119] and steady state free precession (SSFP) [117]. Current research is also focused on generating more application specific sequences by replacing the rectangular gradient pulses with double wave vector [105], oscillating [193] or arbitrary shaped pulses [58]. However, in this work we focus on the PGSE sequence, and do not consider these alternative pulse sequences.

We now describe how the addition of these two gradient pulses results in the diffusion-weighting of the MR signal. This is also illustrated in figure 2.14. As with the standard SE sequence, the PGSE sequence begins with a $\frac{\pi}{2}$ pulse, after which the spins are precessing in phase in the plane perpendicular

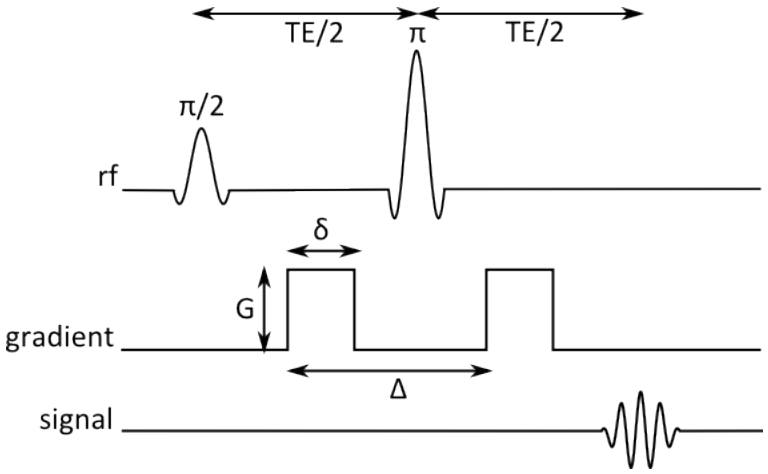

Figure 2.13: The PGSE sequence modifies the SE sequence by adding two magnetic field gradient pulses - one before and one after the π pulse. This pulses add position dependent phases to the spins. If the spins do not move between the two pulses, these phases cancel out; if they do move, the phases do not cancel and the spin has a residual phase which attenuates the spin echo signal.

to $\mathbf{B}_0$. During the first gradient pulse, the local magnetic field experienced by the spins is perturbed, depending on the position **r** of the spin along the gradient direction, so that the magnitude of the magnetic field B is given by

$$B = B_0 + \mathbf{G}(t) \cdot \mathbf{r}(t), \tag{2.32}$$

where $\mathbf{G}(t) = G(t)\hat{\mathbf{G}}$. Due to the change in the local magnetic field strength, the rate of precession ω of the spins is also altered dependent on position along $\hat{\mathbf{G}}$ so that

$$\omega = \gamma \left(B_0 + \mathbf{G}\left(t\right) \cdot \mathbf{r}\left(t\right) \right). \tag{2.33}$$

Because of the change in ω during the pulse, after time δ a phase offset Φ_1 accumulates, which can be calculated as

$$\Phi_1 = -\gamma \int_{t'}^{t'+\delta} \mathbf{G}\left(t\right) \cdot \mathbf{r}\left(t\right) dt. \tag{2.34}$$

For a rectangular pulse, the magnitude of the magnetic field gradient does not vary with time so Φ_1 can be simplified to

$$\Phi_1 = -\gamma \mathbf{G} \cdot \int_{t'}^{t'+\delta} \mathbf{r}\left(t\right) dt. \tag{2.35}$$

The first gradient pulse is followed by a π rephasing pulse and a second, identical gradient pulse. During the second pulse, a second phase offset Φ_2 accumulates which is

$$\Phi_2 = \gamma \mathbf{G} \cdot \int_{t'+\Delta}^{t'+\Delta+\delta} \mathbf{r}\left(t\right) dt, \tag{2.36}$$

which has the opposite sign to the first phase due to the π pulse. If the spins have not moved in the time between the pulses, they experience the same perturbed field during both gradient pulses [145] and thus $\Phi_2 = -\Phi_1$. Therefore at TE, the overall phase change of the spins $\Delta\Phi = \Phi_1 + \Phi_2$ is zero and the spins rephase completely. However, if the spins have diffused randomly, during the second gradient

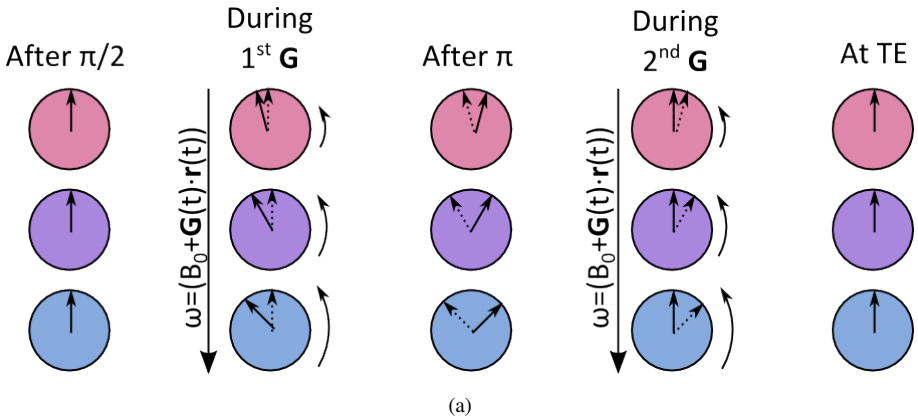

Figure 2.14: In (a) the spins are stationary throughout the PGSE sequence. The phase accumulated during the second pulse is therefore equal and opposite to the phase acquired during the first pulse resulting in full transverse phase coherence at TE. If the spins move along the gradient direction during the pulse, as in (b), the phase acquired during the second pulse no longer negates the effects of the first pulse resulting in a residual phase at TE which attenuates the measured signal.

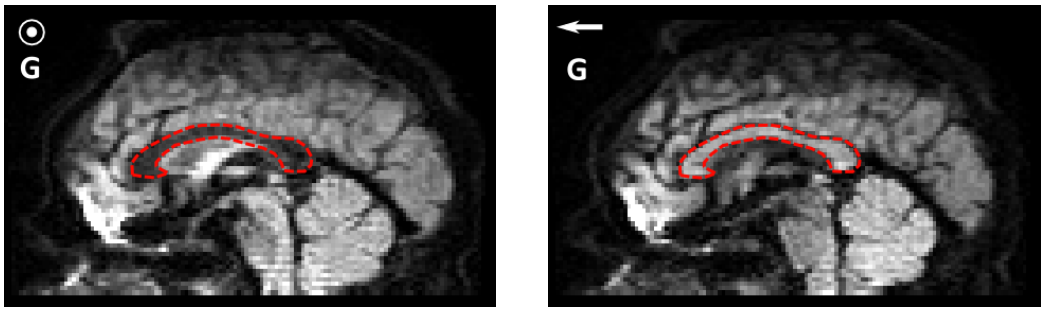

(a) (b)

Figure 2.15: In (a) the diffusion-weighted gradient pulses are along the x direction which aligns with the left-right fibre orientation of the CC (outlined in red). In this case the diffusion-weighted signal is very low, indicating a large amount of signal attenuation, and therefore a high level of diffusion along this direction. In (b), the diffusion-weighted pulses are along the y direction, perpendicular to the CC fibres. Now the diffusion-weighted signal is high, indicating little signal attenuation and thus only a small amount of diffusion along this axis.

pulse the local magnetic field strength is different and therefore $\Phi_2 \neq -\Phi_1$. Instead of rephasing at TE the spins have random phase shifts due to incoherent motion which attenuates the signal [145]. From this signal attenuation we can infer how far the spins have diffused along the gradient direction; the more attenuation there is, the further they have diffused on average.

As discussed in the previous section, the amount of water diffusion within white matter is directionally dependent. This is illustrated in figure 2.15 which shows how the signal attenuation differs depending on the direction we choose for $\hat{\mathbf{G}}$. Therefore, in order to build up a three-dimensional picture of water diffusion, we measure the signal attenuation along multiple gradient directions [21]. These signals allow us to estimate a three-dimensional pdf of spin displacements. Approximations such as the short gradient pulse approximation and the Gaussian phase distribution approximation allow us to do this in a straightforward analytical way. In the next two sections we introduce these approximations and describe them in more depth.

2.3.2 Short gradient pulse approximation

The short gradient pulse (SGP) approximation asserts that the duration of the gradient pulse is infinitely short, i.e. there is no motion of the spins during the pulse itself [145]. Experimentally, this approximation is valid when the duration of the pulses is much shorter than the time between them, i.e. $\delta \ll \Delta$, and for porous systems when $\delta \ll \frac{a^2}{d}$, where a is the size of the pore and d is the diffusivity. Equation 2.35 shows the phase accumulated by a spin during a gradient pulse of finite width. However, under the SGP approximation, the position of the spin $\mathbf{r}$ is constant during the pulse, so we can integrate over the pulse width to find

$$\Phi = \gamma\delta\mathbf{G} \cdot \mathbf{r}. \tag{2.37}$$

Therefore, if a spin is at position $\mathbf{r}_1$ during the first gradient pulse, and diffuses to $\mathbf{r}_2$ in the time Δ before the second pulse, the total phase change for the spin is

$$\Delta\Phi = \gamma\delta\mathbf{G}\cdot(\mathbf{r}_2 - \mathbf{r}_1)\,. \tag{2.38}$$

The signal from a spin diffusing from $\mathbf{r}_1$ to $\mathbf{r}_2$ depends on the initial density of spins at $\mathbf{r}_1$ $\rho(\mathbf{r}_1)$, the diffusion propagator $\Pr(\mathbf{r}_1, \mathbf{r}_2, \Delta)$, which defines the probability of a spin moving from $\mathbf{r}_1$ to $\mathbf{r}_2$ in time Δ, and $\Delta\Phi$ [21]. To calculate the signal at the echo time $S(\mathbf{G}, TE)$ over the whole population of spins, we integrate over all possible values of $\mathbf{r}_1$ and $\mathbf{r}_2$ to give

$$S(\mathbf{G}, TE) = S_0 \iint \rho(\mathbf{r}_1)\Pr(\mathbf{r}_1, \mathbf{r}_2, \Delta)\exp\left(i\gamma\delta\mathbf{G}\cdot(\mathbf{r}_2 - \mathbf{r}_1)\right)d\mathbf{r}_1 d\mathbf{r}_2, \tag{2.39}$$

where S_0 is the signal with no diffusion-weighting. By substituting in the appropriate spin density and propagator functions, the integral can be solved to find an expression for $S(\mathbf{G}, TE)$. For example, in the case of free diffusion, $\rho(\mathbf{r}_1)$ is assumed to be uniform and $\Pr(\mathbf{r}_1, \mathbf{r}_2, \Delta)$ is a trivariate Gaussian function [145]. In this case, equation 2.39 can be solved to give

$$S(\mathbf{G}, TE) = S_0 \exp\left(-\gamma^2\delta^2 G^2 D\Delta\right). \tag{2.40}$$

The particle displacements $\mathbf{R} = \mathbf{r}_2 - \mathbf{r}_1$ can also be used to write the signal in terms of an average diffusion propagator $\hat{\Pr}(\mathbf{R}, \Delta)$ [21] where

$$\hat{\Pr}(\mathbf{R}, \Delta) = \int \rho(\mathbf{r}_1)\Pr(\mathbf{r}_1, \mathbf{r}_1 + \mathbf{R}, \Delta)\,d\mathbf{r}_1. \tag{2.41}$$

Using this average propagator, and defining the wavevector $\mathbf{q}$ as

$$\mathbf{q} = \frac{\gamma\delta\mathbf{G}}{2\pi} \tag{2.42}$$

we can also express the signal as

$$S(\mathbf{q}, TE) = S_0 \int \hat{\Pr}(\mathbf{R}, \Delta)\exp\left(i2\pi\mathbf{q}\cdot\mathbf{R}\right)d\mathbf{R}, \tag{2.43}$$

which shows that there is a Fourier relationship between $\hat{\Pr}(\mathbf{R}, \Delta)$ and $S(\mathbf{q}, TE)$.

Using the SGP approximation, solutions for simple geometries such as infinite planes [18, 36, 173], infinite cylinders [36] and spheres [18, 36, 121] have been calculated.

2.3.3 Gaussian phase distribution approximation

In many cases, δ is the same order of magnitude as Δ, violating the condition required for the SGP approximation. However, there is another useful approximation we can make, the Gaussian phase distribution (GPD) approximation, which is valid for finite δ.

Instead of writing the total signal in terms of spin displacements, $S(\mathbf{G}, TE)$ can also be expressed in terms of the phases Φ accumulated by the spins [38, 57] at TE as

$$S(\mathbf{G}, TE) = S_0 \int_{-\infty}^{\infty} \Pr(\Phi, TE)\cos\Phi d\Phi, \tag{2.44}$$

where $\Pr(\Phi, TE)$ is the probability distribution of phases at the echo time. The GPD approximation, first introduced by Douglass and McCall [57], asserts that if the number of spins within the ensemble is

large compared to the distance travelled by a spin within time TE, the probability distribution of phases is Gaussian. By substituting a Gaussian distribution for $\Pr(\Phi, TE)$ into equation 2.44, Douglass and McCall show that the signal is given by

$$S(\mathbf{G}, TE) = S_0 \exp\left(-\frac{\langle\Phi^2\rangle}{2}\right), \tag{2.45}$$

where $\langle\Phi^2\rangle$ is the mean squared phase change. Therefore, in order to determine $S(\mathbf{G}, TE)$, we need to evaluate $\langle\Phi^2\rangle$. In the case of free diffusion, $\langle\Phi^2\rangle$ can be calculated [57] to give

$$S = S_0 \exp\left(-\gamma^2\delta^2 G^2 D\left(\Delta - \frac{\delta}{3}\right)\right). \tag{2.46}$$

The GPD approximation is only exact in the case of free diffusion [145]. However, it can also be used to derive good approximations to the signal for restricted diffusion in simple geometries. Neuman [128] argues that the GPD approximation is valid for restricted diffusion in both the short and long diffusion time limits. In the short time limit, very few spins have reached and interacted with the boundaries so diffusion is essentially unrestricted; in the long time limit, the positions and phases of the spins become independent of their starting positions, resulting in a Gaussian phase distribution. In the intermediate diffusion time regime, Balinov et al [18] show that analytical signal expressions calculated using the GPD approximation match computer simulations reasonably well, implying that for this case, the distribution of phases is almost Gaussian. As for the SGP approximation, the GPD approximation has been used to calculate expressions for the restricted diffusion signal in simple geometries such as infinite planes [18], infinite cylinders [170, 182] and spheres [125].

2.4 Modelling the diffusion MRI signal in white matter

Analytic models of the diffusion MR signal are fit to the measurements in order to estimate specific tissue parameters from the signals. In this section, we introduce a selection of analytic models of the diffusion MR signal, ranging from simple to complex, and discuss their suitability for studying white matter.

2.4.1 The diffusion tensor

The most widely used diffusion MRI method for estimating microstructural features is diffusion tensor (DT) MRI [20]. DT MRI models the probability density function of the particle displacements $\mathbf{R}$ using a zero mean trivariate Gaussian distribution as

$$\Pr(\mathbf{R}) = \frac{1}{\sqrt{(4\pi t)^3|\mathbf{D}|}} \exp\left(-\frac{\mathbf{R}^T\mathbf{D}^{-1}\mathbf{R}}{4t}\right). \tag{2.47}$$

where $\mathbf{D}$ is the diffusion tensor and $t = \left(\Delta - \frac{\delta}{3}\right)$ is the diffusion time. This extends the Einstein formulation from equation 2.30 by acknowledging that diffusion varies along different directions, i.e. that it can be anisotropic. The diffusion tensor is a 3×3 positive definite symmetric matrix

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{pmatrix}, \tag{2.48}$$

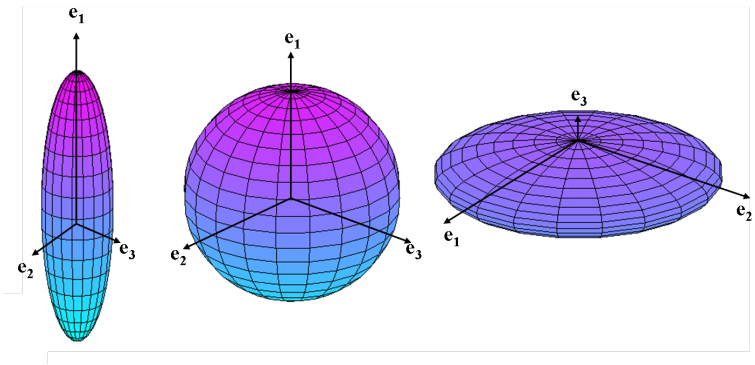

Figure 2.16: From left to right: the prolate tensor represents anisotropic diffusion; the sphere represents isotropic diffusion; the oblate tensor represents diffusion in a plane, which could be due to crossing, bending or fanning fibres.

where D_{xx}, D_{yy} and D_{zz} are the diffusivities in the x, y and z directions in the laboratory reference frame and D_{xy}, D_{xz} and D_{yz} represent the covariances. The measured signal $S(\mathbf{G}, TE)$ is related to the diffusion tensor by

$$S(\mathbf{G}, TE) = S_0 \exp\left(-\gamma^2\delta^2 G^2 t \hat{\mathbf{G}}^T \mathbf{D} \hat{\mathbf{G}}\right) = S_0 \exp -b\hat{\mathbf{G}}^T \mathbf{D}\hat{\mathbf{G}}, \tag{2.49}$$

where $b = \gamma^2\delta^2 G^2 t$ is the strength of the diffusion weighting. In order to estimate the 6 unique elements of the diffusion tensor and the normalising constant S_0 we need at least seven measurements: one $b = 0$ measurement and six non-zero b measurements with non-collinear gradient directions, although simulation studies have shown that in order to reduce the effect of image noise on our estimates at least 20-30 directions are needed [94]. There are many different techniques available for diffusion tensor estimation [104], although it is usually formulated as a least squares problem. This can be solved using linear regression [20], although non-linear optimisation methods are generally less biased [95].

Once the diffusion tensor has been estimated, we can then decompose $\mathbf{D}$ to find the eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ and corresponding eigenvalues $(\lambda_1 \geq \lambda_2 \geq \lambda_3)$. The eigenvectors give the principal directions of diffusion and the eigenvalues are the corresponding diffusivities. The major eigenvector $\mathbf{e}_1$ has the largest diffusivity and provides an estimate of the dominant fibre direction in each voxel. We can visualise the shape of the probability density function using a contour of the Gaussian function, which is an ellipsoid with dimensions determined by the size of the eigenvalues:

- if $\lambda_1 \gg \lambda_2 \approx \lambda_3$, diffusion is anisotropic and we get a prolate tensor
- if $\lambda_1 \approx \lambda_2 \approx \lambda_3$, diffusion is isotropic and we get a sphere
- if $\lambda_1 \approx \lambda_2 \gg \lambda_3$, diffusion takes place in the plane containing $\mathbf{e}_1$ and $\mathbf{e}_2$ and we get an oblate tensor.

We show examples of these tensors in figure 2.16.

We can also derive several rotationally invariant quantities from the eigenvalues to use as indicators of microstructure integrity such as fractional anisotropy (FA) and mean diffusivity (MD) [23]. Fractional

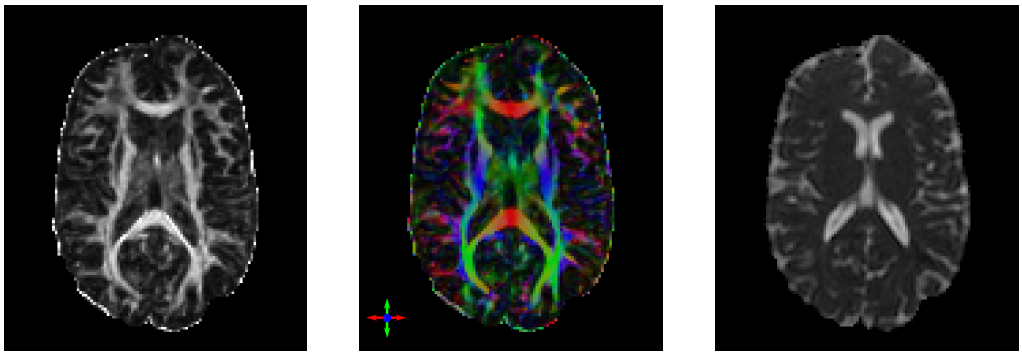

(a) (b) (c)

Figure 2.17: (a) shows an example FA map. The anisotropy is high in the white matter, which appears brightest, and almost zero in the CSF, which appears black. FA maps can also be colour coded according to the direction of the principal eigenvector as shown in (b). In the centre of the brain we can see in red the left-right connections of the CC. (c) shows a map of MD, which displays the opposite trends to FA. Diffusivity is low in the white matter, which appears almost black, whereas it is high in the CSF, which now appears brightest.

anisotropy is defined as

$$\mathrm{FA} = \left(\frac{3}{2} \sum_{i=1}^{3} \frac{\left(\lambda_i - \frac{1}{3}\mathrm{Tr}(\mathbf{D})\right)^2}{\sum_{i=1}^{3} \lambda_i^2} \right)^{\frac{1}{2}} \tag{2.50}$$

and mean diffusivity is defined as

$$\mathrm{MD} = \frac{1}{3} Tr(\mathbf{D}) \tag{2.51}$$

where

$$Tr(\mathbf{D}) = \lambda_1 + \lambda_2 + \lambda_3. \tag{2.52}$$

Fractional anisotropy, which takes values on the interval [0,1], measures the degree of diffusion anisotropy in a voxel whereas MD measures the mean amount of diffusion. Fractional anisotropy is low in the CSF ($\mathrm{FA} \approx 0$) as the water molecules undergo free diffusion and are equally likely to diffuse in all directions. FA is also low in grey matter ($0 < \mathrm{FA} \leq 0.2$), as even though on a microscopic scale grey matter comprises complex microstructure, on the voxel scale at which we make our measurements ($\approx 2\mathrm{mm}^3$) it often appears unordered with little directionality. However in white matter FA is often higher ($0.2 < \mathrm{FA} < 1$). This is because, as described in section 2.1.1, the axons in white matter are arranged in spatially coherent bundles which have definite directionality on the macroscopic as well as the microscopic scale. We see the opposite trend for MD, where diffusivity is high in the CSF. MD is lower in both grey matter and white matter where the motion of water molecules is hindered by the tissue structure, although it is often lower in white matter than grey matter. We show example images of FA and MD in figure 2.17. Other measures that can be derived from the eigenvalues of the DT include axial diffusivity $\lambda_{\parallel} = \lambda_1$ and radial diffusivity $\lambda_{\perp} = \frac{\lambda_2 + \lambda_3}{2}$ [23], as well as linearity $c_L = \frac{\lambda_1 - \lambda_2}{\lambda_1}$ and planarity $c_P = \frac{\lambda_2 - \lambda_3}{\lambda_1}$ [191]. Out of all the indices mentioned so far, FA and MD are the most commonly used measures of tissue microstructure and have been used in numerous studies to quantify changes in

tissue integrity [13, 65, 96, 144]. However there are several drawbacks to the DT. First, the Gaussian diffusion assumption does not accurately describe the motion of water molecules in white matter tissue where water diffusion within the axons is restricted [24]. Second, if a voxel contains two dominant orientations due to fibres crossing, bending or fanning, the DT is oblate and unable to resolve the fibre directions [8]. Third, FA and MD are not specific to underlying changes in tissue structure. Both FA and MD can be affected by changes in tissue properties such as axon density, axon size, myelination, diffusivity, permeability and inflammation [3]. However, it is not necessarily possible to infer the cause of the change and some combinations of structural changes may even leave FA and MD unaltered.

2.4.2 Multi-compartment models of the diffusion MRI signal

In the previous section, we discussed the DT model and highlighted its limitations. To overcome these, more descriptive models of tissue structure have been proposed that allow us to relate parameters such as axon radius and glial cell size directly to the signals we measure. This is done by creating parsimonious representations of white matter tissue structure using a small number of key tissue features. Each feature is represented using a simple geometric compartment whose signal can be calculated using approximations such as those described in sections 2.3.2 and 2.3.3. The total signal is the sum over all compartments. In this section we review a selection of these models and discuss how they have been used to provide estimates of fundamental tissue parameters in both ex vivo and in vivo studies. We focus on models of the diffusion MR signal in white matter tissue, omitting models that describe diffusion within grey matter or other systems.

Stanisz's three pool model

One of the first analytical models of the diffusion MR signal in white matter to employ a realistic tissue model is the three pool model proposed by Stanisz [168] et al in 1997. They model the total signal as a weighted sum of signals from three distinct components: restricted water diffusion within spherical glial cells of size a_S, restricted water diffusion within prolate ellipsoidal axons with dimensions $a_{T\parallel}, a_{T\perp}$, and hindered water diffusion in the extra-cellular space. A schematic diagram of this model is shown in figure 2.18. The diffusivity of water inside the glia and axons is assumed to be the same, although the extra-cellular diffusivity is different. Stanisz et al also consider the effects of permeability by assigning different rates of exchange between the extra-cellular and glial compartments and the extra-cellular and axonal compartments. The model is fit to data acquired from an excised bovine optic nerve and estimates are obtained for the intra- and extra-cellular diffusivities, permeabilities, volume fractions and cell dimensions. Due to the complexity of the model, a very rich and high quality dataset is required. Stanisz et al use NMR spectroscopy to acquire large numbers of measurements covering a wide range of diffusion times and gradient strengths, with the nerve sample carefully positioned both parallel and perpendicular to the direction of the gradient pulses. The parameter estimates obtained for compartment size match the range of values observed during histology whereas estimates for permeability are comparable to those found in the literature. They also test whether the full complexity of the model is required by comparing the full three compartment model to a simplified two compartment model (extra-cellular and axonal compartments only) and a three-compartment model with no permeability. Results show that to fully

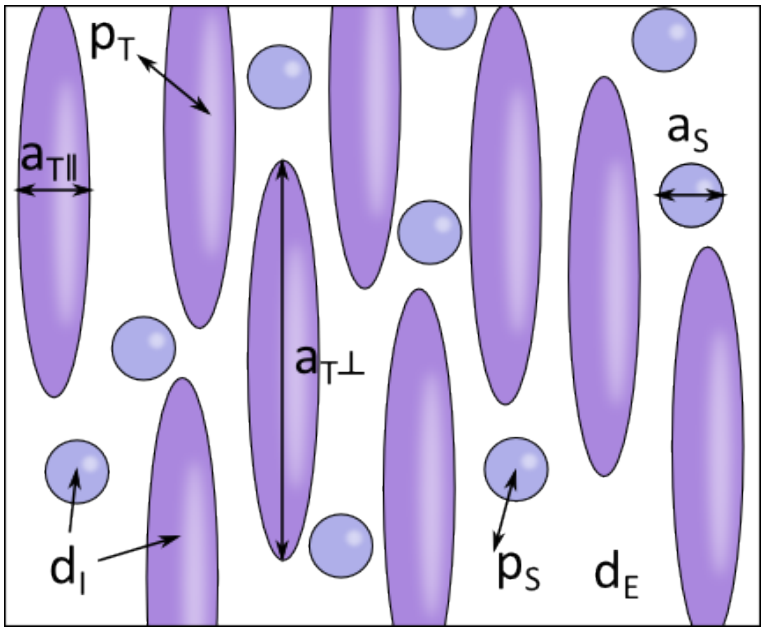

Figure 2.18: A schematic representation of Stanisz's model consisting of prolate axons, spherical glial cells and hindered extra-cellular diffusion. The intra- and extra-cellular diffusivities are d_I and d_E and the permeabilities between the different compartments are p_S and p_T.

capture the behaviour of the measured signals, all three compartments including the permeabilities are needed. This model does provide a thorough description of water diffusion within tissue and incorporates many of the parameters we would like to estimate clinically; however due to both the model complexity and the high data requirements, it is very unlikely that this model could be fit to in vivo measurements obtained using a clinical scanner. We need to consider simpler models that can be fit to clinical data, whilst still modelling the underlying tissue structure as realistically as possible.

Behrens' ball and stick model

The ball and stick model (named following the taxonomy from [134]) introduced by Behrens et al [26] provides a very simple model of the diffusion MR signal that can be readily fit to clinically acquired data. The ball and stick model divides white matter tissue into two distinct compartments: an intra-axonal compartment and an extra-axonal compartment. The total measured signal is a weighted sum of the signals due to each compartment, where the weighting parameter f represents the fraction of the signal due to intra-axonal water (also referred to as the intra-axonal or intra-cellular volume fraction). Behrens models the axons as zero radius cylinders (the stick) with a dominant fibre direction described by the unit vector $\hat{\mathbf{e}}$. Intra-axonal water diffuses anisotropically along the dominant fibre direction with diffusivity d. In the extra-axonal compartment (the ball), the diffusion of water around the fibres is isotropic, also with diffusivity d. A schematic representation of the ball and stick model is shown in figure 2.19. The model for the signal may be expressed as

$$S(\mathbf{G}, TE) = S_0(f \exp(-bd(\hat{\mathbf{e}} \cdot \hat{\mathbf{G}})^2) + (1-f)\exp(-bd)), \tag{2.53}$$

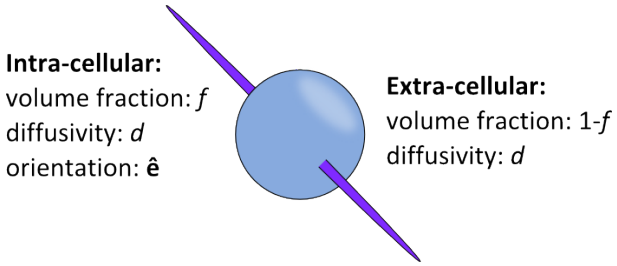

Figure 2.19: A schematic representation of Behrens' ball and stick model. The zero radius axonal compartment, with direction $\hat{\mathbf{e}}$ and diffusivitiy d, contributes a fraction f of the total signal. Extra-axonal water diffusion is isotropic, also with diffusivitiy d.

where $\hat{\mathbf{G}}$ is a unit vector in the direction of the diffusion-weighting gradient. From this model, we can estimate the diffusivity d, the fibre orientation $\hat{\mathbf{e}}$, and the volume fraction f of intra-axonal water in each voxel. Although it is a very simple model, it accounts for restricted diffusion (albeit in a very simple way) and its parameters are more closely related to the underlying tissue properties than FA or MD. The ball and stick model is easy to fit to sparse diffusion-weighted datasets and it has been used in many studies to date [25, 88, 155, 127]. However, it is a very simplistic model and lacks true microstructure indices such as axon radius. More realistic models can be constructed by modelling the axons as cylinders with non-zero radius, and we introduce examples of these models next.

CHARMED and AxCaliber

Like Behrens' model, the composite hindered and restricted model of diffusion (CHARMED) [16] introduced by Assaf et al models white matter tissue using two compartments, one intra-axonal and one extra-axonal. However, unlike Behrens' simple intra-axonal stick model, Assaf represents axons as parallel cylinders whose diameters $2R$ follow a gamma distribution

$$\Pr(2R) = \Gamma(2R; \alpha, \beta) = 2R^{\alpha-1} \frac{\exp\left(-\frac{2R}{\beta}\right)}{\beta^{\alpha}\Gamma(\alpha)} \tag{2.54}$$

where $\Gamma(\alpha) = (\alpha - 1)!$, α is the shape parameter and β is the scale parameter. A gamma distribution is chosen as it approximates histograms of axon radii obtained from histology well [16]. Intra-axonal water molecules undergo restricted diffusion and the signal model for this compartment is given by the expression for restricted diffusion in a cylinder of radius R derived by Neuman [128], which makes use of the GPD approximation discussed in section 2.3.3. Water in the extra-axonal compartment undergoes hindered diffusion, and its signal is modelled using a diffusion tensor whose prinicipal eigenvector is aligned with the orientation of the cylinders. Using the taxonomy of [134], this model can also be referred to as ZeppelinGDRCylinders. The framework may easily extend to include multiple restricted and hindered compartments. In early studies using the CHARMED framework, which focus on resolving crossing fibres, the axon diameter distribution is fixed (based on spinal cord histology studies) during the parameter estimation. However later work extends the CHARMED framework by including the axon diameter distibution as a function to be estimated. This technique is termed AxCaliber [15]. In order to achieve sensitivity to the whole distribution of axon diameters, AxCaliber requires a large

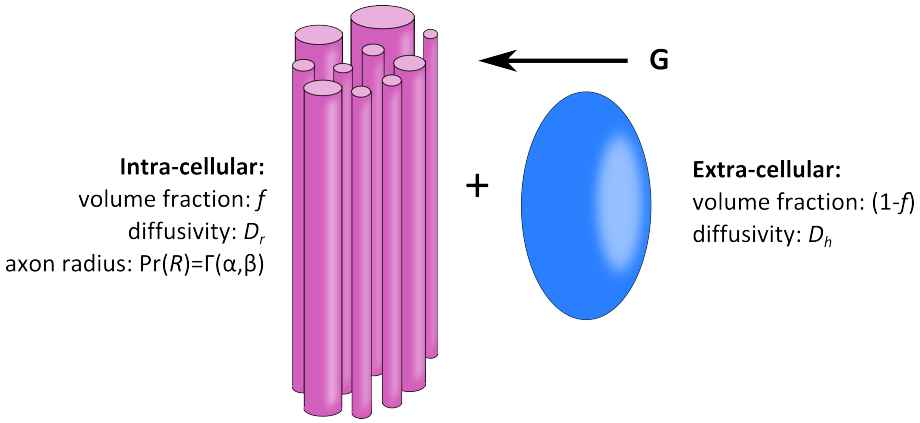

Figure 2.20: A schematic diagram of the AxCaliber tissue model. The intra-axonal compartment comprises axons whose diameters are described by a gamma distribution. Extra-axonal diffusion perpendicular to the fibres is modelled as an isotropic Gaussian distribution.

number of measurements acquired perpendicular to the axons, using high gradient strengths and a range of diffusion times. Because they only consider measurements perpendicular to the fibres, only the perpendicular components of the diffusivities, D_h for the hindered extra-axonal compartment and D_r for the restricted intra-axonal compartment, must be estimated. The diffusivity parallel to the axons is not defined. A schematic representation of the model is shown in figure 2.20. Ex vivo studies performed using fixed sciatic and optic nerve samples show good agreement between histology and the axon diameter distributions estimated using AxCaliber [15]. In further work Barazany uses AxCaliber, although modified to include a CSF partial volume component, to estimate axon diameter distributions in vivo in the rat corpus callosum [19]. Again, the results of this in vivo study agree well with the subsequent histology. Despite the promising results obtained in ex vivo and in vivo studies using CHARMED and AxCaliber, the technique has limited application for in vivo human studies for two reasons: the high gradient strengths and long acquisition times required are not achievable on clinical scanners, and prior knowledge of fibre orientation is required.

Active Imaging and ActiveAx

Alexander [6] attempts to overcome these limitations using two key adaptations. The first is the development of Active Imaging, an optimisation framework that produces orientationally invariant acquisition protocols which minimise the variance on axon radius estimates whilst remaining within the hardware and time constraints of clinical scanners. The second is a simplification of the tissue model proposed by Assaf. Alexander [6] models the fibre population using randomly packed parallel cylinders with a single radius. The intra-axonal signal is given by Stepisnik's expression for the signal from water restricted within a cylinder of radius R [170] (this expression was also independently derived by van Gelderen et al [182]), which uses the GPD assumption described in section 2.3.3. The extra-cellular signal is modelled as a cylindrically symmetric diffusion tensor with principal eigenvector parallel to the fibre direction. Using this model, known as the simplified CHARMED model or ZeppelinCylinder [134], we can esti-

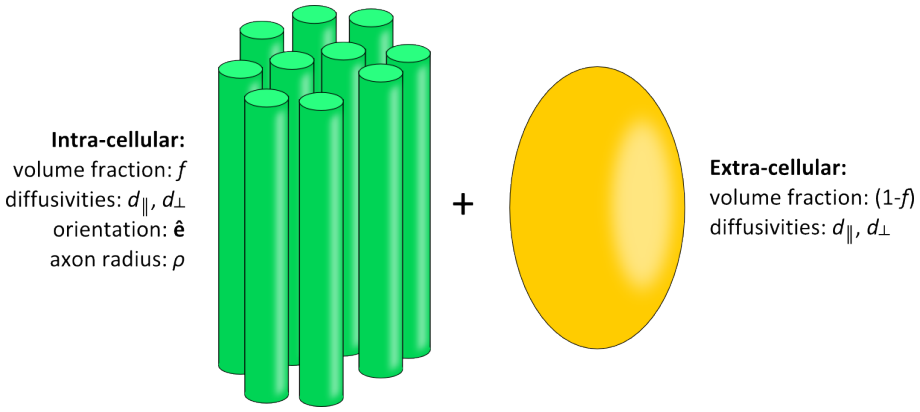

Figure 2.21: A schematic diagram showing the ActiveAx tissue model. The intra-axonal compartment comprises randomly packed parallel cylinders with a single axon radius index ρ. The extra-cellular compartment is a cylindrically symmetric tensor whose principal axis aligns with the fibre orientation.

mate parameters such as an axon radius index ρ, fibre direction $\hat{\mathbf{e}}$ and volume fraction f. The diffusivity parallel to the axons $d_{\parallel}$ is fixed (optimised using a grid search) and a tortuosity relationship [172] is used to relate the perpendicular diffusivity $d_{\perp}$ to $d_{\parallel}$. A diagram showing the compartments of the model is shown in figure 2.21. In recent work Alexander et al [7] demonstrate this technique for axon radius estimation, termed ActiveAx, in the CC in both ex vivo monkey brain ($G_{max} = 300\mathrm{mT/m}$) and in in vivo human subjects ($G_{max} = 60\mathrm{mT/m}$). The simplified CHARMED model is modified to include an isotropic CSF compartment to account for partial volume effects between the CC and CSF in the ventricles. This modified version is known as the minimal model of white matter diffusion (MMWMD), or in the naming framework of [134] ZeppelinCylinderBall. Alexander et al also find that for ex vivo tissue an extra compartment, the dot [134] (which represents completely restricted signal), is needed in the MMWMD to account for trapped water molecules. The expected low-high-low trend of axon radius [2] is recovered in both the monkey and human data, although the human estimates are noisy due to the weak dependence of the signal on ρ. There is also uncertainty in the exact relationship between the axon radius index ρ and the true underlying distribution of axon radii, with Dyrby et al [60] showing that ρ is dependent on G_{max}.

Multi-compartment models of fibre dispersion

The models that we have considered so far assume that all of the axons within a voxel are parallel. Whilst this assumption is valid in white matter tracts with high directional coherence such as the mid-sagittal CC [2] and the spinal cord [53], in other regions of brain white matter, tracts tend to fan out and disperse [8]. In order to model these situations accurately, the orientation dispersion of the axons must be accounted for.

Jespersen et al [91] develop a model for the orientation dispersion of dendrites and axons that can be used to model both grey and white matter. As with the previous models we have described, it comprises two distinct compartments: an intra-cellular compartment, which represents the axons or dendrites, and an extra-cellular compartment, which represents the signal due to unrestricted water molecules. The

extra-cellular compartment is modelled using an isotropic diffusion tensor, although in a later paper [90] they consider cylindrically symmetric tensors instead. Axons and dendrites are modelled using long cylinders, with different diffusivities parallel and perpendicular to the cylinder axis. Using the SGP approximation and the long diffusion time limit, diffusion within the cylinders is assumed to be Gaussian. The orientation dispersion is accounted for by integrating the signal from one cylinder over all possible orientations, using spherical harmonics to capture the dispersion. Using ex vivo monkey brain data, they are able to show that their model explains the data much better than the DT, particularly in grey matter and known regions of white matter fibre dispersion. However, it lacks true microstructure indices such as an axon radius index.

Zhang et al [198, 200] extend the MMWMD to allow for estimation of axon radius in the presence of fibre orientation dispersion. The degree of fibre dispersion is modelled using a Watson distribution [186], which adds an additional parameter κ to the model. The total signal in a voxel is given by integrating the signal from one cylinder of radius R over all possible orientations. The model is fit to in vivo human data in the mid-sagittal and lateral slices of the CC. In the mid-sagittal slice, where fibre dispersion is low, estimates from both the MMWMD and the dispersion model are consistent. However in the lateral slices, where fibre dispersion is higher, the dispersion model is able to provide estimates of the axon radius index ρ consistent with the mid-sagittal results whereas the MMWMD overestimates ρ.

In further work, Zhang et al [201] introduce neurite orientation dispersion and density imaging (NODDI). This technique uses the same model as in [200], although without the axon radius parameter, to allow for fast, practical, whole brain mapping of both orientation dispersion and fibre density. Similar to the work of Jespersen et al [91, 90], this model is applicable to both white and grey matter. Although NODDI lacks a measure of axon size, it is clinically feasible and provides more specific information about the tissue structure than FA. For example, in regions of white matter fibre dispersion, FA is usually lower than in the surrounding, parallel white matter fibre bundles; with NODDI, Zhang et al [201] show that in these regions the fibre density is consistent with that of more coherent regions but with a higher orientation dispersion index.

The Watson distribution utilised by Zhang et al [201] is a symmetric distribution, and thus cannot capture the asymmetric fibre dispersion in fanning tracts such as the corticospinal tract. The orientation dispersion within these tracts can be better modelled using a Bingham distribution, which is a probability distribution on a sphere that allows for anisotropic dispersion about the principal direction. Kaden et al [102] used the Bingham distribution in a parametric spherical deconvolution method for resolving crossing fibre bundles in the presence of fibre dispersion. The Bingham distribution is later used by Sotiropoulos et al [166], who incorporate it into the ball and stick model in order to infer fibre fanning from diffusion-weighted MR data.

General multi-compartment models

All the models described above are specific examples of multi-compartment models. However, a general multi-compartment model can be formed by using any suitable combination of intra-cellular and extra-cellular compartments. This is the approach taken by Panagiotaki et al [133, 134]. They create

a taxonomy of multi-compartment models of the diffusion MR signal in white matter by combining models of the extra-cellular diffusion signal, represented by

- isotropic tensors (ball),
- cylindrically symmetric tensors (zeppelin),
- full tensors (tensor),

with intra-cellular signal models, represented by

- zero radius cylinders (stick),
- single radius cylinders (cylinder)
- gamma-distributed radius cylinders (GDRcylinder) .

They also consider additional compartments to account for glial cells (sphere), trapped water molecules (dot) and fibre dispersion (astrocylinders, astrosticks). Using rat brain data acquired with multiple diffusion times and gradient strengths, Panagiotaki is able to quantitatively compare how well the models fit the data. The best model is the TensorCylinderSphere, which combines a full tensor extra-cellular compartment with a single radius cylinder intra-cellular compartment and a third, non-zero radius spherical compartment to account for water restricted within glial cells. The worst performing model is the DT. Panagiotaki et al also note that whilst the AxCaliber style models with gamma-distributed radius cylinders do fit the data reasonably well, they are also the least stable models to fit and so require the most computational effort.

Summary

In this section, we have described a wide range of analytic models for the diffusion MR signal of varying complexity. Simple models such as the DT [20] and the ball and stick model [26] have low data requirements and are generally robust to noise, but they provide little in the way of true microstructural information. The models described in [6, 15, 134, 200] can provide more detailed tissue measures, although the data requirements are high and the resulting parameter estimates are noisy, particularly when the gradient strength used in the data acquisition is low. However, all the studies mentioned in this section treat each voxel of data independently, and do not exploit the spatial coherence that we know exists in white matter tracts such as the CC [2] and the optic nerve [92]. Using this coherence to guide the fitting procedure may help to reduce the noise on parameters such as the axon radius index ρ and improve the stability of models which aim to estimate the full axon radius distribution. Although simple models such as the DT and the ball and stick model are generally robust to noise, they may also benefit from this natural data pooling when the SNR is low and the number of gradient directions is modest.

2.5 Tractography

One of the key applications to emerge from diffusion MRI is fibre tracking, or tractography. Tractography is commonly used as a method to determine the likelihood of connectivity between different functional

areas in the brain [79], or to identify and segment individual white matter tracts [96]. As many of the shape models we discuss in section 2.7 use tractography as a means of segmentation, we describe the key principles behind fibre tracking briefly.

Starting in a seed voxel (usually manually defined, although there are proposals for completely automatic tractography [40]), tractography uses an estimate of the principal fibre orientation to move along the tract from voxel to voxel, until the whole tract has been reconstructed as a streamline. There are numerous methods available for deciding how to move along the tract. These can be broadly divided into two categories: deterministic and probabilistic.

Deterministic tractography algorithms such as fibre assignment by continuous tracking (FACT) [194] use the first eigenvector $\mathbf{e}_1$ derived from the DT as the principal fibre orientation. A vector field can be constructed over the whole data volume from the estimates of $\mathbf{e}_1$ in each voxel. Starting from the centre of the seed voxel, we can reconstruct the tract by following the local direction of the vector field. In isotropic regions of tissue (e.g. CSF or grey matter) the vector $\mathbf{e}_1$ is often poorly defined, so in order to avoid spurious streamlines, stopping conditions such as FA and angular thresholds are imposed [98]. However, FACT often results in jagged streamlines as its trajectory only changes at the voxel boundaries. This can be ameliorated by defining a step size much smaller than the voxel dimensions and creating a smoother vector field using interpolation [22]. These methods work well when the DT is highly anisotropic [5] and $\mathbf{e}_1$ is well defined. As the DT becomes more oblate (as is the case for crossing fibres) or isotropic, $\mathbf{e}_1$ is often ambiguous and unreliable for fibre tracking. An alternative method [109], known as tensor deflection (TEND), uses the whole DT to determine how the trajectory of the streamline changes as it passes through each voxel. This has the advantage of being able to pass through low anisotropy regions, as an isotropic tensor does not deflect the fibre trajectory, although it still struggles to resolve crossing fibres when the crossing angle is oblique [5]. It is also possible to combine these two approaches and describe the fibre trajectory as a weighted combination of $\mathbf{e}_1$ and the tensor deflection [189, 109].

One of the main disadvantages of deterministic trajectory is that we cannot quantify the uncertainty in the resulting streamlines. Probabilistic tractography algorithms provide a way to incorporate this uncertainty directly into the streamlines. Instead of using a single estimate of fibre orientation in each voxel, probabilistic tracking algorithms utilise pdfs of fibre orientations. These pdfs can be defined by linking properties such as diffusion anistropy to the uncertainty [137] or estimated directly from the noisy diffusion-weighted data using Bayesian [26] or bootstrapping methods [93]. One of the most common methods of probabilistic tracking is to build up a distribution of streamlines [26, 137]. The tracking algorithm is similar to the deterministic method described in [22], although at each voxel the trajectory of the streamline is determined by sampling the fibre orientation pdf. The procedure is repeated many times (preferably thousands) in order to build up a representative distribution of trajectories. By normalising the number of streamlines passing through a voxel by the total number of streamlines generated, we obtain an indication of the likelihood that a voxel is part of the tract. Thresholding is often used to get rid of low probability and thus potentially spurious voxels, although the threshold used is often arbitrary

[135]. Using probabilistic tractography, it is also possible to reconstruct crossing, bending or fanning fibres using fibre orientation distribution functions derived from multiple fibre reconstruction techniques such as spherical deconvolution [174], q-ball [176] or persistent angular structure (PAS) MRI [87].

More recently, there has been a move towards global tractography methods [89, 152, 162] which aim to reconstruct all fibre tracts in the brain simultaneously, instead of only those between two predetermined regions of interest. This requires a lengthy, large-scale optimisation, but has been shown to increase the accuracy of recovered tracts, particularly when prior information about the tracts is incorporated into the optimisation algorithm. For example, MicroTrack, the approach proposed by Sherbondy et al [162] simultaneously performs tractography and estimates microstructure parameters such as an axon radius. Consistency of microstructure along the tracts is enforced by the tractography algorithm, improving the accuracy of the reconstructed tracts, particularly through crossing fibre regions.

2.6 Group studies of white matter

In section 2.4 we introduced a wide range of analytical models of the diffusion MR signal with parameters ranging from simple anisotropy measures to axon radius indices. However, regardless of which model we select, after we fit the model to the data and obtain estimates of the tissue parameters, our key aim is generally the same. As stated in section 2.1.2, we want to investigate how the parameters we estimate are related to tissue changes that occur due to pathology or development, or measures of function or performance. Historically, this has been done using either region of interest (ROI) analysis or voxel-based morphometry (VBM) [14]. Both of these methods test for and detect statistical differences in tissue parameter estimates across populations; however the motivation behind them is different. Region of interest analysis is a hypothesis driven method which aims to test whether there is a change in a given tissue parameter in a specific area of the brain between two groups of subjects, e.g. the ROI analysis in [144] aims to investigate whether there are differences in FA and MD in the splenium and genu of the CC due to schizophrenia. VBM and other voxel-based methods are more exploratory; they test for group differences in tissue indices across all voxels in the brain to see which, if any, voxels are affected. For example, the VBM analysis in [159] aims to investigate whether juggling training over a prolonged period of time induces any changes in FA or MD in any voxels in the brain. Following a training period, they find an increase in FA in the white matter surrounding the intraparietal sulcus, a grey matter region associated with motor coordination. In this section, we give an overview of how these two methods, as well as more recent voxel- or region- based analysis techniques are performed.

During an ROI analysis, the white matter tract under consideration is segmented, either manually as in [144] or using tractography as in [96], and the microstructural parameters of interest are averaged over the whole region. These values can then be compared between subjects or groups. Averaging helps to negate the effects of noise, although if the ROI is large it may overshadow subtle but significant changes. It is also impossible to localise where the changes are occuring within the region. Although this can be improved by selecting smaller, more specific regions, this leads to correspondence problems between subjects. The placement of boundaries between smaller ROIs is also somewhat arbitrary and can overshadow group differences, particularly if reasonably focal differences straddle two regions.

VBM, unlike ROI analysis, analyses changes in tissue parameters across all voxels in the whole brain at once. The data from each subject is registered into a common space and, voxel by voxel, the parameters are compared between the populations. Originally designed to detect grey matter volume changes in structural T_1-weighted images, this approach has since been used in numerous diffusion MRI studies to detect changes in FA and MD [34, 65]. Whilst VBM does remove the need for specific hypotheses and ROI selection, there are several caveats. First, the registration is often imperfect which means that the voxels being compared between subjects are not equivalent. This is especially noticeable when the pathology being studied has significant volume changes. For example, in schizophrenia the enlarged ventricles are often improperly registered and lead researchers to detect significant changes in these regions which are in fact spurious [164]. Second, in order to perform the comparisons, the data must be Gaussian distributed. As shown in [99], a large number of voxels in diffusion MRI datasets are not Gaussian distributed. To reduce the number of non-Gaussian voxels, and thus fulfil the condition, a smoothing kernel is usually applied to the data. The larger the kernel, the more Gaussian-distributed the data becomes, although this increases partial volume effects. The size of kernel has also been shown to influence whether a voxel exhibits significant changes in FA or not [99]. Finally, comparing every voxel within the brain results in a very large number of multiple comparison corrections which could reduce the significance of any findings.

Tract-based spatial statistics (TBSS) [164] is a diffusion MRI specific alternative to VBM for performing whole brain group analysis without the requirements of smoothing or exact alignment. Data produced using TBSS is also Gaussian-distributed [164]. Instead of using all voxels within the brain, TBSS creates a skeleton of the white matter tracts from the mean FA image calculated over all subjects. For each subject, the peak values of the parameter estimates are projected onto the skeleton, and all analysis is performed only on those voxels within the skeleton. This reduces the number of multiple comparison corrections required, however the method does get rid of a lot of potentially useful information within the tracts and could be more susceptible to noise when the SNR of the data is low. Finally, it is not easy to divide the skeleton into anatomically separate tracts [197] or to incorporate other indices such as tract thickness into the analysis.

Another alternative method is to use tract-based analysis which, like ROI analysis, focuses on detecting differences in specific white matter regions. However, the crucial difference is that it can localise the microstructure indices within the tract in a way that reflects its natural geometry. This involves creating a shape model of the tract which incorporates size, shape and thickness information. The information gained from the shape model allows us to normalise distances along the tract, which, when combined with the microstructural information from the diffusion MRI data allows us to generate multiscale white matter models. The profiles of microstructure variation along the tract can then be used in group comparison studies.

2.7 Shape and multiscale models for white matter

There are numerous different methods for creating shape models, many of which originate from computer vision techniques [46, 47]. In this section we describe two methodologies that have been widely

used in diffusion MRI studies.

The shape model proposed by Gerig et al [68] uses deterministic tractography and clustering algorithms to segment specific bundles of white matter fibres. Within the bundle, each streamline is parameterised as a B-spline curve (more detail on B-spline curves is provided in section 2.9.1) and indices of tissue integrity such as FA and MD are mapped along the splines using interpolation. Finally, point to point correspondence is established between the streamlines by comparing local shape measures such as curvature, torsion and arc length. This approach is also used by Corouge et al [49]. However, instead of mapping microstructural indices, they register all the streamlines in the bundle together and average them to create a mean tract. Principal component analysis [29] is then applied to the registered streamlines in order to determine the variability of the streamlines around the mean. A later paper [48] combines both of these approaches in order to create a shape model that describes both microstructure and shape variation. However, these shape models are only demonstrated on individual datasets. Whilst Corouge et al [48] demonstrate reproducibility between healthy controls, in cases of pathology, the number of streamlines generated by the tractography algorithms and their characteristics may be very different. In this case, it is difficult to establish correspondence. In later work by the same group [74], several modifications are used in order to overcome these issues. First, an atlas of tracts is created using deterministic tractography and clustering. The atlas tracts are then registered to each subject's data in order to segment the tracts in their native space and to provide correspondence between subjects. Second, instead of mapping the indices along each individual streamline, one mean curve of microstructure variation along the tract is created per subject, which is also represented using B-splines in order to get rid of excessive variability due to noise. Similar approaches are also proposed by Gong et al [73], Maddah et al [115], Goldsmith et al [71], Clayden [41] and Colby et al [43], although Maddah uses a pointwise representation of the parameter variation instead of B-splines. One drawback of this approach is that when creating the mean curve, the microstructure indices are integrated over the cross-section of the tract. Therefore we lose information about the parameter variation in this plane, which is arguably more important than the variation over the arc length, along which those parameters are expected to remain more consistent. As discussed in section 2.1.2, once an axon has been damaged it decays along its whole length, destroying the connection. It often does not matter where the damage occurs along the length of the axon, however it does matter where it occurs within the cross-section of the tract as this determines which connection is affected. For example, tract-based analysis is often used to examine changes within the splenium [115, 71]. The anterior splenium contains the interhemispheric connections between the auditory system whilst the posterior splenium connects the visual cortices [85]. Examining microstructure along the entire arc length detects whether differences exist between populations, but without the cross-sectional information we cannot determine which of the functional systems is affected. This method is also only able to characterise thin, tubular tracts, such as the fasciculi or the genu and splenium of the CC, and is not able to represent larger, sheet-like tracts such as the whole of the CC, or the corticospinal tract.

An alternative method for representing the shape of tracts is via medial representation. Medial models were first introduced by Blum [31] in order to represent complex two- and three-dimensional

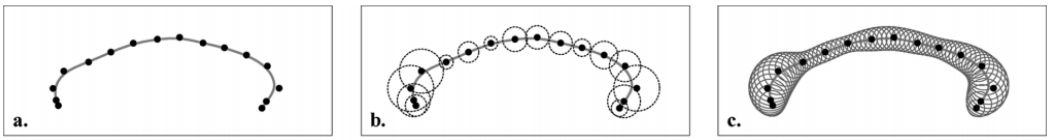

Figure 2.22: This figure, taken from [195], illustrates the process of inverse skeletonisation. In (a) the medial axis of the CC is represented using a B-spline curve. The associated radial function is used to determine the thickness of the tract at every point along the axis. This can be illustrated by drawing circles of the corresponding radius at each point on the axis, as shown in (b). If the medial axis and radial functions are continuous, this can be applied continuously along the whole curve in order to give a boundary, shown in (c).

shapes by simple skeletons, also called medial axes or surfaces. For white matter, this can be done by 'thinning' the tract in the direction perpendicular to its boundary, subject to distance constraints which require that the medial axis is always equidistant from the boundaries. However, dependent on the initial shape, this can result in unwanted branching which is not necessarily representative of the underlying structure. Golland et al [72] demonstrate this approach in two dimensions in the mid-sagittal CC, however the resulting piece-wise skeletons are not smooth across the whole curve.

Continuous-medial representation (cm-rep) has been introduced in order to provide smooth, parametric skeletons without spurious branches. Instead of deriving the skeleton from the boundaries of the object, cm-rep starts with the medial axis and uses inverse skeletonisation to determine the appropriate boundary [195]. Yushkevich et al [195] have demonstrated that the boundary of an object can be uniquely determined by its medial representation and a radial function determining the thickness of the tract at each point. Figure 2.22, taken from [195], shows an example of inverse skeletonisation for the mid-sagittal CC. In three-dimensions the skeleton is a medial surface which is parameterised using a triangular mesh. In order to fit the cm-rep to a tract (which is typically segmented using deterministic tractography algorithms such as FACT [194]), a boundary is generated from initial mesh and thickness functions. The overlap between the true and predicted boundary can be calculated using a Dice overlap score [55] and the parameters of both the mesh and thickness functions are updated until the Dice score is maximised. Continuous medial representation has also been formulated in two dimensions [171], using B-splines to parameterise the medial axis. One advantage of the cm-rep formulation is that the skeleton and boundaries provide an alternative coordinate system in which to analyse the tracts which naturally reflects the geometry of the white matter structures. This shape-based normalisation allows us to establish correspondence between the subjects in a natural way. Microstructural indices such as FA or MD can be sampled within this geometry to produce pointwise functions of parameter variation along the skeleton [197, 171]. The radial thickness function also provides information about the shape of the tracts which can be combined with microstructural information in multivariate group comparisons, as in the Tract Specific Analysis framework introduced by Zhang et al [199].

Despite the complex shape models that have been developed, when it comes to the microstructure

parameters, the methods described above use only the simplest diffusion models and they sample parameter estimates along the curves by fitting models of the diffusion MR signal to each voxel independently. This ignores the inherent spatial and structural coherence that exists in a large number of white matter tracts that would naturally compliment such white matter shape models. Whilst this may not be an issue when fitting simple diffusion models to high quality data, when the data is sparse and noisy or the model being fitted is complex, estimates obtained via independent voxel fitting may be unreliable. Incorporating local spatial information could greatly improve the accuracy and precision of our estimates, allowing us to perform reliable group comparisons for parameters such as axon radius.

2.8 Spatial regularisation

As mentioned in both sections 2.4 and 2.7, incorporating prior information about the spatial coherence of tissue parameters can improve our estimations by reducing the influence of noise. This can be done using spatial regularisation, a technique that is used to enforce similarity between tissue measures or signals over a neighbourhood of voxels.

