Quantifying Brain Microstructure with Diffusion MRI: An Assessment of Microscopic Anisotropy Imaging

Leevi Kerkelä

Doctor of Philosophy

Great Ormond Street Institute of Child Health

University College London

2020

Declaration

I declare that this thesis was written by myself, that the work contained herein is my own except where explicitly stated otherwise in the text.

Leevi Kerkelä

Abstract

Diffusion-weighted magnetic resonance imaging is routinely used for quantifying microstructural properties of brain tissue in both health and disease for its ability to be sensitive to the displacements of water molecules on a microscopic level. Significant effort has been put into the development of methods that provide more information on tissue microstructure than conventional diffusion tensor imaging. Multidimensional diffusion encoding methods render the signal sensitive to the displacements of water molecules that occur along two or three dimensions and can resolve some degeneracies in data acquired with single diffusion encoding methods that measure diffusion along a single dimension. The aim of this thesis is to study the state-of-the-art microstructural imaging methods and to assess their robustness in estimating microscopic diffusion anisotropy, i.e., the average anisotropy of the microscopic diffusion environments irrespective of their orientation dispersion, prior to their adoption in the wider neuroscience research community and possible deployment in clinical studies. First, a massively parallel Monte Carlo random walk simulator is presented. Second, the reproducibility of three commonly used microstructural models is quantified and the shortcomings of such single diffusion encoding methods in estimating microscopic diffusion anisotropy are addressed. Third, the challenges of estimating microscopic diffusion anisotropy in the human brain using double diffusion encoding are addressed using animal imaging experiments and simulations. The results support the feasibility of double diffusion encoding in human neuroimaging but raise hitherto overlooked precision issues when measuring microscopic diffusion anisotropy. Fourth, the accuracy and precision of microscopic diffusion anisotropy estimation using q-space trajectory encoding, a multidimensional diffusion encoding method specifically developed with the limitations of clinical whole-body scanners in mind, are assessed using imaging experiments and simulations. The results suggest that although broken model assumptions and time-dependent diffusion may bias the estimates, the effect of time-dependent diffusion on the estimated microscopic diffusion anisotropy is small in human white matter.

Impact statement

Since its invention in the 1990s, diffusion-weighted magnetic resonance imaging has become a valuable tool in neuroscience for safely and noninvasively studying microstructural properties of the human brain. The work presented in this thesis deepens our understanding of some of the most advanced methods for acquiring and analyzing diffusion-weighted magnetic resonance images. Specifically, the estimation of microscopic diffusion anisotropy, i.e., the average anisotropy of the microscopic diffusion environments irrespective of their orientation dispersion, is studied in detail. Microscopic diffusion anisotropy imaging is a novel method with promising results in imaging brain tumours and lesions. The results presented in this thesis provide valuable information on the robustness of microscopic diffusion anisotropy estimation methods and facilitate the planning of future studies that may benefit from applying these methods. Furthermore, the simulator developed by the PhD candidate takes advantage of the recent developments in both the software and hardware of general-purpose graphical processing unit computing to enable large amounts of synthetic diffusion-weighted magnetic resonance signals to be generated on standard desktop and laptop computers without needing to access high-performance computing clusters. The simulator is written in Python, a programming language known for its readibility and clarity, making it approachable and easily extensible to researchers and students with little or no prior experience in massively parallel computing. This contribution to the growing library of scientific Python packages facilitates the adoption of computing practices that utilize the astonishing computational power of modern graphical processing units.

Acknowledgements

I would like to express my gratitude to Chris Clark for welcoming me to his research group, to Matt Hall for providing guidance throughout this project, and to Jonathan Clayden for his valuable help. I would also like to thank NIHR GOSH BRC for providing funding. I owe my thanks to Noam Shemesh and Rafael Henriques with whom I had the pleasure of working at the Champalimaud Centre for the Unknown in beautiful Lisbon where I learned to operate pre-clinical scanners and surfboards. I would also like say thanks to Fabio, Hanne, Luis, Seb, Laura, Chloe, Melanie, and everyone else at DIBS for making the long, lonely days in front of the computer somewhat bearable, and to Noemi for the peer support in my native language. Also thanks to Lorena. Special thanks to Jessica for always being there for me. Finally, I would like to express my immense gratitude to my parents for having always supported me.

Table of contents

List of figures

List of tables

List of abbreviations

Abbreviation	Explanation
AD	Axial diffusivity
AF	Arcuate fasciculus
AK	Axial kurtosis
ADC	Apparent diffusion coefficient
BRC	Biomedical Research Centre
ConFiG	Contextual fibre growth
CPU	Central processing unit
CSD	Constrained spherical deconvolution
CST	Corticospinal tract
CSF	Cerebrospinal fluid
CT	Computed tomography
CV	Coefficient of variation
CTI	Correlation tensor imaging
dMRI	Diffusion-weighted magnetic resonance imaging
DDE	Double diffusion encoding
DKI	Diffusion kurtosis imaging
DSI	Diffusion spectrum imaging
DTI	Diffusion tensor imaging
δ	Diffusion encoding time
Δ	Diffusion time
EPI	Echo planar imaging
FA	Fractional anisotropy
FID	Free induction decay
FOV	Field of view
GOSH	Great Ormond Street Hospital
GPU	Graphical processing unit
HARDI	High angular resolution diffusion imaging

Abbreviation	Explanation
LTE	Linear tensor encoding
λ^{intra}	Intra-neurite diffusivity
λ^{extra}	Extra-neurite diffusivity
MD	Mean diffusivity
MDE	Multidimensional diffusion encoding
MK	Mean kurtosis
MRI	Magnetic resonance imaging
MSD	Mean squared displacement
MSK	Mean signal kurtosis
µA	Microscopic diffusion anisotropy
µAD	Microscopic axial diffusivity
µFA	Microscopic fractional anisotropy
µMD	Microscopic mean diffusivity
µRD	Microscopic radial diffusivity
NIHR	National Institute for Health Research
NMR	Nuclear magnetic resonance
NODDI	Neurite orientation dispersion and density imaging
OC	Orientation coherence
OD	Orientation dispersion index
ODF	Orientation distribution function
OR	Optic radiation
PET	Positron emission tomography
PGSE	Pulsed gradient spin echo
PTE	Planar tensor encoding
QTE	q-space trajectory encoding
QTI	q-space trajectory imaging
RMSD	Root mean squared displacement
RD	Radial diffusivity
RF	Radio frequency
RK	Radial kurtosis
ROI	Region of interest
SDE	Single diffusion encoding
SMT	Spherical mean technique
SMT1	Single-compartment spherical mean technique
SMT2	Two-compartment spherical mean technique

Abbreviation	**Explanation**
SNR	Signal-to-noise ratio
STE	Spherical tensor encoding
TDE	Triple diffusion encoding
TE	Echo time
TR	Repetition time
τ	Mixing time
UCL	University College London
v_{ic}	Intra-cellular volume fraction
v_{intra}	Intra-neurite volume fraction
v_{iso}	Extra-cellular volume fraction

List of publications

The following is a list of the peer-reviewed published work that the PhD candidate contributed in during the course of the doctorate. The contents of the articles and abstracts marked with an asterisk (*) are included in this thesis. The results presented in chapters 4, 6, and 7 have been published.

Journal articles

- **Leevi Kerkelä**, Fabio Nery, Ross Callaghan, Fenglei Zhou, Noemi G. Gyori, Filip Szczepankiewicz, Marco Palombo, Geoff J.M. Parker, Hui Zhang, Matt G. Hall, and Chris A. Clark. Comparative analysis of signal models for microscopic fractional anisotropy estimation using q-space trajectory encoding. *NeuroImage*, 242, 118445, 2021. DOI:10.1016/j.neuroimage.2021.118445*

- **Leevi Kerkelä**, Fabio Nery, Matt G. Hall, and Chris A. Clark. Disimpy: A massively parallel Monte Carlo simulator for generating diffusion-weighted MRI data in Python. *The Journal of Open Source Software*, 5(52), 2527, 2020. DOI:10.21105/joss.02527*

- **Leevi Kerkelä**, Rafael N. Henriques, Matt G. Hall, Chris A. Clark, and Noam Shemesh. Validation and noise robustness assessment of microscopic anisotropy estimation with clinically feasible double diffusion encoding MRI. *Magnetic Resonance in Medicine*, 83(5), 1698-1710, 2019. DOI:10.1002/mrm.28048*

- Fabio Nery, Filip Szczepankiewicz, **Leevi Kerkelä**, Matt G. Hall, Enrico Kaden, Isky Gordon, David L. Thomas, and Chris A. Clark. In vivo demonstration of microscopic anisotropy in the human kidney using multidimensional diffusion MRI. *Magnetic Resonance in Medicine*, 82(6), 2160-2168, 2019. DOI:10.1002/mrm.27869

Conference abstracts

- **Leevi Kerkelä**, Fabio Nery, Fenglei Zhou, Geoff J.M. Parker, Filip Szczepankiewicz, Matt G. Hall, Chris A. Clark. Accuracy and precision of microscopy anisotropy estimation using q-space trajectory encoding - a model comparison study. *The 27th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, virtual meeting, 2020.*

- Kan Yan Chloe Li, **Leevi Kerkelä**, Jonathan D. Clayden, Kiran K. Seunarine, Matt G. Hall, Chris A. Clark. Assessing scan-rescan reproducibility of surgically relevant white matter tractography reconstructions. *The 27th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, virtual meeting, 2020.

- Jiyoon Yoo, **Leevi Kerkelä**, Patrick W. Hales, Kiran K. Seunarine, Noemi G. Gyori, Enrico Kaden, and Chris A. Clark. High-resolution microscopic diffusion anisotropy imaging in the human hippocampus using multidimensional diffusion encoding. *The 27th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, virtual meeting, 2020.

- Pierre-Louis Bazin, Julia Huntenburg, Julia M. Huck, **Leevi Kerkelä**, Hoang Dung Do, Tristan Glatard, and Christopher Steele. Nighres: a Python toolbox for high-resolution neuroimaging. *2020 Organization for Human Brain Mapping Annual Meeting*, virtual meeting, 2020.

- **Leevi Kerkelä**, Rafael N. Henriques, Matt G. Hall, Chris A. Clark, and Noam Shemesh. Experimental validation and SNR analysis of a clinical double diffusion encoding sequence. *The 26th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, Montreal, Canada, 2019.*

- **Leevi Kerkelä**, Rafael N. Henriques, Matt G. Hall , Noam Shemesh, and Chris A. Clark. Precision of microscopic fractional anisotropy estimation using double diffusion encoding. *The 26th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, Montreal, Canada, 2019.*

Chapter 1

Introduction

1.1 Overview

The human brain is the central organ in the nervous system regulating most of the activities in the body based on the information it receives via the sensory nervous system. Its hundred billion or so neurons[1] form a densely connected network whose hierarchical organization spans length scales from individual synapses to the largest white matter fibre bundles connecting cortical regions[2]. In the 19th century, seminal scientific theories of brain function were inspired by the study of brain anatomy using post-mortem dissections in neurotypical patients and in those with brain lesions[3–5]. Comparing the lesions with the symptoms the patients had shown revealed that the typical functioning of the brain could be disrupted by white matter lesions causing disconnection between cortical areas. In modern neuroscience, noninvasive neuroimaging methods such as magnetic resonance imaging (MRI) are routinely used for measuring properties of the brain *in vivo*, enabling longitudinal studies investigating changes in the brain as a result of aging[6], degeneration[7], and learning[8], to name a few examples.

Since its invention in the 1970s[9,10], MRI has become an important tool both in research and clinical settings for its versatility and high spatial resolution. MRI is based on manipulating the nuclear magnetizations of atomic nuclei in a strong magnetic field and since it does not involve the use of ionizing radiation, it is relatively safe to apply on humans and is free from the harmful side effects of techniques such as computed tomography (CT) or positron emission tomography (PET). Typically, MRI experiments probe the concentration of hydrogen atoms, which are abundant in biological tissues, and the signal magnitude is

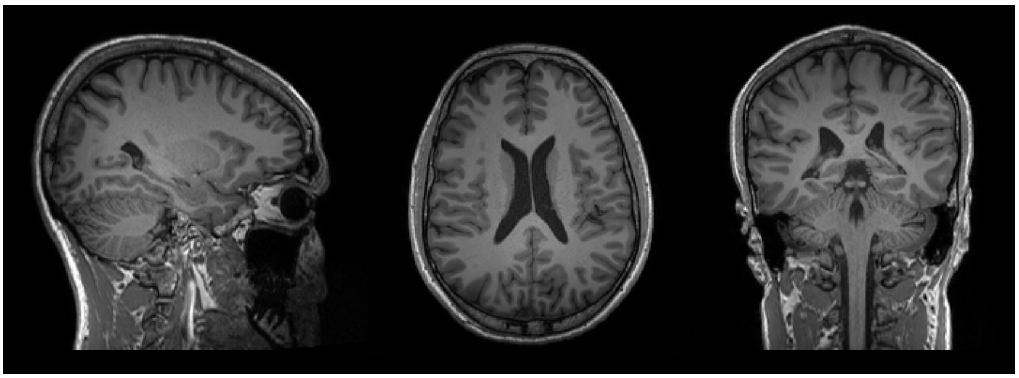

Figure 1.1: *Sagittal, axial, and coronal slices of a T_1-weighted magnetic resonance image of the human brain.*

affected by the used contrast mechanism. For example, slices of a T_1-weighted magnetic resonance image are shown in Figure 1.1. The contrast between white matter, gray matter, and cerebrospinal fluid (CSF) is very clear, reflecting the relaxation properties of the nuclei and the differences in water and lipid concentrations[11]. A general description of nuclear magnetic resonance (NMR) and image formation is given in Chapter 2.

Building on the pioneering work of measuring molecular displacements using NMR[12–14], diffusion-weighted MRI (dMRI) was introduced and applied in human neuroimaging in the 1980s[15–18]. The diffusion of water in biological tissues is restricted by microscopic obstacles such as cell organelles, myelin, and macromolecules. Because the displacements of water molecules during a typical dMRI experiment are in the micrometre range, the measured diffusion propagator is determined by the cellular level microstructure of the surrounding tissue[19,20]. Thus, important microstructural properties of the brain can be studied using dMRI. An overview of neural tissue microstructure is given in Section 3.1 and the basics of the physics of diffusion are described in Section 3.2. Section 3.3 describes how the NMR signal can be made sensitive to diffusion by the application of diffusion encoding magnetic field gradients.

The ability to localize measurements of microscopic diffusion in neural tissue *in vivo* in the form of maps at the millimetre resolution opened a new window to study the brain. The clinical value of dMRI was quickly highlighted by studies showing that the progression of acute stroke can be followed using maps of apparent diffusion coefficient (ADC), allowing therapeutic intervention on tissue that can still be salvaged and saving patients from permanent brain damage[21–23].

Since it had been shown that ADC depends on the direction of the applied diffusion encoding gradient[24] and thus a single scalar is not sufficient for characterizing anisotropic diffusion in white matter, diffusion tensor imaging (DTI) was introduced in 1994 as a way to quantify the magnitude and anisotropy of diffusion in the imaging voxel[25]. The anisotropy of the diffusion tensor, most often quantified with fractional anisotropy (FA), has been shown to be sensitive to degenerative processes in white matter such as axonal loss and demyelination[26–28]. DTI is described in Section 3.4.

Using the orientational information in the dMRI signal in white matter, it also became possible to estimate the positions of white matter tracts *in vivo*[29]. Since then, many techincal improvements have been made in the field of tractography[30–33], which has become a popular method in *in vivo* neuroanatomy[34], the study of brain connectivity[35,36], and presurgical planning[37,38], for instance.

Despite its proven utility, DTI suffers from various well-recognized limitations that have been addressed by the subsequent advances in the field. Because the tensor representation of diffusion confounds isotropic diffusion with orientation dispersion of anisotropic neurites, techniques like constrained spherical deconvolution (CSD)[39–41], q-ball imaging[42], and diffusion spectrum imaging (DSI)[43] were developed to obtain a more accurate estimate of the fibre orientation distribution function (ODF). Such methods accelerated the adoption of high angular resolution diffusion imaging (HARDI) acquisition methods, which involve sampling the signal along a large number of distinct diffusion encoding directions with the aim of characterizing orientational features of the dMRI signal[44].

Another major limitation of DTI is that it represents voxel-level diffusion as Gaussian and thus does not provide a good fit to data in experiments with moderate to high levels of diffusion-weighting (roughly $b > 1$ μm^2/ms in the brain) when the diffusion-attenuated signal decay clearly deviates from a monoexponential curve, revealing that the voxel-level average diffusion propagator is not Gaussian, especially in white matter. This becomes evident in experiments with multi-shell acquisition protocols, i.e., the acquisitions are performed with more than one level of non-zero diffusion-weighting. More complicated signal models such as the biexponential tensor model[45], diffusion kurtosis imaging (DKI)[46], and effective medium theory[47] have been developed to obtain a better fit to data and additional information on tissue microstructure. Furthermore, a large number of microstructural models have been developed with the aim of

extracting biophysically relevant properties of tissue microstructure by imposing strict *a priori* constraints on microstructure geometry[48–51]. Some methods beyond DTI for analyzing HARDI and multi-shell data are discussed in sections 3.5 and 3.6.

All the methods mentioned above belong to a wider class of single diffusion encoding (SDE) methods as they are based on acquisitions measuring the displacements of water molecules along a single dimension[52]. Despite being sensitive to several clinically and scientifically relevant changes in the brain microstructure, SDE acquisitions are fundamentally limited by their lack of specificity. They confound two major sources of voxel-level non-Gaussian diffusion, the orientation dispersion of anisotropic diffusion and the size variance of microscopic diffusion environments. Multidimensional diffusion encoding (MDE), on the other hand, render the dMRI signal sensitive to the displacements of water molecules along two or three dimensions, enabling the correlations of water molecule displacements along different directions to be measured. Over recent years, MDE methods, such as double diffusion encoding (DDE) and q-space trajectory encoding (QTE), have gained significant attention for their ability to resolve the fundamental degeneracy in data acquired with conventional SDE methods, potentially capturing clinically relevant information about tissue microstructure[53–61]. An overview of MDE methods is provided in Section 3.7.

Overall, an abundance of sophisticated mathematical and computational methods have been developed to estimate microstructural properties of neural tissue using dMRI. However, it can be argued that the translation of novel microstructural imaging methods from the dMRI research community to the wider neuroscience community and to the clinical setting has not kept up with the exponentially growing number of introduced methods[62]. Therefore, instead of developing yet another model, the focus of this thesis is on assessing the robustness of existing methods in order to facilitate their adoption and optimal use.

1.2 Research problem

The aim of this PhD thesis is to study the state-of-the-art of quantifying microstructural properties of brain tissue using dMRI and to focus on the estimation of microscopic diffusion anisotropy (μA), i.e., the average anisotropy of

the microscopic diffusion environments in the voxel irrespective of their orientation dispersion[52]. Since μA is not affected by orientation dispersion like the anisotropy indices calculated from DTI, it may be a suitable metric for quantifying white matter integrity. Indeed, promising results of μA imaging have been reported in imaging multiple sclerosis lesions[56,60] and white matter degeneration in Parkinson's disease patients[63]. It has also been applied to quantify the average cell shape in brain tumours[64,55] and microstructural properties of brains of patients with schizophrenia[65] and epilepsy[59]. The studies presented in this thesis assess the robustness of μA estimation methods and thus facilitate their deployment in the wider research community and possible future clinical studies. In particular, this thesis investigates microscopic fractional anisotropy (μFA), a normalized measure of μA that is equivalent to conventional FA in voxels where the microscopic diffusion environments are perfectly aligned. The central focus is on the different methods' abilities to accurately and precisely provide information on neural tissue microstructure within the time and hardware constraints that are typically present in clinical and research settings, imposing limits on the number of acquisitions and signal-to-noise ratio (SNR).

The results presented in this thesis, supported by both imaging experiments and simulations, show that μFA can be robustly estimated in white matter using MDE. However, existing methods can not reliabily map μFA in gray matter or other brain regions where its ground-truth value is low. In particular, the results show that precise measurement of μFA values less than 0.5 require SNR values that are unattainable using modern-day clinical whole-body scanners. Furthermore, the simulation experiments required the development of new software that enables complex dMRI simulations to be performed in massively parallel and that will be maintained and improved in as it is used in future studies.

1.3 Thesis structure

This thesis begins with this introductory chapter followed by an overview of NMR and MRI (Chapter 2). In Chapter 3, a comprehensive review of dMRI in the context of microstructural imaging is provided. The subsequent four chapters describe the original contributions by the PhD candidate. In Chapter 4, simulations of dMRI experiments are discussed and a massively parallel Monte Carlo random walk simulator is presented. In Chapter 5, the reproducibility

of DKI, neurite orientation dispersion and density imaging (NODDI)[66], and spherical mean technique (SMT)[67,68] is quantified using an established clinical SDE protocol and scan-rescan experiments. Furthermore, the limitations of estimating μA using such SDE methods are discussed. In Chapter 6, the limitations and challenges of applying DDE in human neuroimaging are discussed and a clinically feasible DDE acquisition protocol is validated using animal imaging experiments. Furthermore, the precision of μFA estimation is quantified using simulations. The results support the feasibility of DDE in clinical settings but raise hitherto overlooked precision issues when measuring μFA. In Chapter 7, the accuracy and precision of μFA estimation using QTE and different signal models are assessed using imaging experiments and simulations. The results show that broken model assumptions and time-dependent diffusion may introduce bias into the estimates, factors that must be carefully considered when μFA is applied in characterizing brain microstructure. However, the simulations suggest that the bias caused by time-dependent diffusion is small in human white matter. Chapter 8 finishes by providing conclusions, discussing limitations of the presented results, and describing possible directions for future research.

1.4 On notation

This thesis contains a large number of mathematical formulas in which the following conventions have generally been followed. Scalars are represented by Latin and Greek letters written in italics, e.g., x and α. Abbreviations of microstructural parameters like FA are written without italics. Tensors of all ranks are denoted by boldface upright letters in both lower and upper case, e.g., $\mathbf{b}$ and $\mathbf{D}$, with the exception of the quantum mechanical eigenstate vectors which are written in bra-ket notation[69], e.g., $|S\rangle$. Also, in Chapter 7, rank-4 tensors are denoted by upper case Latin letters written in blackboard bold, e.g., $\mathbb{C}$. However, blackboard bold is generally used to denote sets of numbers, e.g., $\mathbb{R}$ for the set of real numbers. When referring to one-dimensional vectors, $\mathbf{a}$ denotes a column vector and its transpose $\mathbf{a}^{\mathrm{T}}$ is a row vector. $\mathbf{ab}$ denotes a standard matrix multiplication between $\mathbf{a}$ and $\mathbf{b}$. $\cdot$ denotes an inner product and $\otimes$ denotes an outer product. : denotes a generalized scalar product, e.g., $\mathbf{A} : \mathbf{B} = \sum_i^N \sum_j^M A_{ij} B_{ij}$ for rank-2 tensors $\mathbf{A}$ and $\mathbf{B}$ with the same size $N \times M$. $\langle\,\rangle$ denotes averaging.

Chapter 2

Magnetic resonance imaging

Sections 2.1 and 2.2 provide a simplified summary of the physics that make MRI possible. A more detailed description of the theory presented here can be found in the textbooks by Griffiths[70] and Levitt[71].

2.1 Spin in a magnetic field

Despite the name "spin", which is a result of early theories explaining spin as a consequence of the rotation of a charged particle around its own axis, spin is now understood as an intrinsic quantum-mechanical property of a particle which can not be completely understood using classical concepts alone. This is illustrated by the fact that a measurement of spin is fundamentally non-deterministic. According to what is widely known as the Copenhagen interpretation of quantum mechanics, introduced by Bohr and Heisenberg, quantum systems do not have strictly defined properties prior to being measured[72]. Instead, quantum mechanics can describe solely the probabilities of different measurement outcomes. A quantum system, such as a particle with spin, is described by its wavefunction which prior to a measurement is in a superposition of eigenstates. Upon measurement, the superposition is reduced to an individual eigenstate through a phenomenon known as wave function collapse.

All atomic nuclei consisting of an odd number of protons and neutrons have non-zero total spin determined by the spin quantum number of the nucleus s. The spin is a quantized property, meaning that the measurement values are *a priori* constrained to take on one of a discontinuous set of values. The number of possible values is given by $2s + 1$. In this chapter, we will focus on the spin properties of the nucleus of hydrogen ^{1}H, i.e., a proton, for which $s = \frac{1}{2}$,

because it is the most common target in MRI.

Using the bra-ket notation introduced by Dirac[69], the general spin state $|S\rangle$ of a proton placed in a magnetic field aligned with the z-axis can be written as

$$|S\rangle = \alpha\,|\uparrow\rangle + \beta\,|\downarrow\rangle\,, \tag{2.1.1}$$

where α and β are the probability amplitudes corresponding to the eigenstates spin up $|\uparrow\rangle$ and spin down $|\downarrow\rangle$, respectively. The probability amplitudes are complex numbers whose modulus squared represents the probability of observing the spin in the corresponding eigenstate and which fulfil the condition $|\alpha|^2 + |\beta|^2 = 1$.

In quantum mechanics, observables, i.e., physical quantities that can be measured, are represented by linear operators whose eigenvalues are real numbers corresponding to the possible outcomes of a measurement. For the proton in the magnetic field aligned with the z-axis, the spin operator S_z has eigenvalues $\pm\frac{\hbar}{2}$ where $\hbar \approx 1.05 \cdot 10^{-34}$ J · s is the reduced Planck constant:

$$S_z\,|\uparrow\rangle = \frac{\hbar}{2}\,|\uparrow\rangle\,, \qquad\qquad S_z\,|\downarrow\rangle = -\frac{\hbar}{2}\,|\downarrow\rangle\,. \tag{2.1.2}$$

Equation 2.1.1 can be rewritten using the polar form of the probability amplitude:

$$\begin{aligned} |S\rangle &= r_\alpha e^{i\phi_\alpha}\,|\uparrow\rangle + r_\beta e^{i\phi_\beta}\,|\downarrow\rangle \\ &= e^{i\phi_\alpha}\left(r_\alpha\,|\uparrow\rangle + r_\beta e^{i\phi_\beta - \phi_\alpha}\,|\downarrow\rangle\right), \end{aligned} \tag{2.1.3}$$

which can be further simplified by taking into account that quantum states that only differ by a phase factor $(e^{i\phi_\alpha})$ are the same state. This gives

$$|S\rangle = r_\alpha\,|\uparrow\rangle + r_\beta e^{i\phi_\beta - \phi_\alpha}\,|\downarrow\rangle\,. \tag{2.1.4}$$

Now we can define angles θ and ϕ so that

$$|S\rangle = \cos\left(\frac{\theta}{2}\right)|\uparrow\rangle + \sin\left(\frac{\theta}{2}\right)e^{i\phi}\,|\downarrow\rangle\,, \tag{2.1.5}$$

where $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$ because of the normalization condition on the probability amplitudes and because $\cos^2(x) + \sin^2(x) = 1$ for all $x \in \mathbb{R}$. Equation 2.1.5 provides an analogy between a general superposition of spin-$\frac{1}{2}$ eigenstates and spherical coordinates, providing an intuitive way of visualizing Equation 2.1.1, known as a Bloch sphere which is shown in Figure 2.1.

The time evolution of the spin state is given given by the time-dependent Schrödinger equation:

$$i\hbar\frac{\partial}{\partial t}\,|S(t)\rangle = H\,|S(t)\rangle\,, \tag{2.1.6}$$

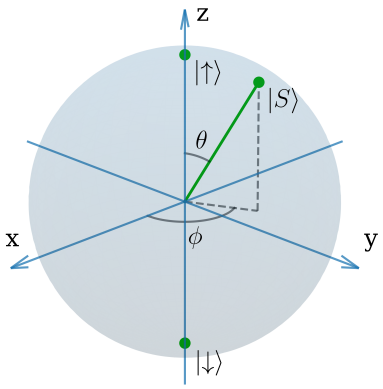

Figure 2.1: *Bloch sphere representation of a general spin-$\frac{1}{2}$ state $|S\rangle$ of a proton in a magnetic field aligned with the z-axis. The eigenstates of S_z, $|\uparrow\rangle$ and $|\downarrow\rangle$, are aligned with the z-axis, and any superposition of the eigenstates is on the surface of the sphere as defined by angles θ and ϕ.*

where $H = -\gamma B_0 S_z$ is the Hamiltonian operator where $\gamma \approx 2.68 \cdot 10^8$ rad$\cdot$s$^{-1}\cdot$ T^{-1} is the gyromagnetic ratio of ^{1}H and B_0 is the magnitude of the external magnetic field. The gyromagnetic ratio is a constant that relates the magnitude of the spin of a nucleus to its magnetic moment. Its value is different for different nuclei. The Hamiltonian gives the total energy of the system. Applying the Hamiltonian on the spin eigenstates gives

$$H\,|\uparrow\rangle = -\frac{\hbar}{2}\gamma B_0\,|\uparrow\rangle\,, \qquad\qquad H\,|\downarrow\rangle = \frac{\hbar}{2}\gamma B_0\,|\downarrow\rangle\,. \tag{2.1.7}$$

The spin up state has lower energy than the spin down state and the energy difference between the eigenstates, $\Delta E = \gamma\hbar B_0$, linearly depends on the external magnetic field magnitude.

By integrating the time-dependent Schrödinger equation over time, we get the following expression for a general spin state over time:

$$|S(t)\rangle = e^{i\gamma B_0 t S_z/\hbar}\,|S(0)\rangle\,. \tag{2.1.8}$$

If we consider a spin state defined by Equation 2.1.5 at time zero, we get

$$\begin{aligned} |S(t)\rangle &= e^{i\gamma B_0 t S_z/\hbar}\cos\left(\frac{\theta}{2}\right)|\uparrow\rangle + e^{i\gamma B_0 t S_z/\hbar + i\phi}\sin\left(\frac{\theta}{2}\right)|\downarrow\rangle \\ &= e^{i\gamma B_0 t/2}\left[\cos\left(\frac{\theta}{2}\right)|\uparrow\rangle + \sin\left(\frac{\theta}{2}\right)e^{i(\phi-\gamma B_0 t)}\,|\downarrow\rangle\right], \end{aligned} \tag{2.1.9}$$

where θ and ϕ define the initial state. Equation 2.1.9 shows that any spin that is not in either one of the eigenstates precesses about the z-axis at an angular

frequency $\omega = \gamma B_0$. This frequency that depends on the field strength and the gyromagnetic ratio of the nucleus is known as the Larmor frequency. In fact, this behaviour of the spin state is the same as the precession of a magnetic dipole in an external magnetic field, as described by classical electromagnetism.

If we consider an ensemble of protons in a thermal equilibrium, the probability for an indvidual proton to occupy the spin up state is slightly higher than the spin down state. The probabilities of observing the spins in one of the two eigenstates follow the Boltzmann distribution:

$$\frac{P_\uparrow}{P_\downarrow} = e^{\frac{\gamma\hbar B_0}{k_B T}}, \tag{2.1.10}$$

where $P_\uparrow$ is the probability to observe the protons in the spin up state, $P_\downarrow$ is the probability to observe the proton in the spin down state, $k_B \approx 1.38{\cdot}10^{-23}\,\mathrm{J{\cdot}K^{-1}}$ is the Boltzmann constant, and T is the temperature. The asymmetry between the states is small at magnetic field strengths generated by most MRI scanners. For instance, in a magnetic field of 3 T and at body temperature, $P_\uparrow \approx 0.5 + 10^{-5}$ and $P_\downarrow \approx 0.5 - 10^{-5}$. The asymmetry between the states creates a net magnetization

$$\mathbf{M} = \gamma \langle \mathbf{S} \rangle, \tag{2.1.11}$$

where $\langle \mathbf{S} \rangle$ is the ensemble average spin vector whose magnitude is equal to the ensemble average spin and whose direction is equal to the average direction of the spins as defined by Equation 2.1.5. In a thermal equilibrium, $\mathbf{M}$ is aligned with the external magnetic field.

2.2 Excitation and decay

The previous section provided a quantum mechanical overview of how net magnetization arises from an ensemble of spins in an external magnetic field. This section provides a classical description of how net magnetization enables NMR.

Let us continue to consider an ensemble of protons in an external magnetic field, $\mathbf{B}_0$, aligned with the z-axis. The field magnitude determines the Larmor frequency, which is typically in the radio frequency (RF) range. By applying an additional time-dependent magnetic field, $\mathbf{B}_{\mathrm{RF}}$, oscillating at the Larmor frequency in the xy-plane, it is possible to perturb the spin system so that the net magnetization vector $\mathbf{M}$ gets tipped towards the xy-plane by an angle

$$\theta_{\mathrm{RF}} = \gamma \mathrm{B}_{\mathrm{RF}} \tau_{\mathrm{RF}}, \tag{2.2.1}$$

where B_{RF} and τ_{RF} are the magnitude and duration of the RF pulse, respectively. A pulse with a duration $\pi/(2\gamma B_{RF})$ is known as a 90° pulse as it tips the magnetization vector into the xy-plane. After the RF pulse, $\mathbf{M}$ begins to precess about the z-axis, creating a time-dependent magnetic field that induces a current in the receiver coil placed around the spin system in the xy-plane. This is the NMR signal, which decays over time as the spin system returns to its equilibrium state.

Charged particles in a given proton's proximity are in constant motion due to thermal energy, creating a complex magnetic field, referred to as the lattice field, which is in constant interaction with the proton. If the interaction happens close to the Larmor frequency, the proton in a spin down state can move to a spin up state, emitting the energy difference between the states to its environment. This phenomenon is known as spin-lattice relaxation, as a consequence of which the longitudinal component of the net magnetization is recovered after an RF pulse. The expected lifetime of a proton in the higher energy state is characterized by a time constant T_1 which depends on the properties of the environment.

The decay of the transverse component of $\mathbf{M}$ happens as the excited spins' accumulated phases deviate from their expected value as a consequence of local fluctuations in the magnetic field and interactions between spins. This process, often referred to as spin-spin relaxation, occurs even in ideal systems, where the local environment of all spins is the same, and is characterized by a time constant T_2. In realistic systems, additional dephasing also occurs because of inhomogeneities in the magnetic field $\mathbf{B}_0$, which leads to dispersion in resonance frequencies. These inhomogeneities can be caused by the magnet hardware or by susceptibility-induced distortions produced by other matter, such as biomolecules, within the field. The combined signal loss from both spin-spin relaxation and magnetic field inhomogeneities is referred to as T_2-star (T_2^*).

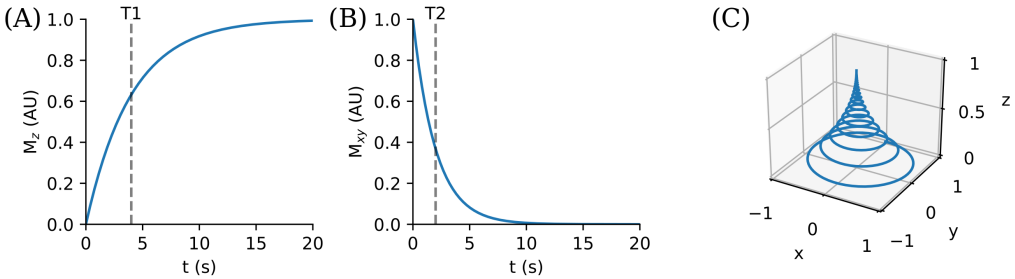

Figure 2.2: *Time evolution of the net magnetization* $\mathbf{M}$ *of an ideal spin system after a 90° RF pulse, as described by the Bloch equation. The recovery of the longitudinal magnetization (A) and the decay of the transverse magnetization (B) are characterized by the time constants* T_1 *and* T_2 *depicted by the dashed lines. In (C) the evolution of* $\mathbf{M}$ *over time is shown in 3D space. The results are calculated using* T_1 *= 4 s and* T_2 *= 2 s, values which are close to experimentally observed values in cerebrospinal fluid at 1.5 T*[73,74].

The time evolution of $\mathbf{M}$ after an RF pulse is summarised by the Bloch equation[75]:

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} - \frac{M_x\mathbf{i} - M_y\mathbf{j}}{T_2} - \frac{M_z - M_0}{T_1}\mathbf{k}, \tag{2.2.2}$$

where M_x, M_y, and M_z are the Cartesian components of $\mathbf{M}$, $\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ are unit vectors defining the axes of the coordinate system in the laboratory frame, M_0 is the initial magnitude of $\mathbf{M}$, T_1 is the decay constant for the recovery of the longitudinal component of $\mathbf{M}$, and T_2 is the decay constant for the the transverse component of $\mathbf{M}$.

The signal measured immediately after the excitation pulse is called free induction decay (FID), which oscillates at the Larmor frequency and follows the T_2^* decay. In an inhomogeneous sample, the Larmor frequencies of the spins depend on their local environment and the FID is a superposition of the Larmor frequencies of all the spins in the sample. By applying another RF pulse before the signal readout, it is possible to reverse much of the T_2^*-related dephasing and regenerate the signal using a spin echo pulse sequence[12]. The refocusing pulse effectively reverses the dephasing direction and the phases of the faster precessing spins will catch up to the slower precessing spins at echo time (TE), generating a measurable signal. An example of a spin echo sequence is depicted in Figure 2.3, where a 180° RF pulse is applied on the spin system at time TE/2. The magnitude of the spin echo is limited by the T_2 decay

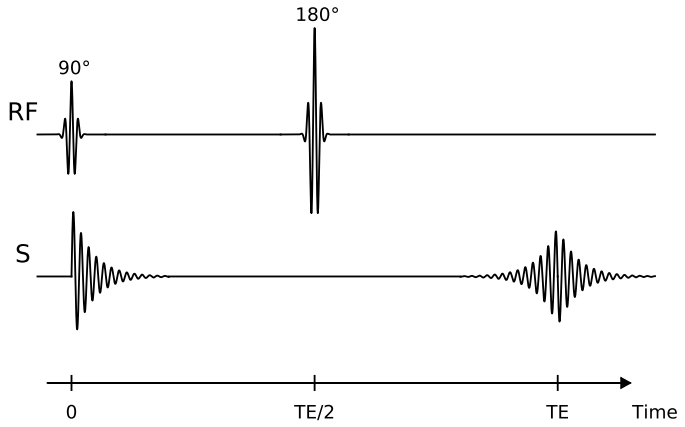

Figure 2.3: *A schematic representation of the spin echo sequence. The excitation pulse (90°) generates a free induction decay in signal. The addition of the refocusing pulse (180°) at TE/2 results in a signal echo at TE.*

because the random fluctuations resulting in T_2-related signal loss can not be reversed[71].

2.3 Image formation

In MRI, a spatial map of net magnetization is reconstructed from the NMR signal by manipulating the phases of spins with additional magnetic fields that vary linearly over space, simply referred to as gradients. Because the Larmor frequency depends linearly on the magnitude of the external magnetic field, a specific slice can be excited by combining the excitation pulse with a magnetic field gradient whose direction is perpendicular to the slice. Let us denote the direction perpendicular to the excited slice with z and the directions defining the plane of the excited slice with x and y. If additional magnetic field gradients are applied along x and y, the precession frequency of an excited spin depends on its location in the excited plane. Thus, the location of the spin is encoded in the accumulated phase shift during the application of the gradients.

Spatial frequencies in k-space can be defined as[20]

$$\begin{aligned} k_x(t) &= \gamma \int_0^t G_x(x,t')dt', \\ k_y(t) &= \gamma \int_0^t G_y(y,t')dt', \end{aligned} \tag{2.3.1}$$

where G_x and G_y are the gradient field magnitudes along x and y. Now the

measured signal S in the receiver coil can be written as

$$S(t) = \int\int M(x,y)e^{i2\pi(k_x x + k_y y)}dxdy, \tag{2.3.2}$$

where $M(x,y)$ is the net magnetization in the excited plane and integration is performed over the excited plane. $M(x,y)$ can be recovered by sampling k-space and calculating the Fourier transform of the resulting signal. The resolution and the field of view (FOV) of the image is determined by the extent of k-space which is sampled and the sampling density, respectively.