When we fit a model of the diffusion MR signal to our measurements, we aim to find the set of parameters $\boldsymbol{\Theta}$ that maximises the likelihood of obtaining those measurements. This is typically done by minimising an objective function $\Omega(\boldsymbol{\Theta})$, usually the sum of squared differences between the signals predicted by the model for a given $\boldsymbol{\Theta}$ and the signals that were measured. The sum of squared differences, or equivalent measure, is often referred to as the data fidelity term $\Psi(\boldsymbol{\Theta})$ [39, 138]. However, as the measurements are affected by noise, the optimised set of parameters $\hat{\boldsymbol{\Theta}}$ may not correspond to the ground truth. We can address this by modifying the objective function so that $\Omega(\boldsymbol{\Theta}) = \Psi(\boldsymbol{\Theta}) + \chi(\lambda, \boldsymbol{\Theta})$, where χ is the regularisation term and λ is a parameter than controls its weight. The regularisation term χ penalises the objective function based on features of the parameter vector, such as its size or gradient, thus favouring certain types of solutions.

Within diffusion MRI, spatial regularisation has been widely used to improve the estimation of tensor fields [50, 185, 175] and subsequent fibre tracking [142, 103, 150]. Diffusion MRI studies which utilise spatial regularisation typically do so in one of three main ways:

1. Regularisation of diffusion-weighted images prior to model fitting [136, 56].
2. Regularisation of tensor or eigenvector fields after model fitting [50, 142, 175, 83]
3. Regularisation of parameter estimates during model fitting [185, 184, 150, 103, 54].

In this section, we introduce some of the key methods for regularisation and discuss how they have been used in diffusion MRI studies.

One of the most widely used techniques, generalised Tikhonov regularisation [28], defines χ as

$$\chi = \lambda \int_{\Omega} |f(\boldsymbol{\Theta})|^2 d\boldsymbol{\Theta}, \tag{2.55}$$

where λ controls the strength of the regularisation term, which is integrated over the whole of the image domain Ω. A common choice is $f(\boldsymbol{\Theta}) = \nabla\boldsymbol{\Theta}$, which minimises sharp transitions in $\boldsymbol{\Theta}$ across the

image domain, or $f(\boldsymbol{\Theta}) = \Delta\boldsymbol{\Theta}$, which penalises the divergence of the image gradient, and is a good measure of smoothness over a sphere. The former representation for χ is particularly favoured when fitting splines or curves to data, such as in [62, 107], whilst the latter has been used to regularise fibre orientation dispersion functions [150] within a Bayesian model fitting framework. However, this form of regularisation enforces smoothness across the whole image domain and cannot capture discontinuities, such as those that occur at white matter fibre crossings or white matter/grey matter interfaces, and thus is not commonly used in diffusion MRI studies.

Due to the lack of edge preservation with Tikhonov regularisation, when considering whole brain images, total variation regularisation [183] is often preferred, which defines χ as

$$\chi = \lambda \int_{\Omega} |\nabla\boldsymbol{\Theta}| d\boldsymbol{\Theta}. \tag{2.56}$$

Unlike Tikhonov regularisation, total variation does not require $\boldsymbol{\Theta}$ to be continuous over Ω, therefore the sharp boundaries between different tissue types in the brain can be recovered [183]. This method has been widely used in the post tensor fitting regularisation of both tensor [175] and principal eigenvector fields [50]. It has also been used in a study by Wang et al that incorporates spatial regularisation directly into tensor estimation [185]. They compare their method to post model fitting regularisation techniques and show improved accuracy in the reconstructed tensor fields. More recently, a variational regularisation framework has been used by Pasternak et al [138] to estimate the amount of free water contamination within the brain. Again, the use of the variational framework to preserve sharp boundaries is important, as regions of oedema can be well localised [138].

Another common form of regularisation for smoothing tensor fields is anisotropic smoothing [136, 56]. This form of regularisation uses information from the gradient of the image intensity to smooth homogeneous regions isotropically, but to smooth boundary regions anisotropically. Specifically, smoothing is favoured in the direction tangential to the boundary rather than along the direction normal to the boundary. This is done using a structure tensor [56]. The structure tensor uses information from the image intensity gradient, but rescales it so that the principal direction of the structure tensor corresponds to the direction of smallest intensity change. This method has been used by Parker et al [136] and Ding et al [56] to smooth raw diffusion-weighted images prior to fitting the diffusion tensor model. Van Hecke et al [83] have also used this approach to smooth FA maps in a VBM population analysis. Using simulated DTI data, they compare anisotropic and isotropic smoothing and find that the anisotropic method is more sensitive and specific to pathological regions.

For Tikhonov, total variation regularisation and anisotropic smoothing, the magnitude of the χ is controlled by a regularisation parameter λ, which must be tuned in order to find the optimal solution. Methods for doing this vary widely. Some studies formulate an error metric specific to the problem at hand [50] and find the value of λ that minimises it. More generic methods of determining λ include generalised cross validation [28], which balances the residual error against the effective degrees of freedom, and L-curve [27], which finds the best trade-off between minimising the residual error and maximising the smoothness. However, determining the appropriate regularisation parameter is a very difficult problem and there is no general solution.

An alternative but similar concept is the use of Markov random fields (MRF) to spatially regularise the model fitting procedure. An MRF is a probability field defined on a graph [147]. In the case of diffusion MR images, the vertices V of the graph represent the individual voxels, and the edges E connect voxels within the same local neighbourhood. The probability of observing parameters $\boldsymbol{\Theta}_i$ in voxel v_i is conditionally dependent upon only those parameters within its local neighbourhood. These dependencies are very often formulated as Gaussian distributions, resulting in a Gaussian MRF [184]. The strength of the dependencies is governed by the precision matrix $\boldsymbol{\Pi}$, similar to the parameter λ described above, which must be determined during the regularisation procedure. However, MRFs can be formulated as Bayesian hierarchical models [103], allowing $\boldsymbol{\Pi}$ to be estimated directly during the regularisation, rather than estimated separately as discussed previously. It is also possible to incorporate hidden variables within the MRF, allowing dependencies between voxels to be 'switched' on and off, thus allowing edges to be conserved [29]. In the majority of cases, connections are only defined between voxels than are immediate neighbours. It is possible to increase the order of the MRF in order to increase the spatial extent of the dependencies, however this often overcomplicates the model making it difficult to find a computationally efficient solution [147]. In the context of diffusion MRI, Gaussian MRFs are most commonly used to regularise and smooth fibre orientation or tensor fields in order to improve tractography [103, 142] or resolve crossing fibres [54]. However MRFs can also be used to regularise non-directional tissue parameters, as in the recent study by Walker-Samuel et al [184], which uses a Gaussian MRF to improve estimation of the apparent diffusion coefficient in low SNR data.

2.9 Basis functions

As stated in chapter 1, the key contribution of this thesis is to introduce the regional variation model framework for estimating the trends of tissue parameters across tracts, using information about spatial coherence to guide and constrain the fitting. These trends are best represented using curves. However the shapes of these trends may be quite complex and variable across subjects; thus they cannot be represented by a single mathematical function. Instead we can use basis functions, which provide a set of building blocks from which we can construct a wide range of functions by varying the relative weights of the basis elements. There are numerous basis systems available, including polynomials, Chebyshev polynomials, Fourier basis functions, splines and wavelets. In this thesis we use B-spline basis functions to represent the trends in tissue parameters.

2.9.1 B-splines

B-splines are one of the most commonly used set of basis functions for representing curves. We have already briefly mentioned B-splines in section 2.7 with regards to shape models of white matter tracts; here we describe their properties and construction in greater detail.

Splines are piecewise-polynomial functions that are commonly used to construct curves in computer graphics [64], represent displacement fields in image registration algorithms [156] and to fit curves to data points in regression problems [177]. One of the most common and convenient representations of a spline function is the basis spline, or B-spline form. B-splines are composed of a weighted sum of k^{th}

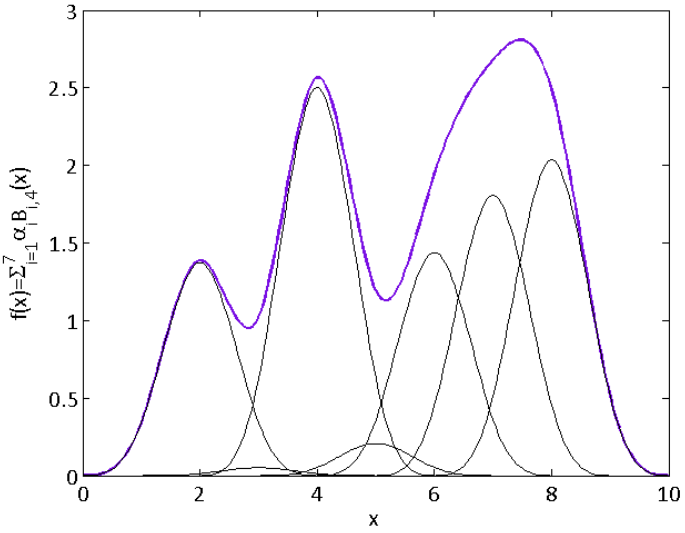

Figure 2.23: An example of a B-spline function created from seven cubic basis functions. In this case, the knot separation is uniform $(t_{j+1} - t_j = 1)$, therefore the shape of the underlying functions is the same. The overall value of the function $f(x)$ is each point is determined by the set of weights $\{\alpha\}_{i=1}^{7}$.

order polynomial basis functions $B_{i,k}$ defined over a set of knot points $t = \{t_j\}_{j=1}^{m}$ where $t_j \leq t_{j+1}$ [160]. We show an example of a B-spline function and its constituent basis functions in figure 2.23. The general form of the spline function $f(x)$ can be expressed as

$$f(x) = \sum_{i=1}^{N=m-k} \alpha_i B_{i,k}(x), \tag{2.57}$$

where α_i is the weight associated with the i^{th} basis function. As we can see from equation 2.57, the number of basis functions forming the curve depends on both the order k and the number of knot points m. The form of the individual basis functions $B_{i,k}$ is determined by a recursive relationship [32]. Starting with the 1^{st} order basis functions $B_{i,1}$, which are defined as

$$B_{i,1}(x) = \begin{cases} 1 & \text{if } t_i \leq x < t_{i+1} \\ 0 & \text{otherwise} \end{cases} \tag{2.58}$$

higher order basis functions are calculated according to the rule

$$B_{i,k}(x) = \frac{x - t_i}{t_{i+k-1} - t_i} B_{i,k-1}(x) + \frac{t_{i+k} - x}{t_{i+k} - t_{i+1}} B_{i+1,k-1}(x). \tag{2.59}$$

From the definition of $B_{i,1}$ and equation 2.59, we can deduce several important points about the shape of the basis functions. First, a k^{th} order basis function is only defined on the interval $[t_i, t_{i+k})$ and is zero everywhere else, i.e. the basis functions have local support. Therefore, unlike polynomials or other parametric representations such as Bezier curves [42], altering the weight α_i only alters the shape of the curve locally around $B_{i,k}$, not globally. This is a particularly useful property when constructing or fitting curves as we have much more control when manipulating the functions. Second, a k^{th} order basis function is composed of k pieces, each of which is a polynomial of order k. Therefore a 1^{st} order basis

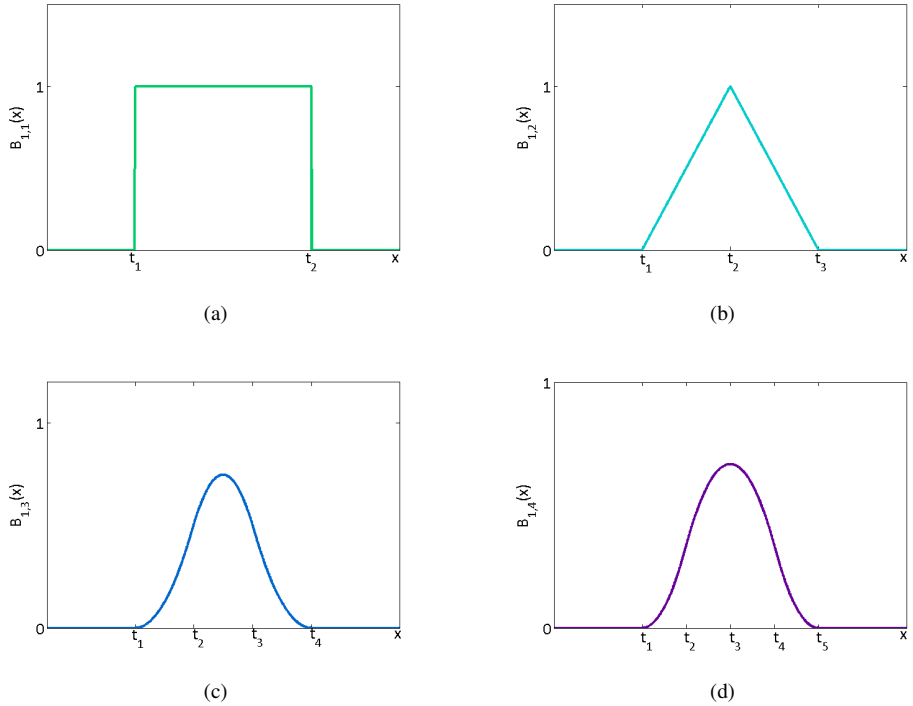

Figure 2.24: Individual basis functions for (a) $k = 1$, (b) $k = 2$, (c) $k = 3$, (d) $k = 4$. The number of knots required to support a single basis function is $k + 1$, and outside the range of the knots $[t_k, t_{k+1})$, the basis function is zero everywhere.

function consists of one 1^{st} order polynomial (i.e. a constant), a 2^{nd} order basis function comprises two linear polynomials, a 3^{rd} order comprises three quadratic polynomials, and so on. Therefore, as the order of the B-spline increases, the basis functions generally become smoother. We show examples of 1^{st} to 4^{th} order basis functions in figure 2.24. Third, given that a k^{th} order basis function can be written as a sum of $k-1^{th}$ order basis functions, its derivative can be written in terms of $k-2^{th}$ order basis functions [160]. When the knot sequence is strictly increasing, i.e. $t_i < t_{i+1}$, then $B_{i,k}$ is $k-1$ times differentiable [177]. However, it is also possible to have multiply defined knots such that $t_i = \cdots = t_{i+p}$ where p is the multiplicity of the knot. In this case, $B_{i,k}$ is only $k - p$ times differentiable. Knot multiplicity can be beneficial as it forces the B-spline to pass through a particular point, such as a specified start or end point [42]; however it can also introduce discontinuities into the curve. Finally, the overall shape of the basis functions depends not only on k, but the relative spacing between the knot points. In the special case that the knot vector is uniform, i.e. $t_{i+1} - t_i = c$ where c is a constant, the recursion relation defined in equation 2.59 can be simplified [32] to

$$B_{i,k}(x) = \frac{x - t_i}{c(k-1)} B_{i,k-1}(x) + \frac{t_{i+k} - x}{c(k-1)} B_{i+1,k-1}(x). \tag{2.60}$$

In this case, all basis functions of order k have the same general form [32], but are shifted relative to each other by c.

As we use B-splines to fit curves to observed data points, we now discuss some general techniques for spline regression. When fitting B-spline curves to our observations we aim to find the set of basis

functions $B_{i,k}$ and their associated weights α_i that best describe the data. There are two key strategies to do this [111]. The first advocates choosing a sparse set of knot points whose basis functions reflect the underlying trends within the data. This requires us to determine both the knot vector t and the weight vector $\alpha = \{\alpha_i\}_{i=1}^{N=m-k}$ simultaneously. The second strategy suggests choosing a large uniform knot vector a priori and using penalty terms to regularise the solution. Examples of this method include penalised splines (also called P-splines) [62] and smoothing splines [32]. These two approaches are related, although the formulation of the penalty is different: P-splines only consider the differences between adjacent α_i, whereas smoothing splines require differentiation and integration along the whole curve. Therefore from a computational point of view, P-splines are more efficient to evaluate. When using this type of method, only the weight vector α and a regularisation parameter λ must be found.

2.10 Bayesian model fitting

Whether we are fitting a diffusion MR signal model, a shape model or fitting a spline to data points, the key goal of model fitting is the same. We want to find the model parameter vector $\boldsymbol{\Theta}$ that most likely explains the data we observe. There are several approaches to model fitting, including maximum likelihood (ML), maximum a posterior (MAP) and full Bayesian inference [147].

ML estimation aims to find the parameter vector $\hat{\boldsymbol{\Theta}}$ that maximises the likelihood (or equivalently and more commonly, the minimum of the negative logarithm of the likelihood) of the parameters given the data $\Pr(D|\boldsymbol{\Theta})$ [114], where $D = \{d_i\}_{i=1}^{N}$ is the set of all data points d_i we observe. Depending on the exact form of $\Pr(D|\boldsymbol{\Theta})$ it may be possible to evaluate $\hat{\boldsymbol{\Theta}}$ analytically. When this is not possible non-linear optimisation techniques such as gradient descent methods [143] or the Levenberg-Marquardt algorithm [116] may be used. However, there are problems associated with ML estimation, such as local maxima (or minima in the negative log likelihood) in the objective function which prevent us from reaching the global maximum likelihood.

This problem can be potentially alleviated by considering the posterior probability of $\boldsymbol{\Theta}$ instead of the likelihood function. The posterior probability $\Pr(\boldsymbol{\Theta}|D)$ is the probability distribution over the unknown parameters $\boldsymbol{\Theta}$, conditional on the evidence we gain from observing the data points D. From Bayes' theorem [110] we can evaluate the posterior probability as

$$\Pr(\boldsymbol{\Theta}|D) = \frac{\Pr(\boldsymbol{\Theta}, D)}{\Pr(D)} = \frac{\Pr(\boldsymbol{\Theta})\Pr(D|\boldsymbol{\Theta})}{\int \Pr(\boldsymbol{\Theta})\Pr(D|\boldsymbol{\Theta})\, d\boldsymbol{\Theta}} \tag{2.61}$$

where $\Pr(\boldsymbol{\Theta}, D)$ is the joint distribution of $\boldsymbol{\Theta}$ and D, which can also be written in terms of the likelihood $\Pr(D|\boldsymbol{\Theta})$ and a prior distribution $\Pr(\boldsymbol{\Theta})$. The prior probability distribution expresses the amount of uncertainty we have about the parameter vector $\boldsymbol{\Theta}$ before we make any observations, or as discussed in section 2.8, it could represent a regularisation term that represents the similarity between neighbouring voxels. The term $\Pr(D)$ is the evidence we gain from our observations, which can be calculated by integrating the joint distribution over all possible values of the parameter vector $\boldsymbol{\Theta}$. We can find $\hat{\boldsymbol{\Theta}}$ that maximises $\Pr(\boldsymbol{\Theta}|D)$ using MAP estimation. Very often, it is more convenient to work with the logarithm of the posterior distribution, which is equivalent to the sum of the log of the likelihood function and the log of the prior. As with maximum likelihood estimation, it may be possible to calculate $\hat{\boldsymbol{\Theta}}$ an-

alytically; otherwise non-linear optimisation methods can be used. The advantage of this method is that $\Pr(D)$ is not dependent on $\boldsymbol{\Theta}$ therefore we do not need to calculate it explicitly [114]. However, MAP estimation can also be affected by local peaks in $\Pr(\boldsymbol{\Theta}|D)$ and the maximum may not be representative of the overall shape of the posterior distribution [114].

An alternative method is Bayesian inference which calculates the full posterior distribution from equation 2.61. This gives us more detailed information about the range of parameters $\boldsymbol{\Theta}$ that would explain the observed data [126]. However, it does require us to evaluate $\Pr(D)$ by integrating over the whole of the parameter space. Depending on the complexity of the likelihood and prior distribution, this may be intractable resulting in no closed form solutions for $\Pr(\boldsymbol{\Theta}|D)$. In this case we can approximate the posterior distribution using Markov chain Monte Carlo. This allows us to draw samples from $\Pr(\boldsymbol{\Theta}|D)$. As we generate more and more samples from $\Pr(\boldsymbol{\Theta}|D)$, our approximation converges towards the true distribution.

Before discussing Markov chain Monte Carlo in more detail, we briefly discuss the choice of $\Pr(\boldsymbol{\Theta})$. Prior distributions are often chosen so that they are the conjugate of the likelihood function [147]. This means that when multiplied together, the resulting posterior distribution has a closed form solution which is also a standard probability distribution [114]. The shape of the prior distribution is often controlled by a small number of hyperparameters which are chosen a priori. These are often chosen so that the prior distribution is uninformative and does not overly bias the form of the posterior distribution [126], although if there is strong information to suggest a particular form for the prior distribution, this is often used instead.

2.10.1 Markov chain Monte Carlo methods

Markov chain Monte Carlo (MCMC) is a technique for drawing samples from a target probability distribution which we cannot sample directly. Typically, the distributions we consider are those which we can only evaluate up to a normalising constant [10, 114]. This is the case for the posterior distribution defined in equation 2.61, for which we can evaluate $\Pr(\boldsymbol{\Theta})$ and $\Pr(D|\boldsymbol{\Theta})$, but not the normalising constant $\Pr(D)$. Once we have obtained samples from the distribution, we can use Monte Carlo integration to estimate moments of the posterior distribution, such as its mean and variance. Due to the law of large numbers [126], if we draw enough samples, our Monte Carlo estimates for the moments converge towards the true values.

MCMC uses Markov chains to explore the parameter space and draws samples from the generated chain. A Markov chain is a series of random variables $X = \{X_1, X_2, \ldots, X_N\}$ for which the current state of the chain X_t depends only on the previous state X_{t-1}. The probability of moving from state X_{t-1} to state X_t is defined by a transition kernel $T(X_t|X_{t-1})$ [126]. Because the current state of the chain only depends upon the previous state, given a suitably large number of transitions, the chain 'forgets' its initial state and converges towards a stationary distribution $\pi(X)$ [70]. Once the chain reaches this distribution, all future states of the chain are drawn from $\pi(X)$ [126]. If the chain is irreducible, i.e. if it is able to reach any state in parameter space in a finite number of updates, the stationary distribution is unique [126]. Therefore, if we can construct an irreducible Markov chain

whose stationary distribution is equal to the target distribution, then the samples we draw from the Markov chain are also samples from the target distribution. This can be done by choosing an appropriate transition kernel. There are numerous choices we can make for T, however for the rest of this section we focus on only one choice of T which results in the Metropolis-Hastings algorithm [82], one of the most popular and most general MCMC algorithms available.

Starting with an initial parameter vector $\boldsymbol{\Theta}_t$, the Metropolis-Hastings algorithm converges towards the target distribution using the following steps:

1. Generate a new set of parameters $\boldsymbol{\Theta}'$ using a proposal distribution $q\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$. The proposal distribution, and thus the new state generated from it, depends only on the current state $\boldsymbol{\Theta}_t$.

2. Accept the new state with probability $\alpha\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$, where

$$\alpha\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right) = \min\left\{1, \frac{\Pr\left(\boldsymbol{\Theta}'|D\right) q\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)}{\Pr\left(\boldsymbol{\Theta}_t|D\right) q\left(\boldsymbol{\Theta}_t|\boldsymbol{\Theta}'\right)}\right\}. \tag{2.62}$$

   The probability of moving from $\boldsymbol{\Theta}_t$ to $\boldsymbol{\Theta}'$ is $q\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$ and the corresponding probability of moving from $\boldsymbol{\Theta}'$ to $\boldsymbol{\Theta}_t$ is $q\left(\boldsymbol{\Theta}_t|\boldsymbol{\Theta}'\right)$. If the move is accepted, we set $\boldsymbol{\Theta}_{t+1} = \boldsymbol{\Theta}'$; else $\boldsymbol{\Theta}_{t+1} = \boldsymbol{\Theta}_t$, i.e. the chain remains in the same state.

3. Repeat the previous two steps until the Markov chain has converged to the stationary distribution, which is also the posterior distribution $\Pr\left(\boldsymbol{\Theta}|D\right)$.

In this case, the total transition kernel T is a mixed distribution of the probability of proposing a move $q\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$ and accepting it with probability $\alpha\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$, and the probability of rejecting the move, which both depend on the form of α and q.

The acceptance probability α relies on the ratio of the posterior distribution at the proposed state $\boldsymbol{\Theta}'$ and the current state $\boldsymbol{\Theta}_t$. This eliminates the intractable normalising constant $\Pr\left(D\right)$ completely, making α easy to evaluate. This form of the acceptance probability also favours proposals for which $\Pr\left(\boldsymbol{\Theta}'|D\right)$ is higher than $\Pr\left(\boldsymbol{\Theta}_t|D\right)$, allowing the Markov chain to converge towards the shape of the posterior distribution.

The choice of proposal distribution $q\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$ is key, as although the Markov chain converges for any q [70], the rate of convergence can vary dramatically. We only want to draw samples from the chain once it has converged; therefore we ignore all the steps in the chain prior to convergence. This is known as the burn-in phase [66]. In order to keep this to a minimum, we would like the chain to converge as quickly as possible. One of the most common choices for q is a multivariate Gaussian function with mean $\boldsymbol{\Theta}_t$ and covariance matrix $\boldsymbol{\Sigma}$. In this case q is symmetric, meaning that $q\left(\boldsymbol{\Theta}_t|\boldsymbol{\Theta}'\right) = q\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right)$. This simplifies the acceptance probability to

$$\alpha\left(\boldsymbol{\Theta}'|\boldsymbol{\Theta}_t\right) = \min\left\{1, \frac{\Pr\left(\boldsymbol{\Theta}'|D\right)}{\Pr\left(\boldsymbol{\Theta}_t|D\right)}\right\}. \tag{2.63}$$

This is also known as the Metropolis algorithm [120] and is a special case of the more general Metropolis-Hastings algorithm. If the elements of $\boldsymbol{\Sigma}$ are small, the proposal distribution is narrow and all of the

update steps $\boldsymbol{\Theta}'$ are close to the current state $\boldsymbol{\Theta}_t$. The posterior distribution is generally very similar at both states resulting in high acceptance probabilities, however because the steps are so small it takes a very long time for the Markov chain to explore the full parameter space. On the other hand, if the proposal distribution is too wide, the Markov chain may miss large regions of parameter space for which the posterior probability is high. This generally leads to very low values of α and a 'stationary' Markov chain. In both of these cases, the mixing of the chain is said to be poor [10]. For the chain to mix well, and therefore converge quickly, it is important to consider the scale of the parameter space when choosing q. Alternatively we can use adaptive MCMC algorithms [154] which monitor the acceptance rate α during burn-in and adjust the covariance matrix of q accordingly. There is no universally accepted range for α, however several authors [66] have suggested that α should be between 20% and 50%. When using adaptive algorithms, it is generally a good idea to divide the vector $\boldsymbol{\Theta}_t$ into n blocks of parameters with similar sizes and scales. Instead of updating the whole of $\boldsymbol{\Theta}_t$ at once, only one of the n blocks is updated at a time using a Gaussian distribution with a covariance matrix $\boldsymbol{\Sigma}_n = \sigma_n \mathbf{I}$. The acceptance probability for this block can then be considered independently and its associated variance σ_n modified accordingly.

However, the following questions still remain: how many samples should be included in the burn-in phase before the chain has converged? and how many samples should we draw from the chain? There is no definitive answer to either question. In terms of convergence, most suggested methods rely on plotting out the evolution of the Markov chain (known as a trace plot) [66] or analysing the correlation between the earlier and later stages of the chain [69]. Other diagnostic tests involve running multiple chains and comparing the inter- and intra-variability between them [67]. However, it is important to remember that none of these methods can prove convergence [66]. Even if that were possible, because MCMC is a random (or at least a pseudorandom) process, the convergence of one chain would not guarantee the convergence of any other, independent chains. In terms of the number of samples, the more we draw the more accurate our estimates of the moments of $\Pr(\boldsymbol{\Theta}|D)$ are. In practice, the number of samples may be limited by time or storage constraints. Because not every proposed update is accepted, there is often correlation between adjacent states in the chain. In order to reduce the correlation between samples, it is standard practice to 'thin' the chain and only sample every i^{th} state [66, 70]. The number of states to skip between samples can be determined by running the chain and evaluating the correlation between adjacent states [70].

2.11 Summary

We began this chapter by describing the overall structure and function of white matter, noting the underlying spatial coherence and patterns that exist within the tracts. We discussed how this structure and function can be altered due to pathology and stressed the importance of detecting and localising the changes that occur in order to improve the diagnosis and prognosis of various white matter disorders. Next we discussed how MRI can be used to image brain structure in vivo and, more specifically, how diffusion-weighted contrast can be used to determine the microstructural properties of white matter tissue. In order to provide specific markers of tissue change we need to be able to relate the diffusion-

weighted MR signals we measure directly to the underlying tissue changes. To this end, numerous signal models have been developed that relate the signals to a wide range of tissue parameters such as axon radius, anisotropy and diffusivity. We provided a comprehensive review of a selection of these models. In general, we find that the models with the most specific tissue parameters (such as glial cell size or axon radius) are the most difficult models to fit and that parameter estimates obtained from them can be noisy and unreliable. However, we noted that this could be improved by using the underlying spatial coherence of the tract to guide the model fitting. We then discussed how these parameter estimates are used in group studies to detect population changes and introduced multiscale models of white matter structure, which incorporate both spatial and microstructure information. However, most of the techniques we discussed, whilst providing a detailed representation of tract shape, focused only on the simplest microstructure indices and did not consider using the local neighbourhood information to inform and improve microstructure estimates. This could be achieved by incorporating one of the spatial regularisation methods mentioned. Finally we discussed some mathematical techniques, B-splines and Bayesian model fitting, that we use in chapter 3 to develop our overall methodology.

As discussed in chapter 1, our overall aim is to develop multiscale models of microstructure variation across white matter regions that fully exploit the underlying local spatial coherence of tracts in order to maximise the accuracy and precision of our parameter estimates. As we have already mentioned throughout this chapter, this would be particularly beneficial when estimating parameters such as axon radius, which could provide very specific markers for measuring white matter tissue changes but are currently too difficult to estimate from clinical data reliably. However, we have also noted that this approach could also improve the estimation of simpler tissue parameters such as diffusivities from noisy, sparse data sets. As demonstrated in this chapter, the concepts of multiscale white matter models and spatial regularisation are well known; however they have not yet been combined in a way that allows us to estimate subtle trends of parameter variation directly from the rich set of diffusion MR signals that we measure. In the next chapter, we propose and describe in depth a novel framework for achieving this.

Chapter 3

The regional variation model framework

3.1 Motivation

In chapter 2, we discussed the role of white matter tissue within the brain and demonstrated the need for specific and sensitive white matter biomarkers that not only quantify the tissue changes due to pathology, but localise them as well. We discussed diffusion MRI and showed how the parameters of diffusion MR signal models can be used as indices of tissue integrity, allowing us to relate the measured signals to the underlying white matter microstructure. We also described several studies in which these indices have been combined with shape models to form multiscale models of white matter tracts. However, we noted several limitations of the approaches taken so far, with regards to both the diffusion MR signal models and the multiscale representations. First, the most informative diffusion MR signal models are often the most difficult to fit to data due to the weak dependence of the signal on the parameters of interest. This results in noisy parameter estimates, making these indices unreliable and therefore unfeasible as biomarkers, particularly when estimated voxel by voxel. Second, when representing microstructure variation within white matter regions, many current approaches estimate the tissue indices in each voxel independently before sampling the values along the tract length to produce a curve, typically described by an underlying mathematical function, for each parameter. By fitting the diffusion signal model in this way, they do not exploit the underlying spatial coherence in the data and therefore, the parameter estimates are more susceptible to noise. By estimating trends of variation for each parameter independently, these methods decouple the spatial dependence of the various microstructure parameters within the tract, meaning that whilst the individual parameter curves may capture the variability of the voxelwise parameter estimates well, the combined curves may not be representative of the measured signals. Both of these steps could be particularly problematic for sparse or noisy data, when even simple models of the diffusion MR signal are difficult to estimate robustly.

In this chapter we propose a novel framework, the regional variation model (RVM), which combines microstructure and shape information in order to provide a biomarker that both accurately estimates and localises white matter tissue indices. The RVM fits functions describing the regional variation of the tissue model indices directly to the diffusion-weighted MR signals, using shape models to map the trends within the natural geometry of the white matter tracts. The RVM exploits the underlying spatial coherence of the tissue microstructure, and thus the measured diffusion MR signals, thereby

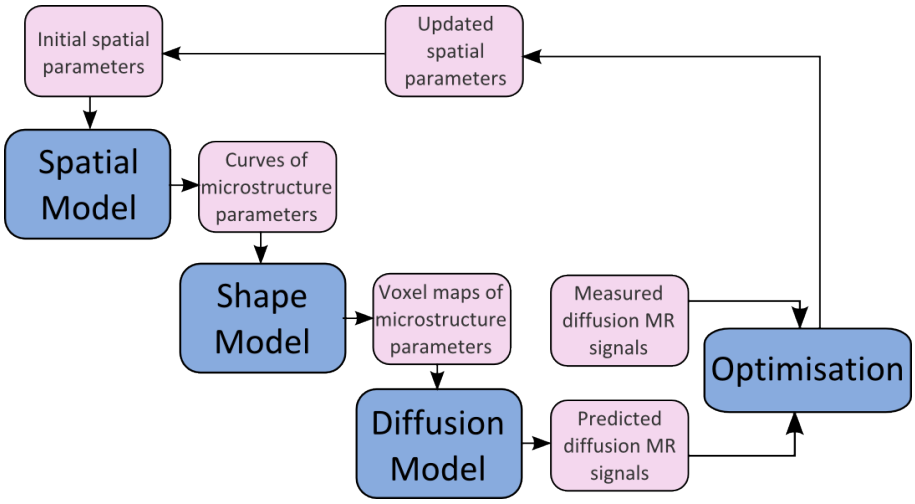

Figure 3.1: A flow diagram illustrating the multi-level forward model used to relate the spatial model parameters to the predicted diffusion MR signal.

pooling the data in a natural way. This increases the effective SNR of the data, reducing the effects of noise on our parameter estimates. This allows us to obtain more accurate and precise values of white matter microstructure indices, which should improve the estimation of weakly sensitive parameters, such as an axon radius index, as well as more common diffusion MRI parameters, such as diffusivities, particularly from noisy data. The RVM estimates the regional variation of all the tissue parameters under consideration simultaneously, which preserves any spatial dependencies that may exist between the parameters. This allows us to map the trends of parameter variation that best represent the measured signals. In the rest of this chapter we

- provide a general overview of the regional variation model,
- discuss the key components of the framework in depth, focusing on the specific implementations we use within this thesis,
- describe optimisation techniques for fitting the RVM to diffusion-weighted data,
- discuss possible extensions of the framework.

This chapter focuses solely on the theory and the implementation of the framework, however it provides the basis for the experimental work in chapters 4, 5 and 6, in which we validate the method and present several key applications of the RVM. The ideas and methods in this chapter have previously been presented and discussed in [123] and [124].

3.2 Overview

As stated above, the goal of the RVM framework is to estimate the trends of white matter microstructure variation directly from the diffusion-weighted MR signals we measure, using a shape model to map these trends within the natural geometry of the ROI. In this context, a ROI can refer to either a large

white matter tract or a small manually defined region of interest. Instead of fitting a model of the diffusion MR signal to each voxel independently, we fit the variation model to all voxels in the ROI simultaneously via a multi-level forward model, using functions to capture the underlying microstructure variation. These functions pool data from the local voxel neighbourhood, exploiting the spatial coherence of the underlying tissue microstructure and informing the parameter estimation. The framework, which is shown in figure 3.1, comprises three key components: the shape model, the spatial model and the diffusion model.

The spatial model comprises the set of functions which represents the regional variation of the microstructure parameters across the ROI. These curves are controlled by a small number of parameters, called the spatial parameters. The shape model provides a projection mapping between the continuous functions of the spatial model and the discrete voxels in the ROI, assigning a set of diffusion model parameters to each individual voxel, depending on its position within the ROI. The diffusion model is used to calculate the expected diffusion-weighted MR signals in each voxel of the ROI from the predicted diffusion model parameters. The optimal regional variation model is determined by finding the most likely set of spatial parameters, and therefore the most likely curves of parameter variation, given the data. This is illustrated in more depth in the schematic diagram shown in figure 3.2.

The framework itself is independent of the choice of shape model, spatial model, diffusion model and optimisation strategy. In the next sections, we provide more detail on the specific implementations used within this thesis.

3.3 Spatial model

The spatial model consists of a set of functions: one to represent the variation of each diffusion signal model parameter over the region of interest. As discussed in section 2.9, the variation of microstructure parameters within white matter tracts or ROIs is often complex and varies among subjects; therefore basis functions provide an ideal means for representing these trends. We can construct complex functions that approximate the data well using only a small number of suitable basis functions controlled by a corresponding set of weights. By using a basis function representation, the problem of estimating the trends of multiple tissue parameters simultaneously is reduced to estimating the most likely set of weights for the parameters. There is a wide range of basis systems that can be used for the spatial model, including polynomials, Fourier series and splines, several of which have been implemented within the RVM framework already (see section 3.7 for details). In this section, we focus on the particular representation we use within this thesis in order to introduce the mathematical concepts and notations that are used in the rest of this work.

We represent the regional variation of the diffusion model parameters using Bayesian penalised B-splines (Bayesian P-splines) [107] because, as mentioned in section 2.9.1, B-splines are particularly suited to capturing local variation within data. In this section we present the theory for one-dimensional Bayesian P-splines, which can be used to model microstructure variation along one-dimensional medial axes, although this approach can easily be extended to n-dimensional Bayesian P-splines as demonstrated in [107].

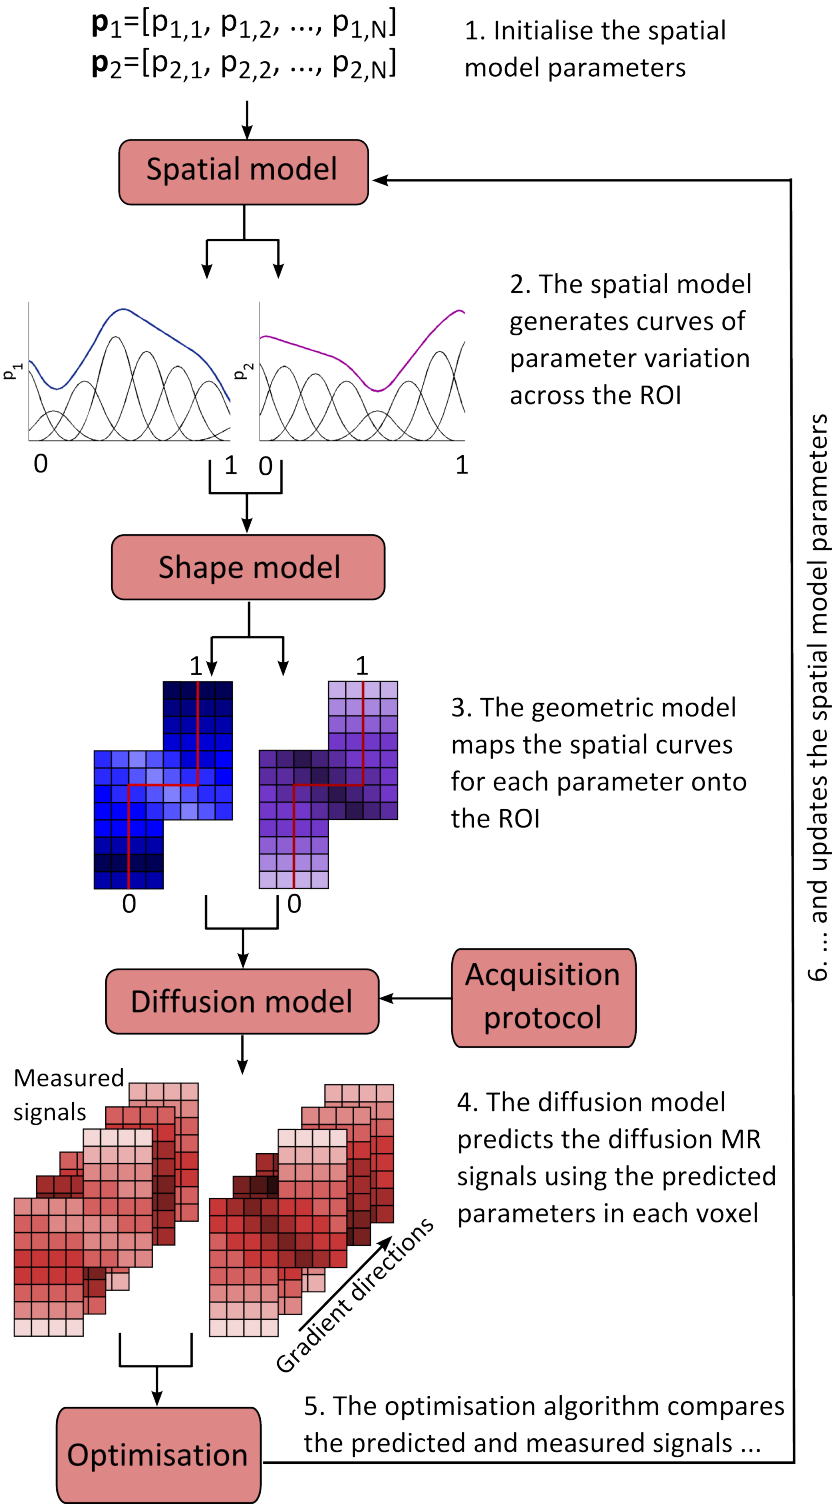

Figure 3.2: The spatial model comprises a set of curves describing the variation of diffusion model parameters across the tract axis. These curves predict local diffusion parameters in every voxel of the tract. The diffusion model predicts the MR signals, given the predicted diffusion parameters. The spatial model is controlled by a small set of parameters, and its goodness of fit can be calculated using an objective function based on predicted and measured MR signals. The spatial model is optimised by iteratively minimising the objective function.

Considering the regional variation of Q diffusion model parameters, we assume that the variation of each parameter $p \in \{p_1, p_2, \ldots, p_Q\}$ across the ROI can be represented as the weighted sum of N k^{th} order B-spline basis functions with corresponding weight vector $\mathbf{a}_p = [a_{1,p}, a_{2,p}, \ldots a_{N,p}]$. We denote the resulting curve for each parameter by y_p. If the position of the i^{th} voxel in the ROI or tract (determined from the shape model) is x_i, we can use the curves to predict the diffusion model parameters in this voxel by evaluating

$$y_p(x_i, \mathbf{a}_p) = \sum_{n=1}^{N} a_{n,p} B_{n,k}(x_i) \tag{3.1}$$

for each parameter p.

As mentioned in section 2.9.1, penalised splines use a large number of equally spaced knots to define the basis functions and add a penalty term to prevent overfitting to the noise. In the Bayesian P-spline approach, the penalty term is formulated as a prior probability on the basis function weights $a_{n,p}$ such that

$$\Pr\left(a_{n,p}|\Lambda_p\right) \sim \mathcal{N}\left(\frac{1}{2}\left(a_{n-1,p} + a_{n+1,p}\right), \Lambda_p\right), \tag{3.2}$$

where $\mathcal{N}$ is a Gaussian distribution with mean $\frac{1}{2}\left(a_{n-1,p} + a_{n+1,p}\right)$ and unknown precision Λ_p, where $\Lambda_p = \frac{1}{\nu_p^2}$ is the reciprocal of the variance ν_p^2. This means that a given value of $a_{n,p}$ is more probable if it is similar to the adjacent basis function weights, enforcing the smoothness of the curve. The precision Λ_p for each p is fit as part of the model. The prior probability distribution on Λ_p is a gamma distribution so that

$$\Pr\left(\Lambda_p\right) \sim \Gamma\left(\Lambda_p; \alpha, \beta\right), \tag{3.3}$$

where α and β are the shape and scale parameters respectively. A gamma distribution is chosen for $\Pr\left(\Lambda_p\right)$ as it is the conjugate distribution to $\Pr\left(a_{n,p}|\Lambda_p\right)$, a Gaussian distribution with a known mean but unknown precision [114]. The hyperparameters α and β determine the strength of the smoothness penalty and are set prior to fitting the RVM model to the data. We discuss possible choices for α and β further in chapter 4. The full prior distribution for $\mathbf{a}_p$ is the product of $\Pr\left(a_{n,p}|\Lambda_p\right)$ over all N basis functions and $\Pr\left(\Lambda_p\right)$ so that

$$\Pr\left(\mathbf{a}_p\right) = \Pr\left(\mathbf{a}_p|\Lambda_p\right)\Pr\left(\Lambda_p\right) = \left(\prod_{n=1}^{N} \Pr\left(a_{n,p}|\Lambda_p\right)\right)\Pr\left(\Lambda_p\right). \tag{3.4}$$

In total, we determine a vector of basis function weights $\mathbf{a}_p$ and a precision parameter Λ_p for each of the Q diffusion model parameters. We represent the total set of spatial model parameters by $\mathbf{a} = \{\mathbf{a}_{p_1}, \mathbf{a}_{p_2}, \ldots, \mathbf{a}_{p_Q}\}$ and precision parameters by $\mathbf{\Lambda} = \{\Lambda_{p_1}, \Lambda_{p_2}, \ldots, \Lambda_{p_Q}\}$. Therefore, the total prior distribution over all the spatial model parameters $\Pr\left(\mathbf{a}\right)$ is the product of the priors over the individual parameters defined in equation 3.4, so that

$$\Pr\left(\mathbf{a}\right) = \prod_{q=1}^{Q} \Pr\left(\mathbf{a}_{p_q}|\Lambda_{p_q}\right)\Pr\left(\Lambda_{p_q}\right) = \prod_{q=1}^{Q}\left(\left(\prod_{n=1}^{N} \Pr\left(a_{n,p_q}|\Lambda_{p_q}\right)\right)\Pr\Lambda_{p_q}\right). \tag{3.5}$$

3.4 Shape model

The shape model provides a mapping between the spatial model and the ROI or tract, given the underlying geometry of the tract itself, as well as enforcing correspondence between subjects. This mapping

allows us to evaluate the spatial model at each voxel, according to equation 3.1, and thus provide a prediction of the diffusion MR signals which we can then compare against the measured signals. It is important to choose an appropriate shape model for the white matter region under consideration as the regional variation of the tissue parameters are modelled along this geometric representation; the axes of the shape model must reflect the key changes in the underlying tissue variation. For example, the microstructure of the CC is known to vary along the medial axis of the midsagittal slice whilst remaining coherent perpendicular to it [2]. Therefore, the medial axis is a suitable representation for the CC. In section 2.7 we discussed several shape models, including streamline-based [68, 74] and medial representations [72, 195]. We choose to use the cm-rep approach as it is able to model a wide range of tract shapes in both two- [171] and three-dimensions [195]. In section 2.7 we briefly discussed how the cm-rep model is fit to a segmented tract via inverse skeletonisation; here we provide more detail on how the fitted cm-rep model is used within the RVM framework.

As shown in equation 3.1, in order to evaluate the spatial model in each voxel, we need to know the position x_i of the voxel within the ROI relative to our reference frame i.e. the shape model. When we fit the two-dimensional cm-rep shape model [171] to an ROI, it provides us with the medial axis represented as a list of 101 equally spaced points $x^m = \{x_l^m\}_{l=1}^{101}$. The first point x_1^m corresponds to the start of the medial axis whereas the final point x_{101}^m corresponds to the end of the medial axis. The cm-rep also provides corresponding boundary points and thickness values at each medial axis point. In order to determine the position x_i of the i^{th} voxel, we use a nearest neighbourhood approach to find the closest medial axis point $x_{l,i}^m$ to the i^{th} voxel. We then assign $x_i = x_{l,i}^m$. Values of x_i are normalised on the interval $[0, 1]$ in order to provide correspondence between multiple subjects [171], so that 0 corresponds to the start of the medial axis and 1 corresponds to the end of the medial axis for all subjects. This process is easily generalisable to two dimensions, in which we find the position $\mathbf{x}_i = (x_i, y_i)$ of the i^{th} voxel by projecting it onto the two-dimensional medial surface and finding the nearest surface point.

3.5 Diffusion model

Using the values of the diffusion model parameters obtained by evaluating the spatial model at each voxel, we use the diffusion model to predict the diffusion MR signals given the current spatial parameters $\mathbf{a}$. In this chapter and figures 3.1 and 3.2, we refer to this as the diffusion model as all our exemplar applications of the RVM in this thesis are for diffusion MRI data. However, as discussed in chapter 1, the RVM framework could also be used to model trends in parameters such as relaxivity or susceptibility, and therefore this component of the RVM could be more generally referred to as the signal model. Within this thesis we consider three different diffusion models. Using the naming conventions of [134], these are the BallStick model [26], the ZeppelinCylinder model [6] and the ZeppelinCylinderDot model [7]. These models have been described briefly in section 2.4 and we introduce full mathematical descriptions, including precise notation and parameters, of these models within the relevant experimental chapters. In this section we treat the diffusion model in more general terms in order to demonstrate how it fits into the RVM framework.

The diffusion MR signal in the i^{th} voxel depends on the spatial model parameters $\mathbf{a}$ and the position

of the voxel within the shape model x_i. It also depends on the acquisition parameters used, for example in the PGSE sequence the signal depends on the gradient strength along the j^{th} gradient direction $\mathbf{G}_j$, the length of the gradient pulses δ and the time between the pulses Δ. Therefore the normalised diffusion MR signal $\tilde{A}_{ij}$ in the i^{th} voxel and along the j^{th} gradient direction can be modelled as

$$\tilde{A}_{ij} = f\left(\mathbf{a}, x_i; \mathbf{G}_j, \delta, \Delta\right), \tag{3.6}$$

where f is a function that relates $\mathbf{a}$ to $\tilde{A}_{ij}$ for a given choice of diffusion model.

3.6 Optimisation

The purpose of the optimisation is to find the most likely set of spatial parameters $\mathbf{a}$ given the data, so that the diffusion MR signals $\tilde{A}_{ij}$ predicted from $\mathbf{a}$ best match the measured signals A_{ij} in all voxels and along all gradient directions. As discussed in section 2.10, there are numerous optimisation methods available, including maximum likelihood and Bayesian optimisation. There are also numerous algorithms available for each of these methods, such as the gradient descent, Gauss-Newton and Levenberg-Marquardt algorithms for maximum likelihood estimation. All of these optimisation techniques are compatible with the RVM and, as we mention in section 3.7, many of them are already implemented within the framework. In this work, we optimise the spatial parameters of the RVM using Bayesian model fitting and estimate the full posterior probability distribution of the model parameters $\Pr\left(\mathbf{a}, \mathbf{\Lambda}\right)$. We choose this over the ML or MAP approaches discussed in section 2.10 as it allows us to quantify the uncertainty of our parameter estimation by considering the moments and confidence intervals of the posterior distribution. This is important as it allows us to determine range of plausible parameters over which the model fits and whether we are sensitive to the parameters under consideration. We now discuss the exact details of our optimisation strategy in more depth.

As shown in section 2.10, we can use Bayes' theorem to express $\Pr\left(\mathbf{a}, \mathbf{\Lambda}\right)$ in terms of the likelihood $\Pr\left(A|\tilde{A}\left(\mathbf{a}\right)\right)$ of the true measurements A given the predictions $\tilde{A}\left(\mathbf{a}\right)$, and the prior probability distributions over the parameters $\Pr\left(\mathbf{a}\right)$. We use the Rician likelihood function [153] to evaluate $\Pr\left(A|\tilde{A}\left(\mathbf{a}\right)\right)$, as the noise on single channel MR magnitude images is known to be Rice distributed [77] (For multi channel MR data, the noise follows a more general non-central chi distribution [4], which can be accounted for in the RVM by substituting the appropriate likelihood function). Therefore we can write likelihood that the data A is described by the spatial parameters **a** as

$$\Pr\left(A|\tilde{A}\left(\mathbf{a}\right)\right) = \prod_{i=1}^{I}\prod_{j=1}^{J} \frac{\tilde{A}_{ij}\left(\mathbf{a}\right)}{\sigma^2} \exp\left(-\frac{\tilde{A}_{ij}\left(\mathbf{a}\right)^2 + A_{ij}^2}{2\sigma^2}\right) I_0\left(\frac{\tilde{A}_{ij}\left(\mathbf{a}\right)^2 A_{i,j}^2}{\sigma^2}\right), \tag{3.7}$$

where σ^2 is the variance of the image noise and I_0 is the zeroth order modified Bessel function of the first kind. In equation 3.7, we have taken the product of the likelihoods over all I voxels and J gradients to give the likelihood term for the whole ROI. The posterior probability may then be written as

$$\Pr\left(\mathbf{a}, \mathbf{\Lambda}\right) \propto \Pr\left(A_{ij}|\tilde{A}_{ij}\left(\mathbf{a}\right)\right) \Pr\left(\mathbf{a}\right), \tag{3.8}$$

where $\Pr\left(A|\tilde{A}\left(\mathbf{a}\right)\right)$ is given in equation 3.7 and $\Pr\left(\mathbf{a}\right)$ is defined in equation 3.5.

Due to the complex form of the Rician likelihood distribution and the non-linear relationship between the spatial model parameters and the measured signals, we cannot evaluate equation 3.8 analytically. Therefore we use MCMC [70] to sample from the posterior distribution on $\mathbf{a}$. Specifically we use an adaptive Metropolis-Hastings sampler with zero mean Gaussian proposal distributions. By using an adaptive algorithm, we are able to monitor the acceptance rate during the burn-in phase and adjust the variances of the proposal distribution to ensure good mixing of the chain. As discussed in section 2.10.1, the length of the burn-in phase and the sampling rate are dependent on the convergence of the chain to the stationary distribution. Therefore we determine the optimal burn-in and sampling rates depending on the specific RVM implementation we use. Further details are provided in appendix A.

3.7 Implementation

The RVM framework is currently implemented using MATLAB® (2011a, MathWorks, Natick, MA, USA). The structure of the code is extremely flexible, allowing the shape, spatial and diffusion models, as well as the optimisation algorithms to be easily interchanged. There is one main program that performs the fitting and calls the spatial and diffusion models and optimisation method. Therefore, it is very easy to alter the form of the RVM by simply changing the spatial, diffusion and optimisation functions that are called. Alternative spatial models that are currently available within the framework include piecewise linear and polynomial models [122]. The shape modelling is done as a pre-processing step and the mapping between the voxels in the ROI and the shape model is fed into the RVM as an input parameter. Therefore, it is possible to use any shape model in conjunction with the RVM framework. Although we only consider three diffusion MR signal models within this thesis, the model taxonomy in [134] has been implemented within the RVM framework, providing a wide range of potential diffusion MR signal models. It is also possible to modify the optimisation procedure to specify Gaussian rather than Rician noise, and to fit the RVM using ML or MAP estimation instead of the fully Bayesian optimisation.

3.8 Summary

In this chapter we have introduced the regional variation model, a novel method for estimating the trends of white matter microstructure variation across large regions of interest. The key improvement of the RVM over currently state of the art methods is that it estimates the trends directly from the diffusion-weighted MR signals. This allows the RVM to capture subtle spatial dependencies that may exist between the various tissue indices we estimate. It also exploits the underlying coherence of the white matter tracts, pooling the data in a natural way to increase the effective SNR of the data and improve estimates of weak but specific tissue parameters such as axon radius.

The RVM consists of three key components: the shape model, the spatial model and the signal model. We have discussed the role of all three components within the framework and provided a detailed explanation of the specific implementations we use within this thesis. We have also discussed methods for optimising the RVM and described alternative diffusion models, spatial models and optimisation algorithms.

In the following chapters we validate the RVM framework and demonstrate two key applications.

More specifically, in chapter 4, we demonstrate that the RVM is able to recover underlying trends in microstructure variation accurately and precisely from both synthetic and 'silver standard' ex vivo monkey brain data. We also investigate the effect of the spatial parameters settings, such as α and β, on our solutions. In chapter 5 we show how the RVM can be used within population studies to detect subtle group differences and compare its discriminative ability against other widely used group comparison methods. Finally, in chapter 6 we demonstrate that the RVM is able to improve estimates of axon radius, even when the data is acquired at low gradient strengths.

Chapter 4

Validation of the RVM

4.1 Motivation

In chapter 3, we introduced the general framework for the regional variation model. The purpose of this framework is to estimate the variation of white matter microstructure indices within ROIs or tracts, by fitting functions capturing the underlying variational trends directly to the diffusion-weighted data. As discussed in chapter 3, this exploits the underlying spatial coherence of the tissue parameters and allows us to pool the data in a natural way, thus increasing the effective SNR of the data and our fitted parameter maps. This also allows us to estimate weak but potentially invaluable parameters such as an index of axon radius with increased accuracy and precision.

The aim of this chapter is to validate the RVM framework, starting with the simple ball and stick model of diffusion, using both simulated and ex vivo diffusion MR data. The experiment with simulated data has two key aims. The first is to investigate the effect of the P-splines parameters described in section 3.3, such as the number of knots and smoothing priors, on the fitted RVMs. The second is to quantitatively compare the performance of the RVM against independent voxelwise parameter estimation, probably the most commonly used fitting technique, as well as trends extracted from these voxelwise estimates using splines or averaging, over a range of data acquisition parameters and noise levels. The aim of the second experiment, performed using high quality ex vivo monkey brain data, is to show that the conclusions drawn from the synthetic data experiments are applicable in real white matter brain tissue, an environment that is much more complex than the idealised, but computationally tractable, simulations.

4.2 Experiment 1: Validation using simulated data

The first experiment uses simulated data, synthesised using the simple ball and stick model [26], to validate the RVM framework. This experiment has two key aims. The first is to investigate the effect of the spatial model settings, i.e. the choice of the knot points t and hyperparameters α and β, on results obtained using the RVM framework. The second aim is to compare the performance of the RVM against the current state of the art methods for extracting trends of microstructure parameter variation.

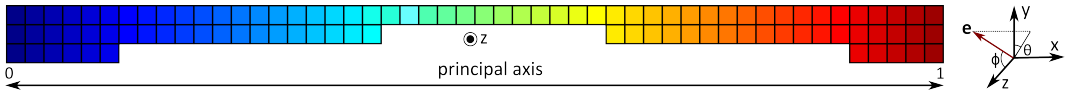

Figure 4.1: The synthetic ground truth phantom comprises one slice with 100 voxels. In this figure, the voxels are colour coded according to their distance along the principal axis, normalised on the interval $[0, 1]$.

4.2.1 Hypotheses

With regards to these specific aims, we hypothesise that

1. the RVM can accurately and precisely capture the underlying variation of microstructure parameters within white matter tracts. We hypothesise that results obtained using the RVM are robust to the choice of spatial model settings.

2. the RVM, which fits trends of microstructure variation directly to the diffusion-weighted MR signals, outperforms current state of the art methods that fit the trends to independent voxelwise parameter estimates. This is likely to be particularly pronounced when the data is sparse and/or noisy, resulting in large uncertainty in the independent voxelwise parameter estimates.

4.2.2 Methods

In this section we describe the diffusion model and ground truth phantom used in this experiment, and discuss the acquisition protocols and SNR values used to simulate noisy data. We give details of the model fitting procedures for the RVM as well as the alternative parameter estimation techniques we consider in this chapter: independent voxelwise estimation, mean trends and smooth trends estimated via B-spline fitting. Finally we discuss the statistical methods used to evaluate the spatial model settings for the RVM and to compare the various parameter estimation methods.