Since the rate at which k-space can be sampled is limited by the scanner hardware, several sampling methods have been developed. In echo planar imaging (EPI), an entire 2D planar image is acquired after a small number of excitations, making it particularly relevant for dMRI because it decreases motion artefacts and enables very fast acquisitions[20].

The measured net magnetization, and therefore the image intensity, depends on the used TE and repetition time (TR) between the excitation pulses as

$$M(x,y) = \rho(x,y)\alpha\left(1 - e^{-TR/T_1}\right)e^{-TE/T_2}, \tag{2.3.3}$$

where $\rho(x,y)$ roughly represents proton density and α is a constant that accounts for spin dephasing caused by other factors than T_1 and T_2^* relaxation, such as diffusion, perfusion, and bulk flow. Equation 2.3.3 shows that image contrast can be altered by changing TR and TE, because T_1 and T_2 vary over tissue types. For example, at 3 T, T_1 in gray matter and white matter are approximately 1.3 and 0.8 s, respectively, and T_2 in gray matter and white matter are approximately 0.08 and 0.11, respectively[76].

2.4 Summary

This chapter provided an overview of MRI, describing how the NMR signal arises and how an image can be formed by sampling k-space. Many important details were omitted intentionally because they are outside the scope of this thesis. The next chapter focuses on adding diffusion-weighting gradients to the pulse sequence and the data analysis of the resulting MR images.

Chapter 3

Microstructural imaging using diffusion-weighted MRI

This chapter provides an overview of dMRI and how it enables the properties of neural tissue microstructure to be studied. The chapter begins by describing the microstructure of neural tissue and the basic physics of diffusion, followed by a description of dMRI in the context of microstructural imaging.

3.1 Overview of neural tissue

The human brain can be loosely divided into four parts: the brainstem, diencephalon, cerebellum, and cerebrum[77]. The brainstem connects the rest of the brain to the spinal cord and is associated with many functions critical to life, such as respiratory control and cardiovascular system control. The brainstem connects to the diencephalon which consists of the thalamus, hypothalamus, epithalamus, and subthalamus. The cerebellum, located inferior to the occipital lobe of the cerebrum, plays an important role in motor control and some cognitive processes[78]. The cerebrum is the largest part of the brain and can be further divided into four distinct lobes: the frontal lobe, parietal lobe, temporal lobe, and occipital lobe. The two cerebral hemispheres are covered by the cerebral cortex which consists of folded gray matter containing a relatively large amount of cell bodies of neurons. White matter, located underneath the cortex, contains mostly densely packed axons connecting neurons in cortical and subcortical gray matter[79]. Subcortical structures are islets of gray matter deep within the brain which include the diencephalon, pituitary gland, limbic structures, and the basal ganglia.

On a microscopic level, brain tissue consists mostly of neurons, glial cells, and blood vessels[80]. In the highly interconnected and complex network of neurons information is transmitted between neurons via synapses which connect the axons of pre-synaptic neurons to the dendrites or somas of post-synaptic neurons. Axons are projections of neurons with lengths ranging from less than a millimetre in intra-cortical neurons to over one metre in the sciatic nerve[81]. Dendrites are branched extensions of neural somas, i.e., the cell bodies. Glial cells, which support neuronal activity, do not produce electrical impulses like neurons. Glial cells can be classified into various categories of which the most common are oligodendrocytes, astrocytes, and microglia[80].

The cells in the brain are surrounded by extra-cellular space which consists of a solution that closely resembles cerebrospinal fluid with additional extracellular matrix molecules[82]. The intra-cellular space contains the cell nucleus and cytoplasm which mostly consists of cytosol, a majority of which is water, and cell organelles[83]. The cell membrane restricts the diffusion of water but it is not completely impermeable to water. Instead, water slowly diffuses through the lipid membrane and is actively transported across the membrane by water-specific membrane channel proteins known as aquaporins[84].

The cortical gray matter of an adult human brain contains approximately 12 billion neurons and even more non-neuronal cells which are mostly glial cells[1]. Pyramidal cells and interneurons make up approximately 70% and 30% of the neurons of the cortex, respectively[85]. Oligodendrocytes, astrocytes, and microglia make up approximately 75%, 19%, and 6% of the glial cells of the cortex, respectively[86]. In the cortex, the densely connected dendritic trees extending from the neural somas are mostly isotropic, branching and extending equally in all directions, whereas in subcortical gray matter the dendritic trees can be highly anisotropic to increase inter-layer connectivity[87,88]. The diameters of the dendrites typically range from 0.1 to 0.3 μm and they extend up to a few millimetres away from the soma[87]. The sizes of the neural somas typically range from approximately 5 to 50 μm across, depending on the type and function of the neuron. The somas of astrocytes and microglia are approximately 10 μm in diameter with the somas oligodendrocyte being slightly larger[89,90].

In human white matter, up to 90% of the space is occupied by axons[79] which serve as conductors of action potentials. The diameters of axons typically vary from 0.1 to 10 μm with larger axons providing the benefit of transmitting action potentials more quickly[91], although at the cost of consuming more energy and

occupying more space[92]. Oligodendrocytes extend their processes to several axons and wrap around them, creating an insulating layer of myelin, known as an internode. The myelin sheaths, made up of roughly 80% lipids and 20% proteins, insulate the axon, form a nearly impermeable barrier to water and increase the conduction speed of action potentials[91,93]. With the exception of small axons with a diameter of less than 1 µm, most axons are myelinated[91], and demyelination may occur as a consequence of various neurological conditions, such as multiple sclerosis[94].

The brain's arteries, veins, and capillaries make up the vascular system which provides the cells with oxygen, glucose, and other necessary nutrients. The diameters of of the vascular system's components range from micrometres to millimetres. The capillaries are more prevalent in gray matter than white matter and lack orientational coherence[95], which may lead the circulating blood to exhibit displacements similar to diffusion, although more rapid[96].

3.2 Physics of diffusion

Diffusion is the stochastic translational motion of particles in a fluid driven by the internal kinetic energy of the system. The unpredictable motion of microscopic particles was first reported by Brown[97] who observed that pollen particles suspended in water move in a seemingly random manner, a phenomenon nowadays known as Brownian motion.

Bulk diffusion, the process of net displacement of particles in a fluid from a region of high concentration towards a region of lower concentration, was extensively studied in the 19th century and summarized in Fick's laws of diffusion:

$$\mathbf{J} = -D\nabla C, \tag{3.2.1}$$

$$\frac{\partial C}{\partial t} = D\nabla^2 C, \tag{3.2.2}$$

where $\mathbf{J}$ is the diffusion flux, D is the diffusion coefficient expressing the rate of diffusion in the medium, and C is the local concentration of particles[98]. Here, D is assumed to be constant across the volume. Fick's laws of diffusion enable the estimation of the diffusion coefficient from experimental data and describe how concentration C across space evolves over time. Fick's macroscopic description of the system states that there is no net displacement of particles if their concentration is constant across the volume.

On a microscopic level, particles in a fluid are in constant thermal motion, continuously mixing and colliding into each other, which leads to self-diffusion that occurs even in the absence of a concentration gradient. Einstein was the first to connect the diffusion coefficient to microscopic Brownian motion[99] by showing that the probability of a freely diffusing particle to undergo a displacement $\mathbf{r}$ during time t is given by

$$P(\mathbf{r}, t) = \frac{1}{(4\pi Dt)^{3/2}} \exp\left(-\frac{|\mathbf{r}|^2}{4Dt}\right), \tag{3.2.3}$$

which is a symmetric three-dimensional normal, i.e., Gaussian, distribution with the following mean and covariance matrix:

$$\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} 2Dt & 0 & 0 \\ 0 & 2Dt & 0 \\ 0 & 0 & 2Dt \end{bmatrix}. \tag{3.2.4}$$

The displacement probability distribution is also referred to as the average diffusion propagator. The Gaussian form of the displacement distribution of particles undergoing free diffusion can be understood by considering the displacement of a particle as a sum of a very large number of elementary steps arising from collisions with other particles[19]. The central limit theorem implies that the displacement distribution follows a normal distribution because the elementary steps become practically uncorrelated after picoseconds[100,101].

The diffusion coefficient is related to the intrinsic properties of the medium by the Sutherland-Einstein equation[99]:

$$D = \frac{k_B T}{6\pi\eta\alpha}, \tag{3.2.5}$$

where k_B is the Boltzmann constant, T is the temperature, η is the viscosity of the medium, and α is the radius of the diffusing particles which are assumed to be spheres in the derivation of the equation.

It follows from Equation 3.2.3 that the mean squared displacement (MSD) of an ensemble of N particles freely diffusing in 3D depends linearly on the diffusion coefficient and diffusion time:

$$\text{MSD} = \left\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \right\rangle = 6Dt, \tag{3.2.6}$$

which is valid for $N \to \infty$.

In general, diffusion is not isotropic. For example, water molecules inside an axon can move much more freely along the axon than along a direction perpendicular to the axon. The displacement distribution of particles undergoing

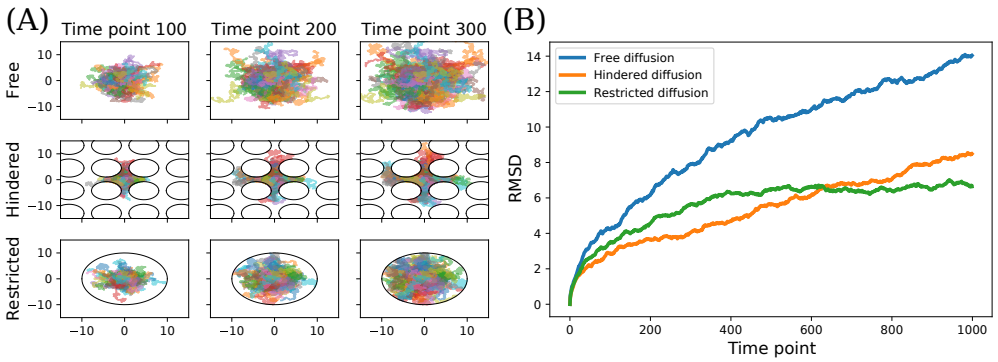

Figure 3.1: *(A) Simulated trajectories of 100 two-dimensional random walkers undergoing free, hindered, and restricted diffusion. The columns correspond to snapshots of the simulation at different time points and the rows correspond to different diffusion environments. (B) The root mean squared displacement (RMSD) of the walkers over time. The initial position of each walker is at the origin and the step length is 0.5. In typical neuroimaging experiments, RMSD is roughly around 10 μm.*

anisotropic Gaussian diffusion has zero mean and the covariance matrix

$$\boldsymbol{\Sigma} = 2t\mathbf{D} = 2t \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix}, \tag{3.2.7}$$

where $\mathbf{D}$ is the diffusion tensor which is symmetric positive definite. The diffusion tensor's eigenvectors $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3\}$ are orthogonal and its eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$ are equal to the diffusion coefficients in the direction of the corresponding eigenvector. Since the diffusion tensor can only represent Gaussian diffusion, it is not an accurate representation of diffusion in neural tissue where the average diffusion propagator at the voxel level is not Gaussian[102].

Diffusion can be characterized as free, hindered, or restricted depending on the environment in which the particles are (Figure 3.1). In free diffusion, a diffusing particle's movement is only affected by other particles and its expected squared displacement follows the description by Equation 3.2.6. In restricted diffusion, the particles are confined to a volume and their displacements are constrained by the size of the restricting volume. In hindered diffusion, particles collide with obstacles but they aren't confined to a volume, so their mean squared displacement increases with time without an upper bound, although at a slower rate than in free diffusion.

In the special case of free diffusion in homogeneous media, the diffusion

coefficient D is independent of time: $D(t) = D$ for all t. However, this does not generalize to most situations in which the diffusion coefficient is time dependent. In the short diffusion time limit ($t \to 0$), a diffusing particle is only interacting with its immediate surroundings and diffusion is characterized by the intrinsic diffusion coefficient of the medium D_0. In Figure 3.1B, the root mean squared displacement initially grows at the same rate for all diffusion types, determined by D_0. As the diffusion time increases, the particles explore more of their surroundings and their movement is affected by the obstacles in their environment. For diffusion times that result in particle displacements that are greater than the length scale of the obstacles, hindered diffusion resembles free diffusion with a diminished value of D, referred to as the effective diffusion coefficient: $D_{\text{effective}} < D_0$. In restricted diffusion, the root mean squared displacement is limited by the size of the restricting volume and thus an increased diffusion time results in a lower value of effective diffusion coefficient if the diffusion coefficient is calculated from the displacements during a given experimental diffusion time.

In some situations, it is helpful to describe diffusion in terms of velocities of particles instead of their displacements. This formalism enables the dMRI signal to be analyzed in the diffusion frequency domain[103]. By expressing the displacement of a particle diffusing in 3D as a time integral of its velocity $\mathbf{v}(t)$, a relationship between the diffusion coefficient and particle velocities can be derived from Equation 3.2.6[104]:

$$\begin{aligned} D &= \lim_{t\to\infty} \frac{1}{6}\frac{\partial}{\partial t}\left\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \right\rangle \\ &= \lim_{t\to\infty} \frac{1}{6}\frac{\partial}{\partial t}\left\langle \left(\int_0^t \mathbf{v}(t')dt'\right) \cdot \left(\int_0^t \mathbf{v}(t')dt'\right) \right\rangle \\ &= \lim_{t\to\infty} \frac{1}{3}\int_0^t \left\langle \mathbf{v}(t') \cdot \mathbf{v}(t) \right\rangle dt' \\ &= \lim_{t\to\infty} \frac{1}{3}\int_0^t \left\langle \mathbf{v}(0) \cdot \mathbf{v}(t-t') \right\rangle dt' \\ &= \frac{1}{3}\int_0^\infty \left\langle \mathbf{v}(0) \cdot \mathbf{v}(t) \right\rangle dt \end{aligned} \quad , \tag{3.2.8}$$

where $\langle \mathbf{v}(t') \cdot \mathbf{v}(t) \rangle$ is the velocity autocorrelation function describing how correlated the particle velocities are at time points t' and t, and which is invariant under time translation[105]. From the velocity autocorrelation function, it is possible to calculate correlation time t_c, the time period over which the

molecular velocities becomes decorrelated[104], as

$$t_c = \int_0^\infty \frac{\langle \mathbf{v}(0) \cdot \mathbf{v}(t) \rangle}{\langle v^2 \rangle} dt, \tag{3.2.9}$$

where the mean squared velocity $\langle v^2 \rangle$ is taken to be constant. Velocities of freely diffusing particles are autocorrelated only over very short time intervals, but collisions with boundaries induce negative autocorrelations over longer time scales in hindered and restricted diffusion. The diffusion spectrum is the Fourier transform of the velocity autocorrelation function[103]:

$$D(\omega) = \frac{1}{2} \int_{-\infty}^{\infty} \langle \mathbf{v}(0) \cdot \mathbf{v}(t) \rangle \, e^{i\omega t} dt, \tag{3.2.10}$$

where ω is the frequency. The diffusion spectrum of free diffusion has equal velocity contributions from all frequencies. However, in hindered and restricted diffusion, collisions with the boundaries result the diffusion spectrum to deviate from a uniform distribution.

3.3 Diffusion-weighted signal

The effects of diffusion on the NMR signal in a spin echo experiment were first observed and reported in the 1950s by Hahn[12] and discussed in more detail by Carr and Purcell[13]. In the 1960s, Stejskal and Tanner introduced the pulsed gradient spin echo (PGSE) sequence (Figure 3.2) in which the diffusion encoding time is separated from the experimental diffusion time, facilitating data analysis. Although other diffusion encoding methods, such as oscillating gradients[106], have also been developed, PGSE EPI remains as the standard pulse sequence in modern dMRI.

In diffusion-weighted NMR and MRI experiments, the spins are exposed to additional time-dependent magnetic fields that vary linearly over space, simply referred to as gradients. The phase shift experienced by a spin is given by

$$\phi(t) = \gamma \int_0^t \mathbf{B}_0 + \mathbf{G}(t') \cdot \mathbf{r}(t') dt', \tag{3.3.1}$$

where γ is the gyromagnetic ratio of the nucleus, $\mathbf{B}_0$ is the static main magnetic field, $\mathbf{G}$ is the gradient, and $\mathbf{r}$ is the spin's position. $\mathbf{B}_0$ and $\mathbf{G}$ are effective fields which change sign after the application of the 180° RF pulse. Since $\mathbf{B}_0$ is uniform over space, it has the same effect on all spins and therefore does not cause phase dispersion. The application of $\mathbf{G}$, on the other hand, results in a location-dependent phase shift for each spin.

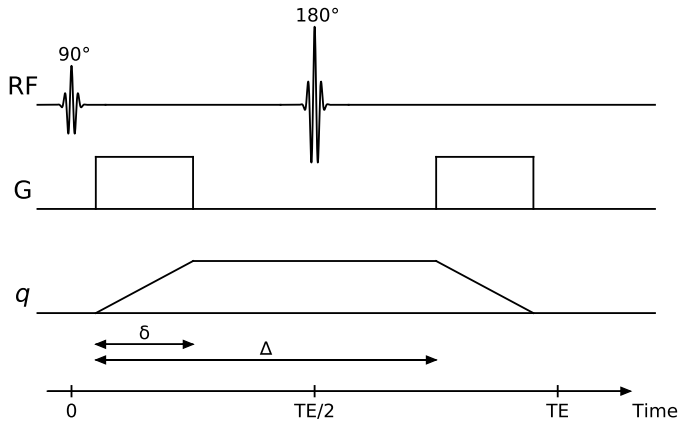

Figure 3.2: *A schematic representation of a pulsed gradient spin echo sequence. Spins are excited with a 90° RF pulse and refocused with a 180° RF pulse which occurs at TE/2. Two gradient pulses are applied on both sides of the refocusing pulse. The duration of the gradient pulse is denoted by δ and the difference between the onsets of the gradient pulses is denoted by Δ. The wave vector magnitude q is proportional to the integral of gradient magnitude over time.*

Gradient waveforms, i.e., the diffusion encoding gradient over time, are defined so that the total phase shift by TE is zero for stationary spins, e.g., the second gradient pulse in a PGSE experiment reverses the phase shift caused by the first one. Diffusing spins experience non-zero phase shifts as a consequence of their random displacements during the experiment, resulting in phase dispersion and a diffusion-attenuated signal

$$S = S_0 \int_{-\infty}^{\infty} P(\phi) \exp(i\phi)\, d\phi, \tag{3.3.2}$$

where S_0 is the measured signal without diffusion-weighing but with other acquisition parameters being equal and P is the phase distribution of the spins in the ensemble. Generally, the real component of the signal is the quantity of interest. To factor out other signal loss mechanisms, it is useful to define the normalized diffusion-weighted signal as

$$E = \frac{S}{S_0}. \tag{3.3.3}$$

It is also helpful to define the wave vector $\mathbf{q}$ as

$$\mathbf{q}(t) = \gamma \int_0^t \mathbf{G}(t')dt'. \tag{3.3.4}$$

By applying integration by parts on Equation 3.3.1, the acquired phase can be related to the spin's velocity as

$$\phi(t) = \gamma \int_0^t \mathbf{B}_0 - \mathbf{q}(t') \cdot \mathbf{v}(t') dt'. \tag{3.3.5}$$

The magnitude of diffusion-weighting depends on both the diffusion time and the magnitude of the applied gradient. The contributions from these two factors are summarized by the b-value[107] which is equal to the trace of the b-tensor $\mathbf{b}$, which is also often referred to as the b-matrix or the measurement tensor. It is a symmetric rank-2 tensor that relates to the wave vector as follows:

$$\mathbf{b} = \int_0^{TE} \mathbf{q}(t)\mathbf{q}^{\mathrm{T}}(t) dt. \tag{3.3.6}$$

For accurate estimation of diffusion from signal, the b-tensor should contain the effects of all imaging and diffusion-weighting gradients.

The theoretical analysis of a PGSE experiment can be greatly simplified by assuming the duration of the gradient pulses to be very short ($\delta \to 0$), meaning that diffusion of spins during the gradient pulses does not need to be considered. In this case, the spins acquire phase at two points of their trajectories and the wave vector can be represented by a point in q-space: $\mathbf{q} = \gamma\delta\mathbf{G}$. The normalized signal at echo time is given by

$$E = \int_{\mathbb{R}^3} \rho(\mathbf{r}_1) \int_{\mathbb{R}^3} P(\mathbf{r}_1, \mathbf{r}_2, \Delta) \exp\left[-i\mathbf{q} \cdot (\mathbf{r}_2 - \mathbf{r}_1)\right] d\mathbf{r}_2 d\mathbf{r}_1, \tag{3.3.7}$$

where $\rho(\mathbf{r}_1)$ is the spin density at $\mathbf{r}_1$ during the first gradient pulse, and $P(\mathbf{r}_1, \mathbf{r}_1, \Delta)$ is the diffusion propagator, i.e., the probability of the spin to displace from $\mathbf{r}_1$ to $\mathbf{r}_2$ during Δ. In practice, it is convenient to introduce a displacement variable $\mathbf{r} = \mathbf{r}_2 - \mathbf{r}_1$ and integrate over all initial positions to get

$$E = \int_{\mathbb{R}^3} P(\mathbf{r}, \Delta) \exp\left(-i\mathbf{q} \cdot \mathbf{r}\right) d\mathbf{r}, \tag{3.3.8}$$

where $P(\mathbf{r}, \Delta)$ is the average propagator representing the probability for a spin to experience a displacement $\mathbf{r}$ during Δ.

Equation 3.3.8 shows that under the short pulse assumption, the normalized signal is the Fourier transform of the average propagator. In q-space imaging, the average propagator is estimated by applying a Fourier transform on data acquired by densely sampling q-space[108–110]. The well-known properties of the Fourier transform naturally apply, namely, the resolution of the estimated propagator is determined by the magnitude of the applied gradient and the propagator field of view is determined by the sampling density in q-space.

In real experiments it is impossible to fulfil the short pulse length condition, which leads to an underestimation of diffusion. Instead of measuring spin displacements during Δ, experiments with finite pulse lengths measure the differences in spins' locations between their average positions during the applications of the two gradient pulses[111]. This general result applies to all PGSE experiments.

Since the displacements of water molecules in the brain are constrained by their microscopic environments, the diffusion-weighted signal contains valuable information on tissue microstructure. The central problem in dMRI is to infer relevant properties of tissue microstructure from the diffusion-attenuated signal. However, the diffusion propagator is related to the tissue microstructure in a complex and time-dependent way. Furthermore, the relationship between the diffusion propagator and the measured signal is highly dependent on the chosen acquisition parameters, and the signal often changes very little as the microstructure changes, making sensitivity an issue. Given these limitations, characterizing neural tissue microstructure with dMRI is a very challenging inverse problem.

3.4 Diffusion tensor imaging

DTI quantifies the diffusion of water in a voxel using the diffusion tensor (Equation 3.2.7). Since its introduction in the 1990s[25], it has become the standard dMRI method in neuroscience[112]. The voxel-level diffusion tensor can be connected to the dMRI signal in a PGSE experiment by considering the diffusion of net magnetization[14,113], as described by the Bloch-Torrey equation[114], an extension of the Bloch equation[75] (Equation 2.2.2):

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} - \frac{M_x \mathbf{i} - M_y \mathbf{j}}{T_2} - \frac{M_z - M_0}{T_1}\mathbf{k} + \nabla \cdot (\mathbf{D} \nabla \mathbf{M}) \,, \tag{3.4.1}$$

where $\mathbf{B}(\mathbf{r}, t) = \mathbf{B}_0 + \mathbf{G}(t) \cdot \mathbf{r}$, the frame of reference is defined so that $\mathbf{B}_0$ is aligned with the z-axis, M_x, M_y, and M_z are the Cartesian components of net magnetization $\mathbf{M}$, $\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ are unit vectors defining the axes of the coordinate system, M_0 is the initial magnitude of $\mathbf{M}$, T_1 and T_2 are the magnetization decay constants, and $\mathbf{D}$ is the diffusion tensor. Here, $\mathbf{r}$ represents a point in space. As shown by Stejskal and Tanner[14], it is helpful to define the complex transverse magnetization as

$$M_x(\mathbf{r}, t) + iM_y(\mathbf{r}, t) = \psi(\mathbf{r}, t) \exp\left[-(i\gamma B_0 + 1/T_2)t\right] , \tag{3.4.2}$$

where ψ is the transverse magnetization magnitude independent of T_2 decay. From Equations 3.4.1 and 3.4.2 it follows that the time evolution of ψ is given by

$$\frac{\partial \psi(\mathbf{r},t)}{\partial t} = -i\gamma \left[\mathbf{r}\cdot\mathbf{G}(t)\right]\psi(\mathbf{r},t) + \nabla\cdot\left[\mathbf{D}\nabla\psi(\mathbf{r},t)\right], \tag{3.4.3}$$

which can be further simplified by dividing the equation into real and imaginary parts. The solution to the imaginary part represents the system without diffusion, and the real part represents the solution to the system with diffusion-related attenuation in transverse magnetization[14]. The real part can be solved by starting from

$$\psi(\mathbf{r},t) = M(t)\exp\left[-i\mathbf{r}\cdot\mathbf{q}(t)\right]. \tag{3.4.4}$$

From Equation 3.4.3 we get

$$\frac{\partial M(t)}{\partial t} = -M(t)\mathbf{q}^{\mathrm{T}}(t)\mathbf{D}\mathbf{q}(t), \tag{3.4.5}$$

which has the following solution for Gaussian diffusion in a homogeneous infinite medium[113]

$$M(t) = M(0)\exp\left(-\int_0^t \mathbf{q}^{\mathrm{T}}(t')\mathbf{D}\mathbf{q}(t')dt'\right), \tag{3.4.6}$$

where $M(0)$ is the transverse magnetization magnitude immediately after the 90° excitation pulse. Now the normalized diffusion-weighted signal can be expressed as

$$E = \exp\left(-\int_0^{TE} \mathbf{q}^{\mathrm{T}}(t')\mathbf{D}\mathbf{q}(t')dt'\right). \tag{3.4.7}$$

In the case of Gaussian diffusion, the diffusion-weighted signal attenuation depends only on the b-value of the applied gradient and not on the exact trajectory in q-space[115]. Therefore, the normalized signal can be expressed as

$$E = \exp\left(-\mathbf{b}:\mathbf{D}\right) = \exp\left(-\sum_i^3\sum_j^3 b_{ij}D_{ij}\right). \tag{3.4.8}$$

By representing the symmetric tensors $\mathbf{b}$ and $\mathbf{D}$ as one-dimensional vectors in Voigt notation, the signals from all acquisitions can be written as a system of linear equations:

$$\mathbf{E} = \mathbf{X}\mathbf{D}, \tag{3.4.9}$$

where

$$\mathbf{E} = \begin{pmatrix} \ln E^1 \\ \vdots \\ \ln E^N \end{pmatrix}, \tag{3.4.10}$$

$$\mathbf{X} = \begin{pmatrix} b_{xx}^1 & \sqrt{2}b_{xy}^1 & \sqrt{2}b_{xz}^1 & b_{yy}^1 & \sqrt{2}b_{yz}^1 & b_{zz}^1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ b_{xx}^N & \sqrt{2}b_{xy}^N & \sqrt{2}b_{xz}^N & b_{yy}^N & \sqrt{2}b_{yz}^N & b_{zz}^N \end{pmatrix}, \tag{3.4.11}$$

$$\mathbf{D} = \begin{pmatrix} D_{xx} & \sqrt{2}D_{xy} & \sqrt{2}D_{xz} & D_{yy} & \sqrt{2}D_{yz} & D_{zz} \end{pmatrix}^T, \tag{3.4.12}$$

where N is the number of acquisitions and the superscript refers to the acquisition number. Given that a minimum of six diffusion-weighted acquisitions with different directions and one acquisition without diffusion-weighting were performed, $\mathbf{D}$ can be estimated by inverting $\mathbf{X}$. However, the application of the logarithmic transformation on the signal introduces heteroscedasticity into the data, for which weighted least squares should be applied in case of multi-shell data[116,117]. It is also possible to fit the signal model (Equation 3.4.8) directly to data using non-linear least squares methods.

After the diffusion tensor has been estimated, it is straightforward to quantify its magnitude and anisotropy using the metrics of mean diffusivity (MD), axial diffusivity (AD), radial diffusivity (RD), and fractional anisotropy (FA)[118]:

$$\text{MD} = \frac{\text{Tr}\left(\mathbf{D}\right)}{3} = \frac{1}{3}\left(\lambda_1 + \lambda_2 + \lambda_3\right), \tag{3.4.13}$$

$$\text{AD} = \lambda_1, \tag{3.4.14}$$

$$\text{RD} = \frac{1}{2}\left(\lambda_2 + \lambda_3\right), \tag{3.4.15}$$

$$\text{FA} = \sqrt{\frac{3}{2}\frac{\text{Var}_\lambda(\mathbf{D})}{\text{Var}_\lambda(\mathbf{D}) + \text{MD}^2}}, \tag{3.4.16}$$

where $\text{Tr}(\,)$ denotes calculating the trace of a matrix, the diffusion tensor eigenvalues are sorted ($\lambda_1 >= \lambda_2 >= \lambda_3$), and $\text{Var}_\lambda(\,)$ denotes an operator that calculates the eigenvalue variance:

$$\text{Var}_\lambda(\mathbf{D}) = \frac{1}{3}\sum_{i=1}^{3}\left(\lambda_i - \frac{\text{Tr}\left(\mathbf{D}\right)}{3}\right)^2. \tag{3.4.17}$$

Changes in neural tissue microstructure may result in changes in the diffusion propagator of water that can be detected with DTI. For example, demyelination of axons can lead to a reduction in FA and an increase in RD in voxels containing tightly packed and aligned axons. The DTI parameter maps, illustrated in Figure 3.3, have been shown to be sensitive to microstructural changes associated with, for instance, development[119], aging[120,121], and lesions associated with neurological conditions[122,123]. However, DTI confounds the orientation dispersion of anisotropic neurites with isotropic diffusion, which can

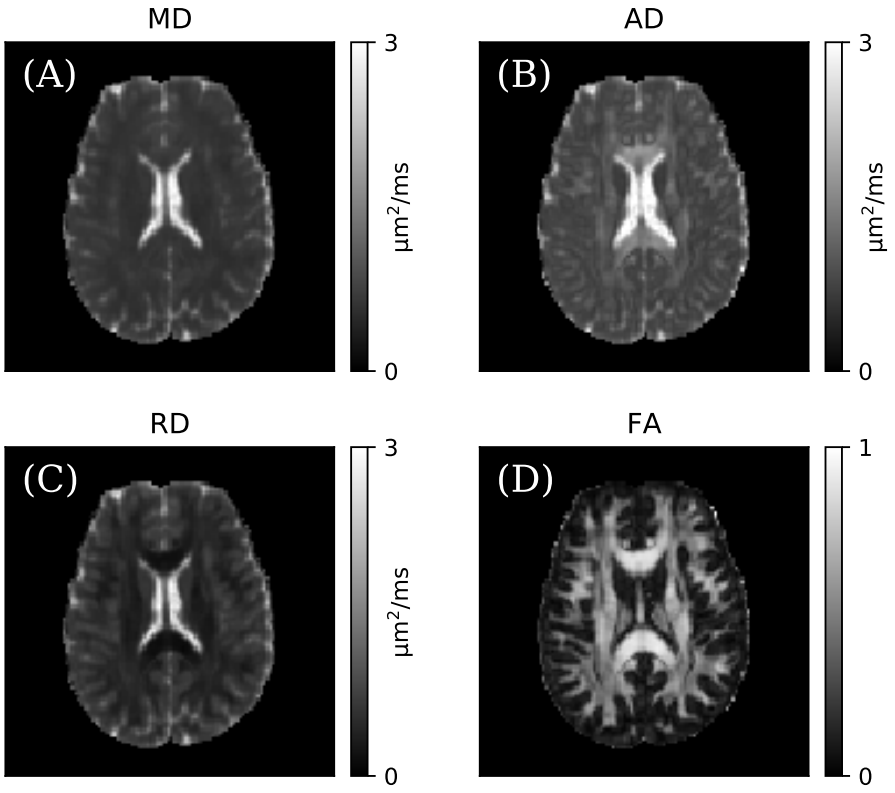

Figure 3.3: *Examples of DTI parameter maps in a healthy volunteer. (A) Mean diffusivity. (B) Axial diffusivity. (C) Radial diffusivity. (D) Fractional anisotropy.*

be seen in voxels containing crossing axonal bundles. Although diffusion inside individual axons is highly anisotropic, the voxel-level diffusion tensor is less anisotropic as a consequence of orientation dispersion. Thus, FA and RD are not similarly sensitive to axonal degeneration in such voxels. Furthermore, DTI is limited by the fact that it is not able to capture non-Gaussian diffusion and thus does not provide a good fit to data in experiments with moderate to high levels of diffusion-weighting (roughly $b > 1$ µm^2/ms in the brain) when the diffusion-attenuated signal decay clearly deviates from a monoexponential curve, revealing that the voxel-level average diffusion propagator is not Gaussian, especially in white matter. The following sections describe some methods that have been developed to address these limitations.

3.5 Signal cumulant expansion

The diffusion-weighted signal in Equation 3.3.2 is the average of an exponential function and can thus be expressed as a cumulant expansion[124]. The signal cumulant expansion is a powerful signal model in which the natural logarithm of the signal is expanded in powers of the applied gradient. The coefficients of the cumulant expansion, expressed in terms of correlations of molecular velocity, quantify properties of the diffusion propagator. This section briefly reviews the signal cumulant expansion and DKI as a natural extension of DTI for quantifying non-Gaussian diffusion at the voxel level.

Consider a real-valued random variable X with probability density f. Its nth raw moment is

$$\mu_n = \mathrm{E}\left(X^n\right) = \int_{-\infty}^{\infty} x^n f(x) dx, \tag{3.5.1}$$

where E denotes the expected value. Given that the moment-generating function $M_X(t)$ exists, it can be expanded about the origin:

$$M_X(t) = \mathrm{E}\left(e^{tX}\right) = \sum_{n=0}^{\infty} \frac{t^n}{n!} \mu_n. \tag{3.5.2}$$

The cumulant-generating function $K_X(t)$ is the natural logarithm of the moment-generating function. The cumulants κ_n are the coefficients of its Taylor expansion about the origin:

$$K_X(t) = \ln\left[\mathrm{E}\left(e^{tX}\right)\right] = \sum_{n=1}^{\infty} \frac{t^n}{n!} \kappa_n. \tag{3.5.3}$$

The first three cumulants are related to the raw moments as

$$\begin{aligned} \kappa_1 &= \mu_1, \\ \kappa_2 &= \mu_2 - \mu_1^2, \\ \kappa_3 &= \mu_3 - 3\mu_2\mu_1 + 2\mu_1^3, \end{aligned} \tag{3.5.4}$$

from which we can see that the first cumulant is equal to the mean and the second cumulant is equal to the variance of X.

By expressing the gradient in terms of wave vector magnitude, the natural logarithm of the normalized diffusion-weighted signal (Equations 3.3.2 and 3.3.3) can be expresed as

$$\ln(E) = \sum_{n=1}^{\infty} \frac{i^n}{n!} \int_0^{TE} \kappa_n(t_1, ..., t_n) q(t_1) ... q(t_n) dt_1 ... dt_n, \tag{3.5.5}$$

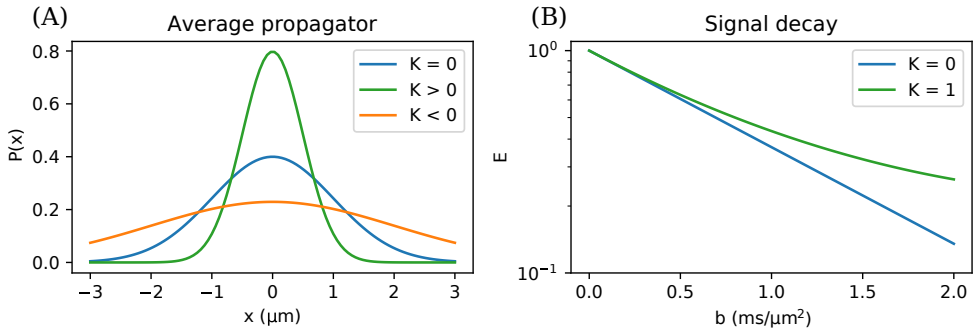

Figure 3.4: *(A) Examples of mesokurtic (K = 0), leptokurtic (K > 0), and platykurtic (K < 0) distributions. (B) An example of diffusion-weighted signal decay corresponding to a mesokurtic (K = 0) and leptokurtic (K = 1) diffusion propagators.*

where κ_n denotes the nth cumulant of molecular velocity and the integration is performed over all time variables. The derivation is omitted here but can be found in a book chapter by Kiselev[124]. Equation 3.5.5 is the general one-dimensional signal cumulant expansion which is strictly valid at $b \to 0$. In the absence of bulk flow, only the even terms contribute to diffusion-weighted signal attenuation because of the symmetry of the diffusion propagator. Since the first terms dominate at diffusion-weighting levels typically used in human neuroimaging and the validity of the truncated expansion diminishes with increasing gradient strength, the cumulant expansion is often truncated after the fourth term, corresponding to the second order in terms of b-value.

In 2005, Jensen et al. introduced DKI by showing that the truncated cumulant expansion in a PGSE experiment can be written as

$$\ln(E) \approx -bD + \frac{1}{6}b^2D^2K, \tag{3.5.6}$$

where b is the b-value, D is the apparent diffusion coefficient, and K is the apparent kurtosis coefficient. The kurtosis coefficient corresponds to the variance of apparent diffusivities and quantifies the magnitude of non-Gaussian diffusion at the voxel level[46]. Both D and K are called apparent coefficients because they depend on Δ and δ. In the context of dMRI, kurtosis generally refers to excess kurtosis which is negative for distributions that have longer tails than a Gaussian distribution (platykurtic) and positive for distributions that have shorter tails than a Gaussian distribution (leptokurtic) (Figure 3.4A). The voxel-level diffusion propagator is generally leptokurtic in neural tissue which consists of various microscopic diffusion environments, resulting in variance

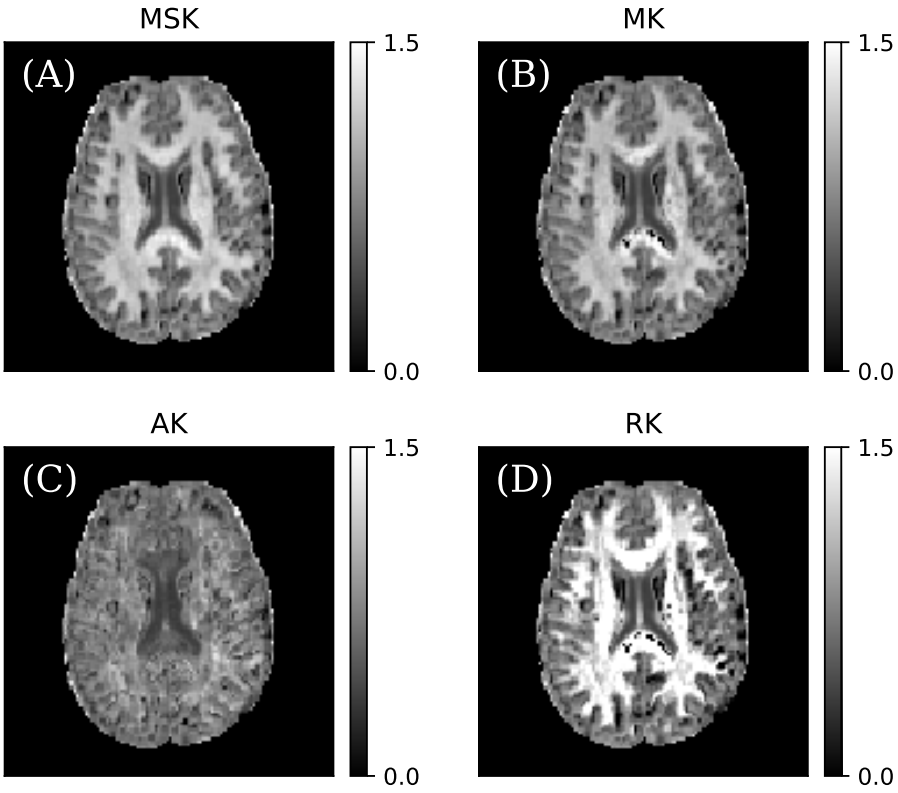

Figure 3.5: *Examples of DKI parameter maps in a healthy volunteer. (A) Mean signal kurtosis. (B) Mean kurtosis. (C) Axial kurtosis. (D) Radial kurtosis.*

of apparent diffusivities and a deviation from a monoexponential signal decay (Figure 3.4B).

In the brain, both D and K are orientationally variant, for which a tensor representation of kurtosis was introduced[46]:

$$\begin{aligned}\ln(E) \approx & -b\sum_{i=1}^{3}\sum_{j=1}^{3} n_i n_j D_{ij} \\ & + \frac{1}{6}b^2(\mathrm{MD})^2\sum_{i=1}^{3}\sum_{j=1}^{3}\sum_{k=1}^{3}\sum_{l=1}^{3} n_i n_j n_k n_l K_{ijkl},\end{aligned} \tag{3.5.7}$$

where n_i is the ith component of a unit vector representing the diffusion encoding direction, and D_{ij} and K_{ijkl} are the elements of the diffusion tensor $\mathbf{D}$ and kurtosis tensor $\mathbf{K}$, respectively. $\mathbf{K}$ is a symmetric rank-4 tensor with 15 unique elements. Estimating $\mathbf{D}$ and $\mathbf{K}$ requires a minimum of 22 acquisitions consisting of at least 15 unique gradient directions and three different levels of diffusion-weighting (including $b = 0$). The metrics of mean kurtosis (MK), axial kurtosis (AK), and radial kurtosis (RK), illustrated in Figure 3.5, quantify the

magnitude $\mathbf{K}$ with respect to the eigenvector corresponding to the largest diffusion tensor eigenvalue. Mean signal kurtosis (MSK) is equal to the apparent diffusion coefficient calculated by fitting Equation 3.5.6 to data that has been averaged over all the acquired diffusion encoding directions.

The measured diffusion kurtosis is caused by two main factors, the variance in microscopic diffusion environments and restricted diffusion[46,125]. For instance, consider a voxel containing tightly packed axons. Variance in the diameters or orientations of the axons results in a non-Gaussian diffusion propagator at the voxel-level even if diffusion inside each axon appears Gaussian because the sum of different normal distributions is not a normal distribution. This is referred to as inter-compartmental kurtosis. Intra-compartmental kurtosis, on the other hand, is caused by restricted diffusion within microscopic diffusion environments. Inter-compartmental kurtosis is positive but intra-compartmental kurtosis can be negative[46,125]. DKI can be sensitive to the different sources of kurtosis but its major limitation is the lack of specificity.

3.6 Microstructural modelling

Microstructural models aim to relate the measured diffusion-weighted signal to the biophysical properties of the tissue using descriptive metrics such as mean axon diameter and intra-neurite volume fraction. These models require the construction of a tissue model that is able to capture the relevant properties of tissue.

Most microstructural models represent the tissue as a set of non-exchanging microscopic diffusion environments, referred to as compartments. The normalized diffusion-weighted signal is expressed as the sum of signals from individual compartments:

$$E = \sum_{i=1}^{N} f_i E_i, \tag{3.6.1}$$

where N is the number of compartments in the voxel, f_i is the signal fraction of the ith compartment, and E_i is the normalized diffusion-weighted signal from the ith compartment. The signal fractions are normalized so that $\sum_{i=1}^{N} f_i = 1$ and $0 \leq f_i \leq 1$ for all i. The compartments are usually modelled as simple geometries that enable the signal to be expressed analytically with a small number of parameters to facilitate fitting the model. The simplest models are relatively straightforward to fit with standard non-linear fitting algorithms and

can be independent of the starting position[126]. The more complex models however, are more challenging to fit due to the flatness of the fitting landscape, local minima, and their sensitivity to changes in the starting position[127,126]. Exploring optimal ways of fitting the models continues to be an active area of research[117,128].

Compartment models with different fitting constraints have been succesfully applied in estimating mean axon diameter[129,51] and neurite density[49,130] in animal experiments with very high magnetic field and gradient strengths. However, their application in human neuroimaging has been less succesful. It has been shown that axon diameter estimates have been inaccurate by up to an order of magnitude[131] and that different fit constraints result in contradictory estimates[132,133]. Such issues have been connected to inadequate models and signal degeneracies[62]. Despite the accuracy issues, microstructural models may be sensitive to scientifically and clinically valuable changes in the brain microstructure. Some microstructural models are described in more detail in Chapter 5.

3.7 Multidimensional diffusion encoding

All the methods discussed in the previous sections of this chapter belong to a class of SDE methods as they are based on acquisitions measuring the displacements of water molecules along a single dimension[52]. Despite being sensitive to several scientifically and clinically relevant microstructural properties, SDE methods are fundamentally limited because they confound two major sources of non-Gaussian diffusion at the voxel level, namely, the orientation dispersion of anisotropic diffusion and the size variance of microscopic diffusion environments, resulting in a lack of specificity. For example, consider the hypothetical voxels shown in Figure 3.6 that contain sets of non-exchanging microscopic diffusion environments that are fully characterized by microscopic diffusion tensors. The sizes of the spherical microenvironments correspond to the projections of the anisotropic microenvironments. Since the measured signal depends on the distribution of apparent diffusivities along the direction of the applied gradient, the two voxels result in nearly indistinguishable signal decay curves. This simple toy model illustrates the fundamental degeneracy of SDE acquisitions[115].

Consider a more realistic voxel that can be divided into separate micro-

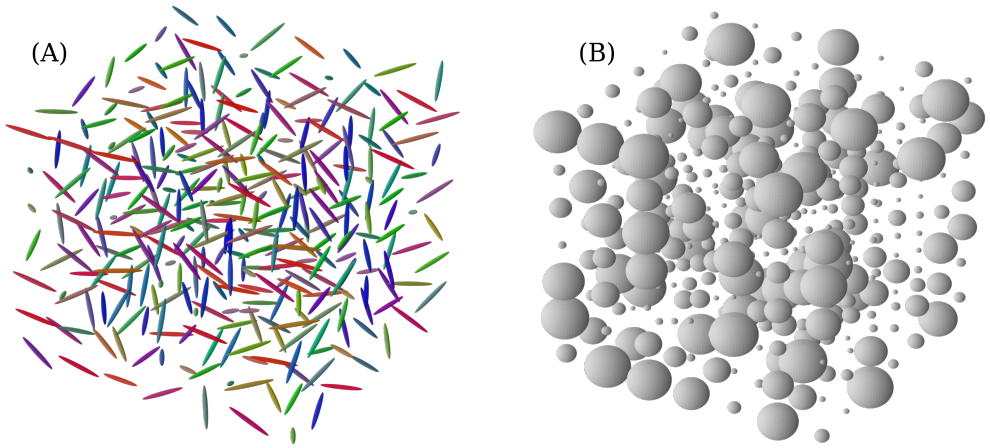

Figure 3.6: *Two hypothetical voxels that result in nearly indistinguishable signal decay curves using SDE. The voxel in (A) consists of highly anisotropic microscopic diffusion environments with a high degree of orientation dispersion. The voxel in (B) consists of isotropic microscopic diffusion environments with high size variance. The sizes of the microenvironments in (B) correspond to the sizes of the projections of the microenvironments in (A). It has been shown that meningiomas and glioblastomas can result in similar signal decay curves using SDE, although at the microscopic level meningiomas resemble (A) and glioblastomas (B)[55]. In a healthy brain, white matter consists mainly of highly anisotropic microenvirionments, whereas isotropic size variance can be more easily observed in gray matter.*

scopic diffusion environments. If the exchange of water between the microenvironments is negligible, the imaged tissue can be represented by a set of time-dependent microscopic diffusion propagators. Given that the diffusion time is long enough for the water molecules to thoroughly explore their environment, each propagator beomes Gaussian, as explained by the central limit theorem[100]. Therefore, in the long diffusion time limit ($\Delta \to \infty$), diffusion in each microenvironment can be characterized by a microscopic diffusion tensor $\mathbf{D}_\mu$, where the subscript μ is used to distinguish it from the voxel-level diffusion tensor $\mathbf{D}$, which relates to the microscopic diffusion tensors as

$$\mathbf{D} = \langle \mathbf{D}_\mu \rangle, \tag{3.7.1}$$

where the averaging operator $\langle\ \rangle$ represents averaging over all microenvironments in the voxel.

A multidimensional b-tensor, as defined by Equation 3.3.6, can be defined by applying diffusion encoding along two or three dimensions between the spin

excitation and signal readout. For example, the following b-tensors correspond to one-, two-, and three-dimensional acquisitions:

$$\mathbf{b}_{\mathrm{LTE}} = \begin{bmatrix} b & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \tag{3.7.2}$$

$$\mathbf{b}_{\mathrm{PTE}} = \frac{1}{2}\begin{bmatrix} 0 & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & b \end{bmatrix}, \tag{3.7.3}$$

$$\mathbf{b}_{\mathrm{STE}} = \frac{1}{3}\begin{bmatrix} b & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & b \end{bmatrix}, \tag{3.7.4}$$

where b is the b-value, and the subscripts LTE, PTE, and STE stand for linear, planar, and spherical tensor encoding, respectively. The tensors are rotated according to the diffusion encoding direction.

By combining acquisitions with different b-tensor shapes it is possible to disentangle the orientation dispersion of anisotropic diffusion from the size variance of microscopic diffusion environments[115]. Figure 3.7 shows simulated signal decay curves of acquisitions using linear and spherical with b-tensors. The deviation between the signal decay curves acquired with the two b-tensor shapes indicates that the diffusion is anisotropic at the microscopic level, and the non-monoexponential signal decay in STE acquisitions indicates size variance of microscopic diffusion environments[115].

According to the conventions and nomenclature by Shemesh et al.[52], microscopic diffusion anisotropy (μA) is defined so that it squared is proportional to the average eigenvalue variance over the microscopic diffusion tensors in the voxel:

$$\mu\mathrm{A}^2 \propto \langle \mathrm{Var}_\lambda(\mathbf{D}_\mu) \rangle. \tag{3.7.5}$$

Since μA depends on both the shape and magnitude of the microscopic diffusion tensors, microscopic fractional anisotropy (μFA) has been introduced to quantify microscopic diffusion anisotropy irrespective of the orientation dispersion and the magnitude of diffusion[134,135]:

$$\mu\mathrm{FA} = \sqrt{\frac{3}{2}\frac{\langle \mathrm{Var}_\lambda(\mathbf{D}_\mu) \rangle}{\langle \mathrm{Var}_\lambda(\mathbf{D}_\mu) \rangle + \langle [\mathrm{Tr}(\mathbf{D}_\mu)/3]^2 \rangle}}. \tag{3.7.6}$$

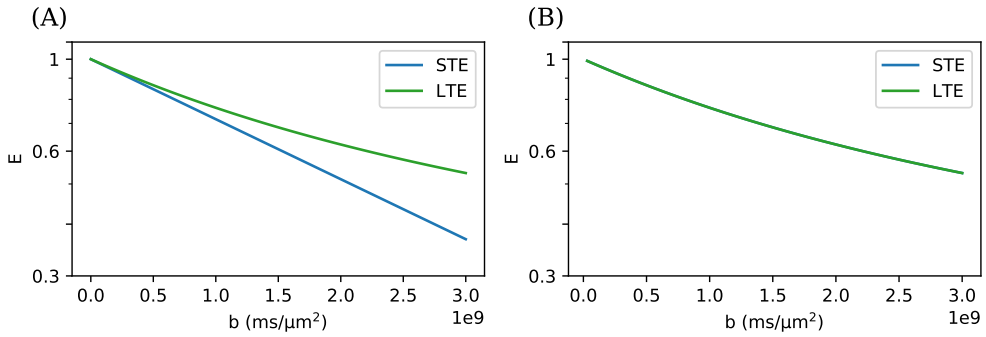

Figure 3.7: *Example of resolving the fundamental degeneracy of SDE acquisitions by combining acquisitions with linear tensor encoding (LTE) and spherical tensor encoding (STE). Similarly to Figure 3.6, the simulated voxel in (A) consists of highly anisotropic microscopic diffusion environments with a high degree of orientation dispersion, and the simulated voxel in (B) consists of isotropic microscopic diffusion environments with high size variance.*

μFA takes values between 0 and 1 and is equal to conventional FA if the microscopic diffusion tensors are aligned. A simple example highlighting the differences between FA and μFA is provided in Figure 3.8. Since μA and μFA do not depend on the orientation dispersion of axons, they may be clinically valuable biomarkers for studying axonal degeneration. Indeed, promising results of microscopic anisotropy imaging have been reported in imaging multiple sclerosis lesions[56,60] and white matter degeneration in Parkinson's disease patients[63]. Furthermore, microscopic anisotropy imaging has been applied to quantify the average cell shape in brain tumours[64,55] and microstructural properties of brains of patients with schizophrenia[65] and epilepsy[59]. μFA estimation with DDE and QTE will be discussed in detail in Chapter 6 and Chapter 7, respectively.

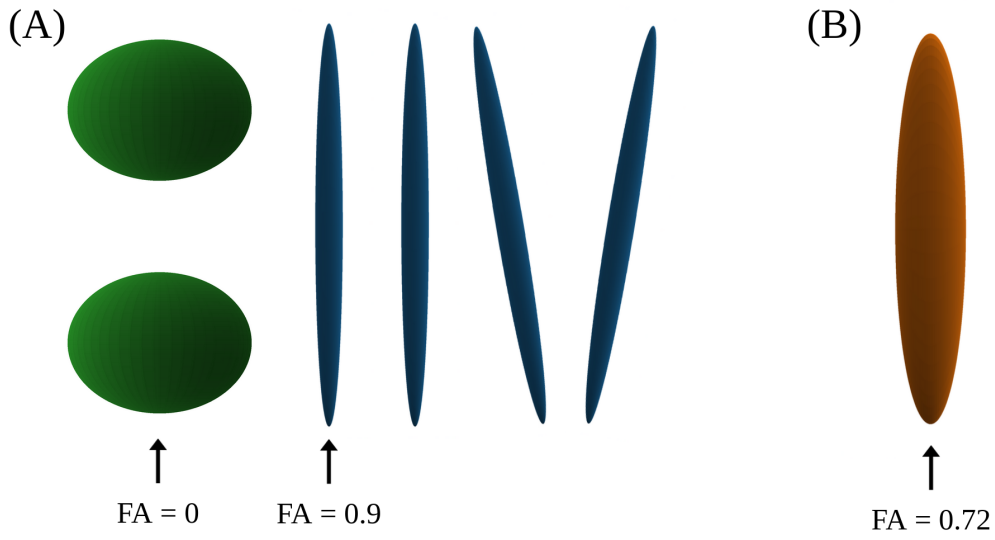

Figure 3.8: *Example of a hypothetical voxel consisting of the microscopic diffusion tensors shown in (A) that correspond to a voxel-level diffusion tensor shown in (B). FA of each tensor is shown. MD of each tensor is 1 $\mu m^2/ms$. In this case, μA = 0.88 $\mu m^2/ms$ and μFA = 0.84. The voxel-level FA is lower than μFA because of the orientation dispersion of the anisotropic microscopic diffusion tensors and the presence of isotropic microscopic diffusion tensors.*

3.8 Summary

This chapter provided a concise overview of dMRI, describing how the MR signal can be made sensitive to diffusion of water molecules, which in turn can be used to study the microstructure of neural tissue. The following chapters describe the original contributions of the PhD candidate.

Chapter 4

GPU-accelerated Monte Carlo random walk simulator

4.1 Introduction

Microstructural parameters calculated from dMRI data have been successfully applied to observe microstructural changes in neural tissue during normal development and in pathologies. For instance, it has been shown that FA in white matter tends to decline with aging[120] and elevated MK values have been observed in the basal ganglia of patients with Parkinson's disease[136]. Despite their potential utility as biomarkers, connecting the values of the estimated parameters to the biophysical properties of tissue is challenging because the ground truth is usually unknown in *in vivo* experiments and the observed changes in the metrics' values can be driven by multiple factors. Thus, the validation of dMRI methods and the resulting parameter maps is an important problem.

Estimated parameter values can be compared to the ground truth, for example, by imaging phantoms with a known microstructure[137] or by performing histology experiments on the imaged tissue[138]. However, such experiments are expensive and time-consuming, especially if the acquisition is repeated using several acquisition protocols or microstructural configurations. On the other hand, simulations with a well-defined ground truth do not require the use of expensive scanner time and provide a powerful tool for assessing the accuracy and precision of parameter estimates. Generally, dMRI simulations are based either on modelling diffusion in some geometry under a set of assumptions about the diffusion process and the pulse sequence to obtain an approximate solution to the diffusion equation, from which an expression for the diffusion-

weighted signal can be derived[139], or on a more generalizable Monte Carlo random walk approach[140].

Some of the simplest dMRI simulations model the tissue as a combination of non-exchanging microscopic diffusion environments that can be characterized by microscopic diffusion tensors. In this case, the normalized dMRI signal can be expressed as

$$E = \frac{1}{N} \sum_{i=1}^{N} \exp\left(-\mathbf{b} : \mathbf{D}_{\mu i}\right), \tag{4.1.1}$$

where N is the number of compartments, $\mathbf{b}$ is the b-tensor used in the simulated acquisition, $\mathbf{D}_{\mu i}$ is the ith microscopic diffusion tensor, and : denotes the generalized scalar product between the two tensors. Equation 4.1.1 provides a computationally inexpensive method for assessing the robustness of microstructural parameter estimation methods against various tissue models defined by distributions of microscopic diffusion tensors[141]. However, this method fails to capture time-dependent diffusion and thus can not be applied to study, for instance, the effect of gradient pulse duration or diffusion time on the parameter estimates, for which more complicated simulations are needed.

Experiments of restricted diffusion in simple geometries, such as spheres and cylinders, can be simulated using analytical signal expressions that can be derived for gradient sequences that fulfil a set of assumptions, e.g., infinitesimal gradient pulses[48]. Simulation experiments of restricted diffusion in simple geometries with arbitrary gradient sequences can be performed using the matrix formalism, which represents the gradient by a succession of sharp gradient impulses, enabling the signal attenuation to be expressed as a computationally inexpensive product of matrix operations[142,143]. Furthermore, finite difference methods, finite element methods, and the lattice Boltzmann method can be used to numerically find an approximate solution to the Bloch-Torrey equation in more complicated geometries[144–146]. The Monte Carlo random walk method, which models the diffusion process as an ensemble of random walks, is another method that is easily generalizable to arbitrary geometries and gradient sequences[147–149].

The Monte Carlo method involves repeated random sampling to obtain numerical solutions to the studied problem[150,151]. In dMRI research, the Monte Carlo method is often applied to simulate diffusion by generating a large number of random walks from which the simulated signal can be calculated. Such simulations have been used since the 1990s[152,140], and several open source software packages, such as Camino[153,147], written in Java, and MC/DC[154], writ-

ten in C++, have been developed to facilitate performing simulation experiments. Since the random walks are independent and generated from independent and identically distributed pseudorandom numbers, they can be generated in parallel, leading to a significant reduction in simulation runtime. Recent developments in the hardware and software of general-purpose graphical processing unit (GPU) computing have facilitated running scientific programs in massively parallel[155], i.e., simultaneously on at least hundreds of cores. GPU-acceleration has been applied to random walk dMRI simulations[156], however, in many cases, the source code of GPU-accelerated simulations has not been shared or is not easy to use or modify.

This chapter presents Disimpy[157], a GPU-accelerated random walk dMRI simulator developed by the PhD candidate during the course of the doctorate. Disimpy is written in Python (Python Software Foundation), a programming language know for its readability and clarity, which makes the source code very approachable and easily extensible to researchers and students with little or no prior experience in general-purpose GPU computing. The GPU-acceleration results in a significant performance gain compared to serial computation, enabling simulation experiments to be performed on standard laptop and desktop computers without needing to access high-performance computing clusters.

4.2 Simulator details

Disimpy is a massively parallel Monte Carlo random walk dMRI simulator that generates random walks on Nvidia's (Nvidia Corporation, Santa Clara, California, United States) CUDA-capable GPUs with compute capability 2.0 or higher[155]. Disimpy provides a simple user interface for performing simulations and useful helper functions for generating and manipulating gradient arrays. The source code is openly available on Github at https://github.com/kerkelae/disimpy/ under the MIT license. The documentation, tutorial, and contributing guidelines are provided on the simulator's home page at https://disimpy.readthedocs.io/. The simulator requires the user to define the simulated microstructure, gradient waveforms, and at least the following parameter values: number of random walkers, number of time steps, and diffusivity. The output of the simulator is the simulated signal for each waveform. Optionally, the trajectories of the random walkers and the signal from each individual walker can also be saved. To generate synthetic dMRI signals, Disimpy follows

the framework established by Hall et al.[147] with some differences. The details of the simulator are described in the following sections.

Dynamics

The diffusion process is modelled as an ensemble of three-dimensional random walks over discrete time. Each random walker represents a diffusing spin. The initial positions of the random walkers are drawn from a uniform distribution across the diffusion environment or defined manually by the user. The steps are uniformly distributed over the surface of a sphere and have a fixed length

$$l = \sqrt{6 \cdot D \cdot dt}, \tag{4.2.1}$$

where D is the diffusion coefficient and dt is the duration of the time step. dt is determined by the duration of the simulated dynamics and the number of time points in the simulation. A single step is calculated as

$$\mathbf{s} = \frac{l}{x^2 + y^2 + z^2} \begin{bmatrix} x \\ y \\ z \end{bmatrix}, \tag{4.2.2}$$

where x, y, and z are normally distributed pseudorandom numbers. The pseudorandom numbers are generated by randomly sampling from a uniform distribution on the interval $[0, 1]$ using the xoroshiro128+ pseudorandom number generator[158]. The Box-Muller transform is then applied to obtain normally distributed pseudorandom numbers. This method produces a uniform density of steps across the surface of a sphere[159].

Diffusion can be simulated without restrictions, inside simple analytically defined geometries (cylinders, spheres, ellipsoids), and in arbitrary geometries defined by triangular meshes[160–162]. Periodic and reflective boundary conditions are supported for triangular meshes. If the step of a random walker intersect with a surface, it is elastically reflected off the intersection point in such a way that the random walker's total path length during dt is equal to l. The intersection check is repeated after the reflection to check for multiple collisions during dt. By default, the surfaces are impermeable. Permeable surfaces can be simulated by enabling the walkers to pass through the surface with a given probability. An example of restricted diffusion is shown in Figure 4.1.

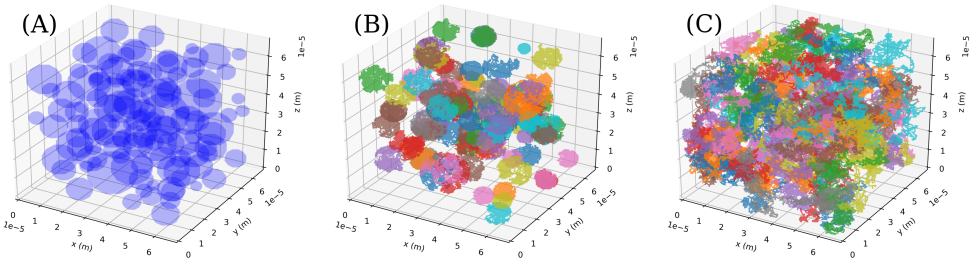

Figure 4.1: *Example of restricted diffusion in an environment defined by a triangular mesh. (A) Example mesh of* 10^4 *triangles defining the synthetic voxel consisting of 100 spheres with gamma-distributed radii. (B) Example trajectories of 100 random walkers whose initial positions were randomly positioned inside the spheres. Some spheres contain more than one walker. (C) Example trajectories of 100 random walkers outside the spheres. In this example, the boundaries of the voxel are impermeable surfaces.*

Signal generation

The simulated signal is generated according to equations 3.3.1 and 3.3.2. Several signals corresponding to different diffusion encoding gradients can be generated from the same simulated dynamics. The gradient waveforms are represented as three-dimensional arrays where the dimensions correspond to the different measurements, time points, and gradient magnitude along different directions. At every time point t in the simulation each random walker accumulates phase

$$d\phi = \gamma \mathbf{G}(t) \cdot \mathbf{r}(t) dt, \tag{4.2.3}$$

where γ is the gyromagnetic ratio of hydrogen ^{1}H, $\mathbf{G}$ is the diffusion encoding gradient, and $\mathbf{r}$ is the location of the random walker. At the end of the simulated dynamics, the signal is calculated as the sum of the real parts of the signals from all random walkers:

$$S = \sum_{j=1}^{N} \mathrm{Re}\left(\exp\left(i\phi_j\right)\right), \tag{4.2.4}$$

where ϕ_j is the phase shift accumulated by the jth random walker and N is the number of random walkers. The normalized diffusion-weighted signal is simply

$$E = \frac{S}{N}. \tag{4.2.5}$$

Implementation

Disimpy uses efficient data structures and numerical methods from Numpy[163] and Scipy[164]. Numba[165] is used to compile Python code into CUDA kernels and device functions[155] which are executed on the GPU. All computation is performed using double-precision floating-point numbers. The random walks are generated in massively parallel on individual threads of one-dimensional CUDA blocks, resulting in a performance gain of over an order of magnitude when compared to Camino[147] which runs serially on the central processing unit (CPU) (Figure 4.2). The maximum number of parallel threads is determined by the GPU but is generally in the tens of thousands[155].

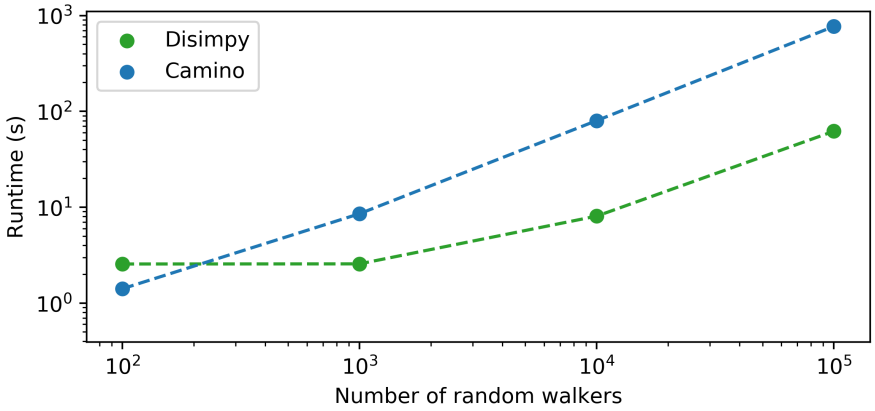

Figure 4.2: *Performance comparison between Disimpy and Camino[147], a popular dMRI simulator that runs single-threaded on the CPU. The comparison was performed on a desktop computer with an Intel Xeon E5-1620 v3 3.50 GHz x 8 CPU and an Nvidia Quadro K620 GPU. The simulations were performed using a mesh consisting of* 10^4 *triangles, shown in Figure 4.1. Significantly improved Disimpy runtimes can be achieved by using a more modern GPU because the GPU used in this comparison is not particularly well suited for double-precision arithmetic. In the future, GPUs with more parallel threads and faster cores will become increasingly more affordable, and the runtime difference between the GPU-accelerated and serially executed simulations will further increase.*

4.3 Validation

Accuracy and precision of Disimpy has been validated by comparing it against analytical signal expressions and results obtained from other established sim-

ulators, such as Camino[147]. Furthermore, code modularity has been a priority during development and unit tests have been written to test the functionality of individual functions. Two examples of the simulations used for validation in terms of accuracy are shown in Figure 4.3

It is well-known that the results of Monte Carlo simulations are valid only if a sufficiently large amount of random sampling has been performed[151]. To assess the convergence in terms of the number of random walkers, simulations of free and restricted diffusion were repeated with differet pseudorandom number generator seeds and the signal variation was quantified. The results, shown in Figure 4.4, reveal that the required number of spins depends on both the simulated microstructure and the magnitude of the diffusion-weighted signal attenuation. The results suggest that at least 10^4 random walkers should be used for all scenarios. However, since the convergence also depends on the number of steps in the simulation and the simulated gradient sequence, a similar analysis of convergence should be conducted when performing simulation experiments to confirm that sufficient convergence has been achieved.

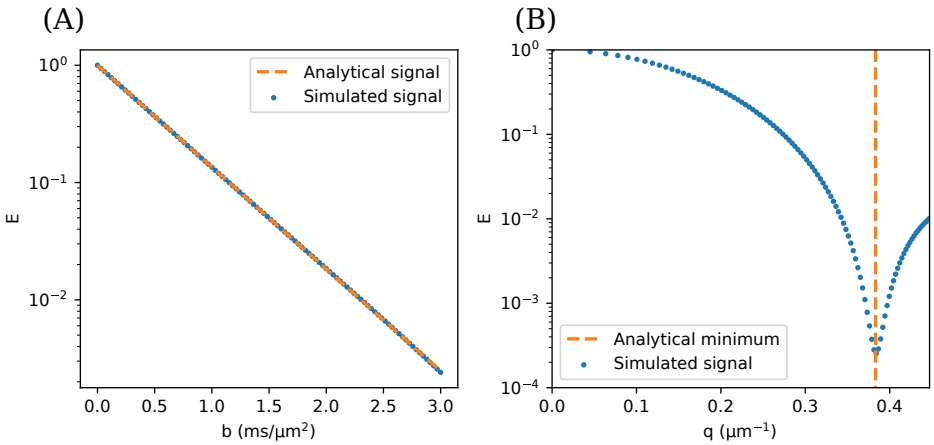

Figure 4.3: *Two examples of the simulations performed to validate Disimpy's output. The simulations were performed with 10^6 random walkers and 10^5 time steps. (A) A comparison between the analytically calculated signal decay (Equation 3.4.8) and Disimpy's simulated signal decay for free diffusion. The simulated gradient sequence was PGSE with Δ = 40 ms and δ = 10 ms. (B) Reproduction of the diffraction pattern that is observable in experiments with diffusion inside cylinders perpendicular to the gradient direction and very high b-values*[166]*. The analytical minimum is equal to $0.61 \cdot 2 \cdot \pi / r$ where r is the cylinder radius*[167]*. Diffusion was simulated inside an infinite cylinder with radius of 10 μm. The simulated gradient sequence was PGSE with Δ = 500 ms and δ = 0.1 ms.*

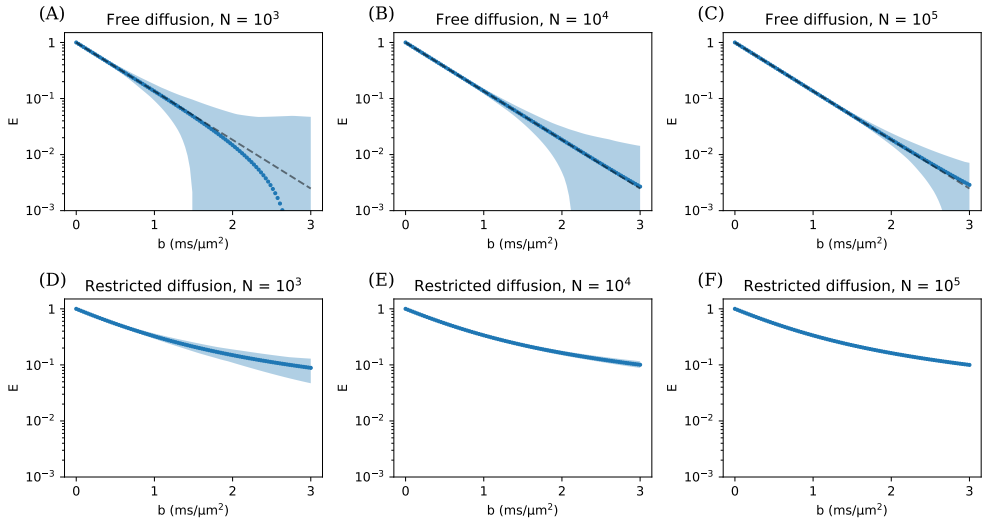

Figure 4.4: *The signal variation over repeated experiments with different numbers of random walkers. (A)-(C) show the results for free diffusion. (D)-(F) show the results for restricted diffusion using the mesh shown in Figure 4.1. The simulations were repeated 100 times. The blue dots represent the mean signal and the shaded area represents the 95% percentile range. The dashed line is the ground truth, which is known for free diffusion but unknown for restricted diffusion. The simulated gradient sequence was PGSE with Δ = 60 ms and δ = 12 ms. The simulations were performed with 10^4 time steps.*

4.4 Discussion

Disimpy facilitates the development and validation of dMRI methods by enabling large amounts of synthetic signals to be generated on standard desktop and laptop computers without needing to access high-performance computing clusters. Disimpy can also easily be run on popular cloud computing platforms, such as Google Colaboratory, that offer access to CUDA-capable GPUs. The way Disimpy generates the simulated signal is very similar to Camino[147]. However, on standard workstation computers, Disimpy is up to several orders of magnitude faster than Camino. Furthermore, the source code of Disimpy is very readable and easily extensible because it is written in Python, making the simulator accessible to other researchers and students. This is an important difference to Camino[153,147], written in Java, and MC/DC[154], written in C++, which are more difficult to modify. Although the performance is sufficient for this thesis, the implementation of some of the algorithms is suboptimal in terms of computational complexity and memory management. The long-term vision is to optimize the code and add more features while Disimpy is used in future research. However, there are several important issues that must be mentioned here and considered prior to using Disimpy in research.

It is crucial to choose the simulated microstructure so that the simulation is representative of the studied problem. The simple geometries included in the software, such as infinite cylinders, are most often not sufficiently representative of neural tissue microstructure. Thus, the use of more realistic models of tissue microstructure, often referred to as substrates, is encouraged. For example, contextual fibre growth (ConFiG) is a recently developed method for generating axon-mimicking fibres with geometrical properties that closely resemble those of real axons[168]. The generation of realistic substrates is outside the scope of this thesis, but Disimpy can be used with ConFiG, as is done in Chapter 7. The size of the simulated substrate is limited by the available memory on the GPU.

To obtain an accurate simulated signal, the parameter values have to be chosen so that sufficient convergence is achieved. As demonstrated by the simulation results in Figure 4.4, which are in agreement with previously reported results[147,154], 10^5 random walkers and 10^4 time steps can be considered to be the minimum as a rule of thumb. However, since the convergence of the results depends on the simulated microstructure and gradient sequence, a similar analysis should be repeated to confirm that the simulation results are accurate.

In terms of performance, it is important that Disimpy's runtime does not linearly depend on the number of random walkers until it is in the thousands or tens of thousands (Figure 4.2), depending on the GPU.

Finally, there are several limitations that apply to Monte Carlo random walk dMRI simulations in general. For example, the simulations do not consider any other signal decay mechanisms besides diffusion and assumes that the magnetization relaxation properties are uniform across the simulated volume. Furthermore, a simple elastic collision off the simulated surface is probably not a realistic representation of the interaction between the water molecules and the cell membranes. However, despite the limitations, the simulations can be a powerful tool for assessing the robustness of dMRI methods.

4.5 Conclusion

This chapter presented the dMRI simulator developed by the PhD candidate. The software enables efficient simulations of time-dependent diffusion and is used in Chapter 7 to study the accuracy and precision of μFA estimates.

Chapter 5

Reproducibility of microscopic anisotropy estimation using microstructural modelling

5.1 Introduction

As an established method for analyzing dMRI data, DTI is regularly used in both research and clinical settings[112]. There is a general consensus in the field about the utility of DTI and its use is facilitated by its inclusion in the scanner software by several scanner manufacturers. However, as explained in Chapter 3, DTI is unable to capture non-Gaussian diffusion, making it a suboptimal method, given the acquisition protocols supported by modern scanners. Modern diffusion-weighted imaging protocols are often multi-shell PGSE, i.e., the acquisitions are performed with more than one non-zero b-value. Such protocols enable the quantification of the signal deviation from a monoexponential decay, a signature of non-Gaussian diffusion. Even though a plenitude of methods have been developed for analyzing multi-shell data, none of them have yet become established standards like DTI. As a consequence, different studies have applied different methods to quantify the same microstructural properties of interest and as a result it can be challenging to interpret reported results in the literature. For a consensus to be achieved on the preferred method for analyzing multi-shell data for a given situation, the accuracy and precision of the proposed methods have to be carefully assessed.

As a logical extension of DTI, DKI offers a model-free way for quantifying the rotationally variant signal deviation from monoexponential decay without

a priori assumptions about tissue microstructure. However, DKI is not specific to the source of non-Gaussian diffusion which can be caused, at the voxel level, by the orientation dispersion of anisotropic diffusion, size variance between microscopic diffusion environments, and time-dependent diffusion. Microstructural models, such as NODDI[66] and SMT[67,68], on the other hand, aim to characterize the imaged microstructure using descriptive metrics, such as intra-neurite volume fraction and mean axon radius. The purpose of the work presented in this chapter is to discuss the validity of these models methods in estimating microscopic diffusion anisotropy and to quantify their reproducibility using a scan-rescan dataset acquired with a standard clinical multi-shell protocol.

In the so-called standard model of dMRI in neural tissue, tissue is represented as a combination of non-exchanging Gaussian compartments and the diffusion-weighted signal from each compartment is expressed as a convolution between the ODF and response kernel defined by microstructural parameters[126]. Due to the flat fitting landscape of the standard model in the b-value range achievable with typical human scanners, strict constraints have to be imposed for efficient and reproducible parameter estimation[169]. In fact, many microstructural modelling methods known by their own acronyms can be seen as iterations of the standard model with specific compartments and fit constraints[49,66–68]. In this study, the focus is on NODDI[66] and SMT[67,68], two versions of the standard model that have been developed so that they are possible to fit to data acquired with a relatively small number of acquisitions and modest gradient strengths, making them clinically feasible.

Since its introduction in 2012, NODDI has become a popular method that has been applied to measure microstructural changes in the brain associated with a wide variety of conditions, such as Alzheimer's disease[170], Parkinson's disease[171], multiple sclerosis[172], and sickle cell disease[123]. In NODDI, the values of the diffusion coefficients and the shape of the ODF are fixed *a priori* to facilitate fitting the model. NODDI enables the estimation of the so-called volume fractions of intra-neurite diffusion, extra-neurite diffusion, and isotropic diffusion as well as the degree of orientation dispersion of neurites.

SMT is a more recent method, introduced in 2015, that enables the estimation of microstructural parameters irrespective of orientation dispersion and which has been applied in studying microstructural changes associated with multiple sclerosis[173,174] and congenital hypothyroidism[175], for example.

In SMT, the model is fit to powder-averaged data, i.e., data averaged over all the acquired diffusion encoding directions, to exclude the effects of orientation dispersion.

Despite the popularity of these models, as well as other versions of the standard model, both NODDI and SMT have been shown to produce biased and contradictory estimates which will be discussed more in the theory section of this chapter[132,133]. However, given the popularity of these methods in neuroscience research, it is important to assess their reproducibility. Furthermore, even an inaccurate microstructural model may be valuable if it is sensitive to relevant microstructural changes in the brain. Thus, in the work presented in this chapter, the reproducibility of the microstructural metrics extracted from NODDI and SMT is evaluated and the results are compared to those obtained from DTI and DKI. Specifically, this study focuses on the microstructural metrics quantifying microscopic diffusion anisotropy, i.e., intra-neurite volume fraction and microscopic fractional anisotropy.