Diffusion model

We use the simple ball and stick model of diffusion described in section 2.4.2 and defined in equation 2.53. The model contains four independent parameters which are the intra-axonal volume fraction f, the diffusivity of water both inside and outside the axons d, and the fibre orientation defined by the angles θ and ϕ. The fibre orientation vector $\mathbf{e}$ is related to these angles via the equation

$$\mathbf{e} = \begin{bmatrix} \cos\theta \sin\phi & \sin\theta \sin\phi & \cos\phi \end{bmatrix}. \tag{4.1}$$

Therefore the spatial model within the RVM consists of four functions, one to describe the spatial variation of each parameter.

Synthetic phantom construction

The ground truth phantom is designed to mimic the mid-sagittal CC and comprises a single slice of data containing 100 voxels, arranged as shown in figure 4.1. The length of the phantom is 51 voxels; therefore when the ROI length is normalised on the interval $[0, 1]$ the spacing between adjacent voxels

is 0.02. Its width varies between one and three voxels, mimicking the variable thickness of tracts such as the CC and the corticospinal tract (CST). This also allows us to test how well the RVM performs in regions where there is less data available to pool. As we are validating the RVM using the ball and stick model, we assign values of the ball and stick model parameters to each voxel which represent the ground truth values. We do this by creating functions which represent the variation of f, d, θ and ϕ along the principal axis of the ROI. The range of values used in these functions is representative of typical white matter indices, although the shapes of the f and d functions are artificial, allowing us to test the sensitivity of the RVM to both subtle and sudden changes in microstructure variation. We set $\theta = \phi = 0$ everywhere, corresponding to fibre orientation vector $\mathbf{e} = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}$. A set of parameters $\{f_i, d_i, \theta_i, \phi_i\}$ is assigned to the i^{th} voxel based on these functions and the distance of the voxel x_i along the ROI axis. At some locations within the ROI there are multiple voxels with the same position x_i which would therefore be assigned exactly the same set of model parameters. In order to introduce a small amount of variability along the minor axis, instead of mapping the parameter variation functions directly onto the ROI, we assign parameters to each voxel by sampling from a Gaussian distribution with mean equal to the value of the parameter function evaluated at x_i and variance chosen so that 95% of values lie within 5% of the mean. The consistency of microstructure perpendicular to the medial axis is confirmed in histology studies [2]. We show scatter plots of the ground truth values for all parameters along with the corresponding parameter maps in figure 4.2.

Data synthesis

We synthesise the noisy diffusion MR signal in every voxel using the ball and stick model given in equation 2.53 (we set $S_0 = 1$ so that the signals are normalised), the ground truth parameters discussed above, four distinct acquisition protocols and five different SNRs. The acquisition protocols we consider range from the simplest DTI protocol with only 6 unique diffusion-encoding directions to high angular resolution diffusion imaging (HARDI) protocols with thirty or more diffusion-encoding directions. For all protocols, $b = 1000$ s mm^{-2}, which is the standard clinical value, and the gradient directions are chosen so that they are equally distributed over the unit sphere [97]. Details of all four protocols are given in table 4.1. The five SNRs we consider are 5, 10, 15, 20 and 25. Noise is added to the simulated signals A according to

$$A_{noisy} = \sqrt{(A + n_r)^2 + n_i^2} \tag{4.2}$$

where n_r and n_i are the real and imaginary components of the noise. $n_r, n_i \sim \mathcal{N}(0, \sigma)$, where σ corresponds to the required SNR in the unweighted signals. This results in Rician rather than Gaussian noise [153], and thus is more representative of real MR images. We generate $N_{noise} = 25$ different noise realisations from the same underlying signals for each combination of SNR and protocol in order to investigate the robustness of the RVM and other parameter estimation methods to noise.

Model fitting

We fit the ball and stick model to all of the simulated datasets (20 SNR/protocol combinations, each with 25 distinct noise realisations) and estimate trends in the model parameters using four different methods. The first method we consider is the RVM framework introduced in chapter 3, in order to fit the trends of

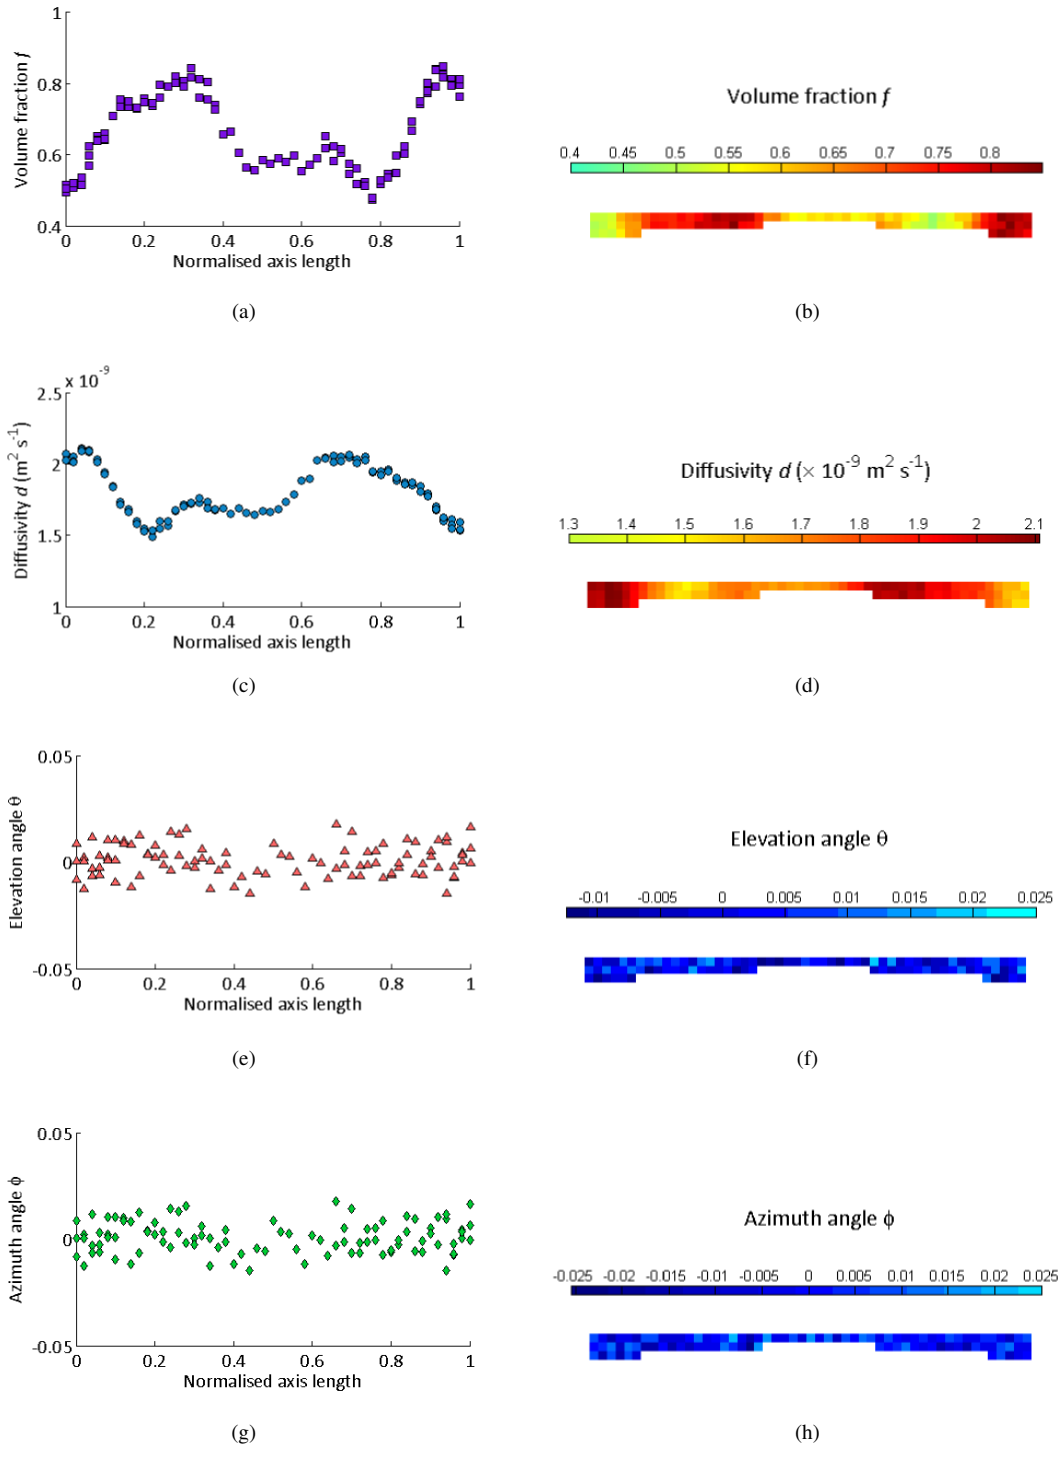

Figure 4.2: Scatter plots and parameter maps for the ground truth parameters. (a) shows a scatter plot of the variation of f along the ROI axis and (b) maps these values onto the ROI itself. (c) and (d) show the corresponding plots for diffusivity d; (e) and (f) for elevation angle θ; (g) and (h) for azimuth angle ϕ.

Protocol	No. of $b = 0$ images	No. of diffusion-encoding gradients	b value (s mm^{-2})
P1	1	6	1000
P2	1	15	1000
P3	3	30	1000
P4	6	60	1000

Table 4.1: Full details of the acquisition protocols used to synthesise data, showing the number of diffusion-encoding directions, the b value and the number of $b = 0$ images simulated. As stated in the text, these protocols range from very simple clinical DTI protocols (P1 and P2) to more research-oriented high angular resolution protocols (P3 and P4).

parameter variation across the ROI directly to the diffusion-weighted data. When fitting the model with the RVM, we estimate a small number of weights for each parameter, which determines the shape of a B-spline function describing the regional variation of that parameter. The second estimation method we use is independent voxelwise estimation, which fits the model to each voxel of data assuming that the signals measured in that voxel are independent of all other signals across the ROI. When using independent voxelwise fitting, we estimate a set of ball and stick model parameters, i.e. f, d, θ and ϕ, in every voxel. This form of parameter estimation is very commonly used in diffusion MRI studies [7, 60, 65]; therefore it is important to compare the RVM against it in order to determine the benefits of utilising the spatial coherence of parameters during fitting. We refer to this as voxel fitting. The final two methods we consider both extract trends of parameter variation from the independent voxelwise estimates. The first of these methods calculates mean trends of parameter variation across the medial axis by averaging the parameter estimates in voxels with the same spatial locations, and is similar to the method used to determine parameter trends by Sun et al [171]. This results in a list of averaged parameter values at every location along the medial axis, and we refer to these as *mean trends*. The second method extracts trends from the independent voxelwise parameter estimates by fitting splines to them, similar to the work of [71, 74]. In this case, as for the RVM fitting, we estimate a set of weights for each diffusion model parameter that determine the shape of a B-spline function that represents the regional variation of that parameter. We refer to these trends as *smooth trends*. Both the mean trends and smooth trends assume some degree of spatial coherence in the model parameters across the ROI, although this is introduced after model fitting. Therefore it is important to compare the RVM against these techniques in order to determine the benefits of exploiting the spatial coherence of the data during, rather than post, parameter estimation.

In the following paragraphs we provide a more detailed description of each model fitting procedure. This is particularly important for the RVM fitting, as we not only compare the RVM to the other trend fitting techniques; the aim of the first experiment is to investigate the dependence of the fitted RVM on the spatial model settings such as the number of splines and prior distributions. Therefore we carefully

introduce the exact notation we use during the experiments as well as the range of priors and knot vectors we consider.

RVM fitting When fitting the ball and stick model to the data, the RVM fits four functions describing the regional variation of the model parameters as follows:

- volume fraction f is modelled by N_f 4^{th} order (i.e. cubic) B-spline basis functions defined over $N_f + 4$ knots spanning the medial axis of the ROI. The basis function weights are represented by a $1 \times N_f$ vector $\mathbf{a}_f = \{a_{i,f}\}_{i=1}^{N_f}$ and the smoothness of the curve is represented by the precision Λ_f, both of which we optimise. Initially all N_f weights are set to 0.7, a physically plausible value for white matter [149], so that the curve is inititally constant everywhere within the ROI with a value of 0.7. The precision parameter is initialised so that $\Lambda_f = 10$, corresponding to weak smoothing. The hyperparameters of $\Pr(\Lambda_f)$ are α_f and β_f,

- diffusivity d is modelled as for f above. However, the basis function weight vector $\mathbf{a}_d = \{a_{i,d}\}_{i=1}^{N_d}$ is initialised to $1.5 \times 10^{-9}\mathrm{m}^2\mathrm{s}^{-1}$, resulting in a constant diffusivity of $1.5 \times 10^{-9}\mathrm{m}^2\mathrm{s}^{-1}$ in all voxels within the ROI. The precision parameter Λ_d, which determines the smoothness of the diffusivity trend, is also initialised as $\Lambda_d = 10$. The hyperparameters of $\Pr(\Lambda_d)$ are α_d and β_d,

- the fibre orientation parameters θ and ϕ are modelled using one 1^{st} order B-spline basis function each, i.e. both angles are assumed to be constant across the ROI. These basis functions are defined over the knot vector $\begin{bmatrix} 0 & 1 \end{bmatrix}$, i.e. one knot at each end of the ROI axis. The weight assigned to the basis function is controlled by one scalar value: a_θ for θ and a_ϕ for ϕ. a_θ and a_ϕ are initialised from the average principal DT direction across all voxels within the ROI.

We use 4^{th} order, or cubic, B-spline basis functions as they are continuous and are non-zero over only 4 knot positions, which means they are defined only very locally. Cubic B-splines are also the most commonly used order of B-spline functions within the literature [41, 74, 71, 157]. The posterior distribution over the RVM model parameters is estimated using MCMC as discussed in sections 2.10.1 and 3.6. We discard the first 5000 iterations during the burn-in phase, after which we draw 500 samples at intervals of 20 iterations. The methods used to determine these values are described in appendix A. To briefly summarise the procedure, we generate multiple Markov chains using different initialisations and determine how many iterations are required for convergence. The variance of the Rician likelihood function, which is required to evaluate the posterior distribution $\Pr(\mathbf{a}, \mathbf{\Lambda})$ defined in equation 3.8, is fixed at the value σ^2 used to synthesise the data. We can then characterise the posterior distribution over $\mathbf{a}$ using its mean and standard deviation, which provides us with useful information about the plausible range of the parameters.

Using the RVM we optimise the following parameters: $\mathbf{a}_f$, Λ_f, $\mathbf{a}_d$, Λ_d, a_θ and a_ϕ. However, the RVM also has several settings, i.e. the number of knots N_f and N_d, and the prior distribution hyperparameters α_f, β_f, α_d and β_d that need to be chosen a priori. The first aim of this experiment is to investigate how the choice of these settings affects the overall fit of the RVM and to identify a range of

appropriate settings. We now discuss the range of knot vectors and hyperparameters that we use in this part of the experiment.

We consider three different knot vectors t with different N, corresponding to a small, medium and large number of knots, in order to determine the effect of t on the fitted parameter variation curves. These are

- $t_1 = \{-0.2 : 0.1 : 1.2\}$. This knot vector comprises 15 knots, with separation 0.1. When using 4^{th} order cubic basis functions, this knot vector supports $N = 11$ basis functions.
- $t_2 = \{-0.1 : 0.05 : 1.1\}$. This knot vector comprises 25 knots, with separation 0.05. When using 4^{th} order cubic basis functions, this knot vector supports $N = 21$ basis functions.
- $t_3 = \{-0.04 : 0.02 : 1.04\}$. This knot vector comprises 55 knots, with separation 0.02. When using 4^{th} order cubic basis functions, this knot vector supports $N = 51$ basis functions.

In the notation $t = \{t_{beg} : \Delta t : t_{end}\}$, t_{beg} is the position of the first knot point, t_{end} is the position of the final knot point and Δt is the interval between adjacent knot points. Although it is possible to construct RVMs in which the number of basis functions is different for each parameter being modelled, we only consider spatial models in which $N_f = N_d$ in order to limit the number of RVMs we fit to each data set. Within the Bayesian P-spline literature, little attention has been paid to the effect of the number of basis functions on the fitted curves, and the general criteria is for N to be large, but still much smaller than the number of data points [107]. We have chosen these three knot vectors based on an initial experiment, presented in appendix B, which shows that the performance of the RVM remains reasonably consistent over limited ranges of knot vectors.

We consider different combinations of the hyperparameters, $\alpha_f, \beta_f, \alpha_d, \beta_d$, that control the smoothness of the parameter curves. Specifically we consider four distinct gamma distributions for $\Pr(\Lambda_p)$ which correspond to uninformative, light, medium and heavy smoothing. The gamma distributions are determined by a pair of hyperparameters (α_p, β_p). These are

- $\mathrm{S}_1 = \Gamma(\Lambda_p|\alpha_1, \beta_1)$, where $\alpha_1 \to 0$, $\beta_1 \to \infty$. This corresponds to an uninformative gamma distribution [114].
- $\mathrm{S}_2 = \Gamma(\Lambda_p|\alpha_2, \beta_2)$, where $\alpha_2 = 4$, $\beta_2 = 2.5$. This corresponds to a gamma distribution with mean $\mu_\Gamma = \alpha\beta = 10$ and variance $\sigma_\Gamma = \alpha\beta^2 = 25$. This favours small values of Λ_p resulting in weak smoothing.
- $\mathrm{S}_3 = \Gamma(\Lambda_p|\alpha_3, \beta_3)$, where $\alpha_3 = 4$, $\beta_3 = 25$. This corresponds to a gamma distribution with mean $\mu_\Gamma = \alpha\beta = 100$ and variance $\sigma_\Gamma = \alpha\beta^2 = 2500$. This favours medium values of Λ_p resulting in moderate smoothing.
- $\mathrm{S}_4 = \Gamma(\Lambda_p|\alpha_4, \beta_4)$, where $\alpha_4 = 4$, $\beta_4 = 250$. This corresponds to a gamma distribution with mean $\mu_\Gamma = \alpha\beta = 1000$ and variance $\sigma_\Gamma = \alpha\beta^2 = 250000$. This favours large values of Λ_p resulting in heavy smoothing.

$\Pr(\Lambda_f)$ \ $\Pr(\Lambda_d)$	S_1	S_2	S_3	S_4
S_1	✓			
S_2		✓	✓	✓
S_3		✓	✓	✓
S_4		✓	✓	✓

Table 4.2: Table showing the different combinations of prior distribution parameters on Λ_f and Λ_d, which control the smoothness of the f and d curves respectively.

We combine these prior distributions to form 10 different combinations of $\Pr(\Lambda_f)$ and $\Pr(\Lambda_d)$: one uses the uniformative prior S_1 for both Λ_f and Λ_d; the other 9 are combinations of the stronger smoothing priors S_2, S_3 and S_4. Table 4.2 summarises the different combinations of smoothing priors we use. Unlike the knot vectors, we consider different priors for Λ_f and Λ_d. There are numerous way in which we could choose the parameters for the smoothing prior α and β. Several approaches [17, 107] advocate using uninformative priors on the smoothness parameters so that $\Pr(\mathbf{\Lambda})$ does not heavily bias the shape of the posterior distribution. However a more recent study [101] has shown that the choice of prior distribution can influence the fitted P-splines, causing both under- and oversmoothing when the SNR is low and the data points are sparse. We restrict our investigation to a small selection of priors; however there is enough overlap between the different prior distributions to ensure that a wide range of smoothing parameters has a non-zero prior probability, greater than we expect to observe in the data.

In total, considering both the knot vectors and prior distributions, there are 30 different combinations of the spatial model settings $\Upsilon = \{t, \Pr(\Lambda_f), \Pr(\Lambda_d)\}$ for the RVM that we fit to each synthesised data set.

Independent voxelwise fitting The ball and stick model is fit to each voxel of data independently using Camino [44], an open source diffusion MRI toolkit. In order to allow for direct comparisons between the voxelwise (voxel) and RVM results, we assume Rician noise and use an MCMC optimisation algorithm when fitting the model to each voxel. The burn-in and sampling parameters match those used when fitting the RVM.

Mean trend estimation We estimate mean trends of f and d across the ROI by calculating the mean value of the microstructure parameters at each position along the medial axis, i.e. in this case at the 51 unique voxel locations across the principal axis shown in figure 4.1. We also calculate one average value for θ and ϕ over all voxels.

Smooth trend estimation using splines We estimate smooth trends of f and d by fitting splines to the independent voxelwise estimates of the ball and stick model parameters. Specifically, we fit Bayesian P-splines to describe the variation of f, d, θ and ϕ over the ROI, similar to the method used in [71]. When fitting the Bayesian P-splines to the voxelwise parameter estimates, we use the same knot vector, smoothing priors and MCMC parameters as for the RVM, as this allows us to directly compare how the

methods perform under identical settings. However, because we are now fitting splines to the diffusion model parameters themselves rather than the MR signals, the form of the posterior distribution is altered. We now estimate the posterior distribution $\Pr\left(\mathbf{a}'_p, \Lambda'_p\right)$ over the spatial parameters $\mathbf{a}'_p$ for each parameter p individually (when fitting Bayesian P-splines to the voxelwise parameter estimates we denote the parameters that we optimise as $\mathbf{a}'_p$ and Λ'_p so as to avoid confusion with the RVM spatial parameters $\mathbf{a}_p$ and Λ_p). The posterior distribution is given by

$$\Pr\left(\mathbf{a}'_p, \Lambda'_p\right) \propto \Pr\left(p'|\mathbf{a}'_p\right) \Pr\left(\mathbf{a}'_p\right) \tag{4.3}$$

where $\Pr\left(p'|\mathbf{a}'_p\right)$ is the likelihood of observing the fitted voxelwise parameter estimates $p' = \{p'_i\}_{i=1}^{I}$ given $\mathbf{a}'_p$ and $\Pr\left(\mathbf{a}'_p\right)$ is the prior probability of $\mathbf{a}'_p$. Instead of the Rician distribution used as the noise model for the MR signals in the RVM, we use a Gaussian distribution to represent the likelihood $\Pr\left(p'|\mathbf{a}'_p\right)$ so that

$$\Pr\left(p'|\mathbf{a}'_p\right) = \prod_{i=1}^{I} \frac{1}{\sqrt{2\pi\sigma_p^2}} \exp\left(\frac{-\left[y'_p\left(x_i, \mathbf{a}'_p\right) - p'_i\right]^2}{2\sigma_p^2}\right), \tag{4.4}$$

where we have taken the product of the likelihoods over all I voxels, and the curves $y'_p\left(x_i, \mathbf{a}'_p\right)$ are defined as in equation 3.1. We use the Gaussian noise model for p as we used a Gaussian distribution in order to generate the underlying ground truth trends described in the section on phantom construction. The variance of the likelihood function σ_p^2 is estimated from the variance of the parameter estimates p'_i over all voxels, and so reflects the natural variation within the parameters. The prior distribution over the spatial parameters $\Pr\left(\mathbf{a}'_p\right)$ is identical to the RVM, as in equation 3.4. Once again, we estimate the posterior distribution over the spatial model parameters $\mathbf{a}'_p$, which we then characterise using the mean of the posterior.

Statistical analysis

In the following paragraphs, we describe the statistical techniques used to determine the optimal number of knots and prior distributions for RVM fitting, and to compare how well the different parameter estimation methods recover the ground truth parameter trends.

Optimising the RVM settings We determine how well each combination of RVM settings Υ (i.e. each combination of N, $\Pr\left(\Lambda_f|\alpha_f, \beta_f\right)$ and $\Pr\left(\Lambda_d|\alpha_d, \beta_d\right)$) captures the underlying ground truth variation by calculating the squared differences between the fitted RVM functions and the ground truth values of f and d for each voxel. We sum this over all voxels to get the total squared difference over the ROI, and average this over the N_{noise} different noisy realisations of the data set, for each of the 20 data sets generated using different SNRs and protocols. This results in average mean squared errors for f (AMSE$_f$) and d (AMSE$_d$) for each combination of settings Υ. We then look for Υ that minimises both AMSE$_f$ and AMSE$_d$. The AMSE for f and d, given Υ, are calculated as

$$\mathrm{AMSE}_{\mathrm{f}}^{\Upsilon} = \frac{1}{\mathrm{N_{noise}}} \sum_{\mathrm{n}=1}^{\mathrm{N_{noise}}} \left(\frac{1}{\mathrm{I}} \sum_{\mathrm{i}=1}^{\mathrm{I}} \left[\mathrm{f_i} - \mathrm{y_f}\left(\mathrm{x_i}, \mathbf{a}_{\mathrm{f}}^{\Upsilon}\right)\right]^2\right) \tag{4.5}$$

and

$$\mathrm{AMSE}_{\mathrm{d}}^{\Upsilon} = \frac{1}{\mathrm{N_{noise}}} \sum_{\mathrm{n}=1}^{\mathrm{N_{noise}}} \left(\frac{1}{\mathrm{I}} \sum_{\mathrm{i}=1}^{\mathrm{I}} \left[\mathrm{d_i} - \mathrm{y_d}\left(\mathrm{x_i}, \mathbf{a}_{\mathrm{f}}^{\Upsilon}\right)\right]^2\right), \tag{4.6}$$

where $\mathbf{a}_f^\Upsilon$ and $\mathbf{a}_d^\Upsilon$ are the best fit spatial parameters for f and d for the given combination Υ, and f_i and d_i and the ground truth parameters. This method follows the approach of Ruppert et al [157], who use the squared differences between fitted splines and a known ground truth in order to find the minimum number of basis functions required to describe the data.

Comparison of parameter estimation methods We compare how well the different methods for estimating regional microstructural trends (RVM, voxel, mean, smooth) capture the underlying variation of the ground truth parameters by comparing the bias η and variance ζ along the medial axis length for each method, calculated for each model parameter over all N_{noise} different noise realisations for every protocol/SNR combination. For volume fraction f the bias η_i^f for the i^{th} voxel is

$$\eta_i^f = \frac{1}{N_{noise}} \sum_{n=1}^{N_{noise}} \left(\frac{\hat{f}_{i,n} - f_i}{f_i} \right), \tag{4.7}$$

where $\hat{f}_{i,n}$ is the estimated volume fraction in the i^{th} voxel from the n^{th} data set (for any estimation technique) and f_i is the ground truth value. The corresponding variance ζ_i^f is given by

$$\zeta_i^f = \frac{1}{N_{noise}} \sum_{n=1}^{N_{noise}} \left(\frac{\hat{f}_{i,n} - f_i}{f_i} \right)^2. \tag{4.8}$$

The expressions for bias and variance in d are defined in a similar way. For the fibre orientation parameters, we consider the bias and variance of the estimated direction vectors $\hat{\mathbf{e}}_i$ in the i^{th} voxel with respect to the ground truth orientation $\mathbf{e}_i$ rather than the bias and variance of θ and ϕ directly. This is because the periodicity of the angles can introduce large errors into the calculations. Therefore we define the bias $\eta_i^{\mathbf{e}}$ as

$$\eta_i^{\mathbf{e}} = \frac{1}{N_{noise}} \sum_{n=1}^{N_{noise}} (1 - |\hat{\mathbf{e}}_{i,n} \cdot \mathbf{e}_i|) \tag{4.9}$$

and the variance $\zeta_i^{\mathbf{e}}$ as

$$\zeta_i^{\mathbf{e}} = \frac{1}{N_{noise}} \sum_{n=1}^{N_{noise}} (1 - |\hat{\mathbf{e}}_{i,n} \cdot \mathbf{e}_i|)^2. \tag{4.10}$$

When $\hat{\mathbf{e}}_{i,n}$ is close to $\mathbf{e}_i$, $|\hat{\mathbf{e}}_{i,n} \cdot \mathbf{e}_i|$ is almost to 1, and therefore the bias is close to zero. We use the absolute value of the dot product between the estimated and ground truth vectors as, due to the symmetry of the diffusion process, we cannot differentiate between parallel and antiparallel fibre orientations.

4.2.3 Results

RVM spatial model settings

In this section we present results to test our first hypothesis, i.e. to test the effect of the spatial model settings Υ on the fitted RVM curves. Based on these results, we also aim to identify a working set of knot vectors and prior distributions that work well over all the SNRs and acquisition protocols tested during the experiment, for use in future experiments.

Figures 4.3 and 4.4 show example plots of AMSE for both f and d for acquisition protocols P4 and P2 respectively. In each figure we show a subset of results covering both high (SNR$=5$) and low (SNR=20) levels of noise, although the trends we observe are similar for all protocols and noise levels.

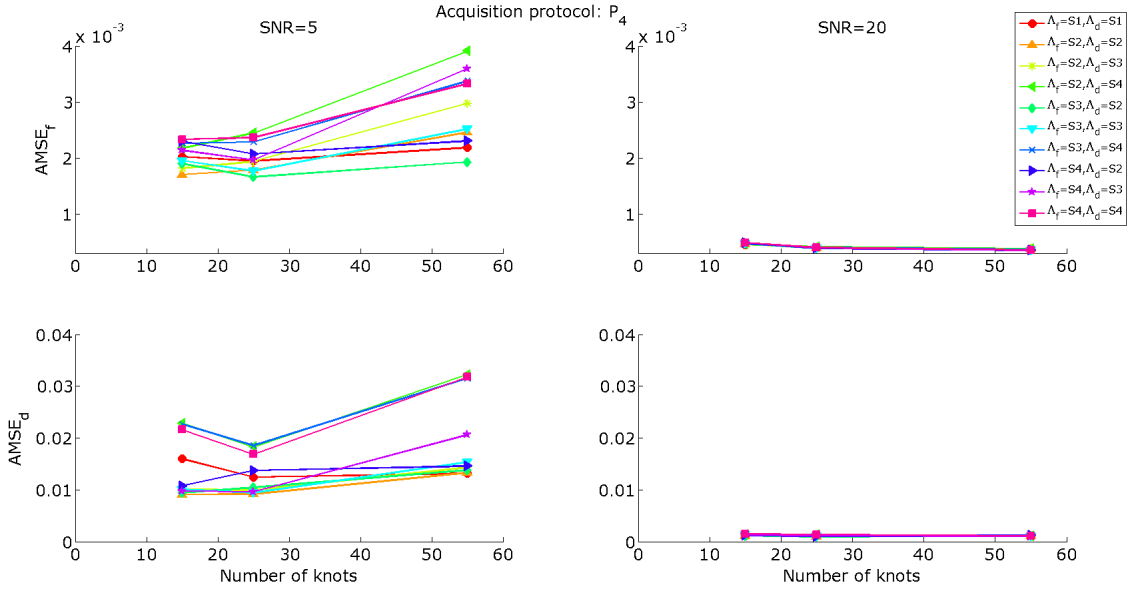

Figure 4.3: Plots of AMSE$_f$ and AMSE$_d$ for data simulated using protocol P$_4$ and SNR=5 and 20. The values of AMSE are shown for both parameters for all 30 combinations of knot vectors and smoothing prior distributions.

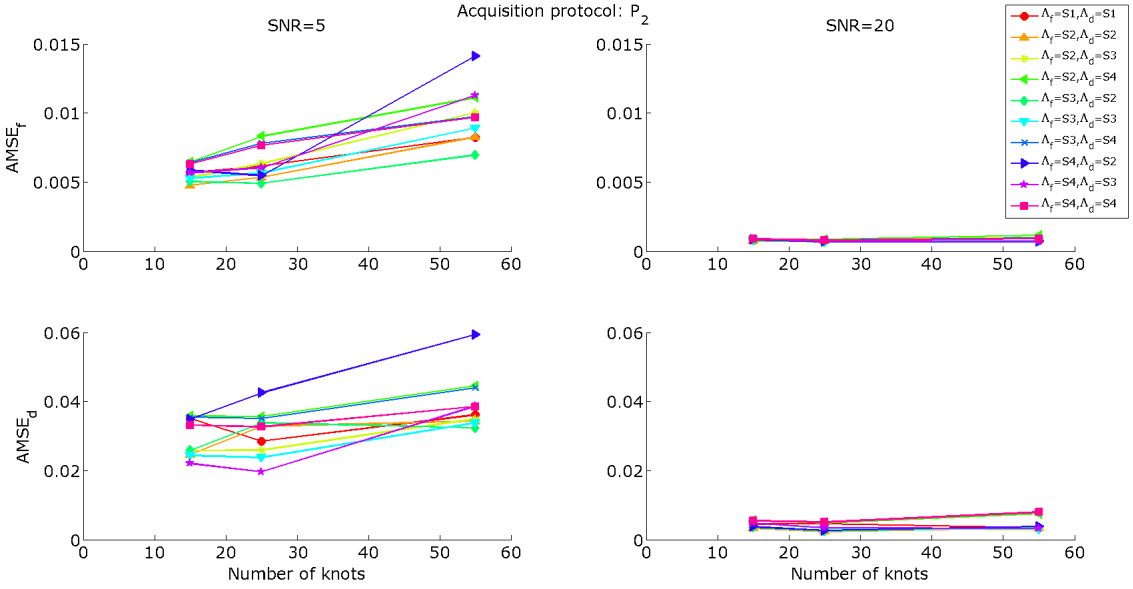

Figure 4.4: As for figure 4.4, but for data simulated using protocol P$_2$ and SNR=5 and 20.

The number of knots is displayed along the x-axis whilst the various combinations of prior distributions are represented using different colours.

In figure 4.3 we see that for a large number of gradient directions, when the SNR is high (SNR=20) the values of AMSE for all combinations of knot vectors and smoothing priors overlap significantly for both f and d. This shows that all combinations perform well, implying that in this case the RVM is robust to the choice of spatial model settings. As the noise increases (SNR=5), the plots of AMSE start to diverge, particularly for d, showing that for noisy data the choice of knots and smoothing does affect the estimated curves of parameter variation. We see the same trends for acquisition protocol P2 in figure 4.4, although in this case we note that the spatial model settings affect the estimated spline curves for d even when the SNR is high, although the changes in AMSE are very small.

This partially contradicts our first hypothesis, which stated the RVM is robust to the choice of spatial model settings. Whilst this is true when the SNR is high (i.e. 20 or above) and the number of gradient directions is large (i.e. 30 or above), in many experimental settings the quality of the data does not meet these criteria and the spatial model settings must be tuned in order to obtain the optimal results. As can be seen from figures 4.3 and 4.4, finding the combination of spatial model settings Υ that simultaneously minimises both AMSE_d and AMSE_f over a wide range of protocols and SNRs is challenging. It is also important to note that the AMSE does not tell the whole story, as it only provides one number to represent the whole curve, and thus does not provide information about how well the fitted RVMs reflect the trends in f and d. Therefore before choosing the knot vector and prior distributions it is important to inspect the curves to ensure that they accurately reflect the ground truth values. For example, in many cases it appears that the uninformative prior ($\Pr(\Lambda_f) = S_1, \Pr(\Lambda_d) = S_1$) performs reasonably well. However, when we plot the fitted curves, they are almost flat for d and do not capture the underlying changes in the ground truth values as the smoothing parameter becomes too strong. Therefore, even though the AMSE is reasonably low, the uniformative prior is not suitable when fitting the RVM to data with low SNR. This is illustrated in figure 4.5. In this figure, we also show that when the smoothing parameter is too weak (i.e. $\Pr(\Lambda_d) = S_2$, favouring weak smoothing), the functions for d estimated using the RVM fit to the noise, rather than the underlying ground truth trends.

Based on the values of AMSE observed for all simulations and the underlying shapes of the fitted curves, we try to provide some general guidelines for choosing the knot vector and smoothing priors. For all protocols and noise levels, we see a distinct pattern in AMSE for both f and d as a function of number of knots. The AMSE generally decreases as the number of knots increases from 15 to 25, then increases again as the number of knots rises to 55. This implies that the medium number of knots, defined on the knot vector t_2 best captures the underlying variation in the model parameters. We also note that AMSE_d is generally higher than AMSE_f, both in terms of magnitude and percentage error, as the signal is less sensitive to changes in d than in f. Therefore when considering the combination of smoothing priors to use, we focus on the settings that minimise AMSE_d. In general, we see that for all protocols and SNRs, AMSE_d is lowest when a medium level of smoothing is used, i.e. $\Pr(\Lambda_d) = \text{S}_3$. For this subset of smoothing prior combinations the choice of $\Pr(\Lambda_f)$ does not significantly affect the

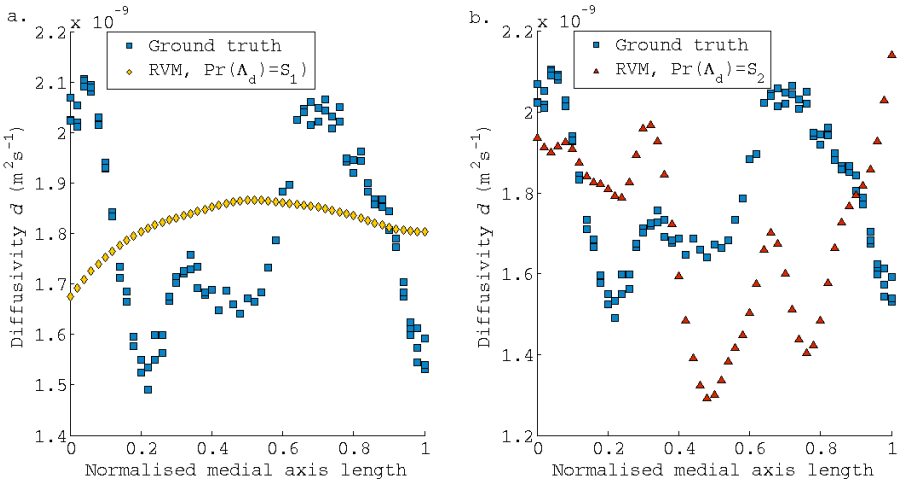

Figure 4.5: Figure a. shows the trend in d estimated using the RVM for one of the noisy realisations. In this case, the knot vector is t_2 and the prior distribution on the smoothness of d is uninformative, i.e. $\Pr(\Lambda_d) = S_1$. In this case, the smoothing parameter is allowed to become too large and the trend is oversmoothed. In b., for the same noisy data set and knot vector but with $\Pr(\Lambda_d) = S_2$, i.e. enforced weak smoothing, the smoothing parameter is too small and the estimated trend fits to noisy outliers.

value of AMSE$_f$, implying that the fitted RVM curves for volume fraction are still reasonably robust. Based on these observations, for the remaining experiments within this chapter, we fix the spatial model settings so that $t = t_2$, $\Pr(\Lambda_f) = \mathrm{S}_4$ and $\Pr(\Lambda_d) = \mathrm{S}_3$.

Comparison of parameter estimation techniques

Figures 4.6 and 4.7 show examples of bias and variance plots for f, d and $\mathbf{e}$ using the four different estimation methods described in section 4.2.2. Figure 4.6 shows the results for protocol P3 and SNR=5, whereas figure 4.7 shows the results for protocol P2 and SNR=20. Ideally the bias should be zero, or at least equally distributed about zero, as this shows that over the 25 noise realisations, the parameter estimates are neither overestimated or underestimated systematically relative to the ground truth data. Similarly, the variance should be zero, or close to zero, as this shows that over the noise trials, the parameter estimation methods are able to consistently estimate the same value, showing robustness to noise.

In figure 4.6, we see that there is a slight negative bias in f, which is consistent for all parameter estimation techniques. For the RVM and smooth trends, this is particularly pronounced at distances of 0.2, 0.6 and 0.9 along the axis. At these positions, there are sharp changes in the ground truth curves, implying that at low SNRs (i.e. SNR< 10), these methods may struggle to pick up sudden changes. The variance in f is large when we estimate the microstructure parameters independently in each voxel. When we estimate the trend in f naively by calculating the mean trend the variance is reduced slightly, provided that there are multiple voxels at each location to average over. However, if we consider the RVM or the smooth trend, the variance in f is almost zero, showing that these techniques are robust to noise when estimating volume fraction. For d, there is a significant positive bias in the independent

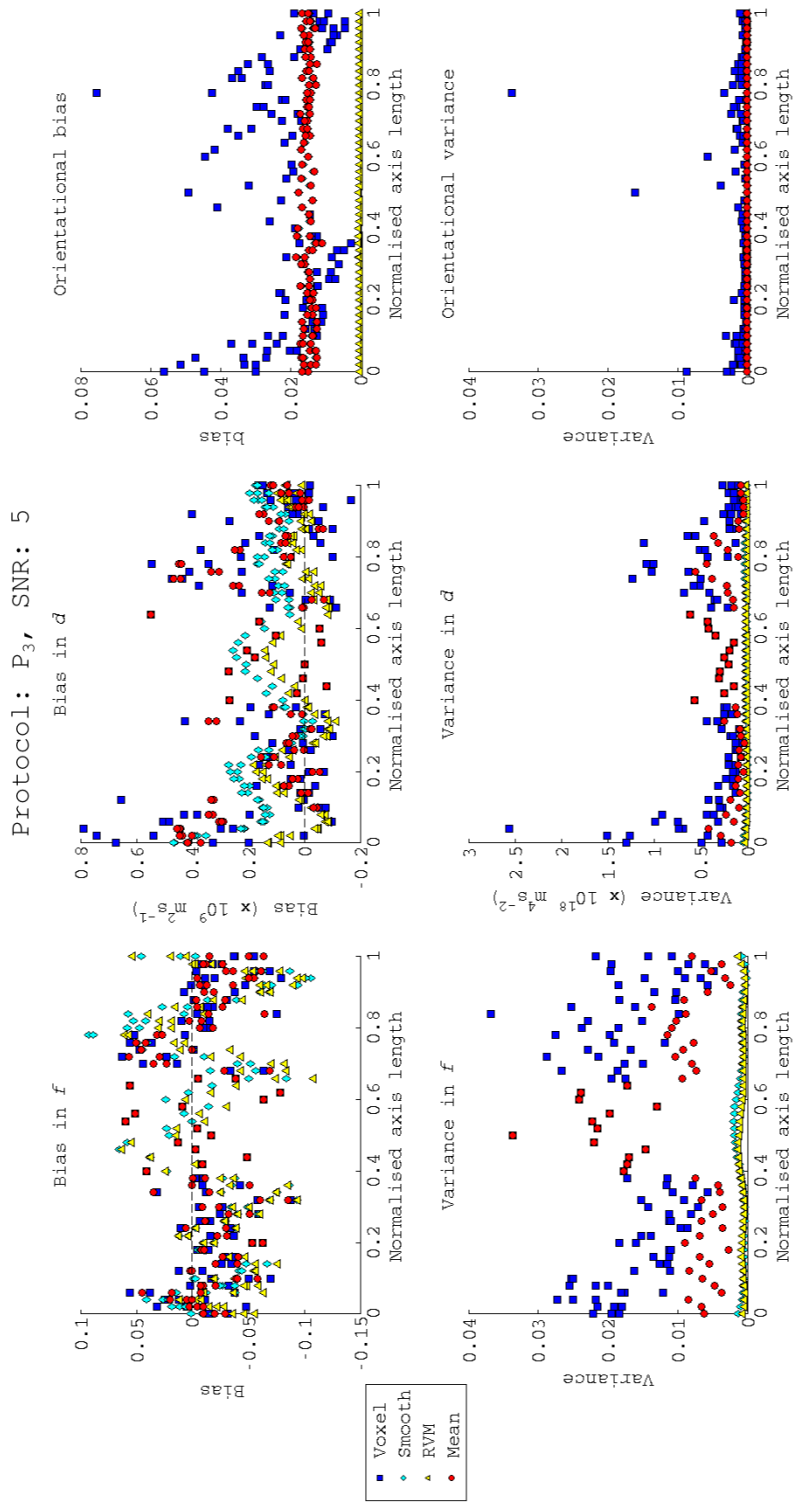

Figure 4.6: Plots of bias and variance for all parameter estimation techniques for data simulated using protocol P3 and SNR=5.

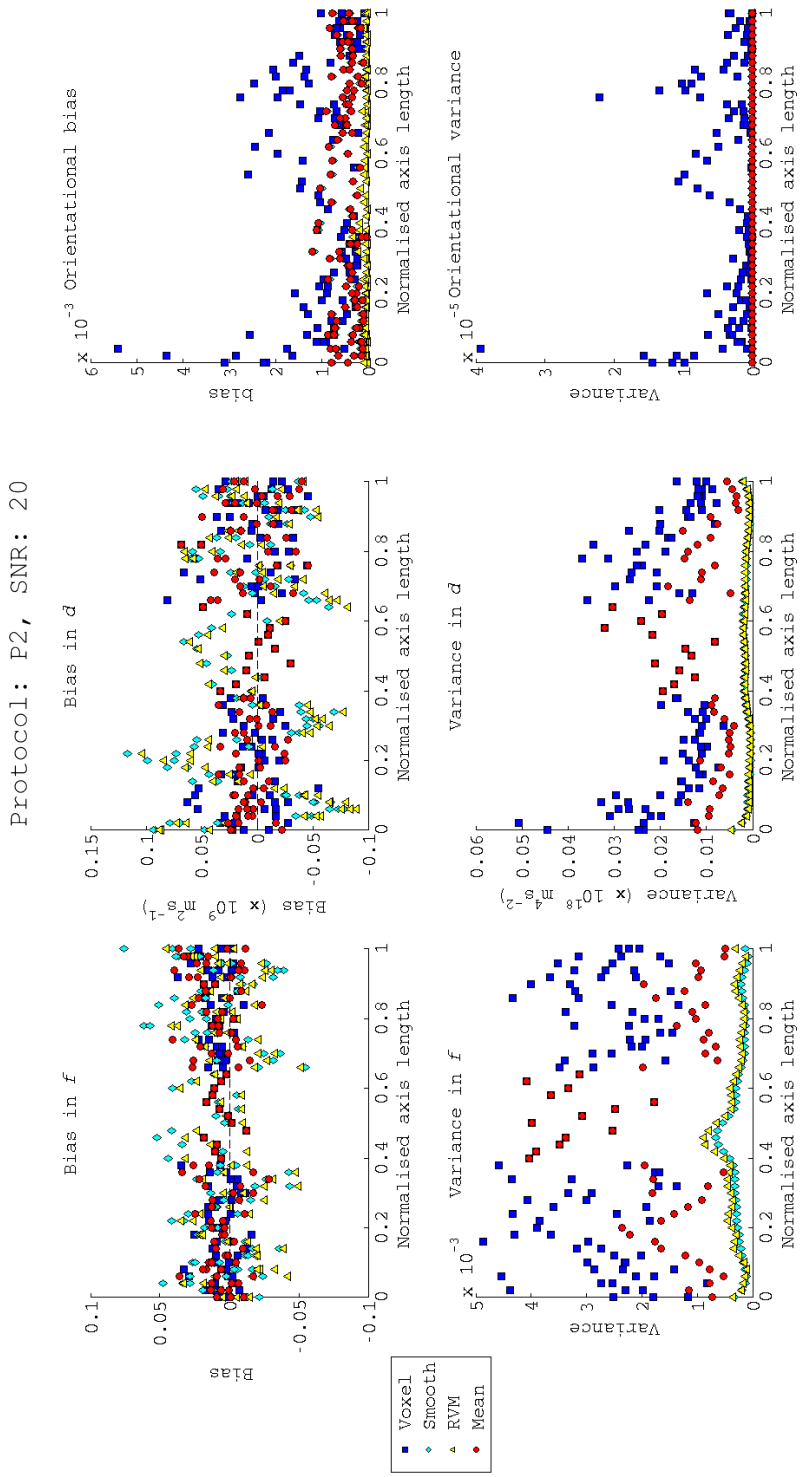

Figure 4.7: As for figure 4.6, but for data simulated using protocol P2 and SNR=20.

voxelwise estimates which is most prominent at a distance of $0 - 0.2$ along the normalised axis length. This introduces a positive bias in the mean and the smooth trends for d, as they depend directly on the voxelwise estimates. We do not observe this large positive bias when the parameter trends are estimated using the RVM. This shows the benefit of fitting the trends in microstructure variation directly to the MR signals, instead of estimating them from voxelwise estimates. However, the bias for both the RVM and smooth trends does oscillate about zero, with the peaks and troughs corresponding to sharp changes in the ground truth parameters. As for f, this indicates that these techniques struggle to capture steep changes in microstructure. The variance in d is much smaller for the RVM and the smooth trends than for the voxelwise estimates and the corresponding mean trend, demonstrating that these techniques are more robust to noise. However, despite the fact that the smooth trends are reasonably robust to noise, they can still be biased. Finally, estimates of fibre orientation obtained from fitting the model to each voxel independently also have a positive bias, and this bias is also present in the mean and smooth trends which rely on the voxelwise estimates. There is minimal bias in fibre orientation when using the RVM. We observe similar results when considering the other acquisiton protocols (P1, P2, P4) and low SNRs ($\mathrm{SNR} = 5, 10$), although the positive bias in d and $\mathbf{e}$ for the voxelwise estimates, mean and smooth trends becomes even more extreme as the number of diffusion-encoding directions decreases. In these cases, the bias for d and $\mathbf{e}$ remains low when using the RVM.

Figure 4.7 shows the same bias and variance plots, using a more modest acquisition protocol (P2) but with SNR=20. In this case we notice that the bias in f and d is more tightly clustered around zero for the voxelwise estimates as well as the mean trends. The values of bias for both the RVM and the smooth trends are greater than those obtained using voxelwise estimation, ranging between -0.05 and 0.05 for f and between -0.05 and $0.05\times10^{-9}\mathrm{m}^2\mathrm{s}^{-1}$ for d. However, over the length of the curve the values for f are still evenly distributed about zero, implying that there is no systematic bias present. For d the bias oscillates about zero, implying that it is systematically under- or overestimated at certain positions along the curve. This is most likely due to the weak dependence of the signals of d, resulting in some oversmoothing. The variance in both f and d over the whole curve is again much smaller for the RVM and smooth trends than for the voxelwise and mean estimates, although all of the values are much smaller than those shown in figure 4.6. Finally, there is still a slight orientation bias when estimating θ and ϕ independently in each voxel, although it is small and does not significantly impact the mean and smooth estimates of fibre orientations. Similar trends are observed for the other acquisition protocols (P1, P3 and P4) when the SNR is high (SNR=15, 20, 25).

As the previous figures have shown, provided that the SNR is not too low (i.e SNR$\geq$ 10), the performance of the RVM and the smoothing trends, with respect to bias and variance, is very similar. Therefore, the question remains: is there any advantage in estimating the trends of parameter variation directly from the data with the RVM, rather than fitting smooth splines to the voxelwise parameter estimates? In figures 4.8 and 4.9 we show curves of volume fraction and diffusivity, estimated with both the RVM and smooth trends. In these figures we show trends for f and d, calculated from the mean of the posterior distribution, along with the 95% confidence intervals. The width of the confidence intervals

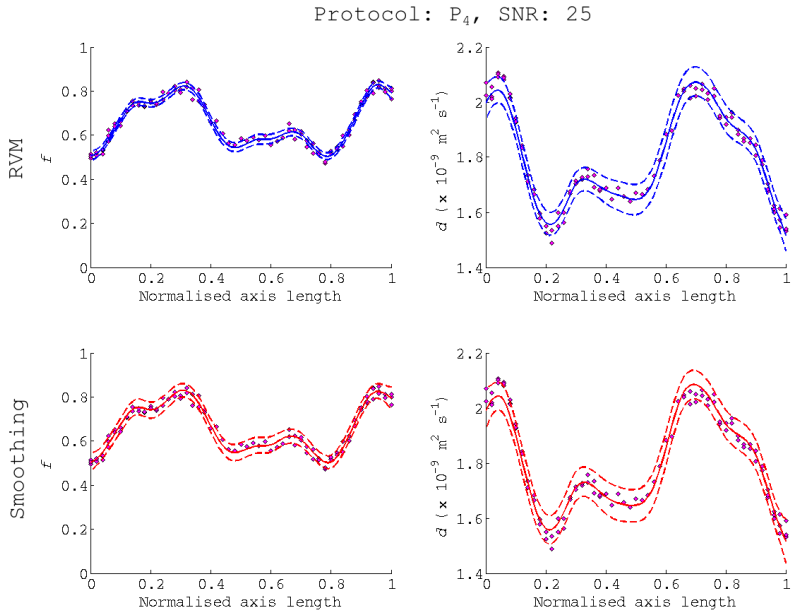

Figure 4.8: Example curves of f and d estimated using both the RVM and smoothing curves, for one data set simulated using P4 and SNR=25. For f the confidence intervals are narrower along the whole curve when using the RVM, implying that it estimates the variation in volume fraction more precisely. The confidence intervals over d are narrower at a distance of $0.3 - 0.5$ along the axis length.

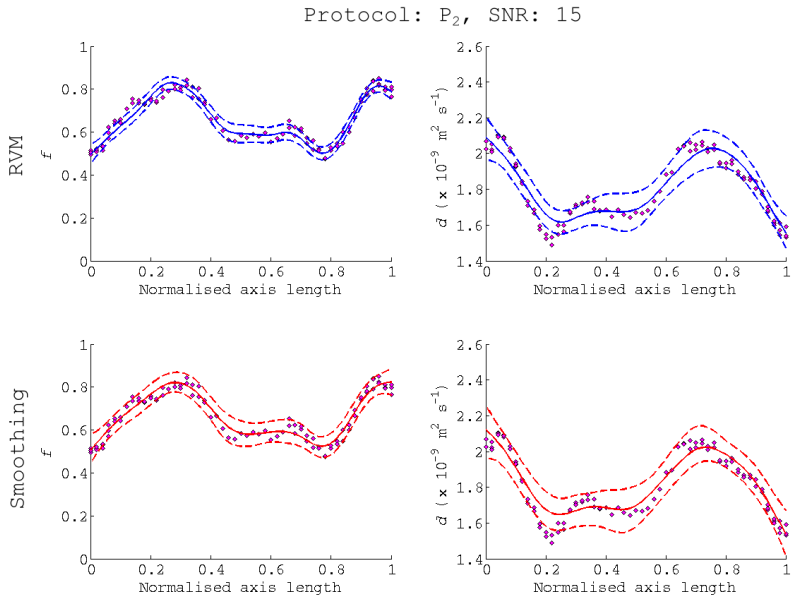

Figure 4.9: As for figure 4.8, but using protocol P2 and SNR=15.

indicates how precise our estimates of the mean parameter trends are. As we can see in both figures, the curves for f and d are almost identical for both methods, implying that the accuracy of the techniques is comparable. However, regardless of the protocol or SNR, the confidence intervals for f over the whole curve are narrower when using the RVM compared to the smooth trends. For example, in figure 4.9, at a distance of 0.58 along the medial axis, the confidence interval for the RVM is 0.55-0.62, whereas it is 0.53-0.64 for the smooth trend. For d, the differences in the confidence intervals are not quite as clear, although in figure 4.8, using protocol P4 and with SNR=25, the confidence intervals using the RVM are slightly narrower at a distance of $0.3 - 0.5$ along the normalised axis length. For example, at a distance of 0.4 along the medial axis, the 95% confidence interval ranges from 1.63-1.71$\times 10^{-9}$m^{2}s^{-1} for the RVM compared to 1.61-1.73$\times 10^{-9}$m^{2}s^{-1} for the smooth trend. In figure 4.9, using protocol P2 and with SNR=15, the confidence intervals for the RVM are narrower at a distance of $0.1 - 0.4$ along the normalised axis length. At a distance of 0.25 along the medial axis, the confidence interval for the RVM is 1.55-1.68$\times 10^{-9}$m^{2}s^{-1} whereas for the smooth trend the range is 1.55-1.74$\times 10^{-9}$m^{2}s^{-1}.

4.2.4 Discussion

The results of the first simulation experiment show that, contrary to hypothesis 1, curves of parameter variation estimated using the RVM are only robust to the choice of spatial model settings when the SNR is high, i.e. 20 or 25. Even then, when the number of diffusion-encoding directions is low, we notice slight differences in the AMSE for different combinations of knots and smoothing priors, particularly for d. We ideally want to use the RVM to estimate microstructural parameter trends without having to fine tune the spatial model settings for every data set. In section 4.2.3, we try to provide some general guidelines that are applicable for a wide range of protocols and SNRs. In general, we find that the AMSE is minimised for both f and d when knot vector t_2 is used. We also note that the AMSE is generally higher for d than for f, most likely because the dependence of the signal on d is weaker than the dependence on f, making it more difficult to distinguish between different values of d. When using the medium number of knots, AMSE$_d$ is consistently low for all protocol/SNR combinations when the amount of smoothing is moderate. If the smoothing is too low, the diffusivity curves tend to overfit to the noise, whereas if the smoothing is too high or the prior is uninformative, the penalty term due to the smoothing prior becomes too large, resulting in flat curves, as shown in figure 4.5. The choice of smoothing parameter for f is less important. We identify one combination of spatial model settings that works well for a wide range of data, however it is unlikely to be optimal. Therefore in future, it may be necessary to optimise the spatial model settings simultaneous with the basis function weights $\mathbf{a}$ during the RVM fitting to ensure that we obtain the best solution for each data set.

We also note in section 4.2.3, that there are other limitations to this method for optimising the spatial model settings. The AMSE only provides a single number to determine how well the RVM matches the ground truth parameters and does not provide any information about how well the model fits along the entire axis. This can result in situations, as previously mentioned, where a completely flat curve is preferred over a curve that fully captures the variation of the microstructure parameters, albeit with a small systematic bias. Therefore, in future we can improve our criteria for the optimal spatial

model settings by incorporating this information, for example by tracking the values of AMSE along the medial axis or using correlations between the ground truth and estimated parameters. We have not included this here as, given that the model fits the data reasonably well, a single correlation value is unlikely to provide more information than the AMSE. Tracking the AMSE across the medial axis would provide more detailed information, but it is a difficult problem to decide how best to use this additional information to select the best combination of settings Υ and would require careful implementation.

When comparing the parameter estimates obtained using the four different techniques, we find that the RVM and smooth trends are consistently more robust to noise than voxelwise and mean trend estimation for a wide range of acquisition parameters and noise levels. The variance of the model parameters estimated using these techniques is consistently lower, implying that the methods are more precise and robust to noise. The level of bias is similar for all methods, except at low SNRs (i.e SNR< 10) when some of the voxelwise parameters display systematic bias. This directly impacts the smooth and mean trends. Although in these cases, the RVM also displays some systematic bias, it is less severe than for the alternative methods. This shows the advantage of estimating trends of microstructure variation directly from the diffusion-weighted signals, particularly when the data is noisy. At low SNRs we observe some oscillations in the bias curves for both the RVM and smooth trends, particularly for d. The oscillations primarily occur in positions where the underlying ground truth values change suddenly. This indicates that the smoothing prior is penalising the curve too heavily at these points and oversmoothing the trends. This could be overcome by introducing variable smoothing [107], which allows the strength of the smoothing penalty to vary as a function of position, or by using a basis function that allows sudden jumps such as wavelets [75]. For richer acquistion protocols and higher SNRs, the differences between the RVM and the smooth trends are subtle. By considering the posterior distribution over the estimated model parameters, we see that whilst the mean posterior distributions are almost identical for the two methods, the confidence intervals are narrower for the RVM. Therefore, even though the RVM is not necessarily more accurate than the smooth trends, the trends of parameter variation estimated using the RVM are more precise. In applications such as Bayesian group comparisons, which use confidence intervals to determine significance, using the RVM could be advantageous when studying subtle differences between populations. We must also note that we use the combination of knot vector and smoothing priors optimised for the RVM in order to fit the smooth trends. The reason for this is to ensure that we compare the models under identical settings. Although the RVM and smooth trends are fit to different data sets, the model parameters that are being estimated represent the same quantities, therefore it is reasonable to assume that the same spatial model settings perform well. However, we may be able to improve the performance of the smooth trends by optimising the spatial model parameters specifically for this method.

4.3 Experiment 2: Validation using ex vivo 'silver standard' data

The previous experiment demonstrated that the RVM is able to capture the underlying trends of microstructure variation in synthetic data, and that it outperforms alternative methods for estimating microstructural trends, particularly when the SNR is low. Whilst synthetic validation is useful to test that

our methods perform sensibly, ideally we want to test how well they work on real data, for which the underlying tissue structure is much more complex. However, validation against real data is often difficult due to the lack of ground truth. In this experiment we use an ex vivo monkey brain data set acquired using over 100 gradient directions. Parameter estimates obtained from this rich data set are used as a pseudo ground truth. We then subsample the full data set to generate sparser data sets with a more realistic number of gradient directions. By estimating the trends of parameter variation using both the RVM and alternative techniques from these subsampled data sets, we can determine how robust the methods are, relative to the pseudo ground truth.

4.3.1 Hypothesis

Based on the results of our first experiment, we hypothesise that

1. when the angular resolution of sampling directions is high, all estimation methods provide robust estimates of f, d and **e**.

2. as the number of sampling directions is reduced, the RVM outperforms the alternative methods, both in terms of bias and variance of the parameter estimates.

4.3.2 Methods

In this section we provide details of the imaging protocol used to acquire the ex vivo monkey brain data. As we focus on fitting microstructural trends within the mid-sagittal CC, we also discuss the methods we use to segment the ROI and fit the cm-rep model. We describe how we generate sparse data sets from the full set of images before discussing the procedures used for model fitting and statistical analysis.

Data acquisition

The ex vivo monkey data is acquired from the perfusion fixed brain of a 32 month old vervet monkey (*Chlorocebus pygerythrus*), obtained from the Montreal Monkey brain bank, on a 4.7 T Varian scanner and is available to download from dig.drcmr.dk. The acquisition protocol, optimised using the framework in [6] for a maximum gradient strength $G_{max} = 300\mathrm{mT/m}$, comprises three HARDI shells with b values $\{2421, 4609, 11329\}$ s mm^{-2} and $\{102, 105, 82\}$ diffusion-encoding directions respectively. The protocol also includes 75 $b = 0$ images, with 25 unweighted images acquired before each HARDI shell. The echo time $\mathrm{TE} = 71\mathrm{ms}$ and the repetition time $\mathrm{TR} = 2500\mathrm{ms}$. The slice direction is perpendicular to the mid-sagittal plane. A total of 30 sagittal slices are acquired, with isotropic voxels of size $0.5 \times 0.5 \times 0.5\mathrm{mm}^3$. The SNR of the data in white matter regions, calculated using Camino, is approximately 20. The data is kindly provided by Dr Tim Dyrby of the Danish Research Centre for Magnetic Resonance, University of Copenhagen.

For the purposes of this experiment, we use 25 $b = 0$ images and a single HARDI shell with the lowest b value to fit the ball and stick model to the data, as the SNR of the diffusion-weighted signals is highest for this shell. We refer to this as the full data set, which we use to determine the pseudo ground truth parameter values.

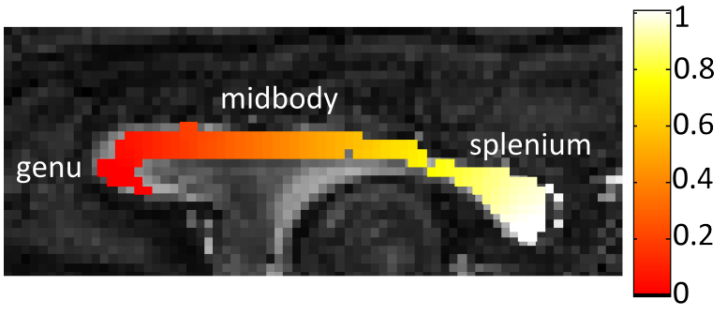

Figure 4.10: The segmented ROI within the mid-sagittal CC. All voxels are colour-coded according to their positions along the medial axis, which is fitted using cm-rep. The length of the medial axis is normalised on the interval $[0, 1]$, where 0 corresponds to the genu, at the anterior end of the CC, and 1 corresponds to the splenium, at the posterior end of the CC.

ROI segmentation

In this experiment we examine trends of microstructure parameter variation in the mid-sagittal CC. We segment this region using the following steps:

1. Extract the mid-sagittal slice. In this case, the slice extraction step can be done manually as the slice direction is perpendicular to the mid-sagittal plane.

2. Calculate maps of linearity c_L and planarity c_P for the mid-sagittal slice, where the expressions for c_L and c_P are given in section 2.4.1.

3. Create a binary mask over the mid-sagittal slice containing only voxels for which $c_L > 0.5$ and $c_P < 0.2$.

4. Perform connected component analysis on the binary mask in order to identify fully-connected regions containing 100 or more voxels.

The CC segmentation often includes several voxels from the fornix. We manually remove these voxels using ITK-SNAP [196]. We also exclude several voxels at the gray matter/white matter boundaries for which, due to large signal drop-outs in the $b = 0$ images, the unweighted signals are significantly lower than the diffusion-weighted signals. This is most likely due to Gibbs ringing [11], which occurs at high contrast boundaries due to undersampling along the phase-encoding direction [118]. The final ROI contains a total of 135 voxels. We then fit the two-dimensional cm-rep model [171] to the ROI mask in order to fit the medial axis. We then use the nearest neighbour approach mentioned in section 3.4 to assign each voxel a distance x along the medial axis. Figure 4.10 shows the final region of interest overlaid on the FA image. Voxels within the ROI are colour-coded according to their position along the medial axis of the CC, normalised on the interval $[0, 1]$.

Data subsampling

We subsample from the full data set to generate sparse data sets with a reduced number of diffusion-encoding directions. Specifically we consider four sparse data protocols. The number of unweighted n_0 and diffusion-encoding directions n_g of these four protocols matches the acquisition parameters given in table 4.1. For each protocol we randomly select n_0 of the 25 $b = 0$ images and n_g of the 102 diffusion-weighted images from the full data set. However, we do not guarantee that the gradient directions are equally distributed over the unit sphere. We repeat this procedure 25 times to provide a representative selection of sparse data sets. By estimating the parameters in each data set we can calculate the bias and variance of the estimates relative to the pseudo ground truth, as in section 4.2. Although there are algorithms such as [45] which optimally divide a set of acquisition directions into subsets, we do not use these here as they do not match the purpose of this experiment. For example, the algorithm in [45] would divide the 102 gradient directions in our full data set into 17 subsets of 6 directions, 6 subsets of 15 directions, 3 subsets of 30 directions and only 1 subset of 60 directions. Therefore, it does not provide us with enough subsampled data sets from which we can calculate meaningful bias or variances.

Model fitting

After performing the subsampling step described above we have five distinct data sets: one full data set comprising 25 unweighted images and 102 images with diffusion weighting $b = 2421$ s mm^2, and four subsampled data sets with 6, 15, 30 and 60 directions respectively, each with 25 replications. In the following paragraphs we describe the model fitting procedures used to estimate the pseudo ground truth and to estimate the fitted parameters using the different fitting methods.

Estimating the pseudo ground truth We fit the ball and stick model to each voxel of the full 102 direction data set independently, in order to obtain estimates of f, d, θ and ϕ in every voxel in the ROI. These values act as our pseudo ground truth when comparing the estimates from the subsampled data sets. We fit the ball and stick model using Camino [44]. We use a Rician noise model, and set the variance of the likelihood function σ according to the value given by the `estimatesnr` function in Camino. We use the MCMC optimisation algorithm to fit the model to the data using the same burn-in and sampling parameters given in section 4.2.2.

Subsampled data sets For all of the subsampled data sets, we fit the ball and stick model using the RVM and independent voxelwise estimation, and derive parameter trends from the voxelwise estimations to create mean and smooth trends as described previously in section 4.2.2. For the RVM and smooth trends, we use the optimal settings determined in section 4.2.3, i.e. the knot vector is $t = \{-0.1 : 0.05 : 1.1\}$ and the prior distributions on the smoothing parameters are $\Pr(\Lambda_f) = \mathrm{S}_4$ and $\Pr(\Lambda_d) = \mathrm{S}_3$. We initialise $\mathbf{a}_d$ (for RVM fitting) and d (for independent voxelwise fitting) to $0.6 \times 10^{-9} m^2 s^{-1}$, which is a more suitable diffusivity for ex vivo data than the values used in the simulation study, and is the value used in previous experiments on this data set [6, 60]. Again, we fix the variance of the Rician likelihood function using the estimate of σ calculated using Camino, and use the same burn-in and sampling rates as in section 4.2.2.

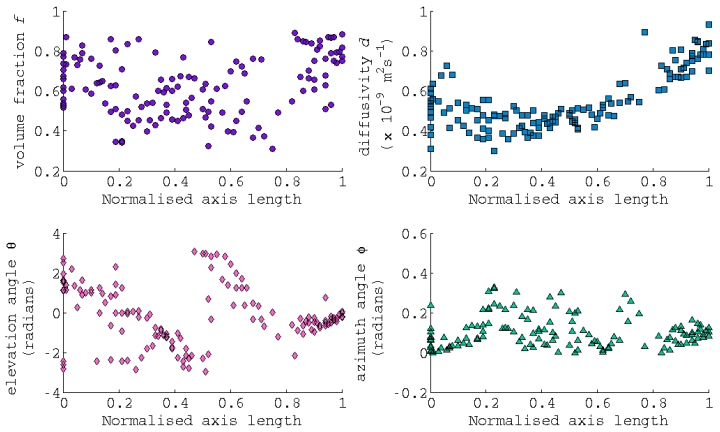

Figure 4.11: The pseudo ground truth parameter values for the ex vivo monkey brain data, estimated by fitting the ball and stick model to the full data set.

Statistical analysis

In order to compare how well the different parameter estimation methods perform for sparse acquisition protocols, we calculate the bias and variance of the parameter estimates in each voxel, $\hat{f}_i$, $\hat{d}_i$ and $\hat{\mathbf{e}}_i$, relative to the ground truth values in each voxel, f_i, d_i and $\mathbf{e}_i$, over all replications in the subsampled data sets for all estimation techniques. The equations used to calculate the bias and variance for all the microstructure parameters are given in section 4.2.2, in equations 4.7, 4.8, 4.9 and 4.10.

4.3.3 Results

Figure 4.11 shows estimates of f, d, θ and ϕ obtained by fitting the ball and stick model to the full data set, which are our pseudo ground truth parameters. There is a high-low-high trend in f across the normalised axis length, however it is a much weaker trend than those used in the simulation experiment with a high amount of variability perpendicular to the medial axis. The trend in d, which increases across the normalised axis length, is much stronger and shows less variability perpendicular to the medial axis than f. The azimuth angle ϕ is close to zero across the whole of the axis, although there is more variability than we would expect, particularly in the midbody of the CC, where a maximum value of $\phi = 0.34$ is observed. This corresponds to an angular difference of 20° compared to when $\phi = 0$. Although this is not completely unfeasible, it does seem large when compared with histology studies such as [2], which observe very coherent, parallel fibres in the mid-sagittal section of the CC. This is unlikely to be because of partial volume effects due to the linearity and planarity thresholds we used when segmenting the ROI, however it is possible that the fixation process has subtly altered some of the tissue properties, including the orientation of the fibres. This may affect the performance of the methods which estimate the trends of microstructure variation as we assume that ϕ is constant across the ROI, although the effect should be small as the z component of $\mathbf{e}$ only varies from 1 to 0.95 as ϕ varies from 0 to 0.3 radians. The elevation angle θ displays a large amount of variation over all voxels. However, given the definition of $\mathbf{e}$, when ϕ is close to zero, $\sin\phi \approx 0$, and therefore θ has negligible effect on $\mathbf{e}$.

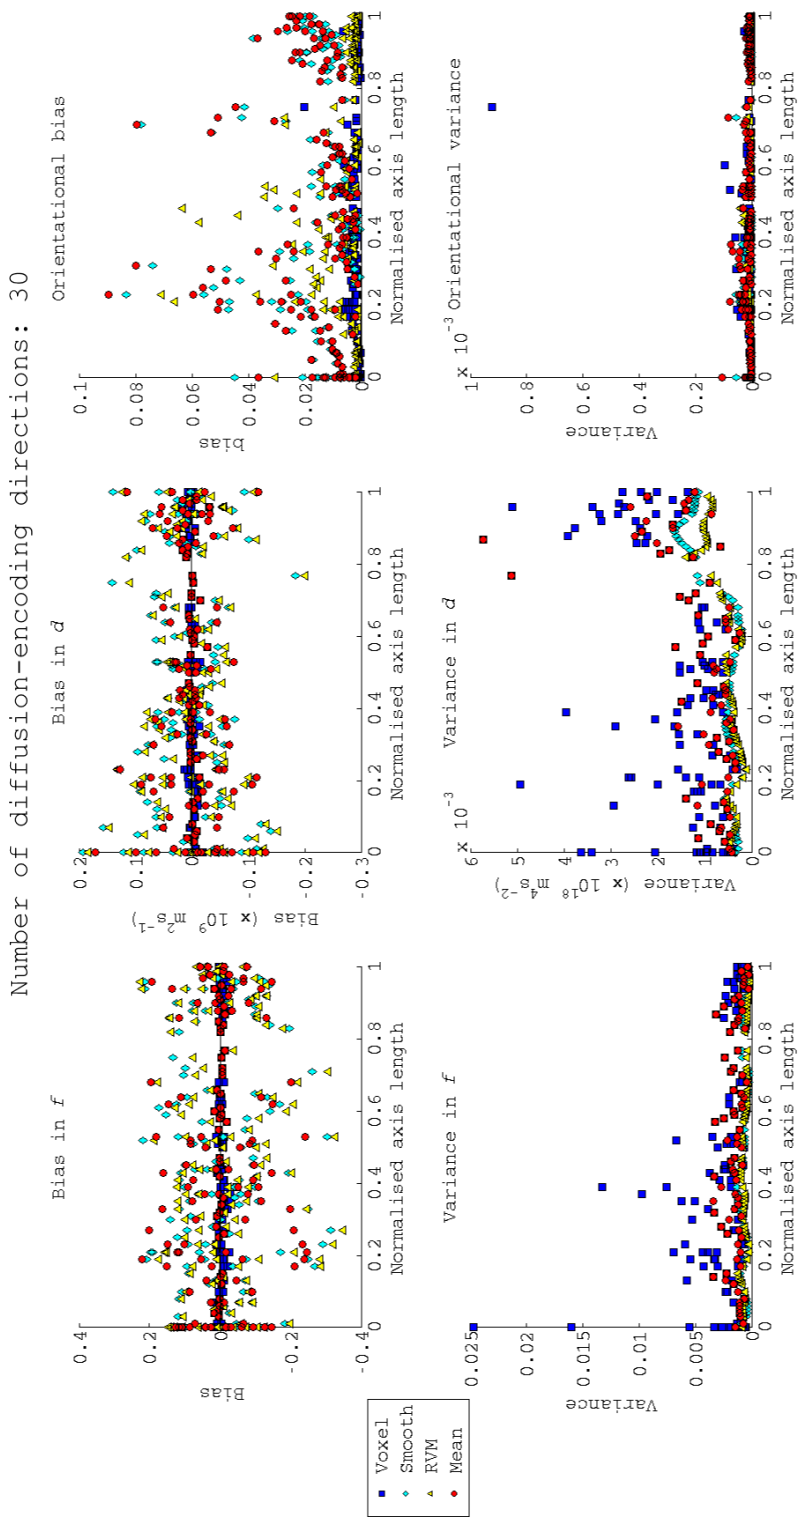

Figure 4.12: Plot of bias and variances for f, d and $\mathbf{e}$ for the four different parameter estimation techniques, estimated from data using 30 subsampled gradient directions.

The results obtained when we fit the models to the data sets with 15, 30 and 60 directions are very similar, therefore we only show the results for one case with 30 diffusion-encoding directions in figure 4.12. As we can see from this figure, the bias in both f and d for the independent voxelwise estimates is very tightly clustered about zero. When we extract the general trends in f and d using either the RVM, mean or smooth trends, the range of values that we observe for bias is much higher, typically between -0.2 and 0.2 for f and between -0.2 and $0.2 \times 10^{-9} m^2 s^{-1}$ for d. However, over the entire length of the CC they are spread evenly about zero, implying that there is no systematic bias in any of the methods. Given the weak trends in the pseudo ground truth parameters shown in figure 4.11, the large values of bias observed in the fitted trends are not necessarily unexpected. Similar to the simulated experiments, the variance of f and d, calculated over the 25 subsampled data sets, when estimated using the RVM or a smooth trend is consistently lower than the voxelwise estimates and slightly lower than the mean trend estimates, particularly when the number of subsampled directions is more modest, i.e. 15 or 30 rather than 60. Due to the variation of ϕ observed in the pseudo ground truth values, all of the methods which estimate a single value of ϕ over the whole of the CC, i.e. the RVM, the mean trend and the smooth trend, display greater orientational bias than the voxelwise estimates.

The results we obtain when fitting the models to the data sets with only 6 directions are noticeably different from those discussed above, and the corresponding figures for bias and variance are shown in figure 4.13. We see a reasonably large negative bias in f and d and a large bias in $\mathbf{e}$ across the whole of the CC when we estimate the model parameters in each voxel independently. This bias is also observed in the mean trends and smooth trends, which rely directly on the voxelwise estimates, however it is not observed in the RVM results for any of the parameters. This highlights the importance of estimating the trends of microstructure variation directly from the diffusion MR signals when the data is sparse, rather than relying on biased voxelwise parameter estimates. The variance in f, d and $\mathbf{e}$ is also large for the voxelwise, mean trend and smooth trends, but once again it is negligible for the RVM results. Figure 4.14 helps to explain the results for the voxelwise estimates, mean trends and smooth trends when only 6 diffusion-encoding directions are sampled. The top row of the figure shows the results for an example data set in which the RVM performs well but the other methods fail, whereas the bottom row shows an example data set in which all methods perform well. In each figure we show the pseudo ground truth parameter values, the voxelwise parameter estimates and the RVM and smooth trends for both f and d. The final panel in each row shows the directions of the subsampled gradients. In the bottom row, we see that all methods perform well when the diffusion-encoding directions are reasonably spread out over the hemisphere. However, when the data set contains two or more diffusion-encoding directions that are roughly collinear, as shown in the top row, independent voxelwise parameter estimation fails as there is not enough independent information from which to infer the parameter values. As the mean and smooth trends are fit directly to the voxelwise parameter estimates, if these estimates fail to capture the underlying parameter variations, so do the estimated curves. In this situation, the RVM still performs well, most likely due to the direct fitting of the curves to the underlying diffusion-weighted signals, which utilises information from the local spatial neighbourhood. Although in practice

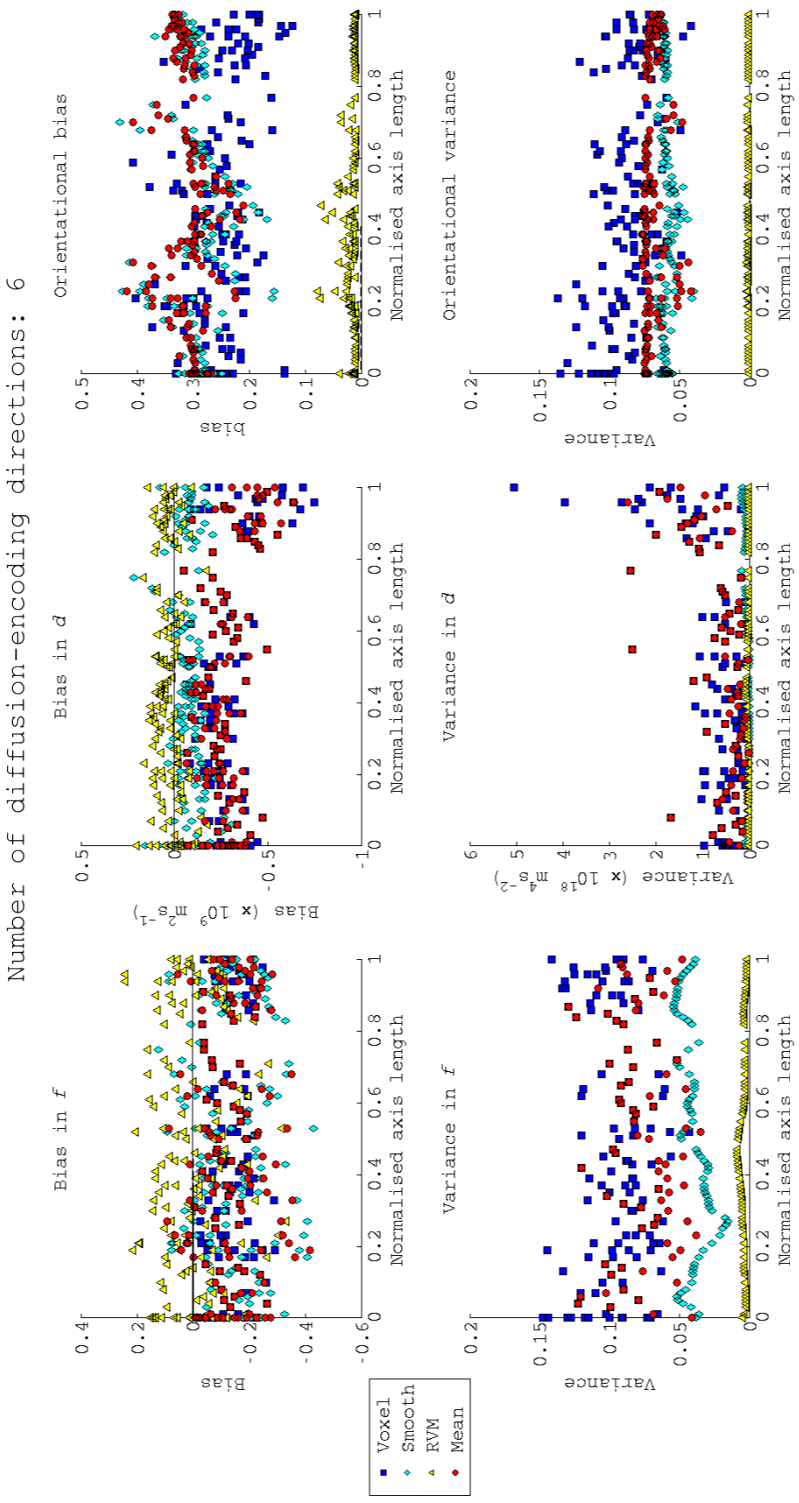

Figure 4.13: As for figure 4.12, but using only 6 subsampled gradient directions.

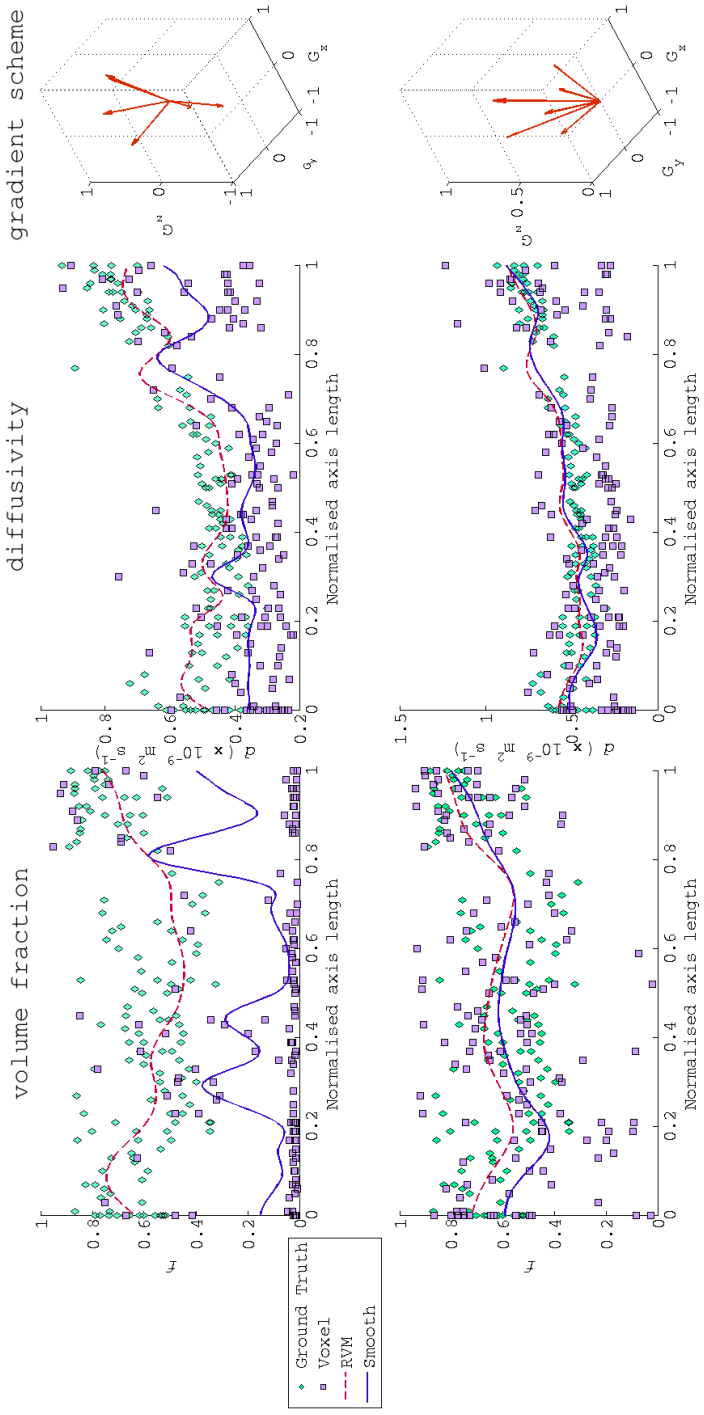

Figure 4.14: When sampling only 6 gradient directions, if two or more directions are almost collinear, as in the top row, voxelwise parameter estimation fails to capture the underlying values of f and d, which negatively impacts the mean and smoothed curves. However, in this scenario the RVM is still able to capture the key parameter trends. When the gradient directions are more spread out, as in the bottom row, all parameter estimation methods produce sensible results.

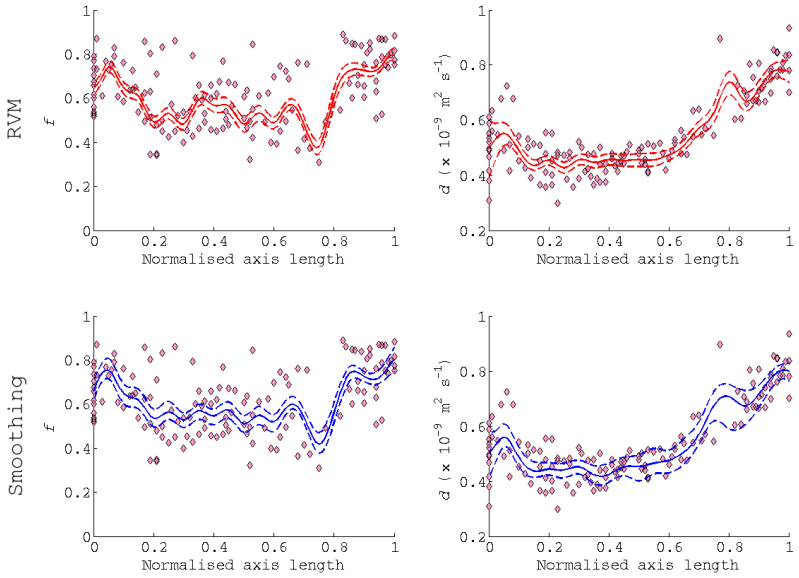

Figure 4.15: Example curves of f and d estimated using both the RVM and smoothing curves, for one subsampled data set with 60 gradient directions. For both f and d, the confidence intervals are narrower when using the RVM, implying that it estimates the variation in both parameters more precisely.

sparse acquisition protocols do not include collinear gradient directions, it is possible that in scans with only 6 directions, one of the images may suffer from artifacts which make it unusable. This information loss would potentially affect the independent voxelwise estimation, whereas the RVM would be able to recover sensible trends.

However, as shown in figure 4.12, provided that we sample enough unique gradient directions, the performance of the RVM and the smooth trends is very similar, so is there an advantage to using the RVM? Figure 4.15 shows the mean trends and the 95% confidence intervals for f and d, estimated using both the RVM and smooth trends, from one of the subsampled data sets with 60 diffusion-encoding gradient directions. As for the simulation experiments we see that the confidences intervals are narrower for both f and d when using the RVM, implying that the curves of parameter variation estimated using the RVM are more precise. For example, in figure 4.15, at a position of 0.25 along the medial axis, the 95% confidence interval on the mean trend for f ranges from 0.51-0.56 for the RVM compared to 0.53-0.61 for the smooth trend. Similarly for d, at a distance of 0.56 along the medial axis the confidence intervals for the mean trend is 0.44-0.48$\times 10^{-9}$ m^{2}s^{-1} for the RVM and 0.43-0.52$\times 10^{-9}$ m^{2}s^{-1} for the smooth trend. This is observed for all subsampled data sets in this experiment, showing that this is a general advantage of the RVM.

4.3.4 Discussion

At the beginning of this experiment, we hypothesised that the RVM would outperform the alternative parameter estimation techniques, particularly when the number of sampled gradient directions is low.

The results support this hypothesis when comparing the RVM to the independent voxelwise estimates and mean trends for all the protocols we consider. However, when comparing the RVM to the smooth spline curves, the improvements are only noticeable when we sample 6 directions from the full data set. In this case, when two or more of the gradient directions are almost collinear, the independent voxelwise estimation struggles to fit the ball and stick model to the data, which is reflected in the shape of the smoothed spline curves. Under these conditions, the RVM performs well and manages to fit sensible trends for f and d to all of the data set, resulting in low bias and variance, suggesting that the RVM may be able to fit trends with fewer gradient directions and thus potentially fewer images. Therefore the RVM could be used for applications such as neonatal imaging, in which scans are often cut short for the comfort of the babies in the MR scanner, resulting in a very limited number of images acquired. As in the previous experiment, when we consider the richer acquisition protocols, the differences between the RVM and the smooth trends are more subtle. By considering the full posterior distributions for both methods, we again notice that the confidence intervals are narrower for the RVM than the smooth trends, indicating greater confidence in the estimated spatial model parameters.

However, there are several important points we must consider. First, we use the spatial model settings identified in the first experiment when fitting both the RVM and the smooth trends. We do this as the number of gradient directions and SNR of the ex vivo monkey data match those used in the simulations and we anticipate that the settings are transferable. However, the trends in the ex vivo monkey data are generally much weaker than those used in the synthetic phantom, which may mean that a different set of knots and smoothing priors would perform better.

Second, both the RVM and the smooth trends only estimate the mean trends of microstructure variation across tracts and do not attempt to quantify the variability perpendicular to the medial axis. As we see in figure 4.11, the variability in this direction can be large and could cause problems when trying to estimate mean trends. This may also be an important indicator of tissue integrity, e.g. we may find that in some pathologies, the mean trend of microstructure variation remains constant although the indices themselves become more variable. In future, we can try to incorporate a measure of this variability into the RVM framework, although given that the number of voxels perpendicular to the tract is often limited, this requires careful consideration to implement in a sensible and meaningful way.

Finally, our assumption of constant fibre orientation may be unsuitable even in the CC. As we see in figure 4.11, the variation in ϕ is larger than we would expect, given histological studies of fibre direction in the CC [2]. In future we could generalise the spatial model further to allow estimation of ϕ and θ as functions of position along the tracts.

4.4 Conclusions

In this chapter we have tested the RVM model using both synthetic and ex vivo monkey brain data. We have investigated the effect of the spatial model settings on the resulting parameter variation curves and provided some general guidelines on choosing suitable knot vectors and smoothing priors. However, this is a difficult problem, given the wide range of knot vectors and prior distributions available, particularly if we want to identify a general set of knot vectors and prior distributions which is suitable for a wide

range of data acquisition parameters. So far, we have used a single estimate of error between our parameter estimates and the ground truth values over the whole ROI. This does work reasonably well for the protocols and SNRs we investigate here, although in future we can modify this procedure to account for the spatial nature of the RVM framework.

We have shown, using both synthetic and ex vivo data as ground truths, that the RVM can estimate plausible trends of microstructure variation for a range of acquisition protocols and noise levels. We have shown that the RVM is more robust to noise than voxelwise estimation and mean trends of variation, in terms of bias and variance calculated over numerous noisy data realisations for all the protocols and SNR levels we have investigated. However, when comparing the RVM to the smooth spline curves proposed by [71, 68], the RVM is only significantly better at estimating microstructural trends robustly when the SNR of the data is particularly low, i.e. 5 or 10, or when the gradient directions used in the acquistions are suboptimal. In these situations, the smooth splines curves do not produce plausible trends due to problems with independent voxelwise estimation. For higher values of SNR and richer acquisition protocols, the accuracy of the RVM and the smooth spline curves is very similar, however trends of microstructure variation estimated using the RVM are slightly more precise.

Chapter 5

Group studies using the RVM

5.1 Motivation

In the previous chapter, we validated the RVM framework using both simulated and ex vivo diffusion MR data. We showed that the RVM can accurately and precisely model the trends in parameter variation across white matter regions and that it outperforms alternative methods for estimating microstructure trends for low SNRs and suboptimal gradient sampling schemes.

Following on from these results, in this chapter we present a key application of the RVM - group studies. As discussed in chapter 2, one of the main aims of medical imaging is to detect and track pathologies using image-derived biomarkers, such as the microstructure indices estimated from diffusion MR signal models. This typically involves comparing these indices across groups of control subjects and patients in order to determine which microstructure indices are affected by the pathology, and to localise where the changes occur. As we mentioned in chapter 2, this is usually done using ROI analysis, VBM, or more recently and specifically for diffusion MR studies, TBSS and tract-based studies. However there are problems inherent in all of these methods. ROI analysis is heavily dependent on the choice and segmentation of white matter regions and the localisation is coarse. TBSS only uses a small number of voxels with the peak parameter values from each tract during the analysis, discarding a large amount of potentially useful information which could make it particularly unreliable at low SNRs. Finally, as shown in the previous chapter, the success of the tract-based methods is dependent on the independent voxelwise estimates, which are often inaccurate when the SNR is low. Therefore, given the precision, accuracy and continuous nature of the parameter variation curves estimated using the RVM, we believe that the framework is ideally suited for performing group comparison studies.

In this chapter we present two experiments to demonstrate the use of the RVM framework in group studies, and to compare it against some of the most widely used population analysis techniques described above. The first experiment is a simulation study, which aims to compare how well the various methods detect group differences relative to a known ground truth for a range of effect sizes, SNRs and acquisition protocols. The second experiment uses in vivo human brain data from 45 subjects in three different age groups, which we hypothesise corresponds to different effect sizes in the underlying microstructure parameters. We use the RVM framework, ROI analysis, TBSS and smooth tract-based methods to detect differences in the underlying parameters between the three groups and provide a comparison of the

techniques.

5.2 Experiment 1: Group study simulation

The goal of this experiment is to quantify how well the parameter curves estimated using the RVM are able to differentiate between distinct subject groups compared to standard population comparison techniques, under a range of conditions including different effect sizes, acquisition protocols and noise levels. As in chapter 4, we use the simple ball and stick diffusion model in order to represent the underlying microstructure. By using simulated data with a known ground truth, we can directly compare the performance of the RVM to ROI analysis, a TBSS-style method and the standard tract-based analysis.

5.2.1 Hypothesis

Given the accuracy and precision of the RVM for estimating trends of parameter variation compared to voxelwise techniques demonstrated in chapter 4, we hypothesise that

1. when the group sizes are large and the quality of the data is high, the power of all methods to detect group differences will be similar.

2. the RVM has greater power to detect subtle changes when the group sizes are small, compared with the alternative methods, particularly when the SNR of the data is low and the number of diffusion-encoding directions used to acquire the data is small.

5.2.2 Methods

In this section we describe the synthetic phantoms for subject groups with small, medium and large effect sizes. We also discuss the data synthesis and model fitting procedures used, as well as the statistical methods we use to detect significant differences between the groups.

Diffusion model

In this experiment we use the ball and stick model, described previously in section 2.4, which models the diffusion-weighted signals as a function of intra-cellular volume fraction f, diffusivity d and fibre orientation angles θ and ϕ.

Synthetic phantom construction

We construct the synthetic phantoms as described in section 4.2.2, using the same ROI comprising one slice of 100 voxels, which mimics the shape and structure of the mid-sagittal CC. However, instead of creating a single ground truth, we create four different ground truth data sets corresponding to the control group, the small effect group, the medium effect group and the large effect group. All of the microstructure differences between the various groups are within the first quarter of the phantom, corresponding to the genu of the CC. We define the three different effect sizes as follows:

- small effect size: 5% increase in d, 5% decrease in f.
- medium effect size: 10% increase in d, 10% decrease in f.
- large effect size: 15% increase in d, 15% decrease in f.

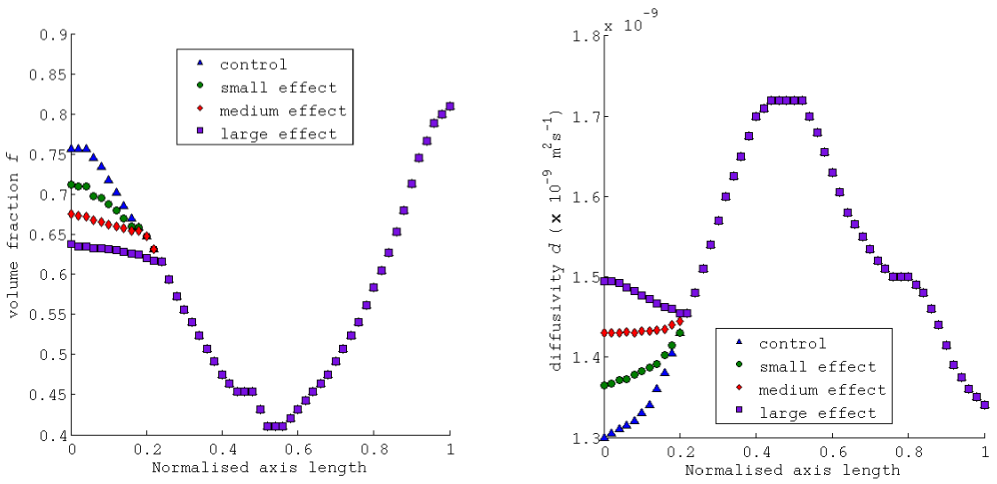

(a) (b)

Figure 5.1: (a) shows the ground truth parameter values for f for the control, small effect, medium effect and large effect groups; (b) shows the corresponding ground truth values for d.

For all effect sizes, the percentage change in the microstructure parameters is calculated relative to the control group. The spatial variation of f and d for each group are shown in figure 5.1. For each group, we create ground truth values for 25 unique subjects. The exact ground truth values for each subject are generated from Gaussian distributions with means equal to the functions f and d evaluated at voxel location x_i, and standard deviations 0.01 and $0.01 \times 10^{-9} m^2 s^{-1}$ respectively. This introduces a small amount of interpopulation variation into the ground truth data, but not enough to cause any overlap between the different groups and thus negate the population differences. As for the simulation experiment described in chapter 4, we set $\theta \approx \phi \approx 0$ in all voxels in the ROI for all effect sizes.

Data synthesis

We synthesise diffusion MR signals for each subject in all four groups using the ball and stick model given in equation 2.53, the ground truth parameters discussed above and two acquisition protocols. These correspond to protocols P2 and P4 described in chapter 4. To recap, P2 comprises 15 non-collinear diffusion-encoding directions and one unweighted signal (i.e. it is a reasonably sparse protocol), whereas P4 comprises 60 non-collinear diffusion-encoding directions and 6 unweighted signals (i.e. it is a high angular resolution protocol).

After synthesising the data for each subject for both acquisition protocols, we add Rician noise to the signals with σ chosen to correspond to SNR=10,25. This allows us to perform the group studies and compare the results under four key sets of conditions:

- low SNR and sparse acquisition,
- low SNR and rich acquisition,
- high SNR and sparse acquisition,
- high SNR and rich acquisition.

Model fitting

We fit the ball and stick model back to the data sets for every subject using both the RVM and independent voxelwise parameter estimation. The RVM is fit to each data set as described previously in section 4.2.2, using the optimal knot vector and smoothing priors determined in chapter 4. Once again, we fit the RVM using MCMC, with 5000 burn-in iterations and 500 samples collected at intervals of 20 iterations. The variance of the Rice distributed noise is fixed to the values of σ used to synthesise the data.

The ball and stick model is fit to each voxel independently using Camino, using MCMC optimisation and assuming Rice distributed noise, as previously described in chapter 4. The burn-in and sampling parameters for the MCMC algorithm are the same as for the RVM, and the variance of the noise is fixed during the optimisation to the value of σ used to simulate the signals.

In order to perform a tract-based group comparison similar to those described in [71, 74], we fit splines to the voxelwise parameter estimates of f and d in order to generate smooth functions of microstructure variation. We fit the splines using the Bayesian P-spline formulation described previously in chapter 4, using the same knot vector and smoothing priors as for the RVM fitting. As in chapter 4, we refer to these as smooth trends.

Group comparisons

In this experiment, we compare four different group comparison techniques: ROI analysis, TBSS, tract-based (smooth) analysis and finally, the RVM-based analysis. In the following paragraphs we provide further details of all these methods, discuss multiple comparison corrections and give an overview of all the comparisons we undertake.

For the ROI analysis, we divide the phantom ROI into segments based upon the position of the voxels along the medial axis. We then calculate average microstructure estimates in each segment and use one-sided t-tests to determine the regions in which the control values are significantly greater or smaller than the effect groups. In this experiment we consider two different ways of segmenting the ROI. The first divides the phantom into three equal regions, which provides very coarse localisation of microstructure changes, but increases the accuracy of the indices we compare by averaging over a large number of voxels. The second divides the phantom into ten equal regions, which improves the localisation at the expense of the accuracy. Schematic representations of these ROIs are shown in figure 5.2.

For the TBSS-style analysis, we skeletonise the phantom ROI and project the peak parameter values onto the skeleton, i.e. we reduce the phantom to a 1×51 voxel region containing the highest microstructure indices found along the tract. This process is shown in figure 5.3. We create a skeleton for each subject in the group and perform one-sided t-tests for every voxel within the skeleton to determine the group differences.

For the smooth tract-based analysis, we use the smooth trends fitted to the independent voxelwise estimates as described above in the Model Fitting section. We then perform pointwise one-sided t-tests at 51 equidistant positions along the medial axis of the ROI to evaluate the differences between the groups.

For the RVM analysis, we use the functions of parameter variation along the ROI obtained by fitting

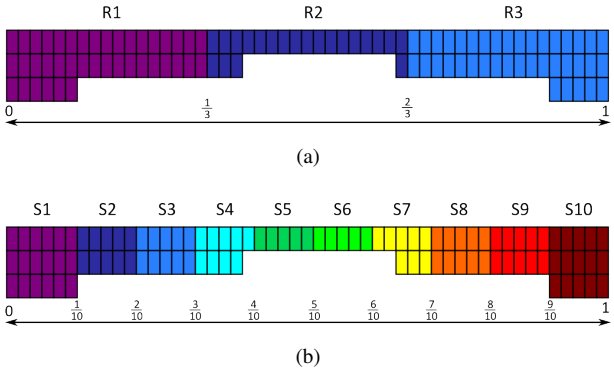

Figure 5.2: (a) shows the division of the phantom ROI into 3 parts of equal length, whereas (b) shows the division into 10 equal segments.

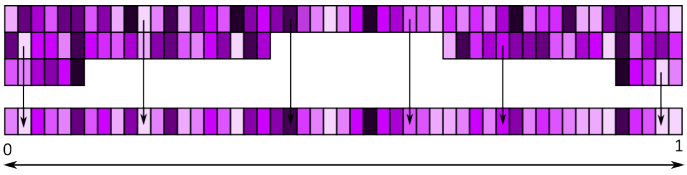

Figure 5.3: The skeleton is created by 'thinning' the ROI perpendicular to the direction of the principal axis. Maps of microstructure indices along the skeleton are generated by finding the maximum parameter values (indicated by lighter colours) perpendicular to the main tract axis for each voxel along the skeleton, and projecting them onto the skeleton.

the RVM to the noisy simulated signals, also described above in the Model Fitting section. As for the smooth tract-based analysis, we look for significant population differences between the RVM functions by performing pointwise one-sided t-tests at 51 equidistant points along the medial axis of the ROI.

For all of the analyses described above, we must correct the raw t-statistics obtained from the t-tests to account for multiple comparisons. In order to do this, we use the Westfall-Young correction [129], which is a permutation based method. The method works by permuting the group labels of the subjects multiple times. For each permutation, t-tests are performed and a permutation distribution is built up from the maximum t-statistics over all test points. The significance of the original raw t-statistics is then determined by evaluating the probability of observing these t-statistics within the permutation distribution. The Westfall-Young correction is the standard method of correcting p-values within TBSS [164]. It is also suitable for pointwise testing along continuous curves, such as those used within the tract-based and RVM group analyses, as it has been shown that as the number of points we test, and therefore the number of multiple comparison increases, the corrected p-values converge towards a continuum limit [51] rather than increasing towards 1.

For each of the four protocol/SNR combinations, we use the ROI, TBSS, smooth and RVM analyses described above to compare f and d between the

- control and small effect groups,
- control and medium effect groups,
- control and large effect groups.

However, as well as performing comparisons for different SNRs and protocols, we also want to investigate the effect of group size on the discriminative ability of the different analysis methods. Therefore, as well as comparing f and d between the groups containing all 25 subjects, we also compare the groups when they contain N_{subj}=20, 15, 10 and 5 subjects only. When using smaller group sizes, as there are numerous ways of choosing a subset of subjects from the full cohort, we investigate 20 possible random combinations for each group size to attempt to overcome the effects of subject selection. For consistency, we use the same subsets of subjects for all comparisons of a given group size.

Statistical analysis

To determine how well each of the methods perform, we compare the results to the ground truth. For each effect and group size, we perform a t-test (using the same multiple comparison corrections described above) on the ground truth values to determine whether the microstructure differences in a particular region or voxel location are significant. Regions or voxel locations that are significantly different are denoted as positives, whereas regions or voxel locations that are not significantly different are negatives. By comparing the p-values obtained from the ROI, TBSS, RVM and smooth analysis to these ground truth values, we can classify the results as either a true positive (TP), a true negative (TN), a false positive (FP) or a false negative (FN). For each method we calculate a sensitivity and specificity score, given the SNR, acquisiton protocol, effect size and group size. The sensitivity, which is the ability of a method to

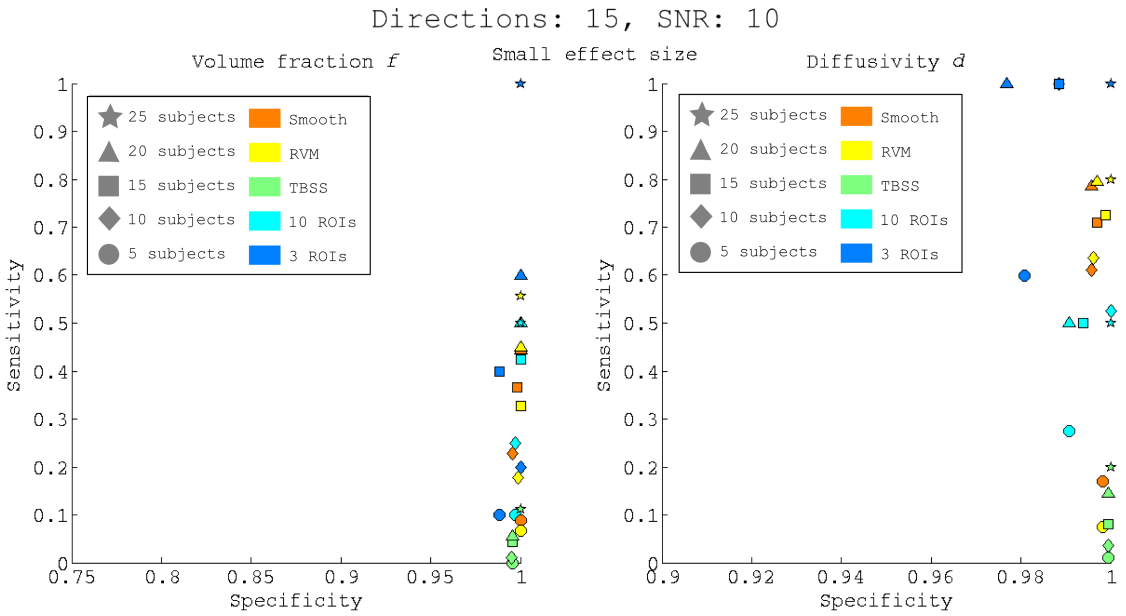

Figure 5.4: Sensitivity-specificity plot for all methods and group sizes, when detecting small effect sizes using an acquistion protocol with 15 gradient directions and an SNR of 10.

detect positive results, is defined as

$$\text{sensitivity} = \frac{n_{\text{TP}}}{n_{\text{TP}} + n_{\text{FN}}} \tag{5.1}$$

where n_{TP} is the number of true positives and n_{FN} is the number of false negatives. The specificity, which is the ability of a method to detect negative results, is defined as

$$\text{specificity} = \frac{n_{\text{TN}}}{n_{\text{TN}} + n_{\text{FP}}} \tag{5.2}$$

where n_{TN} is the number of true negatives and n_{FP} is the number of false positives.

5.2.3 Results

Figure 5.4 shows the sensitivity-specificity plot for detecting small effect sizes when using an acquisition protocol with only 15 gradient directions and an SNR of 10. The sensitivity-specificity values for f and d are displayed in separate plots, although the results for all methods and group sizes are shown simultaneously. As we can see from this figure, when the number of diffusion-encoding directions used is modest and the data are noisy, none of the methods are particularly sensitive to changes in f, even for the larger group sizes we consider. When the group size is 5, the sensitivity of all methods is below 0.1, indicating that none of the techniques are able to identify significant differences between the subjects. As the group size increases, the sensitivity of all methods also increases, apart from TBSS, which never rises above 0.1 even when using all 25 subjects. The sensitivity is highest when the groups contain all 25 subjects, with maximum values of 0.5 when using 10 ROIs, 1 when using 3 ROIs, and 0.56 for both the RVM and smooth analysis. The performance of the 10 ROI, RVM and smooth methods are very similar, however we note that for some group sizes (i.e. N_{subj}=5,10,15), the smooth splines perform slightly better than the RVM, whereas for a group size of 20, the RVM performs slightly better.

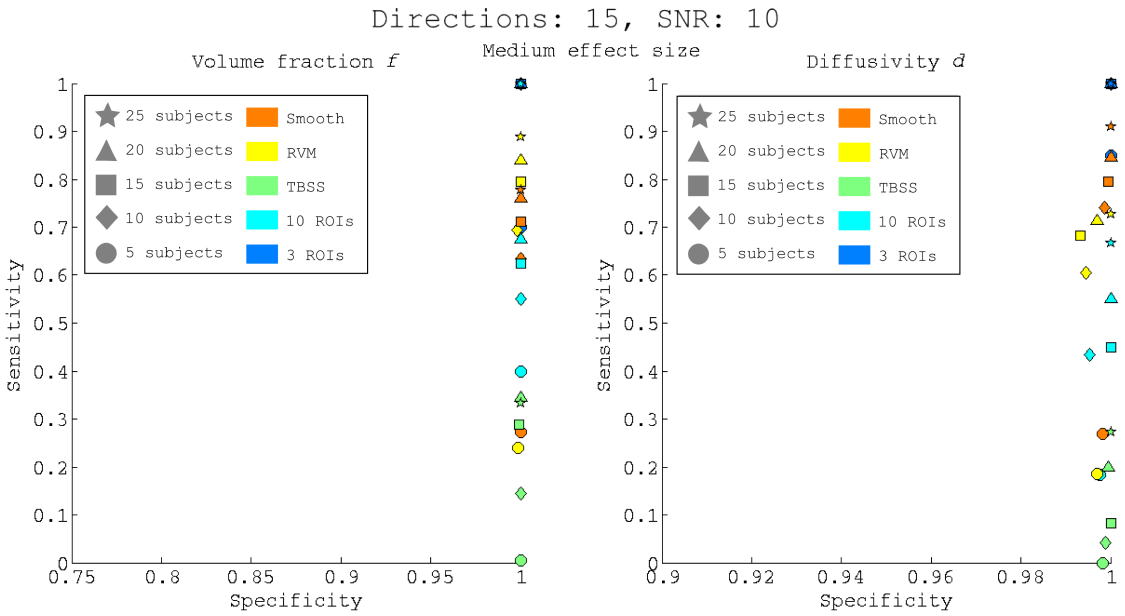

Figure 5.5: As for figure 5.4, but for detecting medium effect sizes.

When analysing changes in d, the sensitivities are higher for all methods. ROI analysis with only three regions is able to correctly identify the significant changes with a sensitivity of 1 for group sizes of 10 and above. TBSS is still the least sensitive method to detect changes between the groups, with the values for sensitivity ranging between $0.01 - 0.2$ as the group size increases from 5 to 25. Except for the group size of 5, the RVM and smooth splines are more sensitive than 10 regions of interest. We also note that, unlike for f, for groups with N_{subj} >5 , the RVM is slightly more sensitive to the underlying changes than the smooth splines, although the performance of the two methods is still very close.

Figure 5.5 shows how the sensitivities change as the effect size increases to the medium effect level. The patterns we observe for d are very similar, however for f the sensitivities increase markedly. Except for the smallest groups, the three region of interest method is able to identify all significant changes. Once again, for both f and d, TBSS is the least sensitive method for all group sizes. Both the RVM and smooth analysis are more sensitive than using 10 regions of interest, except for the smallest group sizes. The effectiveness of the RVM and smooth splines is very similar, however for group sizes larger than 5, the RVM is slightly more sensitive to changes in f, whereas the smooth splines are slightly more sensitive to changes in d.

The sensitivity-specificity plots for both 15 directions, $\mathrm{SNR} = 25$, and 60 directions, $\mathrm{SNR} = 10$, follow very similar trends. We do not display the results here, but briefly describe the key features of the plots. For all effect and group sizes, ROI analysis with only three regions performs well and is generally able to identify all regions of significant differences between the groups. TBSS analysis performs the worst, although the sensitivity increases with both effect and group size for both f and d. For example for 15 gradient directions and $\mathrm{SNR} = 25$, when the effect size is small, the sensitivity ranges from $0.03 - 0.67$ for f and from $0.03 - 0.4$ for d as the group size increases from 5 to 25. When

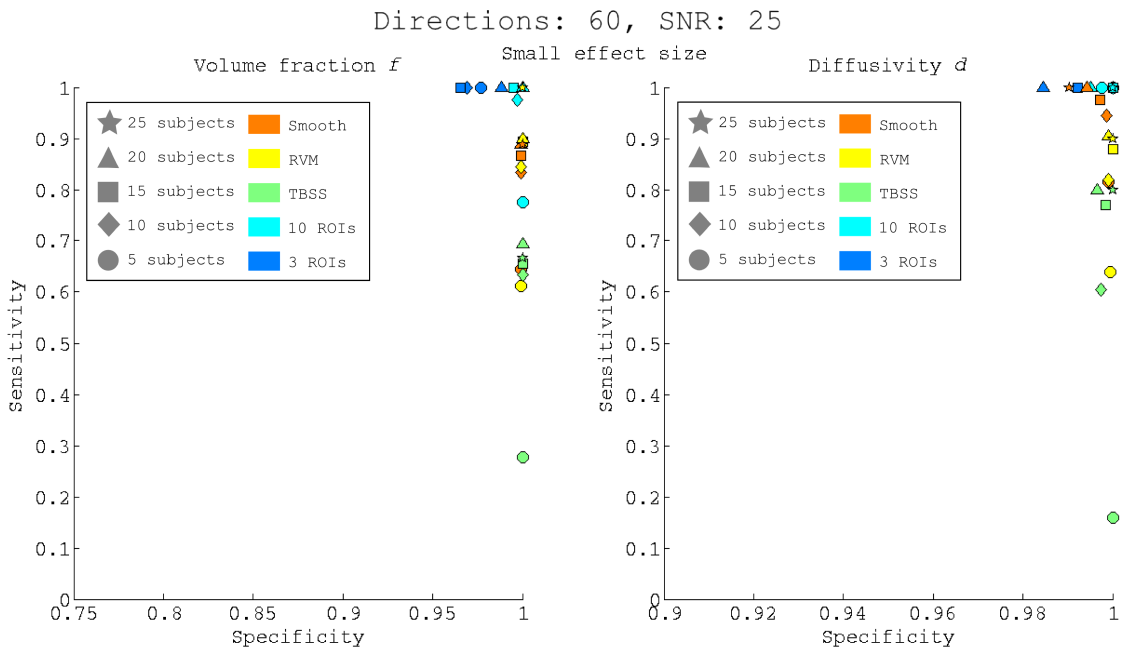

Figure 5.6: As for figure 5.4, but using an acquisition protocol with 60 gradient directions and an SNR of 25.

investigating large effect sizes, the corresponding sensitivity ranges are $0.37 - 0.83$ for f and $0.36 - 0.9$ for d. When the effect size is small, ROI analysis with 10 regions, the RVM and smooth analyses all give similar sensitivities. As the effect size increases, the RVM and smooth methods begin to outperform 10 regions of interests. Once again we note that for a given group size, the RVM outperforms smoothing at identifying changes in f, whereas the opposite is seen when looking for changes in d. However, the performance of both methods is very similar.

Figure 5.6 shows the sensitivity-specificity plots when we investigate small effect sizes using a high number of diffusion-encoding directions with little noise. Unlike the previous figures, the results are more clustered towards the top corner of the plot, implying that provided that the quality of the data is high enough, all methods are sensitive enough to detect subtle changes in f and d, even for small sample sizes. The only anomalous point occurs for the TBSS analysis with a group size of 5, which is still unable to detect many of the significant underlying changes. As the effect size increases, the results become even more densely clustered towards high sensitivities and specificities.