5.2 Theory

NODDI

The NODDI tissue model consists of three non-exchanging Gaussian compartments: intra-cellular space, extra-cellular space, and CSF. Using this tissue model, the normalized signal can be expressed as

$$E_{\text{NODDI}} = (1 - v_{\text{iso}})(v_{\text{ic}}E_{\text{ic}} + (1 - v_{\text{ic}})E_{\text{ec}}) + v_{\text{iso}}E_{\text{iso}}, \tag{5.2.1}$$

where v_{iso} and E_{iso} are the normalized volume fraction and signal of the CSF compartment, v_{ic} and E_{ic} are the normalized volume fraction and signal of the intra-cellular compartment, and v_{ec} and E_{ec} are the normalized volume fraction and signal of the extra-cellular compartment. It would be more appropriate to refer to the volume fractions as signal fractions because the signal contributions of the different compartments also depend on their magnetization relaxation mechanisms instead of only volumes. However, this chapter will follow the terminology used in the publication that first introduced the model[66].

The cells in NODDI are highly anisotropic neurites inside which diffusion is considered to be one-dimensional. A central assumption of NODDI is that

the neurite ODF follows a Watson distribution:

$$\mathrm{ODF}(\mathbf{u}) = M\left(\frac{1}{2}, \frac{3}{2}, \kappa\right)^{-1} \exp\left[\kappa(\boldsymbol{\mu} \cdot \mathbf{u})^2\right], \tag{5.2.2}$$

where $\mathbf{u}$ is the orientation represented by a unit vector, M is the confluent hypergeometric function, $\boldsymbol{\mu}$ is the average orientation represented by a unit vector, and κ is the concentration parameter measuring dispersion about $\boldsymbol{\mu}$. Using the symmetry of the Watson distribution, the dispersion of the neurites is quantified with the orientation dispersion index (OD):

$$\mathrm{OD} = \frac{2}{\pi} \arctan\left(\frac{1}{\kappa}\right). \tag{5.2.3}$$

OD ranges from 0 for aligned neurites to 1 for neurites whose orientations are uniformly distributed across the surface of a sphere. Given the ODF, the normalized intra-cellular signal can be written as

$$E_{\mathrm{ic}} = \int_{\mathbb{S}^2} \mathrm{ODF}(\mathbf{u}) \exp\left[-b d_{||} (\mathbf{n} \cdot \mathbf{u})^2\right] d\mathbf{u}, \tag{5.2.4}$$

where $\mathbf{n}$ is the direction of the applied gradient, b is the applied b-value, $d_{||}$ is the intra-neurite diffusivity, and integration is performed over the surface of a sphere.

The extra-cellular compartment covers the space outside the neurites that includes cell bodies, glial cells, extra-cellular space etc. The extra-cellular signal is expressed as

$$E_{\mathrm{ec}} = \exp\left[-b \mathbf{n}^{\mathrm{T}} \left(\int_{\mathbb{S}^2} \mathrm{ODF}(\mathbf{u}) \mathbf{D}(\mathbf{u}) d\mathbf{u}\right) \mathbf{n}\right], \tag{5.2.5}$$

where b is the b-value, $\mathbf{n}$ is the diffusion encoding direction, and $\mathbf{D}(\mathbf{u})$ is an axially symmetric diffusion tensor with principal diffusion direction $\mathbf{u}$ and with axial and radial components of $d_{||}$ and $d_{||}(1 - v_{\mathrm{ic}})$, respectively.

The CSF compartment, represented by an isotropic diffusion tensor, is introduced to account for partial volume effects from CSF. The normalized signal from CSF is simply

$$E_{\mathrm{iso}} = \exp\left(-b d_{\mathrm{iso}}\right), \tag{5.2.6}$$

where b is the b-value of the applied gradient and d_{iso} is the diffusivity of the isotropic compartment.

The values of $d_{||} = 1.7$ μm^2/ms and $d_{\mathrm{iso}} = 3$ μm^2/ms are fixed *a priori* to reduce the model to the following unknown parameters: v_{ic}, v_{iso}, κ, and $\boldsymbol{\mu}$.

With these constraints the NODDI model can be fit to data acquired with just two HARDI shells.

It is possible to calculate μFA from the NODDI parameter estimates using Equation 3.7.6 by weighting the terms with the corresponding volume fractions when calculating the averages. However, the resulting μFA estimate would be biased for at least two reasons[132]. First, since NODDI is based on SDE acquisitions, the signal suffers from the fundamental degeneracy of SDE acquisitions and thus the size variance of microscopic diffusion environments is confounded with the orientation dispersion of anisotropic diffusion. Second, the assumption that the intra-neurite volume fraction is directly connected to the voxel-level MD predicts a greater contrast in MD maps between white and gray matter than what is measured, raising questions about the validity of the NODDI tissue model[132,169].

SMT

SMT is a technique that consists of fitting the standard model of dMRI in neural tissue to powder-averaged SDE data which is invariant of the ODF, given that the acquisition was performed with diffusion encoding directions that cover the surface of the sphere sufficiently well. In 2015, a single-compartment model of SMT was introduced[67], followed by a two-compartment model of SMT in 2016[67]. From hereafter, SMT1 refers to the single-compartment model and SMT2 refers to the two-compartment model.

SMT1

In SMT1, the tissue is modelled as a combination of identical and axially symmetric microscopic diffusion environments characterized by microscopic diffusion tensors with arbitrary orientations. Under this assumption, the normalized powder-averaged signal can be expressed as

$$E_{\text{SMT1}} = \exp\left(-b\lambda_{\perp}\right) \frac{\sqrt{\pi}\text{erf}\left(\sqrt{b(\lambda_{\parallel} - \lambda_{\perp})}\right)}{\sqrt{b(\lambda_{\parallel} - \lambda_{\perp})}}, \tag{5.2.7}$$

where b is the b-value of the applied gradient, erf is the error function, and $\lambda_{\parallel}$ and $\lambda_{\perp}$ are the axial and radial diffusivities of the microscopic diffusion tensor, respectively.

SMT1 parameters can be estimated by fitting Equation 5.2.7 to the powder-averaged data. The shape and size of the estimated microscopic diffusion tensor can be quantified using microscopic versions of the standard DTI metrics

(Equations 3.4.13 - 3.4.16): μMD, μAD, μRD, and μFA, where μ is used to distinguish them from the voxel-level DTI metrics.

SMT1 is valid if the imaged microstructure can be accurately represented by a distribution of identical microscopic diffusion tensors that only vary in their orientation. However, this assumption does not hold in the human brain where size and shape variance invariably exist between the microscopic diffusion environments, e.g., axons and extra-axonal space. Thus, SMT1 significantly overestimates μFA in the human brain because it interprets all non-Gaussian diffusion to be caused by microscopic diffusion anisotropy[133].

SMT2

SMT2 expanded the SMT framework to two Gaussian compartments representing intra- and extra-neurite diffusion. The diffusivities and volume fractions of the compartments are λ^{intra} and λ^{extra}, and v_{intra} and v_{extra}, respectively. Again, signal fraction would be a more appropriate term than volume fraction but this chapter follows the terminology used in the publication that first presented the method[67].

Similarly to NODDI, several *a priori* constraints are imposed to facilitate fitting the model. First, diffusion inside neurites is considered to be effectively one-dimensional: $\lambda_{\perp}^{\text{intra}} = 0$. Second, axial diffusivities of the intra- and extra-neurite compartments are assumed to be equal: $\lambda^{\text{intra}} = \lambda_{\parallel}^{\text{extra}}$. Third, extra-neurite diffusion is assumed to follow the tortuosity model[140]: $\lambda_{\perp}^{\text{extra}} = (1 - v_{\text{intra}})\lambda_{\parallel}^{\text{extra}}$. Given this model, the normalized powder-averaged signal can be expressed as

$$\begin{aligned} E_{\text{SMT2}} = v_{\text{intra}} \frac{\sqrt{\pi}}{2} \frac{\operatorname{erf}\left(\sqrt{b\lambda^{\text{intra}}}\right)}{\sqrt{b\lambda^{\text{intra}}}} + \\ (1 - v_{\text{intra}}) \exp\left[-b(1 - v_{\text{intra}})\lambda^{\text{intra}}\right] \frac{\sqrt{\pi}}{2} \frac{\operatorname{erf}\left(\sqrt{bv_{\text{intra}}\lambda^{\text{intra}}}\right)}{\sqrt{bv_{\text{intra}}\lambda^{\text{intra}}}}, \end{aligned} \tag{5.2.8}$$

where b is the b-value of the applied gradient and erf is the error function. The model parameters of SMT2 can be estimated by fitting Equation 5.2.8 to the powder-averaged data.

Similarly to NODDI, μFA can be calculated from the SMT2 parameter estimates using Equation 3.7.6 by weighting the terms with the corresponding volume fractions when calculating the averages. However, it has been shown that SMT2 fails to accurately estimate μFA because the two-compartment model is not sufficient to represent neural tissue[133]. For instance, as a two-compartment

model, the model can not capture size variance of microscopic diffusion environments. Furthermore, SMT2 can not resolve the fundamental degeneracy of SDE acquisitions.

5.3 Methods

Volunteers

Eight healthy, right-handed volunteers with no known neurological conditions participated in this study. Each volunteer gave written and informed consent prior to each scan. The volunteers consisted of three females and five males, with ages ranging from 23 to 57 years at the time of the first scan.

Data acquisition

All imaging data was acquired on a Siemens Magnetom Prisma 3T with maximum gradient strength of 80 mT/m, maximum slew rate of 200 T/m/s, and with a 64-channel head coil (Siemens Healthcare, Erlangen, Germany). Each subject was scanned with identical scanning protocols at two time points separated by two weeks. An echo-planar multiband PGSE sequence with the following acquisition parameters was used: Δ = 28.7 ms, δ = 16.7 ms, b-values = 1 and 2.2 ms/μm^2, 60 interleaved diffusion encoding gradient directions distributed uniformly over half a sphere for both b-values, TE = 60 ms, TR = 3050 ms, FOV = 220 x 220 mm, voxel size = 2 x 2 x 2 mm^3, slice gap = 0.1 mm, 66 slices, and phase partial Fourier = 6/8. 14 acquisitions were performed without diffusion-weighting, one of which had the phase encoding direction reversed. Additionally, an anatomical T1-weighted image with isotropic voxel size of 1 mm^3 was acquired. Data acquisition was performed by Cara Foley and Kiran Seunarine at Great Ormond Street Hospital, London, United Kingdom.

Data preprocessing

All data was converted from DICOM to Nifti using TractoR[176]. Using MRtrix3[177], the diffusion-weighted data was denoised with Marchenko-Pastur random matrix denoising[178] and Gibbs ringing artefacts were estimated and corrected for with a sub-voxel shift algorithm[179]. Susceptibility and eddy current induced distortions were estimated and corrected for with topup and eddy[180]

of FSL[181]. The images were visually inspected throughout the preprocessing stages to check for artefacts and to review image quality.

Brain segmentation

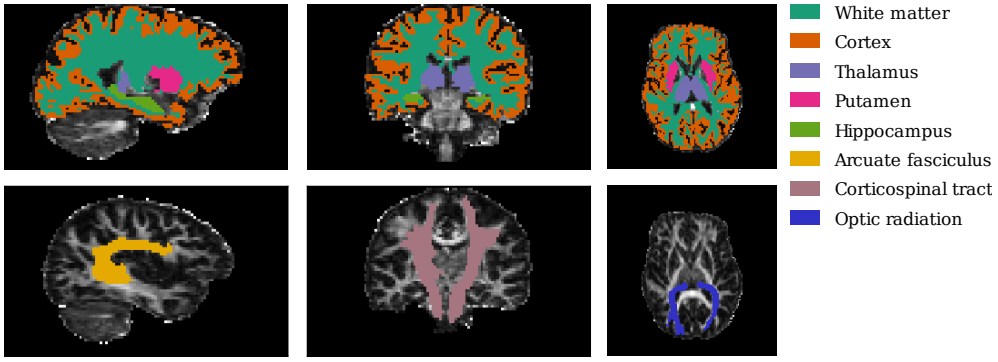

Figure 5.1: *The ROIs used in the study in a representative volunteer. The ROIs calculated using automated brain segmentation are shown in the top row. The white matter tracts defined with manually curated tractography are shown in the bottom row. Only one of the tracts is shown in each slice.*

A region of interest (ROI) based analysis was chosen for this study. The reproducibility of the mean values of the microstructural parameters were assessed in the following ROIs: white matter, cortical gray matter, thalamus, putamen, hippocampus, arcuate fasciculus (AF), corticospinal tract (CST), and optic radiation (OR). The ROIs, shown in Figure 5.1, were defined using automated brain segmentation and manually curated tractography.

White matter, cortex, thalamus, putamen, and hippocampus were delineated in both hemispheres and each subject using Freesurfer by segmenting the T_1-weighted images according to the MNI305 atlas[182,183]. The segmentations were then transformed to the space of the diffusion-weighted data using a rigid registration and resampling.

AF, CST, and OR were delineated in both hemispheres and each subject using a tractography pipeline illustrated in Figure 5.2. The fibre ODFs were estimated using multi-shell multi-tissue CSD[32] with MRtrix3[177]. Seeds, target regions, and inclusion/exclusion criteria were defined in their appropriate positions based on anatomical landmarks. Five million initial streamlines were generated for each tract using a probabilistic algorithm and the final tractograms were filtered to contain 5000 streamlines each. Tract volumes were calculated

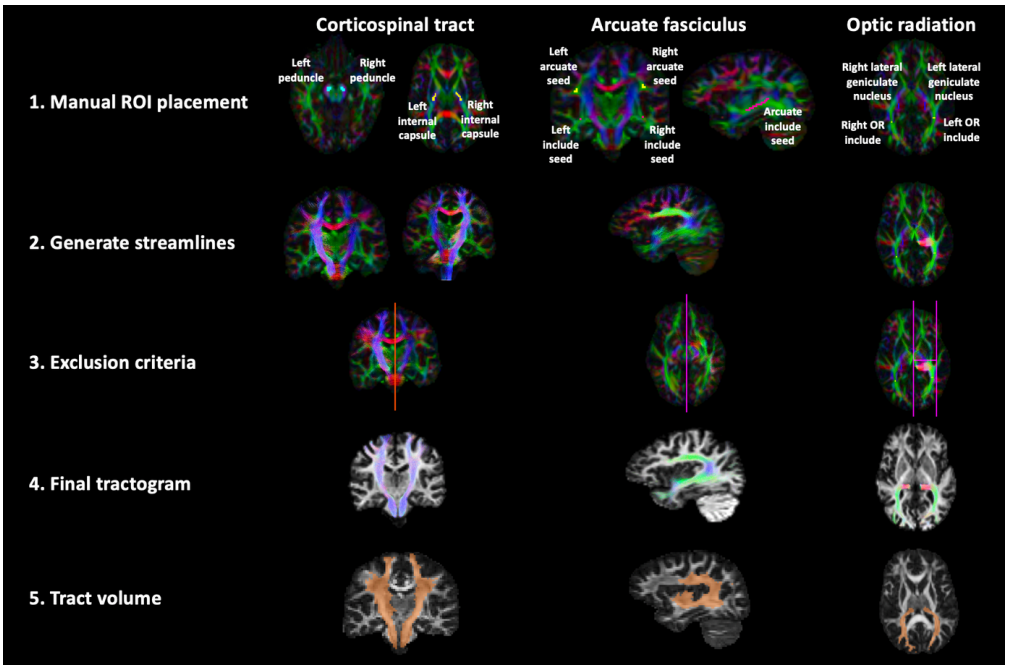

Figure 5.2: *Manual tractography pipeline used for delineating the white matter tracts. After fitting CSD to data, the ROIs were manually drawn as depicted on the colour encoded principal diffusion direction map (1). Seeds, target regions and inclusion/exclusion criteria were defined in their appropriate positions based on anatomical landmarks (1, 3). The sets of streamlines (4) were converted into binary tract maps (5) by thresholding based on streamline density.*

from the tractograms by thresholding the streamline density at 20 streamlines per voxel. The tract segmentation was performed with Kan Yan Chloe Li.

The consistency of the segmentation results was validated by visual inspection and by estimating the within-subject coefficient of variation (CV) in ROI volume by dividing the within-subject standard deviation in ROI volume by the corresponding grand mean.

Parameter estimation

The DTI and DKI parameter maps (i.e., MD, AD, RD, FA, MSK, MK, AK, and RK) were estimated with a weighted least squares algorithm using Dipy[184]. OD, v_{ic}, and v_{iso} were estimated by fitting the NODDI model to the data using the NODDI toolbox[185]. µAD, µRD, µMD, and µFA were estimated by fitting the SMT1 model to the data using the SMT toolbox[186]. v_{intra}, λ^{intra}, and λ^{extra} were estimated by fitting the SMT2 model to the data using the SMT toolbox[186]. Furthermore, SMT-based measures of orientation coherence (OC_{SMT1},

OC_{SMT2}) were calculated as the Kullback-Leibler divergence of the fibre ODF with respect to the uniform spherical distribution as reference measure, where the ODF was estimated using the kernel from SMT1 and SMT2 as described by Kaden et al.[68].

Statistical analysis

To quantify the reproducibility of the parameter estimates, a 1-way random effects model was constructed[187]:

$$x_{ij} = \mu + s_i + e_{ij}, \tag{5.3.1}$$

where x_{ij} is the value of the microstructural metric x in scan j for subject i, μ is the population mean of x, s_i is the subject-level random effect, and e_{ij} is the measurement error. It is assumed that s_j and e_{ij} are mutually independent and normally distributed: $s_j \sim \mathcal{N}(0, \sigma_s^2)$, $e_{ij} \sim \mathcal{N}(0, \sigma_e^2)$. Since the subjects of the study were healthy volunteers who were scanned two weeks apart with no treatment or other significant events taking place between the scans, a scan-level fixed effect was not included in the model.

The model parameters μ, σ_s, and σ_e were estimated as[187,188]

$$\mu \approx \frac{1}{nk} \sum_{i=1}^{n} \sum_{j=1}^{k} x_{ij}, \tag{5.3.2}$$

$$\sigma_s \approx \sqrt{\frac{\text{MSBS} - \text{MSWS}}{k}}, \tag{5.3.3}$$

$$\sigma_e \approx \sqrt{\text{MSWS}}, \tag{5.3.4}$$

where n is the number of subjects, k is the number of scans, and MSBS and MSWS are the mean squares between-subjects and within-subjects, respectively, calculated as follows:

$$\text{MSBS} = \frac{1}{n-1} \sum_{i=1}^{n} \sum_{j=1}^{k} \left(\frac{1}{k} \sum_{j=1}^{k} x_{ij} - \frac{1}{nk} \sum_{i=1}^{n} \sum_{j=1}^{k} x_{ij} \right)^2, \tag{5.3.5}$$

$$\text{MSWS} = \frac{1}{n(k-1)} \sum_{i=1}^{n} \sum_{j=1}^{k} \left(x_{ij} - \frac{1}{k} \sum_{j=1}^{k} x_{ij} \right)^2. \tag{5.3.6}$$

To assess the reproducibility of the parameter estimates, the within-subject coefficient of variation (CV) and intra-class correlation coefficient (ICC) were calculated from the model parameters as

$$\text{CV} = \frac{\sigma_e}{\mu} \cdot 100\%, \tag{5.3.7}$$

$$\text{ICC} = \frac{\sigma_s^2}{\sigma_e^2 + \sigma_s^2}. \tag{5.3.8}$$

CV quantifies the extent of variability relative to the mean value of the population and thus enables the comparison of variability between two metrics that may have different mean values. CV was multiplied by 100% to follow the conventions in the field. ICC is a measure of how much of the total observed variance is explained by the variance of interest which in this case is the between-subjects variance. The definition of ICC used here is widely known as ICC(1, 1) according to the definitions by Shrout and Fleiss[189].

The confidence intervals of CV and ICC were estimated using the confidence intervals of the parameter estimates of a 1-way random effects model[187]. The estimated confidence interval of CV is given by

$$\left[\sqrt{\frac{\text{SSWS}}{\chi^2_{n(k-1),(1-\alpha/2)\cdot 100}}} \cdot \frac{100\%}{\mu}, \sqrt{\frac{\text{SSWS}}{\chi^2_{n(k-1),(\alpha/2)\cdot 100}}} \cdot \frac{100\%}{\mu}\right], \tag{5.3.9}$$

where $\text{SSWS} = \text{MSWS} \cdot n(k-1)$ is the sum of squares within-subjects, $\chi^2_{a,b}$ denotes the bth percentile of a Chi-squared distribution with a degrees of freedom, and $\alpha = 0.05$ is the significance level. The confidence intervals of ICC were estimated as

$$\left[\frac{F/F_U - 1}{k + F/F_U - 1}, \frac{F/F_L - 1}{k + F/F_L - 1}\right], \tag{5.3.10}$$

where $F = \text{MSBS}/\text{MSWS}$, and F_L and F_U denote the $[(\alpha/2) \cdot 100]$th and the $[(1-\alpha/2) \cdot 100]$th percentiles, respectively, of the F distribution with $n-1$ and $n(k-1)$ degrees of freedom.

5.4 Results

A representative example of the brain segmentation in one of the volunteers is shown in Figure 5.1. A visual inspection of the ROIs and the estimated within-subject CV of ROI volumes confirmed that the segmentation process met the criteria for the reproducibility analysis. The within-subject CV of ROI volumes ranged from 0.1% to 8.0% with the mean CV calculated over all subjects and ROIs being 2.1%. The greatest variability in volume was observed in white matter, for which the within-subject CV of averaged over all subjects was 2.6%, whereas the lowest variability in volume was observed in the corticospinal tract, for which the within-subject CV averaged over all subjects was 0.8%.

Representative axial slices of the parameter maps in one of the volunteers are shown in Figure 5.3. DKI results in highly noisy RK estimates in the voxels in the corpus callosum that contain tightly packed and aligned axons (Figure 5.3G). This issue also affects MK but is not visible in MSK maps (Figure 5.3H). The maps also reveal that NODDI and SMT2 result in strikingly different and contradictory estimates of intra-neurite volume fraction (Figures 5.3J and 5.3S). The Spearman correlation coefficient between NODDI's v_{ic} and SMT2's v_{intra} calculated over all voxels and all subjects was 0.51. As expected based on NODDI's model assumptions that connect v_{ic} to MD, the intra-neurite volume fraction map calculated using NODDI shows less contrast between white matter and gray matter than SMT2. Interestingly, the contrast in v_{intra} maps closely resembles that of MSK maps with the Spearman correlation coefficient calculated over all voxels and all subjects between v_{intra} and MSK being 0.97. Furthermore, SMT1 results in very high μFA estimates due to the model conflating all non-Gaussian diffusion with microscopic diffusion anisotropy (Figure 5.3O).

The point estimates and confidence intervals of the within-subject CV of each microstructural metric mean value in every ROI are shown in Figure 5.4 and reported in Table 5.1. The within-subject variability of each metric's mean value is dependent on the ROI in which it is calculated. Of the studied metrics, NODDI's v_{iso} suffered from the greatest within-subject variability. For the microstructural metrics quantifying microscopic diffusion anisotropy, averaged across all ROIs, the mean CV was 1.5% for NODDI's v_{ic}, 1.1% for SMT1's μFA, and 2.7% for SMT2's v_{intras}. These values are comparable to the average CV of the DTI metrics which ranged from 1.6% for AD to 2.9% for FA.

The point estimates and confidence intervals of the within-subject ICC of each microstructural metric mean value in every ROI are shown in Figure 5.5 and reported in Table 5.2. Some of the ICC values are incorrectly estimated to be less than or equal to zero which is a consequence of the modest amount of data acquired for the study and an expected result of the used statistical model[187]. The ICC of each metric's mean depends on the ROI in which it is calculated. For the microstructural metrics quantifying microscopic diffusion anisotropy, averaged across all ROIs, the mean ICC was 0.59 for NODDI's v_{ic}, 0.39 for SMT1's μFA, and 0.70 for SMT2's v_{intra}. For reference, the average ICC values of the DTI metrics ranged from 0.68 for FA to 0.81 for RD. The relatively low ICC of SMT1's μFA reveals that its low CV does not imply great specificity but rather low intra-subject variance.

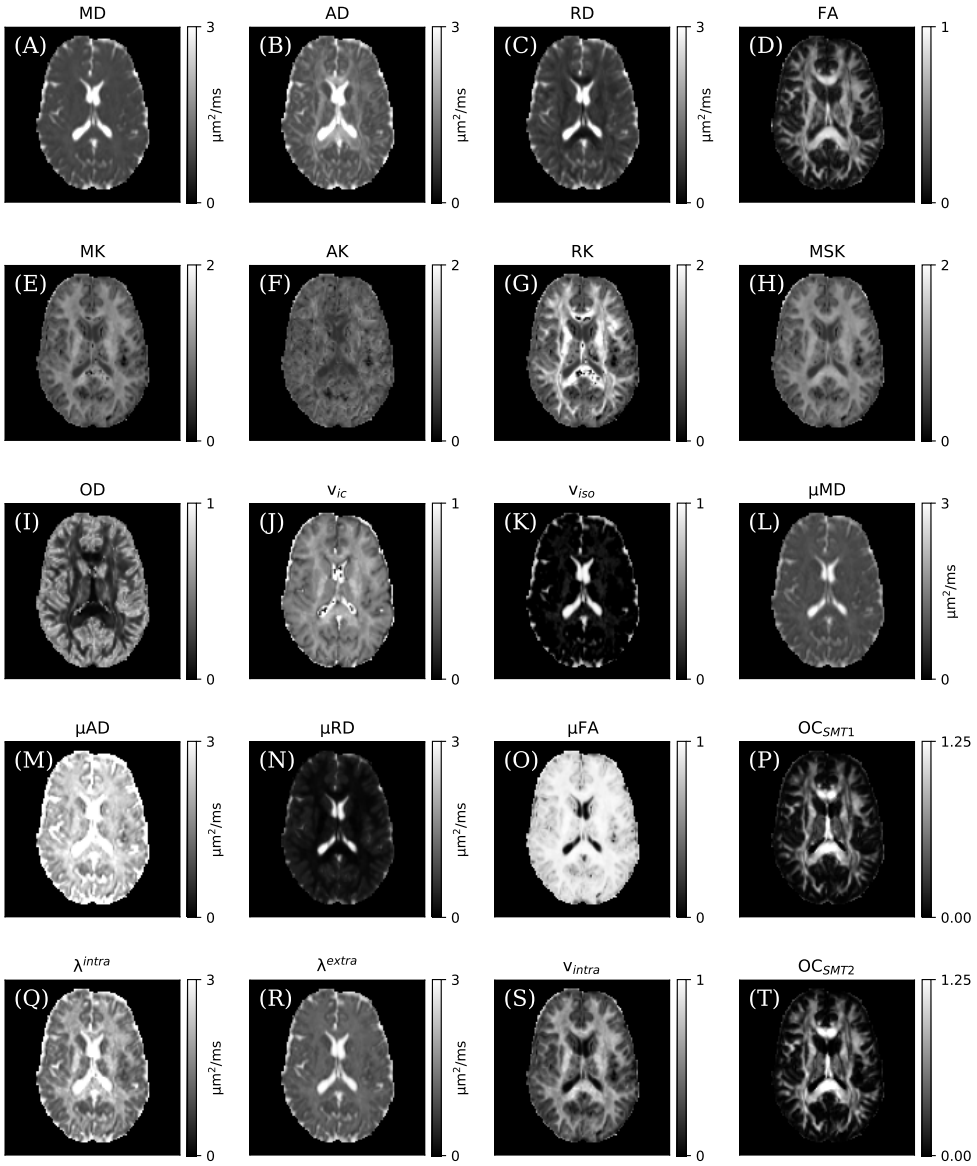

Figure 5.3: *Representative axial slices of the microstructural maps used in this study. (A)-(D) are from DTI. (E)-(H) are from DKI. (I)-(K) are from NODDI. (L)-(P) are from SMT1. (Q)-(T) are from SMT2.*

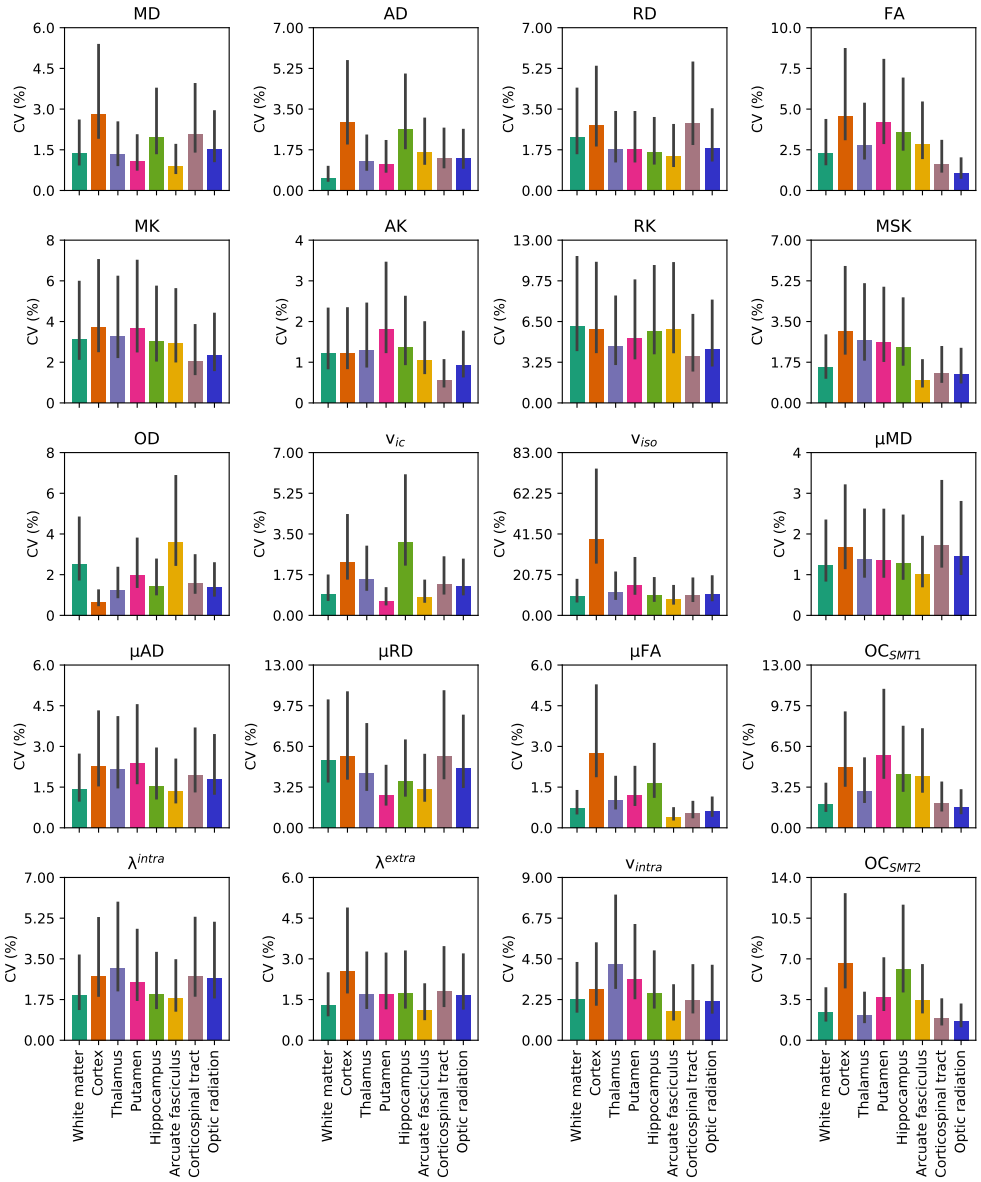

Figure 5.4: *The estimated within-subject coefficient of variation (CV) of the mean microstructural metrics in each ROI. The black lines denote the 95% confidence intervals.*

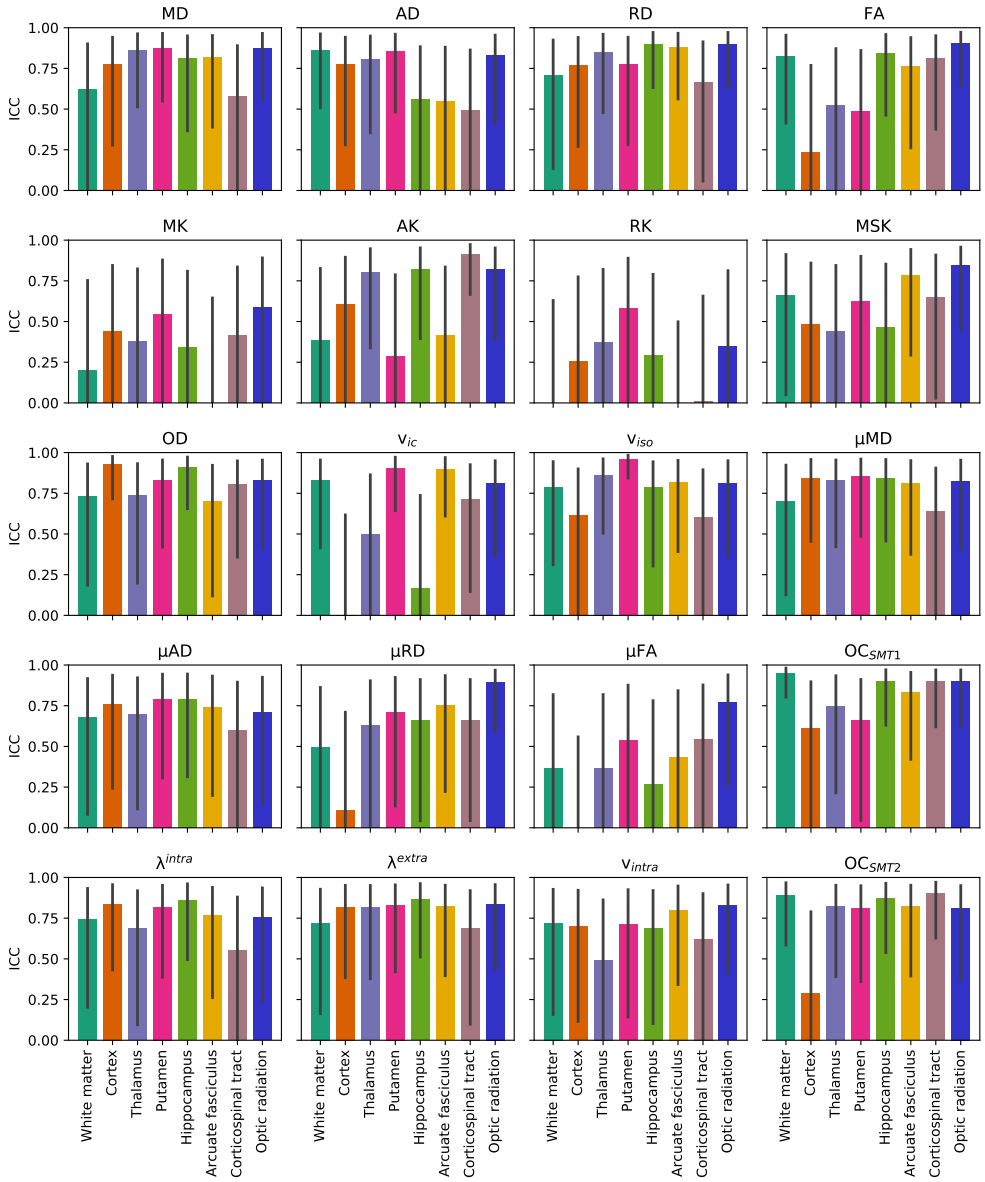

Figure 5.5: *The estimated intra-class correlation coefficients (ICC) of the mean microstructural metrics in each ROI. The black lines denote the 95% confidence intervals.*

Metric	White matter	Cortex	Thalamus	Putamen	Hippocampus	Arcuate fasciculus	Corticospinal tract	Optic radiation
MD	1.4 [0.9, 2.6]	2.8 [1.9, 5.4]	1.3 [0.9, 2.5]	1.1 [0.7, 2.1]	2.0 [1.3, 3.8]	0.9 [0.6, 1.7]	2.1 [1.4, 4.0]	1.5 [1.0, 3.0]
AD	0.6 [0.4, 1.1]	2.9 [2.0, 5.6]	1.3 [0.8, 2.4]	1.1 [0.8, 2.2]	2.6 [1.8, 5.0]	1.6 [1.1, 3.1]	1.4 [1.0, 2.7]	1.4 [0.9, 2.7]
RD	2.3 [1.6, 4.4]	2.8 [1.9, 5.4]	1.8 [1.2, 3.4]	1.8 [1.2, 3.4]	1.6 [1.1, 3.2]	1.5 [1.0, 2.9]	2.9 [2.0, 5.5]	1.8 [1.2, 3.5]
FA	2.3 [1.5, 4.4]	4.6 [3.1, 8.8]	2.8 [1.9, 5.4]	4.2 [2.9, 8.1]	3.6 [2.4, 6.9]	2.9 [1.9, 5.5]	1.6 [1.1, 3.1]	1.1 [0.7, 2.0]
MK	3.1 [2.1, 6.0]	3.7 [2.5, 7.1]	3.3 [2.2, 6.3]	3.7 [2.5, 7.0]	3.0 [2.0, 5.8]	2.9 [2.0, 5.6]	2.0 [1.4, 3.9]	2.3 [1.6, 4.4]
AK	1.2 [0.8, 2.3]	1.2 [0.8, 2.4]	1.3 [0.9, 2.5]	1.8 [1.2, 3.5]	1.4 [0.9, 2.6]	1.0 [0.7, 2.0]	0.6 [0.4, 1.1]	0.9 [0.6, 1.8]
RK	6.1 [4.1, 11.7]	5.9 [4.0, 11.3]	4.5 [3.0, 8.6]	5.1 [3.5, 9.9]	5.7 [3.9, 11.0]	5.9 [4.0, 11.2]	3.7 [2.5, 7.1]	4.3 [2.9, 8.3]
MSK	1.5 [1.0, 2.9]	3.1 [2.1, 5.9]	2.7 [1.8, 5.1]	2.6 [1.8, 5.0]	2.4 [1.6, 4.5]	1.0 [0.7, 1.9]	1.3 [0.9, 2.4]	1.2 [0.8, 2.4]
ODI	2.5 [1.7, 4.9]	0.7 [0.5, 1.3]	1.2 [0.8, 2.4]	2.0 [1.3, 3.8]	1.5 [1.0, 2.8]	3.6 [2.4, 6.9]	1.6 [1.1, 3.0]	1.4 [0.9, 2.6]
v_{ic}	0.9 [0.6, 1.8]	2.3 [1.5, 4.4]	1.6 [1.1, 3.0]	0.6 [0.4, 1.2]	3.2 [2.1, 6.1]	0.8 [0.5, 1.5]	1.3 [0.9, 2.5]	1.3 [0.9, 2.4]
v_{iso}	9.7 [6.6, 18.6]	39.1 [26.4, 74.8]	11.7 [7.9, 22.4]	15.5 [10.5, 29.8]	10.2 [6.9, 19.5]	8.1 [5.5, 15.5]	10.1 [6.8, 19.3]	10.7 [7.2, 20.4]
μMD	1.2 [0.8, 2.4]	1.7 [1.1, 3.2]	1.4 [0.9, 2.6]	1.4 [0.9, 2.6]	1.3 [0.9, 2.5]	1.0 [0.7, 2.0]	1.7 [1.2, 3.3]	1.5 [1.0, 2.8]
μAD	1.4 [1.0, 2.7]	2.3 [1.5, 4.3]	2.1 [1.5, 4.1]	2.4 [1.6, 4.6]	1.5 [1.0, 3.0]	1.3 [0.9, 2.6]	1.9 [1.3, 3.7]	1.8 [1.2, 3.5]
μRD	5.4 [3.6, 10.3]	5.7 [3.8, 10.9]	4.4 [2.9, 8.4]	2.6 [1.8, 5.0]	3.7 [2.5, 7.1]	3.1 [2.1, 5.9]	5.7 [3.9, 11.0]	4.7 [3.2, 9.0]
μFA	0.7 [0.5, 1.4]	2.8 [1.9, 5.3]	1.0 [0.7, 1.9]	1.2 [0.8, 2.3]	1.6 [1.1, 3.1]	0.4 [0.3, 0.8]	0.5 [0.4, 1.0]	0.6 [0.4, 1.2]
OC_{SMT1}	1.9 [1.3, 3.6]	4.9 [3.3, 9.3]	2.9 [2.0, 5.6]	5.8 [3.9, 11.1]	4.3 [2.9, 8.2]	4.2 [2.8, 8.0]	1.9 [1.3, 3.7]	1.6 [1.1, 3.1]
λ^{intra}	1.9 [1.3, 3.7]	2.8 [1.9, 5.3]	3.1 [2.1, 6.0]	2.5 [1.7, 4.8]	2.0 [1.3, 3.8]	1.8 [1.2, 3.5]	2.8 [1.9, 5.3]	2.7 [1.8, 5.1]
λ^{extra}	1.3 [0.9, 2.5]	2.6 [1.7, 4.9]	1.7 [1.2, 3.3]	1.7 [1.1, 3.2]	1.7 [1.2, 3.3]	1.1 [0.7, 2.1]	1.8 [1.2, 3.5]	1.7 [1.1, 3.2]
v_{int}	2.3 [1.5, 4.3]	2.8 [1.9, 5.4]	4.2 [2.8, 8.1]	3.4 [2.3, 6.4]	2.6 [1.8, 5.0]	1.6 [1.1, 3.1]	2.2 [1.5, 4.2]	2.2 [1.5, 4.2]
OC_{SMT2}	2.4 [1.6, 4.6]	6.6 [4.5, 12.7]	2.2 [1.5, 4.2]	3.7 [2.5, 7.1]	6.1 [4.1, 11.7]	3.4 [2.3, 6.6]	1.9 [1.3, 3.6]	1.6 [1.1, 3.2]

Table 5.1: *The estimated within-subject coefficient of variation (CV) of the mean microstructural metrics in each ROI. The values in the square brackets are the estimated 95% confidence interval.*

Metric	White matter	Cortex	Thalamus	Putamen	Hippocampus	Arcuate fasciculus	Corticospinal tract	Optic radiation
MD	0.62 [-0.03, 0.91]	0.77 [0.27, 0.95]	0.86 [0.51, 0.97]	0.88 [0.54, 0.97]	0.81 [0.36, 0.96]	0.82 [0.38, 0.96]	0.58 [-0.09, 0.90]	0.88 [0.54, 0.97]
AD	0.86 [0.50, 0.97]	0.78 [0.27, 0.95]	0.81 [0.35, 0.96]	0.85 [0.47, 0.97]	0.56 [-0.12, 0.89]	0.55 [-0.13, 0.89]	0.50 [-0.21, 0.87]	0.83 [0.41, 0.96]
RD	0.71 [0.13, 0.93]	0.77 [0.26, 0.95]	0.85 [0.47, 0.97]	0.78 [0.27, 0.95]	0.90 [0.62, 0.98]	0.88 [0.55, 0.97]	0.67 [0.05, 0.92]	0.90 [0.62, 0.98]
FA	0.83 [0.41, 0.96]	0.24 [-0.47, 0.78]	0.52 [-0.17, 0.88]	0.49 [-0.22, 0.87]	0.85 [0.45, 0.97]	0.77 [0.25, 0.95]	0.81 [0.37, 0.96]	0.91 [0.64, 0.98]
MK	0.20 [-0.50, 0.76]	0.44 [-0.28, 0.85]	0.38 [-0.34, 0.83]	0.55 [-0.14, 0.89]	0.34 [-0.38, 0.82]	-0.01 [-0.65, 0.65]	0.41 [-0.31, 0.84]	0.59 [-0.08, 0.90]
AK	0.39 [-0.33, 0.83]	0.60 [-0.06, 0.90]	0.80 [0.33, 0.96]	0.28 [-0.43, 0.80]	0.82 [0.39, 0.96]	0.41 [-0.31, 0.84]	0.91 [0.66, 0.98]	0.82 [0.38, 0.96]
RK	-0.04 [-0.66, 0.64]	0.25 [-0.46, 0.78]	0.37 [-0.35, 0.83]	0.58 [-0.09, 0.90]	0.29 [-0.43, 0.80]	-0.23 [-0.76, 0.51]	0.01 [-0.63, 0.66]	0.35 [-0.37, 0.82]
MSK	0.66 [0.04, 0.92]	0.48 [-0.22, 0.87]	0.44 [-0.28, 0.85]	0.62 [-0.03, 0.91]	0.46 [-0.25, 0.86]	0.78 [0.28, 0.95]	0.65 [0.02, 0.92]	0.84 [0.44, 0.97]
ODI	0.73 [0.18, 0.94]	0.93 [0.71, 0.98]	0.74 [0.19, 0.94]	0.83 [0.41, 0.96]	0.91 [0.65, 0.98]	0.70 [0.11, 0.93]	0.81 [0.35, 0.96]	0.83 [0.40, 0.96]
v_{ic}	0.83 [0.41, 0.96]	-0.06 [-0.67, 0.63]	0.50 [-0.21, 0.87]	0.91 [0.63, 0.98]	0.17 [-0.53, 0.75]	0.90 [0.60, 0.98]	0.71 [0.14, 0.93]	0.81 [0.36, 0.96]
v_{iso}	0.79 [0.30, 0.95]	0.62 [-0.03, 0.91]	0.86 [0.50, 0.97]	0.96 [0.84, 0.99]	0.79 [0.29, 0.95]	0.82 [0.38, 0.96]	0.60 [-0.06, 0.90]	0.81 [0.36, 0.96]
μMD	0.70 [0.12, 0.93]	0.84 [0.45, 0.97]	0.83 [0.41, 0.96]	0.85 [0.48, 0.97]	0.84 [0.45, 0.97]	0.81 [0.37, 0.96]	0.64 [-0.00, 0.91]	0.83 [0.40, 0.96]
μAD	0.68 [0.07, 0.93]	0.76 [0.23, 0.95]	0.70 [0.11, 0.93]	0.79 [0.30, 0.95]	0.79 [0.31, 0.95]	0.74 [0.19, 0.94]	0.60 [-0.06, 0.90]	0.71 [0.13, 0.93]
μRD	0.49 [-0.21, 0.87]	0.11 [-0.57, 0.72]	0.63 [-0.01, 0.91]	0.71 [0.13, 0.93]	0.66 [0.03, 0.92]	0.75 [0.21, 0.94]	0.66 [0.04, 0.92]	0.89 [0.59, 0.98]
μFA	0.37 [-0.36, 0.83]	-0.15 [-0.72, 0.57]	0.37 [-0.36, 0.83]	0.54 [-0.15, 0.88]	0.27 [-0.45, 0.79]	0.43 [-0.28, 0.85]	0.54 [-0.14, 0.89]	0.77 [0.25, 0.95]
OC_{SMT1}	0.95 [0.79, 0.99]	0.61 [-0.05, 0.91]	0.75 [0.21, 0.94]	0.66 [0.04, 0.92]	0.90 [0.62, 0.98]	0.83 [0.41, 0.96]	0.90 [0.61, 0.98]	0.90 [0.61, 0.98]
λ^{intra}	0.74 [0.19, 0.94]	0.84 [0.42, 0.96]	0.69 [0.09, 0.93]	0.82 [0.38, 0.96]	0.86 [0.49, 0.97]	0.77 [0.25, 0.95]	0.55 [-0.14, 0.89]	0.76 [0.23, 0.94]
λ^{extra}	0.72 [0.16, 0.94]	0.82 [0.38, 0.96]	0.82 [0.37, 0.96]	0.83 [0.41, 0.96]	0.86 [0.50, 0.97]	0.82 [0.39, 0.96]	0.69 [0.09, 0.93]	0.84 [0.43, 0.96]
v_{int}	0.72 [0.15, 0.94]	0.70 [0.11, 0.93]	0.49 [-0.21, 0.87]	0.71 [0.13, 0.93]	0.69 [0.09, 0.93]	0.80 [0.33, 0.96]	0.62 [-0.03, 0.91]	0.83 [0.40, 0.96]
OC_{SMT2}	0.89 [0.58, 0.98]	0.29 [-0.43, 0.80]	0.82 [0.38, 0.96]	0.81 [0.35, 0.96]	0.87 [0.53, 0.97]	0.82 [0.39, 0.96]	0.90 [0.62, 0.98]	0.81 [0.36, 0.96]

Table 5.2: *The estimated intra-class correlation coefficients (ICC) of the mean microstructural metrics in each ROI. The values in the square brackets are the estimated 95% confidence interval.*

5.5 Discussion

In this chapter, microscopic diffusion anisotropy estimation using three microstructural models was discussed and the reproducibility of the model parameter estimates was quantified. The reproducibility of the parameter estimates derived from DTI and DKI were provided for reference. The analysis focused on the mean values in several scientifically and clinically relevant ROIs which included all of white matter and the cerebral cortex as well as individual subcortical gray matter structures and major white matter tracts. Estimates of within-subject variability were provided and they were related to between-subjects variability using ICC.

The reproducibility of the kurtosis measures was observed to be lower than the diffusivity measures, which was expected as the second order effects in terms of the applied b-value are more susceptible to noise than the first order effects. These results are in agreement with previous reports in the literature. For instance, in a study by Palacios et al.[190], the within-subject CV estimates of DTI metrics averaged over large ROIs were found to be consistently below 5%, and in a study by Kasa et al.[191], the kurtosis measures were found to be less reproducible than the diffusitivity measures. RK's high sensitivity to noise in voxels containing tightly packed and aligned axons is a consequence of low radial diffusivity and has been observed before. It has lead to some authors to recommend smoothing the diffusion-weighted data prior to the DKI fit or using MSK for quantifying kurtosis instead[46,192]. Indeed, similar problems were not observed in the MSK maps. NODDI's and SMT's metrics measuring microscopic diffusion anisotropy did not suffer from this issue either.

When DTI and DKI were used as a reference, the reproducibility of the metrics quantifying microscopic diffusion anisotropy was found to be good. Of the studied methods, SMT1 produced the most reproducible measure of microscopic diffusion anisotropy when measured using the within-subject CV only. However, when the within-subjects variability was related to between-subjects variability using ICC, it was found that SMT2's v_{intra} was the most reproducible measure of microscopic diffusion anisotropy and SMT1's μFA the least reproducible. SMT1's μFA's low within-subjects variability is due to the model conflating all non-Gaussian diffusion with microscopic diffusion anisotropy[133], resulting in μFA estimates that are consistently very close to 1. Based on this analysis, SMT2's v_{intra} is expected to be the most sensitive metric to microstructural differences between groups of over time.

Although the focus of this work was on the estimation of microscopic diffusion anisotropy, the analysis was also repeated for other model parameters. Overall, the reproducibility of NODDI parameters was found to be comparable to DTI with the exception of v_{iso}. These results are in good agreement with previous studies that have reported precision issues with NODDI's v_{iso} [193,194]. The overall reproducibility of SMT1 and SMT2 was found to be comparable to DTI and DKI parameters across all the model parameters.

This study is limited by the modest number of subjects and scans which resulted in some negative point estimates of ICC which are clearly meaningless. However, given the number of microstructural metrics and ROIs included in the study, this was an expected property of the applied statistical model [187]. The provided confidence intervals facilitate the interpretation of the results. The presented results may also change if the ROIs were defined differently, such as using a different tractography pipeline or parameter choices for binary mask generation. Furthermore, the number of volunteers was sufficient for the purpose of this study which was to relate the reproducibility of NODDI and SMT to that of DTI and DKI. It is also important to mention that the reproducibility of the studied methods also depends on the scanner hardware and sequence parameters [193]. Overall, the results presented in this chapter support using NODDI and SMT in situations where they offer the greatest sensitivity to microstuctural differences across groups or over time. However, due to the invalid model assumptions, the parameter estimates are biased and thus it is crucial not to take the parameter names literally and to interpret them with caution.

5.6 Conclusion

The results presented in this chapter show that despite the accuracy issues of microscopic diffusion anisotropy estimation using SDE acquisitions and microstructural modelling, the reproducibility of most NODDI and SMT parameter estimates is comparable to those derived from DKI and DTI, possibly making them valuable in research and clinical settings. However, the accuracy issues discussed in the theory section of this chapter must be taken into account when interpreting the results. The following chapters discuss more accurate methods for estimating microscopic diffusion anisotropy using MDE.

Chapter 6

Validation and noise robustness of clinically feasible microscopic anisotropy estimation using double diffusion encoding

6.1 Introduction

Despite being sensitive to several clinically and scientifically relevant microstructural changes in the brain[112], SDE acquisitions are fundamentally limited because they confound two major sources of voxel-level non-Gaussian diffusion, namely, the orientation dispersion of anisotropic diffusion and the size variance of microscopic diffusion environments, resulting in a lack of specificity. As explained in Section 3.7, this fundamental degeneracy of SDE acquisitions can not be resolved without substantial *a priori* information about tissue microstructure. Over recent years, MDE methods, such as DDE, have gained significant attention for their ability to disentangle the two sources of voxel-level non-Gaussian diffusion, allowing microscopic diffusion anisotropy to be estimated in a model-free way[54,134,135,64,55–61,63].

DDE sequences contain two trapezoidal diffusion encoding pulse pairs separated by a mixing time (τ) as shown in Figure 6.1. The two periods of diffusion encoding allow the correlations of spin displacements along different directions to be studied. In the context of estimating microstructural properties, DDE was first proposed by Cory et al. in 1990 for measuring the anisotropy of randomly oriented yeast cells[195]. It was soon shown theoretically that in

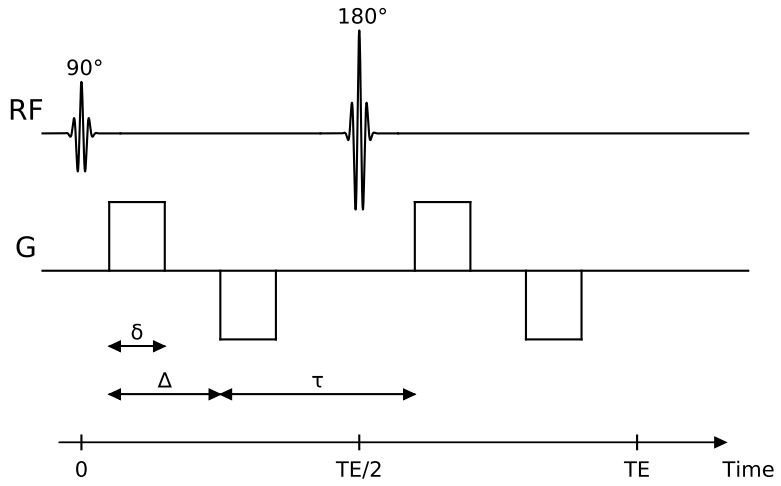

Figure 6.1: *A schematic representation of a DDE gradient waveform where two trapezoidal gradient pulse pairs of equal magnitude and duration are separated by a mixing time τ and a 180° RF pulse.*

angular DDE experiments, i.e., experiments in which the angle between the two gradient pulse pairs is varied, the amplitude of the signal modulation as a function of the angle between the pulse pair directions depends on μA (Equation 3.7.5) in powder-averaged systems if the mixing time is long enough for the positions of the spins during the two gradient pulse pairs to be uncorrelated[196]. In this context, a powder-averaged system refers to a sample that contains a very large number of similar structures with random orientations and that would therefore appear isotropic in DTI experiments. Following these findings, several studies then applied DDE to measure μA in liquid crystals[197], porous media[198], and *ex vivo* animal neural tissue[199], for example.

Despite its potential, applying DDE in human neuroimaging has been challenging, and two major issues have most likely contributed to its slow adoption. First, the theory behind μA estimation using DDE assumes the mixing time to be long enough that the spins have time to completely sample their microscopic environment[200]. Considering that DDE sequences consists of two diffusion times and a mixing time, this requirement invariably leads to long echo times and a diminished signal. Second, neural tissue is not a powder-averaged system and therefore an orientationally invariant acquisition scheme is required[201]. Along with poor vendor availability and long echo times, the large number of acquisitions required for obtaining an orientationally invariant signal has likely impeded DDE from becoming more common in human

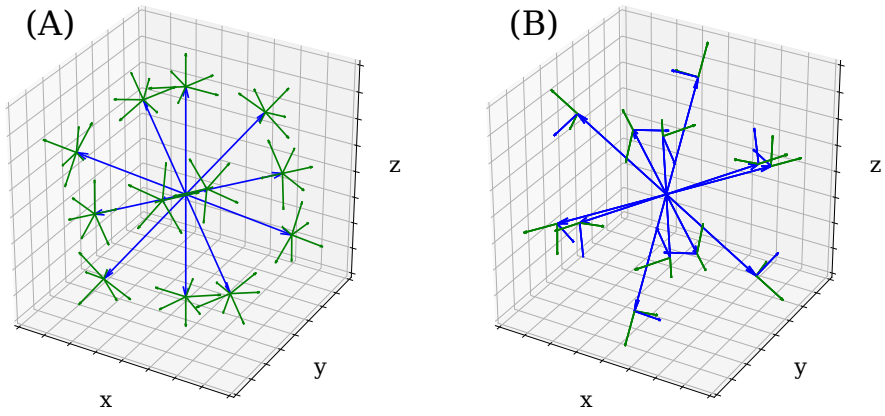

Figure 6.2: *Diffusion encoding directions of the 5-design (A) and a symmetric version of the minimal design (B). The directions of the first gradient pulse pair are shown in blue and the directions of the second gradient pulse pair are shown in green. The directions of the first pulse pair pointing towards the vertices of the icosahedron are represented by longer arrows.*

neuroimaging.

In 2013, Jespersen et al. introduced a theoretically justified orientationally invariant acquisition scheme termed the 5-design[134] which is a set of 72 DDE wave vector rotations that can be applied to estimate the powder-averaged signal, i.e., the signal averaged over all possible diffusion encoding directions, which is orientationally invariant. In 2017, Yang et al. proposed that, if diffusion in the microscopic diffusion environments appears Gaussian, as few as 12 acquisitions could be sufficient for an orientationally invariant acquisition scheme to map μA[56]. This set of directions is hereafter defined as the minimal design. Both the 5-design and the minimal design are illustrated in Figure 6.2. This simplification is based on the idea that, if diffusion in every microscopic diffusion environment can be fully characterized by a microscopic diffusion tensor, some acquisitions with the orthogonal gradient pulse pair directions in the 5-design are redundant as they correspond to the same in-plane trace of the diffusion tensor. Additionally, asymmetric sampling of diffusion-weighting directions covering half a sphere is sufficient in the absence of bulk flow. This realization allows the 5-design to be simplified from 12 + 60 acquisitions to 6 + 15 acquisitions for parallel and orthogonal gradient pulse pairs, respectively. Using simulations of infinite cylinders with orientation dispersion and *in vivo* imaging experiments using a clinical human whole-body scanner, Yang

et al. experimentally showed that the design can be further reduced to 6 + 6 without significant differences in the observed value of μA. Importantly, such a dramatic reduction in directions, and hence in acquisition time, has strong implications for the application of DDE in human neuroimaging.

However, the new minimal design has not been directly compared with the more theoretically robust 5-design. Clearly, validation is important because new μA estimation methods have been shown to produce biased results[132,133], and bias due to an insufficient number of directions in estimating the powder-averaged signal has been reported[202]. This study sought to experimentally validate the proposed approach by directly comparing it to the 5-design over a range of relevant b-values in animal imaging experiments. Additionally, simulation experiments were performed to assess the precision of μFA, a normalized measure of μA, estimated from DDE data, a problem which has received little attention thus far.

6.2 Theory

This section provides a concise summary of the theory behind μA estimation using DDE. A detailed description and derivation of the presented equations is provided by Jespersen et al.[203,200,134].

Consider a DDE sequence with two diffusion encoding gradient pulse pairs with the same pulse length (δ), diffusion time (Δ), and gradient magnitude separated by a mixing time (τ) (Figure 6.1). Under the short pulse assumption ($\delta \to 0$), the gradient pulse pairs can be represented by points in q-space: $\mathbf{q}_1$ and $\mathbf{q}_2$. Breaking the short pulse assumption has the same effect as in regular PGSE experiments, i.e., the acquisition is sensitive to the differences in the water molecules' average positions during the applied pulses[111,200]. Using the cumulant expansion[124] up to the fourth order in terms of the applied wave vector magnitude, the natural logarithm of the normalized DDE signal from a powder-averaged system can be expressed as[200]

$$\begin{aligned}\ln(E) \approx & -q^2\left[2\Delta(\mathrm{MD}) - Q\cos(\theta)\right] + \frac{1}{3}q^4\Delta^2(\mathrm{MD})^2 K_{zzzz} \\ & + \frac{1}{4}q^4\left[\cos^2(\theta)\left(Z_{zzzz} - Z_{zzxx}\right) + Z_{zzxx}\right] \\ & - \frac{1}{3}q^4\cos(\theta)S_{zzzz},\end{aligned} \tag{6.2.1}$$

where $q = |\mathbf{q}_1| = |\mathbf{q}_2|$, MD is the mean diffusivity, Q is the trace of the rank-2 correlation tensor $\mathbf{Q}$[203], θ is the angle between $\mathbf{q}_1$ and $\mathbf{q}_2$, K_{zzzz} is an element

of the kurtosis tensor $\mathbf{K}$, and Z_{zzzz}, Z_{zzxx} and S_{zzzz} are elements of the rank-4 displacement correlation tensors $\mathbf{Z}$ and $\mathbf{S}$.

In the long mixing time limit ($\tau \rightarrow \infty$), the displacements of the spins during the two gradient pulse pairs become uncorrelated and the elements of the tensors $\mathbf{Q}$ and $\mathbf{S}$ tend towards zero[200]. In the long diffusion time limit ($\Delta \rightarrow \infty$), if the imaged microstructure can be characterized by a set of non-exchanging microscopic diffusion environments, referred to as compartments hereafter, the tensor $\mathbf{Z}$ becomes a measure of the heterogeneity of the average compartment shape and orientation[200]. Under these assumptions, the above expression of the powder-averaged signal can be reduced to[134]

$$\begin{aligned} \ln(E) \approx = & -q^2 2\Delta(\mathrm{MD}) + \frac{1}{3}q^4\Delta^2(\mathrm{MD})^2 K_{zzzz} \\ & + \frac{1}{4}q^4\left[\cos^2(\theta)\left(Z_{zzzz} - Z_{zzxx}\right) + Z_{zzxx}\right], \end{aligned} \tag{6.2.2}$$

where the amplitude of the $q^4\cos^2(\theta)$ term is proportional to μA^2. Since the other terms do not depend on θ, μA can be directly measured by comparing powder-averaged DDE signals acquired with two different angles between the gradient pulse pair directions.

In systems with an arbitrary orientation distribution, obtaining a powder-averaged signal would entail sampling over all possible orientations of the gradient pulse pair directions:

$$S^{\mathrm{PA}} = \frac{1}{8\pi^2}\int_{\mathrm{SO(3)}} S(\mathbf{L}q\hat{\mathbf{e}}_1, \mathbf{L}q\hat{\mathbf{e}}_2)d\mathbf{L}, \tag{6.2.3}$$

where the superscript PA stands for the powder-averaged signal, S is the measured signal, q is the magnitude of the wave vector, θ is the angle between the pulse pair directions represented by unit vectors $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2$, $\mathbf{L}$ represents a rotation, and SO(3) is the rotation group. Jespersen et al. showed using quadrature theory that, if S is equal to the signal cumulant expansion up to the fifth order in q, the integral in Equation 6.2.3 can be calculated exactly with a finite number of rotations and the powder-averaged signal can be estimated as[134]

$$S^{\mathrm{PA}} \approx \frac{1}{|\chi|}\sum_{\mathbf{L}\in\chi} S(\mathbf{L}q\hat{\mathbf{e}}_1, \mathbf{L}q\hat{\mathbf{e}}_2)d\mathbf{L}, \tag{6.2.4}$$

where the set of wave vector rotations χ is known as the 5-design consisting of 12 rotations for parallel wave vectors and 60 rotations for non-parallel (typically orthogonal) wave vectors.

Given the powder-averaged data acquired with parallel and orthogonal wave vectors, μA (Equation 3.7.5) can be estimated as

$$\mu\mathrm{A} = \sqrt{\ln\left(\frac{S_{\parallel}^{\mathrm{PA}}}{S_{\perp}^{\mathrm{PA}}}\right) b^{-2}}, \tag{6.2.5}$$

where the subscripts $\parallel$ and $\perp$ denote acquisitions with parallel and orthogonal wave vectors, respectively, and b is the b-value of one of the gradient pulse pairs of the DDE sequence[134]. In realistic voxels, a distribution of microscopic diffusion environments will invariably exist and thus the estimated μA will correspond to the average over the microenvironments in the voxel[134]. Since μA depends on both the shape and size of the microscopic diffusion tensors, μFA is often used instead to quantify microscopic diffusion anisotropy irrespective of the magnitude of diffusion (Equation 3.7.6). μFA can be approximately calculated from μA as

$$\mu\mathrm{FA} = \sqrt{\frac{2}{3}}\sqrt{\frac{\mu\mathrm{A}^2}{\mu\mathrm{A}^2 + \frac{3}{5}(\mathrm{MD})^2}}, \tag{6.2.6}$$

where MD is the mean diffusivity[135,56].

Because the formalism presented above is based on the truncation of the signal cumulant expansion, μA estimates can be corrupted by non-vanishing higher order terms. Introducing the following higher order correction to Equation 6.2.5 has been shown to improve the accuracy of the estimates[204]:

$$\mu\mathrm{A} = \sqrt{\ln\left(\frac{S_{\parallel}^{\mathrm{PA}}}{S_{\perp}^{\mathrm{PA}}}\right) b^{-2} - P_3 b^3}, \tag{6.2.7}$$

where P_3 is non-negative. Furthermore, in terms of precision, since μA estimation is based on observing the deviation between two signal decay curves, precise quantification of μA becomes more difficult as the signal deviation becomes smaller.

6.3 Methods

Imaging experiments

The accuracy of the proposed minimal design in neural tissue was validated by performing three separate animal imaging experiments that were carried out

by the PhD candidate and Rafael Neto Henriques at the Champalimaud Centre for the Unknown, Lisbon, Portugal, according to the European Union Directive 2010/63. Competent institutional and national authorities preapproved the experiments.

Specimen preparation

For the *ex vivo* experiments, a rat brain was extracted through standard transcardial perfusion from a healthy adult animal and was then immersed in 4% paraformaldehyde solution for 24 h, followed by immersion in a phosphate-buffered saline solution for at least 48 h. The extracted brain was inserted into a Fluorinert-filled (Sigma Aldrich, Lisbon, Portugal) 10 mm NMR tube and placed in the scanner at 23° C.

In the *in vivo* experiments, a healthy rat (age = 135 days, weight = 288 g) was anesthetized with isoflurane (4% for induction, 2% maintenance delivered through a nose cone) and placed in the scanner. Respiration rate and rectal temperature were monitored and kept stable over the entire experiment by slightly adjusting isoflurane levels and circulating warm water, respectively.

Data acquisition

The experiments were performed on a Bruker Biospec 9.4T (Bruker Corporation, Karlsruhe, Germany) harnessing an 86 mm volume coil for transmission and 4-element array cryocoil for reception. A DDE EPI pulse sequence written in-house was used with δ = 5 ms and $\Delta = \tau = 15$ ms. Other relevant imaging parameters were: TE = 69 ms in Experiment 1, TE = 65 ms in experiments 2 and 3, voxel size = 0.2 $\times$ 0.2 $\times$ 0.8 mm^3, FOV = 20 $\times$ 20 mm^2, partial Fourier = 4/5, slice gap = 0.5 mm, 3 slices in Experiment 1, and 5 slices in experiments 2 and 3. The diffusion encoding directions are shown in Figure 6.2. Because acquisitions were repeated several times, rather than using six parallel and six orthogonal gradient pulse pair directions distributed over half a sphere as proposed by Yang et al.[56], 12 + 12 directions distributed over the surface of the sphere were used to eliminate the possibility of artefacts arising from cross-terms between the imaging and diffusion encoding gradients[205]. The 5-design was acquired with 12 + 60 directions as usual. The b-values reported here refer to the total diffusion-weighting during the DDE experiment. For every 5-design acquisition, 13 images without diffusion-weighting were acquired for SNR estimation.

In Experiment 1, the *ex vivo* rat brain was scanned with both the minimal design and the 5-design using three b-values in the range typically used in clinical settings: 1, 2, and 3 ms/μm^2. To maximize SNR, the 5-design measurements were averaged over 30 acquisitions and the minimal design measurements over 90 acquisitions to account for the different number of directions. Additionally, the 5-design experiment was repeated twice to measure the inherent variance in μFA due to finite SNR.

In Experiment 2, the b-value dependency of the comparison was investigated by imaging the *ex vivo* brain with ten b-values uniformly distributed between 0.5 and 5 ms/μm^2. The 5-design measurements were averaged over eight acquisitions and the minimal design measurements over 24 acquisitions, accounting for the different number of directions.

In Experiment 3, the comparison of the two methods was repeated *in vivo* to avoid bias due to the fixation process and temperature. The experiment was performed with b-values 1 and 2.5 ms/μm^2. The 5-design measurements were averaged over eight acquisitions and the minimal design measurements over 24 acquisitions, accounting for the different number of directions.

Data preprocessing

The data was denoised using Marchenko-Pastur random matrix denoising[178]. Gibbs ringing artifacts were reduced using a sub-voxel shift algorithm[179]. To correct for motion, the images registered to the first non–diffusion-weighted image using a discrete sub-voxel Fourier transform algorithm[206]. No procedure for eddy current induced artefact correction was applied because no eddy current distortions could be visually observed in any of the images.

Parameter estimation

μFA maps were calculated from the powder-averaged data without the higher order correction (Equation 6.2.7) using equations 6.2.5 and 6.2.6. MD used in calculating μFA was estimated by fitting a diffusion tensor to the data acquired with parallel wave vectors at $b = 1$ ms/μm^2 using DiPy[184]. Equation 6.2.5 was fit to multi-shell data (Experiments 1 and 3) using a non-linear least squares trust region reflective algorithm[207] in Scipy[164]. Voxel-specific SNR was quantified as the mean signal divided by the standard deviation in signal over the images with no diffusion-weighting, and the average SNR was obtained by averaging voxel-specific SNR values over the brain volume.

Simulations

The orientational variance and noise robustness of μFA estimation using the minimal design was studied by generating synthetic data from axially symmetric diffusion tensors of various sizes and shapes. The normalized signal was simulated using the following equation:

$$E = \exp\left(-\mathbf{b} : \mathbf{D}\right), \tag{6.3.1}$$

where $\mathbf{b}$ is the b-tensor used in the acquisition, $\mathbf{D}$ is the diffusion tensor, and : denotes the generalized scalar product between the two tensors[115]. The DDE sequence with parallel and orthogonal wave vectors encode linear (Equation 3.7.2) and planar (Equation 3.7.3) b-tensors, respectively. Gaussian diffusion was simulated instead of restricted diffusion because the minimal design is based on the assumption that diffusion can be approximated as Gaussian in every microscopic diffusion environment. μFA was calculated from the simulated single-shell DDE data according to equations 6.2.5 and 6.2.6 with MD being equal to its ground-truth value.

First, the rotational variance of the 5-design and the minimal design was quantified by generating signals from single diffusion tensors with 100 unique orientations uniformly distributed over half a sphere. The standard deviation of the resulting μFA estimates was calculated to quantify the rotational variance of the two methods for a given diffusion tensor. The simulated b-value was 2.25 ms/μm^2. The ground-truth μFA was varied from 0 to 1 and MD was varied from 0.1 to 3 μm^2/ms. Fifty uniformly distributed values for both μFA and MD were used. Imaginary values of μFA caused by $S_{\perp}^{\mathrm{PA}}$ being greater than $S_{\parallel}^{\mathrm{PA}}$, which may occur as a consequence of noise or an inaccurately estimated powder-averaged signal (Equation 6.2.3), were mapped to 0.

Second, the noise robustness of the minimal design was quantified by adding random noise to the synthetic data and quantifying its effect on the estimated μFA. Noisy signals were generated as

$$E_{\mathrm{noisy}} = \sqrt{\left(E + X\right)^2 + Y^2}, \tag{6.3.2}$$

where E is the synthetic signal without noise and X and Y are randomly sampled from a normal distribution with zero mean and standard deviation σ. Here, SNR refers to the signal without diffusion-weighting divided by σ. Single tensors aligned with the x-axis were used. The ground-truth μFA was varied from 0 to 1, and MD was varied from 0.1 to 3 μm^2/ms. Fifty uniformly distributed

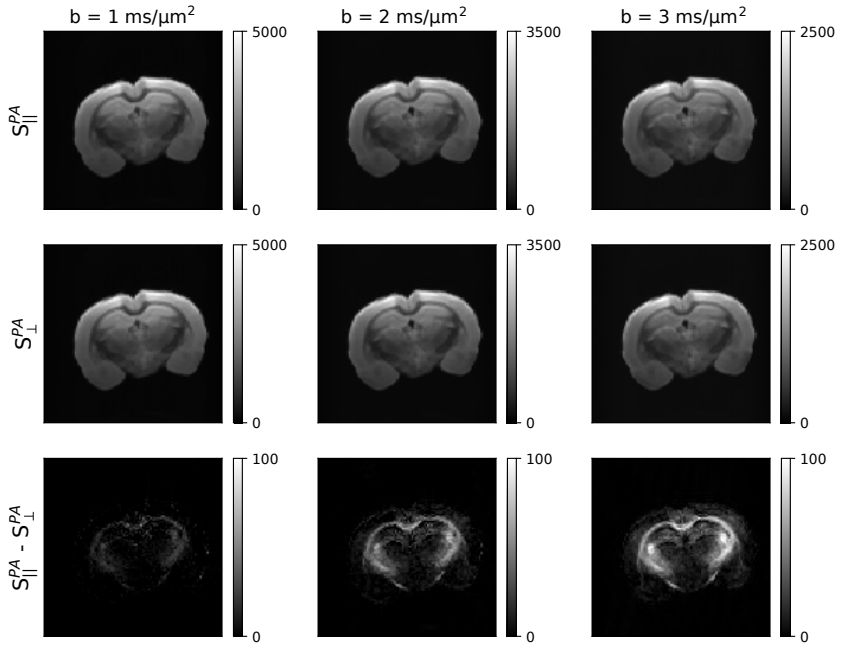

Figure 6.3: *Representative coronal slices of the powder-averaged data acquired with parallel (top row) and orthogonal wave vectors (middle row) with three b-values. μA can be clearly observed by subtracting the powder-averaged data acquired with orthogonal wave vectors from the powder-averaged data acquired with parallel wave vectors (bottom row).*

values for both parameters were used. The simulations were performed using b-values of 1.5, 2.25, and 3 ms/μm^2. The minimum SNR required for measuring μFA within 0.1 from its asymptotic value with 95% confidence was estimated by repeating the noisy simulations 10^3 times over the whole parameter space at 100 levels of SNR uniformly distributed between 1 and 10^3. The asymptotic value of the μFA estimate refers to the μFA value calculated from the simulated single-shell data with infinite SNR without the higher order correction, which leads to an underestimation of μFA that increases with an increasing b-value [204]. Imaginary values of μFA were not mapped to 0.

6.4 Results

Imaging experiments

The results of the first *ex vivo* imaging experiment are illustrated in figures 6.3 and 6.4. Figure 6.3 shows the powder-averaged data for parallel and orthogonal wave vectors from a representative slice and at different b-values. Average SNR was 151. The subtraction between the powder-averaged signals, dependent on μA, shows a b-value dependent pattern increasing in intensity with higher b-values, especially in white matter.

The μFA maps derived using the 5-design and the minimal design are shown in figures 6.4A and 6.4B. Qualitatively, the maps are nearly indistinguishable, and when the maps are subtracted, no clear spatial pattern is observed (Figure 6.4C). When the μFA derived from each method were plotted against each other voxel-wise (Figure 6.4D), a strong and statistically significant correlation was observed (Pearson's R = 0.91, $p < 10^{-3}$) with the data points very close to the unity line. More importantly, only a small bias was observed as the distribution of voxel-wise differences was centered near zero at 0.014 with a standard deviation of 0.07. Figure 6.4E shows the voxel-wise differences plotted against their mean value, revealing larger differences at lower values of μFA, as expected based on the discussion in the theory section.

To test whether the small variance observed was due to the different methods or due to noise, a simple test-retest experiment was performed using the 5-design (Figures 6.4F and 6.4G). The difference image again shows no particular spatial pattern (Figure 6.4H), and when the test-retest μFA values were plotted against each other voxel-wise, a very similar pattern was observed than with the comparison of the minimal design against the 5-design (Figure 6.4I). Indeed, the variance in this test-retest experiment was also similar to that observed before, with the distribution of voxel-wise differences between test and retest exhibiting a mean of 0.004 and a standard deviation of 0.06. The Pearson's correlation coefficient was 0.93 for the μFA maps acquired with 2 repetitions of the 5-design ($p < 10^{-3}$), and the data points closely followed the unity line. The voxel-wise differences are plotted against their mean values in Figure 6.4J, again revealing larger differences at lower values of μFA and showing that the precision of μFA decreases with decreasing μA.

In Experiment 2, the methods were compared over a larger interval of b-values which was more densely sampled to see if the results deviate at higher

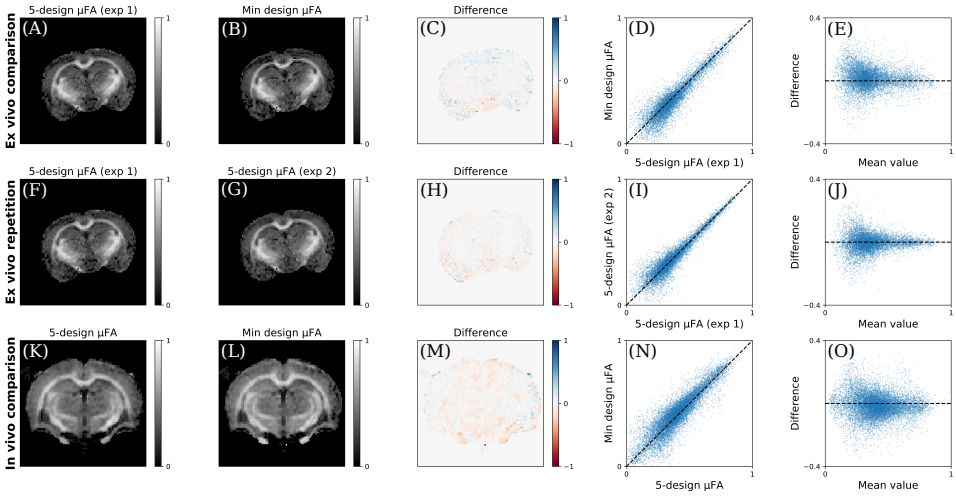

Figure 6.4: *Comparison of μFA maps calculated with the 5-design and the minimal design. The top row shows a comparison between the 5-design and the minimal design (Experiment 1). The middle row shows the same comparison between 2 repetitions of the 5-design experiment (Experiment 1). Bottom row shows the results of the in vivo experiment (Experiment 3). Difference stands for the map obtained by reducing the map in the second column from the map in the first column. The fourth and fifth columns show the voxel-wise comparison of the maps. The dashed lines in the fourth column are identity lines. The dashed lines in the fifth column denote no difference.*

b-values due to differences in how non-Gaussian diffusion affects the estimated powder-averaged signal acquired with the two methods. Results shown in Figure 6.5 suggest that the accuracy of the minimal design is comparable to the 5-design with b-values up to 5 ms/μm^2. Figure 6.5 also shows that increasing diffusion-weighting leads to underestimation of μFA, because the maps were calculated without a higher order correction.