Finally we note that the specificity of all methods is high, regardless of the acquisition protocol, noise level, effect size and group size. The lowest specificities are in the range $0.7 - 0.75$ and are observed for the ROI analysis methods only. The lower specificities arise for all effect sizes, when either the number of gradient directions is low but the SNR is high, or the number of gradient directions is high but the data are noisy.

5.2.4 Discussion

From the results presented above, we can draw the following broad conclusions about the effectiveness of the five different comparison methods over a range of group sizes, effect sizes and data qualities. First, three large regions of interest are generally most sensitive to the underlying changes in f and d, followed by the RVM and smooth splines. These two methods perform very similarly, although we note that the RVM is slightly more sensitive to changes in f whereas the smooth splines are slightly better at detecting changes in d. For nearly all group and effect sizes, using these smooth functions to represent the underlying microstructure is consistently better than using 10 small ROIs, which in turn always outperforms the TBSS-style analysis. These key observations are most obvious when the data quality is low and the effect size is small. However, the general trends can still be seen for high quality data and large effect sizes, although the sensitivity and specificity points are more closely clustered together.

These observations can be understood more thoroughly if we consider how the sensitivity (also referred to as statistical power) is related to the sample size, effect size and variance of the measurements. As a general rule of thumb [180], in order to detect group differences with a sensitivity of 0.8, the sample size n should be

$$n = \frac{16\sigma^2}{\delta^2}, \tag{5.3}$$

where σ^2 is the common variance of the observations and $\delta = \mu_1 - \mu_0$ is the difference between the two population means. Therefore, to detect a fixed effect size δ, the sample size required depends directly on the variability of the observations. In the simulations, the intra-group variance is designed to be very low, therefore the main source of variability in our observations comes from how well we estimate the microstructure parameters from the noisy synthesised data. As shown in chapter 4, the variance of parameter estimates obtained using the RVM or smooth Bayesian P-splines is much lower than those estimated using independent voxelwise model fitting, which forms the basis of the ROI and TBSS analysis methods. Due to this, we would expect these methods to be able to detect changes in smaller group sizes with reasonable power. ROI analysis also performs well as it reduces the effective variance of the microstructure estimates by averaging over a large number of voxels. This is more pronounced when a larger number of voxels is included in the ROI, i.e. when we use three large regions rather than ten small ones, which explains why three regions of interest is more sensitive than ten. Finally, TBSS uses the raw voxelwise estimates. Due to the high variance of the microstructure parameters, particularly when the data quality is low, TBSS struggles to detect the true underlying differences in microstructure particularly when the group sizes are small.

However, there are a few caveats to the conclusion that the 3 ROI analysis performs best. First, despite the sensitivity of the technique to underlying microstructure indices, it is unable to localise the changes or determine exactly how the tissue parameters are changing. Second, the sensitivity of the ROI analysis may be artificially inflated by the overall experiment design. As shown in figure 5.1, the changes in both f and d occur across the whole of the first region, making it potentially easier to identify. If the experiment was repeated with the changes occuring in a much narrower region, we are likely to find that the sensitivity of the ROI analysis is greatly reduced. The spatial resolution of an ROI

analysis can be improved by dividing the region into smaller sections, such as in the 10 ROI scheme. However with fewer voxels to average over, the population variance may be larger, resulting in reduced sensitivities. Both schemes for dividing the ROI, either into 3 large regions or 10 smaller regions, also potentially suffer from edge effects. This can mask underlying parameter differences, particularly for focal parameter differences located on the boundary between two neighbouring ROIs.

As noted above, the sensitivity of the RVM and smooth analysis is very similar across all synthetic data sets, effect sizes and group sizes, although the RVM is slightly better at identifying changes in f whereas the smooth splines are slightly more sensitive to changes in d. In general, for the cases we consider here, we do not identify any significant advantage to using the RVM framework over smoothing the voxelwise parameter estimates. However, as discussed in chapter 4, the accuracy and precision of the smoothed Bayesian P-splines is highly dependent on the quality of the independent voxelwise estimates. These tend to become unreliable when the SNR is lower than 10, so if the experiment was repeated using a wider range of synthetic data sets, we may be able to identify cases in which the RVM is preferable to the smoothed splines. Regardless of whether the RVM or post parameter estimation smoothing method is used, the key benefits of a spline-based approach are that the splines have higher spatial resolution than ROIs and are not affected by edge effects, whilst also retaining the same performance level as ROI analysis in terms of sensitivity and specificity to the underlying changes.

The TBSS-style method that we implement here is unable to detect many of the significant underlying microstructure differences that exist between the populations using a standard voxelwise t-test. However, it is typically recommended to use threshold free cluster enhancement (TFCE) [165] when performing a TBSS analysis. Cluster-based thresholding, which finds groups of significant voxels and assigns a p-value to each cluster based on its overall size, is generally considered to be more sensitive than voxelwise analysis [165]. However, we anticipate that even if we used TFCE to determine p-values for TBSS, the RVM and the smooth splines, the smooth functions would still be more sensitive. This is because, due to the reduced variance of the parameter estimates, the clusters are less likely to be interrupted by noisy voxels, resulting in larger, and therefore more significant clusters. However, this does not mean that TBSS should be dismissed entirely. Provided that the data quality is reasonable and the group size large enough, it is reasonably sensitive and specific to changes in both f and d. Out of all the techniques considered here, TBSS provides the best spatial resolution when localising changes. It is also much more automated than ROI, RVM or tract-based analysis and allows exploratory rather than hypothesis driven studies of white matter tissue changes.

There are several limitations to the experiment, which could be accounted for in future work in order to provide a more thorough assessment of how well these different methods detect population differences. First, the groups of subjects used in the simulations are very homogeneous. This was done to ensure that the significant regions of difference between groups were consistent for different effect and group sizes. This might also account for the high specificities calculated for all methods. However it is not particularly realistic and in future we could ensure higher intra-group variablity.

Second, we only investigated simultaneous changes in both f and d over a large region. As we

discussed earlier, this could lead us to overestimate the sensitivity of some methods, in particular the ROI analysis. We have already mentioned that this could be improved by reducing the range over which the microstructure parameters vary and by considering changes in one parameter only.

The sensitivity and specificity analysis is potentially problematic as we are not necessarily comparing the same number of observations between the different methods. For example, for the ROI analysis, we are only counting the number of true positives or negatives over 3 or 10 regions, whereas the TBSS, RVM and smooth analyses are counted over 51 unique voxel locations. This means that one false positive or negative has a much greater effect on the sensitivity and specificity for the ROI analyses. In order to provide a more balanced assessment of the techniques, in future we can investigate alternative methods for calculating these types of statistics, that ideally normalise the results according to the total number of observations.

Next, we must be careful when comparing p-values, as the different population analysis methods we consider use the data in different ways, meaning that the p-value from one technique is not necessarily directly equivalent to the p-value obtained using a different technique. However, as long as this is taken into consideration when analysing and interpreting the results, we still believe it is valuable to compare the methods in this way. If there is a true underlying microstructure difference between two groups due to either pathology or development, ideally all of these techniques shoud be able to identify it. Therefore, it is important to understand in simulation what the strengths and limitations of these methods are before applying them to real patient data which are used to draw conclusions about the nature of certain pathologies.

Finally, we note that the RVM framework used to fit the microstructure parameter variation functions is Bayesian, and therefore, it would be possible to incorporate the group comparison directly into the framework. We could fit the parameter variation curves over the groups simultaneously, and sample the posterior distribution over the difference in the population means. We chose to perform frequentist t-tests in this experiment in order to provide a direct comparison between existing methods such as ROI analysis or TBSS. However, performing a Bayesian population study may also highlight advantages of the RVM over the smooth splines because, as discussed in chapter 4, the confidence intervals of the posterior distributions tend to be narrower for the RVM than the smoothed splines, even when the quality of the data is high.

5.3 Experiment 2: Investigating healthy ageing in the human corpus callosum

In the previous experiment, we showed in simulation that the parameter curves estimated using the RVM are able to detect and localise microstructural changes between populations, even when the effect is subtle and the quality of the data is low. We also showed that it performs better than current state of the art group comparison techniques such as TBSS, particularly when the sample size of the groups is small. In this experiment, we demonstrate the use of the RVM curves for in vivo human diffusion MRI group studies, specifically to investigate how microstructural indices in the mid-sagittal CC change

during healthy ageing. We also compare the results obtained using the RVM to standard ROI, TBSS and smooth analyses.

As mentioned in chapter 2, the myelination of white matter tracts in the brain begins before birth and continues throughout childhood and adolescence, peaking in early adulthood [187]. However as we age, the myelin is damaged as the sheaths split apart and cavities appear within the layers [139], and neurons and their axonal connections decay leaving debris amongst the tracts. Changes to the white matter tissue during both myelination and ageing affect the dispersion of water molecules within the brain and should be reflected in the microstructural tissue indices we measure with diffusion MRI. Previous studies such as [132] have shown the dependence of DTI parameters such as FA and MD on age by calculating correlations between the indices over a wide spectrum of ages. In this study we aim to detect changes in the ball and stick model parameters, f and d, between three distinct age groups rather than a continuous spectrum of ages. We consider subjects aged $20-29$, $40-49$ and $60-69$. We would expect the differences between the $20-29$ and $60-69$ year groups to be the largest and therefore the easiest to identify from comparing the diffusion MR microstructure indices. However, the changes between ages $20-29$ and $40-49$, and between $40-49$ and $60-69$ are likely to be more subtle due to the shorter timescales for myelin damage to occur.

5.3.1 Hypothesis

Given the results of the previous simulation studies, we hypothesise that

1. when comparing age groups $20-29$ and $60-69$, for which we anticipate the microstructure changes are significant, all methods are capable of identifying group differences.

2. when comparing age groups $20-29$ and $40-49$, and age groups $40-49$ and $60-69$, for which we anticipate more subtle microstructure changes, the RVM and smoothing methods are more suitable for detecting microstructure differences compared with the traditional ROI or TBSS methods.

5.3.2 Methods

In this section we discuss data acquisition and subject selection, the pre-processing steps taken to identify and segment the ROI (in this case, the mid-sagittal CC), as well as the model fitting procedures and statistical analysis methods used.

Data acquisition

The data used in this study were acquired as part of the Information eXtraction from Images (IXI) project (EPSRC GR/S21533/02). The data are freely available at www.brain-development.org. MR images were acquired from a total of 600 healthy subjects, using multiple image contrasts (T_1,T_2, pd, magnetic resonance angiography (MRA) and diffusion-weighting). The images were acquired at three separate sites on three different MR scanners (Philips 3T at Hammersmith Hospital, Philips 1.5T at Guys Hospital, GE 1.5 T at Institute of Psychiatry), although the imaging protocols were identical. We only consider subjects ($N = 175$) whose data were acquired on the Philips 3T scanner at Hammersmith Hospital, as the effects of different scanners (particularly those with different field strengths) could confound results.

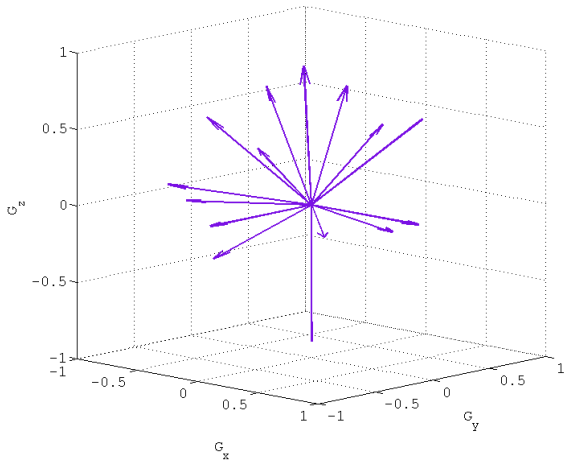

Figure 5.7: Vector plot showing the gradient directions used in the diffusion-weighted images in the IXI study.

.

Although the $T2$-weighted images can be used to detect white matter changes due to healthy ageing, we focus on the diffusion-weighted images only. The diffusion-weighting imaging protocol consists of 1 $b = 0$ image and 15 diffusion-weighted images acquired along non-collinear gradient directions, with $b = 1000$ s mm^{-2}. A vector plot of the applied gradient directions is shown in figure 5.7. As previously stated, we consider three distinct age groups: $20-29$, $40-49$ and $60-69$, which was the oldest group containing a reasonable number of subjects. From the 175 available subjects, there were 42 in the age range $20-29$, 25 in the age range $40-49$, and 39 in the age range $60-69$. However, when inspecting the data for each subject, we found regions of signal dropout in the splenium in the $b = 0$ images for a significant number of the $60-69$ year-old subjects, potentially due to motion, which resulted in the diffusion-weighted signals being significantly larger than the unweighted signals. Because only one $b = 0$ image was acquired as part of the protocol, we could not correct or compensate for this. Thus these subjects were excluded from the analysis, leaving only 15 (5M, 10F) viable data sets in the $60-69$ age group. This problem was also observed in the $20-29$ and $40-49$ age groups, although not to such a large extent. For consistency, we limited the sample size to 15 for all groups, which were randomly selected for both the $20-29$ (6M, 9F) and $40-49$ (5M, 10F) age groups. The SNR of the data in the CC in the $b = 0$ image, which we determined using a region of interest in the background, varied between 5 and 9.

Data pre-processing

The pre-processing steps used to segment the mid-sagittal CC and fit the shape model are similar to those described in section 4.3.2. However, due to the unique way that each individual subject lies whilst in the MR scanner, in most cases the mid-sagittal plane is not perpendicular to the slice direction. Therefore, in order to identify which voxels lie within the mid-sagittal slice, we use the algorithm described by Ardekani et al [12] to automatically detect the mid-sagittal plane. The algorithm does this by finding

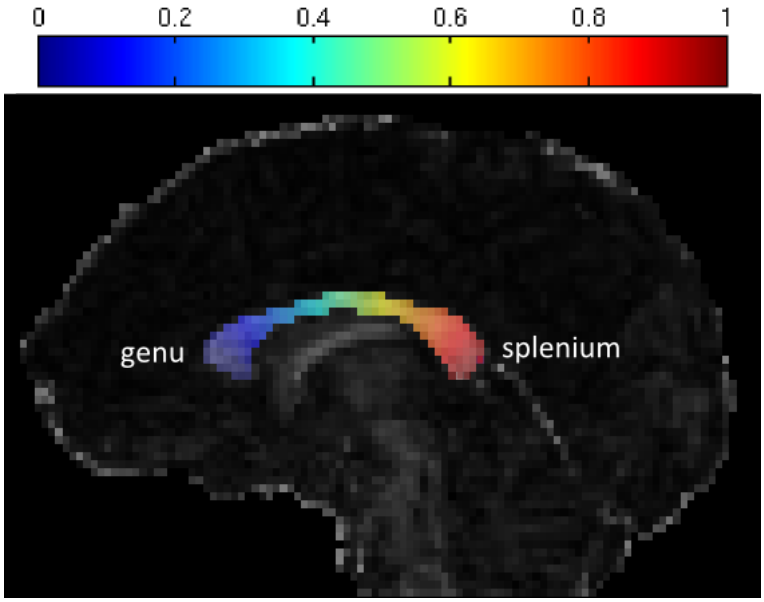

Figure 5.8: The segmented CC from a representative subject. The voxels within the region of interest are colour-coded according to position within the CC, from genu to splenium.

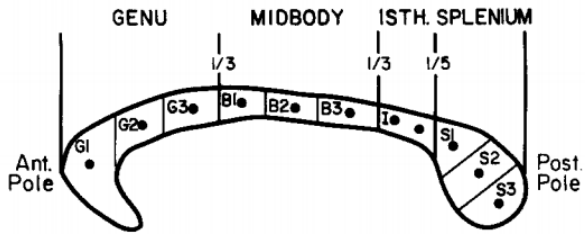

Figure 5.9: Proposed division of the CC into ten anatomically specific regions, taken from [2].

the plane within the image that maximises symmetry, using cross-correlation as a measure of symmetry. Instead of using linearity and planarity to identify the CC, we use the FA image with a threshold of 0.3, before performing connected component analysis and fitting the cm-rep shape model. Figure 5.8 shows a segmented CC, colour-coded according to voxel position along the medial axis, for a representative subject.

Model fitting

We fit the ball and stick model to each subject's data as described above in section 5.2.2, using both the RVM and independent voxelwise parameter estimation. We also fit Bayesian P-spline functions to the voxelwise parameter estimates to use in the smooth group comparison.

Group comparison

We use the same group comparison techniques and statistical tests discussed in section 5.2.2. However when using the ROI analysis with 10 regions, instead of segmenting the CC equally we use the more anatomically specific division proposed by Aboitiz et al [2], shown in figure 5.9. We compare f and d

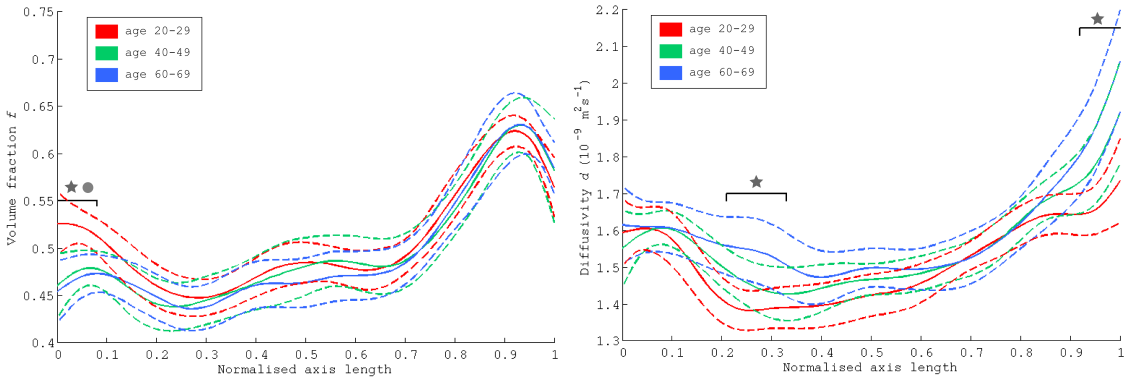

(a) (b)

Figure 5.10: Figure 5.10a shows the mean trends for volume fraction f over the three subject groups fitted using the RVM, with 95% confidence intervals indicated. Regions of significant difference, with p-values corrected for multiple comparisons, are marked: ● represents a significant difference between ages $20-29$ and $40-49$ at the 5% level; ■ represents a significant difference between ages $40-49$ and $60-69$; ★ represents a significant difference between ages $20-29$ and $60-69$. Figure 5.10b shows the mean trends for diffusivity d for the three groups. Regions of significant differences at the 5% level are marked as for f.

between

- subjects aged $20-29$ and subjects aged $40-49$.
- subjects aged $40-49$ and subjects aged $60-69$.
- subjects aged $20-29$ and subjects aged $60-69$.

For each comparison we use one-sided t-tests to determine whether the microstructural indices are greater or smaller in the first group compared to the second group. As described previously, we correct for multiple comparisons using the Westfall-Young permutation test.

5.3.3 Results

Figures 5.10, 5.11, 5.12, 5.13 and 5.14 show the results of the group comparisons between f and d in populations with ages $20-29$, $40-49$ and $60-69$, using the RVM, smooth splines, ROI analyses and TBSS respectively.

From figure 5.10, we see that when using the RVM to fit profiles of f and d to each subject, we find four regions of significant difference between the groups using pointwise testing. In the genu of the CC, at a distance of $0-0.07$ along the normalised medial axis, we find that f is significantly greater in the $20-29$ age group than in both the $40-49$ and $60-69$ age groups. There are no significant differences in f between the $40-49$ and $60-69$ age groups. When comparing diffusivities, we only find differences in d between the $20-29$ and $60-69$ age groups. These differences occur between $0.21-0.33$ (corresponding to the posterior genu) and $0.92-1$ (corresponding to the posterior splenium)

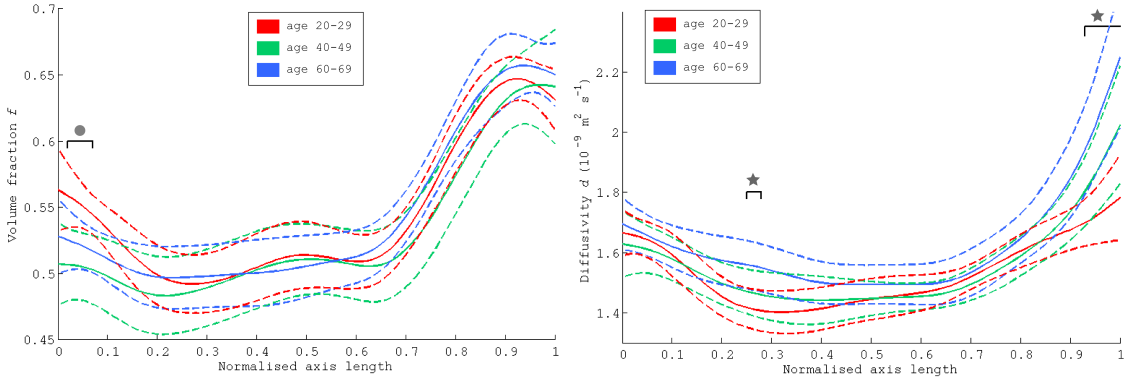

(a) (b)

Figure 5.11: Figure 5.11a shows the mean trends for volume fraction f over the three subject groups fitted using smooth splines, whereas figure 5.11b shows the mean trends for diffusivity d for the three groups. Regions of significant differences at the 5% level are marked as for figure 5.10

along the normalised medial axis length. In both cases d is significantly higher in the $60-69$ group than in the $20-29$ age group.

Figure 5.11 shows the results of pointwise testing after using splines to smooth the independent voxelwise estimates of f and d. The curves for f and d in figure 5.11 exhibit the same general trends as those shown in figure 5.10 for the RVM, although they appear smoother. This is most likely due to the poor SNR of the data, which results in noisy independent voxelwise estimates of both f and d. As discussed in chapter 4, fitting splines to model the trends in independent voxelwise parameters in difficult when the data are noisy, often resulting in some bias or oversmoothing. However, the regions of significant differences between the groups are generally consistent with those found using the RVM, with one exception. We no longer observe significantly higher values of f in the anterior genu of the $20-29$ age group compared with the $60-69$ age group. This region is still significant when comparing the $20-29$ group with the $40-49$ group, although the range of the significant region is slightly reduced, spanning only $0.02-0.07$ of the normalised medial axis length. The significant differences that we find for d are consistent with the results obtained using the RVM, although as for f, the size of the significant regions is reduced. This is particularly noticeable in the posterior genu, where d in the $60-69$ group is only significantly higher than the $20-29$ group at distances of $0.25-0.28$ along the medial axis, but is also seen in the posterior splenium, where the range of significance is slightly reduced to $0.93-1$ along the medial axis length.

Figures 5.12 and 5.13 show the results of ROI analysis of the three groups, using 3 and 10 regions of interest respectively. When using only 3 large regions of interest, we observe three regions of significant differences. We find increased f in the genu of the $20-29$ age group compared to the $40-49$ year olds, which is consistent with the results of the RVM and smooth experiments. We also observe increased diffusivity in the splenium in the $60-69$ group compared with the $20-29$ group, which is also consistent

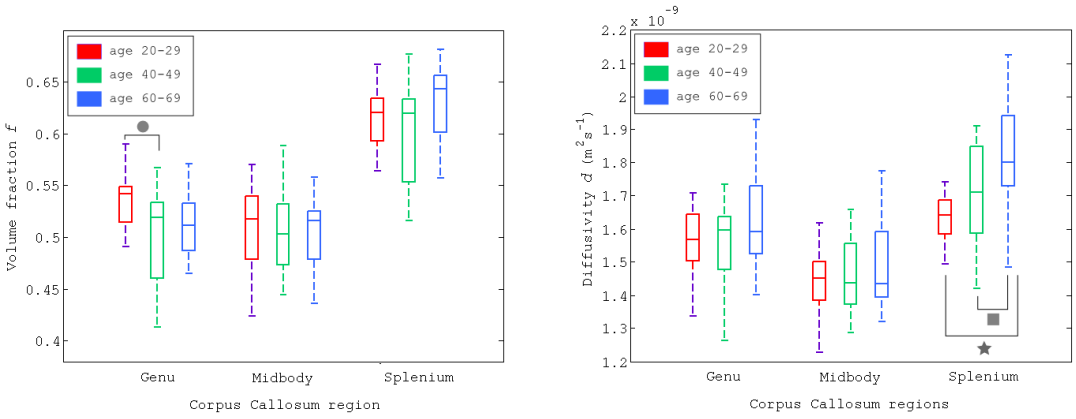

(a) (b)

Figure 5.12: Figure 5.12a shows the mean values for volume fraction f for the three subject groups calculated in 3 ROIs, whereas figure 5.12b shows the values for diffusivity d. Regions of significant differences at the 5% level are marked as for figure 5.10

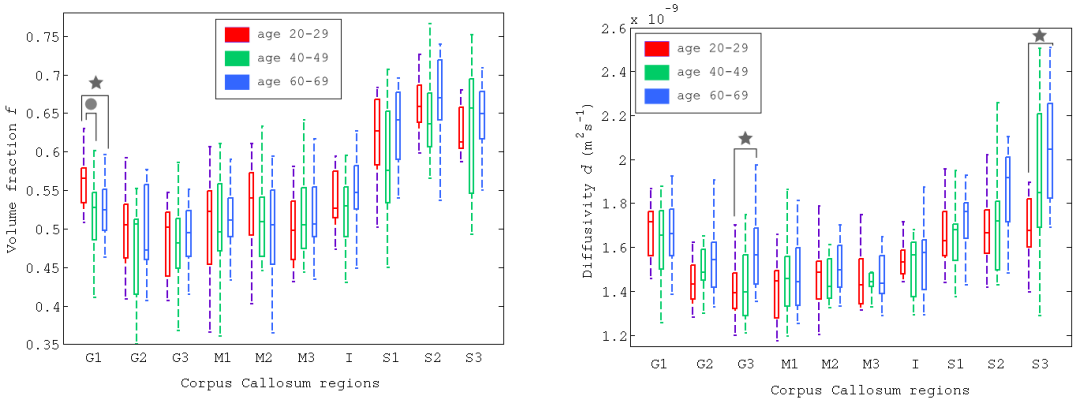

(a) (b)

Figure 5.13: Figure 5.13a shows the mean values for volume fraction f for the three subject groups calculated in 10 ROIs, whereas figure 5.13b shows the values for diffusivity d. Regions of significant differences at the 5% level are marked as for figure 5.10

with the previous experiments. We also find an increased diffusivity in the splenium of the $60-69$ group compared to the $40-49$ group, although this is not observed in any of the other group comparison results. When using 10 smaller regions, we find four regions of significant difference which are consistent with the RVM results. These are significantly higher f in G1 for the $20-29$ group compared to both the $40-49$ and $60-69$ age groups, and significantly higher diffusivity for the $60-69$ age group compared with the $20-29$ group in both G3 and S3.

Finally, figure 5.14 shows the regions of significant difference identified when using TBSS. The rows show the comparisons between ages $20-29$ and $40-49$, $20-29$ and $60-69$, and $40-49$ and $60-69$ respectively. The first column shows the differences in volume fraction f whereas the second column shows changes in diffusivity d. Unlike all of the previous methods, TBSS does not detect any significant changes in f between the $20-29$ and $40-49$ age groups. When comparing ages $20-29$ and $60-69$, TBSS does identify 3 voxels on the skeleton for which the younger group has a higher value of f, although the positioning of the voxels is not necessarily consistent with the other methods which identify the changes at the tip of the genu. TBSS does not identify any significant changes between the $40-49$ and $60-69$ age groups which is consistent with the other group comparison methods. When considering changes in diffusivities, TBSS finds a single voxel in the posterior splenium for which the $40-49$ year old subjects have significantly higher diffusivities than the youngest subject group, which is not observed using the other comparison techniques. The significant differences in diffusivities identified by TBSS between the $20-29$ and $60-69$ year olds are more consistent with the RVM, smooth tract-based and ROI results. A small number of voxels in both the posterior genu and posterior splenium are found for which the diffusivities in the older group are significantly higher than those for the youngest group. This matches the results from the RVM, smooth and 10 ROI analyses, although we also observe some individual significant voxels throughout the midbody which are not observed in the other experiments. Finally, TBSS identifies significant differences in d between the $40-49$ and $60-69$ age groups. As we can see from figure 5.14f, in these voxels, which are scattered around the posterior genu and midbody, d is higher for the older subjects relative to the $40-49$ year group. These changes are not observed in the other experiments. The only other method to observe changes in d between these two groups is the ROI analysis with 3 regions, although in this case, the changes are observed in the splenium only.

5.3.4 Discussion

In this experiment we have tested various methods for comparing microstructure indices across three populations, including the RVM, smooth spline fitting, ROI analysis and TBSS. In some cases, the results we obtained are consistent amongst the methods. For example, all methods find an increased diffusivity in the splenium in the $60-69$ group compared to the $20-29$ group, and many of the methods find increased f in the genu of the $20-29$ compared with the $40-49$ group (RVM, smooth, 3 ROIs, 10 ROIs), as well as increased d in the posterior genu of the $60-69$ age group compared with the youngest subjects (RVM, smooth, 10 ROIs, TBSS). Whilst we did observe consistent differences between ages $20-29$ and $40-49$ and $20-29$ and $60-69$, we did not observe plausible differences between the $40-49$ and $60-69$ age groups, suggesting that either the majority of changes to white matter

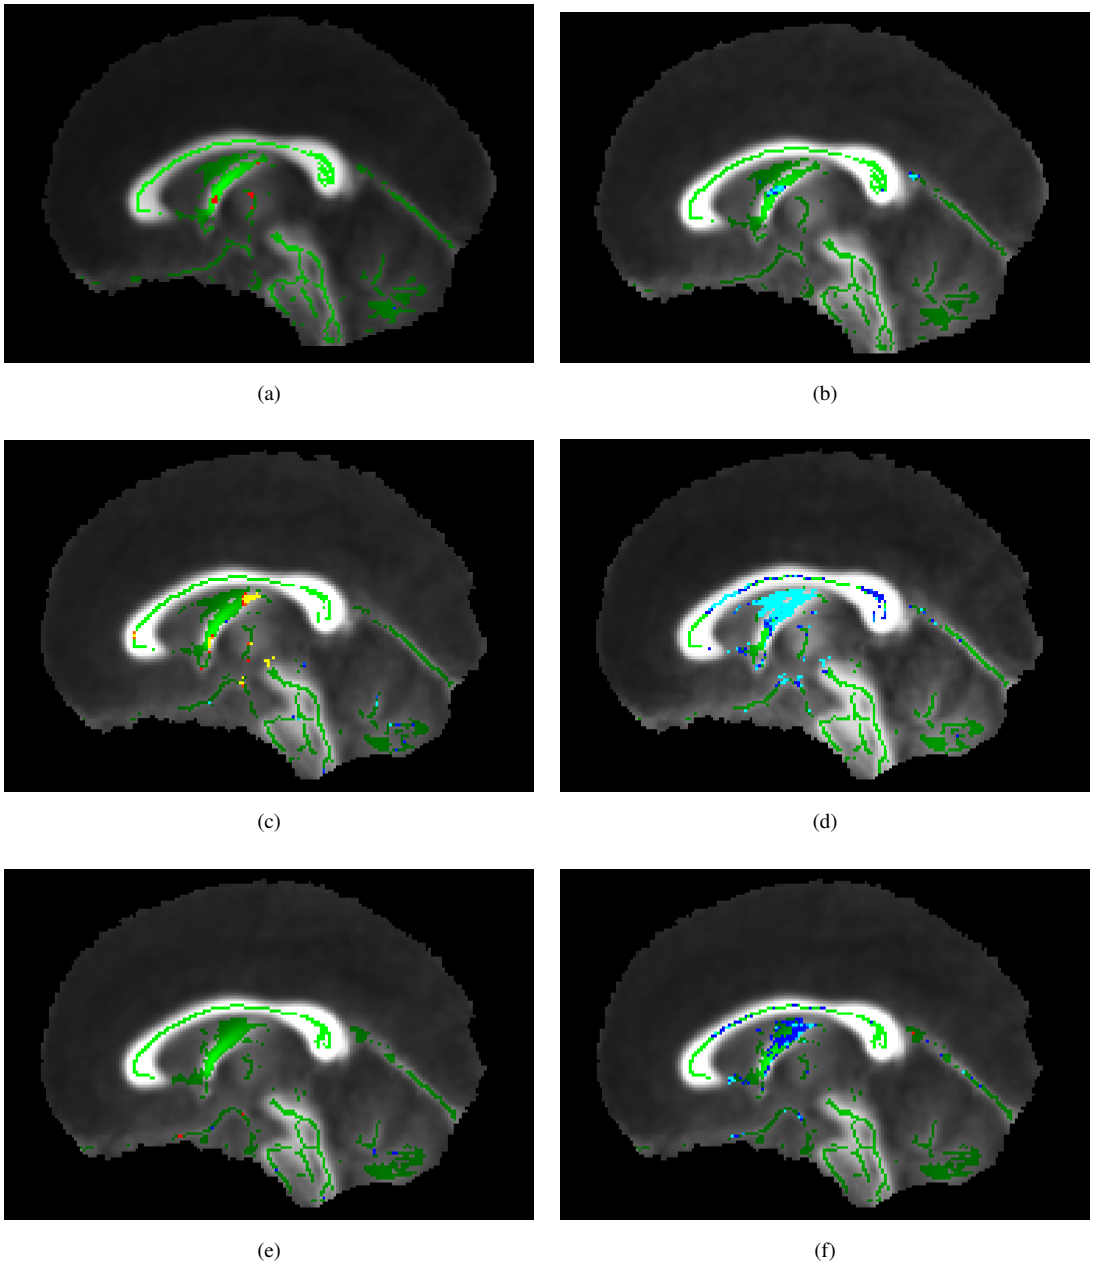

(a) (b)

(c) (d)

(e) (f)

Figure 5.14: Figure 5.14a shows regions of significant difference (at the 5% level) in volume fraction f between ages $20-29$ and $40-49$ using TBSS overlaid on the mean FA images of all subjects. The mean skeleton is shown in green. Red and yellow voxels indicate that f is significantly greater in the younger group; blue voxels indicate that f is significanly lower in the younger group. The darker the voxels, the more significant the change. Figure 5.14b shows the regions of significant difference in diffusivity d between the same groups. Figure 5.14c and 5.14d show the regions of significant difference in f and d between ages $20-29$ and $60-69$ respectively. Figure 5.14e and 5.14f show the regions of significant difference in f and d between ages $40-49$ and $60-69$ respectively. The colour-coding of the significant regions is the same for all figures.

microstructure in the CC have already occured by the age of around 50, or that the changes that do occur are too subtle to be identified in the IXI data sets. Although we have no ground truth to compare these results to, based on previous ageing studies these results are reasonable. For example, [1, 132, 140], which cover similar age ranges to our experiment, all show that FA decreases and MD increases with age in the anterior parts of the CC, i.e. the genu. We also observe an increased diffusivity in the splenium with increasing age. Whilst this is generally not observed in ageing studies in the CC, it has previously been reported in the IXI cohort by Robinson et al [155].

However, there are also several inconsistencies between the different methods. For example, we identify an increase in f in the $20-29$ group compared to the $60-69$ group when using the RVM or 10 ROIs, and we only find increases in diffusivity in the $60-69$ group compared with the $40-49$ group, albeit in different locations within the CC, when using TBSS or 3 ROIs. With no ground truth, we cannot calculate sensitivity and specificity indices as we did in the simulation experiment in order to determine which methods perform best. However, we can assess the plausibility of the microstructure changes by considering how methods support a given result. For example, the RVM and 10 ROI methods identify the same microstructure changes in the same locations within the CC. Out of these four changes, one of them, the increase in f between ages $20-29$ and $40-49$ in the genu, is only found by these two methods. However, it can still be considered a plausible change, given that it is identified by more than one technique. The other three changes are corroborated by at least one other method, meaning that the consensus for these tissue changes is strong. Despite the fact that they do identify the same parameter changes in the CC, the RVM is still advantageous over the 10 ROI approach as it provides more detail about how the microstructure parameters are changing on a finer scale, without the hard boundaries of the ROI approach.

The smooth method identifies three regions of significant change amongst the groups, all of which are consistent with the RVM and 10 ROI approaches. However, it does not identify a change in f between the $20-29$ and $60-69$ groups. Given the similarity of the RVM and smooth methods in the sensitivity-specificity analysis performed in section 5.2, this is unexpected. However, in the simulation experiments, the lowest quality data we synthesised using 15 unique gradient directions and a SNR of 10. Whilst the IXI data sets do use the same number of gradients, the data is actually noisier with SNRs of between 5 and 10. As shown in chapter 4, as the SNR falls to 5, the smooth method tends to become less accurate, which may explain why it did not identify a potentially subtle difference between the groups. If the quality of the data was higher, given the previous simulation results, we would expect the performance of the RVM and smooth methods to be closer. For both the smooth and RVM methods, the extent of the significantly different regions is quite small, only covering about 5% of the total curve length. At the $p = 0.05$ significance level, this could be attributed to random chance. However we do not think this is the case. First, the Westfall-Young multiple comparison correction calculates adjusted p values to take into account the number of comparisons we make. Second, points along smooth curves are not independent of each other and thus a cluster of significant points at the same location is more likely due to real underlying changes rather than chance, which should cause the significant points to be

more randomly distributed along the curves.

The 3 ROI method identifies three changes between the groups, only two of which are supported by alternative techniques. The third, an increase in d in the genu of the $60-69$ group compared to the $40-49$ group, is only identifed by this method, and therefore is less plausible. As we can see from figures 5.13 and 5.12, when we divide the splenium into smaller regions, although the mean diffusivities for the $40-49$ is always lower in each region, the overlap between the range of values is large enough for the difference to be statistically significant. However, when combined into a larger region the difference between the values seems larger. This highlights the importance of observing the changes in microstructure parameters over the region continuously, rather than dividing into arbitrary regions. However, performance may be improved by using a more anatomical parcellation as in [85, 192].

Finally, TBSS identifies 4 clusters of microstructure changes, only 3 of which are supported by other techniques. It also misses other changes, such as the increase in f in the genu of the $20-29$ year olds compared with the $40-49$ group. It does identify one increase in f, although the region it identifies is very small and does not spatially correspond to the changes found by the other methods. This is not necessarily surprising because, as demonstrated by our earlier simulation studies, TBSS appears to be particularly insensitive to changes in f for low quality data, regardless of effect size or group size. These simulation studies also show that TBSS is slightly more sensitive to changes in d, and it does identify a larger number of significant differences in diffusivity. As already mentioned, two of these are consistent with the other group comparison methods. The final significant region TBSS identifies (increased diffusivity in the genu and midbody for the $60-69$ group over the $40-49$ group) is not identified by the other methods. However, this region is not a large cluster but smaller individual voxels, and therefore may be due to noise. As we discussed in chapter 4, TFCE is often the preferred method for statistical analysis within TBSS. Because TFCE reduces the significance of small clusters, it would potentially get rid of the noisy significant voxels.

Improvements

There are several improvements we could make to this group study, some of which could improve all of the methods and some specifically to improve the RVM. First, we could improve the way in which we identify the mid-sagittal slice. This would improve not only the RVM method, but also the smooth and ROI analyses. The algorithm we currently use relies on the reflective symmetry of the brain across the mid-sagittal plane. This symmetry is usually observed in healthy patients, but for subjects who have significant pathology in one hemisphere of the brain e.g. a tumour or who have undergone resection surgery, this pre-processing step would most likely fail. There are two potential methods to overcome this problem. The first involves performing an initial DTI analysis and tractography on the scanner. The tractography visualisation aids in the identification of the mid-sagittal slice, and thus the slice direction can be set perpendicular to this. However, it does require extra time in the scanner which may not be possible, and it cannot be applied retrospectively to previously acquired data sets. The second method would be to use an alternative post-processing algorithm such as [146], which has been designed specifically to identify the mid-sagittal plane in subjects who have undergone brain resection surgery.

Currently the RVM is fit to a single slice of data which limits the number of tracts we can currently study. To overcome this, we would like to extend the RVM to model larger, three-dimensional tracts such as the whole of the CC or the corticospinal tract (CST). The framework is already in place to fit the RVM to larger sheet-like tracts using medial surfaces to model the overall shape of the tracts [195] and two-dimensional spline surface to model the microstructure variation. However, it has not been fully tested and, given the variable width of a tract such as the CST, careful thought would be needed to determine how many basis functions to use to describe the variation of the microstructure parameters in both directions. In the case of a tract like the CST, which fans out over its length from the spinal cord to the cortex, we may also want to alter the model slightly in order to account for the fibre dispersion, using an index of dispersion such as κ used by Zhang et al [200], which we discussed briefly in chapter 2.

Despite these current limitations, we believe that the RVM has future potential for performing group comparison studies, particularly when the parameter changes are subtle and the quality of the data is low. Therefore in future, we would like to use the RVM to further study the effect of healthy ageing by considering a wider range of ages. This will allow us to provide a more detailed description of microstructure changes with age. We also think that the RVM would be useful to investigate pathologies, particularly those with subtle microstructure changes such as schizophrenia. Finally, we believe that the RVM provides a definite advantage over other group comparison techniques when considering weak microstructure parameters such as axon radius, which are very sensitive to the effects of noise. This is investigated more thoroughly in the next chapter.

5.4 Conclusions

In this study we have shown how the RVM can be used to perform group comparison studies. Using simulation studies, we have compared the sensitivity and specificity of the RVM to alternative population study techniques such as ROI analysis, TBSS and smooth splines over a range of effect sizes, group sizes and data qualities. Overall we find that the RVM outperforms 10 small ROIs and TBSS in the majority of cases. Compared to 3 larger ROIs, the RVM is not necessarily as sensitive, although it does provide more detailed information on how the microstructure parameters vary spatially. However, within the range of data and group parameters we have investigated here, there appears to be little advantage in using the RVM over fitting smooth splines to the independent voxelwise parameter estimates. We also demonstrated the RVM in a group study using in vivo human brain data acquired from subjects within three distinct age groups. Compared to the alternative methods of group comparison, the RVM and 10 ROI analysis identify the largest number of consistent changes. In this case, due to the very low SNR of the data, the RVM provides a potential advantage over the smooth splines. The results presented here demonstrate the usefulness of the RVM for performing group studies, particularly when the quality of the data is low, and we hope to use the RVM to determine subtle microstructural changes in future studies into pathologies and brain development.

Chapter 6

Estimating direct microstructure using the RVM

6.1 Motivation

In previous chapters we introduced the regional variation model framework (chapter 3), validated the RVM against simulations and ex vivo monkey data (chapter 4) and demonstrated a potentially important application for the technique, namely population studies (chapter 5). However, in the population study performed, the parameters being compared across the groups, volume fraction and diffusivity, are derived from the simple ball and stick model [26] of diffusion in white matter. As discussed in section 2.4.2, although these indices are affected by numerous underlying tissue changes (such as axon size, axon density, permeability etc.), they are neither sensitive nor specific to any one tissue property. These parameters can still be useful as they allow us to differentiate populations; however they are limited as they cannot explain the specific tissue changes that cause the differences between groups. As noted in section 2.1.2, white matter diseases are primarily caused by damage and/or changes to the glial cells, the axons and the myelin surrounding them. Therefore, to gain a greater understanding of the mechanisms and progression of these diseases, we need to estimate more fundamental, direct tissue properties such as axon radius, axon density and glial cell size. In section 2.4.2, we highlighted a number of models and techniques that estimate an axon radius distribution or index, such as AxCaliber [15] and ActiveAx [7]. The ActiveAx technique can be implemented on clinical scanners, however estimates of the axon radius index in the human corpus callosum estimated from data acquired using G_{max}=60 mT/m are noisy and may not be precise enough to track the progression of white matter diseases.

A recent study using ActiveAx by Dyrby et al [60] showed that the accuracy and precision of estimates of axon radius index R depends on the maximum gradient strength used to acquire the data. When G_{max}=300 mT/m, the standard deviation of estimates of R is low. As G_{max} is decreased to 200 mT/m and 140 mT/m, estimates of R are consistent with those obtained at G_{max}=300 mT/m, but the standard deviation increases. At G_{max}=60 mT/m estimates of R are much greater than those with G_{max}=300 mT/m and the standard deviation is higher still. This dependence of R on gradient strength can be understood (at least partly) by examining profiles of the perpendicular component of the intra-axonal signal as a function of axon radius. Figure 6.1 shows these signal profiles for three different values

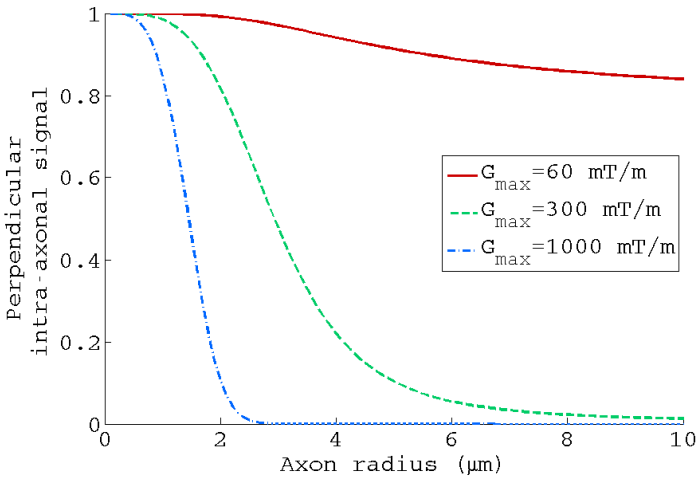

Figure 6.1: The perpendicular component of the intra-axonal signal as a function of axon radius for three different protocols with G_{max}=$\{60, 300, 1000\}$ mT/m, $\Delta = 20$ms and $\delta = 10$ms. For $R = 0 - 5$ μm, the range of axon radii most prevalent in the brain, as the maximum gradient decreases the slope of the signal as a function of R also decreases, making it more difficult to estimate the axon radius index accurately and precisely, particular when the data are noisy.

of G_{max}: 1000 mT/m, 300 mT/m and 60 mT/m, for $\Delta = 20$ms and $\delta = 10$ms. We can see that as the gradient strength decreases the signal curves (as a function of R) become flatter, particularly for the small axon radii which are most prevalent in the brain [2]. This makes it difficult to estimate the axon radius index in an individual voxel, particularly in the presence of noise which can lead to overestimation and large uncertainties. This is the primary motivation for using the RVM to estimate complex microstructure parameters such as *R*. The RVM exploits the underlying spatial coherence that exists in many white matter tracts and pools the data in a natural way that respects the underlying geometry of the tract. For voxels with similar underlying axon radius distributions, if the model of the diffusion MR signal is fit independently in each voxel, estimates of *R* could be very different depending on noise levels. However when using the RVM, the voxels are assumed to have similar values of *R* which negates some of the effect of the noise. In this chapter we perform three experiments to test the suitability of the RVM for estimating direct microstructure parameters. Specificially we

- perform an initial simulation study using a range of acquisition protocols from pre-clinical (G_{max}=300 mT/m) to clinical (G_{max}=40 mT/m). This provides a ground truth, allowing us to check whether the RVM can recover the expected trends in axon radius and to test the stability of the RVM over a range of spatial model settings, i.e. knot vectors and smoothing prior distributions, as in chapter 4,

- use the simulated data to evaluate differences in parameter estimation between the RVM (fitted using the optimal spatial model parameters determined in the previous step) and current state of the art techniques,

- estimate axon radius indices in the CC from high quality ex vivo monkey data, acquired at gradient strengths from G_{max}=300 mT/m to G_{max}=60 mT/m using both the RVM and current state of the art techniques in order to test the results of the simulation study in a realistic tissue substrate,

- test the performance of the RVM when estimating an axon radius index from in vivo human data acquired on typical clinical scanners, with maximum gradient strengths of 60 mT/m and 40 mT/m.

As in previous chapters, we focus on estimating microstructure indices within the corpus callosum.

6.2 Experiment 1: Validation using synthetic data

The purpose of this experiment is two-fold. The first aim is to use simulated data from a known ground truth in order to test the stability of the RVM with regards to the choice of knot vector and smoothing prior distributions when using the ZeppelinCylinder model of diffusion in white matter. As in chapter 4, we also aim to identify a suitable set of general purpose spatial model settings which can be used over a range of acquisition parameter settings and noise levels. This set of spatial model settings can then be used in later experiments. The second aim is to compare values of the axon radius index estimated from the simulated data using the RVM (with the best spatial model settings) to a known ground truth in order to evaluate the method quantitatively. We also compare the results from the RVM to estimates of axon radius index fit in each voxel independently. This allows us to compare the proposed RVM method to current state of the art techniques for axon radius estimation and identify circumstances in which the RVM performs better.

6.2.1 Hypothesis

In terms of the stability of the RVM, based on results obtained in chapter 4 we hypothesise that, due to the sensitivity of the axon radius index to noise, the results for R obtained using the RVM will be sensitive to the choice of spatial model parameters and smoothing priors (defined previously in section 3.3), particularly for low gradient strengths and SNRs, and may need careful tuning. Our results in chapter 4 showed that f is generally robust to the effects of noise, and therefore we expect the choice of spatial model parameters to have less influence on estimated volume fractions.

With regards to the performance of the RVM, we hypothesise that when high gradient strengths are used to acquire the data, estimates of the axon radius index from both the RVM and current state of the art methods should be comparable in terms of accuracy and precision. However, as the maximum gradient strength decreases, the sensitivity of the measured signals to axon radius also decreases. Therefore we hypothesise that for these low gradient data sets, the RVM is able to capture the general trends of axon radius variation with more accuracy and precision than independent voxelwise estimation. We also hypothesise that the RVM performs better than independent voxel estimation as the SNR of the acquired data decreases, even at reasonably high gradient strengths.

6.2.2 Methods

In this section we discuss the construction of the ground truth phantom, the choice of acquisition protocols for testing and data synthesis. We also describe the procedures for parameter estimation and indicate

the range of knot vectors and prior distributions we use when testing the robustness of the RVM.

The diffusion model

The RVM in this experiment uses the ZeppelinCylinder model [134], introduced in section 2.4.2. We now describe the model and present the model parameters and signal equations in full. The model comprises two distinct compartments: one intra-axonal and one extra-axonal compartment. The intra-axonal compartment comprises randomly packed cylinders of radius R and orientation $\mathbf{e}$. As for the ball and stick model, the vector $\mathbf{e}$ can be described according to two independent angles θ and ϕ. The diffusion MR signal due to this compartment A^{ia} can be decomposed into two parts. Parallel to $\mathbf{e}$, we assume that water molecules diffuse freely with diffusivity $d_\parallel$. This component of the signal $A^{ia}_\parallel$ is

$$A^{ia}_\parallel\left(d_\parallel,\theta,\phi\right)=\exp\left(-t\gamma^2\delta^2G^2\left(\mathbf{e}\left(\theta,\phi\right)\cdot\hat{\mathbf{G}}\right)^2 d_\parallel\right). \tag{6.1}$$

Perpendicular to $\mathbf{e}$ the water molecules are restricted by the boundaries of the cylinders. This component of the signal $A^{ia}_\perp$ is given by [170, 182]

$$A^{ia}_\perp\left(d_\parallel,R,\theta,\phi\right)=\exp\left[-2\gamma^2\delta^2G^2\left(1-\left(\mathbf{e}\cdot\hat{\mathbf{G}}\right)^2\right)\times\right.$$
$$\left.\sum_{m=1}^{\infty}\frac{2d_\parallel\alpha_m^2\delta-2+2\exp\left(-d_\parallel\alpha_m^2\delta\right)+2\exp\left(-d_\parallel\alpha_m^2\Delta\right)-\exp\left(-d_\parallel\alpha_m^2\left(\Delta-\delta\right)\right)-\exp\left(-d_\parallel\alpha_m^2\left(\Delta+\delta\right)\right)}{d_\parallel^2\alpha_m^6\left(R^2\alpha_m^2-1\right)}\right], \tag{6.2}$$

where Δ and δ are pulse sequence timing parameters defined in section 2.3.1, and α_m are the roots of the equation $J_1'\left(\alpha_m,R\right)=0$, where J_1' is the derivative of the first order Bessel function of the first kind. Combining equations 6.1 and 6.2 the total intra-axonal signal A^{ia} is

$$A^{ia}\left(d_\parallel,R,\theta,\phi\right)=A^{ia}_\parallel\left(d_\parallel,\theta,\phi\right)A^{ia}_\perp\left(R,d_\parallel,\theta,\phi\right). \tag{6.3}$$

The intra-axonal compartment accounts for a fraction f of the total signal. The remaining signal is due to the hindered diffusion of water molecules around the axons. This extra-axonal signal A^{ea} is modelled as a cylindrically symmetric tensor $\mathbf{D}^{ea}$ whose principal axis aligns with $\mathbf{e}$. We assume that in the extra-axonal space, water molecules diffusing parallel to the axons are unhindered and therefore have diffusivity $d_\parallel$. However, perpendicular to the axons the hindered water molecules have diffusivity $d_\perp$. We use a tortuosity approximation [172] to relate $d_\perp$ to $d_\parallel$ and f so that

$$d_\perp=d_\parallel\left(1-f\right). \tag{6.4}$$

The tensor $\mathbf{D}^{ea}$ can be written as

$$\mathbf{D}^{ea}\left(d_\parallel,f,\theta,\phi\right)=fd_\parallel\mathbf{e}\mathbf{e}^\intercal+\left(1-f\right)d_\parallel\mathbf{I} \tag{6.5}$$

where $\mathbf{I}$ is the identity matrix. The extra-axonal signal A^{ea} is then

$$A^{ea}\left(d_\parallel,f,\theta,\phi\right)=\exp\left(-t\gamma^2\delta^2G^2\hat{\mathbf{G}}\mathbf{D}^{ea}\hat{\mathbf{G}}\right). \tag{6.6}$$

The total signal A is the sum of the signals over the two compartments so that

$$A\left(d_\parallel,R,f,\theta,\phi\right)=fA^{ia}\left(d_\parallel,R,\theta,\phi\right)+\left(1-f\right)A^{ea}\left(d_\parallel,f,\theta,\phi\right). \tag{6.7}$$

The spatial model

The spatial variation of the ZeppelinCylinder model tissue parameters is described using Bayesian penalised B-splines, which we described in full in section 3.3. Each of the five model parameters (f, R, $d_{\parallel}$, θ and ϕ) is characterised by a cubic B-spline function defined using a fixed knot vector t, a set of weighting parameters $\mathbf{a}$ and a smoothing parameter Λ. During parameter estimation we optimise the weighting and smoothing parameters for all tissue indices simultaneously.

Synthetic phantom construction

The ground truth phantom for this experiment is constructed in a similar way to those described in sections 4.2.2 and 5.2.2. As in the previous experiments, the ROI comprises one slice of 100 voxels, which mimics the shape and structure of the mid-sagittal CC. Unlike our previous work, we include an axon radius parameter R alongside the parameters representing volume fraction, diffusivity and fibre orientation for each voxel. The functions describing the variation of R, f, $d_{\parallel}$, θ and ϕ across the ROI are shown in figure 6.2. The synthetic ROI exhibits spatial variation in both R and f, whereas $d_{\parallel}$, θ and ϕ are constant across all voxels. As before, the synthetic values of all parameters are chosen to match realistic values measured in white matter tissue and the trend in R is chosen to match the trend observed in histology [2]. For the parallel diffusivity, $d_{\parallel} = 0.6 \times 10^{-9}\mathrm{m}^2\mathrm{s}^{-1}$ is physically plausible for ex vivo tissue [7, 59]. From electron microscopy studies [2], it is known that there exists a distribution of axon radii in the CC. However, we choose to use a single axon radius parameter in each voxel as the exact relationship between the underlying axon radius distribution and the axon radius index is still unclear. Alexander et al [7] hypothesise that the index corresponds to the mean of the axon radius distribution weighted by axonal volume, however the relationship appears to be more complex. Without knowledge of this relationship we would be unable to fully evaluate how well the parameter estimates correspond to the ground truth values. The fibre direction is set to align with the z axis i.e. $\theta = \phi = 0$. We do not introduce any angular variation as in previous chapters, as we synthesise the data using Monte Carlo simulation (see the Data Synthesis section for details), for which it is not trivial to change the angles of the synthetic axons.

Acquisition protocols

We use ten acquisition protocols to synthesise data, ranging from high ($G_{max} = 300$ mT m^{-1}) to low ($G_{max} = 40$ mT m^{-1}) gradient strengths. All protocols are optimised using the optimisation in [6] for axon radii of 1, 2 and 5 μm. Each protocol comprises 360 unique measurements, roughly the number of measurements that can be acquired on a scanner within an hour [6], divided into three High Angular Resolution Diffusion Imaging (HARDI) shells. Full details of all protocols are presented in table 6.1. The constraints on rise times ($T_{rise} = 200\mu$s) and other pulse sequence timing parameters used during the optimisation are based on the hardware capabilities of a pre-clinical experimental scanner.