Finally, in Experiment 3, the results of the *ex vivo* comparison were repeated *in vivo* to confirm that the *ex vivo* experiments are representative of *in vivo* conditions. The results are shown in Figure 6.4K-O. The very similar μFA maps extracted from data acquired with the 5-design and the minimal design are shown in figures 6.4K and 6.4L. The μFA values extracted from the 2 methods are highly correlated (Pearson's R = 0.86, $p < 10^{-3}$) with the data points following the unity line. The voxel-wise difference distribution's mean is equal to 0.012 with a standard deviation of 0.10.

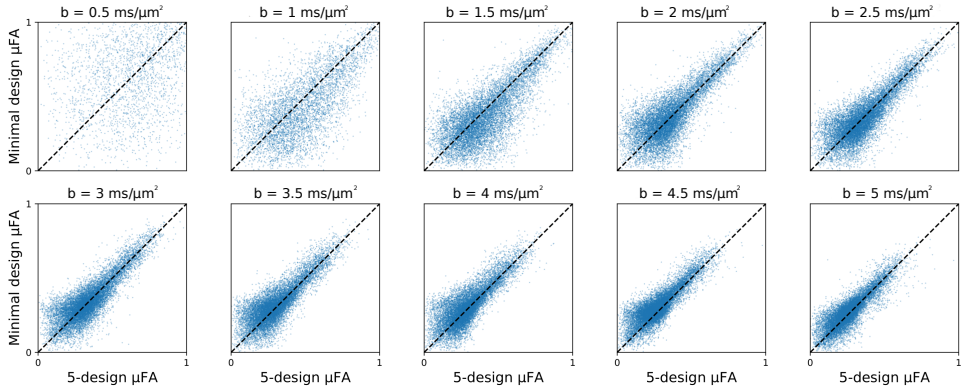

Figure 6.5: *Voxel-wise comparison of μFA maps calculated from single-shell data (Experiment 2) acquired with the 5-design and the minimal design over a range of b-values. The dashed lines are identity lines.*

Simulations

The standard deviations of the μFA estimates calculated from the simulated signals with different tensor orientations are shown in Figure 6.6. These results reveal that the minimal design is more rotationally variant than the 5-design when used for measuring μFA in aligned microscopic diffusion environments with high anisotropy or diffusivity. The standard deviation of the μFA estimates ranges from 0 to 0.1 for the 5-design and 0 to 0.26 for the minimal design with the mean of the standard deviation of μFA estimates over the whole parameter space being 0.03 and 0.06 for the 5-design and the minimal design, respectively.

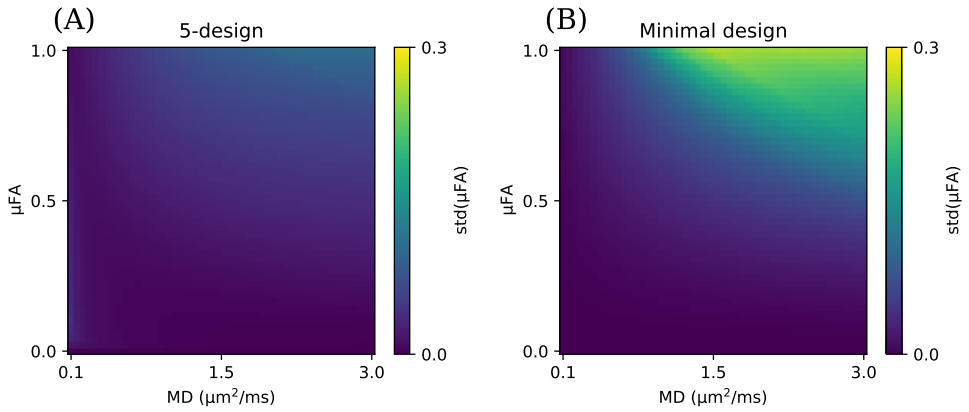

Figure 6.6: *The standard deviations of the μFA estimates from 400 simulation experiments of axially symmetric single diffusion tensors with directions uniformly distributed over half a sphere with the 5-design (A) and the minimal design (B) using a b-value of 2.25 ms/μm^2. No noise was added to the data.*

Figure 6.7 shows the estimated minimum required SNR for measuring μFA within 0.1 margin from its asymptotic value with 95% confidence using the minimal design. The simulation results reveal that the precision of μFA is highly dependent on its value and that particularly low values of MD, 0.25 μm^2/ms or less (depending on the SNR and the b-value), prevent precise measurement of μFA irrespective of its value. These results suggest that an SNR of 50 would enable reliable measurement of μFA values greater than 0.75, and that an SNR of nearly 200 would be necessary for reliably measuring μFA values of less than 0.5. The results also show that, by increasing the magnitude of diffusion-weighting, reliable quantification of μFA becomes possible in smaller compartments. On the other hand, higher level of diffusion-weighting prevents measurements in large compartments due to the signal having hit the noise floor. These results suggest that to maximize precision in human brain tissue, where MD is usually around 0.8 μm^2/ms, it is desirable to use a b-value near 2.25 ms/μm^2. Greater diffusion-weighting increases the precision of the noisy estimate of small compartments' μFA but without enabling precise measurement of low values of μFA.

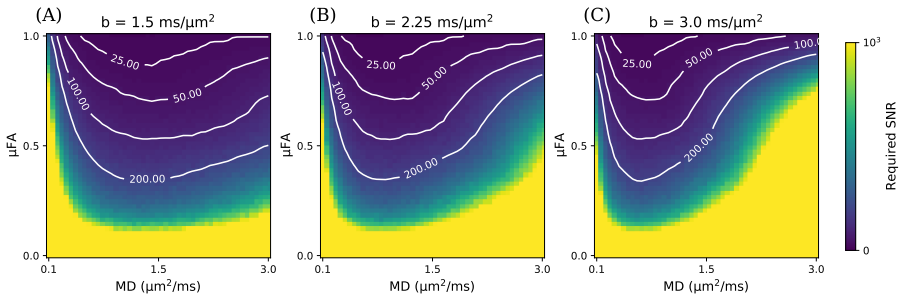

Figure 6.7: *The minimum SNR required for measuring μFA within 0.1 margin from its asymptotic value with 95% confidence with 3 different b-values and varying diffusion tensor size and shape. The figure was generated from 10^3 repetitions of noisy simulations of axially symmetric single diffusion tensors aligned with the x-axis using the minimal design at 100 levels of SNR equally distributed between 1 and 10^3. The contour lines follow data smoothed with a Gaussian filter (σ = 0.8 pixels).*

6.5 Discussion

The 5-design is a well-established method for estimating the powder-averaged DDE signal from which μFA, a clinically and scientificaly relevant microstructural parameter[54,134,135,64,55–61,63], can be calculated without strict *a priori* assumptions about tissue microstructure. The minimal design enables crucial acceleration of the method for human neuroimaging but prior to this study its accuracy had not been validated against the theoretically justified 5-design. Hence, the study presented in this chapter had two main purposes. The first goal of the study was to experimentally validate the minimal design in neural tissue by directly comparing it to the 5-design using a pre-clinical scanner that enables long acquisition times and high SNR. Second, simulation experiments were performed to assess μFA's noise robustness as a function of the ground-truth μFA and MD over a range of relevant b-values.

While the 5-design requires 72 separate acquisitions per b-value, the minimal design requires a minimum of only 12 acquisitions. In this study, a symmetric version of the minimal design was used, consisting of 24 acquisitions to eliminate the possibility of bias arising from cross-terms between the imaging and diffusion encoding gradients[205]. After correcting for the different number of acquisitions in the two protocols, a comparison of the μFA maps acquired with the two methods revealed that the maps are nearly indistinguishable. In fact, μFA's variance in the test-retest of the 5-design was very similar

to the comparison of the 5-design and the minimal design. Thus, the estimated powder-averaged signal calculated from 72 measurements of the DDE 5-design can be closely approximated in neural tissue, including both gray matter and white matter tissues, by the minimal design in the b-value range that was studied here. The results presented here hold for both *ex vivo* and *in vivo* conditions, thereby excluding tissue fixation effects as confounders.

For every orthogonal gradient pulse pair in the 5-design, there are three other gradient pulse pairs encoding the same b-tensor. Thus, in case of Gaussian diffusion and infinite SNR, 15 orthogonal wave vector pairs are sufficient for measuring the exact same powder-averaged signal than what is measured with the 5-design. In voxels with sufficient orientation dispersion of the microscopic diffusion environments, this can be further reduced to 6 as shown by Yang et al.[56]. However, if the microscopic compartments are perfectly aligned, the minimal design suffers from larger rotational variance than the 5-design, as shown by the simulations presented here. This result is in good agreement with previous reports on the minimum number of directions required for estimating the powder-averaged signal using QTE[202].

The underestimation of μFA due to higher order effects can be corrected for by acquiring at least four diffusion-weighting shells (including $b = 0$) and applying the higher order correction[204]. However, every shell consumes more experimental time that is limited in human neuroimaging. In this study, the higher order correction was not applied to avoid the prolonged acquisition and therefore the μFA estimates are somwehat biased. Nevertheless, based on the simulations in this study, the bias induced by the higher order terms is expected to be, on average, less than 0.1 if $b = 2.25$ ms/μm^2 and MD = 1 μm^2/ms. The accuracy of several signal models for estimating μFA is discussed in detail in Chapter 7. It is also worth mentioning that if the biased μFA is precise, it can be highly useful both in neuroscience research and in the clinic.

The investigation of μFA's precision revealed two important features: μFA's precision depends strongly and nonlinearly on MD and on the actual μFA value. When MD is very small, the estimation of μFA is hampered even with very high SNR, and for reasonable values of MD, precise estimates of μFA < 0.5 require very high (> 200) SNR. On the other hand, large values of MD result in very diminished signal, impeding precise measurements. The simulations also suggest that the optimal b-value for measuring μFA in the brain, where MD ≈ 0.8 μm^2/ms, is near 2.25 ms/μm^2, a result which is consistent with previous stud-

ies[56]. Combined, these findings suggest that caution needs to be exercised in interpreting μFA maps, and the SNR should be reported to assess the confidence in the estimates' precision. In particular, μFA maps are more reliable in areas with higher μFA, which should be taken into account upon interpretation of results. Therefore, μFA may be a powerful metric for observing axonal degeneration in white matter voxels with high values of μFA but it lacks the precision to be able to reliably pick up microstructural changes in gray matter where μFA is low. This lack of precision for low values of μFA have been previously reported with the μFA values derived from QTE experiments[135,64]. Also, the large variance in voxels with low μFA can result in negative values of μA^2, leading to imaginary values of μFA. It is important to note that the simulations incorporated diffusion tensors, rather than restricted diffusion, whose effects were not the focus of this study. Still, given the experimental results, it seems that the Gaussian regime was not severely violated in our experiments or the μFA derived from the 5-design and minimal design would have significantly differed.

These results show that the minimal design provides an accelerated method for estimating μFA approximately as robustly as the 5-design in neural tissue. Additionally, it may be possible to use the minimal design to shorten the acquisition time required for calculating μFA maps with the higher order correction, providing more accurate estimates. However, three issues must be controlled for when applying the minimal design. First, it is very challenging to reliably measure values of μFA lower than 0.5, e.g., in gray matter, with SNR typically achievable with clinical whole-body scanners. Second, the rotational variance of the minimal design may result in biased signal powder average in tissues where microenvironments are nearly aligned and highly anisotropic. Third, when using DDE with clinical whole-body scanners, it is important to control for concomitant fields, which can arise from asymmetric gradient waveforms and give rise to substantial signal bias[208]. Furthermore, if the assumption of apparent Gaussian diffusion is correct, then QTE would probably be more suitable for mapping μFA in the human brain due to more efficient diffusion encoding[209]. Microscopic diffusion anisotropy estimation with QTE will be discussed in the next chapter.

6.6 Conclusion

In conclusion, reducing the number of wave vector rotations in calculating the powder-averaged DDE signal, as proposed by Yang et al.[56], does not prevent the accurate quantification of microscopic diffusion anisotropy in neural tissue. However, the SNR requirements for precise quantification of μFA should be carefully considered when applying the minimal design for estimating μFA.

Chapter 7

Comparative analysis of signal models for microscopic anisotropy estimation using q-space trajectory encoding

7.1 Introduction

DDE gradient waveforms consist of two trapezoidal pulse pairs separated by a mixing time (Figure 6.1) and therefore require rather long echo times, resulting in a diminished signal. Asymmetric DDE waveforms can also produce significant concomitant gradients in spin echo experiments, causing artefacts and signal dropout[208]. QTE gradient waveforms, on the other hand, are designed to enable shorter echo times by following an optimized trajectory in q-space to define a b-tensor with the desired shape, magnitude, and orientation (Equation 3.3.6)[209]. This makes it feasible to perform MDE experiments using more limited scanner hardware[202], albeit at the cost of assuming diffusion time dependency to be negligible because QTE gradient waveforms lack a well-defined scalar diffusion time. If this assumption is not valid, a discrepancy between the diffusion times of the gradient waveforms encoding different b-tensor shapes results in biased parameter estimates[210]. Furthermore, isotropic diffusion encoding with QTE may be orientationally variant in the case of significant time-dependent diffusion anisotropy[211,212].

Promising results of the application of μFA (Equation 3.7.6), a normalized measure of microscopic diffusion anisotropy that is equivalent to conventional

FA if orientation dispersion is zero, have been reported in imaging brain tumours[64,55], multiple sclerosis lesions[56,60], white matter degeneration in Parkinson's disease patients[63], and microstructural properties of brains of patients with schizophrenia[65] and epilepsy[59]. However, different studies have applied different methods for estimating μFA and the differences between the methods have received little attention thus far. Understanding how different μFA estimation methods relate to each other is crucial in planning future studies and interpreting reported results in the literature. Furthermore, it is important to reach a consensus on the optimal method for measuring μFA, an essential requirement for future clinical studies.

In this study, imaging experiments and simulations were performed to compare the accuracy and precision of three μFA estimation methods: the truncated cumulant expansion of the powder-averaged signal[134,135,204,56,57], the gamma-distributed apparent diffusivities assumption[135,64,55,60], and q-space trajectory imaging (QTI)[65,59,63], a generalization of the truncated cumulant expansion to individual acquisitions. These methods were chosen as they have each been used in several studies and it is crucial to understand how they relate to each other in order to interpret the reported results. The μFA maps calculated from the same data using the different methods were compared and the precision of the estimates was quantified by repeatedly imaging a microfibre phantom. Furthermore, since the ground-truth value of μFA was unknown in the imaging experiments, Monte Carlo random walk simulations were performed using axon-mimicking fibres for which the ground truth was known. The simulations were also used to assess the effect of time-dependent diffusion on the parameter estimates. Although the work presented in this chapter focuses on QTE, these signal models can be applied to analyze DDE data as well.

7.2 Theory

This section provides a concise summary of the theory behind μFA estimation using QTE. A more detailed description can be found, for example, in the review by Topgaard[115].

Conventional DTI represents diffusion at the voxel level by a symmetric positive definite 3×3 tensor $\mathbf{D}$ that captures Gaussian diffusion. In imaging experiments with moderate to high diffusion-weighting (roughly $b > 1$ ms/μm^2 in the brain), the diffusion-weighted signal attenuation clearly deviates from

a monoexponential decay, revealing that the voxel-level diffusion is not Gaussian, especially in white matter[46]. At the voxel level, non-Gaussian diffusion is caused by restricted diffusion within the microscopic diffusion environments, referred to as compartments hereafter, and varince in diffusion properties across the compartments[125].

If inter-compartmental exchange of water and intra-compartmental non-Gaussianity of diffusion are negligible, the tissue in a voxel can be represented by a distribution of microscopic diffusion tensors. In this case, the signal can be expressed as

$$S = S_0 \int P_{6\mathrm{D}}(\mathbf{D}_\mu) \exp\left(-\mathbf{b} : \mathbf{D}_\mu\right) d\mathbf{D}_\mu, \tag{7.2.1}$$

where S_0 is the signal without diffusion-weighting, $P_{6\mathrm{D}}$ is the normalized distribution of microscopic diffusion tensors in the voxel, $\mathbf{D}_\mu$ is a microscopic diffusion tensor, $\mathbf{b}$ is the b-tensor used in the acquisition, : denotes the generalized scalar product between the two tensors, and integration is performed over all symmetric positive definite 3×3 tensors[115].

Since recovering the six-dimensional microscopic diffusion tensor distribution is an ill-posed problem[115], several methods have been developed for estimating relevant properties of it. One of these is μFA, a normalized measure of the average eigenvalue variance of the microscopic diffusion tensors[135,55]:

$$\mu\mathrm{FA} = \sqrt{\frac{3}{2} \frac{\langle \mathrm{Var}_\lambda(\mathbf{D}_\mu) \rangle}{\langle \mathrm{Var}_\lambda(\mathbf{D}_\mu) \rangle + \langle [\mathrm{Tr}(\mathbf{D}_\mu)/3]^2 \rangle}}, \tag{7.2.2}$$

where $\mathrm{Var}_\lambda(\,)$ denotes an operator that calculates the eigenvalue variance of a tensor, $\mathrm{Tr}(\,)$ denotes the trace of a tensor, and $\langle\,\rangle$ denotes averaging over the microscopic diffusion tensor distribution.

q-space trajectory encoding

QTE is based on the idea that if the time dependency of the microscopic diffusion tensors is negligible, the diffusion-weighted signal attenuation depends only on the applied b-tensor and not on the gradient wave vector's trajectory in q-space. A given b-tensor shape can be efficiently encoded by defining a q-space trajectory $\mathbf{q}(t)$ that maximizes the b-value. An optimized q-space trajectory to be used in a spin echo experiment can be found by solving the following

optimization problem[209]:

$$
\begin{aligned}
\underset{\mathbf{q}}{\text{maximize}} \quad & \mathrm{Tr}(\mathbf{b}) \\
\text{subject to} \quad & \int_0^{TE} \mathbf{q}(t)\mathbf{q}^{\mathrm{T}}(t)dt = \mathbf{b} \\
& \mathbf{b}/\mathrm{Tr}(\mathbf{b}) = \mathbf{M} \\
& \mathbf{q}(t) = \mathbf{q}(TE) = 0 \\
& \left.\frac{d\mathbf{q}}{dt}\right|_{t\in I_{RF}} = 0 \\
& \left\| \frac{d\mathbf{q}}{dt} \right\| \leq \gamma G_{\max} \\
& \left\| \frac{d^2\mathbf{q}}{dt^2} \right\| \leq \gamma R_{\max} \\
& \int_0^{TE} \left(\frac{dq_i}{dt}\right)^2 dt \leq \eta\gamma^2 G_{\max}^2(TE), i = x, y, z,
\end{aligned}
\tag{7.2.3}
$$

where $\mathbf{b}$ is the b-tensor, $\mathbf{M}$ is a tensor defining the b-tensor shape, TE is the echo time, I_{RF} is the time interval during which the gradient amplitude has to be zero because of the application of the refocusing pulse, γ is the gyromagnetic ratio, $G_{\max}$ is the maximum gradient amplitude, $R_{\max}$ is the maximum slew rate, and $\eta \in [0, 1]$ is a parameter controlling the time given for the gradient coil to cool down via heat dissipation.

Powder-averaged signal

The powder-averaged signal, i.e., the signal averaged over all possible diffusion encoding directions using the same b-tensor shape and magnitude, is orientationally invariant and thus enables the dimensionality of Equation 7.2.1 to be reduced:

$$
S = S_0 \int_0^{\infty} P(D_\mu) \exp\left(-bD_\mu\right) dD_\mu, \tag{7.2.4}
$$

where P is the distribution of apparent microscopic diffusivities in the voxel, b is the b-value used in the acquisition, and D_μ is the apparent microscopic diffusivity[135]. P depends on the shape of the b-tensor used in the acquisition. The powder-averaged signal can be estimated by averaging over the acquired gradient directions, given that a sufficient number of gradient directions was used[134,202].

Using the cumulant expansion[124] up to the second order in b, the powder-

averaged signal can be expressed as

$$S \approx S_0 \exp\left(-b\text{MD} + b^2\frac{V}{2}\right), \tag{7.2.5}$$

where $\text{MD} = \langle \text{Tr}(\mathbf{D}_\mu)/3 \rangle$ is the voxel-level mean diffusivity and V is the second central moment, i.e., variance, of P[135]. By combining acquisitions with different b-tensor shapes, V can be decomposed into anisotropic and isotropic variances, V_{aniso} and V_{iso}, respectively. V_{aniso} is proportional to the average eigenvalue variance of the microscopic diffusion tensors and V_{iso} is equal to the variance of mean diffusivities of the microscopic diffusion tensors[115]:

$$V_{\text{aniso}} = \frac{2}{5}\langle \text{Var}_\lambda(\mathbf{D}_\mu)\rangle, \tag{7.2.6}$$

$$V_{\text{iso}} = \langle [\text{Tr}(\mathbf{D}_\mu)/3]^2\rangle - \text{MD}^2. \tag{7.2.7}$$

In experiments performed with linear and spherical b-tensors, as is the case in this study, the total variance relates to V_{aniso} and V_{iso} as

$$V_{\text{LTE}} = V_{\text{aniso}} + V_{\text{iso}}, \tag{7.2.8}$$

$$V_{\text{STE}} = V_{\text{iso}}, \tag{7.2.9}$$

where the subscripts LTE and STE represent acquisitions with linear tensor encoding (Equation 3.7.2) and spherical tensor encoding (Equation 3.7.4), respectively[115].

The limitation of estimating μFA as described above is that Equation 7.2.5 is strictly valid at $b \to 0$. However, to estimate μFA, the applied b-value has to be sufficiently high for the voxel-level non-Gaussian diffusion to be measurable. The accuracy of μFA estimation using the truncated cumulant expansion and DDE has been shown to improve if an additional term is introduced to account for higher order effects[204]:

$$S \approx S_0 \exp\left(-b\text{MD} + b^2\frac{V}{2} - b^3 P_3\right), \tag{7.2.10}$$

where P_3 is non-negative.

Another method for avoiding the invalidity of the cumulant expansion at higher b-values is to *a priori* assume the distribution of apparent microscopic diffusivities to be such that enables Equation 7.2.4 to be analytically evaluated, e.g., gamma-distributed[135]:

$$P(D_\mu; k, \theta) = \frac{(D_\mu)^{k-1} e^{-\frac{D_\mu}{\theta}}}{\theta^k \Gamma(k)}, \tag{7.2.11}$$

where $k \in \mathbb{R}_+$ and $\theta \in \mathbb{R}_+$ are the shape and scale parameters of the gamma distribution, respectively, and Γ is the gamma function. Under this assumption, the powder-averaged signal can be expressed as

$$S \approx S_0 \left(1 + b\frac{V}{\mathrm{MD}}\right)^{-\frac{\mathrm{MD}^2}{V}}, \tag{7.2.12}$$

where $\mathrm{MD} \approx k\theta$ and $V = k\theta^2$. However, the resulting parameter estimates are naturally inaccurate if the distribution of apparent diffusivities is not gamma-distributed.

Covariance tensor

The second order cumulant expansion (Equation 7.2.5) has also been generalized to describe individual acquisitions instead of the powder-averaged signal. In QTI[65], the signal in an individual acquisition is expressed as

$$S \approx S_0 \exp\left(-\mathbf{b} : \mathbf{D} + \frac{1}{2}(\mathbf{b} \otimes \mathbf{b}) : \mathbb{C}\right), \tag{7.2.13}$$

where $\otimes$ denotes the tensor outer product and $\mathbb{C}$ is a $3 \times 3 \times 3 \times 3$ covariance tensor defined as

$$\mathbb{C} = \langle \mathbf{D}_\mu \otimes \mathbf{D}_\mu \rangle - \mathbf{D} \otimes \mathbf{D}. \tag{7.2.14}$$

The covariance tensor contains 21 unique elements and enables the estimation of several relevant properties of the microscopic diffusion tensor distribution. μFA can be calculated as

$$\mu\mathrm{FA} = \sqrt{\frac{3}{2}\frac{\langle \mathbf{D}_\mu \otimes \mathbf{D}_\mu \rangle : \mathbb{E}_{\mathrm{shear}}}{\langle \mathbf{D}_\mu \otimes \mathbf{D}_\mu \rangle : \mathbb{E}_{\mathrm{iso}}}}, \tag{7.2.15}$$

where $\mathbb{E}_{\mathrm{iso}} = \mathbb{I}_6/3$ and $\mathbb{E}_{\mathrm{shear}} = \mathbb{E}_{\mathrm{iso}} - (\mathbb{I}_3/3 \otimes \mathbb{I}_3/3)$, where $\mathbb{I}$ denotes an identity matrix, and the $3 \times 3 \times 3 \times 3$ rank-4 tensors are represented by 6×6 rank-2 tensors as described by Westin et al.[65].

7.3 Methods

MRI data was acquired using a prototype MDE spin echo sequence[213] on a Siemens Magnetom Prisma 3T with maximum gradient strength of 80 mT/m, maximum slew rate of 200 T/m/s, and a 64-channel head coil (Siemens Healthcare, Erlangen, Germany) at Great Ormond Street Hospital, London, United

Kingdom. Images were preprocessed by applying random matrix denoising[178] and Gibbs ringing correction[179] using MRtrix3[177] and distortion correction with FSL's[181] topup and eddy[180].

Volunteer experiments

The brains of three healthy adult volunteers (two females and one male with ages ranging from 27 to 42 years) were imaged with numerically optimized[209] and Maxwell-compensated[208], i.e., corrected for concomitant fields arising from gradient waveform asymmetry, gradient waveforms encoding linear and spherical b-tensors (Figure 7.1A-B). The volunteers gave written and informed consent prior to the scan. To reduce peripheral nerve stimulation and gradient coil heating, the slew rate was constrained to a maximum of 65 T/m/s when calculating the gradient waveforms using the software provided by Sjölund et al.[209].

Under the Gaussian phase approximation[214], the diffusion time of a QTE gradient waveform can be quantified by calculating the power spectrum of the corresponding wave vector[103]:

$$\mathrm{PSD}(\omega) = \left| \int_0^{TE} q(t) \exp\left(-i\omega t\right) dt \right|^2 , \tag{7.3.1}$$

where PSD stands for power spectral density, q is the magnitude of the wave vector along a single direction, and ω is the diffusion frequency. This formalism enables the problem to be analyzed in the diffusion frequency domain (Equation 3.2.10). The power spectra of the Cartesian components of the waveforms used in the imaging experiments are shown in Figure 7.1E-F.

Since the QTE waveform for spherical encoding exhibits spectral anisotropy which may lead to an orientationally variant signal[212], the orientational variance of the acquisition was quantified using FA. DTI was fit separately to the data acquired with linear and spherical tensor encoding using a weighted linear least squares fit in Dipy[184]. If the STE acquisition is orientationally invariant, non-zero values of FA are caused by noise and they are uncorrelated with the FA values calculated from the LTE data.

Twelve diffusion encoding directions were used for b-values 0.1, 0.5, and 1 ms/μm^2 and 32 directions for b-values 1.5 and 2 ms/μm^2. Gradient waveforms for both linear and spherical encoding were rotated[213]. The directions were distributed uniformly around the surface of a sphere by combining the vertex coordinates of the icosahedron and the dodecahedron[65]. The sequence also

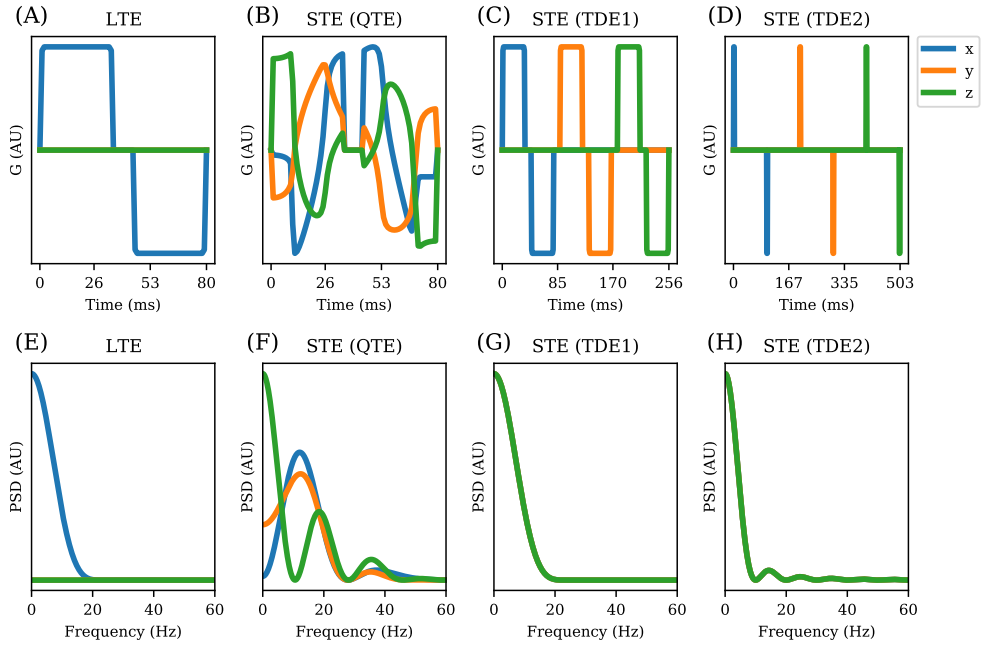

Figure 7.1: *The diffusion encoding gradient waveforms used in the imaging experiments (A-B) and simulations (A-D) and the corresponding power spectra (E-H). The LTE waveforms in the TDE simulations were equal the x-components of the corresponding STE waveforms. It is important to mention that the power spectra of the unrotated TDE waveforms are identical along x, y, and z in the lab frame but not along arbitrary directions.*

included 15 images without diffusion-weighting, one of which had the phase encoding direction reversed. Other relevant imaging parameters were: voxel size = $2 \times 2 \times 2$ mm^3, FOV = 256×256 mm^2, 60 slices, TE = 103 ms, TR = 10 s, partial Fourier = 6/8, and 2 repetitions. The total scan time was approximately 40 minutes per volunteer.

Phantom experiments

A phantom consisting of highly hydrophilic hollow polycaprolactone microfibres[137] was imaged using the same acquisition protocol as the volunteers except the following parameters: TR = 3 s, 10 slices, and 1 repetition.

The phantom contained three ROIs with different fibre configurations: parallel fibres, perpendicularly crossing fibres, and fibres with random orientations. The inner diameter of the fibres was 9.9 ± 1.2 µm (mean $\pm$ standard deviation) in the ROI containing parallel and crossing fibres, and 7.8 ± 0.5 µm in the ROI containing randomly oriented fibres. The randomly oriented fibres

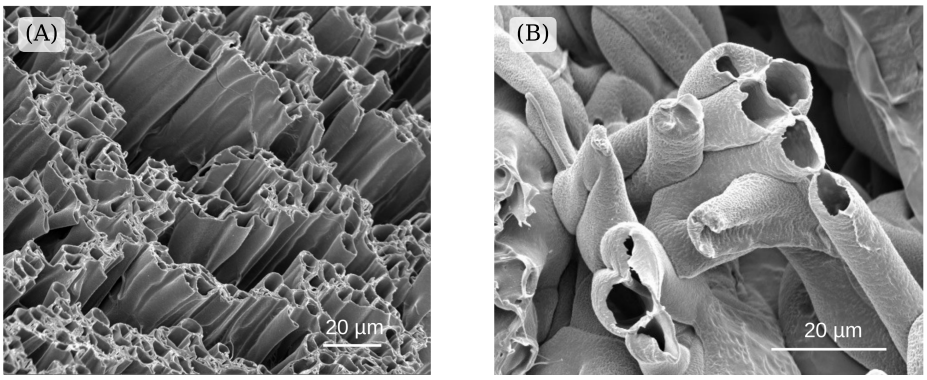

Figure 7.2: *Representative scanning electron microscope images illustrating the microstructure of the parallel microfibre phantom (A) and the randomly oriented microfibre phantom (B). The phantom with crossing microfibres was made from the same sample as the one with parallel microfibres.*

were placed randomly on a 2D plane. Scanning electron microscope images illustrating the microstructure of the phantom are shown in Figure 7.2.

The acquisition was performed 11 times to study the precision of the µFA estimates. The coefficient of variation was calculated at each voxel as

$$\mathrm{CV} = \frac{\sigma}{\mu} \cdot 100\%, \tag{7.3.2}$$

where σ is the standard deviation of µFA estimated using a given method over the repeated acquisitions and μ is the mean µFA averaged over all acquisitions and methods. A single value of μ was used to not penalize the methods that produce lower values of µFA.

The signal anisotropy of the acquisitions with the two waveforms was quantified using the same method as in the volunteer experiments.

Simulation experiments

GPU-accelerated Monte Carlo random walk simulations were performed using Disimpy[157] to compare the µFA estimates of the different methods in a scenario where the ground truth was known. Three simulation experiments with different gradient waveforms were performed.

Simulation 1 was performed using the same waveforms as in the imaging experiments. Then, in simulations 2 and 3, the simulation was repeated using triple diffusion encoding (TDE). The TDE waveforms consist of three gradient pulse pairs with orthogonal directions and with identical diffusion encoding

times (δ), diffusion times (Δ) and gradient magnitudes. Similarly to DDE, the difference between the onset of the subsequent gradient pulses with orthogonal directions is referred to as mixing time (τ). Following the approach proposed by Lundell et al.[210], Simulation 2 was performed with the waveform shown in Figure 7.1C ($\delta = 35$ ms, $\Delta = \tau = 44$ ms) to reduce parameter bias due to spectral anisotropy and the discrepancy in diffusion times between the gradient waveforms encoding linear and spherical b-tensors. The linear waveform in Simulation 2 was equal to the x-component of the spherical waveform. Finally, Simulation 3 was performed using a TDE waveform (Figure 7.1D) with a shorter pulse length ($\delta = 3$ ms), longer diffusion time ($\Delta = 100$ ms), and longer mixing time ($\tau = 100$ ms) to enable the random walkers to more fully sample their environments to minimize the correlations between their positions during the different gradient pulse pairs. The linear waveform in Simulation 3 was a TDE waveform with the same direction for the three pulse pairs and with the same pulse length, diffusion time, and mixing time as the spherical waveform. It is important to mention that the power spectra of the unrotated TDE waveforms are identical along x, y, and z in the lab frame but not along arbitrary directions.

The simulated voxel, shown in Figure 7.3A, contained 381 axon-mimicking irregular and impermeable fibres generated using ConFiG[168]. A histogram of the simulated axon radii is shown in Figure 7.3B. The thickness of the fibre surface was zero. The simulated voxel had periodic boundary conditions, i.e., the random walkers that left the voxel encountered repeating identical copies of the shown microstructure. The volume of the simulated voxel was 39 $\times$ 39 $\times$ 32 μm^3 of which 60% was occupied by the synthetic axons. $3 \cdot 10^6$ random walkers were used. $1.095 \cdot 10^6$ random walkers were randomly placed in the extra-axonal space and $5 \cdot 10^3$ inside each axon. Simulations 1, 2, and 3 were performed with 10^4, 31720, and 62195 time steps, respectively, to keep the step length constant (l = 0.31 μm) across the simulations. A uniform diffusivity of 2 μm^2/ms was used.

To study the precision of μFA, SNR of the simulated signals was lowered to 25. Noisy signals were generated as

$$S_{\text{noisy}} = \sqrt{(S + X)^2 + Y^2}, \tag{7.3.3}$$

where S is the simulated signal without noise and X and Y are randomly sampled from a normal distribution with zero mean and standard deviation σ. Here, SNR refers to the signal without diffusion-weighting divided by σ.

The ground-truth μFA in each simulation was calculated according to Equation 7.2.2 by separately fitting diffusion tensors to the LTE signals from each compartment, i.e., the synthetic axons and the extra-axonal space. The diffusion tensors were estimated by fitting diffusion and kurtosis tensors[46] to data using a weighted linear least squares algorithm in Dipy[184]. To vary the ground-truth μFA, three signal fractions for the spins in the intra-axonal space were used: 0.2, 0.6, and 1.

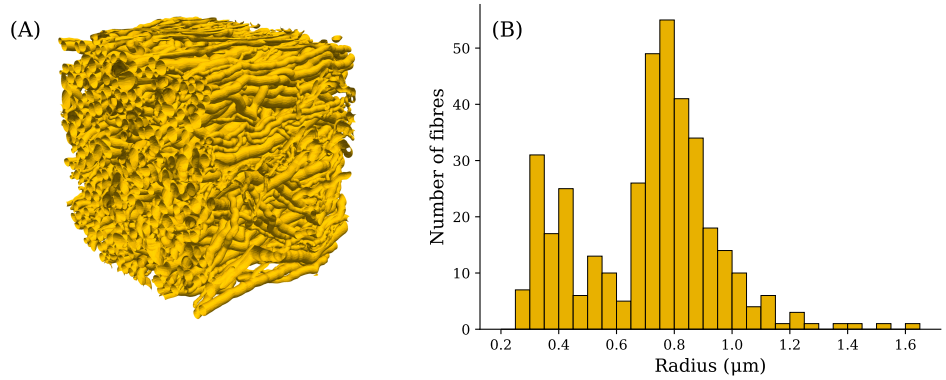

Figure 7.3: *(A) The simulated voxel containing 381 axon-mimicking fibres. (B) A histogram of the fibre radii (mean along the fibre length).*

μFA estimation

Four μFA estimation methods were compared: regression of the cumulant expansion of the powder-averaged signal up to the second order in b and with a higher order correction, the gamma-distributed diffusivities assumption, and QTI. For brevity, these signal models will be referred to hereafter as the cumulant model, higher order model, gamma model, and QTI, respectively.

The cumulant model, higher order model, and gamma model were fit to the powder-averaged data using a non-linear least squares trust region reflective algorithm[207] in Scipy[164] and equations 7.2.5, 7.2.10, and 7.2.12, respectively. S_0, MD, V_{iso}, V_{aniso}, and P_3 were used as fit parameters from which μFA was estimated using equations 7.2.2, 7.2.6, and 7.2.7. P_3 was included in the equation for the signal acquired with linear b-tensors. The fit parameters were constrained to be non-negative real numbers. The initial values of the fit parameters were: S_0 = average signal over images with no diffusion-weighting, MD $= 1$ μm^2/ms, $V_{\text{iso}} = 0.1$ μm^4/ms^2, $V_{\text{aniso}} = 0.1$ μm^4/ms^2, and P_3 = 0.

QTI was fit to data using the linear equation

$$\boldsymbol{\beta} = \left(\mathbf{X}^{\mathrm{T}}\mathbf{C}\mathbf{X}\right)^{-1}\left(\mathbf{X}^{\mathrm{T}}\mathbf{C}\right)\mathbf{S}, \tag{7.3.4}$$

where

$$\hat{\boldsymbol{\beta}} = \begin{pmatrix} \ln S_0 & \mathbf{D} & \mathbb{C} \end{pmatrix}^{\mathrm{T}}, \tag{7.3.5}$$

$$\mathbf{X} = \begin{pmatrix} 1 & -\mathbf{b}_1^{\mathrm{T}} & \frac{1}{2}(\mathbf{b}_1 \otimes \mathbf{b}_1)^{\mathrm{T}} \\ \vdots & \vdots & \vdots \\ 1 & -\mathbf{b}_n^{\mathrm{T}} & \frac{1}{2}(\mathbf{b}_n \otimes \mathbf{b}_n)^{\mathrm{T}} \end{pmatrix}, \tag{7.3.6}$$

$$\mathbf{S} = \begin{pmatrix} \log S_1 \\ \vdots \\ \log S_n \end{pmatrix}, \tag{7.3.7}$$

where n is the number of acquisitions and $\mathbf{D}$, $\mathbb{C}$, $\mathbf{b}_i$, and $(\mathbf{b}_i \otimes \mathbf{b}_i)$ are column vectors in Voigt notation[65]. The diagonal matrix $\mathbf{C}$ with elements $C_{ii} = S_i$ was used to correct for heteroscedasticity in the log-transformed data[116]. The matrix inversion in Equation 7.3.4 was performed using the Moore-Penrose pseudoinversion[215] in Numpy[163]. μFA was calculated from the estimated values of $\mathbf{D}$ and $\mathbb{C}$ according to equations 7.2.14 and 7.2.15.