Data synthesis

For each acquisition protocol, we synthesise data from all 100 voxels in the ROI using the Monte Carlo simulation framework of Hall and Alexander [80], implemented in Camino [44]. The substrate used

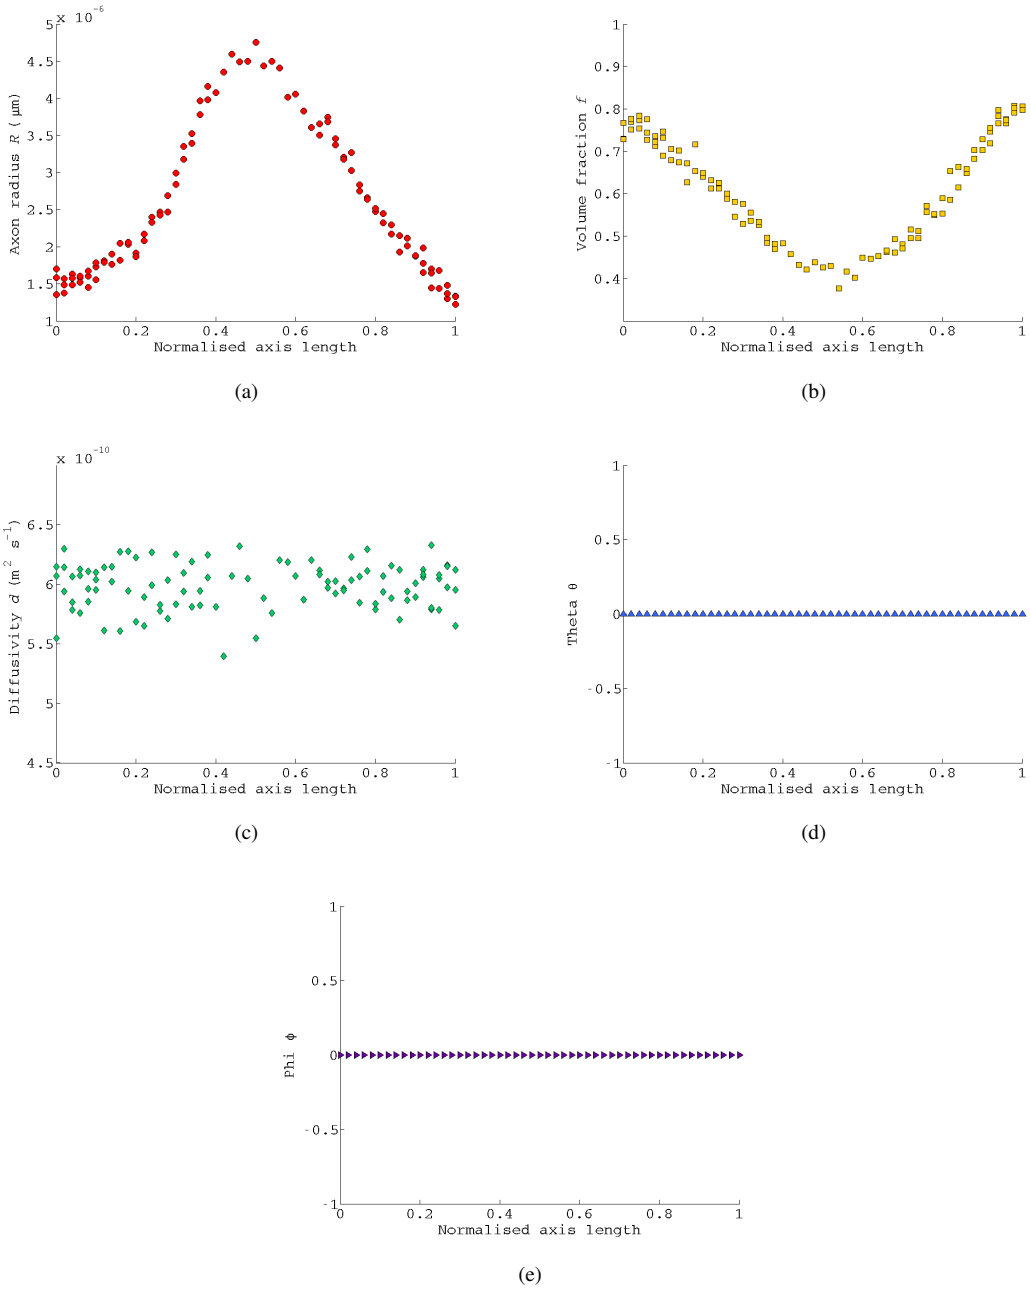

Figure 6.2: Ground truth parameter values as a function of distance along the normalised axis of the synthetic phantom for (a) R, (b) f, (c) $d_{\|}$, (d) θ and (e) ϕ.

G_{max} (mT m^{-1})	Shell	$N_{b=0}$	N_{meas}	b (s mm^{-2})	G (mT m^{-1})	Δ (ms)	δ (ms)
40	3	22	87	968	40	27.0	20.5
	2	22	94	1848	29.7	46.8	27.9
	1	22	113	3891	40	40.6	34.1
60	1	23	91	1299	60	23.4	16.9
	2	23	99	2333	42.8	39.8	23.6
	3	22	102	5322	60	35.0	28.5
80	3	23	94	1540	80	20.9	14.4
	1	23	102	2669	54.8	35.2	21.0
	2	23	95	6515	80	31.3	24.8
100	3	23	95	1718	100	19.2	12.7
	2	23	104	2936	66.5	32.0	19.0
	1	24	91	7565	100	28.8	22.3
120	2	23	97	1852	120	17.8	11.3
	1	24	105	3150	77.7	29.5	17.5
	3	23	88	8496	120	26.8	20.3
140	1	24	98	1951	140	16.7	10.2
	2	23	105	3318	88.9	27.5	16.3
	3	23	87	9296	140	25.2	18.7
160	1	24	99	2029	160	15.8	9.3
	2	24	105	3457	99.5	25.9	15.3
	3	23	85	9987	160	23.9	17.4
180	1	24	100	2099	180	15.1	8.6
	3	23	105	3599	110	24.5	14.5
	2	24	84	10671	180	22.8	16.3
200	1	24	100	2158	200	14.5	8.0
	2	24	105	3723	121	23.4	13.8
	3	23	84	11254	200	21.8	21.8
300	1	24	102	2336	300	12.4	5.9
	2	24	105	4164	170	19.3	11.4
	3	23	82	13663	300	18.6	12.1

Table 6.1: Acquisition protocols used to synthesise data. $N_{b=0}$ is the number of $b = 0$ s mm^{-2} acquired and N_{meas} is the number of diffusion-encoding gradient directions used in each shell of the protocol. The acquisition parameters b, G, Δ and δ are previously defined. The highlighted rows show the acquisition parameters used later in experiment 2.

in the simulation comprises impermeable, hexagonally packed, parallel cylinders with a single radius corresponding to the ground truth value of R for each voxel. The spacing between the cylinders is calculated so that the volume fraction f matches the specified ground truth value. The diffusivity of the walkers, which are the particles undergoing the random walk in the MC simulation, for each voxel of data is set to the corresponding value of $d_{\parallel}$. The simulation is run using 100,000 walkers and 1000 timesteps, the lower limit suggested by Hall and Alexander [80]. After data synthesis, we add Rician noise to each voxel of data as described in section 4.2.2. We consider two different noise levels: low noise data, corresponding to SNR $= 25$, and high noise data, corresponding to SNR $= 10$. For each combination of acquisition protocol and noise level we generate 25 independent, noisy data sets, as this allows us to evaluate the bias and variance of our parameter estimates, relative to the ground truth, for each parameter estimation method. This results in a total of 500 synthetic data sets, simulated from 10 unique acquisition protocols and two noise levels.

Independent voxelwise parameter estimation

We fit the ZeppelinCylinder model to each of the independent voxelwise estimation following the method described by Alexander et al [7]. Parameter estimation consists of three separate stages:

1. Grid search. Based on physically plausible values for R and f, a grid of all possible combinations is constructed. Estimates of θ and ϕ are obtained from fitting the DT to the voxel data. The parallel diffusivity $d_{\parallel}$ is fixed to a physically plausible value; for these simulation experiments we fix $d_{\parallel}$ to its true value of $0.6\times10^{-9}\mathrm{m}^2\mathrm{s}^{-1}$. For each combination of R (R takes values of $0.25, 0.5, 1, 1.5, 2$ μm) and f (f ranges from 0.5 to 0.9 in steps of 0.1), together with the fixed values of $d_{\parallel}$, θ and ϕ, we evaluate the likelihood of these model parameters given the data using the Rician likelihood function. The most likely combination of R and f is the one that maximises the Rician likelihood.

2. Gradient descent algorithm. The output of the grid search is used to initialise a gradient descent search of the parameter space in order to provide an esimate of the most likely tissue parameters. As for the grid search, $d_{\parallel}$ remains fixed in order to ensure stability and the objective function is the Rician likelihood function.

3. MCMC sampling. The final step in the parameter estimation is MCMC sampling. The maximum likelihood estimate obtained during the gradient descent search is used as the starting point for MCMC. Only two of the parameters, R and f, are updated during the sampling, whilst $d_{\parallel}$, θ and ϕ remain fixed. We discard the first 5000 iterations during the burn in phase, after which we draw 500 samples at intervals of 20 iterations. As for the previous two stages, we use the Rician likelihood function to determine whether to accept or reject a parameter update step. The final parameter estimates are calculated as the mean of the MCMC samples.

For all three stages of parameter estimation, the variance of the Rician distribution is set to the value used to generate the noisy data.

Parameter estimation using the RVM

In chapters 4 and 5, we used the same optimisation steps regardless of whether we were estimating the tissue parameters in each voxel independently or using the RVM. However, when fitting the Zeppelin-Cylinder model two extra stages are incorporated into the voxelwise parameter estimation which are not particularly feasible to implement when fitting the RVM. First, performing a grid search over all of the RVM spatial model parameters would be very time consuming. Second, due to the weak dependence of the diffusion MR signal on R, the gradient descent algorithm encounters many local minima and, very frequently, the maximum likelihood estimate is identical to the starting point. Therefore, we estimate the posterior distribution over the RVM model parameters using only MCMC (It would also be possible to fit the ZeppelinCylinder model to each voxel independently using just MCMC. However, we choose to follow the three stage approach of Alexander in order to compare our results to those obtained in [7] and [60]). Unlike the independent voxelwise fitting procedure, we allow θ and ϕ to be updated along with R and f, as the fibre orientation can vary slightly over the ROI and we want our posterior distributions to reflect this variation. To provide a direct comparison between the methods, $d_{\|}$ remains fixed throughout. For consistency with independent voxelwise estimation, the burn in phase comprises the first 5000 iterations, after which we collect 500 samples at at interval of 20 iterations (further discussion of these parameters is provided in appendix A). The variance of the Rician distribution is fixed to the value used to generate the noisy data.

When fitting the RVM to the data, the variation of the ZeppelinCylinder diffusion model parameters is initialised as follows:

- axon radius *R* is modelled by N_R 4^{th} order (i.e. cubic) B-spline basis functions defined over N_R+4 knots spanning the medial axis of the ROI. The basis function weights are determined by a $1\times N_R$ vector $\mathbf{a}_R = \{a_{i,R}\}_{i=1}^{N_R}$ and the smoothness of the curve is determined by the precision Λ_R, both of which we optimise. Initially all N_R weights are set to 3μm, a physically plausible value for white matter [2] which is in the middle of the range of values used for R in the grid search in [7], although we test a range of starting points from $R = 2$ to $R = 5\mu$m. The precision parameter is initialised so that $\Lambda_R = 100$, corresponding to medium smoothing. The hyperparameters of $\Pr(\Lambda_R)$ are α_R and β_R.

- volume fraction f is modelled by N_f 4^{th} order (i.e. cubic) B-spline basis functions defined over N_f+4 knots spanning the medial axis of the ROI. The basis function weights are determined by a $1\times N_f$ vector $\mathbf{a}_f = \{a_{i,f}\}_{i=1}^{N_f}$ and the smoothness of the curve is determined by the precision Λ_f, both of which we optimise. Initially all N_f weights are set to 0.7, a physically plausible value for white matter [149], so that the curve is initially constant everywhere within the ROI with a value of 0.7. Again, we test the stability of the model to a range of starting points, from $f = 0.4 - 0.8$. The precision parameter is initialised so that $\Lambda_f = 100$, corresponding to medium smoothing. The hyperparameters of $\Pr(\Lambda_f)$ are α_f and β_f.

- the fibre orientation parameters θ and ϕ are modelled using one 1^{st} order B-spline basis function

each, i.e. both angles are assumed to be constant across the ROI. These basis functions are defined over the knot vector $\begin{bmatrix}0 & 1\end{bmatrix}$, i.e. one knot at each end of the ROI axis. The weight assigned to the basis function is controlled by one scalar value - a_θ for θ and a_ϕ for ϕ. The weights a_θ and a_ϕ are initialised from the average principal DT direction across all voxels within the ROI.

- parallel diffusivity is fixed to its ground truth value of $0.6 \times 10^{-9}\mathrm{m}^2\mathrm{s}^{-1}$. This is done in order to provide a direct comparison with the results from independent voxelwise fitting, although initial experiments (results not shown here) indicate that the RVM fitting is stable if we choose instead to fit a constant value of $d_\parallel$ over the ROI.

As in chapter 4, in order to investigate the stability of the RVM over different spatial model settings, we consider three different knot vectors with different N, corresponding to a small, medium and large number of knots, in order to determine the effect of t on the fitted microstructural variation curves. These are

- $t_1 = \{-0.2 : 0.1 : 1.2\}$. This knot vector comprises 15 knots, with separation 0.1. When using 4^{th} order cubic basis functions, this knot vector supports $N = 11$ basis functions.

- $t_2 = \{-0.1 : 0.05 : 1.1\}$. This knot vector comprises 25 knots, with separation 0.05. When using 4^{th} order cubic basis functions, this knot vector supports $N = 21$ basis functions.

- $t_3 = \{-0.04 : 0.02 : 1.04\}$. This knot vector comprises 55 knots, with separation 0.02. When using 4^{th} order cubic basis functions, this knot vector supports $N = 51$ basis functions.

We only consider spatial models in which $N_R = N_f$ in order to limit the number of RVMs we fit to each data set. Further discussion of these knot vectors is given in appendix B

We also consider a range of different smoothing priors for both R and f. The smoothing priors are defined as in chapter 4; however, given the results observed in the validation experiments in chapter 4 we do not consider the uninformative prior here as it performed poorly when fitting the simple ball and stick model, and the ZeppelinCylinder model is more sensitive to the effects of noise. Because of this, we also replace the weak smoothing prior (which we hypothesise will be overly influenced by any noise in the data) with a stronger smoothing prior. Therefore, the smoothing priors we use in this chapter are:

- $\mathrm{S}_1 = \Gamma(\Lambda_p|\alpha_1, \beta_1)$, where $\alpha_1 = 4$, $\beta_1 = 25$. This corresponds to a gamma distribution with mean $\mu_\Gamma = \alpha\beta = 100$ and variance $\sigma_\Gamma = \alpha\beta^2 = 2500$. This favours small values of Λ_p resulting in moderate smoothing.

- $\mathrm{S}_2 = \Gamma(\Lambda_p|\alpha_2, \beta_2) = \Gamma(\Lambda_p|\alpha_2, \beta_2)$, where $\alpha_2 = 4$, $\beta_2 = 250$. This corresponds to a gamma distribution with mean $\mu_\Gamma = \alpha\beta = 1000$ and variance $\sigma_\Gamma = \alpha\beta^2 = 250000$. This favours medium values of Λ_p resulting in heavy smoothing.

- $\mathrm{S}_3 = \Gamma(\Lambda_p|\alpha_3, \beta_3)$, where $\alpha_3 = 4$, $\beta_3 = 2500$. This corresponds to a gamma distribution with mean $\mu_\Gamma = \alpha\beta = 10000$ and variance $\sigma_\Gamma = \alpha\beta^2 = 25000000$. This favours large values of Λ_p resulting in very heavy smoothing.

The smoothing priors may be different for R and f, and we consider all 9 possible combinations when fitting the RVMs. Although the number of smoothing priors we use is limited, as in chapter 4, the distributions are wide and there is enough overlap between them so that they cover a large range of smoothing parameters.

Parameter estimation using smooth curves

As previously demonstrated in chapters 4 and 5, and studies such as [74], it is also possible to fit trends of microstructure variation from the independent voxelwise parameter estimates themselves. When fitting these smooth curves, we again use the Bayesian penalised B-splines framework as described in chapter 4, section 4.2.2. As described previously, in order to perform a direct comparison between the RVM and the smoothed curves, we use the optimal knot vectors and smoothing priors determined during the initial RVM experiments when fitting the smoothed spline functions for both R and f. We do not calculate mean curves because, as shown in chapter 4, they do not provide a significant advantage over voxelwise fitting, particularly in narrow regions of tracts.

Statistical analysis

The first aim of this experiment is to investigate the effect of the spatial model settings on the estimated trends of R and f using the RVM, and to determine the optimal number of knots and smoothing priors. In order to do this we consider two metrics. The first metric is the mean squared error (MSE) between the ground truth and estimated tissue parameters averaged over all voxels in the ROI. The best combination of knots and smoothing priors should minimise the MSE for both R and f simultaneously. For a full definition of these metrics, see chapter 4, section 4.2.2. The second metric is the correlation between the ground truth and estimated tissue parameters. The best combination of knots and smoothing priors should maximise the correlation coefficients for both R and f simultaneously. We include the correlation metric here, as Dyrby et al [60] show that at lower gradient strengths the axon radius index is overestimated even though the correct trend is still recovered. When this happens the MSE is not necessarily the best metric to use as there is a large discrepancy between estimated and ground truth parameters leading to a large MSE, even though the trend is still useful. In this case, the correlation is a better measure of the RVM performance.

Unlike chapter 4, in which we determined the optimal spatial model settings for each protocol/SNR combination after averaging these metrics over all 25 independent noisy sets, in this experiment we calculate the metrics for one noisy dataset only. These metrics are used to determine the best number of knots and smoothing priors to use for each protocol/SNR combination, which are then used to fit the RVM to the remaining 24 independent noisy data sets. The main reason for this change is computational time. Fitting the RVM using the ZeppelinCylinder model takes on the order of hours, compared to minutes for the BallStick model. Therefore we decided that it would be too time consuming to fit 27 different representations of the RVM to all 500 synthetic data sets for the purposes of this initial experiment.

The second aim of this experiment is to compare how well the parameter estimates from the different fitting methods match the ground truth values. To do this, we calculate the bias and variance of

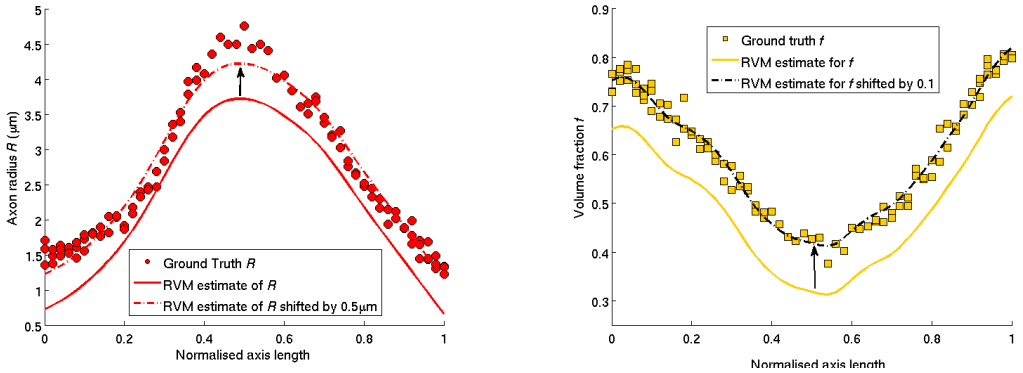

(a) (b)

Figure 6.3: Initial estimates of (a) R and (b) f using the RVM from the simulated G_{max}=300 mT/m, SNR=25 data. In both figures we show the ground truth values. For this dataset, which is the highest quality synthetic data generated, we would expect the RVM to recover the ground truth well. However, there is a slight offset for both parameters, which is also seen in the voxelwise fits (not shown here). This suggests that the ZeppelinCylinder model with the tortuosity constraint may not fully characterise the signal from the MC simulations.

the parameter estimates over all 25 independent noise trials, relative to the ground truth, for each protocol/SNR combination for each fitting technique. The general formulae for calculating bias and variance are presented in chapter 4, section 4.2.2. For each protocol/SNR combination, we can then compare these values for independent voxelwise estimation, the RVM and smooth curves in order to determine which methods perform best under different data acquisition conditions.

6.2.3 Results

Figure 6.3 shows curves of R and f estimated using the RVM from data synthesised using $G_{max} = 300$ mT/m and SNR=25, along with the ground truth values. For data of this quality, we would expect that the recovered parameters would match the ground truth values well, however there is a systematic underestimation of both R and f along the whole of the ROI for all initialisations of R and f. Although we do not show the results here, we observe the same effect in voxelwise fits, indicating that the problem is unlikely to be due to the RVM itself. Given that $d_{||}$ is fixed at its ground truth value during all the synthetic data experiments, and the estimated values of θ and ϕ are both correct, the shift in R and f is not due to the incorrect estimation of the other model variables. The reason for this mismatch is due to the ZeppelinCylinder model itself and its inability to capture the characteristics of the simulated diffusion MR signal fully, potentially due to the tortuosity approximation we make. To test this, we performed another simulation experiment (results not shown here), but synthesised the diffusion signals for the ZeppelinCylinder model itself, instead of using MC simulations. In this case, the RVM was able to recover the underlying trends in both R and f. However, given that the trends in parameter variation across the ROI are correct, the ZeppelinCylinder model with the simplifying tortuosity approximation is still a useful model. The potentially systematic bias in parameter approximation does mean that the

proposed MSE metric may not be the best way of assessing the optimal number of knots and smoothing. Instead, because the fitted trends match the ground truth trend well, we use the correlation coefficient between estimated and ground truth values to assess the RVM knot and smoothing parameters.

Table 6.2 shows the best choices of knots and smoothing parameters for R and f for all gradient strengths and noise levels, based on the finding the maximum correlation coefficients between fitted and ground truth values of R and f. From these values, we can note several patterns to guide our choice of knot and smoothing parameters in subsequent experiments. First, the knot vector t_2, with $N = 25$ knots is the best for all acquisition protocols and noise levels. Second, for nearly all protocols, the best value of smoothing for the axon radius index S_R depends on the SNR of the data. With the exception of $G_{max} = 80, 100$ mT/m, when the SNR is 10 the weakest smoothing prior S_1 is chosen whereas for SNR=25 the medium smoothing prior S_2 is chosen. This is because the dependence of the signal on R is very weak. Therefore when the data are noisy, the stronger smoothing priors are liable to oversmooth and the underlying axon radius trend will not be recovered. For higher SNR data, the medium smoothing prior does not flatten out the trend and is preferable to the weaker one as it is less likely to fit to noisy fluctuations in the data. To test whether an even weaker prior would be preferable for the low SNR data, we fit the RVM to these data sets using a weaker prior on R. However, when this very weak prior is used, the RVM fits trends to the noise, rather than the ground truth values of R. Third, the volume fraction trend is best fitted using stronger smoothing, i.e. either $S_f = S_2$ or S_3. The signal has a linear dependence on f which is much stronger than its dependence on R; therefore the stronger priors do not flatten out the trends whilst also recovering smooth trends that are robust to noisy fluctuations. Finally we note that, based on the full correlation results for all possible prior and knot combinations, f is fairly robust to the choice of RVM model parameters. Whilst the values given in table 6.2 result in the highest correlation between fitted and ground truth trends, any combination of knots and smoothing prior can be used and still result in a correlation coefficient greater than 0.8. The same is not true for trends in R, which are much more dependent on the exact knot vector and smoothing prior used. In general good trends can be recovered for all protocols and SNR when either $N = 25$ and $S_R = S_1$ or $S_R = S_2$, or when $N = 15$ and $S_R = S_1$. If the strongest smoothing prior is used, the trend is flattened out regardless of the number of knots; if the weak or medium prior is used with $N = 55$ knots, the trend tries to fit to the noise rather than matching the ground truth values.

The second objective of this experiment is to compare the results obtained with the optimal RVM settings to voxelwise and smoothed estimates of the trends in R and f over 25 independent noise trials to determine how robust the methods are to noise for different acquisition protocols and SNR levels. Although we compare the different methods for all 20 combinations of acquistion protocols and SNR values, we only present a selection of results here that illustrate the key differences or changes between the fitting techniques.

Figure 6.4 shows plots of bias and variance in the estimated trends for R and f for voxelwise, smooth curve and RVM fitting. This is the highest quality dataset, which should therefore result in good fits using all estimation techniques. As noted earlier, and shown in figure 6.3, all methods consistently

G_{max} (mT m^{-1})	SNR	optimal t	optimal S_R	correlation R	optimal S_f	correlation f
40	25	25	S_1	0.9783	S_3	0.9887
	10	25	S_1	0.8310	S_3	0.9838
60	25	25	S_2	0.9870	S_3	0.9883
	10	25	S_1	0.9639	S_3	0.9861
80	25	25	S_2	0.9885	S_3	0.9886
	10	25	S_2	0.9893	S_2	0.9883
100	25	25	S_2	0.9893	S_2	0.9883
	10	25	S_2	0.9864	S_2	0.9860
120	25	25	S_2	0.9902	S_2	0.9885
	10	25	S_1	0.9891	S_2	0.9862
140	25	25	S_2	0.9830	S_3	0.9885
	10	25	S_1	0.9875	S_2	0.9890
160	25	25	S_2	0.9867	S_3	0.9869
	10	25	S_1	0.9875	S_2	0.9890
180	25	25	S_2	0.9895	S_3	0.9885
	10	25	S_1	0.9862	S_3	0.9885
200	25	25	S_2	0.9911	S_2	0.9868
	10	25	S_1	0.9871	S_2	0.9866
300	25	25	S_2	0.9891	S_3	0.9872
	10	25	S_1	0.9831	S_2	0.9864

Table 6.2: Table showing the optimal number of knots and smoothing factors to use in the RVM for a range of gradient strengths and noise levels, determined by finding the best correlations between the ground truth and estimated parameters. We also show the resulting correlation coefficients for both R and f.

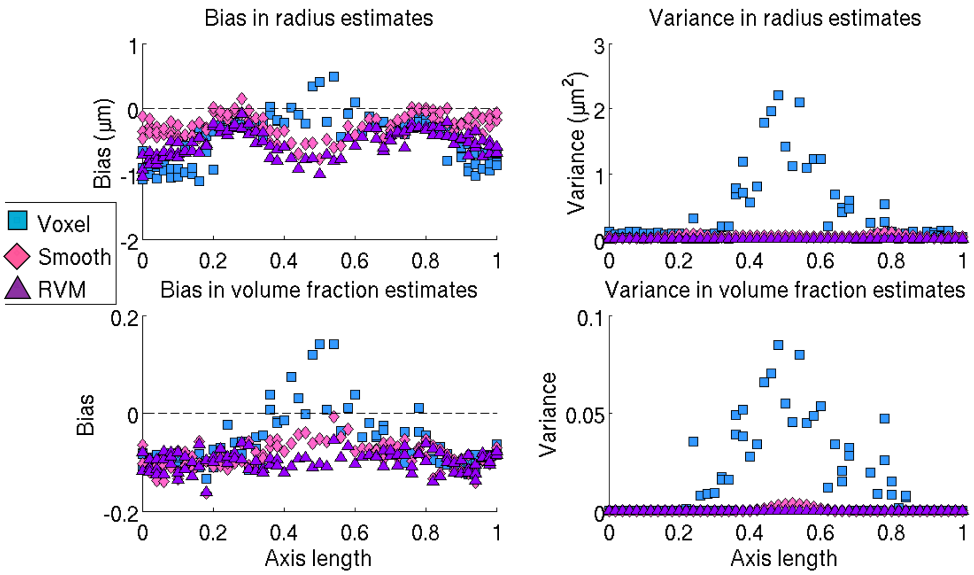

Figure 6.4: Plots of bias and variance of R and f estimated using voxelwise, smoothing curve and RVM fitting, as a function of distance along the medial axis of the ROI, for synthetic data generated using $G_{max} = 300$ mT/m and SNR=25.

underestimate both axon radius and volume fraction parameters over the whole region, because the ZeppelinCylinder model, with a tortuosity approximation, is unable to fully capture the characteristics of the simulated signals. When comparing the variance of estimates over multiple noisy trials, both the smooth curves and the RVM perform very well for both R and f. Surprisingly, given the high gradient strength and SNR of the data, there are several voxels for which the independent voxelwise estimation of both f and R is very variable, particularly in voxels in the centre of the ROI. A reason for this is that the axon radius index is reasonably high at this location. As we see from figure 6.1, for an acquisition protocol with $G_{max} = 300$ mT/m, the gradient of the signal curve as a function of R is steep between approximately 2 and 4 μm. However, after 4 μm the signal flattens out making it harder to differentiate between larger axons. As our ground truth values for R are $4 - 5\mu$m in the centre of the ROI, this may explain the high variance associated with R, as it is difficult to correctly estimate in the presence of noise. We also note from figure 6.1 that for large axon radii, not only is the signal curve flat but it is also close to zero and therefore most likely within the noise floor, even for high SNR data. This low signal could potentially cause problems for the more stable volume fraction parameter, which may help to explain the higher than expected variances for f in this region.

Figure 6.5 shows the bias and variance plots for data synthesised with $G_{max} = 160$ mT/m and SNR=25. For this dataset, and all those with lower gradient strengths, we start to see the voxelwise estimation being outperformed by the smoothed and RVM estimates, at least in terms of the variability of the trends in R over the independent noise trials. We see that the variance of independent voxelwise estimates of R is consistently higher than estimates obtained with either smooth curves or the RVM, particularly at either end of the medial axis. At these points, as we see from the ground truth plots in

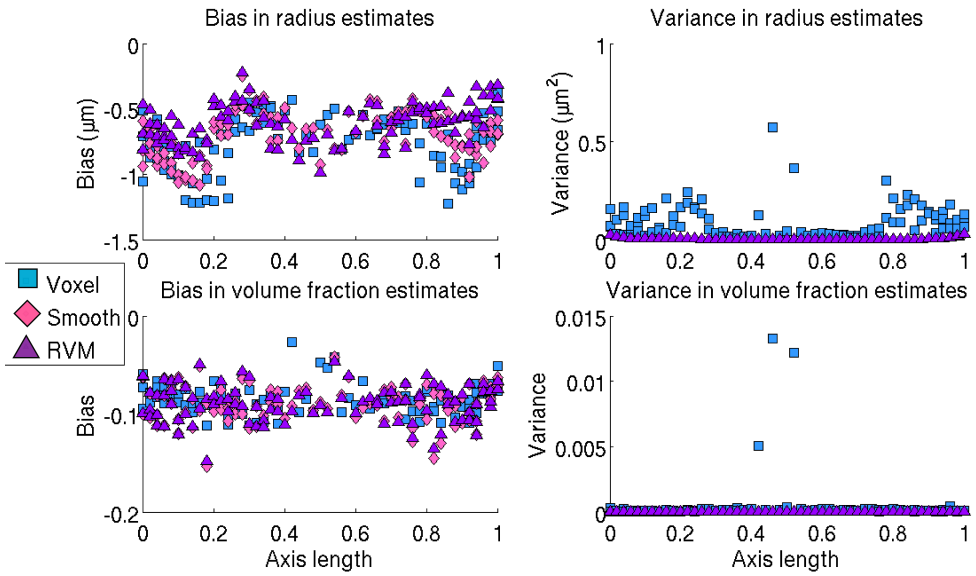

Figure 6.5: As 6.4, but for data generated with $G_{max} = 160$ mT/m and SNR=25.

figure 6.2, the axon radius index is smallest. Smaller axon radii of approximately 0-2 μm are particularly difficult to estimate as the signal curve as a function of R is again very flat in this range, making it hard to distinguish the correct value. Over multiple noise trials, this can lead to a lot of variability in our estimates. Both the RVM and smoothing methods help to ameliorate this. Variance in f is low for all methods, although there are a couple of voxels for which the independent voxelwise parameter estimation struggles. Again, all methods display a systematic bias in both R and f which, as discussed above, we believe is caused by the simplifying tortuosity approximation that we employ in the ZeppelinCylinder model.

Figure 6.6 shows the bias and variance plots for data synthesised with $G_{max} = 80$ mT/m and SNR=10. At this gradient strength and noise level, the RVM provides an advantage over both independent voxelwise estimation and the smooth functions for extracting regional trends in R. The variance in voxelwise estimates of R varies across the ROI from 0.5 to 3.5 μm^2, with the largest variance occuring in the middle of the ROI. Because the synthetic phantom is designed so that it is thinnest at this point, the variability of the voxelwise estimates influence the smooth curves fitted to them, and so we see an increase in variance in the smooth curves compared to higher gradient strength data sets, particularly in the centre of the ROI. In contrast, the variance of RVM estimates over the noise trials is low across the whole ROI. Again, the variance for f is generally low, although it is now higher for the voxelwise estimates than the smooth or RVM estimates. There is a systematic negative bias in f for all methods, as observed in previous datasets. However, the bias for R is slightly different. We know see that for all methods, R is overestimated at the edge of the ROI (particularly when using the RVM), but underestimated in the centre.

Finally, figure 6.7 shows the resulting bias and variance plots for the lowest quality dataset, synthesised using $G_{max} = 40$ mT/m and with added noise so that the SNR of the unweighted signals is 10. In

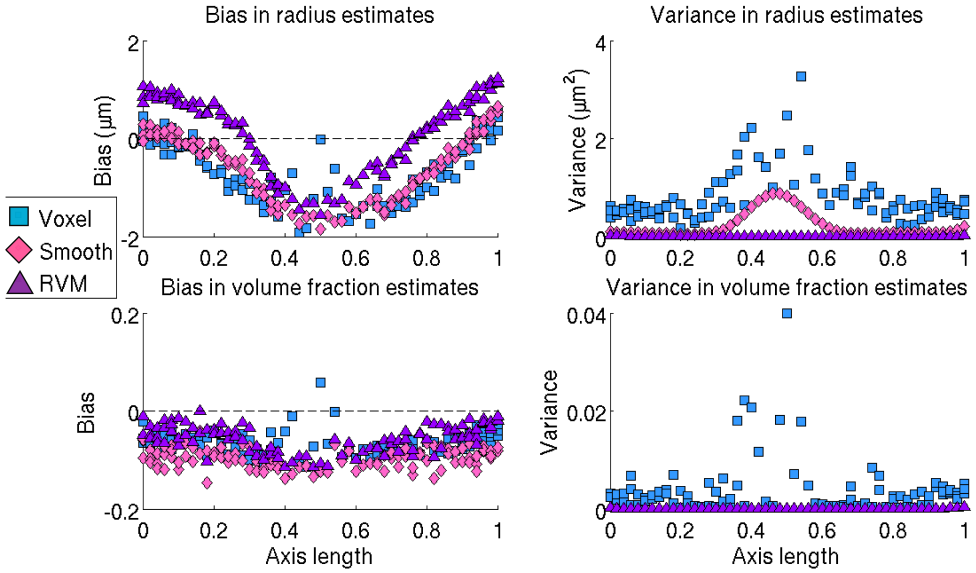

Figure 6.6: As 6.4, but for data generated with $G_{max} = 80$ mT/m and SNR=10.

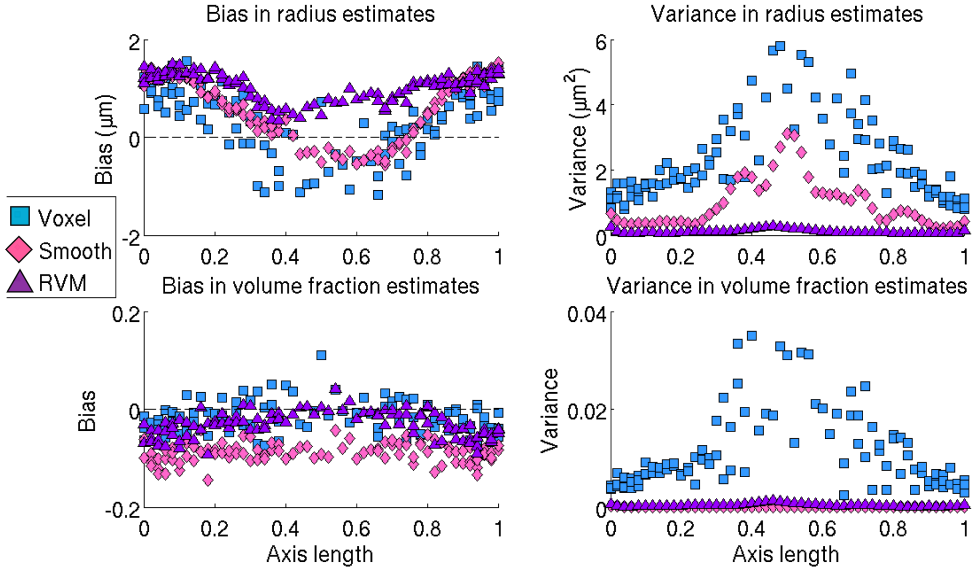

Figure 6.7: As 6.4, but for data generated with $G_{max} = 40$ mT/m and SNR=10.

this case we see very clearly that the RVM is more robust to noise than independent voxelwise estimation and smooth curves when extracting regional trends in R. Over the 25 trials, the variance in estimates of R obtained with independent voxelwise estimation is large, ranging between 1 and 6 μm^2 over the length of the ROI. The variability in the underlying voxelwise estimates influence the smooth curves, which are derived from them, leading variances of up to 3 μm^2 in some voxels in the middle of the ROI. Although the voxels in the centre have the largest axon radii, which should be easier to estimate, because the synthetic region is thinnest at this point, the smooth curves are more susceptible to noisy voxelwise estimates here. The RVM is still able to extract reliable trends in R over multiple noisy instances of the same datasets despite these issues, because we fit the trends directly to the signals over the whole ROI simultaneously. We still see a systematic bias in both f and R. There is still a slight underestimation of f for all methods, although the bias in R is more variable. The RVM now consistently overestimates R by approximately 1 μm, whereas the voxelwise estimations, and thus subsequently derived smooth curves, both overestimate the small axon radii at the edges of the ROI and underestimate the larger axons in the centre.

6.2.4 Discussion

In this experiment, we used synthetic data generated from a known ground truth to investigate the optimal number of knots and smoothing parameters that can be used to accurately recover regional trends in axon radius R and volume fraction f. We then compared the trends extracted using the RVM (with the optimal parameters) to those estimated using independent voxelwise fitting and smoothing, for a range of gradient strengths and noise levels, to determine the accuracy and robustness of the different fitting techniques.

The first part of the experiment, designed to determine the appropriate number of knots and smoothing priors for each diffusion model parameter, again highlights one of the weaknesses of the RVM, namely that the settings for weak parameters such as axon radius require careful tuning to ensure that the estimated trends really do represent the underlying ground truth values. However, using the correlation coefficients between ground truth and estimated parameter values (which we used due to a systematic bias which may make the MSE inappropriate), we can identify useful patterns between the data properties and the best RVM parameters to use. For example we see that because of its strong influence on the diffusion signal, we can use any number of knots and any smoothing prior when estimating volume fraction and still fit sensible trends. Despite this, in order to achieve the best trends, the medium knot vector t_2 with $N = 25$ knots and either the medium or strong smoothing priors are preferred in order to avoid fitting to noisy fluctuations in the data. In contrast, in order to fit the appropriate trend in R, a limited range of smoothing parameters and knot vectors must be used. For example, if we use $N = 10$ or $N = 50$ knots, the trends are either oversmoothed or too noisy. Therefore, we always choose $N = 25$ knots. Regardless of the number of knots, the strongest smoothing prior S_3 always flattens out the trends. There is also a link between the SNR of the data and the appropriate smoothing prior for R. When the SNR is high we can use the medium prior S_2 to ensure smooth, noise-free trends; however for low SNR data we must use the weaker smoothing prior S_1 to prevent oversmoothing. The only exception is for the lowest gradient strength protocol, with $G_{max} = 40$ mT/m, for which the sensivity of the signal to R

is weak enough that we must use the weakest smoothing prior to best fit the trend, regardless of SNR. We use these guidelines to pick appropriate knot vectors and smoothing priors for the experiments in the rest of this chapter. However, we acknowledge that any differences between the simulations and real data experiments in terms of acquisition protocols, noise levels and underlying ground truth values may affect how well our choice of parameters works. In future, we would like to improve how the knots and smoothing priors are chosen, potentially by incorporating it into the RVM fitting procedure.

The second part of this experiment compared the trends obtained using the optimal RVM settings with the trends obtained by independent voxelwise fitting and the smooth curves over 25 independent noise trials. Overall, we note that at high gradient strengths (G_{max} >160 mT/m), all methods are able to extract trends in R and f robustly, albeit with some systematic bias. For data acquired with a maximum gradient strength of 160 mT/m or below, we start to see advantages of the RVM and smooth curves over independent voxelwise estimation for robustly estimating regional trends of axon radius. Finally at gradient strengths of 80 mT/m, we start to see that the RVM is able to estimate R more robustly than both independent voxelwise and the smooth curves. Given that the maximum available gradient strength on most clinical scanners is currently 60 mT/m, this suggests that the RVM could provide a real advantage for estimating parameters such as the axon radius index clinically compared to the current state of the art methods. Finally we note that for the diffusion MR signal model we currently use, there is a bias in our estimates of R and f for all the fitting methods we consider. This bias does not necessarily mean that the model we are using is inappropriate. Provided that the bias is constant regardless of the parameter value, we can still use our estimates to determine whether a parameter changes significantly over time. This appears to be the case for f as for all acquisition protocols, the bias appears to be the same regardless of the ground truth volume fraction. For the axon radius however, there are some examples where the bias we introduce into our estimate seems to depend on the ground truth value of R. If this is the case, our current axon radius index parameter may be less useful. In future, we can try to alter the diffusion MR signal model to reduce or eliminate this bias, for example by removing the tortuosity approximation in the ZeppelinCylinder model and incorporating $d_\perp$ as a function to be estimated. Due to the data pooling feature of the RVM, removing this approximation should not significantly affect the stability of the fitting, although we would need to carefully consider how to incorporate an independent perpendicular diffusivity $d_\perp$ parameter into the model. However, we may find that the tortuosity approximation is appropriate for modelling real brain tissue, in which there is a distribution of axon radii. In future, we can test this in simulation, by synthesising signals from a tissue substrate with a distribution of axon radii.

6.3 Experiment 2: Investigating the effect of G_{max} on ex vivo axon radius estimation

The aim of this experiment is to compare estimates of the axon radius index obtained from ex vivo monkey CC using the RVM and current state of the art techniques. We use the ex vivo monkey brain data acquired by Dyrby et al [60] at a range of gradient strengths. This allows us to test whether the

conclusions drawn from the simulated data in section 6.2 are also valid from real brain tissue, which is much more complex than the simple computational substrates investigated previously.

6.3.1 Hypothesis

Using the results obtained in section 6.2 as a guide, we hypothesise that for data acquired using maximum gradient strengths of 200 and 300 mT/m, the precision of estimated axon radius indices is comparable for both the RVM and independent voxelwise estimation. However, when estimating the axon radius index from data acquired using maximum gradient strengths of 60 and 140 mT/m, we hypothesise that the RVM performs better than independent voxelwise estimation in terms of precision.

6.3.2 Methods

In this section, we outline the combination of diffusion model used to build the specific RVM for this experiment. We also present details of the diffusion-weighted data acquisition and discuss the region of interest selection and parameter estimation steps.

The diffusion model

The diffusion model used for this experiment is the ZeppelinCylinderDot model introduced by Panagiotaki et al [134]. This model differs slightly from the ZeppelinCylinder model described in section 6.2.2 and used in the simulation experiments in 6.2 due to the inclusion of a third compartment known as the Dot. The Dot represents the signal A_{dot} due to stationary water trapped within or on white matter tissue membranes. We include this additional compartment in the diffusion signal model for this experiment as both [7] and [134] found that it was necessary to fully capture the characteristics of the diffusion-weighted MR signal in fixed tissue. Because the water molecules described by the dot are trapped, there is no attenuation of the signal and

$$A_{dot} = 1. \tag{6.8}$$

Given that the fraction of trapped water molecules is g, the modified diffusion signal model is then

$$A\left(d_{\|}, R, f, \theta, \phi, g\right) = fA^{ia}\left(d_{\|}, R, \theta, \phi\right) + (1 - f - g)\, A^{ea}\left(d_{\|}, f, \theta, \phi\right) + gA_{dot}, \tag{6.9}$$

where the equations describing the signals due to both the intra- and extra-axonal compartments are as described above in section 6.2.2.

In both [7] and [60], the diffusion model incorporates a fourth compartment to model partial volume effects due to CSF contamination. We do not include this compartment in this experiment for two reasons. First, unlike the other parameters in the RVM, we do not expect the amount of CSF contamination to vary smoothly across white matter regions making it more problematic to include in the RVM framework. Second, due to the constraints on linearity and planarity used in ROI selection, which we describe in more detail below, we do not expect any of the voxels in the ROI to suffer from significant contamination. However, for white matter regions in which CSF contamination is more significant, it is possible to estimate the amount of free water contamination using a method such as [138] and use this as prior information during the RVM parameter estimation.

Data acquisition

We summarise the data acquisition procedure here, which is described in full in [60]. The perfusion fixed brain of a female, 32 month old vervet monkey (*Chlorocebus aethiops*) was obtained from the Montreal Monkey Brain Bank, and prepared for scanning as described by Dyrby et al [59]. Diffusion-weighted MR data is acquired using a 4.7 T Varian scanner using four acquisition protocols with maximum gradient strengths $G_{max} = \{60, 140, 200, 300\}$ mT/m, which are optimised using the Active Imaging framework described in [6]. Full details are given in the highlighted rows in table 6.1. All protocols consist of 360 measurements; the remaining measurements are $b = 0$ s mm^{-2} images. A total of 30 sagittal slices are acquired, covering the whole of the CC. The voxel size is 0.5 mm isotropic. The minimum echo time possible is used for each protocol resulting in TE=$\{71.5, 52, 46, 39.1\}$ms and the repetition time TR$= 2.5$s. The total scanning time is 168 hours, during which the ex vivo brain is kept at a constant temperature of $20\pm1^{\circ}$C. We estimate the SNR of white matter in the $b = 0$ images of the 4 datasets using the `estimatesnr` command in camino [44]. Due to the increase in TE as the gradient strength decreases, the SNR also decreases so that SNR=$\{13, 15, 17, 22\}$ for the four acquisitions.

Region of interest selection

As in [60], we focus on estimating the axon radius index within the MSCC. The method for segmenting the MSCC and identifying the medial axis is as described in section 4.3.2.

Voxelwise parameter estimation

The ZeppelinCylinderDot model is fit to each voxel of data in the MSCC using the three stage parameter estimation method described in section 6.2.2. As in the synthetic data experiment in section 6.2, the parallel diffusivity is fixed during parameter estimation to ensure the stability of the model fitting. To determine the appropriate diffusivity, we perform a preliminary experiment in which we fit the model to each voxel of data using camino [7, 44] with a range of diffusivities, from $d_{\|} = 0.35$ to 0.95×10^{-9} m^2 s^{-1} in steps of 0.01 m^2 s^{-1}. We calculate the mean Rician log likelihood over the whole ROI in order to determine the value of $d_{\|}$ that minimises the likelihood function.

RVM parameter estimation

The RVM is fit to the data as described in section 6.2.2, using the optimal number of knots N and prior distributions S_R and S_f suggested in section 6.2.3 to model the variation of R and f across the MSCC. For the G_{max}=200, 300 mT/m data, for which the SNR is closer to 25 than 10, we use $N = 25$, $S_R = S_2$ and $S_f = S_3$. We note that although the optimal value of $S_f = S_2$ for $G_{max} = 200$ and SNR=25, altering the smoothing for f has little to no effect on the fitting, and so we use the same values for consistency. For the G_{max}=60, 140 mT/m data, for which the SNR is closer to 10 than 25, we use $N = 25$, $S_R = S_1$ and $S_f = S_3$.

Preliminary parameter maps for the trapped water fraction obtained by Alexander et al [7] suggest that fraction of water in the Dot compartment g varies across the CC. Therefore we also model the variation of g as a function of medial axis length using N cubic B-spline basis functions. As g is the same order of magnitude as f, we set the prior distribution S_g for g to be equal to S_f. As in section

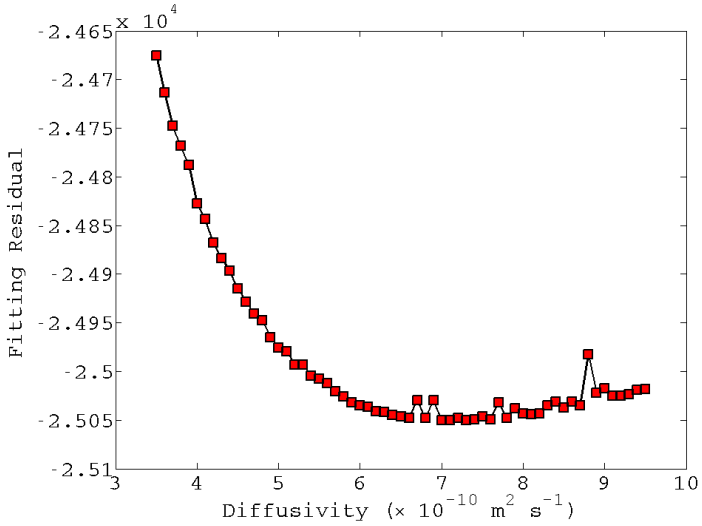

Figure 6.8: The ZeppelinCylinderDot model is fit to each data of voxel independently, using a range of fixed diffusivities between $d_{\parallel} = 0.35$ to 0.95×10^{-9} m^2 s^{-1}. The objective function is averaged over all voxels and the value $d_{\parallel}$ that minimises the objective function is determined.

6.2.2, we model the fibre orientation as constant across the MSCC and fix $d_{\parallel}$ at the value used during the voxelwise parameter estimation to ensure a fair comparison between the two methods, although preliminary experiments show that the RVM is stable when $d_{\parallel}$ is not fixed.

In this experiment, we do not fit smooth curves to the voxelwise estimates as in previous experiments. Instead we focus on a comparison between the current state of the art voxelwise estimation and the RVM. However, based on the results obtained in section 6.2.3, we expect that the performance of the smooth curves would be similar to the RVM for the datasets acquired with $G_{max} = 140$, 200, 300 mT/m, but the performance would perform worse than the RVM for data acquired with $G_{max} = 60$ mT/m, when the independent voxelwise estimates from which the smooth curves are derived are noisy.

6.3.3 Results

Figure 6.8 shows the likelihood of the ZeppelinCylinderDot model fit over the whole ROI for a range of different values of $d_{\parallel}$ for ex vivo data acquired using $G_{max} = 300$ mT/m. For each diffusivity $d_{\parallel}$ we plot the mean of the objective function, in this case the negative Rician likelihood function, over all voxels. The objective function is minimised when $d_{\parallel} = 0.71 \times 10^{-9}$ m^2 s^{-1}, and therefore we use this as our fixed diffusivity value in the remaining ex vivo experiments for both independent voxelwise and RVM parameter estimation. The value we determine is slightly larger than the fixed diffusivity of 0.6×10^{-9} m^2 s^{-1} used in [7]. This is most likely because we do not include a CSF compartment, with an intrinsically higher diffusivity, in our model of the diffusion MR signal. However, we note that the objective function remains relatively flat over a large range of diffusivities, from $d_{\parallel} = 0.62 \times 10^{-9}$ m^2 s^{-1} to $d_{\parallel} = 0.80 \times 10^{-9}$ m^2 s^{-1}, indicating that within this range the model fitting is relatively insensitive to changes in diffusivity. As $d_{\parallel}$ increases, the estimated values of R, f and g are affected. However, within the stable range of diffusivities from $d_{\parallel} = 0.62$ to 0.8 m^2 s^{-1}, this effect is a small

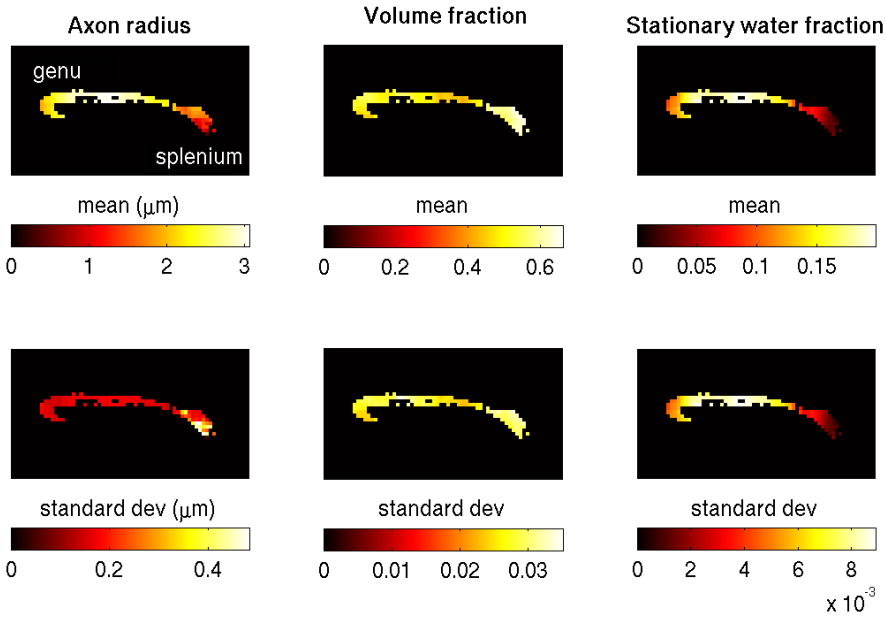

Figure 6.9: Voxelwise parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 300$ mT/m ex vivo monkey data.

positive shift in the estimated values of R, f, and g which does not affect the overall trends over the ROI.

Figures 6.9 and 6.10 show parameter and standard deviation maps for the three key parameters, axon radius R, intra-axonal volume fraction f, and stationary water fraction g, that we model the variation of across the MSCC for the ex vivo monkey data acquired with a maximum gradient strength of 300 mT/m. Figures 6.11 and 6.12 show the same parameter maps for the $G_{max} = 200$ mT/m data, figures 6.13 and 6.14 for the $G_{max} = 140$ mT/m data and finally, 6.15 and 6.16 for the $G_{max} = 60$ mT/m data. The parameter maps obtained are stable regardless of initial values of R and f. Note that the scale bars are not the same in all figures. Although we do not show parameter maps for the fibre orientation angles θ and ϕ here, we note that for all datasets, both independent voxelwise and RVM estimation are able to capture the expected left-right fibre direction in the corpus callosum.

When a maximum gradient strength of 300 mT/m is used to acquire the data, independent voxelwise estimation of axon radius R, volume fraction f and stationary water fraction g results in smooth parameter maps, as demonstrated in figure 6.9. We recover the expected low-high-low trend in axon radius in the CC [2, 7, 60], and the estimated values of R between 1 and 3 μm are anatomically plausible [2]. The volume fraction of intracellular water is approximately constant across the whole ROI, varying between 0.5 and 0.6. If, using these values, we followed the approach of Alexander et al [7] and calculated an axon density based upon f and R, we would also recover the expected high-low-high trend in axon density [2, 7]. The stationary water fraction, which represents water trapped on or within membranes, also exhibits a low-high-low trend across the CC. Whilst we do not have any data from histology regarding the possible distribution of trapped water molecules within the CC, this trend agains seems plausible when compared to the parameter maps obtained by Alexander et al [7]. The parameter

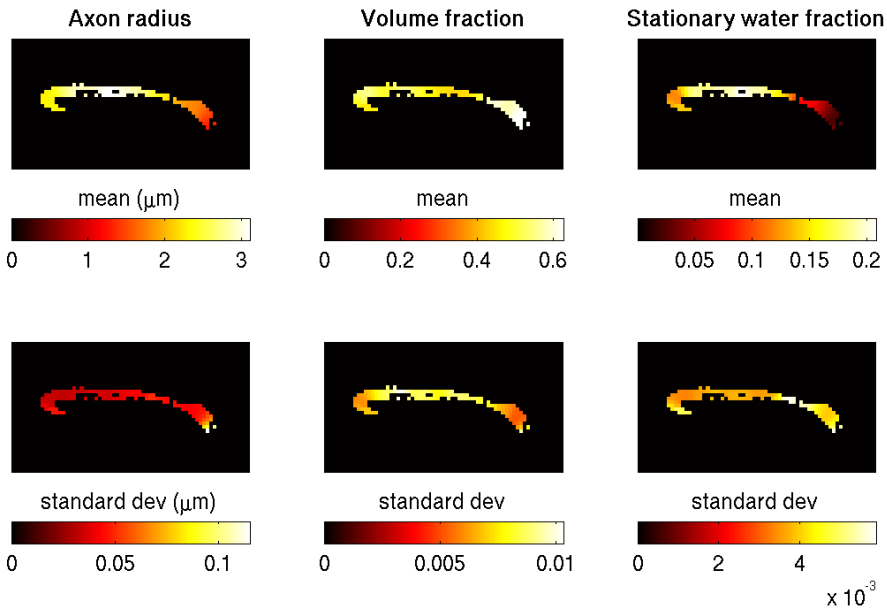

Figure 6.10: RVM parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 300$ mT/m ex vivo monkey data.

maps in figure 6.10, which use the RVM to estimate R, f and g, match these trends and values well. The standard deviation of the estimates, calculated using the posterior distributions generated during MCMC, do show that the RVM is able to estimate these indices more precisely. Across the whole ROI, the standard deviation for R estimated using voxelwise fitting is $\sigma_R = 0.2\mu$m opposed to 0.05 μm using the RVM. For f, the standard deviation using voxelwise fitting is $\sigma_f = 0.03$ compared to 0.007 for the RVM. For the stationary water fraction, the standard deviation for both voxelwise and RVM fitting is the same at $\sigma_g = 5 \times 10^{-3}$. However, despite the small improvement of the RVM over the current state of the art in terms of standard deviations, values of σ_R, σ_f and σ_g are still low enough at $G_{max} = 300$ mT/m to provide confident estimates of R, f and g, regardless of the estimation method.

When the maximum gradient strength is reduced to $G_{max} = 200$ or 140 mT/m, the parameter maps for R, f and g still show the same trends for both independent voxelwise estimation and the RVM. However, for voxelwise estimation the parameter maps for R start to look noisier, particularly in the splenium, where in some voxels the radius is estimated to be zero. At these gradient strengths, the key difference between voxelwise estimation and the RVM is exhibited in the standard deviation, and therefore the precision, of our estimates. Using voxelwise estimation, $\sigma_R \approx 0.3 - 0.5$ μm, $\sigma_f \approx 0.03$ and $\sigma_g \approx 0.005 - 0.01$. In contrast, when using the RVM we obtain $\sigma_R \approx 0.05 - 0.15$ μm, $\sigma_f \approx 0.0075$ and $\sigma_g \approx 0.005$. The difference between these standard deviations is even greater when $G_{max} = 140$ mT/m. In this case, when using voxelwise estimation, $\sigma_R \approx 0.4 - 0.7$ μm, $\sigma_f \approx 0.03 - 0.05$ and $\sigma_g \approx 0.005 - 0.01$. The corresponding values obtained using the RVM we obtain $\sigma_R \approx 0.05 - 0.2$ μm, $\sigma_f \approx 0.005 - 0.01$ and $\sigma_g \approx 0.006$.

The trends for R estimated using the data acquired with a maximum gradient strength of 60 mT/m

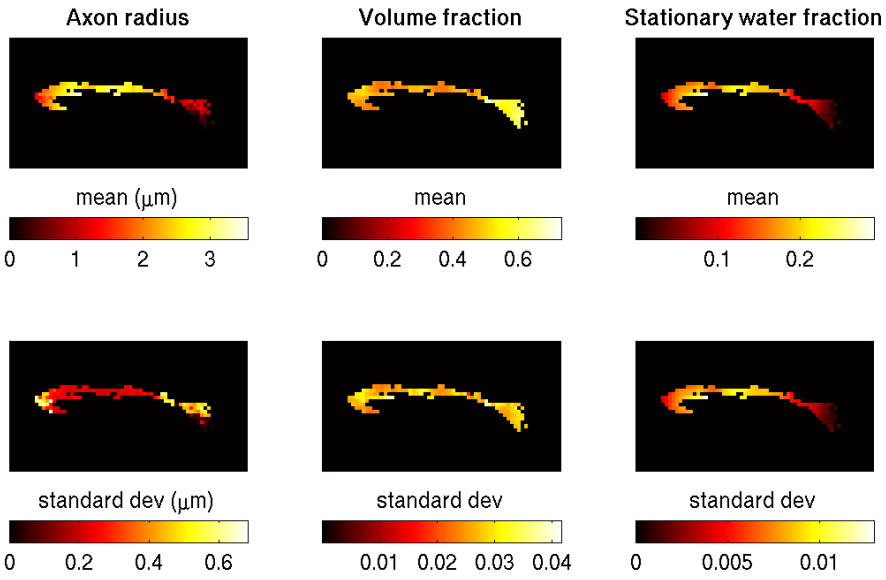

Figure 6.11: Voxelwise parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 200$ mT/m ex vivo monkey data.

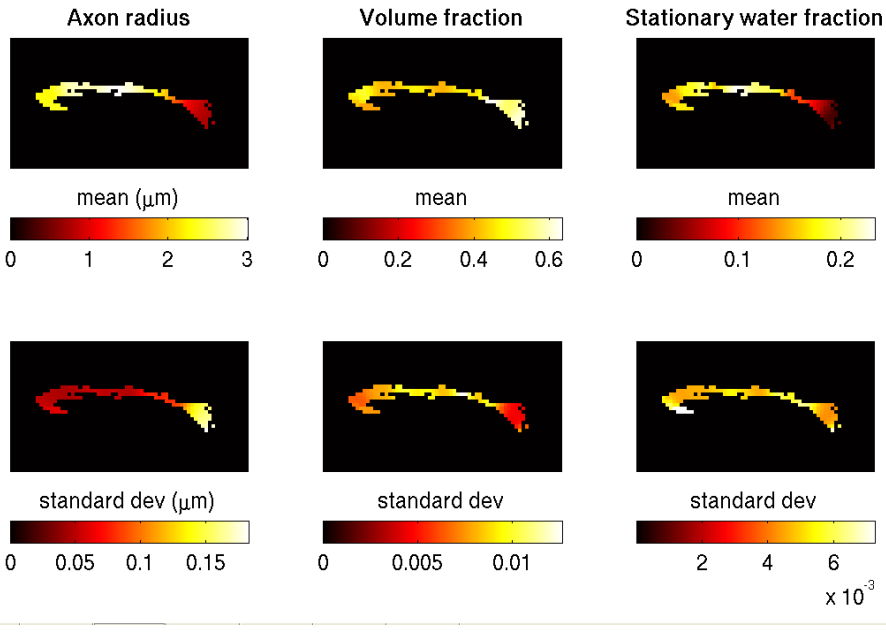

Figure 6.12: RVM parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 200$ mT/m ex vivo monkey data.

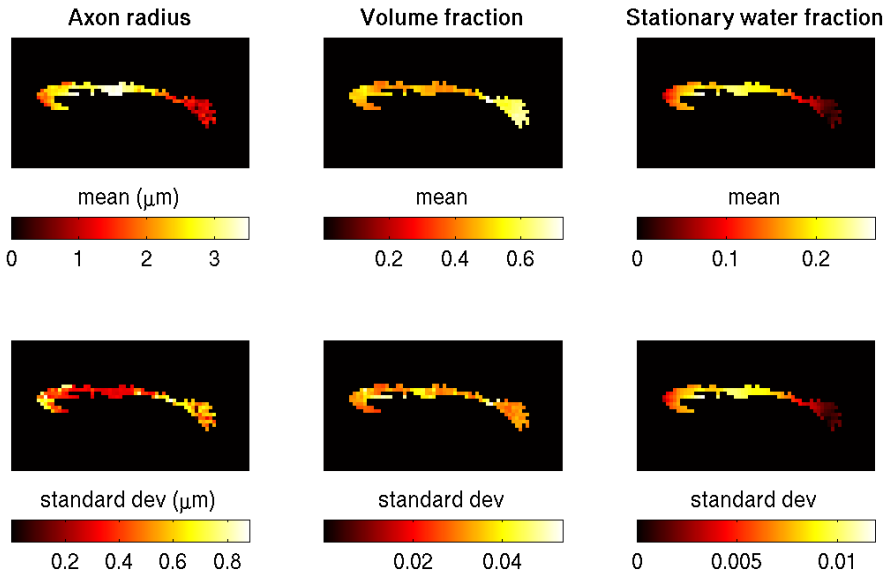

Figure 6.13: Voxelwise parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 140$ mT/m ex vivo monkey data.

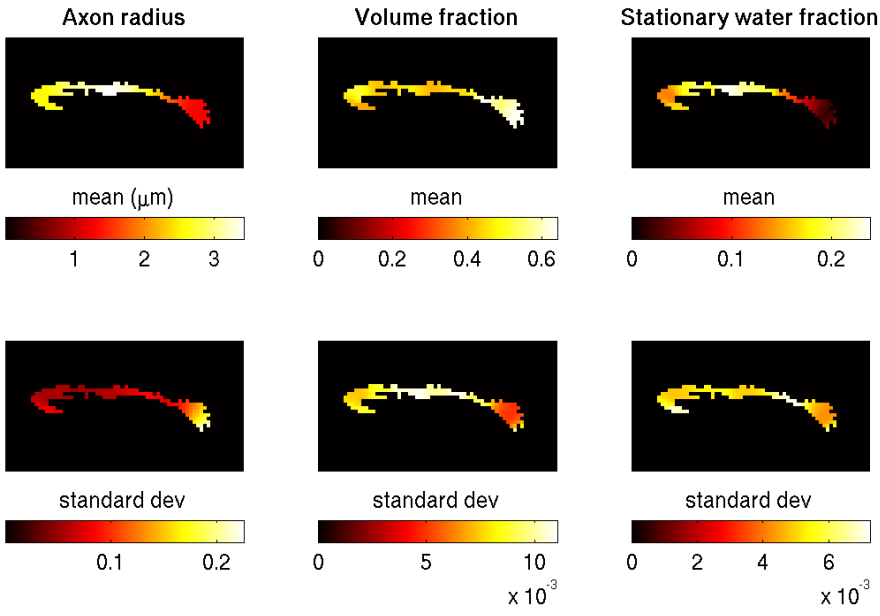

Figure 6.14: RVM parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 140$ mT/m ex vivo monkey data.

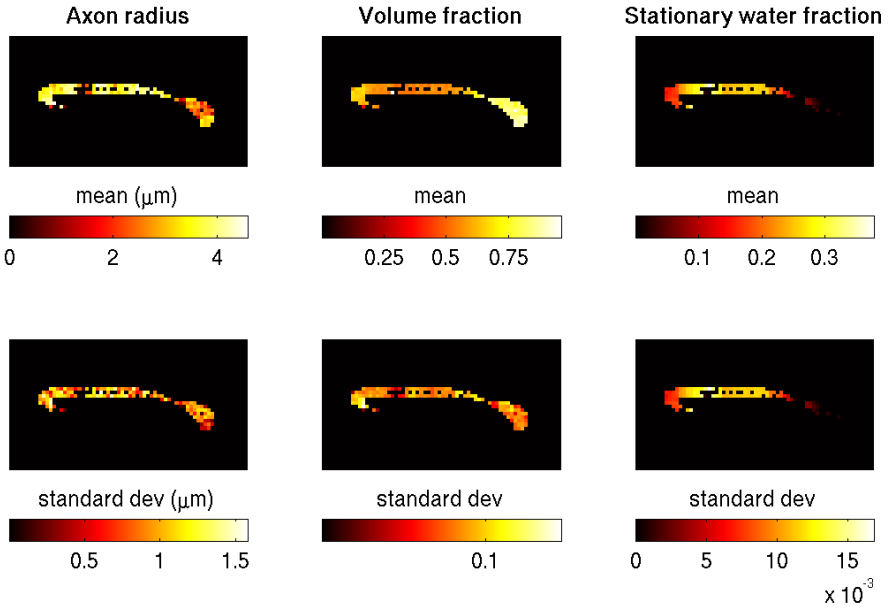

Figure 6.15: Voxelwise parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 60$ mT/m ex vivo monkey data.

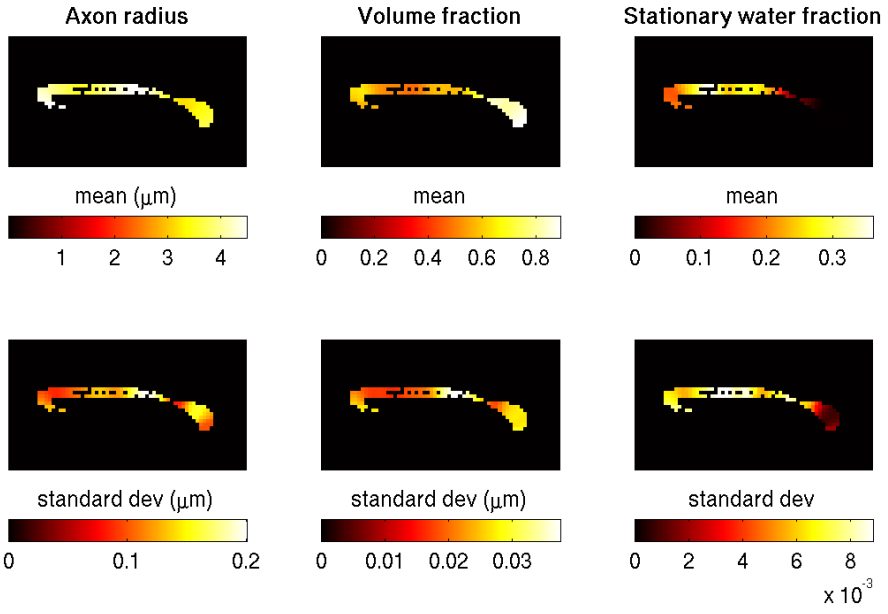

Figure 6.16: Voxelwise parameter maps of the mean and standard deviation of R, f and g from $G_{max} = 60$ mT/m ex vivo monkey data.

do not match those obtained for the higher gradient strengths for either parameter estimation method. Instead of observing the expected low-high-low trend, we instead see a high-low trend in which axon radius is high ($R \approx 4\mu$m) in both the genu and midbody of the CC before getting smaller in the splenium ($R \approx 2\mu$m for voxelwise, $R \approx 3$ μm for the RVM). The trend is noisier when estimated independently in each voxel, with a standard deviation of $\sigma_R = 0.5 - 1.5\mu$m compared to $\sigma_R = 0.1 - 0.2$ μm for the RVM. The trends obtained for f and g at $G_{max} = 60$ mT/m are more similar to those obtained at higher gradient strengths, although a much higher value of $f \approx 0.8$ is estimated in voxels in the splenium. The standard deviation of the volume fraction is again much higher for voxelwise estimation ($\sigma_f \approx 0.1$) compared to the RVM ($\sigma_f \approx 0.02 - 0.03$). The trends for g are very similar for both methods as are the calculated standard deviations of $\sigma_g = 0.005 - 0.01$.

6.3.4 Discussion

The goal of this experiment was to compare estimates of axon radius index obtained from ex vivo monkey brain data, acquired at a range of gradient strengths, using both the RVM and independent voxelwise estimation. Our initial hypothesis was that both methods would perform comparably when the data was acquired at high gradient strengths, i.e. $G_{max} = 200, 300$ mT/m, but that for data acquired with lower gradient strengths, i.e. $G_{max} = 60, 140$ mT/m, the precision of parameters would be better when estimated with the RVM.

For data acquired with $G_{max} = 300$ mT/m, both voxelwise estimation and the RVM recover the expected low-high-low trend of axon radius across the CC, and the mean and standard deviation over the parameter posterior distributions is very similar, showing that at high gradient strengths, both methods perform well. However, we do find that for data acquired at $G_{max} = 200$ mT/m, as well as at $G_{max} = 140$ mT/m, the standard deviations over the posterior distribution is better using the RVM compared with independent voxelwise estimation. This is slightly surprising, as based on simulation experiments, we anticipated that independent voxelwise estimation for data acquired with $G_{max} = 200$ mT/m would be more robust, although the simulation experiments were considering variance over multiple noise trials instead of posterior distributions, which is what we investigate here. At these gradient strengths, both methods recover the expected low-high-low trend in axon radius index. For the lowest gradient strength, $G_{max} = 60$ mT/m, neither method is able to capture this trend. This is due to artifacts that are present in this particular data set, and is most notable in voxels in the splenium of the CC, which exhibit large fluctuations in the signal intensities in the $b = 0$ images. However, the RVM estimates are still more precise for both R and f. Overall, regardless of gradient strength, the RVM does consistently reduce the standard deviation of fitted values for R and f, increasing the precision of our estimates even at low gradient strengths. This is unsurprising, as one of the key features of the RVM is the pooling of data from voxels with similar features to increase the effective SNR of the estimated parameters. This can have important clinical consequences as being able to measure over time are due to real underlying changes in the tissue structure rather than due to noise.

Dyrby et al have previously fitted the MMWD, i.e. the ZeppelinCylinderDotBall model to the same data sets; therefore we can compare our results. For gradient strengths of 140, 200 and 300 mT/m,

the trends we obtain for R are similar and we obtain very similar axon radius indices, indicating that the lack of CSF compartment in our model does not affect the overall results. However, for the data acquired at $G_{max} = 60$ mT/m, we obtain a high-low trend across the CC, with $R = 4\mu$m in the genu and midbody and $R = 2 - 3\mu$m in the splenium, whereas Dyrby et al still obtain a low-high-low trend with $R \approx 4\mu$m in the genu, $R \approx 6\mu$m in the midbody and $R \approx 2\mu$m in the splenium. This discrepancy is mainly due to the difference in the diffusion MR signal model, but also affected by different parallel diffusivities. Dyrby et al include a CSF compartment, whereas our ZeppelinCylinderDot model does not. After comparing the results, we refit the diffusion model to each voxel of data independently, but including the CSF compartment. Using this modified model, our results match those of Dyrby et al more closely (although our values for R are slightly and systematically greater by about 1μm), although the CSF volume fractions we estimate are unrealistic. In all voxels in the splenium, the CSF fraction is approximately 0.5. This is probably due to artifacts, potentially Gibbs ringing, which are present in the $G_{max} = 60$ mT/m data, causing large fluctuations in the $b = 0$ images. Dyrby et al use $d_{\|} = 0.6 \times 10^{-9}$ m^2 s^{-1} compared to our value of $d_{\|} = 0.71 \times 10^{-9}$ m^2 s^{-1}. This causes shifts in estimates of R, f and g, although from our experiment to select $d_{\|}$, we found that lowering the diffusivities shifts R towards lower values, rather than the increased axon radius indices observed in Dyrby's results. Therefore, we think the key factor causing these discrepancies is the CSF compartment.

One of the key changes we can make in future to improve our results, particularly at low gradient strengths, is to include the diffusivity as a parameter to be estimated during fitting. The diffusivity estimates we obtain from diffusion MRI are only apparent, as they depend on the timing parameters of the pulse sequence. Therefore, although we fix $d_{\|} = 0.71 \times 10^{-9}$ m^2 s^{-1} for all datasets here, this is probably not true due to the different values of δ, Δ and TE used in the acquisitions. We fix it here to match the methodology of Alexander et al [7] and Dyrby et al [60] as closely as possible, who fix the value to increase the stability of the model. The RVM is stable when $d_{\|}$ is included as a parameter, according to preliminary tests, and this might improve our estimated trends in R, particularly for data acquired at low gradient strengths.

6.4 Experiment 3: Estimating axon radius from in vivo human data

The previous experiments in this chapter have focused on estimating an index of axon radius using the RVM from either synthetic data or high quality ex vivo data acquired using long acquisition times and high gradient strengths. However one of the key goals in medical imaging is to make indices such as axon radius viable as clinical biomarkers. Currently most 3T clinical MR scanners have a maximum gradient strength of 40 mT/m, although some systems with a maximum gradient strength of 60 mT/m are available. Alexander et al previously showed [7], using optimised protocols with G_{max}=60 mT/m and one hour of scanning time, that we can estimate a potentially useful axon radius index in the mid-sagittal corpus callosum and recover the expected low-high-low trend [2]. However the estimates obtained are noisy, and are unlikely to be precise enough to distinguish subtle microstructural changes due to pathol-

ogy. At the lower gradient strength of 40 mT/m, this will be even more difficult. In this experiment we aim to investigate whether the RVM can improve the precision of axon radius indices estimated from in vivo human data acquired using clinical scanners (G_{max}= 40, 60 mT/m) and therefore determine whether the framework can improve the viability of an axon radius index as a clinical biomarker.

6.4.1 Hypothesis

We hypothesise that, due to its natural pooling of data with white matter tracts, the RVM will improve the precision of axon radius indices estimated from clinically acquired, in vivo human data, compared to the current state of the art method, which estimates the index in each voxel of data independently.

6.4.2 Methods

In this section, we briefly describe the diffusion model used in this experiment as well as the data acquisition, ROI segmentation and model fitting procedures for both voxelwise and RVM parameter estimation.

The diffusion model

As we are dealing with in vivo data in this experiment, we do not need to use the third Dot compartment to account for the signal from stationary water molecules [7]; therefore we revert to using the simpler ZeppelinCylinder model introduced in section 6.2 which parameterises the signal using the axon radius R, intra-axonal volume fraction f, parallel diffusivity $d_{\|}$ and fibre orientation angles θ and ϕ.

Data acquisition

We acquire a total of 6 data sets for this experiment; two subjects are scanned using a maximum gradient strength of 60 mT/m whereas the other four subjects are scanned using a maximum gradient strength of 40 mT/m. The acquisition protocols used to acquire the data are optimised as in [6] and full details of both protocols are shown in table 6.3.

The $G_{max} = 60$ mT/m protocol is set up on a 3T Philips Achieva scanner. Data is acquired from two healthy volunteers (32 year old male, 25 year old female), using a cardiac-gated EPI sequence. Ten sagittal slices, centred on the mid-sagittal plane, are acquired with an in-plane resolution of 2 mm $\times$ 2 mm and slice thickness of 5 mm. The repetition time TR is 7RR, and the echo time TE$= 100$ ms. The data has previously been described by Schneider et al [158].

The $G_{max} = 40$ mT/m protocol is set up on a 3T Siemens Tim Trio scanner. Data is acquired from four healthy volunteers (3M 1F, age range: $26 - 35$), using a cardiac-gated EPI sequence. A total of 24 sagittal slices, centred on the mid-sagittal plane, are acquired with an in-plane resolution of 2 mm $\times$ 2 mm and slice thickness of 3 mm. The repetition time TR is 6RR, and the echo time TE$= 110$ ms.

Data acquired with a gradient strength of 40 mT/m are acquired at the GlaxoSmithKline Clinical Imaging Centre. As this scanner has a maximum achievable gradient strength of 40 mT/m, in order to perform the 60 mT/m experiment, we used previously acquired data on a different scanner. Due to the fact that different scanners are used to acquire the data, it is more difficult to directly compare the results as hardware differences can have confounding effects and this must be taken into account when interpreting results.

For all datasets, we estimate the SNR of white matter in the $b = 0$ image using the `estimatesnr`

G_{max} (mT m^{-1})	Shell	$N_{b=0}$	N_{meas}	b (s mm^{-2})	G (mT m^{-1})	Δ (ms)	δ (ms)
40	3	8	76	630	38	87.2	8.5
	1	8	70	651	38	26.8	17.3
	2	6	94	2740	38	76.2	19.5
60	3	22	101	540	60.0	19.2	11.7
	2	24	107	870	47.8	38.2	12.5
	1	25	81	2634	60.0	29.1	21.6

Table 6.3: Acquisition protocols used for in vivo human scanning. $N_{b=0}$ is the number of $b = 0$ s mm^{-2} acquired and N_{meas} is the number of diffusion-encoding gradient directions used in each shell of the protocol. The acquisition parameters b, G, Δ and δ are previously defined.

command in camino. For the $G_{max} = 60$ mT/m data the SNR is approximately 13; for the $G_{max} = 40$ mT/m data, the SNR is approximately 11.

Region of interest selection

Once again, we focus on estimating the ZeppelinCylinder diffusion model parameters in the MSCC. We identify and segment the MSCC as described in section 4.3.2. However, for some subjects the mid-sagittal plane does not align with the slice direction. In section 5.3, we were able to automatically detect the mid-sagittal plane using the algorithm described by Ardekani et al [12]. However, because we only acquire data from a limited number of slices centred on the mid-sagittal CC in order to limit the scanning time, the algorithm fails to identify the sagittal plane correctly. In these cases, we identify the appropriate transformation manually.

Voxelwise parameter estimation

The model is fit to each voxel of data independently as described in section 6.3. For each data set we initially fit the model for a range of $d_{\|}$, from 1.5×10^{-9} m^2 s^{-1} to 2.3×10^{-9} m^2 s^{-1} in intervals of 0.01×10^{-9} m^2 s^{-1} for the $G_{max} = 60$mT/m data sets and from 1.4×10^{-9} m^2 s^{-1} to 2.3×10^{-9} m^2 s^{-1} in intervals of 0.01×10^{-9} m^2 s^{-1} for the $G_{max} = 40$mT/m data sets. For each value of $d_{\|}$ we evaluate the likelihood function over the whole ROI in order to find the best value. In this experiment we perform this step for each subject individually as anatomical differences may result in different ideal values of $d_{\|}$.

RVM parameter estimation

The RVM is fit to the data as described in section 6.2.2, using the optimal number of knots N and smoothing parameters S_R and S_f suggested in section 6.2.3 to model the variation of R and f across the MSCC. For both the $G_{max} = 60$ mT/m data and the $G_{max} = 40$ mT/m data, for which the SNR is close to 10, we use $N = 25$, $S_R = S_1$ and $S_f = S_3$. We fix $d_{\|}$ at the optimal value determined in our preliminary experiments.

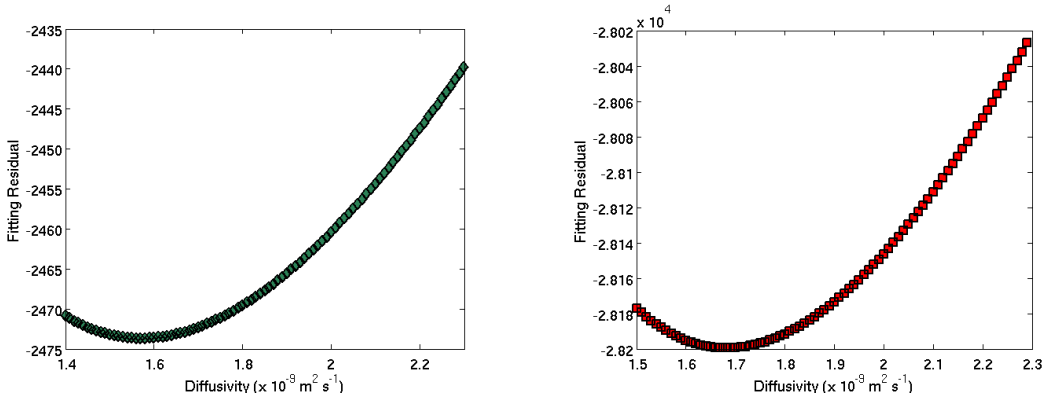

(a) (b)

Figure 6.17: Plots showing the mean negative likelihood function over the ROI as a function of $d_{\|}$ for (a) one subject imaged using a maximum gradient strength of $G_{\max} = 40$ mT/m, and (b) f for one subject imaged using a maximum gradient strength of $G_{max} = 60$ mT/m.

As in section 6.3, we focus on a comparison between the current state of the art voxelwise estimation and the RVM. We note that, based on the results obtained in section 6.2.3, if the smooth curves were used, they would be more variable than the RVM, given that $G_{max} = 40$, 60 mT/m.

6.4.3 Results

From our preliminary experiment to choose the most appropriate values for $d_{\|}$, we find that for the two subjects imaged using a maximum gradient strength of 60 mT/m, the optimal values are $d_{\|} = 1.82, 1.68\times\ 10^{-9}$ m^{2}s^{-1} respectively. For the four subjects imaged using a maximum gradient strength of 40 mT/m, the optimal values are $d_{\|} = 1.58, 1.74, 1.58, 1.51\times\ 10^{-9}$ m^{2}s^{-1}. These diffusivities are then used to fit the ZeppelinCylinder model to the data using both independent voxelwise and the RVM. As for the ex vivo data sets we find that around each minima there is a range of diffusivities, typically extending to $0.05\times\ 10^{-9}$ m^{2}s^{-1} either side of the minimum, for which the likelihood function is reasonably stable. In figure 6.17, we show example plots of the likelihood as a function of $d_{\|}$ for one subject imaged with $G_{max} = 40$ mT/m and one subject imaged at $G_{max} = 60$ mT/m.

Figures 6.18 and 6.19 show parameter maps for R and f, along with corresponding standard deviation maps, for the first subject (female, age: 25) acquired using $G_{max} = 60$ mT/m using voxelwise estimation and the RVM respectively. In this case, the voxelwise estimation shows a low-high-low-high trend in R instead of the low-high-low trend expected from histology [2], with a reasonable number of voxels in the splenium and a couple of voxels in the genu showing relatively high ($\approx\ 4\mu$m) axon radius indices. Mean values of R range from $2 - 3\mu$m in the genu, 5μm in the midbody and $3 - 4\mu$m in the genu. In the isthmus, the thinnest part of the CC between the midbody and the splenium, R is approximately 2μm. Overall, the parameter map is reasonably noisy with several voxels, particularly in the genu and splenium indicating axon radius indices close to zero, i.e. $R \approx 0.1 - 0.2\mu$m. From the corresponding standard deviation map, we see that there is high uncertainty in these estimates, ranging from $\sigma_R = 0.5 - 1.25\mu$m. The RVM also picks up this low-high-low-high trend in R, although the

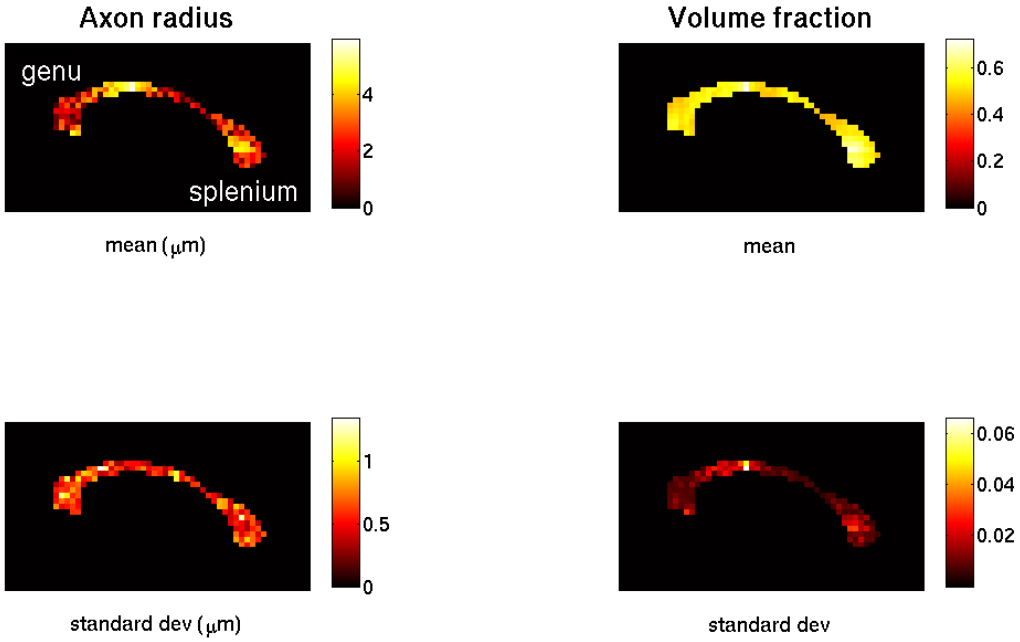

Figure 6.18: Voxelwise parameter maps of the mean and standard deviation of R, f for subject 1 from $G_{max} = 60$ mT/m in vivo human brain data.

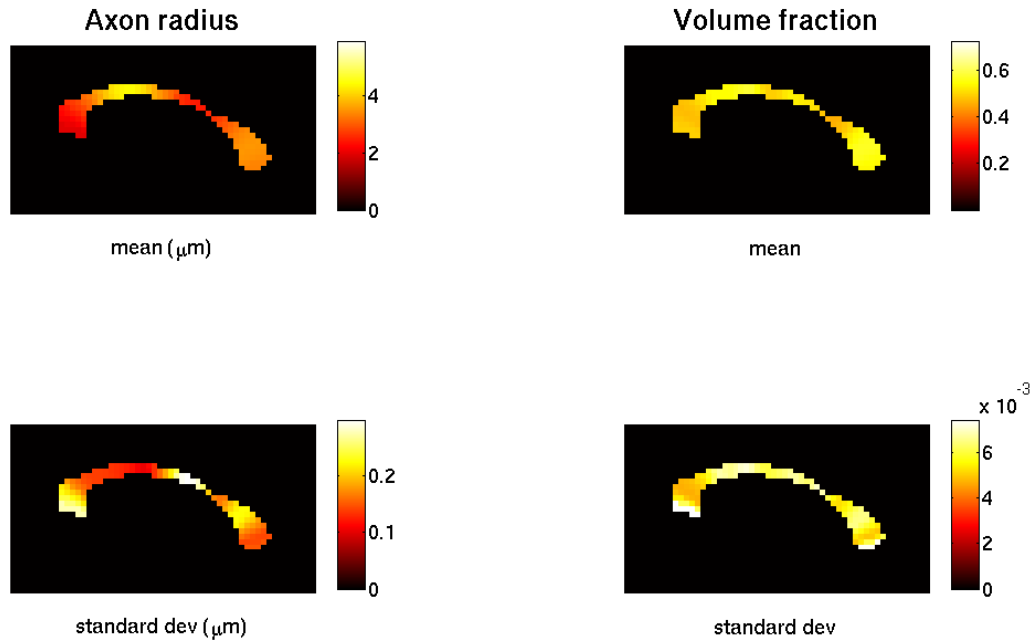

Figure 6.19: RVM parameter maps of the mean and standard deviation of R, f for subject 1 from $G_{max} = 60$ mT/m in vivo human brain data.

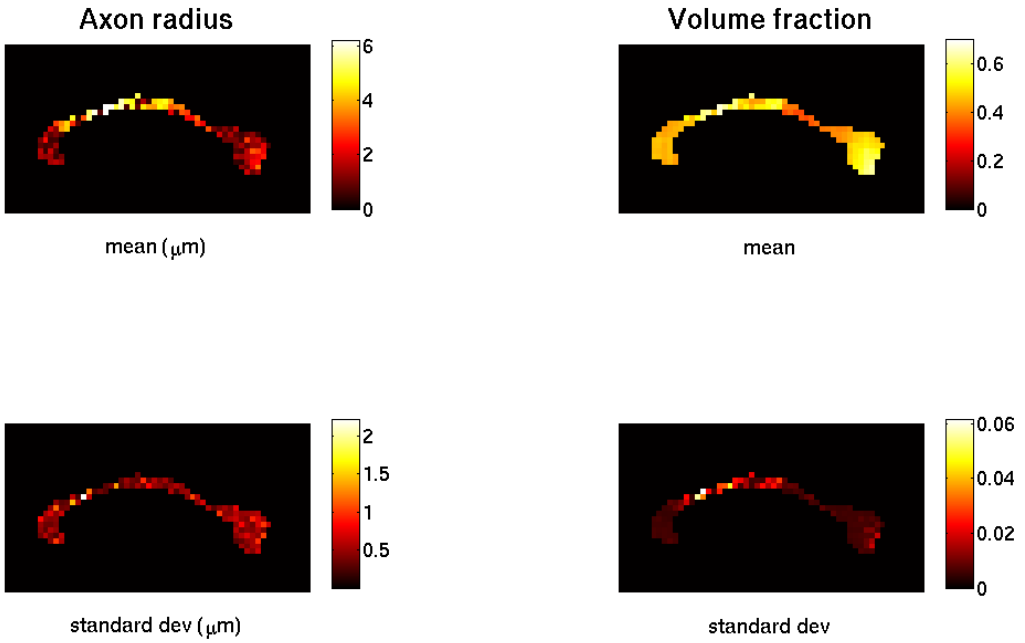

Figure 6.20: Voxelwise parameter maps of the mean and standard deviation of R, f for subject 2 from $G_{max} = 60$ mT/m in vivo human brain data.

standard deviation in the axon radius index trend is much tighter, ranging from $\sigma_R = 0.1 - 0.3\mu$m. The regions of higher uncertainty corresponds to regions with lower axon radius index, for which the signal is less sensitive to changes in R. Both methods show very similar results for volume fraction which is approximately constant with $f = 0.5$. Once again, the standard deviation of the estimated RVM trend is much lower than the standard deviation of the voxelwise estimates, with values of $\sigma_f = 0.005 - 0.006$ compared to $\sigma_f = 0.01 - 0.02$ for independent voxelwise fitting.

We see similar results for the second subject (male, age: 32) acquired using $G_{max} = 60$ mT/m, whose parameter maps for voxelwise and RVM estimation are shown in figures 6.20 and 6.21, although this time both methods capture the expected low-high-low trend for R. Both independent voxelwise estimation and the RVM estimate axon radius indices of approximately 2μm in the genu, rising to 5μm in the midbody and falling to $2 - 3\mu$m in the splenium. These values compare favourably with previous results in the mid-sagittal CC obtained from different datasets also acquired using $G_{max} = 60$ mT/m [7, 200], although they are much higher than the values measured from histology [2]. As for the previous subject however, the maps of R are much noisier for voxelwise estimation than for the RVM and the uncertainties in the estimates are also much higher; $\sigma_R = 0.5 - 2\mu$m for voxelwise fitting compared to $\sigma_R = 0.1 - 0.3\mu$m for the RVM. For the volume fraction estimates, the results are again similar to those obtained for subject 1. Both methods estimate f at an approximately constant value of $0.4 - 0.6$. Once again, the standard deviation for the RVM is lower than for independent voxelwise estimation, with $\sigma_f = 0.004 - 0.012$ compared to $\sigma_f = 0.01 - 0.06$.

Figures 6.22 and 6.23 show parameter and standard deviation maps for R and f for a subject (fe-

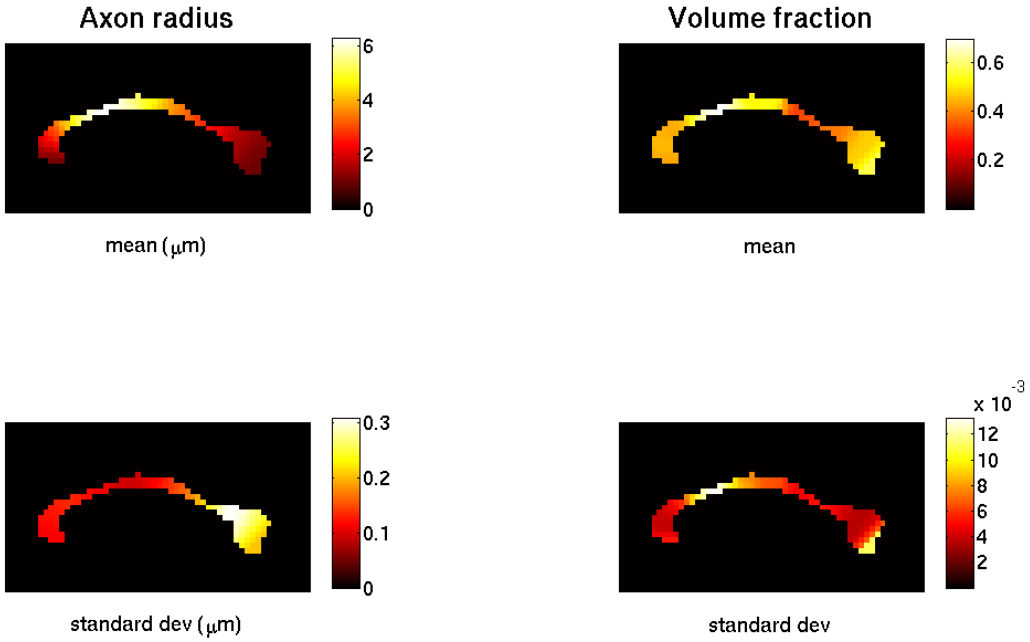

Figure 6.21: RVM parameter maps of the mean and standard deviation of R, f for subject 2 from $G_{max} = 60$ mT/m in vivo human brain data.

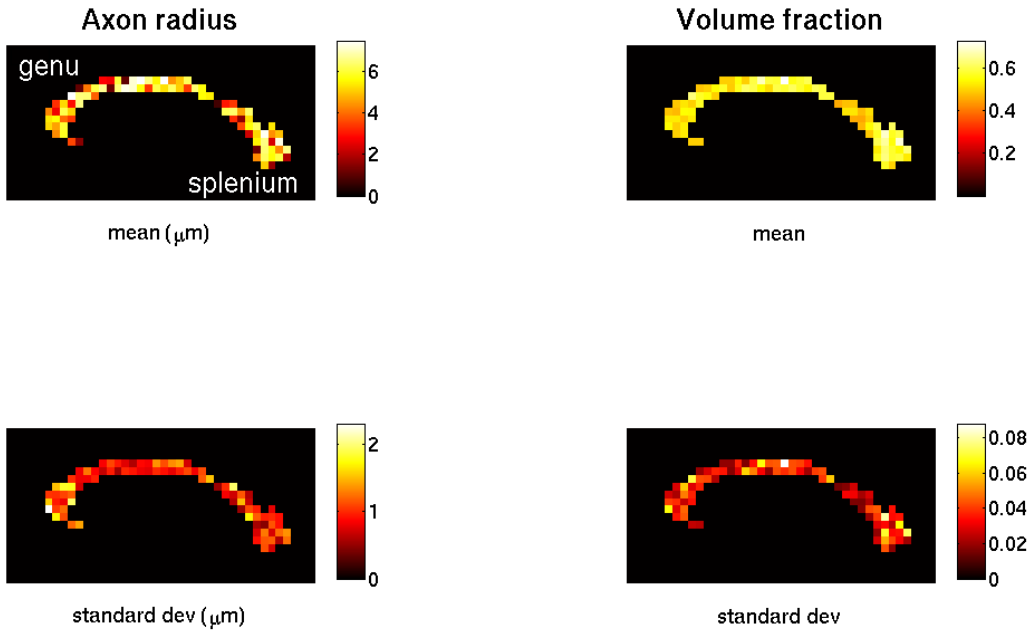

Figure 6.22: Voxelwise parameter maps of the mean and standard deviation of R, f for subject 1 (out of 4) from $G_{max} = 40$ mT/m in vivo human brain data.

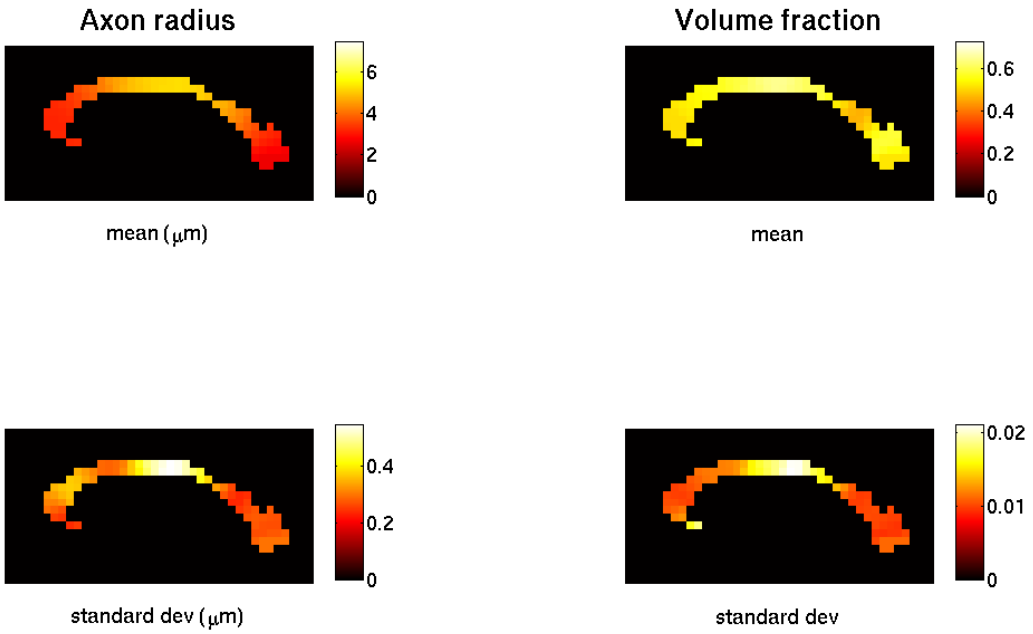

Figure 6.23: RVM parameter maps of the mean and standard deviation of R, f for subject 1 (out of 4) from $G_{max} = 40$ mT/m in vivo human brain data.

male, age: 26) acquired using a maximum gradient strength of 40 mT/m using voxelwise and RVM estimation. Voxelwise parameter maps of the axon radius index are very noisy, with values of R ranging from $1 - 7\mu$m across the CC, and we cannot identify any meaningful pattern in how R varies across the region of interest. The uncertainty in these measurements is also very high, with $\sigma_R = 1 - 2\mu$m. From the RVM estimates, we do see a low-high-low trend in R, with values of approximately 3μm in the genu and splenium and 5μm in the midbody, and standard deviations ranging from $\sigma_R = 0.3 - 0.5\mu$m. As for the subjects imaged with $G_{max} = 60$mT/m, the volume fraction maps are consistent between methods ($f = 0.4 - 0.6$), although once again, the standard deviation using the RVM ($\sigma_f = 0.01 - 0.02$) is much lower than for voxelwise estimation ($\sigma_f = 0.02 - 0.08$).

Figures 6.24 and 6.25 show the corresponding figures for another subject (male, age: 35), which is more representative of the remaining three subjects. The voxelwise maps of R are again very noisy, with a high standard deviation ($\sigma_R = 1 - 2\mu$m) and it is difficult to identify any pattern in how R varies across the region. Unlike the first subject, in which we observe the expected low-high-low trend in R, we do not see this trend here. Instead, the axon radius index is relatively high in the genu ($R \approx 5\mu$m) and decreases across the CC, with $R \approx 2\mu$m in the splenium. For one of the remaining subjects, the RVM estimates a similar high-low trend; for the remaining subject, the RVM estimates a low-high trend across the CC. For both methods, the maps of f are consistent across methods and subjects, although once again, the RVM results in lower standard deviations.

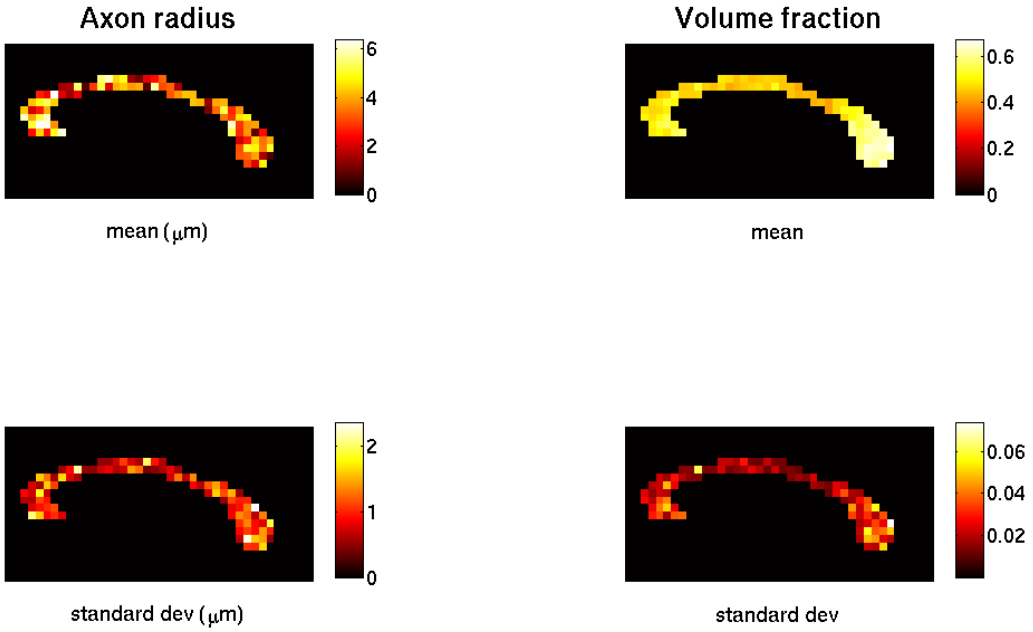

Figure 6.24: Voxelwise parameter maps of the mean and standard deviation of R, f for subject 2 (out of 4) from $G_{max} = 40$ mT/m in vivo human brain data.

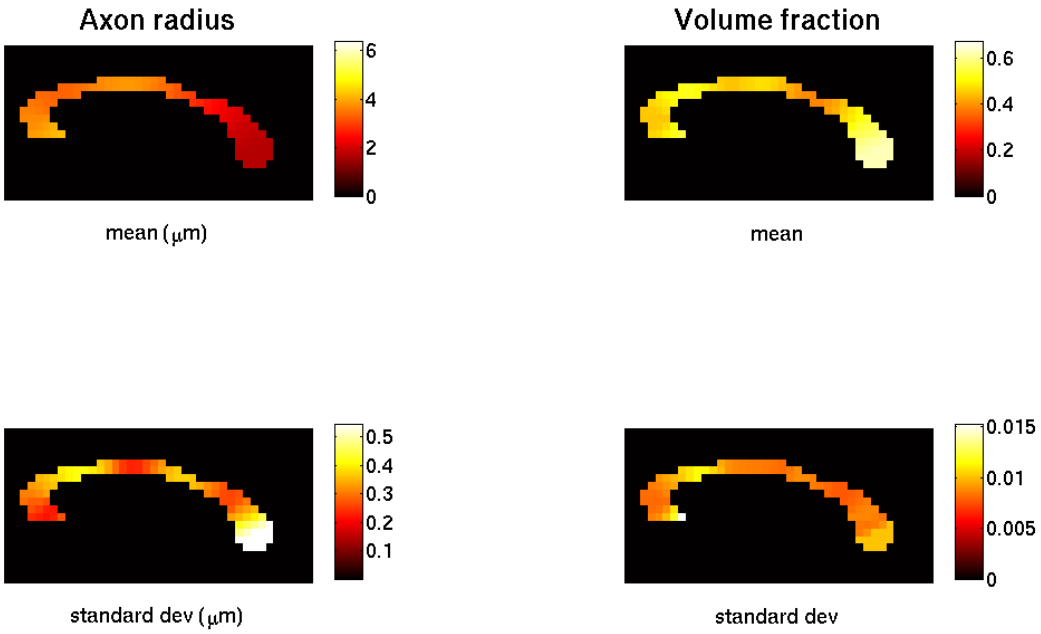

Figure 6.25: RVM parameter maps of the mean and standard deviation of R, f for subject 2 (out of 4) from $G_{max} = 40$ mT/m in vivo human brain data.

6.4.4 Discussion

At the beginning of this section, we hypothesised that the RVM would improve the precision of the axon radius compared with the current state of the art voxelwise estimation techniques. From the results obtained here, for data obtained using $G_{max} = 60$ mT/m this appears to be true. For both subjects, both methods were able to extract anatomically plausible trends of axon radius index and volume fraction variation across the mid-sagittal CC; however the standard deviation associated with these estimates was much smaller for the RVM compared with independent voxelwise estimation. Therefore, at gradient strengths of at least $G_{max} = 60$ mT/m, we believe that the RVM could be used to extract clinically useful trends of axon radius index.

At $G_{max} = 40$ mT/m, whilst the standard deviation of R and f are lower for the RVM than the equivalent values obtained in each voxel independently, the estimated trends in R are not anatomically realistic for three of the four subjects analysed here. For the first subject, the RVM extracts a low-high-low trend across the CC, with values of R ranging from 3μm in the genu and splenium to R=5μm. However, for the other three subjects the RVM estimates are either relatively flat or low-high trends across the CC. Whilst our initial simulations in section 6.2 suggested that the RVM would be able to robustly extract trends in axon radius from data acquired at $G_{max} = 40$ mT/m and with SNR=10, we must remember that such simulations are very idealised. In reality, white matter tissue is much more complex and thus the sensitivity to R weaker. Therefore, the results of this experiment suggest that at a gradient strength of $G_{max} = 40$ mT/m, the axon radius index is not sensitive to the underlying trend in axon radius over the ROI; we only begin to gain sensitivity at the slightly higher gradient strength of 60 mT/m.

6.5 Conclusions

In this chapter we modified the RVM for use in estimating direct microstructure parameters such as the axon radius index. We hypothesised that, compared to independent voxelwise estimation and the smoothing techniques previously discussed, the RVM would be able to improve the precision of axon radius estimates, particularly at clinical gradient strengths of 40 or 60 mT/m.

Initial simulation experiments were used to determine the optimal number of knots and smoothing priors to use in the RVM for different data characteristics such as gradient strength and noise level. Overall we found that, out of the diffusion parameters modelled with the RVM, the trends in volume fraction were robust to the RVM spatial settings, whereas the priors and knots for estimating trends in R need to be more finely tuned. Using the simulated data and optimal RVM settings, we compared trends estimated using the RVM compared to those estimated using voxelwise estimation and smooth curves. In terms of robustness to noise over multiple noise trials, the RVM provides an advantage for estimating R over simple voxelwise estimation for gradient strengths of $G_{max} = 160$ mT/m or lower. The RVM provides an advantage over smooth curves for gradient strengths of 80 mT/m or lower. Given that this is the current operating range for clinical scanners, this implies that the RVM could provide real advantages for estimating axon radius indices precisely in a clinical environment. This is supported by ex vivo data

experiments. For all gradient strengths, the standard deviations of R and f are lower when estimated using the RVM instead of voxelwise estimation. Due to artifacts in the data, we are unable to extract the expected trends in R at $G_{max} = 60$ mT/m data; we do recover the expected anatomical trends for the higher gradient strengths. In vivo human data experiments conducted using a maximum gradient strength of 60 mT/m support use of the RVM over the current state of the art voxelwise estimation. However at $G_{max} = 40$ mT/m, which is on the upper limit of the current standard gradient strength for clinical scanners (ranging between 20 and 40 mT/m, depending on the age of the scanner), the axon radius index is not sensitive to the underlying trend in axon radius. Thus a maximum gradient strength $G_{max} = 40$ mT/m may be too low for estimating a useful axon radius index, although to confirm this in future we would like to repeat the experiment at 40 mT/m on the Philips scanner in order to quantify the effect of scanner manufacturer on our results.

Chapter 7

Conclusions

At the beginning of this thesis, we identified the following key problem for the medical imaging community. Due to the inherent noise present in image data, it is often difficult to obtain accurate and precise estimates of informative tissue properties on a voxel by voxel basis. This is particularly problematic for diffusion MRI, which attenuates the MR signal due to the displacement of water molecules, lowering the SNR of the data. This affects our estimation of all diffusion MR tissue parameters, in particular those parameters that only weakly influence the diffusion MR signal. This includes direct microstructure parameters such as axon radius and density indices, which provide much more specific information about the underlying tissue and are potentially very useful for studying white matter diseases. However, our current inability to estimate these measures robustly hinders their clinical impact and utility.

In this work, we have addressed this problem by introducing a parameter estimation framework, the RVM, that exploits the underlying spatial coherence of the tissue microstructure in order to improve the estimation and mapping of diffusion MR tissue parameters. Throughout this thesis, we have tested the RVM using both simulations and real brain data and demonstrated that our approach provides some clear advantages over standard voxelwise techniques, particularly with regards to estimating direct microstructure indices and population analyses.

In chapter 3, we introduced the regional variation model framework, the key theoretical contribution of this thesis. This overview of the RVM was an important step in addressing the problem as it described how to exploit the underlying spatial coherence of the tissue in an appropriate way, using prior knowledge of the shape of the tract and the location of each voxel within the tract. This allowed us to pool measurements from voxels that should have similar underlying tissue properties whilst keeping the data from voxels with different underlying tissue properties separate. We also described how the functions describing the spatial variation of each parameter are linked, via a multi-level forward model that uses both shape information and an appropriate MR signal model, to the measured signals. In this work, we focused on diffusion MR signal models as our aim was to improve parameter estimation of diffusion MR tissue properties. However the framework itself is quite general and signal models for other MR contrasts or imaging modalities could be easily substituted into the framework. Therefore, as well as contributing to the project aims by providing a formulation of the RVM for diffusion MR, we believe this chapter provided a more general contribution to the wider medical imaging community by

introducing a new technique for estimating regional trends in medical imaging biomarkers.

The second key contribution of the thesis was to provide an implementation of the RVM framework that is both flexible and easy to use. In chapter 3, we described the implementation of the RVM in Matlab which currently includes a range of spatial models (polynomials, B-splines) and the full hierarchy of models in [134]. In chapters 4 and 6 we presented and utilised three specific implementations of the RVM with diffusion MR signal models of varying complexity and demonstrated, in both simulation and real brain data, that the RVM is consistently able to extract consistently stable and robust parameter trends, regardless of the diffusion MR signal model used.

The final key aim and contribution of this work was to demonstrate exemplar applications of the RVM framework, and to compare the RVM against current state of the art techniques. We have focused on two key applications: group studies and the estimation of direct microstructure parameters.

In chapter 5, we showed how the parameter trends estimated using the RVM can be used in population analyses in order to detect tissue changes due to pathology or tissue development. In simulation we showed that, over a wide range of group sizes, effect sizes and noise levels, the RVM provides improved spatial localisation compared to ROI analysis and is more robust to noise than TBSS. These general conclusions are supported by experiments using real brain data. This is an important contribution to the project aims, but also has potentially useful implications for the diffusion MR community in general. Diffusion MRI is a powerful tool for investigating and quantifying the structure of brain white matter tissue due to its sensitivity to the displacement of water molecules, and thus diffusion MRI parameters are often used in group studies in order to study the progression of white matter diseases such as schizophrenia and MS. In order to provide useful information for diagnosis, prognosis and treatment, the methods used in the group study must be both sensitive and precise. Given the localisation and robustness of the RVM compared with current state of the art methods, methods like the RVM, which exploit the spatial coherence of the data, provide a very promising avenue of research for future group studies.

Finally, in chapter 6 we showed how the RVM can be used to improve estimates of direct microstructure parameters such as an axon radius index. Using both simulations and real brain data, we showed that at high gradient strengths ($G_{max} = 200 - 300$ mT/m) available on pre-clinical scanners, both the RVM and voxelwise estimation perform similarly. At $G_{max} = 60$mT/m, which is on the upper limits of what is currently achievable on clinical scanners, the RVM showed clear advantages over independent voxelwise estimation. The trends in axon radius index estimated were similar using both methods, but the RVM was more robust to noise (over multiple noise trials) and estimated the trend with higher precision. However at a lower gradient strength of $G_{max} = 40$mT/m, we find that, even when using the RVM to fully exploit the underlying coherence of the white matter microstructure, the axon radius index we estimated is no longer sensitive enough to the true underlying trend in axon radius. This has important implications as estimating specific microstructure indices such as axon radius is one of the key goals of the diffusion MRI community. In particular, we would like to be able to estimate these indices within the clinic in order to obtain more detailed information about white matter microstructure changes due to disease or development. These results suggest that in order to do this, scanners with

gradient strengths of at least 60 mT/m are needed. Furthermore, in order to estimate these indices with enough precision to detect significant differences in axon radius due to pathology, techniques such as the RVM are needed.

7.1 Future work

Further to the contributions presented in this thesis, there are improvements and modifications that can be made to the current RVM framework in order to address the problem set out in chapter 1. In this section, we provide an overview of some of the key avenues for future work, some of which we have already done preliminary work towards.

The current implementation of the RVM estimates the mean trend of the diffusion MR indices without taking into account the variability perpendicular to the medial axis. One of our initial assumptions was that indices would be identical along this direction; in reality they will not be the same, only very similar. Although optimisation with MCMC does provide us with confidence intervals, these indicate the limits on where the mean could lie, rather than how variable the tissue properties are about the mean. Therefore, it would be useful to separate this variability from the image noise and estimate a variance trend as well. The variance could be a useful biomarker itself, particularly in diseases in which the degeneration of axons might leave the mean radius unchanged but reduce the variance. We think this would be a unique feature of the RVM which would set the method apart from voxelwise estimation and subsequent smoothing. Because the RVM fits regional trends directly to the data, we can consider the effects of image noise and 'noise' due to parameter variability simultaneously and attempt to separate the two. This is not possible when fitting splines to voxelwise estimates, in which we cannot separate the two effects. We have already done some preliminary work to investigate this approach, using error propagation techniques such as the unscented Kalman filter [100] in order to determine how variability in the ground truth tissue parameters introduces variability in the signals. Using this approach on simulated data, we have been able to estimate the variance due to noise and separate parameter variance curves for volume fraction and diffusivity simultaneously. To do this well we need enough voxels perpendicular to the medial axis. For thin tracts such as the CC, we often do not have enough data in a single; in this case we can only estimate a single parameter variance across the whole tract. However, this can be improved when extending the RVM to three dimensions as discussed below, and we think that a variance trend would be a useful addition to the RVM framework and deserves further work in order to formulate and implement it properly.

Another very important area of future work is to extend the RVM from two dimensions to three dimensions. So far, our work has focused on the two-dimensional mid-sagittal slice of the CC. However, to maximise the potential applications and usage for the RVM, we need to ensure that it works in three dimensions so that it is able to estimate trends in all tracts within the brain. Again, we have done some initial work towards this aim and the correct code exists within the current framework to fit B-spline surfaces instead of curves. However, we have only done some simple testing using CC-like ground truth phantoms and the ball and stick diffusion model. This required us to model the variation of the fibre orientation angles θ and ϕ as well as the diffusivity and volume fraction. Preliminary results suggest this

can be done with a modest number of gradient directions and some added noise; in some cases it seems that the extra data from voxels along the third dimension can help to stabilise the fitting. This could be particularly useful for estimating axon radius indices, which shouldn't change along the fibres, as we can include extra constraints. However, as we begin to consider fibre regions that begin to bend and fan, we need alternative models such as [200] in order to capture the fibre dispersion. Once this has been fully investigated using simulated data, we need to demonstrate that it produces sensible results from in vivo data on a range of tracts. For structures such as the CST, the width of which changes along the length of the tract, special care will need to be taken when deciding how to represent the surfaces.

The Bayesian P-spline representation used here can be easily replaced by alternative basis functions. Although we originally chose the Bayesian P-spline approach due to its limited support (i.e. each basis function is non-zero over only a small range) and flexibility, for noisy data this can be a hindrance and may lead us to oversmooth or undersmooth. Alternative basis functions such as Chebyshev polynomials, or in the case of periodic data, Fourier basis functions are defined over the whole axis, and might help to relieve some of these problems. This would also remove the need to optimise knot vectors and smoothing priors, although the number of basis functions to use still remains an issue.

The RVM can also be used in conjunction with other constrast techniques or imaging modalities. Although the current framework is focused on diffusion-weighted MRI, the diffusion model can be easily replaced with mathematical models relating tissue properties e.g. indices for relaxivities or rates of chemical uptake, to measured signals without any disruption or modification to the shape model or spatial model components.