7.4 Results

Volunteer experiments

Representative axial slices of the μFA maps in one of the volunteers are shown in Figure 7.4. A visual inspection suggests that the gamma model results in higher μFA values than the other methods and that the higher order model results in noisier maps than the other methods.

The Bland-Altman plots in Figure 7.5 show the voxel-wise differences between the methods' μFA estimates against their mean values across all volunteers. The good agreement between the cumulant model and QTI can be seen in the small mean difference (0.01) and the small width of the 95% central range of the differences (0.17). The gamma model produced the highest values of μFA, 0.1 greater than QTI when averaged over all data. The plots also highlight the heteroscedasticity of μFA.

An ordinary least-squares linear regression of the differences on the averages revealed that the differences between the estimates are proportional to

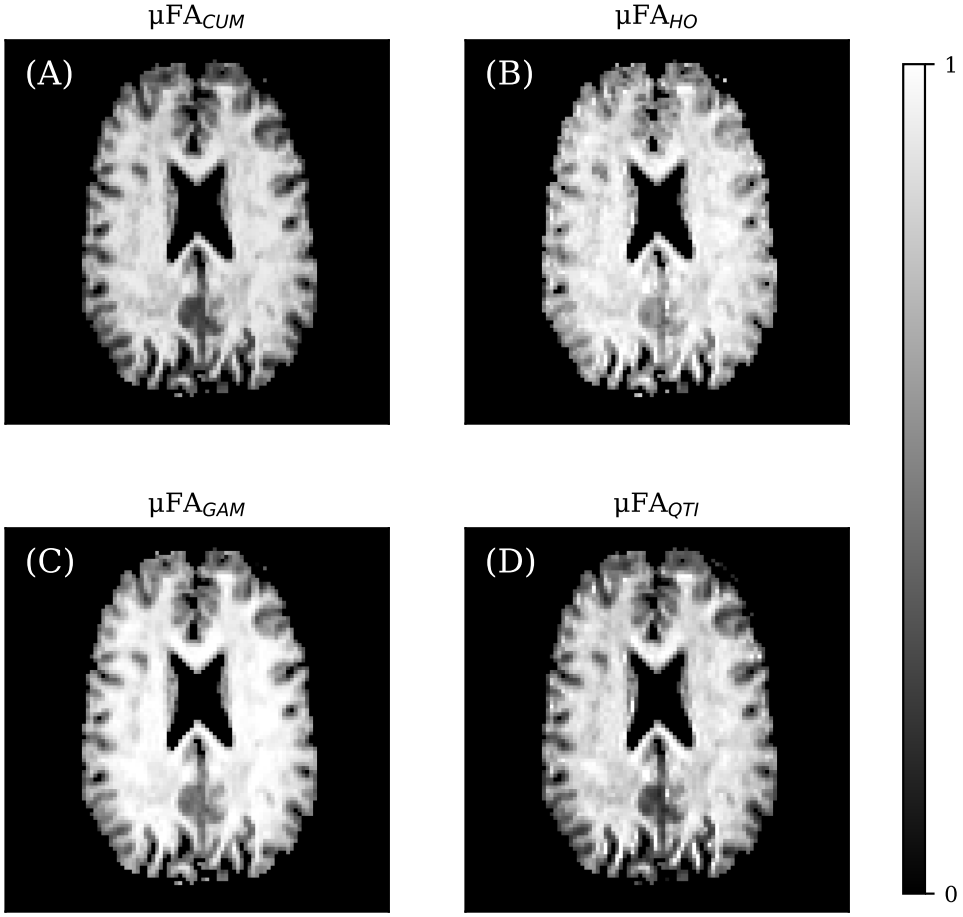

Figure 7.4: *Representative axial slices of the μFA maps in one of the volunteers calculated using the cumulant model (A), higher order model (B), gamma model (C), and QTI (D).*

the mean value of μFA for all pairs of methods ($p < 10^{-3}$). However, since the relationship between the methods' estimates is non-linear, a linear regression does not provide a satisfactory fit to data. Specifically, the higher order model's relationship to the other methods is strikingly non-linear. The width of the 95% central range of the differences was largest in the comparisons involving the higher order model

A significant correlation (Pearson's $R = 0.22, p < 10^{-3}$) was found between the FA maps calculated from the data acquired with linear and spherical b-tensors, revealing that the acquisition with isotropic diffusion-weighting was moderately orientationally variant.

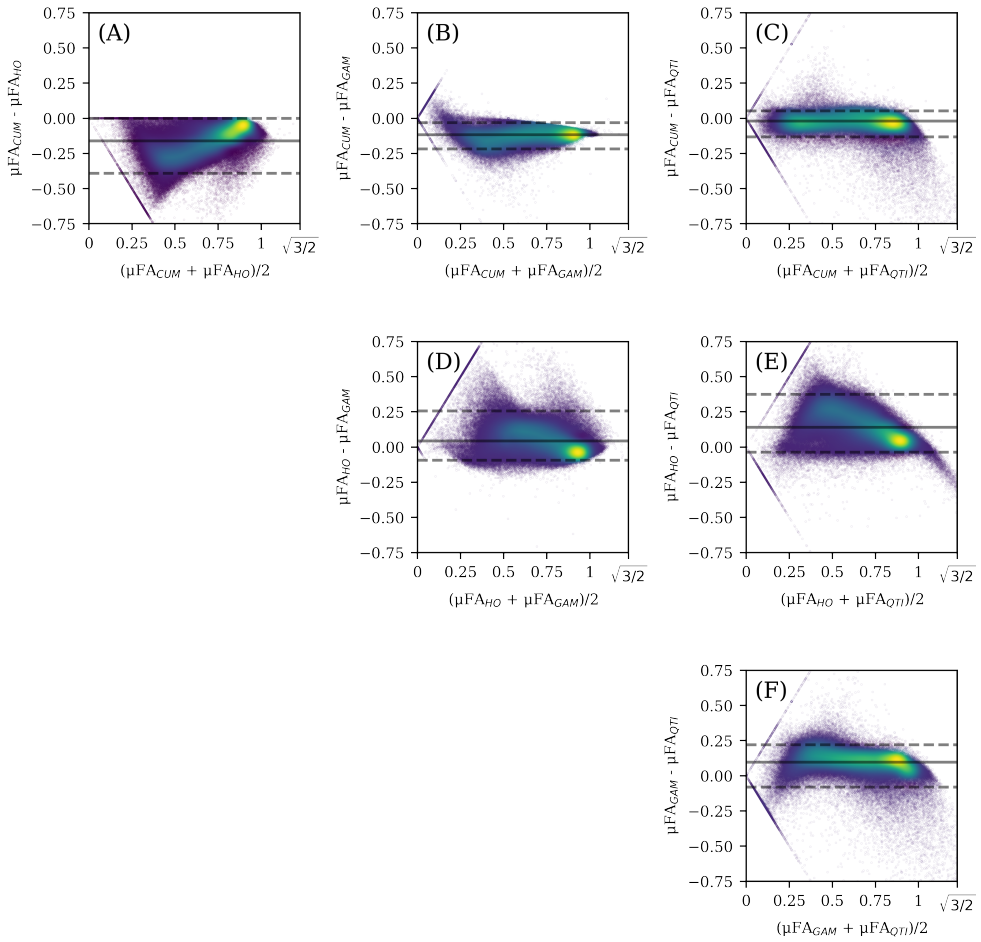

Figure 7.5: *Voxel-wise comparison of the cumulant model (μFA_{CUM}), higher order model (μFA_{HO}), gamma model (μFA_{GAM}), and QTI (μFA_{QTI}). The differences in μFA are plotted against their mean value. The dashed lines show the 2.5th and the 97.5th percentiles of the distribution of the differences. The solid black lines depict the mean difference. Colour represents data point density. The figure was generated using data from all volunteers after excluding the voxels in the cerebrospinal fluid using a mask that was genereated by thresholding the MD map with a threshold of 2 $\mu m^2/ms$.*

Phantom experiments

μFA values averaged over the phantom ROIs from the repeated acquisitions are shown in Figure 7.6A-C. The gamma model consistently resulted in the greatest values of μFA, followed by the higher-order model, cumulant model, and QTI. The differences in mean μFA calculated using the different signal models were statistically significant in all pair-wise comparisons using an independent samples t-test ($p < 10^{-3}$). Despite the smaller fibre radius, μFA was the lowest in the ROI with randomly oriented fibres due to the lower packing density of the microfibres.

The distributions of the voxel-level CVs are shown in Figure 7.6D-F. The median CVs were well below 2%, indicating good repeatability in voxels with high μFA. The higher-order model stood out for having the lowest repeatability, especially in the phantom ROIs with crossing or randomly oriented fibres. A significant correlation ($R^2 = 0.11, p < 10^{-3}$) was found between the FA maps calculated from the data acquired with waveforms encoding linear and spherical b-tensors.

The average SNR was 123 as quantified by the mean signal divided by the standard deviation over the images without diffusion-weighting.

Simulation experiments

The results of the simulation experiments are presented in Figure 7.7. The ground-truth μFA was 0.34 for f_{intra} = 0.2, 0.59 for f_{intra} = 0.6, and 0.97 for f_{intra} = 1.

In Simulation 1, all methods overestimated μFA and the bias was greater at lower values of ground-truth μFA (Figure 7.7A-C). In Simulation 2, the application of the TDE waveform whose power spectrum along three orthogonal directions matches that of the linear waveform reduced the overestimation of μFA for all methods (Figure 7.7D-F). In Simulation 3, the application of the TDE waveforms with short pulses, long diffusion times, and long mixing times further reduced the bias (Figure 7.7G-I). These results show that both the discrepancy between the diffusion times of the waveforms encoding different b-tensor shapes, and the correlations between the spins' positions during the diffusion encoding along different directions can result in overestimation of μFA. However, the simulations suggest that the bias is small in white matter, where μFA is high. The average positive bias caused by time-dependent diffusion, i.e., the difference between the expected μFA in simulations 1 and 3, was 0.11 for f_{intra}

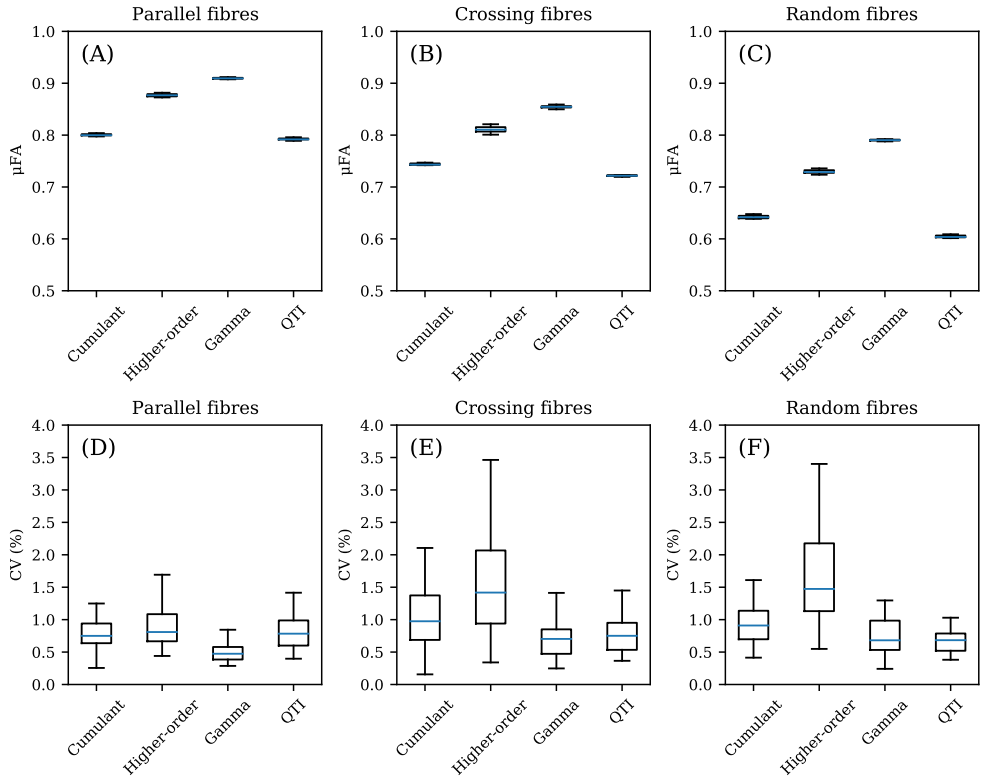

Figure 7.6: *The results of the phantom experiments. (A-C) μFA values averaged over the phantom ROIs from the repeated acquisitions. (D-F) The distributions of the voxel-level coefficients of variation. The blue line denotes the median. The box spans the interquartile range. The whiskers extend to the furthest data point within 1.5 · IQR from the 1st or the 3rd quartile.*

= 0.2, 0.09 for f_{intra} = 0.6, and 0.04 for f_{intra} = 1.

The accuracy of the signal models for μFA estimation was assessed using the results from Simulation 3 which are not confounded by the bias induced by time-dependent diffusion. The mean squared error (MSE) between the expected value of μFA and the ground-truth value was $1.8 \cdot 10^{-3}$ for the cumulant model, $5.4 \cdot 10^{-3}$ for the higher-order model, $12.7 \cdot 10^{-3}$ for the gamma model and $1.6 \cdot 10^{-3}$ for QTI. In terms of precision, the average CV of μFA was 6.6% for the cumulant model, 10.6% for the higher-order model, 6.6% for the gamma model, and 8.5% for QTI. However, μFA is heteroscedastic and the precision of each method strongly depended on the ground-truth μFA.

The higher-order model is motivated by the b-value-dependent underestimation of μFA using the cumulant model. Indeed, the higher-order model produced more accurate estimates of μFA when the cumulant model under-

estimated μFA. However, because P_3 was constrained to be non-negative, the higher-order model could not produce lower values of μFA and thus it provided even more inaccurate estimates in situations where the cumulant model overestimated μFA. Allowing P_3 to take negative values improved the accuracy (MSE = $2.0 \cdot 10^{-3}$.) but greatly destabilized the fit (mean CV = 24%). In most cases, at the b-values used in the simulations, the underestimation due to the invalidity of the cumulant expansion of the powder-averaged signal was not strong enough to cancel the overestimation.

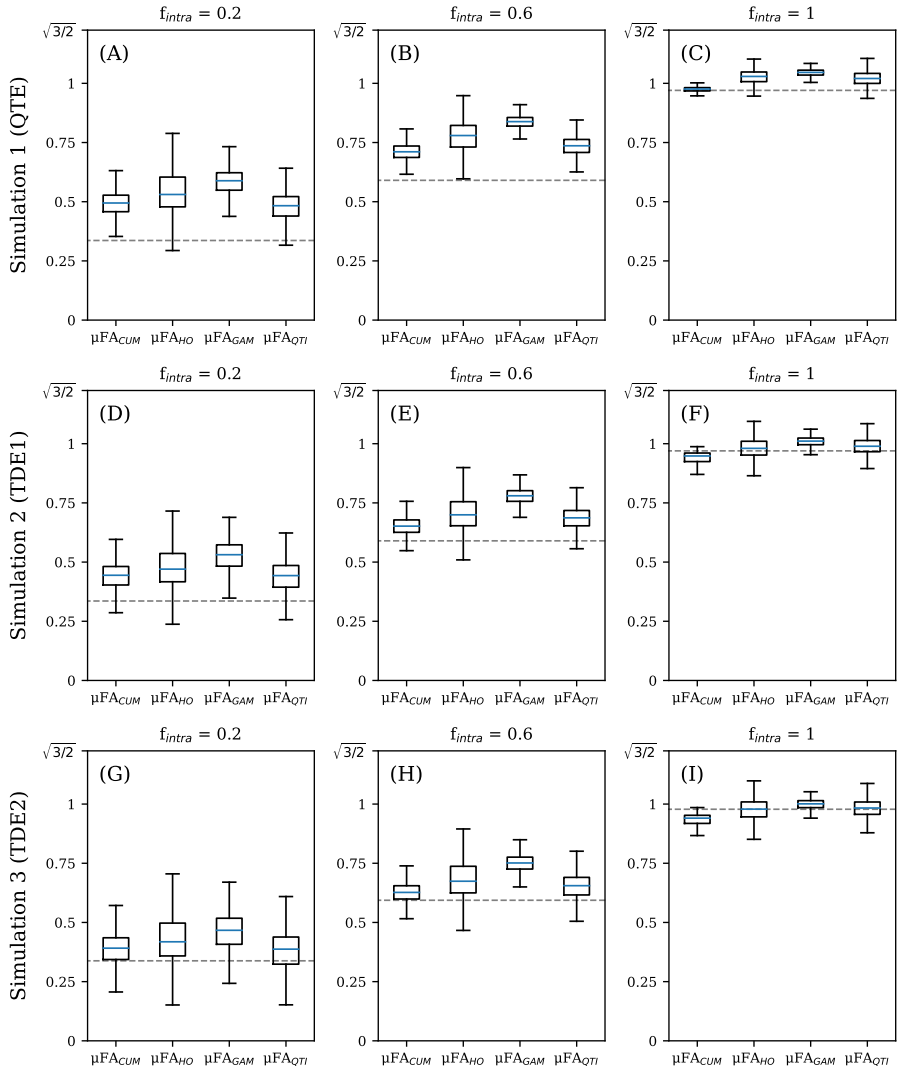

Figure 7.7: *The distributions of the μFA estimates calculated from the simulated data after* 10^4 *repetitions of Rician noise addition to lower SNR to 25 using the cumulant model (*μFA_{CUM}*), higher order model (*μFA_{HO}*), gamma model (*μFA_{GAM}*), and QTI (*μFA_{QTI}*).* f_{intra} *is the signal fraction of the intra-axonal space. (A)-(C) Results of Simulation 1. (D)-(F) Results of Simulation 2. (G)-(I) Results of Simulation 3. The dashed line denotes the ground-truth μFA. The blue line denotes the median. The box spans the interquartile range (IQR). The whiskers extend to the furthest data point within 1.5 · IQR from the 1st or the 3rd quartile.*

7.5 Discussion

Over recent years, several microscopic anisotropy estimation methods have been introduced [134,135,65,216,204,125], yet the differences between the methods have received little attention thus far. The purpose of this study was to assess the accuracy and precision of μFA estimates calculated using different signal models. μFA was chosen as the metric of interest for its connection to conventional FA that is well known in the neuroscience community. The cumulant model, gamma model, and QTI were chosen as the signal models of interest for their numerous recent applications in neuroimaging studies [64,55,65,56,60,59,63]. The higher order model was included in the study for it being a simple extension to the cumulant model that has been shown to improve the accuracy of μFA estimation in animal studies using DDE [204]. Until now, a study by Reymbaut et al. [217] is the only one to have explicitly compared μFA estimation methods. However, they focused on computationally sampling the parameter space by modelling diffusion as a distribution of microscopic diffusion tensors using Equation 7.2.1. In another relevant study, Martin et al. [218] estimated the covariance of the microscopic diffusion tensor distribution using QTI and investigated how μFA relates to other microscopic anisotropy indices in terms of contrast-to-noise ratio. In this study, four μFA estimation methods were applied on MRI data acquired using a clinical whole-body scanner and on synthetic data generated by simulating time-dependent diffusion.

In this study, the gamma model suffered from the greatest bias, consistently producing greater values than the other methods in the volunteer experiments, phantom experiments, and simulations. The disagereement between the gamma model and the other methods can be understood by the fact that the gamma model assumes the distribution of apparent diffusivities to follow a gamma distribution whereas the other methods are all based on the truncated cumulant expansion. However, Reymbaut et al. [217] have reported that the gamma model produces biased estimates when the signal strongly deviates from a monoexponential decay, as is the case in white matter, and that QTI results in biased estimates in the case of significant size variance of microsocpic diffusion environments. Therefore, given the microstructure-dependent bias, the results presented here can not be generalized to all situations.

In terms of precision, the higher order model stood out by being the most sensitive to noise whereas the other three methods were more similar. This is most likely explained by the larger number of fit parameters compared to data

points and the third order polynomial fit's sensitivity to signal variation. In the phantom experiments, the voxel-specific CV of μFA from all methods was low with the majority CV values being well below 2%, indicating very good repeatability in voxels with high microscopic diffusion anisotropy. Since the phantom contained highly anisotropic microfibres, this result was expected based on previous reports of how the precision of μFA is high in voxels with high anisotropy but quickly diminishes with a decreasing value of μFA[135,64,202,219]. Because of this property of μFA and the possible overestimation of μFA by the gamma model, it may be more sensitive to microstructural changes than the other methods in some situations.

The simulation experiments showed that the combination of trapezoidal linear gradient waveforms and isotropic encoding with QTE can lead to a significant overestimation of μFA, a result that has been discussed in detail by Lundell et al.[210]. As a solution, they proposed designing the linear waveform so that its power spectrum matches that of the spherical waveform at least along one direction. Indeed, in the simulations, the accuracy of μFA estimates was improved when the power spectra were matched. However, the simulations also showed that this is not sufficient for accurate estimation of μFA. The accuracy of μFA further improved when a TDE waveform with short pulses, long diffusion time, and long mixing time was applied, eliminating the correlations between the positions of the random walkers during the applications of the gradient pulse pairs, an important assumption in the theory of μFA estimation with DDE[134]. However, the simulated TDE waveforms are completely unrealistic in today's human neuroimaging experiments because of their duration and gradient strength requirements. Considering that the bias in μFA induced by time-dependent diffusion was small (< 0.1), and the effects of time-dependent diffusion are easier to observe in Monte Carlo random walk simulations than in the brain[220], these results support using QTE for μFA estimation in the human brain.

This study has several limitations. Most importantly, the accuracy and precision of μFA estimation depend on the details of the acquisition protocol which was kept constant here. For instance, the effect of b-values and diffusion encoding directions was outside the scope of this study. Also, the problem of designing optimal gradient waveforms for MDE experiments, discussed in detail by Szczepankiewicz et al.[221], was outside the scope of this study. Furthermore, the acquisitions with the lowest b-value may be affected by perfusion

effects that may cause parameter bias. However, the focus of this study was to assess the performance of the signal models relative to each other when applied on the same data. Furthermore, the microfibre phantom used for studying the precision of μFA contains fibres that are too large to be representative of white matter microstructure. Finally, in the simulations, the effects of orientation dispersion, distribution of axon radii, and packing density were not assessed because the simulated geometry was kept constant across the simulations. Nevertheless, the simulated fibre geometry, generated using the modern algorithm by Callaghan et al.[168], resemble real axons and enabled the μFA estimates to be compared to the ground truth. Future studies should address the mentioned issues to facilitate the optimal use of μFA.

7.6 Conclusion

This study demonstrated that the different signal models yield significantly different μFA estimates when calculated from the same data. Considering that QTI provided the most accurate estimates and also enables several other microstructural paremeters to be estimated, this study supports using it over the other included methods. Furthermore, the simulations showed that matching the power spectrum of the linear waveform to the power spectrum of a component of the spherical waveform is not sufficient for accurate estimation of μFA. Nevertheless, the simulations suggest that the bias in μFA caused by time-dependent diffusion is small in human white matter.

Chapter 8

Conclusions

The aim of this thesis was to study the state-of-the-art dMRI methods for estimating microscopic diffusion anisotropy in the human brain and to assess their robustness to facilitate their adoption in the wider neuroscience research community and possible deployment in future clinical studies. The thesis consists of a comprehensive review of the relevant literature followed by four independent research projects.

Chapter 4 presented Disimpy, a massively parallel dMRI simulator and a contribution to the growing library of scientific Python packages. The simulator takes advantage of the recent developments in both the software and hardware of general-purpose GPU computing to enable large amounts of synthetic signals to be generated on standard desktop and laptop computers without needing to access high-performance computing clusters. The simulator is written in Python, a programming language known for its readibility and clarity, making it very approachable and easily extensible to researchers and students with little or no prior experience in massively parallel computing. Although it was sufficient for the study presented in Chapter 7, the implementation of some of the algorithms in Disimpy is suboptimal in terms of computational complexity and memory management. The long-term vision is to continue the development of Disimpy by implementing more features, such as flow, and optimizing performance while it is used in future studies.

In Chapter 5, the reproducibility of three commonly used microstructural models for estimating μFA or intra-neurite volume fraction, a microstructural metric closely related to μFA, was quantified and the shortcomings of such SDE methods in quantifying microscopic diffusion anisotropy were discussed. The results of the study show that despite the accuracy issues highlighted in

the theory section, the reproducibility of most of the microstructural metrics calculated using both NODDI and SMT is comparable to the reproducibility of DTI and DKI parameters. Therefore, NODDI and SMT may be valuable tools in situations where they offer the greatest sensitivity to relevant differences in the brain microstructure across groups or over time. However, due to the invalid model assumptions, the resulting parameter estimates are biased and thus the parameter names should not be taken literally.

In Chapter 6, the challenges and limitations in estimating microscopic diffusion anisotropy using DDE in the human brain were discussed and a clinically feasible acquisition protocol was validated using animal imaging experiments. Furthermore, the precision of the resulting μFA estimates was quantified using simulations. The results of the study support using DDE in human neuroimaging but raise hitherto overlooked precision issues when measuring low values of μFA using DDE. Finally, although preliminary DDE experiment using the proposed clinically feasible protocol was performed with a clinical whole-body scanner, unfortunately, it was not possible to acquire more scans as a consequence of the Covid-19 pandemic. An interesting direction for future research would be to optimize the used acquisition protocol, minimize the signal bias caused by concomitant gradients[208], and apply the recently developed correlation tensor imaging (CTI) in the human brain[125]. CTI is a recently developed method that enables intra-compartmental kurtosis to be estimated from DDE data. Furthermore, a comprehensive comparison between DDE and QTE should be performed to understand which diffusion encoding method is more suitable in a given situation.

In Chapter 7, μFA estimation using QTE, an MDE method specifically developed with the hardware and time limitations of human neuroimaging in mind, was discussed and four previously introduced signal models were compared in terms of accuracy and precision. It is worth mentioning that the signal models discussed in Chapter 7 can be applied to estimate μFA from data acquired with both DDE and QTE. The results show that the different methods result in significantly different values of μFA when calculated from the same data. Furthermore, the simulations of time-dependent diffusion showed that all the studied methods may overestimate μFA if there is a discrepancy between the diffusion times of the gradient waveforms encoding different b-tensor shapes or if there is significant time-dependent microscopic diffusion anisotropy. Nevertheless, the results suggest that the resulting bias in μFA is less than 0.1 in

human white matter. Considering that QTI provided the most accurate μFA estimates and also enables several other microstructural parameters to be calculated from the estimated covariance tensor, this study supports using it instead of the other included signal models. An interesting direction for future research would be to design optimal gradient waveforms for minimizing the parameter bias and maximizing the sensitivity to the relevant microstructural length scales. The development of more realistic and diverse simulated tissue microstructures would also be useful in validating the signal models. Disimpy could be used in these studies to simulate the signal with which the models can be validated. Furthermore, since QTI is based on a standard weighted least squares fit, future studies should optimize it by applying more advanced methods.

Overall, the work presented in this thesis shows that microscopic diffusion anisotropy can be robustly estimated in white matter of the human brain using multidimensional diffusion encoding to solve the fundamental degeneracy in data acquired with single diffusion encoding. However, the current methods lack the precision to reliably map microscopic diffusion anisotropy in gray matter or other brain regions with low microscopic diffusion anisotropy. As discussed in Section 3.7, the measurement of microscopic diffusion anisotropy without significant *a priori* information about tissue microstructure requires combining acquisitions with at least two different b-tensor shapes. DDE waveforms can be used to generate linear and planar b-tensors with well-defined diffusion times and thus the data can be analyzed without assuming the effects of time-dependent diffusion to be negligible. QTE waveforms, on the other hand, enable more efficient diffusion encoding for arbitrary b-tensor shapes, although the resulting waveforms lack a well-defined scalar diffusion time, a factor that must be taken into account when analyzing the data. The choice between DDE and QTE waveforms depends on the application and the available scanner hardware. It is important that future studies investigate the optimal choice of waveform type for a given scenario.

In conclusion, since μFA is not affected by the orientation dispersion of axons and its precision is high in white matter, μFA can be very useful in measuring white matter integrity and axonal degeneration associated with a variety of neurological conditions.

Bibliography

[1] F. A. Azevedo, L. R. Carvalho, L. T. Grinberg, J. M. Farfel, R. E. Ferretti, R. E. Leite, W. J. Filho, R. Lent, and S. Herculano-Houzel, "Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain," *Journal of Comparative Neurology*, vol. 513, no. 5, pp. 532–541, 2009.

[2] E. R. Kandel, J. H. Schwartz, T. M. Jessell, D. of Biochemistry, M. B. T. Jessell, S. Siegelbaum, and A. Hudspeth, *Principles of neural science*, vol. 4. McGraw-hill New York, 2000.

[3] P. Broca, "Remarques sur le siège de la faculté du langage articulé, suivies d'une observation d'aphémie (perte de la parole)," *Bulletin et Memoires de la Societe anatomique de Paris*, vol. 6, pp. 330–357, 1861.

[4] C. Wernicke, *Der aphasische Symptomencomplex: eine psychologische Studie auf anatomischer Basis*. Cohn., 1874.

[5] L. Lichtheim, "On aphasia," *Brain*, vol. 7, pp. 433–484, 1885.

[6] C. E. Sexton, K. B. Walhovd, A. B. Storsve, C. K. Tamnes, L. T. Westlye, H. Johansen-Berg, and A. M. Fjell, "Accelerated changes in white matter microstructure during aging: a longitudinal diffusion tensor imaging study," *Journal of Neuroscience*, vol. 34, no. 46, pp. 15425–15436, 2014.

[7] S. Khoury, C. Guttmann, E. Orav, M. Hohol, S. Ahn, L. Hsu, R. Kikinis, G. Mackin, F. Jolesz, and H. Weiner, "Longitudinal mri in multiple sclerosis: correlation between disability and lesion burden," *Neurology*, vol. 44, no. 11, pp. 2120–2120, 1994.

[8] A. G. Thomas, S. Marrett, Z. S. Saad, D. A. Ruff, A. Martin, and P. A. Bandettini, "Functional but not structural changes associated with learning:

an exploration of longitudinal voxel-based morphometry (vbm)," *Neuroimage*, vol. 48, no. 1, pp. 117–125, 2009.

[9] P. C. Lauterbur, "Image formation by induced local interactions: examples employing nuclear magnetic resonance," *nature*, vol. 242, no. 5394, pp. 190–191, 1973.

[10] P. Mansfield, "Multi-planar image formation using nmr spin echoes," *Journal of Physics C: Solid State Physics*, vol. 10, no. 3, p. L55, 1977.

[11] W. Nitz and P. Reimer, "Contrast mechanisms in mr imaging," *European radiology*, vol. 9, no. 6, pp. 1032–1046, 1999.

[12] E. L. Hahn, "Spin echoes," *Physical review*, vol. 80, no. 4, p. 580, 1950.

[13] H. Y. Carr and E. M. Purcell, "Effects of diffusion on free precession in nuclear magnetic resonance experiments," *Physical review*, vol. 94, no. 3, p. 630, 1954.

[14] E. O. Stejskal and J. E. Tanner, "Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient," *The journal of chemical physics*, vol. 42, no. 1, pp. 288–292, 1965.

[15] D. Le Bihan and E. Breton, "Imagerie de diffusion in vivo par résonance magnétique nucléaire," *Comptes rendus de l'Académie des sciences. Série 2, Mécanique, Physique, Chimie, Sciences de l'univers, Sciences de la Terre*, vol. 301, no. 15, pp. 1109–1112, 1985.

[16] K.-D. Merboldt, W. Hanicke, and J. Frahm, "Self-diffusion nmr imaging using stimulated echoes," *Journal of Magnetic Resonance (1969)*, vol. 64, no. 3, pp. 479–486, 1985.

[17] D. Taylor and M. Bushell, "The spatial mapping of translational diffusion coefficients by the nmr imaging technique," *Physics in Medicine & Biology*, vol. 30, no. 4, p. 345, 1985.

[18] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet, "Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders.," *Radiology*, vol. 161, no. 2, pp. 401–407, 1986.

[19] H. C. Berg, *Random walks in biology*. Princeton University Press, 1993.

[20] H. Johansen-Berg and T. E. Behrens, *Diffusion MRI: from quantitative measurement to in vivo neuroanatomy*. Academic Press, 2013.

[21] M. Moseley, J. Mintorovitch, Y. Cohen, H. Asgari, N. Derugin, D. Norman, and J. Kucharczyk, "Early detection of ischemic injury: comparison of spectroscopy, diffusion-, t2-, and magnetic susceptibility-weighted mri in cats," in *Brain Edema VIII*, pp. 207–209, Springer, 1990.

[22] S. Warach, D. Chien, W. Li, M. Ronthal, and R. Edelman, "Fast magnetic resonance diffusion-weighted imaging of acute human stroke," *Neurology*, vol. 42, no. 9, pp. 1717–1717, 1992.

[23] S. Warach, J. Gaa, B. Siewert, P. Wielopolski, and R. R. Edelman, "Acute human stroke studied by whole brain echo planar diffusion-weighted magnetic resonance imaging," *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, vol. 37, no. 2, pp. 231–241, 1995.

[24] T. L. Chenevert, J. A. Brunberg, and J. G. Pipe, "Anisotropic diffusion in human white matter: demonstration with mr techniques in vivo.," *Radiology*, vol. 177, no. 2, pp. 401–405, 1990.

[25] P. J. Basser, J. Mattiello, and D. LeBihan, "Mr diffusion tensor spectroscopy and imaging," *Biophysical journal*, vol. 66, no. 1, pp. 259–267, 1994.

[26] S.-K. Song, S.-W. Sun, M. J. Ramsbottom, C. Chang, J. Russell, and A. H. Cross, "Dysmyelination revealed through mri as increased radial (but unchanged axial) diffusion of water," *Neuroimage*, vol. 17, no. 3, pp. 1429–1436, 2002.

[27] M. Rovaris and M. Filippi, "Diffusion tensor mri in multiple sclerosis," *Journal of Neuroimaging*, vol. 17, pp. 27S–30S, 2007.

[28] C. J. Cochrane and K. P. Ebmeier, "Diffusion tensor imaging in parkinsonian syndromes: a systematic review and meta-analysis," *Neurology*, vol. 80, no. 9, pp. 857–864, 2013.

[29] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi, "In vivo fiber tractography using dt-mri data," *Magnetic resonance in medicine*, vol. 44, no. 4, pp. 625–632, 2000.

[30] T. E. Behrens, H. J. Berg, S. Jbabdi, M. F. Rushworth, and M. W. Woolrich, "Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?," *Neuroimage*, vol. 34, no. 1, pp. 144–155, 2007.

[31] R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly, "Anatomically-constrained tractography: improved diffusion mri streamlines tractography through effective use of anatomical information," *Neuroimage*, vol. 62, no. 3, pp. 1924–1938, 2012.

[32] B. Jeurissen, J.-D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers, "Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion mri data," *NeuroImage*, vol. 103, pp. 411–426, 2014.

[33] O. Ciccarelli, M. Catani, H. Johansen-Berg, C. Clark, and A. Thompson, "Diffusion-based tractography in neurological disorders: concepts, applications, and future developments," *The Lancet Neurology*, vol. 7, no. 8, pp. 715–727, 2008.

[34] M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones, "Virtual in vivo interactive dissection of white matter fasciculi in the human brain," *Neuroimage*, vol. 17, no. 1, pp. 77–94, 2002.

[35] D. K. Jones, A. Simmons, S. C. Williams, and M. A. Horsfield, "Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 42, no. 1, pp. 37–41, 1999.

[36] D. K. Jones, "Challenges and limitations of quantifying brain connectivity in vivo with diffusion mri," *Imaging in Medicine*, vol. 2, no. 3, pp. 341–355, 2010.

[37] C. A. Clark, T. R. Barrick, M. M. Murphy, and B. A. Bell, "White matter fiber tracking in patients with space-occupying lesions of the brain: a new technique for neurosurgical planning?," *Neuroimage*, vol. 20, no. 3, pp. 1601–1608, 2003.

[38] J. Berman, "Diffusion mr tractography as a tool for surgical planning," *Magnetic resonance imaging clinics of North America*, vol. 17, no. 2, pp. 205–214, 2009.

[39] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly, "Direct estimation of the fiber orientation density function from diffusion-weighted mri data using spherical deconvolution," *NeuroImage*, vol. 23, no. 3, pp. 1176–1185, 2004.

[40] J.-D. Tournier, F. Calamante, and A. Connelly, "Robust determination of the fibre orientation distribution in diffusion mri: non-negativity constrained super-resolved spherical deconvolution," *Neuroimage*, vol. 35, no. 4, pp. 1459–1472, 2007.

[41] J.-D. Tournier, C.-H. Yeh, F. Calamante, K.-H. Cho, A. Connelly, and C.-P. Lin, "Resolving crossing fibres using constrained spherical deconvolution: validation using diffusion-weighted imaging phantom data," *Neuroimage*, vol. 42, no. 2, pp. 617–625, 2008.

[42] D. S. Tuch, "Q-ball imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 52, no. 6, pp. 1358–1372, 2004.

[43] V. J. Wedeen, R. Wang, J. D. Schmahmann, T. Benner, W.-Y. I. Tseng, G. Dai, D. Pandya, P. Hagmann, H. D'Arceuil, and A. J. de Crespigny, "Diffusion spectrum magnetic resonance imaging (dsi) tractography of crossing fibers," *Neuroimage*, vol. 41, no. 4, pp. 1267–1277, 2008.

[44] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen, "High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 48, no. 4, pp. 577–582, 2002.

[45] C. A. Clark, M. Hedehus, and M. E. Moseley, "In vivo mapping of the fast and slow diffusion tensors in human brain," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 47, no. 4, pp. 623–628, 2002.

[46] J. H. Jensen, J. A. Helpern, A. Ramani, H. Lu, and K. Kaczynski, "Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 53, no. 6, pp. 1432–1440, 2005.

[47] D. S. Novikov and V. G. Kiselev, "Effective medium theory of a diffusion-weighted signal," *NMR in Biomedicine*, vol. 23, no. 7, pp. 682–697, 2010.

[48] 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*, vol. 37, no. 1, pp. 103–111, 1997.

[49] S. N. Jespersen, C. D. Kroenke, L. Østergaard, J. J. Ackerman, and D. A. Yablonskiy, "Modeling dendrite density from magnetic resonance diffusion measurements," *Neuroimage*, vol. 34, no. 4, pp. 1473–1486, 2007.

[50] 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: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 59, no. 6, pp. 1347–1354, 2008.

[51] 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*, vol. 52, no. 4, pp. 1374–1389, 2010.

[52] N. Shemesh, S. N. Jespersen, D. C. Alexander, Y. Cohen, I. Drobnjak, T. B. Dyrby, J. Finsterbusch, M. A. Koch, T. Kuder, F. Laun, *et al.*, "Conventions and nomenclature for double diffusion encoding nmr and mri," *Magnetic resonance in medicine*, vol. 75, no. 1, pp. 82–87, 2016.

[53] E. Özarslan, "Compartment shape anisotropy (csa) revealed by double pulsed field gradient mr," *Journal of Magnetic Resonance*, vol. 199, no. 1, pp. 56–67, 2009.

[54] S. Eriksson, S. Lasic, and D. Topgaard, "Isotropic diffusion weighting in pgse nmr by magic-angle spinning of the q-vector," *Journal of Magnetic Resonance*, vol. 226, pp. 13–18, 2013.

[55] F. Szczepankiewicz, D. van Westen, E. Englund, C.-F. Westin, F. Ståhlberg, J. Lätt, P. C. Sundgren, and M. Nilsson, "The link between diffusion mri and tumor heterogeneity: mapping cell eccentricity and density by diffusional variance decomposition (divide)," *NeuroImage*, vol. 142, pp. 522–532, 2016.

[56] G. Yang, Q. Tian, C. Leuze, M. Wintermark, and J. A. McNab, "Double diffusion encoding mri for the clinic," *Magnetic resonance in medicine*, vol. 80, no. 2, pp. 507–520, 2018.

[57] F. Nery, F. Szczepankiewicz, L. Kerkelä, M. G. Hall, E. Kaden, I. Gordon, D. L. Thomas, and C. A. Clark, "In vivo demonstration of microscopic anisotropy in the human kidney using multidimensional diffusion mri," *Magnetic resonance in medicine*, vol. 82, no. 6, pp. 2160–2168, 2019.

[58] S. Coelho, J. M. Pozo, S. N. Jespersen, D. K. Jones, and A. F. Frangi, "Resolving degeneracy in diffusion mri biophysical model parameter estimation using double diffusion encoding," *Magnetic resonance in medicine*, vol. 82, no. 1, pp. 395–410, 2019.

[59] B. Lampinen, A. Zampeli, I. M. Björkman-Burtscher, F. Szczepankiewicz, K. Källén, M. Compagno Strandberg, and M. Nilsson, "Tensor-valued diffusion mri differentiates cortex and white matter in malformations of cortical development associated with epilepsy," *Epilepsia*, vol. 61, no. 8, pp. 1701–1713, 2020.

[60] K. W. Andersen, S. Lasič, H. Lundell, M. Nilsson, D. Topgaard, F. Sellebjerg, F. Szczepankiewicz, H. R. Siebner, M. Blinkenberg, and T. B. Dyrby, "Disentangling white-matter damage from physiological fiber orientation dispersion in multiple sclerosis," *Brain Communications*, 2020.

[61] S. Lasič, F. Szczepankiewicz, E. Dall'Armellina, A. Das, C. Kelly, S. Plein, J. E. Schneider, M. Nilsson, and I. Teh, "Motion-compensated b-tensor encoding for in vivo cardiac diffusion-weighted imaging," *NMR in Biomedicine*, vol. 33, no. 2, p. e4213, 2020.

[62] D. S. Novikov, V. G. Kiselev, and S. N. Jespersen, "On modeling," *Magnetic resonance in medicine*, vol. 79, no. 6, pp. 3172–3193, 2018.

[63] K. Kamiya, K. Kamagata, K. Ogaki, T. Hatano, T. Ogawa, H. Takeshige-Amano, S. Murata, C. Andica, K. Murata, T. Feiweier, *et al.*, "Brain white-matter degeneration due to aging and parkinson disease as revealed by double diffusion encoding," *Frontiers in Neuroscience*, vol. 14, p. 1091, 2020.

[64] F. Szczepankiewicz, S. Lasič, D. van Westen, P. C. Sundgren, E. Englund, C.-F. Westin, F. Ståhlberg, J. Lätt, D. Topgaard, and M. Nilsson, "Quantification of microscopic diffusion anisotropy disentangles effects of orientation dispersion from microstructure: applications in healthy volunteers and in brain tumors," *NeuroImage*, vol. 104, pp. 241–252, 2015.

[65] C.-F. Westin, H. Knutsson, O. Pasternak, F. Szczepankiewicz, E. Özarslan, D. van Westen, C. Mattisson, M. Bogren, L. J. O'Donnell, M. Kubicki, *et al.*, "Q-space trajectory imaging for multidimensional diffusion mri of the human brain," *Neuroimage*, vol. 135, pp. 345–362, 2016.

[66] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander, "Noddi: practical in vivo neurite orientation dispersion and density imaging of the human brain," *Neuroimage*, vol. 61, no. 4, pp. 1000–1016, 2012.

[67] E. Kaden, F. Kruggel, and D. C. Alexander, "Quantitative mapping of the per-axon diffusion coefficients in brain white matter," *Magnetic resonance in medicine*, vol. 75, no. 4, pp. 1752–1763, 2016.

[68] E. Kaden, N. D. Kelm, R. P. Carson, M. D. Does, and D. C. Alexander, "Multi-compartment microscopic diffusion imaging," *NeuroImage*, vol. 139, pp. 346–359, 2016.

[69] P. A. M. Dirac, *The principles of quantum mechanics*. No. 27, Oxford university press, 1981.

[70] D. J. Griffiths and D. F. Schroeter, *Introduction to quantum mechanics*. Cambridge University Press, 2018.

[71] M. H. Levitt, *Spin dynamics: basics of nuclear magnetic resonance*. John Wiley & Sons, 2001.

[72] J. E. Baggott, J. Baggott, and E. O. L. M. J. Baggott, *Beyond measure: Modern physics, philosophy, and the meaning of quantum theory*. Oxford University Press on Demand, 2004.

[73] A. Hopkins, H. Yeung, and C. Bratton, "Multiple field strength in vivo t1 and t2 for cerebrospinal fluid protons," *Magnetic resonance in medicine*, vol. 3, no. 2, pp. 303–311, 1986.

[74] W. D. Rooney, G. Johnson, X. Li, E. R. Cohen, S.-G. Kim, K. Ugurbil, and C. S. Springer Jr, "Magnetic field and tissue dependencies of human brain longitudinal 1h2o relaxation in vivo," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 57, no. 2, pp. 308–318, 2007.

[75] F. Bloch, "Nuclear induction," *Physical review*, vol. 70, no. 7-8, p. 460, 1946.

[76] J. P. Wansapura, S. K. Holland, R. S. Dunn, and W. S. Ball Jr, "Nmr relaxation times in the human brain at 3.0 tesla," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 9, no. 4, pp. 531–538, 1999.

[77] D. Purves, G. J. Augustine, D. Fitzpatrick, W. C. Hall, A.-S. LaMantia, J. O. McNamara, and S. M. Williams, "Neuroscience," *Sunderland, MA: Sinauer Associates*, vol. 773, 2004.

[78] J. D. Schmahmann and D. Caplan, "Cognition, emotion and the cerebellum," *Brain*, vol. 129, no. 2, pp. 290–292, 2006.

[79] S. S.-H. Wang, J. R. Shultz, M. J. Burish, K. H. Harrison, P. R. Hof, L. C. Towns, M. W. Wagers, and K. D. Wyatt, "Functional trade-offs in white matter axonal scaling," *Journal of neuroscience*, vol. 28, no. 15, pp. 4047–4056, 2008.

[80] T. Paus, "Imaging microstructure in the living human brain: a viewpoint," *NeuroImage*, vol. 182, pp. 3–7, 2018.

[81] A. Schüz and V. Braitenberg, "The human cortical white matter: quantitative aspects of cortico-cortical long-range connectivity," *Cortical areas: Unity and diversity*, pp. 377–385, 2002.

[82] C. Nicholson and S. Hrabětová, "Brain extracellular space: the final frontier of neuroscience," *Biophysical journal*, vol. 113, no. 10, pp. 2133–2142, 2017.

[83] V. A. Shepherd, "The cytomatrix as a cooperative system of macromolecular and water networks," *Current topics in developmental biology*, vol. 75, pp. 171–223, 2006.

[84] B. Barbara, "Aquaporin biology and nervous system," *Current neuropharmacology*, vol. 8, no. 2, pp. 97–104, 2010.

[85] R. Druga, "Neocortical inhibitory system," *Folia Biol (Praha)*, vol. 55, no. 6, pp. 201–17, 2009.

[86] D. P. Pelvig, H. Pakkenberg, A. K. Stark, and B. Pakkenberg, "Neocortical glial cell numbers in human brains," *Neurobiology of aging*, vol. 29, no. 11, pp. 1754–1762, 2008.

[87] J. C. Fiala and K. M. Harris, "Dendrite structure," *Dendrites*, vol. 2, pp. 1–11, 1999.

[88] D. C. Alexander, T. B. Dyrby, M. Nilsson, and H. Zhang, "Imaging brain microstructure with diffusion mri: practicality and applications," *NMR in Biomedicine*, vol. 32, no. 4, p. e3841, 2019.

[89] Y. Bakiri, R. Káradóttir, L. Cossell, and D. Attwell, "Morphological and electrical properties of oligodendrocytes in the white matter of the corpus callosum and cerebellum," *The Journal of physiology*, vol. 589, no. 3, pp. 559–573, 2011.

[90] N. A. Oberheim, T. Takano, X. Han, W. He, J. H. Lin, F. Wang, Q. Xu, J. D. Wyatt, W. Pilcher, J. G. Ojemann, *et al.*, "Uniquely hominid features of adult human astrocytes," *Journal of Neuroscience*, vol. 29, no. 10, pp. 3276–3287, 2009.

[91] S. Waxman and M. V. Bennett, "Relative conduction velocities of small myelinated and non-myelinated fibres in the central nervous system," *Nature New Biology*, vol. 238, no. 85, pp. 217–219, 1972.

[92] J. A. Perge, K. Koch, R. Miller, P. Sterling, and V. Balasubramanian, "How the optic nerve allocates space, energy capacity, and information," *Journal of Neuroscience*, vol. 29, no. 24, pp. 7917–7928, 2009.

[93] R. S. Smith and Z. J. Koles, "Myelinated nerve fibers: computed effect of myelin thickness on conduction velocity," *American Journal of Physiology-Legacy Content*, vol. 219, no. 5, pp. 1256–1258, 1970.

[94] C. Lucchinetti, W. Brück, J. Parisi, B. Scheithauer, M. Rodriguez, and H. Lassmann, "Heterogeneity of multiple sclerosis lesions: implications

for the pathogenesis of demyelination," *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, vol. 47, no. 6, pp. 707–717, 2000.

[95] B. Weber, A. L. Keller, J. Reichold, and N. K. Logothetis, "The microvascular system of the striate and extrastriate visual cortex of the macaque," *Cerebral cortex*, vol. 18, no. 10, pp. 2318–2330, 2008.

[96] D. Le Bihan, E. Breton, D. Lallemand, M. Aubin, J. Vignaud, and M. Laval-Jeantet, "Separation of diffusion and perfusion in intravoxel incoherent motion mr imaging.," *Radiology*, vol. 168, no. 2, pp. 497–505, 1988.

[97] R. Brown, "Xxvii. a brief account of microscopical observations made in the months of june, july and august 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies," *The Philosophical Magazine*, vol. 4, no. 21, pp. 161–173, 1828.

[98] A. Fick, "Ueber diffusion," *Annalen der Physik*, vol. 170, no. 1, pp. 59–86, 1855.

[99] A. Einstein, "Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen," *Annalen der physik*, vol. 322, no. 8, pp. 549–560, 1905.

[100] V. G. Kiselev, "Fundamentals of diffusion mri physics," *NMR in Biomedicine*, vol. 30, no. 3, p. e3602, 2017.

[101] N. Bloembergen, E. M. Purcell, and R. V. Pound, "Relaxation effects in nuclear magnetic resonance absorption," *Physical review*, vol. 73, no. 7, p. 679, 1948.

[102] S. De Santis, A. Gabrielli, M. Palombo, B. Maraviglia, and S. Capuani, "Non-gaussian diffusion imaging: a brief practical review," *Magnetic resonance imaging*, vol. 29, no. 10, pp. 1410–1416, 2011.

[103] J. Stepišnik, "Time-dependent self-diffusion by nmr spin-echo," *Physica B: Condensed Matter*, vol. 183, no. 4, pp. 343–350, 1993.

[104] P. T. Callaghan, "Physics of diffusion," *Diffusion MRI: Theory, Methods and Applications*, pp. 45–56, 2010.

[105] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*, vol. 1. Elsevier, 2001.

[106] C. A. Baron and C. Beaulieu, "Oscillating gradient spin-echo (ogse) diffusion tensor imaging of the human brain," *Magnetic resonance in medicine*, vol. 72, no. 3, pp. 726–736, 2014.

[107] D. B. Le, "Molecular diffusion nuclear magnetic resonance imaging.," *Magnetic resonance quarterly*, vol. 7, no. 1, pp. 1–30, 1991.

[108] P. T. Callaghan, D. MacGowan, K. Packer, and F. Zelaya, "High-resolution q-space imaging in porous structures," *Journal of Magnetic Resonance (1969)*, vol. 90, no. 1, pp. 177–182, 1990.

[109] M. D. King, J. Houseman, S. A. Roussel, N. Van Bruggen, S. R. Williams, and D. G. Gadian, "q-space imaging of the brain," *Magnetic resonance in medicine*, vol. 32, no. 6, pp. 707–713, 1994.

[110] V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff, "Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging," *Magnetic resonance in medicine*, vol. 54, no. 6, pp. 1377–1386, 2005.

[111] P. P. Mitra and B. I. Halperin, "Effects of finite gradient-pulse widths in pulsed-field-gradient diffusion measurements," *Journal of Magnetic Resonance, Series A*, vol. 113, no. 1, pp. 94–101, 1995.

[112] Y. Assaf, H. Johansen-Berg, and M. Thiebaut de Schotten, "The role of diffusion mri in neuroscience," *NMR in Biomedicine*, vol. 32, no. 4, p. e3762, 2019.

[113] P. J. Basser and E. Ozarslan, "Anisotropic diffusion: from the apparent diffusion coefficient to the apparent diffusion tensor," *Diffusion MRI: Theory, Methods, and Applica tions. Oxford University Press, Oxford*, pp. 79–91, 2010.

[114] H. C. Torrey, "Bloch equations with diffusion terms," *Physical review*, vol. 104, no. 3, p. 563, 1956.

[115] D. Topgaard, "Multidimensional diffusion mri," *Journal of Magnetic Resonance*, vol. 275, pp. 98–113, 2017.

[116] D. K. Jones and M. Cercignani, "Twenty-five pitfalls in the analysis of diffusion mri data," *NMR in Biomedicine*, vol. 23, no. 7, pp. 803–820, 2010.

[117] J. Veraart, J. Sijbers, S. Sunaert, A. Leemans, and B. Jeurissen, "Weighted linear least squares estimation of diffusion mri parameters: strengths, limitations, and pitfalls," *Neuroimage*, vol. 81, pp. 335–346, 2013.

[118] P. J. Basser and C. Pierpaoli, "Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor mri," *Journal of magnetic resonance*, vol. 111, no. 2, pp. 209–219, 1996.

[119] P. S. Huppi and J. Dubois, "Diffusion tensor imaging of brain development," in *Seminars in Fetal and Neonatal Medicine*, vol. 11, pp. 489–497, Elsevier, 2006.

[120] E. V. Sullivan and A. Pfefferbaum, "Diffusion tensor imaging and aging," *Neuroscience & Biobehavioral Reviews*, vol. 30, no. 6, pp. 749–761, 2006.

[121] D. J. Madden, I. J. Bennett, A. Burzynska, G. G. Potter, N.-k. Chen, and A. W. Song, "Diffusion tensor imaging of cerebral white matter integrity in cognitive aging," *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease*, vol. 1822, no. 3, pp. 386–400, 2012.

[122] D. Werring, C. Clark, G. Barker, A. Thompson, and D. Miller, "Diffusion tensor imaging of lesions and normal-appearing white matter in multiple sclerosis," *Neurology*, vol. 52, no. 8, pp. 1626–1626, 1999.

[123] H. Stotesbury, F. J. Kirkham, M. Kölbel, P. Balfour, J. D. Clayden, S. Sahota, S. Sakaria, D. E. Saunders, J. Howard, R. Kesse-Adu, *et al.*, "White matter integrity and processing speed in sickle cell anemia," *Neurology*, vol. 90, no. 23, pp. e2042–e2050, 2018.

[124] V. G. Kiselev, "The cumulant expansion: an overarching mathematical framework for understanding diffusion nmr," *Diffusion MRI: theory, methods, and applications*, pp. 152–168, 2010.

[125] R. N. Henriques, S. N. Jespersen, and N. Shemesh, "Correlation tensor magnetic resonance imaging," *NeuroImage*, vol. 211, p. 116605, 2020.

[126] D. S. Novikov, E. Fieremans, S. N. Jespersen, and V. G. Kiselev, "Quantifying brain microstructure with diffusion mri: Theory and parameter estimation," *NMR in Biomedicine*, vol. 32, no. 4, p. e3998, 2019.

[127] I. O. Jelescu, J. Veraart, E. Fieremans, and D. S. Novikov, "Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue," *NMR in Biomedicine*, vol. 29, no. 1, pp. 33–47, 2016.

[128] A. Daducci, E. J. Canales-Rodríguez, H. Zhang, T. B. Dyrby, D. C. Alexander, and J.-P. Thiran, "Accelerated microstructure imaging via convex optimization (amico) from diffusion mri data," *NeuroImage*, vol. 105, pp. 32–44, 2015.

[129] D. Barazany, P. J. Basser, and Y. Assaf, "In vivo measurement of axon diameter distribution in the corpus callosum of rat brain," *Brain*, vol. 132, no. 5, pp. 1210–1220, 2009.

[130] S. N. Jespersen, C. R. Bjarkam, J. R. Nyengaard, M. M. Chakravarty, B. Hansen, T. Vosegaard, L. Østergaard, D. Yablonskiy, N. C. Nielsen, and P. Vestergaard-Poulsen, "Neurite density from magnetic resonance diffusion measurements at ultrahigh field: comparison with light microscopy and electron microscopy," *Neuroimage*, vol. 49, no. 1, pp. 205–216, 2010.

[131] G. M. Innocenti, R. Caminiti, and F. Aboitiz, "Comments on the paper by horowitz et al.(2014)," *Brain Structure & Function*, vol. 220, no. ARTICLE, pp. 1789–1790, 2015.

[132] B. Lampinen, F. Szczepankiewicz, J. Mårtensson, D. van Westen, P. C. Sundgren, and M. Nilsson, "Neurite density imaging versus imaging of microscopic anisotropy in diffusion mri: A model comparison using spherical tensor encoding," *Neuroimage*, vol. 147, pp. 517–531, 2017.

[133] R. N. Henriques, S. N. Jespersen, and N. Shemesh, "Microscopic anisotropy misestimation in spherical-mean single diffusion encoding mri," *Magnetic resonance in medicine*, vol. 81, no. 5, pp. 3245–3261, 2019.

[134] S. N. Jespersen, H. Lundell, C. K. Sønderby, and T. B. Dyrby, "Orientationally invariant metrics of apparent compartment eccentricity from double pulsed field gradient diffusion experiments," *NMR in Biomedicine*, vol. 26, no. 12, pp. 1647–1662, 2013.

[135] S. Lasič, F. Szczepankiewicz, S. Eriksson, M. Nilsson, and D. Topgaard, "Microanisotropy imaging: quantification of microscopic diffusion anisotropy and orientational order parameter by diffusion mri with

magic-angle spinning of the q-vector," *Frontiers in Physics*, vol. 2, p. 11, 2014.

[136] J.-J. Wang, W.-Y. Lin, C.-S. Lu, Y.-H. Weng, S.-H. Ng, C.-H. Wang, H.-L. Liu, R.-H. Hsieh, Y.-L. Wan, and Y.-Y. Wai, "Parkinson disease: diagnostic utility of diffusion kurtosis imaging," *Radiology*, vol. 261, no. 1, pp. 210–217, 2011.

[137] F.-L. Zhou, Z. Li, J. E. Gough, P. L. H. Cristinacce, and G. J. Parker, "Axon mimicking hydrophilic hollow polycaprolactone microfibres for diffusion magnetic resonance imaging," *Materials & design*, vol. 137, pp. 394–403, 2018.

[138] K. G. Schilling, V. Janve, Y. Gao, I. Stepniewska, B. A. Landman, and A. W. Anderson, "Histological validation of diffusion mri fiber orientation distributions and dispersion," *Neuroimage*, vol. 165, pp. 200–221, 2018.

[139] A. V. Barzykin, "Exact solution of the torrey-bloch equation for a spin echo in restricted geometries," *Physical Review B*, vol. 58, no. 21, p. 14171, 1998.

[140] A. Szafer, J. Zhong, and J. C. Gore, "Theoretical model for water diffusion in tissues," *Magnetic resonance in medicine*, vol. 33, no. 5, pp. 697–712, 1995.

[141] T. G. Close, J.-D. Tournier, F. Calamante, L. A. Johnston, I. Mareels, and A. Connelly, "A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms," *NeuroImage*, vol. 47, no. 4, pp. 1288–1300, 2009.

[142] P. T. Callaghan, "A simple matrix formalism for spin echo analysis of restricted diffusion under generalized gradient waveforms," *Journal of Magnetic Resonance*, vol. 129, no. 1, pp. 74–84, 1997.

[143] A. Ianuş, D. C. Alexander, and I. Drobnjak, "Microstructure imaging sequence simulation toolbox," in *International Workshop on Simulation and Synthesis in Medical Imaging*, pp. 34–44, Springer, 2016.

[144] C.-L. Chin, F. W. Wehrli, S. N. Hwang, M. Takahashi, and D. B. Hackney, "Biexponential diffusion attenuation in the rat spinal cord: computer simulations based on anatomic images of axonal architecture," *Magnetic*

Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine, vol. 47, no. 3, pp. 455–460, 2002.

[145] J.-R. Li, V.-D. Nguyen, T. N. Tran, J. Valdman, C.-B. Trang, K. Van Nguyen, D. T. S. Vu, H. A. Tran, H. T. A. Tran, and T. M. P. Nguyen, "Spindoctor: A matlab toolbox for diffusion mri simulation," *NeuroImage*, vol. 202, p. 116120, 2019.

[146] N. M. Naughton, C. G. Tennyson, and J. G. Georgiadis, "Lattice boltzmann method for simulation of diffusion magnetic resonance imaging physics in multiphase tissue models," *Physical Review E*, vol. 102, no. 4, p. 043305, 2020.

[147] M. G. Hall and D. C. Alexander, "Convergence and parameter choice for monte-carlo simulations of diffusion mri," *IEEE transactions on medical imaging*, vol. 28, no. 9, pp. 1354–1364, 2009.

[148] G. T. Balls and L. R. Frank, "A simulation environment for diffusion weighted mr experiments in complex media," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 62, no. 3, pp. 771–778, 2009.

[149] M. Palombo, D. C. Alexander, and H. Zhang, "A generative model of realistic brain cells with application to numerical simulation of the diffusion-weighted mr signal," *NeuroImage*, vol. 188, pp. 391–402, 2019.

[150] N. Metropolis and S. Ulam, "The monte carlo method," *Journal of the American statistical association*, vol. 44, no. 247, pp. 335–341, 1949.

[151] R. Y. Rubinstein and D. P. Kroese, *Simulation and the Monte Carlo method*, vol. 10. John Wiley & Sons, 2016.

[152] H.-G. Lipinski, "Monte carlo simulation of extracellular diffusion in brain tissues," *Physics in Medicine & Biology*, vol. 35, no. 3, p. 441, 1990.

[153] P. Cook, Y. Bai, S. Nedjati-Gilani, K. Seunarine, M. Hall, G. Parker, and D. C. Alexander, "Camino: open-source diffusion-mri reconstruction and processing," in *14th scientific meeting of the international society for magnetic resonance in medicine*, vol. 2759, p. 2759, Seattle WA, USA, 2006.

[154] J. Rafael-Patino, D. Romascano, A. Ramirez-Manzanares, E. J. Canales-Rodríguez, G. Girard, and J.-P. Thiran, "Robust monte-carlo simulations in diffusion-mri: Effect of the substrate complexity and parameter choice on the reproducibility of results," *Frontiers in Neuroinformatics*, vol. 14, p. 8, 2020.

[155] J. Nickolls, I. Buck, M. Garland, and K. Skadron, "Scalable parallel programming with cuda," *Queue*, vol. 6, no. 2, pp. 40–53, 2008.

[156] K.-V. Nguyen, E. Hernández-Garzón, and J. Valette, "Efficient gpu-based monte-carlo simulation of diffusion in real astrocytes reconstructed from confocal microscopy," *Journal of Magnetic Resonance*, vol. 296, pp. 188–199, 2018.

[157] L. Kerkelä, F. Nery, M. G. Hall, and C. A. Clark, "Disimpy: A massively parallel monte carlo simulator for generating diffusion-weighted mri data in python," *Journal of Open Source Software*, vol. 5, no. 52, p. 2527, 2020.

[158] D. Blackman and S. Vigna, "Scrambled linear pseudorandom number generators," *arXiv preprint arXiv:1805.01407*, 2018.

[159] M. E. Muller, "A note on a method for generating points uniformly on n-dimensional spheres," *Communications of the ACM*, vol. 2, no. 4, pp. 19–20, 1959.

[160] T. Möller and B. Trumbore, "Fast, minimum storage ray-triangle intersection," *Journal of graphics tools*, vol. 2, no. 1, pp. 21–28, 1997.

[161] M. G. Hall, G. Nedjati-Gilani, and D. C. Alexander, "Realistic voxel sizes and reduced signal variation in monte-carlo simulation for diffusion mr data synthesis," *arXiv preprint arXiv:1701.03634*, 2017.

[162] E. Panagiotaki, M. G. Hall, H. Zhang, B. Siow, M. F. Lythgoe, and D. C. Alexander, "High-fidelity meshes from tissue samples for diffusion mri simulations," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 404–411, Springer, 2010.

[163] S. v. d. Walt, S. C. Colbert, and G. Varoquaux, "The numpy array: a structure for efficient numerical computation," *Computing in Science & Engineering*, vol. 13, no. 2, pp. 22–30, 2011.

[164] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, *et al.*, "Scipy 1.0: fundamental algorithms for scientific computing in python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.

[165] S. K. Lam, A. Pitrou, and S. Seibert, "Numba: A llvm-based python jit compiler," in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 1–6, 2015.

[166] L. Avram, E. Özarslan, Y. Assaf, A. Bar-Shir, Y. Cohen, and P. J. Basser, "Three-dimensional water diffusion in impermeable cylindrical tubes: theory versus experiments," *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In vivo*, vol. 21, no. 8, pp. 888–898, 2008.

[167] P. T. Callaghan, *Principles of nuclear magnetic resonance microscopy*. Oxford University Press on Demand, 1993.

[168] R. Callaghan, D. C. Alexander, M. Palombo, and H. Zhang, "Config: Contextual fibre growth to generate realistic axonal packing for diffusion mri simulation," *Neuroimage*, vol. 220, p. 117107, 2020.

[169] D. S. Novikov, J. Veraart, I. O. Jelescu, and E. Fieremans, "Rotationally-invariant mapping of scalar and orientational metrics of neuronal microstructure with diffusion mri," *NeuroImage*, vol. 174, pp. 518–538, 2018.

[170] N. Colgan, B. Siow, J. M. O'Callaghan, I. F. Harrison, J. A. Wells, H. E. Holmes, O. Ismail, S. Richardson, D. C. Alexander, E. C. Collins, *et al.*, "Application of neurite orientation dispersion and density imaging (noddi) to a tau pathology model of alzheimer's disease," *Neuroimage*, vol. 125, pp. 739–744, 2016.

[171] K. Kamagata, T. Hatano, A. Okuzumi, Y. Motoi, O. Abe, K. Shimoji, K. Kamiya, M. Suzuki, M. Hori, K. K. Kumamaru, *et al.*, "Neurite orientation dispersion and density imaging in the substantia nigra in idiopathic parkinson disease," *European radiology*, vol. 26, no. 8, pp. 2567–2577, 2016.

[172] T. Schneider, W. Brownlee, H. Zhang, O. Ciccarelli, D. H. Miller, and C. G. Wheeler-Kingshott, "Sensitivity of multi-shell noddi to multiple sclero-

sis white matter changes: a pilot study," *Functional neurology*, vol. 32, no. 2, p. 97, 2017.

[173] S. By, J. Xu, B. A. Box, F. R. Bagnato, and S. A. Smith, "Multi-compartmental diffusion characterization of the human cervical spinal cord in vivo using the spherical mean technique," *NMR in Biomedicine*, vol. 31, no. 4, p. e3894, 2018.

[174] F. Bagnato, G. Franco, H. Li, E. Kaden, F. Ye, R. Fan, A. Chen, D. C. Alexander, S. A. Smith, R. Dortch, *et al.*, "Probing axons using multi-compartmental diffusion in multiple sclerosis," *Annals of Clinical and Translational Neurology*, vol. 6, no. 9, pp. 1595–1605, 2019.

[175] H. E. Cooper, E. Kaden, L. F. Halliday, D.-E. Bamiou, K. Mankad, C. Peters, and C. A. Clark, "White matter microstructural abnormalities in children with severe congenital hypothyroidism," *NeuroImage: Clinical*, vol. 24, p. 101980, 2019.

[176] J. D. Clayden, S. M. Maniega, A. J. Storkey, M. D. King, M. E. Bastin, C. A. Clark, *et al.*, "Tractor: magnetic resonance imaging and tractography with r," *Journal of Statistical Software*, vol. 44, no. 8, pp. 1–18, 2011.

[177] J.-D. Tournier, R. Smith, D. Raffelt, R. Tabbara, T. Dhollander, M. Pietsch, D. Christiaens, B. Jeurissen, C.-H. Yeh, and A. Connelly, "Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation," *NeuroImage*, p. 116137, 2019.

[178] J. Veraart, D. S. Novikov, D. Christiaens, B. Ades-Aron, J. Sijbers, and E. Fieremans, "Denoising of diffusion mri using random matrix theory," *Neuroimage*, vol. 142, pp. 394–406, 2016.

[179] E. Kellner, B. Dhital, V. G. Kiselev, and M. Reisert, "Gibbs-ringing artifact removal based on local subvoxel-shifts," *Magnetic resonance in medicine*, vol. 76, no. 5, pp. 1574–1581, 2016.

[180] J. L. Andersson and S. N. Sotiropoulos, "An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging," *Neuroimage*, vol. 125, pp. 1063–1078, 2016.

[181] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith, "Fsl," *Neuroimage*, vol. 62, no. 2, pp. 782–790, 2012.

[182] B. Fischl, "Freesurfer," *Neuroimage*, vol. 62, no. 2, pp. 774–781, 2012.

[183] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans, "Automatic 3d intersubject registration of mr volumetric data in standardized talairach space.," *Journal of computer assisted tomography*, vol. 18, no. 2, pp. 192–205, 1994.

[184] E. Garyfallidis, M. Brett, B. Amirbekian, A. Rokem, S. Van Der Walt, M. Descoteaux, and I. Nimmo-Smith, "Dipy, a library for the analysis of diffusion mri data," *Frontiers in neuroinformatics*, vol. 8, p. 8, 2014.

[185] G. H. Zhang, "Noddi matlab toolbox v1.0.1." https://www.nitrc.org/projects/noddi_toolbox, 2012-2018.

[186] E. Kaden, "smt v0.4." https://github.com/ekaden/smt, 2016-2019.

[187] S. R. Searle, G. Casella, and C. E. McCulloch, *Variance components*, vol. 391. John Wiley & Sons, 2009.

[188] D. Liljequist, B. Elfving, and K. S. Roaldsen, "Intraclass correlation–a discussion and demonstration of basic features," *PloS one*, vol. 14, no. 7, 2019.

[189] P. E. Shrout and J. L. Fleiss, "Intraclass correlations: uses in assessing rater reliability.," *Psychological bulletin*, vol. 86, no. 2, p. 420, 1979.

[190] E. M. Palacios, A. J. Martin, M. A. Boss, F. Ezekiel, Y. S. Chang, E. L. Yuh, M. J. Vassar, D. M. Schnyer, C. L. MacDonald, K. L. Crawford, *et al.*, "Toward precision and reproducibility of diffusion tensor imaging: a multicenter diffusion phantom and traveling volunteer study," *American Journal of Neuroradiology*, vol. 38, no. 3, pp. 537–545, 2017.

[191] L. W. Kasa, R. A. Haast, T. K. Kuehn, F. N. Mushtaha, C. A. Baron, T. Peters, and A. R. Khan, "Evaluating high spatial resolution diffusion kurtosis imaging at 3t: Reproducibility and quality of fit," *Journal of Magnetic Resonance Imaging*, 2020.

[192] R. Neto Henriques, *Advanced Methods for Diffusion MRI Data Analysis and their Application to the Healthy Ageing Brain.* PhD thesis, University of Cambridge, 2018.

[193] A. W. Chung, K. K. Seunarine, and C. A. Clark, "Noddi reproducibility and variability with magnetic field strength: A comparison between 1.5 t and 3 t," *Human brain mapping*, vol. 37, no. 12, pp. 4550–4565, 2016.

[194] P. McCunn, K. M. Gilbert, P. Zeman, A. X. Li, M. J. Strong, A. R. Khan, and R. Bartha, "Reproducibility of neurite orientation dispersion and density imaging (noddi) in rats at 9.4 tesla," *PloS one*, vol. 14, no. 4, p. e0215974, 2019.

[195] D. Cory, A. Garroway, and J. Miller, "Applications of spin transport as a probe of local geometry," in *Abstracts of Papers of the American Chemical Society*, vol. 199, pp. 105–POLY, AMER CHEMICAL SOC 1155 16TH ST, NW, WASHINGTON, DC 20036, 1990.

[196] P. P. Mitra, "Multiple wave-vector extensions of the nmr pulsed-field-gradient spin-echo diffusion measurement," *Physical Review B*, vol. 51, no. 21, p. 15074, 1995.

[197] P. Callaghan and I. Furo, "Diffusion-diffusion correlation and exchange as a signature for local order and dynamics," *The Journal of chemical physics*, vol. 120, no. 8, pp. 4032–4038, 2004.

[198] N. Shemesh, E. Özarslan, T. Adiri, P. J. Basser, and Y. Cohen, "Noninvasive bipolar double-pulsed-field-gradient nmr reveals signatures for pore size and shape in polydisperse, randomly oriented, inhomogeneous porous media," *The Journal of chemical physics*, vol. 133, no. 4, p. 044705, 2010.

[199] N. Shemesh and Y. Cohen, "Microscopic and compartment shape anisotropies in gray and white matter revealed by angular bipolar double-pfg mr," *Magnetic resonance in medicine*, vol. 65, no. 5, pp. 1216–1227, 2011.

[200] S. N. Jespersen, "Equivalence of double and single wave vector diffusion contrast at low diffusion weighting," *NMR in Biomedicine*, vol. 25, no. 6, pp. 813–818, 2012.

[201] M. Lawrenz, M. A. Koch, and J. Finsterbusch, "A tensor model and measures of microscopic anisotropy for double-wave-vector diffusion-weighting experiments with long mixing times," *Journal of Magnetic Resonance*, vol. 202, no. 1, pp. 43–56, 2010.

[202] F. Szczepankiewicz, J. Sjölund, F. Ståhlberg, J. Lätt, and M. Nilsson, "Tensor-valued diffusion encoding for diffusional variance decompo-

sition (divide): Technical feasibility in clinical mri systems," *PloS one*, vol. 14, no. 3, 2019.

[203] S. N. Jespersen and N. Buhl, "The displacement correlation tensor: microstructure, ensemble anisotropy and curving fibers," *Journal of Magnetic Resonance*, vol. 208, no. 1, pp. 34–43, 2011.

[204] A. Ianuş, S. N. Jespersen, T. S. Duarte, D. C. Alexander, I. Drobnjak, and N. Shemesh, "Accurate estimation of microscopic diffusion anisotropy and its time dependence in the mouse brain," *Neuroimage*, vol. 183, pp. 934–949, 2018.

[205] M. Neeman, J. P. Freyer, and L. O. Sillerud, "A simple method for obtaining cross-term-free images for diffusion anisotropy studies in nmr microimaging," *Magnetic Resonance in Medicine*, vol. 21, no. 1, pp. 138–143, 1991.

[206] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, "Efficient subpixel image registration algorithms," *Optics letters*, vol. 33, no. 2, pp. 156–158, 2008.

[207] M. A. Branch, T. F. Coleman, and Y. Li, "A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems," *SIAM Journal on Scientific Computing*, vol. 21, no. 1, pp. 1–23, 1999.

[208] F. Szczepankiewicz, C.-F. Westin, and M. Nilsson, "Maxwell-compensated design of asymmetric gradient waveforms for tensor-valued diffusion encoding," *Magnetic resonance in medicine*, vol. 82, no. 4, pp. 1424–1437, 2019.

[209] J. Sjölund, F. Szczepankiewicz, M. Nilsson, D. Topgaard, C.-F. Westin, and H. Knutsson, "Constrained optimization of gradient waveforms for generalized diffusion encoding," *Journal of magnetic resonance*, vol. 261, pp. 157–168, 2015.

[210] H. Lundell, M. Nilsson, T. Dyrby, G. Parker, P. H. Cristinacce, F.-L. Zhou, D. Topgaard, and S. Lasič, "Multidimensional diffusion mri with spectrally modulated gradients reveals unprecedented microstructural detail," *Scientific reports*, vol. 9, no. 1, pp. 1–12, 2019.

[211] T. M. de Swiet and P. P. Mitra, "Possible systematic errors in single-shot measurements of the trace of the diffusion tensor," *Journal of Magnetic Resonance, Series B*, vol. 111, no. 1, pp. 15–22, 1996.

[212] S. N. Jespersen, J. L. Olesen, A. Ianuş, and N. Shemesh, "Effects of nongaussian diffusion on "isotropic diffusion" measurements: An ex-vivo microimaging and simulation study," *Journal of Magnetic Resonance*, vol. 300, pp. 84–94, 2019.

[213] F. Szczepankiewicz, S. Hoge, and C.-F. Westin, "Linear, planar and spherical tensor-valued diffusion mri data by free waveform encoding in healthy brain, water, oil and liquid crystals," *Data in brief*, vol. 25, p. 104208, 2019.

[214] J. Stepišnik, "Validity limits of gaussian approximation in cumulant expansion for diffusion attenuation of spin echo," *Physica B: Condensed Matter*, vol. 270, no. 1-2, pp. 110–117, 1999.

[215] G. Strang and K. Borre, *Linear algebra, geodesy, and GPS*. Siam, 1997.

[216] J. P. de Almeida Martins and D. Topgaard, "Two-dimensional correlation of isotropic and directional diffusion using nmr," *Physical review letters*, vol. 116, no. 8, p. 087601, 2016.

[217] A. Reymbaut, P. Mezzani, J. P. de A lmeida Martins, and D. Topgaard, "Accuracy and precision of statistical descriptors obtained from multidimensional diffusion signal inversion algorithms," *NMR in Biomedicine*, p. e4267, 2020.

[218] J. Martin, S. Endt, A. Wetscherek, T. A. Kuder, A. Doerfler, M. Uder, B. Hensel, and F. B. Laun, "Contrast-to-noise ratio analysis of microscopic diffusion anisotropy indices in q-space trajectory imaging," *Zeitschrift für Medizinische Physik*, vol. 30, no. 1, pp. 4–16, 2020.

[219] L. Kerkelä, R. N. Henriques, M. G. Hall, C. A. Clark, and N. Shemesh, "Validation and noise robustness assessment of microscopic anisotropy estimation with clinically feasible double diffusion encoding mri," *Magnetic Resonance in Medicine*, vol. 83, no. 5, pp. 1698–1710, 2020.

[220] N. Gyori, M. G. Hall, C. A. Clark, D. C. Alexander, and E. Kaden, "Discrepancy between in-vivo measurements and monte-carlo simulations

with spherical structures in b-tensor encoding," in *International Society for Magnetic Resonance in Medicine and Society for MR Radiographers & Technologists Virtual Conference & Exhibition*, 2020.

[221] F. Szczepankiewicz, C.-F. Westin, and M. Nilsson, "Gradient waveform design for tensor-valued encoding in diffusion mri," *Journal of Neuroscience Methods*, p. 109007, 2020.