Finally, as demonstrated in chapters 4 and 5, the RVM has potential for use in large, multimodal group studies such as ADNI, where the quality of the DTI data acquired is low, due to the small number of gradient directions used and no repeat measurements due to lack to time in the acquisition protocol, and may mask the underlying tissue changes between control and patient groups. Once the methodology behind the RVM is improved further, we would like to apply it to large scale, freely available data sets such as ADNI, in order to see whether the RVM can probe deeper into the data and identify microstructure changes unnoticed by current techniques. However, it is important to consider when the RVM is suitable for group studies. The current implementation of the RVM, using Bayesian P-splines, is particularly suitable for detecting diffuse changes along tracts. However, the RVM may not be suitable for detecting very focal changes, which could be smoothed out. In these cases methods such as TBSS, which compare the data voxel by voxel, are likely to perform better. However, the RVM can be modified to deal with this by changing the basis functions or the way in which they are smoothed. It is still important to exploit the spatial coherence of the tissue microstructure throughout the rest of the tract, but at the focal change the RVM must allow sharp changes in the parameter trend. This could be done by incorporating variable smoothing [107] in the Bayesian P-spline representation, or by using alternative basis functions such as wavelets which naturally allow sharp edges.

Appendix A

Markov chain Monte Carlo convergence

In section 2.10.1 we discussed Markov chain Monte Carlo (MCMC) methods for estimating the posterior distribution over our model parameters. To recap, MCMC methods are stochastic algorithms that allow us to draw samples from probability distributions. A random walker exploring the parameter space are more likely to move towards high density regions of the probability distribution, thus allowing us to sample the likely parameter values from the resulting Markov chain. However, the starting position of the walker may not be close to the high density regions of the distribution, and so we must ignore the initial iterations until the chain has converged. This initial period is known as the burn in period and determining its length is key to the success of the MCMC algorithm. In this appendix, we describe and present results from some of our preliminary experiments to determine the appropriate burn-in for the RVM.

A.1 Methods

A.1.1 Synthetic Data

We test the convergence of the of the Metropolis-Hasting algorithm using synthetic data sets described in chapters 4 and 6, in order to determine the appropriate burn in intervals when fitting both the BallStick and the ZeppelinCylinder models. Details of the region of interest, ground truth parameter values and data synthesis can be found in the relevant chapters. For each model, we consider one low quality data set and one high quality data set. To test the convergence for the BallStick model we use one data set with $N = 60$, SNR=25 and one data set with $N = 15$, SNR=10, where N is the number of gradient directions. For the ZeppelinCylinder model, we use one data set with $G_{max} = 200$mT/m, SNR=25 and one data set with $G_{max} = 60$mT/m, SNR=10.

A.1.2 MCMC fitting

For each data set, we fit the model using different starting points in order to generate multiple chains. For the BallStick model, we generate four chains using four unique starting points. This initial values are:

1. $f = 0.3$, $d = 1.4 \times 10^{-9}$ m^2 s^{-1}, $\theta = -\frac{\pi}{2}$, $\phi = -\frac{\pi}{2}$.
2. $f = 0.5$, $d = 1.8 \times 10^{-9}$ m^2 s^{-1}, $\theta = 0$, $\phi = 0$.

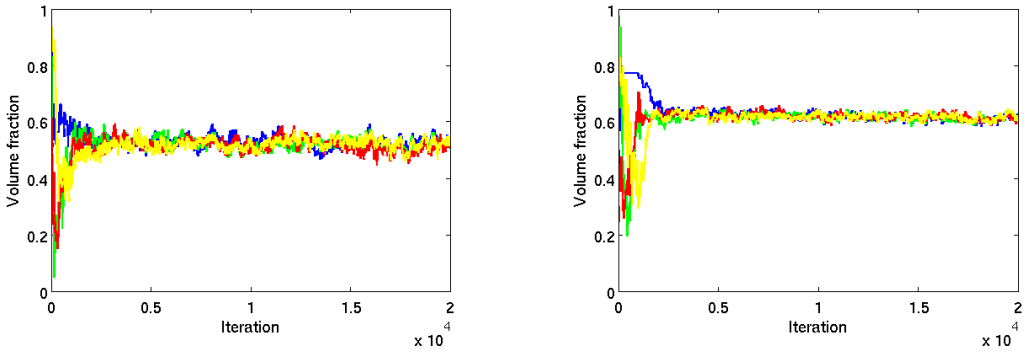

(a) (b)

Figure A.1: MCMC chains generated using different starting points for a representative spatial coefficient for f when (a) $N = 15$, SNR=10, and (b) $N = 60$, SNR=25.

3. $f = 0.7$, $d = 1.2 \times 10^{-9}$ m^2 s^{-1}, $\theta = \frac{\pi}{2}$, $\phi = \frac{\pi}{2}$.

4. $f = 0.9$, $d = 2 \times 10^{-9}$ m^2 s^{-1}, $\theta = \frac{\pi}{3}$, $\phi = \frac{\pi}{3}$.

Because the ZeppelinCylinder is more computationally complex and takes longer to run, we generate three chains with three unique starting points. These initial values are:

1. $R = 2\mu$m, $f = 0.3$, $\theta = -\frac{\pi}{2}$, $\phi = -\frac{\pi}{2}$.

2. $R = 3\mu$m, $f = 0.5$, $\theta = 0$, $\phi = 0$.

3. $R = 4\mu$m, $f = 0.7$, $\theta = \frac{\pi}{2}$, $\phi = \frac{\pi}{2}$.

When fitting the ZeppelinCylinder model, we keep the parallel diffusivity fixed throughout, as in chapter 6. We let each chain run for 20000 iterations, collecting every sample. We generate chains using a range of knot vectors and smoothing priors, but results here will focus on the optimal parameters determined in chapters 4 and 6. We also monitor the acceptance rate of the samples, adjusting the size of the steps the walker takes after every 1000 iterations in order to ensure acceptance rates between 15 and 30% as suggested in [70].

A.2 Results

Figure A.1 shows the evolution of the MCMC chains for a representative coefficient within the parameter vector $\mathbf{a}_f$ over all 20000 iterations for both low and high quality data. When $N = 15$, SNR=10, the chains converge after approximately 1500-2000 iterations. The number of iterations for convergence is similar for $N = 60$, SNR=25. After convergence the chains are noisier for the lower quality data set than for the higher quality data set.

Figure A.2 shows shows the evolution of the MCMC chains for a representative coefficient within the parameter vector $\mathbf{a}_d$ over all 20000 iterations for both low and high quality data. As we see for f, the chains converge after a similar number of iterations, approximately 1500, when $N = 60$, SNR=25 and

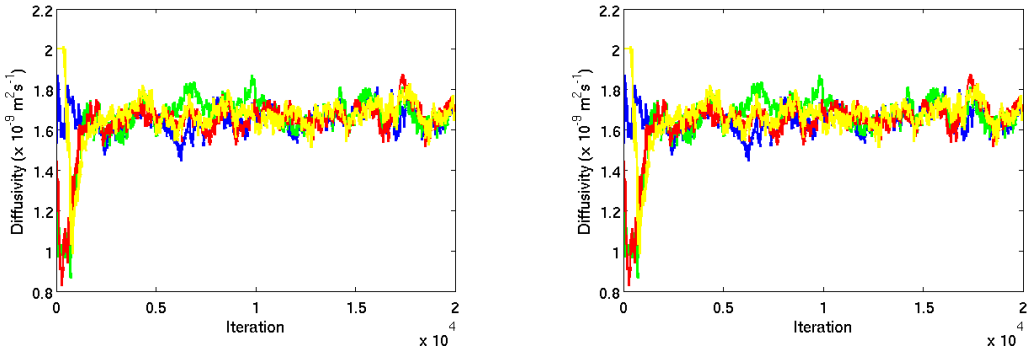

(a) (b)

Figure A.2: MCMC chains generated using different starting points for a representative coefficient for d when (a) $N = 15$, SNR=10, and (b) $N = 60$, SNR=25.

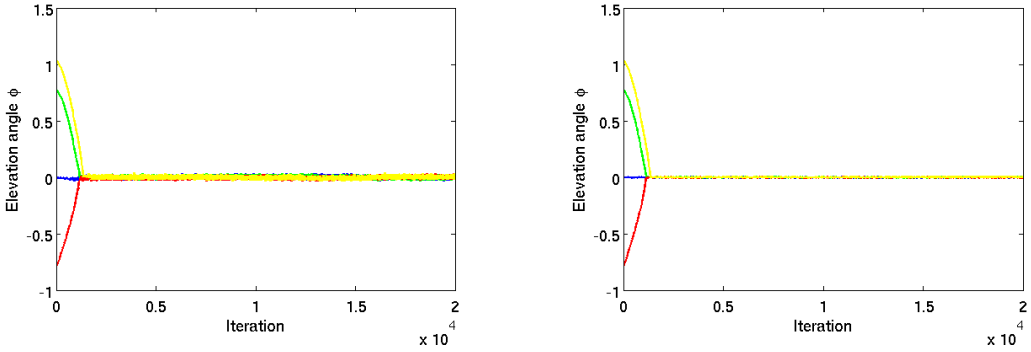

(a) (b)

Figure A.3: MCMC chains generated using different starting points for a representative coefficient for ϕ when (a) $N = 15$, SNR=10, and (b) $N = 60$, SNR=25.

when $N = 15$, SNR=10. Again, after convergence the chains are noisier for the lower quality data set than for the higher quality data set.

Figure A.3 shows corresponding plots for the evolution of the MCMC chains for a_ϕ. Again, regardless of the number of directions or SNR, the chains converge after roughly the same number of iterations, i.e. 1000-1500 iterations. After convergence, the chains are slightly noisier for the lower quality than higher quality data set, but it is minimal compared with the noise observed for f and d. Because $\phi = 0$, the other fibre orientation angle θ has no effect on the overall fibre direction. Therefore, it does not converge over the 20000 iterations, but this does not affect the overall fit of the model, and we do not show example plots of a_θ over all iterations.

Finally, figure A.4 corresponding plots for the evolution of the MCMC chains for a representative coefficient within the parameter vector $\mathbf{a}_R$. When the data is synthesised using a high gradient strength and high SNR the chains converge quickly, after 1000-2000 iterations. When the data is synthesised using a lower gradient strength and low SNR, the chain converges much slower, after 4000-5000 iterations.

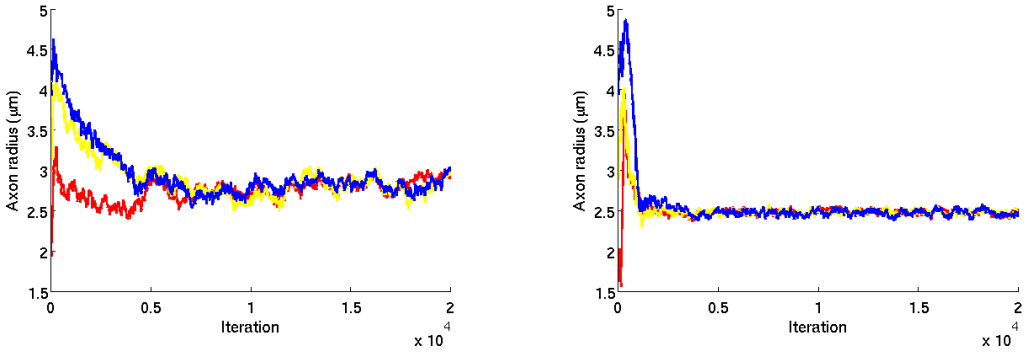

(a) (b)

Figure A.4: MCMC chains generated using different starting points for a representative coefficient for R when (a) $G_{max} = 60$mT/m, SNR=10, and (b) $G_{max} = 200$mT/m, SNR=25.

Again, the chains are noisier for the lower quality data.

We monitor the acceptance rates of all parameters over the 20000 iterations, and after roughly 3000-5000 iterations the acceptance rates stabilise at around 15-30%. Therefore on average we sample a new parameter vector over parameter space every 3 to 6 iterations. Therefore, to ensure that there is variation within our samples, our sampling interval should be greater than this. Figure A.5 shows posterior distributions for a representative example from $\mathbf{a}_f$, $\mathbf{a}_d$, a_ϕ and $\mathbf{a}_R$ comprising 500 samples drawn at intervals of 20 iterations using the lower quality data sets. All distributions are evenly centred about the mean, and indicate that 500 samples are enough to capture the shape of the posterior distributions.

A.3 Conclusions

In conclusion, we find that a burn in of 5000 iteration is enough to ensure convergence for all parameters and models considered here, even when the quality of the data is low. Therefore, we set our burn in to be equal to 5000 iterations for all experiments in this thesis, after which we obtain 500 samples at an interval of 20, which should ensure that we do not repeatedly sample the same point on the posterior distribution on consecutive samples.

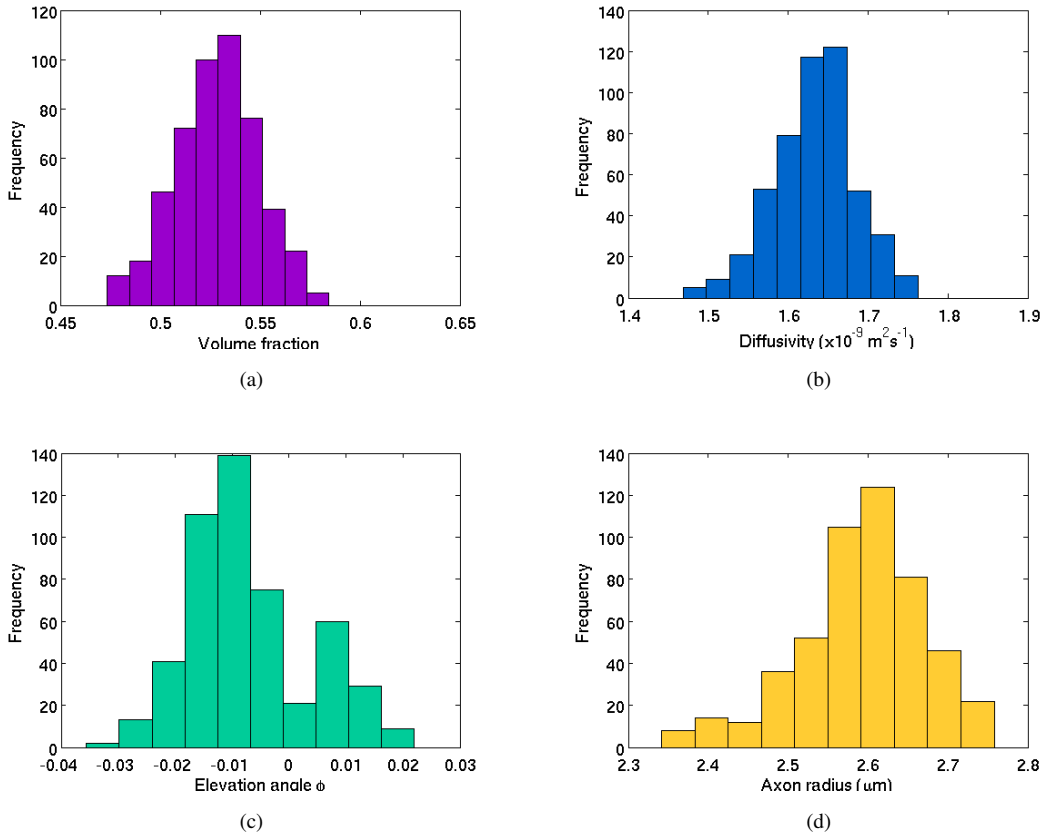

Figure A.5: Example posterior distributions consisting of 500 samples drawn at intervals of 20 iterations for (a) a coefficient from $\mathbf{a}_f$, (b) a coefficient from $\mathbf{a}_d$, (c) a_ϕ, and (d) a coefficient from $\mathbf{a}_R$. All distributions are generated from fitting the models to the lower quality data sets.

Appendix B

Choosing knot vectors

In section 2.9.1, we discussed methods for fitting B-splines to data. There are two main strategies for doing this. The first uses a small, but carefully positioned number of knots in order to reflect the underlying shape of the data. The second uses a large number of equally spaced knots, but penalises the weights of the basis functions in order to prevent overfitting to the noise. Within the current RVM framework, we use Bayesian penalised B-splines, which uses the latter approach. However, there is very little guidance on how many knots are enough. In their paper introducing Bayesian P-splines, Lang and Brezger [107] use a fixed knot vector with 40 knots for all experiments, although no justification is given. In earlier work on non-Bayesian P-splines, Eilers and Marx [62] suggest using one basis function per 4 data points, although this guideline is not rigorously tested. Ruppert et al [157] introduce a technique for determining the best number of knots using noisy data generated from a known ground truth. The P-spline is fit to the data using a wide range of knot vectors and the one that minimises the mean squared error (MSE) between the fitted spline and the ground truth values is found.

Unlike work by Ruppert et al [157], the RVM fits multiple splines to the data simultaneously, and we are trying to find the optimal smoothing priors as well as the best knot vector for a wide range of data sets generated with different acquisition protocols and SNRs. Therefore, it would be unfeasible to test a large number of knot vectors for every combination of protocol, SNR and priors. The aim of this experiment is to find a reduced set of knot vectors to use in the rest of our experiments which are representative of the much wider range of possible knot vectors, but allow us to reduce the amount of computational time required.

B.1 Methods

B.1.1 Simulated data

We use the ground truth parameter values described in chapters 4 and 6 in order to generate synthetic data sets. However, unlike previous data sets, in order to simplify the experiment only one of the parameters exhibits regional variation at a time. We create three ground truths overall: one in which f varies, one in which d varies and one in which R varies. In each ground truth, all other parameters are constant across the ROI. From each of these ground truths we generate low and high quality data sets using the BallStick model for the variable f and d ground truths, and the ZeppelinCylinder model for the variable

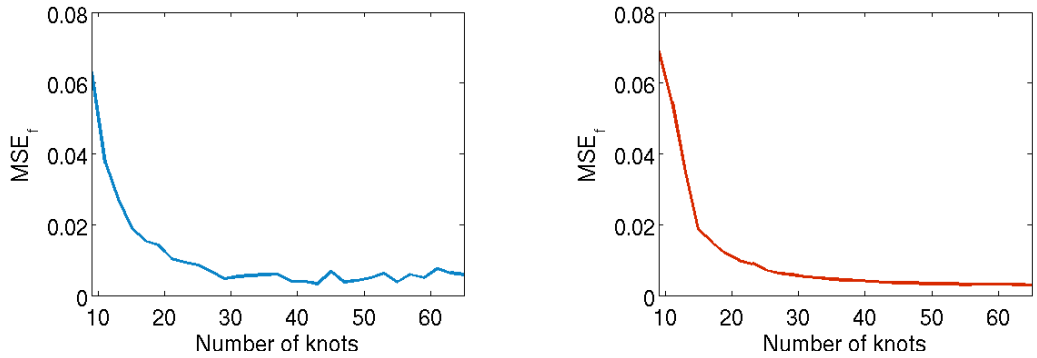

Figure B.1: Mean squared error in f as a function of number of knots calculated from data with (a) $N = 6$, SNR=10, and (b) $N = 60$, SNR=20.

R ground truth. When simulating data using the BallStick model, we generate the low quality data using $N = 6$ gradient directions and SNR=10 and the high quality data using $N = 60$ gradient directions and SNR=20. For the ZeppelinCylinder model, we synthesise low quality data usin $G_{max} = 40$mT/m and SNR=10 and the high quality data using $G_{max} = 200$mT/m and SNR=20. Full details of the protocols are presented in chapters 4 and 6.

B.1.2 Fitting the RVM

We fit the diffusion models to the data sets using the RVM as described previously, however we only use B-splines to model the variation of one model parameter at a time. The rest of the model parameters are fit as constants across the whole ROI. The Gamma-distributed smoothing prior is fixed throughout, with hyperparameters $\alpha = 4$ and $\beta = 25$, corresponding to moderate smoothing. For each of the data sets, we fit the RVM with knot vectors containing 9 knots (corresponding to 5 cubic B-spline basis functions covering the ROI) to 65 knots, in steps of 2 knots.

B.1.3 Statistical analysis

After fitting the basis function weights to each data set, we calculate the mean squared error between the spline fit and the ground truth parameter over all voxels in the ROI as described in chapter 4.

B.2 Results

Figures B.1, B.2 and B.3 show the MSE as a function of knot vector for f, d and R respectively, for both low and high quality data. When the models are fit to the high quality data, the trend in MSE is similar for all parameters is very similar. The MSE reduces sharply as the number of knots increases from 9 to 17. Between 17-39 knots, the MSE decreases still although the slope is much flatter. Above 39 knots, the MSE is almost flat. When the models are fit to the lower quality data, we still see a sharp decrease in MSE as the knots increases from 9 to 17, although it is a lot flatter between 21 and 45 knots. As the number of knots increases beyond 45, we even see a slight increase in MSE, particularly for R and d.

This suggests that the behaviour of a wide range of knot vectors could be represented by a small

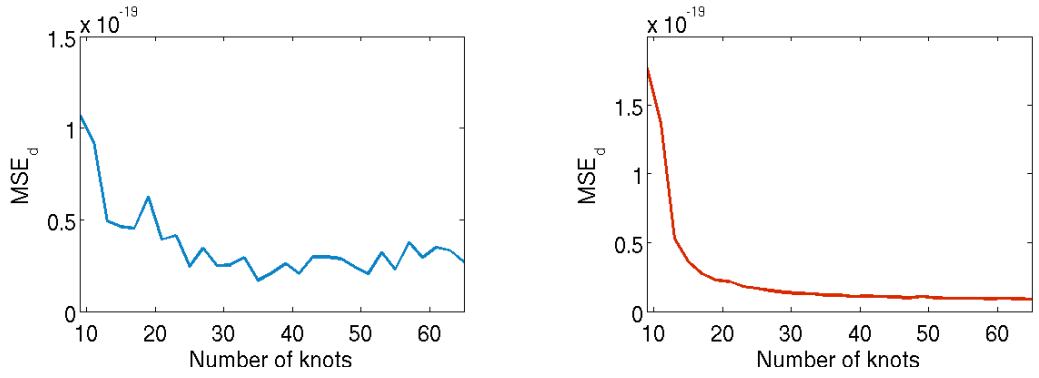

Figure B.2: Mean squared error in d as a function of number of knots calculated from data with (a) $N = 6$, SNR=10, and (b) $N = 60$, SNR=20.

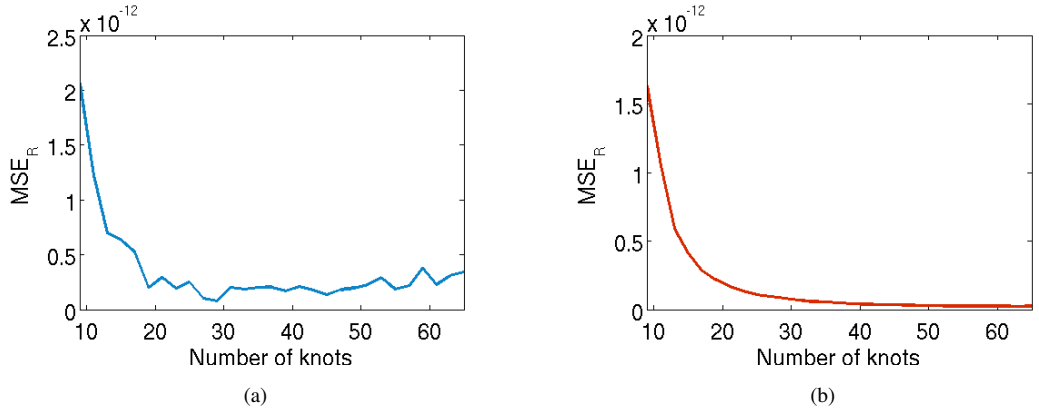

Figure B.3: Mean squared error in R as a function of number of knots calculated from data with (a) $G_{max} = 40$mT/m, SNR=10, and (b) $G_{max} = 200$mT/m, SNR=20.

number of knot vectors, e.g. one in the range 11-17 knots (9 knots performs particularly badly so we exclude it from consideration), one in the range 17-45 knots and one in the range 45-65 knots.

B.3 Conclusions

For f, d and R, the plots of MSE as a function of number of knots are very similar, for both high and low quality data. The plots suggest that knot vectors with similar number of knots behave similarly when fitting the RVM to the data, and thus we can use a smaller selection of knot vectors throughout the experiments which are representative of three main regimes. Therefore, throughout the rest of the experiments within this thesis we consider knot vectors with 15, 25 and 55 knots.

Although the knot vector with 15 knots performs reasonably poorly here, we think it is worth including as it is reasonably representative of the small number of knots, and may perform better when multiple splines are being fit or weaker smoothing priors are used. It is important to remember that this is a simplified experiment in which only one spline is fit to model the variation of one parameter, whilst the rest are estimated as constants. When multiple splines are being fit as part of the RVM simultaneously and with more variable smoothing priors, it is likely that the MSE will be worse, particular for knot vectors with large numbers of knots.

Bibliography

[1] O. Abe, S. Aoki, N. Hayashi, H. Yamada, A. Kunimatsu, H. Mori, T. Yoshikawa, T. Okubo, and K. Ohtomo. Normal aging in the central nervous system: quantitative mr diffusion-tensor analysis. *Neurobiology of aging*, 23(3):433–441, 2002.

[2] F. Aboitiz, A. B. Schiebel, R. S. Fisher, and E. Zaidel. Fiber composition of the human corpus callosum. *Brain Research*, 598:143–153, 1992.

[3] J. J. H. Ackerman and J. J. Neil. Biophysics of diffusion in cells. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[4] S. Aja-Fernandez, A. Tristan-Vega, and C. Alberola-Lopez. Noise estimation in single- and multiple-coil magnetic resonance data based on statistical models. *Magnetic Resonance Imaging*, 27(10):1397–1409, 2009.

[5] A. L. Alexander. Deterministic white matter tractography. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[6] D. C. Alexander. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine*, 60:439–448, 2008.

[7] D. C. Alexander, P. L. Hubbard, M. G. Hall, E. A. Moore, M. Ptito, G. J. Parker, and T. B. Dyrby. Orientationally invariant indices of axon diameter and density from diffusion MRI. *NeuroImage*, 52(4):1374–1389, 2010.

[8] D. C. Alexander and K. K. Seunarine. Mathematics of crossing fibers. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[9] C. Andersen. In vivo estimation of water content in cerebral white matter of brain tumour patients and normal individuals: towards a quantitative brain oedema definition. *Acta Neurochirurgica*, 139:249–256, 1997.

[10] C. Andrieu, N. De Freitas, A. Doucet, and M. I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50:5–43, 2003.

[11] R. Archibald and A. Gelb. A method to reduce the Gibbs ringing artifact in MRI scans while keeping tissue boundary integrity. *IEEE Transactions on Medical Imaging*, 21(4), 2002.

[12] B. A. Ardekani, J. Kershaw, M. Braun, and I. Kanno. Automatic detecton of the mid-sagittal plane in 3D brain images. *IEEE Transactions on Medical Imaging*, 16:947–952, 1997.

[13] B. A. Ardekani, J. Nierenberg, M. J. Hoptman, D. C. Javitt, and K. O. Lim. MRI study of white matter diffusion anisotropy in schizophrenia. *NeuroReport*, 14:2025–2029, 2003.

[14] J. Ashburner and K. J. Friston. Voxel-based morphometry - the methods. *NeuroImage*, 11:805–821, 2000.

[15] Y. Assaf, T. Blumenfeld-Katzir, Y. Yovel, and P. J. Basser. AxCaliber: A method for measuring axon diameter distribution from diffusion MRI. *Magnetic Resonance in Medicine*, 59:1347–1354, 2008.

[16] Y. Assaf, R. Z. Freidlin, G. K. Rohde, and P. J. Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52:965–978, 2004.

[17] V. Baladandayuthapani, B. K. Mallick, and R. J. Carroll. Spatially adaptive bayesian penalized regression splines (p-splines). *Journal of Computational and Graphical Statistics*, 14(2):378–394, 2005.

[18] B. Balinov, B. Jonsson, P. Linse, and O. Sodermann. The NMR self-diffusion method applied to restricted geometries: simulation of echo attenuation from molecules in spheres and between planes. *Journal of Magnetic Resonance, Series A*, 104:17–25, 1993.

[19] D. Barazany, P. J. Basser, and Y. Assaf. In vivo measurement of axon diameter distribution in the corpus callosum of rat brain. *Brain*, 132:1210–1220, 2009.

[20] P. J. Basser, J. Matiello, and D. Le Bihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.

[21] P. J. Basser and E. Ozarslan. Introduction to diffusion MR. In H. Johansen-Berg and T. E. J. Behrens, editors, *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*. Academic Press, 2009.

[22] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.

[23] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance Series B*, 111:209–219, 1996.

[24] C. Beaulieu. What makes diffusion anisotropic in the nervous system? In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[25] T. E. J. Behrens, H. Johansen-Berg, M. W. Woolrich, S. M. Smith, C. A. M. Wheeler-Kingshott, P. A. Boulby, G. J. Barker, E. L. Sillery, K. Sheehan, O. Ciccarelli, A. J. Thompson, J. M. Brady,

and P. M. Matthews. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nature Neuroscience*, 6(7):750–757, 2003.

[26] T. E. J. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion weighted MR imaging. *Magnetic Resonance in Medicine*, 50:1077–1088, 2003.

[27] M Belge, M. E. Kilmer, and E. L. Miller. Efficient determination of multiple regularization parameters in a generalized l-curve framework. *Inverse Problems*, 18:1161–1183, 2002.

[28] M. Bertero and P. Boccacci. *Introduction to Inverse Problems in Imaging*. Institute of Physics, 1998.

[29] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[30] F. Bloch. Nuclear induction. *Physical Review*, 70:460–474, 1946.

[31] H. Blum. Biological shape and visual science. *Journal of Theoretical Biology*, 38:205–287, 1973.

[32] C. De Boor. *Practical Guide to Splines*. University College London, 1978.

[33] M. Brown and R. Semelka. *MRI – Basic Principles*. Wiley, 2003.

[34] J. Burns, M. E. Bastin, H. Whalley, T. MacGillivray, and E. C. Johnstone. Structural disconnectivity in schizophrenia: a diffusion tensor magnetic resonance imaging study. *British Journal of Psychiatry*, 182:439–443, 2003.

[35] F. Calamante, D. L. Thomas, G. S. Pell, J. Wiersma, and R. Turner. Measuring cerebral blood flow using magnetic resonance imaging technqiues. *Journal of Cerebral Blood Flow and Metabolism*, 19:701–735, 1999.

[36] P. T. Callaghan. Pulsed-gradient spin-echo NMR for planar, cylindrical and spherical pores under conditions of wall relaxation. *Journal of Magnetic Resonance, Series A*, 113(1):53–59, 1995.

[37] P. T. Callaghan. Physics of diffusion. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[38] H. Y. Carr and E. M. Purcell. Effects on diffusion on free precession in nuclear magnetic resonance experiments. *Physical Review*, 94(3):629–638, 1954.

[39] O. Christiansen, T.-M. Lee, J. Lie, U. Sinha, and T. F. Chan. Total variation regularization ofmatrix-valued images. *International Journal of Biomedical Imaging*, 2007(3):753–768, 2007.

[40] P. L. Clatworthy, G. B Williams, J. Acosta-Cabronero, S. P. Jones, S. G. Harding, H. Johansen-Berg, and J. C. Baron. Probabilistic tractography of the optic radiations - an automated method and anatomical validation. *NeuroImage*, 49(3):2001–2012, 2009.

[41] J. D. Clayden, S. Zhang, S. Correia, and D. H. Laidlaw. Finely-grained comparison of anisotropy differences between groups of white matter tracts. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 82, 2007.

[42] E. Cohen, R. F. Riesenfeld, and G. Elber. *Geometric Modeling with Splines*. A K Peters, 1st edition, 2001.

[43] J. Colby, L. Soderberg, C. Lebel, I. D. Dinov, P. M. Thompson, and E. R. Sowell. Along-tract statistics allow for enhanced tractography analysis. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 2024, 2011.

[44] P. A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 2579, 2006.

[45] P. A. Cook, P. A. Boulby, M. R. Symms, and D. C. Alexander. Optimal acquisition order of diffusion-weighted measurements on a sphere. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 1303, 2005.

[46] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:681–685, 2001.

[47] T. F. Cootes and C. J. Taylor. Statistical models of appearance for medical image analysis and computer vision. In *Proceedings of SPIE Medical Imaging*, pages 236–248, 2001.

[48] I. Corouge, P. T. Fletcher, S. Joshi, S. Gouttard, and G. Gerig. Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Analysis*, 10:786–798, 2006.

[49] I. Corouge, S. Gouttard, and G. Gerig. A statistical shape model of individual fiber tracts extracted from diffusion tensor MRI. In *Proceedings of Medical Image Computing and Computer-assisted Intervention*, pages 571–579, 2004.

[50] O. Coulon, D. C. Alexander, and S. R. Arridge. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis*, 8(1):47–67, 2004.

[51] D. D. Cox and J. S. Lee. Pointwise testing with functional data using the westfall-young randomization method. *Biometrika*, 95:621–634, 2008.

[52] J. Crank. *The mathematics of diffusion*. Clarendon Press, 2nd edition, 1975.

[53] F. de Assis Aquino Gondim. Topographic and functional anatomy of the spinal cord. *http://emedicine.medscape.com/article/1148570-overview*, 2011.

[54] C. Demiralp and D. H. Laidlaw. Generalizing diffusion tensor model using probabilistic inference in Markov random fields. In *Proceedings of the computational diffusion MRI workshop, Medical image computing and computer-assisted intervention*, 2011.

[55] L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.

[56] Z. Ding, J. C. Gore, and A. W. Anderson. Reduction of noise in diffusion tensor images using anisotropic smoothing. *Magnetic Resonance in Medicine*, 53(2):485–490, 2005.

[57] D. C. Douglass and D. W. McCall. Diffusion in paraffin hydrocarbons. *Journal of Physical Chemistry*, 62:1102–1107, 1958.

[58] I. Drobnjak, B. Siow, and D. C. Alexander. Optimizing gradient waveforms for microstructure sensitivity in diffusion-weighted MR. *Journal of Magnetic Resonance*, 206:41–51, 2010.

[59] T. B. Dyrby, W. F. Barre, D. C. Alexander, J. Jelsing, E. Garde, and L. V. Sogaard. Pipeline for producing high-quality and high-resolution diffusion-weighted imaging datasets. *Human Brain Mapping*, 32(4):544–563, 2011.

[60] T. B. Dyrby, P. L. Hubbard, M. Ptito, M. G. Hall, and D. C. Alexander. Dependence of axon diameter index on maximum gradient strength. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 576, 2010.

[61] J. M. Edgar and I. R. Griffiths. White matter structure: A microscopist's view. In H. Johansen-Berg and T. E. J. Behrens, editors, *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*. Academic Press, 2009.

[62] P. H. C. Eilers and B. D. Marx. Flexible smoothing with B-splines and penalties. *Statistical Science*, 11:89–121, 1996.

[63] C. M. Filley. Neurobiology of white matter disorders. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[64] J. D. Foley. *Computer Graphics: principles and practice*. Addison-Wesley, 1995.

[65] J. Foong, M. R. Symms, G. J. Barker, M. Maier, D. H. Miller, and M. A. Ron. Investigating regional white matter in schizophrenia using diffusion tensor imaging. *NeuroReport*, 13:333–336, 2002.

[66] D. Gameran. *Markov Chain Monte Carlo*. Chapman and Hall, 1997.

[67] A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992.

[68] G. Gerig, S. Gouttard, and I. Corouge. Analysis of brain white matter via fibre tract modeling. In *Proceedings of the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 4421–4424, 2004.

[69] J. Geweke. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In *Bayesian Statistics*, pages 169–193. University Press, 1992.

[70] W. Gilks, S. Richardson, and D. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. Chapman and Hall, 1996.

[71] J. Goldsmith, C. M. Crainiceanu, B. S. Caffo, and D. S. Reich. Penalized functional regression analysis of white-matter tract profiles in multiple sclerosis. *NeuroImage*, 57:431–439, 2011.

[72] P. Golland, W. E. L. Grimson, and R. Kikinis. Statistical shape analysis using fixed topology skeletons: Corpus callosum study. In *Proceedings of Information Processing in Medical Imaging*, pages 382–287, 1999.

[73] G. Gong, T. Jiang, C. Zhu, Y. Zang, F. Wang, S. Xie, J. Xiao, and X. Guo. Assymetry analysis of cingulum based on scale-invariant parameterization by diffusion tensor imaging. *Human Brain Mapping*, 24:92–98, 2005.

[74] C. B. Goodlett, P. T. Fletcher, J. H. Gilmore, and G. Gerig. Group analysis of DTI fiber tract statistics with application to neurodevelopment. *NeuroImage*, 45:S133–S142, 2009.

[75] A. Graps. An introduction to wavelets. *Computational Science and Engineering*, 2(2):50–61, 1995.

[76] M. J. Graves. Magnetic resonance angiography. *British Journal of Radiology*, 70:6–28, 1997.

[77] H. Gudbjartsson and S S. Samuel Patz. The rician distribution of noisy mri data. *Magnetic Resonance in Medicine*, 34:910–914, 1995.

[78] E. M. Haacke, Y. Xu, Y.-C. N. Cheng, and J. R. Reichenbach. Susceptibility weighted imaging (SWI). *Magnetic Resonance in Medicine*, 52:612–618, 2004.

[79] P. Hagmann, M. Kurant, X. Gigandet, P. Thiran, V. J. Wedeen, R. Meuli, and J.-P. Thiran. Mapping human whole-brain structural networks with diffusion MRI. *PLoS One*, 2(7):230–238, 2007.

[80] M. G. Hall and D. C. Alexander. Convergence and parameter choice for monte-carlo simulations of diffusion MRI. *IEEE Transactions on Medical Imaging*, 28:1354–1364, 2009.

[81] R. H. Hashemi, W. G. Bradley, and C. J. Lisanti. *MRI: The basics*. Lippincott, Williams & Wilkins, 2nd edition, 2004.

[82] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[83] W. Van Hecke, A. Leemans, S. De Backer, B. Jeurissen, P. M. Parizel, and J. Sijbers. Comparing isotropic and anisotropic smoothing for voxel-based dti analyses: A simulation study. *Human Brain Mapping*, 31(1):98–114, 2010.

[84] J. R. Highley, M. M. Esiri, B. McDonald, M. Cortina-Borja, B. M. Herron, and T. J. Crow. The size and fibre composition of the corpus callosum with respect to gender and schizophrenia: a post-mortem study. *Brain*, 122:99–110, 1999.

[85] S. Hofer and J. Frahm. Topography of the human corpus callosum revisited: comprehensive fiber tractography using diffusion tensor magnetic resonance imaging. *NeuroImage*, 32:989–994, 2006.

[86] Y. Hu, I. Doudevski, D. Wood, M. Moscarello, C. Husted, C. Genain, J. A. Zasadzinski, and J. Israelachvili. Synergistic interactions of of lipids and myelin basic protein. *Proceedings of the National Academy of Science*, 101(37):13466–13471, 2004.

[87] K. M. Jansons and D. C. Alexander. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Problems*, 19:1031–1046, 2003.

[88] S. Jbabdi, T. E. J. Behrens, and S. Smith. Crossing fibres in tract-based spatial statistics. *NeuroImage*, 49:249–256, 2010.

[89] S. Jbabdi, M. W. Woolrich, and J. L. R. Anderssonand T. E. J. Behrens. A bayesian framework for global tractography. *NeuroImage*, 37(1):116–129, 2007.

[90] S. N. Jespersen, C. R. Bjarkam, J. R. Nyengaard, M. M. Chakravarty, B. Hansen, T. Vosegaard, L. Ostergaard, D. Yablonskiy, N. C. Nielsen, and P. Vestergaard-Poulsen. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Problems*, 19:1031–1046, 2003.

[91] S. N. Jespersen, C. D. Kroenke, L. Ostergaard, J. J. H Ackerman, and D. A. Yablonskiy. Modeling dendrite density from magnetic resonance diffusion measurements. *NeuroImage*, 34:1473–1486, 2007.

[92] J. B. Jonas, J. A. Muller-Burgh, U. M. Schlotzer-Schrehardt, and G. O. H. Naumann. Histomorphometry of the human optic nerve. *Investigative Opthamology and Visual Science*, 31(4):736–744, 1990.

[93] D. K. Jones. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49(1):7–12, 2003.

[94] D. K. Jones. The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study. *Magnetic Resonance in Medicine*, 51(4):807–815, 2004.

[95] D. K. Jones and P. J. Basser. Squashing peanuts and smashing pumpkins: How noise distorts diffusion-weighted MR data. *Magnetic Resonance in Medicine*, 52:979–993, 2004.

[96] D. K. Jones, M. Catani, C. Pierpaoli, and S. J. Reeves. Age effects on diffusion tensor magnetic resonance imaging tractography measures of frontal cortex connections in schizophrenia. *Human Brain Mapping*, 27:230–238, 2006.

[97] D. K. Jones, M. A. Horsfield, and A. Simmons. Optimal strategies for measuring diffusion in anisotropic systems by MRI. *Magnetic resonance in medicine*, 42(3):515–525, 1999.

[98] D. K. Jones, A. Simmons, S. C. Williams, and M. A. Horsfield. Non-invasive assessment of axonal fibre connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42:37–41, 1999.

[99] D. K. Jones, M. R. Symms, M. Cercignani, and R. J. Howard. The effect of filter size on VBM analyses of DT-MRI data. *NeuroImage*, 26(2):546–554, 2005.

[100] S. Julier and J. Uhlmann. Proceedings of the IEEE. *Magnetic Resonance in Medicine*, 94(3):401–422, 2004.

[101] A. Jullion and P. Lambert. Robust specification of the roughness penalty prior distribution in spatially adaptive bayesian p-spline models. *Computational Statistics & Data Analysis*, 51(5):2542–2558, 2007.

[102] E. Kaden, T. R. Knosche, and A. Anwander. Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion mr imaging. *NeuroImage*, 37(2):474–488, 2007.

[103] M. D. King, D. G. Gadian, and C. A. Clark. A random effects modelling approach to the crossing-fibre problem in tractography. *NeuroImage*, 44(3):753–768, 2009.

[104] C. G. Koay. Least squares approaches to diffusion tensor estimation. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[105] M. A. Koch and J. Finsterbusch. Compartment size estimation with double wave vector diffusion-weighted imaging. *Magnetic Resonance in Medicine*, 60(1):90–101, 2008.

[106] A-S. Lamantia and P. Rakic. Cytological and quantitative characteristics of four cerebral commissures in the rhesus monkey. *Journal of Comparative Neurology*, 291:520–537, 1990.

[107] S. Lang and A. Brezger. Bayesian p-splines. *Journal of Computational and Graphical Statistics*, 13(1):183–212, 2004.

[108] C. Laule, I. M. Vavasour, S. H. Kolind, D. K. B. Li, T. L. Traboulsee, G. R. W. Moore, and A. L. MacKay. Magnetic resonance imaging of myelin. *Neurotherapeutics*, 4:460–484, 2007.

[109] M. Lazar, D. M. Weinstein, J. S. Tsuruda, K. H. Hasan, K. Arfanakis, M. E. Meyerand, B. Badie, H. A. Rowley, V. Haughton, A. Field, and A. L. Alexander. White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18:306–321, 2003.

[110] P. M. Lee. *Bayesian statistics: an introduction*. Wiley-Blackwell, 3rd edition, 2004.

[111] T. C. M Lee. On algorithms for ordinary least squares regression spline fitting: a comparative study. *Journal of statistical computation and simulation*, 72:647–663, 2002.

[112] I. B. Levitan and L. K. Kaczmarek. *The neuron: cell and molecular biology*. Oxford University Press, 3rd edition, 2002.

[113] Z-P. Liang and P. Lauterbur. *Principles of Magnetic Resonance Imaging: A signal processing perspective*. IEEE Press, 2000.

[114] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 7th edition, 2008.

[115] M. Maddah, M. Kubicki, W. M. Wells, C.-F. Westin, M. E. Shenton, and W. E. L. Grimson. Findings in schizophrenia by tract-oriented DT-MRI analysis. In *Medical image computing and computer-assisted intervention*, pages 917–924. Springer, 2008.

[116] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11:431–441, 1963.

[117] J. A. McNab and K. L. Miller. Sensitivity of diffusion-weighted steady state free precession to anisotropic diffusion. *Magnetic Resonance in Medicine*, 60(2):405–413, 2008.

[118] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince. *MRI: From picture to proton*. Cambridge University Press, 2nd edition, 2007.

[119] K. D. Merboldt, W. Hanicke, and J. Frahm. Diffusion imaging of the human brain in vivo using high-speed STEAM MRI. *Magnetic Resonance in Medicine*, 23(1):179–192, 1992.

[120] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[121] P. P. Mitra and P. N. Sen. Effects of microgeometry and surface relaxation on NMR pulsed-field-gradient experiments: simple pore geometries. *Physical Review B*, 45(1):143–156, 1992.

[122] G. L. Morgan, R. D. Newbould, B. Whitcher, and D. C. Alexander. Polynomial models of the spatial variation of axon radius in white matter. In *Proceedings of the International Society of Magnetic Resonance in Medicine*, page 1564, 2010.

[123] G. L. Morgan, H. Zhang, B. Whitcher, and D. C. Alexander. A spatial variation model of white matter microstructure. In *Workshop on Computational Diffusion MRI, Medical Image Computing and Computer-assisted Intervention*, pages 165–173, 2010.

[124] G. L. Morgan, H. Zhang, B. Whitcher, and D. C. Alexander. A bayesian framework for modelling the regional variation of white matter microstructure. In *Proceedings of Medical Image Understanding and Analysis*, 2011.

[125] J. S. Murday and R. M. Cotts. Self-diffusion coefficient of liquid lithium. *Journal of Chemical Physics*, 48(11):4938–4945, 1968.

[126] R. M. Neal. Probabilistic inference using Markov chain Monte Carlo methods. Technical Report CRG-TR-931, University of Toronto, 1993.

[127] S. Nedjati-Gilani, G. J. M. Parker, and D. C. Alexander. Regularized super-resolution for diffusion MRI. In *Proceedings of the International Symposium on Biomedical Imaging: from nano to macro*, pages 14–17, 2008.

[128] C. H. Neuman. Spin echo of spins diffusing in a bound medium. *Journal of Chemical Physics*, 60(11):4508–4511, 1974.

[129] T. E. Nichols and A. P. Holmes. Non-parametric permutation tests for functional neuroimaging: a primer with examples. *Human Brain Mapping*, 15:1–25, 2001.

[130] J. Nolte. *The human brain: An introduction to its functional anatomy*. Mosby Elsevier, 6th edition, 2008.

[131] R. Olivares, J. Montiel, and F. Aboitiz. Differences and similarities in the fine structure of the mammalian corpus callosum. *Brain, Behavior and Evolution*, 57:98–105, 2001.

[132] M. Ota, T. Obata, Y. Akine, H. Ito, H. Ikehira, T. Asada, and T. Suhara. Age-related degeneration of corpus callosum measured with diffusion tensor imaging. *NeuroImage*, 31:1445–1452, 2006.

[133] E. Panagiotaki, H. M. Fonteijn, B. Siow, M. G. Hall, A. Price, M. F. Lythgoe, and D. C. Alexander. Two-compartment models of the diffusion MR signal in brain white matter. In *Medical image computing and computer-assisted intervention*, pages 329–336. Springer, 2009.

[134] E. Panagiotaki, T. Schneider, B. Siow, M. G. Hall, M. F. Lythgoe, and D. C. Alexander. Compartment models of the diffusion signal in brain white matter: A taxonomy and comparison. *NeuroImage*, 59(3):2241–2254, 2012.

[135] G. J. M Parker. Probabilistic fiber tracking. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[136] G. J. M. Parker, J. A. Schnabel, M. R. Symms, D. J. Werring, and G. J. Barker. Nonlinear smoothing for reduction of systematic and random errors in diffusion tensor imaging. *Journal of Magnetic Resonance Imaging*, 11(6):702–710, 2000.

[137] G. M. J. Parker, H. A. Haroon, and C. A. Wheeler-Kingshott. A framework for a streamline-based probabilistic index of connectivity (pico) using a structural interpretation of MRI diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18:242–254, 2003.

[138] O. Pasternak, N. Sochen, Y. Gur, N. Intrator, and Y. Assaf. Free water elimination and mapping from diffusion mri. *Magnetic Resonance in Medicine*, 62(3):717–730, 2009.

[139] A. Peters. The effects of normal aging on myelin and nerve fibers: A review. *Journal of Neurocytology*, 31:581–593, 2002.

[140] A. Pfefferbaum, E. V. Sullivan, M. Hedehus, K. O. Li, E. Adalsteinsson, and M. Moseley. Age-related decline in brain white matter anisotropy measured with spatially corrected echo-planar diffusion tensor imaging. *Magnetic Resonance in Medicine*, 44(2):259–268, 2000.

[141] J. Pipe. Pulse sequences for diffusion-weighted MRI. In H. Johansen-Berg and T. E. J. Behrens, editors, *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*. Academic Press, 2009.

[142] C. Poupon, J.-F. Mangin, C. A. Clark, V. Frouin, J. Regis, D. Le Bihan, and I. Bloch. Towards inference of human brain connectivity from MR diffusion tensor data. *Medical Image Analysis*, 5:1–15, 2001.

[143] W. H. Press, S. A Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The art of scientific computing*. Cambridge University Press, 2nd edition, 1995.

[144] G. Price, M. S. Bagary, M. Cercignani, D. R. Altmann, and M. A. Ron. The corpus callosum in first episode schizophrenia: a diffusion tensor imaging study. *Journal of Neurology, Neurosurgery and Psychiatry*, 76:585–587, 2005.

[145] W. S. Price. Pulsed field gradient nuclear magnetic resonance as a tool for studying translational diffusion: Part 1. basic theory. *Magnetic Resonance: An Educational Journal*, 9:299–336, 1997.

[146] S. Prima, S. Ourselin, and N. Ayache. Computation of the mid-sagittal plane in 3D brain images. *IEEE Transactions on Medical Imaging*, 21(2):122–138, 2002.

[147] S. J. D. Prince. *Computer Vision: Models, Learning, and Inference*. Cambridge University Press, 2012.

[148] R. Quester and R. Schroder. The shrinkage of the human brain stem during formalin fixation and embedding in paraffin. *Journal of neuroscience methods*, 75:81–89, 1997.

[149] J. D. Quirk, G. L. Bretthorst, T. Q. Duong, A. Z. Snyder, C. S. Springer, J. J. H. Ackerman, and J. J. Neil. Equilibrium water exchange between the intra- and extracellular spaces of mammalian brain. *Magnetic Resonance in Medicine*, 50:493–499, 2003.

[150] A. Raj, C. Hess, and P. Mukherjee. Spatial HARDI: Improved visualization of complex white matter architecture using bayesian spatial regularization. *NeuroImage*, 54(1):396–409, 2011.

[151] T. G. Reese, O. Heid, R. M. Weisskoff, and V. J. Wedeen. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magnetic Resonance in Medicine*, 49:177–182, 2003.

[152] M. Reisert, I. Mader, C. Anastasopoulos, M. Weigel, S. Schnell, and V. Kiselev. Global fiber reconstruction becomes practical. *NeuroImage*, 54(2):955–962, 2011.

[153] S. O. Rice. Mathematical analysis of random noise. *Bell System Technical Journal*, 24:46–156, 1945.

[154] G. O. Roberts and J. S. Rosenthal. Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.

[155] E. Robinson, A. Hammers, A. Ericsson, A. D. Edwards, and D. Rueckert. Identifying population differences in whole-brain structural networks: a machine learning approach. *NeuroImage*, 50:910–919, 2010.

[156] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J Hawkes. Non-rigid registration using free-form deformations: applications to breast mr images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.

[157] D. Ruppert. Selecting the number of knots for penalised splines. *Journal of Computational and Graphical Statistics*, 11:735–757, 2002.

[158] T. Schneider, C. A. M. Wheeler-Kingshott, and D. C. Alexander. In-vivo estimates of axonal characteristics using optimized diffusion MRI protocols for single fibre orientation. In *Proceedings of Medical Image Computer and Computer-assisted Intervention*, pages 623–630, 2010.

[159] J. Scholz, M. C. Klein, T. E. J. Behrens, and H. Johansen-Berg. Training induces changes in white-matter architecture. *Nature Neuroscience*, 12:1370–1371, 2009.

[160] L. L Schumaker. *Spline Functions: Basic Theory*. Wiley, 1981.

[161] T. M. Shepherd, J. J. Flint, P. E. Thelwall, G. J. Stanisz, T. H. Mareci, A. T. Yachnis, and S. J. Blackband. Postmortem interval alters the water relaxation and diffusion properties of rat nervous tissue – implications for MRI studies of human autopsy samples. *NeuroImage*, 44:820–826, 2009.

[162] A. J. Sherbondy, M. C. Rowe, and D. C. Alexander. Microtrack: An algorithm for concurrent projectome and microstructure estimation. In *Proceedings of Medical Image Computer and Computer-assisted Intervention*, pages 183–190, 2010.

[163] S. Skare and R. Bammer. EPI based pulse sequences for diffusion tensor MRI. In D. K. Jones, editor, *Diffusion MRI: theory, methods and applications*. Oxford University Press, 2011.

[164] S. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. Behrens. Tract-based spatial statistics: Voxelwise analysis of multisubject diffusion data. *NeuroImage*, 31:1487–1505, 2006.

[165] S. M. Smith and T. E. Nichols. Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neur*, 16:947–952, 2009.

[166] S. N. Sotiropoulos, T. E. J. Behrens, and S. Jbabdi. Ball and rackets: Inferring fiber fanning from diffusion-weighted mri. *NeuroImage*, 37(2):474–488, 2007.

[167] S. Standring, editor. *Gray's Anatomy: the anatomical basis of clinical practice*. Elsevier, 39th edition, 2005.

[168] G. J. Stanisz, G. A. Wright, R. M. Henkelman, and A. Szafer. An analytical model of restricted diffusion in bovine optic nerve. *Magnetic Resonance in Medicine*, 37:103–111, 1997.

[169] E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *Journal of Chemical Physics*, 42:288–292, 1965.

[170] J. Stepisnik. Time dependent self-diffusion by NMR spin echo. *Physica B*, 183:343–350, 1993.

[171] H. Sun, P. A. Yushkevich, H. Zhang, P. A. Cook, J. T. Duda, T. J. Simon, and J. C. Gee. Shape-based normalization of the corpus callosum for DTI connectivity analysis. *IEEE Transactions on Medical Imaging*, 26:1166–1178, 2007.

[172] A. Szafer, J. Zhong, and J. C. Gore. Theoretical model for water diffusion in tissues. *Magnetic Resonance in Medicine*, 33:697–712, 1995.

[173] J. E. Tanner and E. O. Stejskal. Restricted self-diffusion of protons in colloidal systems by the pulsed-gradient, spin-echo method. *Journal of Chemical Physics*, 49(4):1768–1777, 1968.

[174] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. Direct estimation of the fiber orientaion density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*, 23:1176–1185, 2004.

[175] C. Tschumperle and R. Deriche. Diffusion tensor regularization with constraints preservation. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition*, pages 948–953. IEEE, 2001.

[176] D. S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[177] M. Unser. Splines: a perfect fit for signal and image processing. *IEEE Signal Processing*, 6:22–38, 2002.

[178] N. A. Uranova, D. Orlovskaya, O. Vikhreva, I. Zimina, N. Kolomeets, V. Vostrikov, and V. Rachmanova. Electron microscopy of oligodendroglia in severe mental illness. *Brain Research Bulletin*, 55:597–610, 2001.

[179] N. A. Uranova, V. M. Vostrikov, D. D. Orlovskaya, and V. I. Rachmanova. Oligodendroglial density in the prefrontal cortex in schizophrenia and mood disorders: a study from the stanley neuropathology consortium. *Schizophrenia Research*, 67:269–275, 2004.

[180] G. van Belle. *Statistical rules of thumb*. Wiley, 2002.

[181] M. S. van der Knapp and J. Valk. *Magnetic resonance of myelination and myelin disorders*. Springer, 3rd edition, 2005.

[182] P. van Gelderen, D. DesPres, P. C. van Zijl, and C. T. Moonen. Evaluation of restricted diffusion in cylinders phosphocreatine in rabbit leg muscle. *Journal of Magnetic Resonance, Series B*, 103:255–260, 1994.

[183] C. R. Vogel and M. E. Oman. Fast, robust, total variation-based reconstruction of noisy, blurred images. *IEEE Transactions on Image Processing*, 7(6):813–824, 1998.

[184] S. Walker-Samuel, M. Orton, J. K. Boult, and S. P. Robinson. Improving apparetn diffusion coefficient estimates and elucidating tumor heterogeneity using bayesian adaptive smoothing. *Magnetic Resonance in Medicine*, 65(2):438–447, 2011.

[185] Z. Wang, B. C. Vermuri, Y. Chen, and T. H. Mareci. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Transactions on Medical Imaging*, 23(8):930–939, 2004.

[186] G. S. Watson. The distribution of the ratio of two quadratic forms. *Australian Journal of Physics*, 8:402–407, 1955.

[187] S. G. Waxman, editor. *The axon: structure, function and pathophysiology*. Oxford University Press, 1995.

[188] A. Webb. *Introduction to Biomedical Imaging*. IEEE Press, 2003.

[189] D. M. Weinstein, G. Kindlmann, and E. Lundberg. Tensorlines: advection-diffusion based propagation through diffusion tensor fields. In *IEEE Visualization*, pages 249–253, 1999.

[190] C. Westbrook and C. Kautner. *MRI in Practice*. Blackwell, 2nd edition, 1998.

[191] C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108, 2002.

[192] S. F. Witelson. Thand and sex differences in the isthmus and genu of the human corpus callosum: a postmortem morphological study. *Brain*, 112:799–835, 1989.

[193] J. Xu, M. D. Does, and J. C. Gore. Sensitivity of MR diffusion measurements to variations in intracellular structure: Effects of nuclear size. *Magnetic Resonance in Medicine*, 61:828–833, 2009.

[194] R. Xue, P. C. Van Zijl, B. J. Crain, M. Solaiyappan, and S. Mori. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45:265–269, 1999.

[195] P. A. Yushkevich. Continuous medial representation for anatomical objects. *IEEE Transactions on Medical Imaging*, 25:1547–1564, 2006.

[196] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3d active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *NeuroImage*, 31:1116–1128, 2006.

[197] P. A. Yushkevich, H. Zhang, T. J. Simon, and J. C. Gee. Structure-specific statistical mapping of white matter tracts. *NeuroImage*, 41:448–461, 2008.

[198] H. Zhang and D. C. Alexander. Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. In *Proceedings of Medical Image Computer and Computer-assisted Intervention*, pages 640–647, 2010.

[199] H. Zhang, S. P. Awate, S. R. Das, J. H. Woo, E. R. Melhem, J. C. Gee, and P. A. Yushkevich. A tract-specific framework for white matter morphometry combining macroscopic and microscopic tract features. *Medical Image Analysis*, 14:666–673, 2010.

[200] H. Zhang, P. L. Hubbard, G. J. M. Parker, and D. C. Alexander. Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. *NeuroImage*, 56(3):1301–1315, 2011.

[201] H. Zhang, T. Schneider, C. A. M. Wheeler-Kingshott, and D. C. Alexander. NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage*, 61(4):1000–1016, 2012.